PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Automatic glomerulus detection in renal histological images

Moreira Cardozo Rehem, Jonathan, Luís Conrado dos Santos, Washington, Duarte, Angelo Amâncio, de Oliveira, Luciano Rebouças, Angelo, Michele Fúlvia

Jonathan Moreira Cardozo Rehem, Washington Luís Conrado dos Santos, Angelo Amâncio Duarte, Luciano Rebouças de Oliveira, Michele Fúlvia Angelo, "Automatic glomerulus detection in renal histological images," Proc. SPIE 11603, Medical Imaging 2021: Digital Pathology, 116030K (15 February 2021); doi: 10.1117/12.2582201

SPIE.

Event: SPIE Medical Imaging, 2021, Online Only

Automatic Glomerulus Detection in Renal Histological Images

Jonathan Moreira Cardozo Rehem[a], Washington Luís Conrado dos Santos[b], Angelo Amâncio Duarte[a], Luciano Rebouças de Oliveira[c] and Michele Fúlvia Angelo[a]
[a]Universidade Estadual de Feira de Santana (UEFS), Feira de Santana, Bahia, Brazil; [b]Centro de Pesquisas Gonçalo Muniz da Fundação Oswaldo Cruz (CpqGM/FIOCRUZ), Salvador, Brazil; [c]IVISION Lab, Universidade Federal da Bahia (UFBA), Salvador, Brazil

ABSTRACT

Glomeruli are microscopic structures of the kidney affected in many renal diseases. The diagnosis of these diseases depends on the study by a pathologist of each glomerulus sampled by renal biopsy. To help pathologists with the image analysis, we propose a glomerulus detection method on renal histological images. For that, we evaluated two state-of-the-art deep-learning techniques: single shot multibox detector with Inception V2 (SI2) and faster region-based convolutional neural network with Inception V2 (FRI2). As a result, we reached: 0.88 of mAP and 0.94 of F1-score, when using SI2, and 0.87 of mAP and 0.97 of F1-score, when using FRI2. On average, to process each image, FRI2 required 30.91s, while SI2 just 0.79s. In our experiments, we found that SI2 model is the best detection method for our task since it is 64% faster in the training stage and 98% faster to detect the glomeruli in each image.

Keywords: Glomerulus, Object Detection, Deep Learning, Computational Pathology.

1. INTRODUCTION

Glomeruli are microscopic renal structures that perform blood filtration. Glomerulopathies are diseases that affect the glomeruli and can lead the patient to depend on permanent treatment, or even death [1]. In 2017, Thomé et al. [2] indicate that 126,583 patients depended on permanent dialysis treatment in Brazil. In the United States, Saran et al. [3] indicated that, in 2016, that 678,883 patients were in this same condition.

The diagnosis and therapeutic planning of these diseases depend on the performance of biopsy that makes it possible to determine the type of lesion and its stage [4]. During the biopsy, a tissue fragment is removed from the patient, chemically treated and placed on histological slides that are examined by pathologists with the aid of a microscope [5]. Currently, these slides can be digitized or photographed, making it possible to manipulate them using computers [6]. This evolution has helped in the development of tools that can assist pathologists in the automation of tasks that can shorten the diagnosis, contributing to the improvement in treatment efficiency, in addition to optimizing the use of analysis time by doctors.

Recently, studies focused on the automatic classification of kidney diseases have gained importance, especially those related to the PathoSpotter system [7-10] - a system dedicated to the automatic classification of elementary lesions, based on renal histological images. PathoSpotter is already able to classify some glomerular lesions present in snapshotted images from renal biopsies and depends on images that contain only one glomerulus to perform this task. Currently, PathoSpotter is to classify proliferative glomerulopathies with an accuracy of 0.88 [8-9], segmental glomerulosclerosis with an accuracy of 0.84 for H&E (Hematoxoline and Eosin) stain and 0.81 for PAS (Periodic Acid Schiff) stain [7], and glomerular hypercellularity with F1-score of 1.0 [10].

Some studies are related to the detection of glomeruli in slides [11-12]. Gallego et al. [11] proposed a method for classification and detection of glomeruli in WSI (Whole Slide Image), which are complete scans of a histological slide, which generates a digital image of high resolution. They trained CNN (Convolutional Neural Networks) to be able to identify two classes of objects, glomeruli and non-glomeruli in order to identify segments that contained regions containing glomeruli. They used Transfer Learning [13] from a pre-trained AlexNet network to adapt it to the proposed problem. The method uses a sliding window algorithm where at each step the CNN network is used to classify whether or not there are glomeruli in the selected segment. They also used pre-processing for color normalization, in trying to reduce color variation in inter-laboratory images. The method achieved 0.881 of precision, 1.00 of recall and 0.937 of F1-score.

Medical Imaging 2021: Digital Pathology, edited by John E. Tomaszewski, Aaron D. Ward, Proc. of SPIE Vol. 11603, 116030K · © 2021 SPIE
CCC code: 1605-7422/21/$21 · doi: 10.1117/12.2582201

Simon et al. [12] proposed the automatic detection of glomeruli in WSI using traditional machine learning techniques: visual descriptor Multi Radial Color Local Binary Pattern (mrcLBP) [14] and SVM (Support Vector Machine) classifier [15]. They used their own data set, composed of images of tissues with glomerulopathy and without glomerulopathy, collected from rats, mice and humans, treated with H&E, Jones, PAS and Gomoni Trichrome stains. The process of preparing tissues, image acquisition and image digitalization was standardized. The digitalization method used was scanning, with a fixed and determined scale of approach. They were able to detect glomeruli in healthy human tissues with precision of 0.917, recall of 0.761 and F1-score of 0.932. In human tissues with glomerulopathies, they achieved precision of 0.904, recall of 0.767 and F1-score of 0.831.

