

Rethinking Histology Slide Digitization Workflows for Low-Resource Settings

Talat Zehra^{1*}, Joseph Marino^{2*}, Wendy Wang^{2*}, Grigoriy Frantsuzov^{2*}, and Saad Nadeem²

¹ Jinnah Sindh Medical University, Karachi, Pakistan

² Memorial Sloan Kettering Cancer Center, New York, USA

*Equal Contribution. Correspondence: nadeems@mskcc.org

Abstract. Histology slide digitization is becoming essential for telepathology (remote consultation), knowledge sharing (education), and using the state-of-the-art artificial intelligence algorithms (augmented/automated end-to-end clinical workflows). However, the cumulative costs of digital multi-slide high-speed brightfield scanners, cloud/on-premises storage, and personnel (IT and technicians) make the current slide digitization workflows out-of-reach for limited-resource settings, further widening the health equity gap; even single-slide manual scanning commercial solutions are costly due to hardware requirements (high-resolution cameras, high-spec PC/workstation, and support for only high-end microscopes). In this work, we present a new cloud slide digitization workflow for creating scanner-quality whole-slide images (WSIs) from uploaded low-quality videos, acquired from cheap and inexpensive microscopes with built-in cameras. Specifically, we present a pipeline to create stitched WSIs while automatically deblurring out-of-focus regions, upsampling input 10X images to 40X resolution, and reducing brightness/contrast and light-source illumination variations. We demonstrate the WSI creation efficacy from our workflow on World Health Organization-declared neglected tropical disease, Cutaneous Leishmaniasis (prevalent only in the poorest regions of the world and only diagnosed by sub-specialist dermatopathologists, rare in poor countries), as well as other common pathologies on core biopsies of breast, liver, duodenum, stomach and lymph node. The code and pretrained models will be accessible via our GitHub (<https://github.com/nadeemlab/DeepLIIF>), and the cloud platform will be available at <https://deepliif.org> for uploading microscope videos and downloading/viewing WSIs with shareable links (no sign-in required) for telepathology and knowledge sharing.

Keywords: Slide digitization · stitching · limited-resource settings · whole slide images · neglected diseases · cloud platform.

1 Introduction

Digitizing histology slides is becoming essential for remote consultation (e.g. rural/remote general pathologists seeking opinion from sub-specialist pathologists located in distant urban centers), knowledge sharing (e.g. fast sharing of

rare/regional disease histology presentation with national and international colleagues for education), and assistive/augmented artificial intelligence analysis (e.g. computational algorithm development/testing as well as for external validation of developed algorithms on cohorts from different demographics/settings). These digitization workflows however are extremely costly to setup/operate (with digital multi-slide scanners, scanning space, digital cloud/on-premises storage, and IT personnel/technicians) and hence out-of-reach for limited-resource hospitals in low- and middle-income countries where two-thirds of the population reside with a declining number of pathologists (pathologist:population ratio as low as 1:10million).

To alleviate some of the costs associated with automatic digital multi-slide scanners (costs $> \$70,000$ [3] and contains precise automatic motor stage and illumination control for capturing/stitching image tiles at 20X or 40X magnification to create whole-slide images [WSIs]), manual single-slide stitching software ($> \$3,500$) have been released commercially for offline stitching from manual slide scans at 10X via a microscope camera’s LIVE video stream. However, these single-slide commercial offline solutions (such as Microvisioneer’s mvSlide, Meyer Instrument’s Panoptiq, and Promicra’s PRO-SCAN) only support specific high-end microscopes ($> \$5,000$) with high-resolution digital cameras ($> \$500$) and require expensive high-spec workstations ($> \$1,000$) to operate. The 10X manual scan constraint is due to requiring a large enough spatial context to avoid local/global stitching drift (common at 20X or 40X [6, 10] because of repeating patterns and sparse features) and relatively fast scanning time (1-2 mins for 10X vs 20mins or more for 20X or 40X).

