

Segmentation of Glomeruli Within Trichrome Images Using Deep Learning



Shruti Kannan¹, Laura A. Morgan², Benjamin Liang², McKenzie G. Cheung², Christopher Q. Lin², Dan Mun³, Ralph G. Nader⁴, Mostafa E. Belghasem⁵, Joel M. Henderson⁵, Jean M. Francis⁴, Vipul C. Chitalia^{4,5,6,7} and Vijaya B. Kolachalama^{1,6,8}

¹Section of Computational Biomedicine, Department of Medicine, Boston University School of Medicine, Boston, Massachusetts, USA; ²College of Engineering, Boston University, Boston, Massachusetts, USA; ³College of Health & Rehabilitation Sciences, Sargent College, Boston University, Boston, Massachusetts, USA; ⁴Renal Section, Department of Medicine, Boston University School of Medicine, Boston, Massachusetts, USA; ⁵Department of Pathology and Laboratory Medicine, Boston University School of Medicine, Boston, Massachusetts, USA; ⁶Whitaker Cardiovascular Institute, Boston University School of Medicine, Boston, Massachusetts, USA; ⁷Veterans Administration Boston Healthcare System, Boston, Massachusetts, USA; and ⁸Hariri Institute for Computing and Computational Science & Engineering, Boston University, Boston, Massachusetts, USA

Introduction: The number of glomeruli and glomerulosclerosis evaluated on kidney biopsy slides constitute standard components of a renal pathology report. Prevailing methods for glomerular assessment remain manual, labor intensive, and nonstandardized. We developed a deep learning framework to accurately identify and segment glomeruli from digitized images of human kidney biopsies.

Methods: Trichrome-stained images ($n = 275$) from renal biopsies of 171 patients with chronic kidney disease treated at the Boston Medical Center from 2009 to 2012 were analyzed. A sliding window operation was defined to crop each original image to smaller images. Each cropped image was then evaluated by at least 3 experts into 3 categories: (i) no glomerulus, (ii) normal or partially sclerosed (NPS) glomerulus, and (iii) globally sclerosed (GS) glomerulus. This led to identification of 751 unique images representing nonglomerular regions, 611 images with NPS glomeruli, and 134 images with GS glomeruli. A convolutional neural network (CNN) was trained with cropped images as inputs and corresponding labels as output. Using this model, an image processing routine was developed to scan the test images to segment the GS glomeruli.

Results: The CNN model was able to accurately discriminate nonglomerular images from NPS and GS images (performance on test data: Accuracy: $92.67\% \pm 2.02\%$ and Kappa: 0.8681 ± 0.0392). The segmentation model that was based on the CNN multilabel classifier accurately marked the GS glomeruli on the test data (Matthews correlation coefficient = 0.628).

Conclusion: This work demonstrates the power of deep learning for assessing complex histologic structures from digitized human kidney biopsies.

Kidney Int Rep (2019) 4, 955–962; <https://doi.org/10.1016/j.ekir.2019.04.008>

KEYWORDS: computational pathology; deep learning; digital pathology; glomerulus; image segmentation; kidney biopsy; trichrome stain

© 2019 International Society of Nephrology. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

See Commentary on Page 914

Glomerular damage is a common manifestation in a spectrum of renal diseases that lead to chronic kidney disease and end-stage renal disease.¹ Morphological and ultrastructural alterations within the glomeruli provide invaluable information on the mechanisms of

renal impairment and facilitate accurate clinical diagnosis.^{1–3} Assessment of this highly relevant structure is therefore integral to histopathological analysis of kidney biopsies. A fundamental morphologic parameter in nephropathology is the quantification of normal and abnormal glomeruli in the light microscopic material. For example, the number of glomeruli is required for the assessment of tissue sufficiency in kidney transplant pathology. Also, histological analysis of glomerular diseases involves careful examination of the entire kidney biopsy slide, and this includes, in part, identification of all the glomeruli, assessment of the state of each glomerulus, and integration of these data with other parameters to pinpoint the

Correspondence: Vijaya B. Kolachalama, Boston University School of Medicine, 72 E. Concord Street, Evans 636, Boston, Massachusetts 02118, USA. E-mail: vkola@bu.edu

Received 24 October 2018; revised 4 April 2019; accepted 8 April 2019; published online 15 March 2019

diagnosis of the glomerular disease.^{4–7} Although this multistep process of counting and assessing all the glomeruli can be handled efficiently at large medical centers under the supervision of an in-house nephropathologist, this expertise is not available at all locations across the globe. In addition, even for the institutions that have the expert nephropathologist, we need approaches that can automatically perform some of these tasks in order to assist clinical practice to maximize their efficiency.