In our work, different from the images used by Gallego et al. [11] and Simon et al. [12], the set of images used were obtained through digital photography, with a digital camera coupled to an optical microscope. Images of healthy human tissues and affected by segmental glomerulosclerosis and membranous glomerulopathy were used, captured in different scales of approximation. In some images, the glomerulus occupies a large part of the image, bringing all the details of the texture, while in others, they appear occupying smaller areas, in groups of some glomeruli or a single glomerulus distant, bringing few details of the texture. The process of preparing tissues, image acquisition, and capture has not been standardized, therefore, images bring variations in size, colors, resolution, and textures, which makes automatic detection of glomeruli difficult.

Thus, the objective of this work is to automatically detect glomeruli in digital histological images of low resolution and non-standard acquisition, to improve the classification capabilities of PathoSpotter lesions.

Object detection combines the tasks of classification and location, that is, the method must be able to locate all target objects in the image, each object must be demarcated with boundary boxes and each box must be classified according to type of object it contains [16] [17].

2. METHODOLOGY

The work was implemented using the Python [18] programming language and the Tensorflow [19] and Tensorflow Object Detection (TOD) [20] frameworks. The LabelImg software [21] was used to note the location of each glomerulus in the images. A remote virtual machine, Google Compute Engine with 8 vCPU (Virtual Computer Processing Unit), was used. Its configuration was: 08 vCPU (Virtual Computer Processing Unit) Intel Xeon, 30 gigabytes of RAM (Random Access Memory), 50 gigabytes of SSD (Solid State Drive) and 01 Nvidia Tesla GPU (Graphics Processing Unit), running the Linux Ubuntu operating system, version 16.04.

The initialization of the weights was done through Transfer Learning with the objective of saving time and due to it requiring less samples for the training of the model, since the new training adjusts only the last layers of the CNN. The framework used was TOD and it offers a library of 35 pre-trained models, which use the CNN that won the main object detection competitions, the ILSVRC (ImageNet Large Scale Visual Recognition Challenge) [22], COCO (Common Objects in Context) [23] and PASCAL VOC (Pascal Visual Object Classes) [24]. For reasons of time constraint and resources availability, it would not be possible to test all the models provided; therefore, it was decided to use two models that presented the best results in the main object recognition competitions (ILSVRC 2014, ILSVRC 2015) [22] and COCO 2015 [23]. The chosen models were SI2 (SSD Inception V2 COCO) and FRI2 (Faster RCNN Inception V2 COCO), both pre-trained with the COCO dataset. The SI2 model uses a combination of the Single Shot Multibox Detector (SSD) network [25] with the Inception V2 network [26], while the FRI2 model uses a combination of the Faster Region-Based Convolutional Neural Networks (RCNN) [27] with the Inception V2 network.

The methodology for developing this work was divided into 4 stages: (A) Creation of data sets; (B) Training configuration; (C) Model training and (D) Validation. The description of each step is made in the subtopics bellow.

2.1 Creation of data sets

In this work, the data set was created from set of 909 images provided by the Centro de Pesquisas Gonçalo Muniz da Fundação Oswaldo Cruz (CpqGM/FIOCRUZ). These are images of human tissue with glomeruli: without glomerulopathy, membranous glomerulopathy and segmental glomerulosclerosis. The images used are from tissue stained with H&E and PAS. The method of acquiring the images was through digital photography, using a digital camera coupled to an optical microscope. The average spatial resolution of the images was 0.798 megapixel, the lowest being

0.30 megapixel and the highest resolution being 3.14 megapixel. The images were stored in the formats: JPEG, GIF and TIFF, all in the RGB color model.

From that data set and with the help of a pathologist, the images were selected, annotated and divided to compose 3 subsets: training, validation and final tests.

The training subset consisted of 509 images, being: 209 (41.06%) containing glomeruli without glomerulopathy and 300 (58.93%) containing glomeruli with glomerulopathy (150 (29.46%) with membranous glomerulopathy and 150 (29.46 %) with segmental glomerulosclerosis).

The validation subset aims to evaluate the performance of the models. This subset of images was composed of 200 images, being: 100 (50%) containing glomeruli without glomerulopathy and 100 (50%) containing glomeruli with glomerulopathy (50 (25%) with membranous glomerulopathy and 50 (25%) with segmental glomerulosclerosis).

The subset of tests aimed at the final assessment of the ability to generalize of the models. It consisted of 200 images, of which: 100 (50%) containing glomeruli without glomerulopathy and 100 (50%) glomeruli affected by glomerulopathy (50 (25%) with membranous glomerulopathy and 50 (25%) with segmental glomerulosclerosis).

2.2 Training configuration

After the creation of the image subsets, they were converted to the Tensorflow format, the TFRecord. In this format, all images and the reference boxes (ground truth) are concatenated in a single file in order to optimize the training time.

The process of configuring the training of the models was done through pipeline files provided by the model library of the Tensorflow Object Detection (TOD) framework [20], used in the pre-training. This file serves as a starting point for the new training of the models.

These files were used with minimal changes, necessary to adapt the process to the available hardware and the new task. The number of classes was changed, which was defined for one since the objective is the detection of a single class (glomeruli) and the number of epochs was set at 200,000 for both models, in order to compare the training time between them. This value of 200,000 training epochs is an indication of the TOD documentation [20], which points to this number as being sufficient for the retraining of most models.

2.3 Models training

The training of the models was monitored by Tensorboard, a visualization tool of the Tensorflow framework [19]. During training, at each end of the training cycle, validation tests were performed and plotted on monitoring charts in order to monitor the model's validation performance.

The total training time of the FRI2 model was 54 hours and 19 minutes and it was necessary to interrupt reaching 100,000 epochs, half of the planned training, to have the model with the best performance. The SI2 model took 38 hours and 54 minutes, completing the planned training program. The SI2 model was 64% faster than the FRI2 in the model training phase.

At the end of the training process, a new execution was performed using the same configuration files, however, using the subset of tests in order to measure performance with new images.

2.4 Validation

The TOD framework supports three evaluation protocols that meet the requirements of the main object detection competitions, ILSVRC [22], COCO [23] and PASCAL VOC [24]. The protocol is a standardized set of metrics for each competition and your choice is made from the configuration file. The protocol used in this work was the COCO, since the models used by this work were previously trained using the data set of this competition.

Then, the metrics used to evaluate the models were: mean average precision (mAP), used to evaluate the general performance of the model, average precision (AP) of the size of objects (small, medium and large), and average recall (AR) of the maximum number of objects per image (max = 1, max = 10, max = 100) for each size of an object (small, medium and large).

The mAP is a number from 0 to 1 (normalized), which the highest number is the best, and represents the median of average precision (AP) among all classes that the model is capable of detecting. To determine whether a detection is correct, an overlap threshold between the ground truth box and the detected box is calculated, also called intersection

over union (IoU). Usually, boxes with an IoU below 0.5 are not considered to be correct detections. In some competitions, the mAP is considered by IoU range, for example mAP@0.50 and mAP@0.75 is the mAP with cut-off threshold by IoU of 0.5 and 0.75, respectively.

Although mAP is the standard for object detection competitions, some studies [11-12] present their results using the common metrics in classification works, they are: Precision, Recall and F1-score. For the purpose of presenting the results using the same metrics as the works presented in the literature, we measure these measures.

3. RESULTS AND DISCUSSIONS

After the training process, an evaluation process was carried out using the subset of tests and all data was collected (see Table 1).

Table 1: Performance of Models.

Metrics	Validation		Final Tests	
	SI2	FRI2	SI2	FRI2
mAP	0,88	0,88	0,88	0,87
mAP@.50	0,99	1,00	1,00	1,00
mAP@.75	0,97	0,99	0,97	0,99
mAP^{small}	-1,00	-1,00	-1,00	-1,00
mAP^{medium}	-1,00	-1,00	-1,00	-1,00
mAP^{large}	0,88	0,88	0,88	0,87
AR@1	0,78	0,78	0,77	0,73
AR@10	0,91	0,99	0,91	0,91
AR@100	0,91	0,92	0,91	0,91
$AR@100^{large}$	0,91	0,92	0,91	0,91
$AR@100^{medium}$	-1,00	-1,00	-1,00	-1,00
$AR@100^{small}$	-1,00	-1,00	-1,00	-1,00

From the analysis of the data in Table 1, it can be seen that there are very small differences between the tests done during training and the final tests. The performance analysis of the final tests shows that the overall performance (mAP) is very similar in the two models, with approximately more than 0.01 point of distance between them.

The SI2 model (0.88) is slightly better than the FRI2 (0.87) in mAP in the final tests, and when observing the mAP through the cut of IoU (intersection over union) mAP@.50 and mAP@.75 (proposed border-box overlaps 50% or 75% of the reference box area, respectively), it is noted that the two models have the same performance in the metric mAP@.50 (SI2 - 1.00 and FRI2 - 1.00). In the metric mAP@.75 shows that the FRI2 model is a little better (SI2 - 0.97 and FRI2 - 0.99).

Performance by scale (mAPsmall, mAPmedium, mAPlarge) indicates that there were no small (<32x32 pixels) and medium (>32x32 and <96x96 pixels) objects noted in the data set. In this way, large (> 96x96 pixels) objects brings the same numbers as the general mAP.

When the recall is observed by the number of objects (AR@1 [up to 1 glomerulus per image], AR@10 [up to 10], AR@100 [up to 100]) it is noted that, in images of a single glomerulus, the SI2 (0.77) is better than FRI2 (0.73). In images with up to 10 glomeruli, FRI2 (0.91) presents a result equal to that of SI2 (0.91), it is also noted that both models make fewer mistakes in images with more than one glomerulus. There were no images with more than 10 glomeruli, so

AR@100 reflects the same numbers as AR@10. Likewise, the AR@100 divided by scale (AR@100large, AR@100medium, AR@100small) reflects the same behavior, since there were only images considered large and with up to 10 glomeruli.

Although the results presented by the two models are similar, there is a large discrepancy in the time required to process each image. FRI2 needs an average of 30.91 seconds to process an image, while SI2 only needs 0.79 seconds. Thus, SI2 is 39.1 times faster than FRI2 in processing an image.

In order to present the results using the same metrics of evaluation of the related works found in the literature (Gallego et al. [11] and Simon et al. [12]), the values of recall, precision and F1-score were obtained for each model (see Table 2).