Machine learning (ML), a powerful technique that is increasingly being used in medicine, has the ability to perform these tasks in an efficient fashion. ML approaches give computers the ability to integrate discrete datasets in an agnostic manner to detect previously indecipherable patterns and generate a disease-specific fingerprint. ML can leverage many images as inputs and correlate patterns and features with clinical outcomes. Building on the advances of ML, scientists recently have developed so-called “deep learning” frameworks, such as CNN for object recognition and classification.⁸ We have previously leveraged these self-learning approaches to quantify key renal pathological features, such as fibrosis and associated with various outcomes of interest.⁹ We now demonstrate a framework to automatically identify and segment NPS glomeruli as well as the glomeruli with GS present within digitized images of human kidney biopsies.

MATERIALS AND METHODS

Data Collection

Anonymized kidney biopsies were obtained and digitized after approval by the Boston University Medical Campus’ institutional review board under waiver of consent (H-26367). Kidney biopsy procedures were performed on selected patients treated at Boston Medical Center between January 2009 and December 2012 (Table 1). In total, 171 kidney biopsy slides were available for subsequent imaging. These biopsy

Table 1. Patient and digitized kidney biopsy characteristics used for this study

Characteristic	Value
Number of patients	171
Age, yr, median (range)	52 (19–86)
Male, %	59.6
Patients per race/ethnicity (white, black, Hispanic, other)	46, 79, 24, 22
Body mass index, kg/m ² , median (range)	28.94 (15–56.2)
Creatinine, mg/dl, median (range)	2.31 (0.54–13.29)
Estimated glomerular filtration rate, ml/min per 1.73 m ² , median (range)	30 (5–163)
Proteinuria, g/g, median (range)	1.79 (0.03–20.5)
Number of unique images	275
Total number of glomeruli	745
Number of normal or partially sclerosed glomeruli	611
Number of globally sclerosed glomeruli	134

samples were obtained from adult patients who had a native or an allograft biopsy, independent of the indication for the biopsy procedure.⁹ The criterion for inclusion was the availability of pathology slides.

Imaging

Biopsy samples were obtained in the form of individual trichrome-stained slides prepared from formalin-fixed, paraffin-embedded core-needle biopsy tissue. A selected core visible on each slide was imaged at $\times 40$ magnification (indicating a $\times 4$ objective and a $\times 10$ eyepiece) using a Nikon Eclipse TE-2000 microscope (Melville, NY; <http://www.bumc.bu.edu/busm/research/cores/>). Images were generated with a special consideration to cover the entirety of the biopsy sample which resulted in multiple $\times 40$ images per patient. All the images were manually focused using the NIS-Elements AR software (Nikon, Tokyo, Japan) that was installed on the computer connected to the microscope. A total of 275 unique images ($\sim 2 \times 40$ images per patient) were used and the average size of each image was approximately $2560 \times 1920 \times 3$ pixels, corresponding to a field of 2.176×1.632 mm², which resulted in a length scale of $0.85 \mu\text{m}/\text{pixel}$. These images were then converted to 8-bit red–green–blue color images in TIFF format.

Glomerular Dataset Generation

Each of the digitized images represent a large portion of the digitized biopsy (Supplementary Figure S1), and the information contained within them had to be filtered to train a glomerular image classifier. Therefore, we created a dataset that was more amenable for CNN model training using the sliding window operator. This operation allowed us to systematically crop the original images ($n = 275$) into smaller images of size $300 \times 300 \times 3$ pixels (Figure 1). As a result, we obtained 745 images with glomeruli and 751 images with non-glomerular regions, each of size $300 \times 300 \times 3$ pixels. A team of clinical experts comprising 3 nephrologists and 1 nephropathologist participated in the review of the cropped images. Each image was first independently reviewed by 2 experts. If there was agreement between them, then this image was selected for further model training. If there was disagreement, then a third review was solicited, and the majority opinion among the 3 experts was considered as final. The experts were in full agreement with all the 751 nonglomerular images. We used the definition of NPS and GS glomeruli based on the guidelines published in D’Agati et al.¹⁰ Briefly, the glomeruli with completely open capillary tufts were considered normal. The glomeruli with incomplete obliteration of their tufts were grouped as those with partial glomerulosclerosis. The glomeruli with complete obliteration of their entire glomerular tuft and loss of

their architecture and normal shape were considered as GS. Both normal and NPS glomeruli were grouped together in one class, and GS glomeruli in another class. All these images were used for model training. See the [Supplementary Appendix S1](#) for more details.

Model Training for Glomerular Classification

We used Google's Inception v3 CNN architecture, which was pretrained on millions of images with 1000 object classes,¹¹ incorporated minor changes to fine-

tune the framework, and trained it to predict the presence or absence of a glomerulus within the cropped trichrome images ([Figure 2](#)). Specifically, we removed the final classification layer from the network and retrained it with our dataset using the 3 output labels (no glomerulus, NPS glomerulus, and GS glomerulus). We then performed fine-tuning of the parameters at all layers. This procedure, known as transfer learning, is optimal, given the amount of data available. See the [Supplementary Appendix S1](#) for more details.




Figure 1. Cropped images. The sliding window operator was used to generate different sets of images to train the convolutional neural network model. The first column (a) contains images with nonglomerular tissue, the second column (b) contains images with either a single normal or partially sclerosed glomerulus, and the third column (c) contains images with a single globally sclerosed glomerulus. Each cropped image is of size $300 \times 300 \times 3$ pixels. Trichrome stain. A single core on the biopsy slide was imaged at original magnification $\times 40$.




Figure 2. Schematic of the deep neural network. Our classification technique is based on leveraging a pretrained convolutional neural network, which was fine-tuned on our dataset (see Methods). The architecture is reprinted with permission (<https://research.googleblog.com/2016/03/train-your-own-image-classifier-with.html>).

The cropped image dataset was randomly split at the patient level. Specifically, to capture intra- and inter-patient variabilities, and to verify whether the CNN model works well on images and image characteristics that it has not been trained on, the patient list was randomly split into 2 parts in a 7:3 ratio (70% training, 30% testing). This resulted in 120 patients in the training set, and 51 patients in the test set. Cropped images belonging to each patient on the list were included in the corresponding dataset (training vs. testing). Also, for consistency, we repeated the process of random splitting 4 times. CNN model training and testing were performed on each split, and average performance values were recorded.

Data Augmentation

Some of the glomeruli on the biopsy images were observed on the edges of the tissue sample. When cropping was performed to capture these cases, a portion of the cropped region had only the background pixels. All these images were used as part of the training data, but they were not in sufficient number to be able to generate a model that could accurately identify the glomeruli present on the edges of the biopsy. We therefore augmented the training data by creating copies ($n = 5$) of each image by randomly whitening a small fraction ($= 0.2$) of the total pixels in the images, resulting in 6 total images per original cropped image (Figure 3). See the [Supplementary Appendix S1](#) for more details.

Image Segmentation

Using the CNN model with the best test performance, we tested an image processing routine to scan the test images to identify and segment the GS glomeruli (Figure 4). The sliding window operation was used again to scan the entire test image of size approximately $2560 \times 1920 \times 3$ pixels in increments of $300 \times 300 \times 3$ pixels. Each cropped image was then processed through the trained CNN model that predicted if there was a GS glomerulus. An output of “0” indicated that the CNN model determined that no glomerulus was

present, whereas an output of “2” indicated a GS glomerulus was detected within that cropped image. Note that an output of “1” was reserved for identifying an NPS glomerulus, but this result was not used for glomerular segmentation. When a GS glomerulus was detected, the pixel coordinates of the 4 corners of that image were stored in an array. This process was repeated as the sliding window operation swept from one end of the corner to the other, which resulted in bright patches that corresponded to the areas that were predicted to contain a GS glomerulus. A heatmap was generated using these corners. The brightness of the patch in the heatmap was found to be directly proportional to how confident the model was in terms of detecting the presence of a GS glomerulus in that area. Every nonbright region (i.e., area with pixel intensity close to 0) on the heatmap then represents all the nonglomerular regions.