Table 2: Performance of models in classical machine learning metrics.

Models	***Precision***	***Recall***	***F1-score***
SI2	0,99	0,90	0,94
FRI2	0,94	0,99	0,97

Table 2 shows that the SI2 model is better in precision, which indicates a lower occurrence of false positives (when the model indicates boxes where there are no glomeruli). Already recalled, the FRI2 model presents a better result, indicating fewer occurrences of false negatives (when the model does not indicate boxes where, in fact, there were glomeruli). The F1-score metric represents the overall performance of the model and it can be seen that the FRI2 has the best performance.

Table 3 shows the performance of the models, but considering glomerulopathies separately. The Model column indicates the grouping of measures for each of the models, the Condition column indicates the condition of the glomerulus for the data in each row, which can be presented as WG (Without Glomerulopathy), SG (Segmental Glomerulosclerosis) and MG (Membranous Glomerulopathy). The remaining columns show performance data (Precision, Recall and F1-score).

Table 3: Performance of the models by condition of the glomerulus.

Models	**Condition**	***Precision***	***Recall***	***F1-score***
SI2	WG	0,99	0,87	0,92
	GS	1,00	0,96	0,98
	GM	1,00	0,89	0,94
FRI2	WG	0,97	0,99	0,98
	GS	0,87	1,00	0,93
	GM	0,96	1,00	0,98

Analyzing Table 3, it is noted that the SI2 model has more difficulty in detecting glomeruli without glomerulopathy (WG), with F1-score of 0.92 and recall of 0.87, indicating the occurrence of cases in which glomeruli were not detected. The SI2 model achieves a better performance with glomeruli affected by membranous glomerulopathy (MG), with F1-score of 0.98. In contrast, the FRI2 model has its worst performance with glomeruli affected by membranous glomerulopathy, with F1-score of 0.93 and precision of 0.87 indicating the occurrence of cases in which other structures were detected as if they were glomeruli. Its best performance was with glomeruli without glomerulopathy, with F1-score of 0.98.

Direct comparison of the performance achieved by this work with related works is not possible, due to differences in the type of image used by Gallego et al. [11] and Simon et al. [12]. In these works, the authors used the WSI, which has a very high resolution, including the entire histological slide, and in addition, the entire process of tissue preparation,

image acquisition and capture was standardized. In our work, the digitalization method was digital photography, the scale of approach was not fixed, nor was the preparation of the slides standardized. Such characteristics make the nature of data sets very different, which prevents a direct comparison between the works.

In the following figures, it is possible to observe some results in the detection of glomeruli by the SI2 and FRI2 models. In each figure, the image on the left refers to a result obtained by the FRI2 model and the image on the right a result obtained by the SI2 model. In these figures the rectangles with yellow borders represent the reference boundary boxes (ground truth) and the rectangles with green border represent the border boxes detected by the respective model that must contain a glomerulus.

In Figure 1, it is possible to observe the result obtained by both models on an image that contains a glomerulus without glomerulopathy. It is observed, in this example, that the green and yellow rectangles are superimposed, which indicates a high IoU and that the detection reached its objective with excellent quality. It should also be noted that, in this example, the models showed almost identical results.

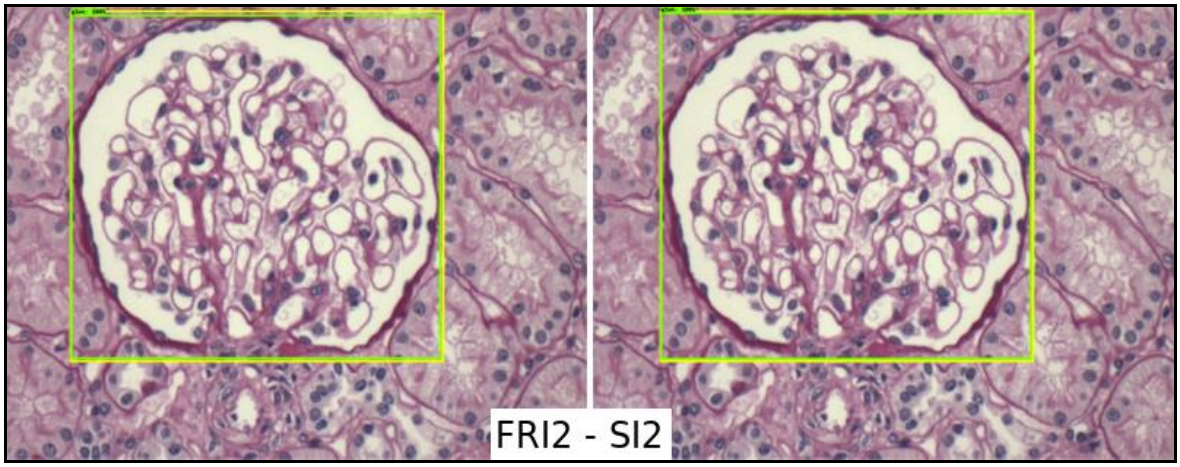

Figure 1 - Result of the detection of a glomerulus by FRI2 and SI2 models.

In Figure 2, the models are able to detect glomeruli even in images of tissues that have been prepared with different stains. Figure 2 (a) shows a glomerulus affected by membranous glomerulopathy. The results were slightly different from each other (the SI2 model detected with a wider box), but both results encompass the entire glomerulus.