Generated heatmaps were processed further to segment the identified GS glomeruli using a simple annotation defined as a “green box” surrounding the GS glomerular region. We performed this task by first binarizing the image using Otsu’s method.¹² Note that the threshold value for binarization was empirically determined ($= 20$), after examining several images. Subsequently, a distance transform was applied on the heatmaps, which simply calculated the distance of each foreground pixel from the nearest background pixel. We then performed watershed segmentation to separate the identified “blobs” in the image. The watershed transformation treats the image it operates on like a topographic map, with the brightness of each pixel representing its height, and finds the lines that run along the tops of ridges.¹³ Finally, a box was automatically placed by the segmentation algorithm to highlight the identified GS glomerulus.

Performance Metrics

Datasets for training and testing were divided such that none of the images belonging to patients in the testing




Figure 3. Whitening transformation used for data augmentation. On each cropped image with a single glomerulus (a,b) or a non-glomerular aspect (c,d) of the kidney biopsy, approximately 20% of the pixels were randomly selected, and a whitening transform was applied. This process generated images that still contained a major portion of the original content that represented either the glomerular or nonglomerular aspects of the kidney biopsy (e–h).

set were available to the model while training. This implies that the training-testing split was done at the patient level as opposed to the image level, and the variability in the test images was not part of the images that the model was trained on. This strategy allowed us to systematically evaluate the model performance on completely new patient data. The random splitting of the train and test data was performed 4 different times and model performances on the test data were averaged across all the runs. The CNN model developed to perform multilabel classification was evaluated by computing overall mean accuracy and mean Cohen's kappa (κ). The segmentation model performance was




Figure 4. Glomerular segmentation pipeline. The trained convolutional neural network (CNN) model was used in conjunction with the sliding window operator to scan a test image (a) that was not used in model training. (b) A heatmap was generated based on how the CNN model detected the presence of globally sclerosed (GS) glomeruli. (c) An Otsu binarization operation was attempted on the heatmap, followed by a distance transform (d) and then watershed segmentation (e), which resulted in segmentation of 2 distinct GS glomeruli (f).

evaluated by computing overall accuracy, sensitivity, and specificity on the test data. We also computed F1-score as a measure of model accuracy that considers both the precision and recall of a test. We also

computed Matthews correlation coefficient, which is a balanced measure of quality for dataset classes of different sizes of a binary classifier.

RESULTS

Patient Population

We examined a patient cohort representative of Boston's inner-city population comprising 46% African American population. Approximately 60% of the patients were male, approximately 84% had hypertension, approximately 75% cardiovascular disease, and approximately 43% had diabetes. Approximately 82% of patients had chronic kidney disease stage 3 to 5; 6% had stage 2 chronic kidney disease, and the rest had stage 1 chronic kidney disease. Approximately 35% of patients had nephrotic-range proteinuria (>3.5 g/d). On the basis of varied genetic background and several comorbidities (described previously), it is worth noting that the dataset that we generated provides a wide range of glomerular morphologies, including cases containing normal glomeruli as well as those that manifest partial or GS.

Glomerular Classification Model

The sliding window operator with a small stride (20 pixels) generated a large number of cropped images, and histogram-based thresholding selected many of the images containing the kidney tissue from the images that contained only the background (Supplementary Figure S1). This thresholding method was time-efficient and it was able to filter more than half of the cropped images. Using the selected data, binary classification models, constructed by fine-tuning a well-known pretrained CNN architecture (Inception V3),¹¹ identified images with a NPS or GS glomerulus with high accuracy across 4 different train-test splits (Table 2). Combining random whitening and data augmentation strategies resulted in better CNN model performance on the testing data as exemplified by model accuracy and kappa (Table 2). The CNN model accuracy on the test data across the 4 training-test splits ranged from 89.66% to 95.06%. Also, kappa for these cases ranged from 0.8079 to 0.9111.

Table 2. Convolutional neural network model performance

Whitening factor	Augmentation factor	Accuracy (%)	Cohen's Kappa (κ)
0	0	80.51 \pm 3.01	0.6569 \pm 0.0454
5	0	90.27 \pm 1.62	0.8238 \pm 0.0293
5	10	92.67 \pm 2.02	0.8681 \pm 0.0392

Three different models were developed to understand the effect of random whitening as well as other data augmentation strategies on the convolutional neural network model performance. Model performance is shown on test data that was not used for model training.

Glomerular Segmentation Model

The glomerular classification model that generated the best classification performance (Accuracy = 95.06%), was used for segmenting the glomeruli on the test images. The image processing routine involving image binarization, distance transform, and the watershed method to segment the GS glomeruli performed well on the test data. Specifically, when the trained CNN model was processed in the form of a sliding window operation on the original images, the model scanned through each region of the image and detected several $300 \times 300 \times 3$ pixel windows as nonglomerular regions (model specificity = 0.999). The image processing routine also was able to identify and mark the GS glomeruli on different test images (model sensitivity = 0.558, F1-score = 0.623, Matthew correlation coefficient = 0.628).

DISCUSSION

Deep learning algorithms are transforming medicine especially the way by which images and other forms of data are analyzed to uncover interesting patterns and facilitate clinical diagnosis and management of patients.¹⁴ This is especially the case in the field of digital pathology whereby we and others are using these powerful techniques to address specific questions in a spectrum of disease scenarios.^{9,15-19} For many of these cases, the clinical workflow is quite similar, that is, a biopsy procedure is performed to extract a tiny portion of the organ, which is then subjected to a series of histological staining processes before the evaluation by a pathologist. In facilities that are equipped with biopsy slide scanners, the tissue slides are digitized to generate images, which then serve as the input data of interest. The deep learning algorithms can read and process these digital signatures to extract relevant quantitative information or associate them with corresponding outputs of interest. Once trained on a sufficient number of cases, these models can have the ability to predict on new test cases that the models have never seen before with remarkable accuracy. Currently, the process of biopsy digitization is not performed at all the centers and thus is not integrated within all clinical practices. However, there is a growing interest in this direction, as the medical community at large is realizing the enormous potential of such a resource. In parallel, efforts should be used on analyzing these images with cutting-edge techniques to extract maximum benefit for the management of patients. Eventually, this approach is anticipated to complement the pathologists to improve their accuracy and workflow.

Assessment of renal pathology slides has several features worthy of consideration. The biopsy report methodically deals with all the components of the slide

with different staining, and clinical correlation is pursued to eventually arrive at a diagnosis. Some of the features (or descriptors) are objective (number of glomeruli), whereas others are descriptive (type of sclerosis). Although the latter item calls on the expertise of a pathologist, the former can be easily automated. Also, features, such as the location (cellular or compartmental) and distribution (focal or diffuse or segmental or global) of the intraglomerular damage, determine the type of glomerular disease (glomerulonephritis vs. glomerulosclerosis) can be automated.^{10,20–29} Availability of digitized images provides an immense resource that has opened up opportunities to leverage different tools to improve the analysis of such descriptors that characterize the disease.

Our CNN model identified the presence of an NPS or a GS glomerulus with high accuracy even when a small number of cropped images were used for training ($n = 1496$) (Table 2). It was possible to overcome this limitation with the transfer learning approach. Also, the process of introducing additional noise/variability using random whitening and other data augmentation strategies has shown to limit model overfitting and increase model generalizability (Table 2). Our image analysis pipeline processed the heatmaps to complete the segmentation process that resulted in output images with highlighted areas of the selected GS glomeruli. Note that although simple binarization involves thresholding an image based on a preselected value, Otsu's binarization assumes that an image contains 2 classes of pixels, and then searches for a threshold that minimizes intraclass variance between the classes.¹² The watershed transformation treats the image it operates on like a topographic map, with the brightness of each point representing its height, and finds the lines that run along the tops of ridges.¹³ Ultimately, the nature and locations of the glomeruli dictated the performance of the image processing pipeline.

Although the current body of work differentiates itself from previous studies,^{30–33} it is also complementary to them. Integration of these types of studies is needed to generate a comprehensive platform of analyzing digitized renal biopsies using ML techniques. For example, Bukowy *et al.*³⁰ used an ML technique on approximately 28,000 mouse glomeruli using the periodic acid-Schiff stain. Although the periodic acid-Schiff stain is used for the assessment of basement membrane, the trichrome stain is suited to evaluate fibrosis and sclerosis. Although the use of mouse kidney biopsies facilitates analysis on a large number of glomeruli, it is imperative to examine and validate ML-based findings on human glomeruli for meaningful translation. Although the latter is ideal, there may be limitation on the number of samples.

Methods such as transfer learning could be used to analyze such small datasets without compromising significantly on the performance of the model. Last, the cohort used for the current work is ethnically and racially diverse, which adds to the significance of these findings.