Figure 2 (b) shows a glomerulus without glomerulopathy. Note that the image has another type of stain, in addition to bringing a glomerulus on an approximation scale, or a different size. This feature makes it difficult to detect objects and is common in the dataset used in this work.

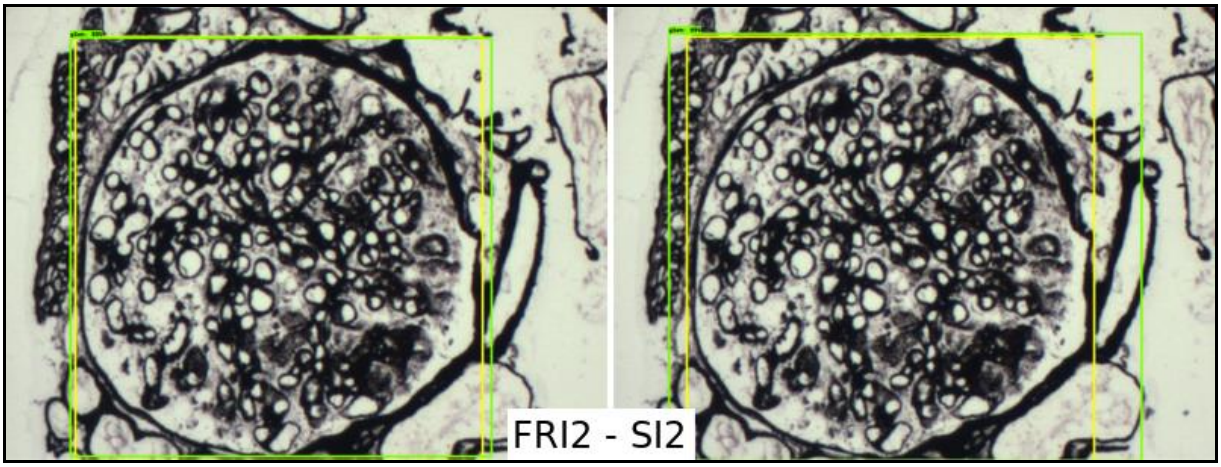

(a)

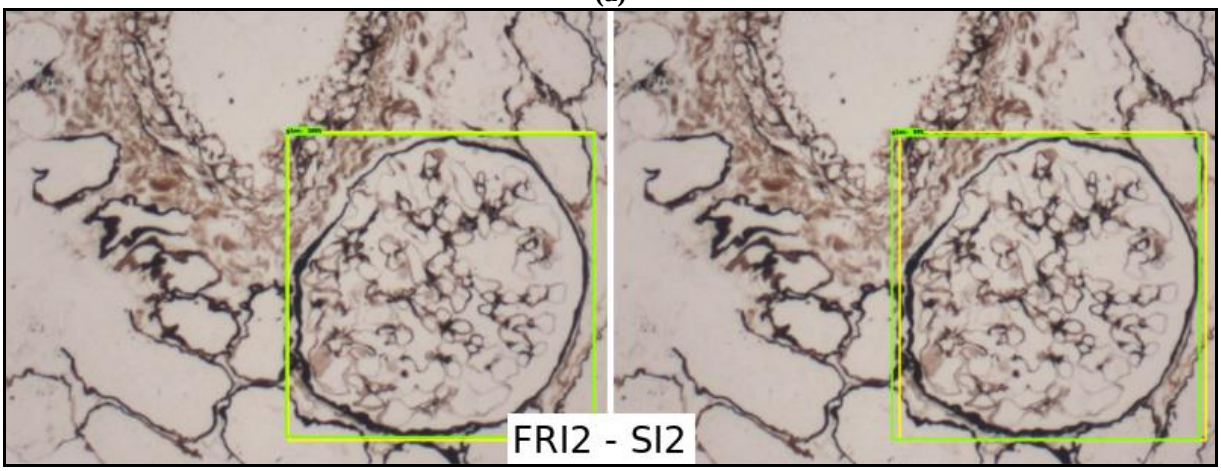

(b)

Figure 2 - Results of the detection of a glomerulus by the FRI2 and SI2 models with stain variation.

Figure 3 shows the results of the models on an image that contains more than one glomerulus, this image contains four glomeruli affected by membranous glomerulopathy.

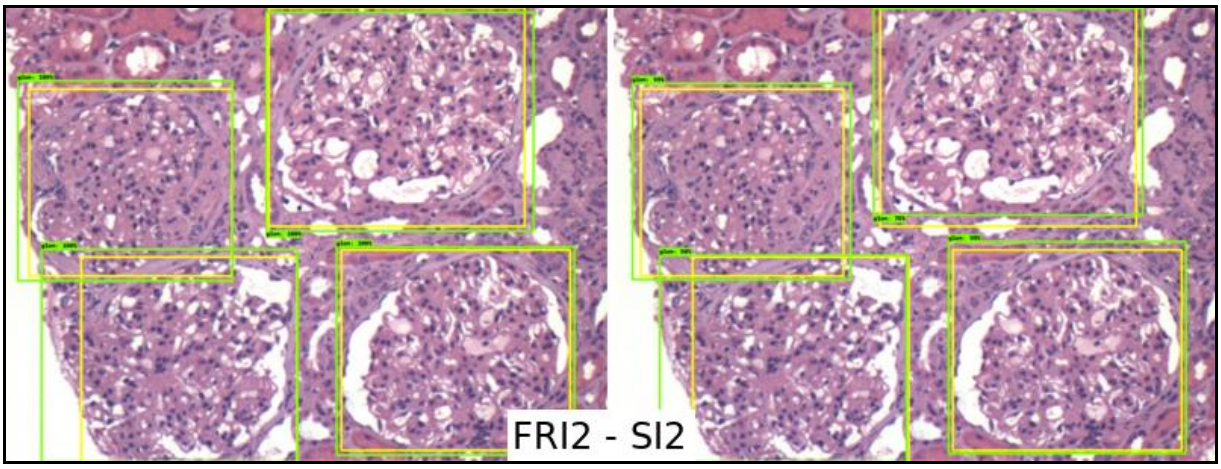

Figure 3 - Result of Detection of four glomeruli by the FRI2 and SI2 models.