Our study has the following limitation. Our model was able to discriminate GS glomeruli from NPS glomeruli. Although discriminating normal from partially sclerosed glomeruli also can be helpful, a routinely assessed pathologic finding is the percentage of global glomerulosclerosis,³² and we believe our current model can estimate this value with reasonable accuracy on a digitized kidney biopsy image. In the future, we plan to develop a model that can distinguish between normal and partially sclerosed glomeruli. This task can be accomplished successfully when we have access to a large collection of NPS and GS glomerular images.

Our paradigm for identifying and segmenting glomeruli is likely to be useful for the processing of the pathology slides and can be easily extended to images generated using other staining protocols. Adoption of such methods without disturbing the workflow of a pathologist can expedite the assessment of slides and serve as a first step toward more comprehensive, automated analysis. Further validation of the deep learning framework along with the image processing operations across different clinical practices and image datasets is necessary to validate this technique across the full distribution and spectrum of lesions encountered in a typical nephropathology service.

DISCLOSURE

All the authors declared no competing interests.

ACKNOWLEDGMENTS

This project was supported in part by the National Center for Advancing Translational Sciences, National Institutes of Health, through BU-CTSI Grant 1UL1TR001430, a Scientist Development Grant (17SDG33670323) from the American Heart Association, Boston University's Digital Health Initiative Research Incubation Award (#2018-02-008) to Vijaya B. Kolachalama, Boston University's Undergraduate Research Opportunities Program funding to Laura A. Morgan, and National Institutes of Health grants (R01-HL132325 and R01-CA175382) to Vipul C. Chitalia.

SUPPLEMENTARY MATERIAL

Appendix S1. Glomerular dataset generation, model training, and data augmentation.

Figure S1. Histogram-based thresholding. A sliding window operator scanned the entire original image of size $2560 \times 1920 \times 3$ pixels and generated cropped images of

size $300 \times 300 \times 3$ pixels. For each cropped image, a histogram based on pixel intensity was generated (a1 for a cropped image representing the background and a2 representing a portion of the kidney biopsy). These histograms were then reordered according to the bin frequency. A threshold value of 100 was empirically selected as a cutoff and median value for the bin frequency was computed. Images with a median value below the cutoff were selected as the ones representing the background (b1), and the ones with a median value above the cutoff were selected as part of the kidney biopsy (b2). Supplementary material is linked to the online version of the paper at www.kireports.org.