Figure 4 shows the results of the models on an image with 3 glomeruli affected by membranous glomerulopathy. In this figure, the difference between the behavior of the models is observed. The SI2 model didn't detect the glomeruli in the peripheral region of the image, a case of false negative.

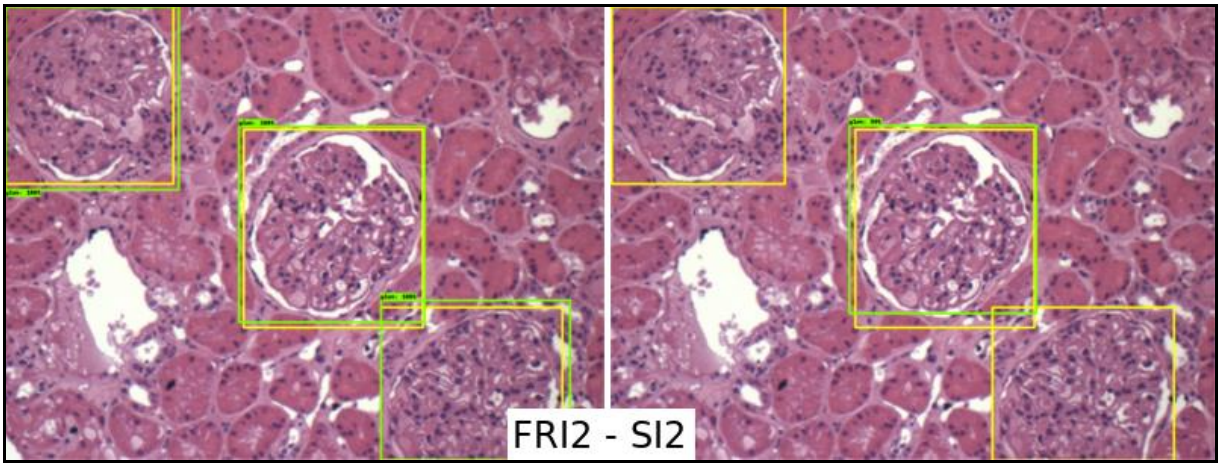

Figure 4 - Differences in detection by FRI2 and SI2 models in image with more than one glomerulus.

Figure 5 shows an example of detecting glomeruli that occupy a smaller relative area than those of the previous images, indicating a smaller approximation scale and showing that the models were still successful. The figure shows glomeruli affected by membranous glomerulopathy.

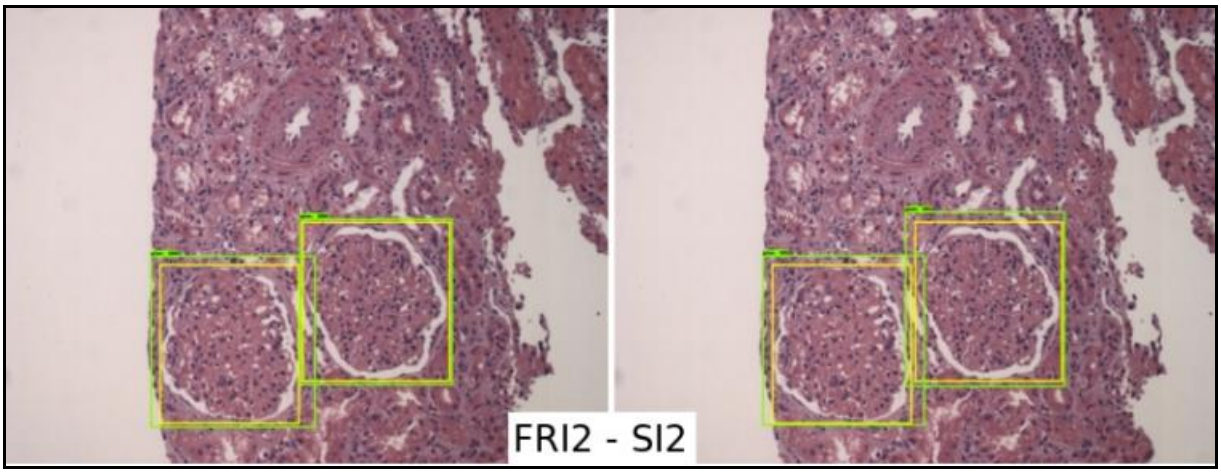

Figure 5 - Result of the detection of glomeruli by the FRI2 and SI2 models in an image of reduce approximation scale.

Figures 6 and 7 show examples of detection in images that contain glomeruli with glomerulopathy and a significant variation in their morphology. Despite this condition, the models are able to detect the glomeruli. Figure 6 contains a glomerulus affected by segmental glomerulosclerosis and Figure 7 contains a glomerulus with membranous glomerulopathy.

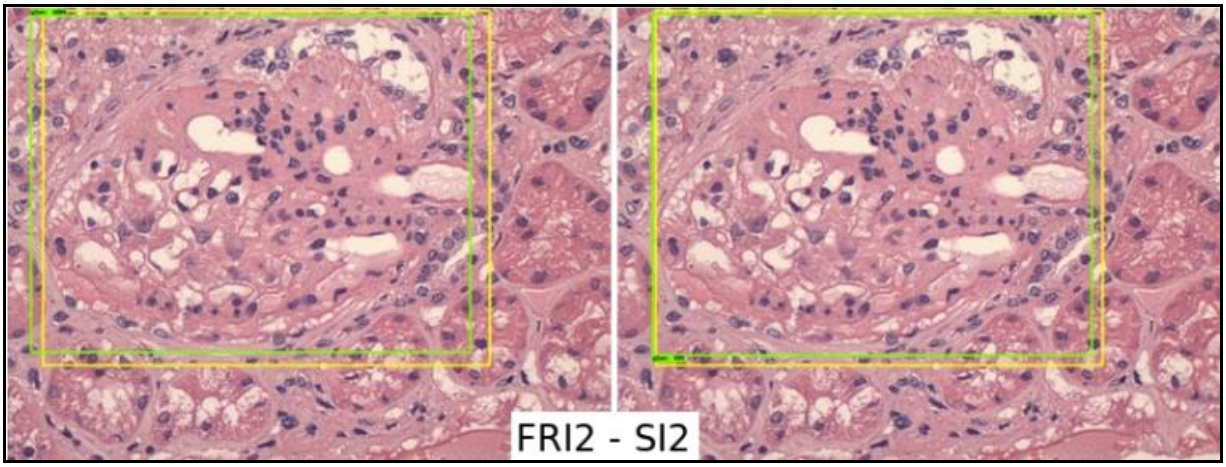

Figure 6: Result of glomerulus detection with segmental glomerulosclerosis by the FRI2 and SI2 models.

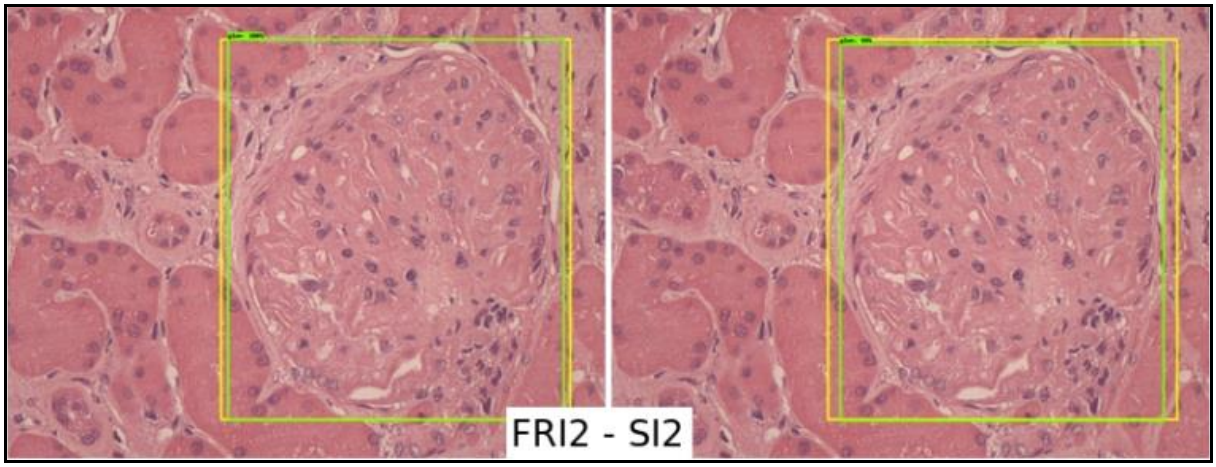

Figure 7: Result of glomerulus detection with membranous glomerulopathy by the FRI2 and SI2 models.

4. CONCLUSION

Using the standard object detection metric, the best result achieved was mAP of 0.88 for the SI2 model, against 0.87 of FRI2, a small difference.

When analyzing the classical machine learning metrics, it was observed that the FRI2 model obtained 0.97 of F1-score, against 0.94 of the SI2 model. Thus, the SI2 was the best model in mAP, while the FRI2 was better in F1-score. It is also noted that the difference in the overall performance of both is small (0.01 of mAP and 0.03 of F1-score). Thus, this alternation can be attributed to the different methodologies of data collection, or even to the random drawing of images for the tests of the models, since it is natural to have a small variation in the numbers for each testing process.

The big difference between the models is the time required for training and processing each image. In training, the SI2 model was 64% faster than the FRI2 model, and in the processing of each image, the SI2 was 98% faster than the FRI2. In view of the great difference in execution and training times, and as time is an important factor to achieve the proposed objective, it is concluded that the use of the SI2 model is the most appropriate.

Finally, we conclude that the results presented demonstrate that the use of deep-learning techniques can be effective in solving the proposed problem, contributing to the development of methodologies that make possible the automatic detection of glomeruli in low resolution renal histological images and without any standardization in their acquisition.