REFERENCES

- Greenberg A, Cheung AK. *Primer on kidney diseases*. 5th ed. Philadelphia, PA: Saunders/Elsevier : National Kidney Foundation; 2009. xvii, 594 p.
- Lees GE, Cianciolo RE, Clubb FJ Jr. Renal biopsy and pathologic evaluation of glomerular disease. *Top Companion Anim Med*. 2011;26:143–153.
- Rayat CS, Joshi K, Dey P, et al. Glomerular morphometry in biopsy evaluation of minimal change disease, membranous glomerulonephritis, thin basement membrane disease and Alport's syndrome. *Anal Quant Cytol Histol*. 2007;29:173–182.
- Puelles VG, Bertram JF. Counting glomeruli and podocytes: rationale and methodologies. *Curr Opin Nephrol Hypertens*. 2015;24:224–230.
- Basgen JM, Steffes MW, Stillman AE, Mauer SM. Estimating glomerular number in situ using magnetic resonance imaging and biopsy. *Kidney Int*. 1994;45:1668–1672.
- Bertram JF. Estimating glomerular number: why we do it and how. *Clin Exp Pharmacol Physiol*. 2013;40:785–788.
- Cullen-McEwen LA, Armitage JA, Nyengaard JR, et al. A design-based method for estimating glomerular number in the developing kidney. *Am J Physiol Renal Physiol*. 2011;300:F1448–F1453.
- LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015;521:436–444.
- Kolachalama VB, Singh P, Lin CQ, et al. Association of pathological fibrosis with renal survival using deep neural networks. *Kidney Int Rep*. 2018;3:464–475.
- D'Agati VD, Fogo AB, Bruijn JA, Jennette JC. Pathologic classification of focal segmental glomerulosclerosis: a working proposal. *Am J Kidney Dis*. 2004;43:368–382.
- Szegedy C, Vanhoucke V, Ioffe S, et al. Rethinking the inception architecture for computer vision. Cornell University; 2015. Available at: <https://arxiv.org/abs/1512.00567>. Accessed August 1, 2018.
- Otsu N. A threshold selection method from gray-level histograms. *IEEE Trans Syst Man Cybern*. 1979;9:62–66.
- Beucher S. Watershed, hierarchical segmentation and waterfall algorithm. *Mathematical Morphology and Its Applications to Image Processing*. 1994;2:69–76.
- Miotto R, Wang F, Wang S, et al. Deep learning for healthcare: review, opportunities and challenges. *Brief Bioinform*. 2018;19:1236–1246.
- Ertosun MG, Rubin DL. Automated grading of gliomas using deep learning in digital pathology images: a modular approach with ensemble of convolutional neural networks. *AMIA Annu Symp Proc*. 2015;2015:1899–1908.
- Janowczyk A, Madabhushi A. Deep learning for digital pathology image analysis: a comprehensive tutorial with selected use cases. *J Pathol Inform*. 2016;7:29.
- Qiu JX, Yoon H-J, Fearn PA, Tourassi GD. Deep learning for automated extraction of primary sites from cancer pathology reports. *IEEE J Biomed Health Inform*. 2018;22:244–251.
- Saltz J, Gupta R, Hou L, et al. Spatial organization and molecular correlation of tumor-infiltrating lymphocytes using deep learning on pathology images. *Cell Rep*. 2018;23:181–193.e7.
- Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*. 2017;542:115–118.
- Brill G, Mendelow H. Inter-capillary glomerulosclerosis; a clinicopathologic study. *J Mt Sinai Hosp N Y*. 1956;23:663–670.
- Conlon PJ, Butterly D, Albers F, et al. Clinical and pathologic features of familial focal segmental glomerulosclerosis. *Am J Kidney Dis*. 1995;26:34–40.
- D'Agati V. Pathologic classification of focal segmental glomerulosclerosis. *Semin Nephrol*. 2003;23:117–134.
- Gupta R, Sharma A, Mahanta PJ, et al. Focal and segmental glomerulosclerosis in renal allograft recipients: a clinicopathologic study of 37 cases. *Saudi J Kidney Dis Transpl*. 2013;24:8–14.
- Iskandar SS, Falk RJ, Jennette JC. Clinical and pathologic features of fibrillary glomerulonephritis. *Kidney Int*. 1992;42:1401–1407.
- Robbins SL, Rogers J, Wollenman OJ Jr. Inter-capillary glomerulosclerosis a clinical and pathologic study. III. A pathologic study of 100 cases. *Am J Med*. 1952;12:700–705.
- Rogers J, Robbins SL. Inter-capillary glomerulosclerosis: a clinical and pathologic study. I. Specificity of the clinical syndrome. *Am J Med*. 1952;12:688–691.
- Rogers J, Robbins SL, Jeghers H. Inter-capillary glomerulosclerosis: a clinical and pathologic study. II. A clinical study of 100 anatomically proven cases. *Am J Med*. 1952;12:692–699.
- Stokes MB, Valeri AM, Markowitz GS, D'Agati VD. Cellular focal segmental glomerulosclerosis: clinical and pathologic features. *Kidney Int*. 2006;70:1783–1792.
- Thomas DB, Franceschini N, Hogan SL, et al. Clinical and pathologic characteristics of focal segmental glomerulosclerosis pathologic variants. *Kidney Int*. 2006;69:920–926.
- Bukowy JD, Dayton A, Cloutier D, et al. Region-based convolutional neural nets for localization of glomeruli in trichrome-stained whole kidney sections. *J Am Soc Nephrol*. 2018;29:2081–2088.
- Kato T, Relator R, Ngouv H, et al. Segmental HOG: new descriptor for glomerulus detection in kidney microscopy image. *BMC Bioinformatics*. 2015;16:316.
- Marsh JN, Matlock MK, Kudose S, et al. Deep learning global glomerulosclerosis in transplant kidney frozen sections. *IEEE Trans Med Imaging*. 2018;37:2718–2728.
- Simon O, Yacoub R, Jain S, et al. Multi-radial LBP features as a tool for rapid glomerular detection and assessment in whole slide histopathology images. *Sci Rep*. 2018;8:2032.