ACKNOWLEDGEMENTS

Authors are grateful to Fundação de Amparo à Pesquisa do Estado da Bahia (FAPESB) grants TO-SUS0031/2018 and TO-PET0008/2015. Washington Santos and Luciano Oliveira have research scholarships from Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq), grants 306779/2017 and 307550/2018-4, respectively.

REFERENCES

[1] D. C. Meldau, "Rim - Anatomia dos Rins - Sistema Excretor," InfoEscola, 2017. [Online]. Available: https://www.infoescola.com/sistema-urinario/rim/. [Accessed: 23-Jul-2019].
[2] F. S. Thomé, R. C. Sesso, A. A. Lopes, J. R. Lugon, and C. T. Martins, "Brazilian chronic dialysis survey 2017," J. Bras. Nefrol., vol. 41, no. 2, pp. 208–214, Jun. 2019.
[3] R. Saran et al., "US Renal Data System 2016 Annual Data Report: Epidemiology of Kidney Disease in the United States," American Journal of Kidney Diseases, vol. 69, no. 3, pp. A7–A8, Mar. 2017.
[4] F. J. V. Veronese, D. D. Morales, E. J. G. Barros, and J. V. Morales, "Síndrome Nefrótica Primária em adultos," Primary Nephrotic Syndrome in adults, 2010.
[5] C. Kelly and J. Landman, "Anatomia do trato urinário," Coleção Netter de ilustrações médicas. 2a ed. Rio de Janeiro: Saunders-Elsevier, pp. 24–7, 2014.
[6] A. D. Belsare and M. M. Mushrif, "Histopathological image analysis using image processing techniques: An overview," Signal & Image Processing, vol. 3, no. 4, p. 23, 2012.
[7] I. C. de Araujo, L. Schnitman, and A. A. Duarte, "Automated Detection of Segmental Glomerulosclerosis in Kidney Histopathology," XIII Brazilian Congress on Computational Intelligence, p. 12, 2017.
[8] G. O. Barros and W. L. Dos-Santos, "PathoSpotter: Um Sistema para Classificação de Glomerulopatias a partir de Imagens Histológicas Renais," SIBGRAPI 2015, 2015.
[9] G. O. Barros, B. Navarro, A. Duarte, and W. L. Dos-Santos, "PathoSpotter-K: A computational tool for the automatic identification of glomerular lesions in histological images of kidneys," Scientific reports, vol. 7, p. 46769, 2017.
[10] P. Chagas, L. Souza, I. Araújo, N. Aldeman, A. Duarte, M. Angelo, W. L. C. dos-Santos and L. Oliveira, "Classification of glomerular hypercellularity using convolutional features and support vector machine," Artificial Intelligence in Medicine, v. 103, p. 101808, 2020.
[11] J. Gallego, A. Pedraza, S. Lopez, G. Steiner, L. Gonzalez, A. Laurinavicius, G. Bueno, "Glomerulus classification and detection based on convolutional neural networks," Journal of Imaging, vol. 4, no. 1, p. 20, 2018.
[12] O. Simon, R. Yacoub, S. Jain, J. E. Tomaszewski, and P. Sarder, "Multi-radial LBP features as a tool for rapid glomerular detection and assessment in whole slide histopathology images" Scientific reports, vol. 8, no. 1, p. 2032, 2018.
[13] L. Y. Pratt and S. Thrun, Machine Learning, vol. 28. Springer, 1997.
[14] G. Schaefer and N. P. Doshi, "Multi-dimensional local binary pattern descriptors for improved texture analysis," in Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012), 2012, pp. 2500–2503.
[15] M. A. Hearst, S. T. Dumais, E. Osuna, J. Platt, and B. Scholkopf, "Support vector machines," IEEE Intelligent Systems and their applications, vol. 13, no. 4, pp. 18–28, 1998.
[16] S. Karagiannakos, "Localization and Object Detection with Deep Learning," 2019. [Online]. Available: https://sergioskar.github.io/Localization_and_Object_Detection/.
[17] R. Parmar, "Detection and Segmentation through ConvNets," Towards Data Science, 2018. [Online]. Available: https://towardsdatascience.com/detection-and-segmentation-through-convnets-47aa42de27ea.
[18] G. V. Rossum, "Python tutorial, Technical Report," Centrum voor Wiskunde en Informatica (CWI), 1995.
[19] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mane, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viegas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," presented at the 12th USENIX Symposium on Operating Systems Design and Implementation, 2015, pp. 265–283.

[20] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, K. Murphy, "Speed/accuracy trade-offs for modern convolutional object detectors," Github, 2017. [Online]. Available: https://github.com/tensorflow/models/tree/master/research/object_detection.
[21] Tzutalin, "LabelImg," Github, 2015. [Online]. Available: https://github.com/tzutalin/labelImg. [Accessed: 25-Jul-2019].
[22] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," International Journal of Computer Vision (IJCV), vol. 115, no. 3, pp. 211–252, 2015.
[23] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. Zitnick, "Microsoft COCO: Common Objects in Context," Lecture Notes in Computer Science, vol. 8693, pp. 740–755, 2014.
[24] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge: A Retrospective," International Journal of Computer Vision (IJCV), vol. 111, pp. 98–136, 2015.
[25] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C. Fu, A. Berg, "SSD: Single Shot MultiBox Detector," arXiv:1512.02325 [cs], vol. 9905, pp. 21–37, 2016.
[26] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, "Going Deeper with Convolutions," arXiv:1409.4842 [cs], Sep. 2014.
[27] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 39, no. 6, pp. 1137–1149, Jun. 2017.