In this paper, we present a first free open-access cloud-based slide digitization workflow to create, view, and share scanner-quality stitched whole-slide images from uploaded 10X manual scan videos, acquired from cheap (\$200) microscopes with built-in digital cameras. We also present a new deep learning algorithm for deblurring out-of-focus regions, upsampling stitched images from 10X to 40X, and reducing stain brightness/contrast and light-source illumination variations. Finally, we demonstrate our WSI creation efficacy on a World Health Organization-declared neglected tropical disease, Cutaneous Leishmaniasis (prevalent only in the poorest regions of the world and diagnosed by subspecialist dermatopathologists, residing in urban centers), as well as on core biopsies from other common pathologies of breast, liver, duodenum, stomach and lymph node.

2 Related Work

There have been prior works [6, 10] exploring whole-slide image creation from videos or image tiles acquired at 20X or 40X but these are focused on automatic whole-slide image creation via a motorized XY stage (e.g. Grundium Ocus20 or Ocus40 automatic single-slide scanners that cost \$14,000) and used video data simulated from scanned images (ignoring all the light-source illumination variations common in real manual scan videos). Manual scanning at 20X or 40X

is not feasible due to the time it takes to complete the scan (15 mins or more depending on the sample size) and the pairwise image registration errors it can accumulate due to the repeated/sparse features with minimal spatial context. In contrast, we create WSIs from 10X with large spatial context and low scanning time while handling blurring and staining/light-source illumination variations as well as recovering 40X details from 10X input.

For deblurring, upsampling, and reducing stain variations, Restore-GAN was recently introduced [11] which used the standard pix2pix architecture with input/output data that again is simulated (added simple Gaussian blur, down-sampling, artificial brightness/contrast variations) from scanned whole-slide images (ignoring light-source illumination and more complex non-Gaussian blurring common in real data). In our work, we use real microscope-derived video data co-registered with scanned whole-slide images to simultaneously deblur real out-of-focus regions, upsample 10X to 40X resolution, and reduce staining brightness/contrast as well as light-source illumination variations in stitched WSIs from manual scan video input.

3 Leishmaniasis and the curated dataset

Leishmaniasis background and histology presentation. The World Health Organization (WHO) has selected cutaneous leishmaniasis as a neglected tropical disease (NTD) with growing, uncontrolled, and ignored infection affecting millions of people every year [14, 9]. It is unfortunately one of the world’s most neglected poverty-related diseases, affecting the poorest people in developing countries and it is associated with risk factors like malnutrition, immune system deficiency, migration, inadequate education, illiteracy, gender inequality, and a shortage of services [1]. Furthermore, only eight countries in the world contribute to 90% of leishmaniasis cases: Afghanistan, Algeria, Brazil, Iran, Pakistan, Peru, Saudi Arabia, and Syria [12]. It is a vector-borne infection caused by the protozoan parasite *Leishmania*. The vector is the female sand-fly. The lesions of cutaneous leishmaniasis vary in presentation ranging from a single self-limited skin lesion to multiple large destructive and ulcerated lesions on mostly the exposed parts of the body such as face, forearms, and lower legs [2]. Skin biopsy is considered the gold standard method for confirmation of diagnosis. The presence of large epithelioid granulomas and small hematoxylinophilic LD bodies (that are round, uniform in appearance, intracytoplasmic and sometimes distributed around the outer rim of the vacuoles) is critical for diagnosis and confirmation of leishmaniasis [4]. Large granulomas are visible at 10X but the smaller hematoxylinophilic bodies (around 3-4mm in size) are usually visible at 40X and can easily be missed if they are few in number, leading to misdiagnosis. Specialized dermatopathologists are the only ones who can diagnose these accurately based on patient history, clinical features, and laboratory diagnosis. It is thus critical to share the digitized slides of these cases from rural poor remote areas with dermatopathologists (mostly residing in urban centers) for quick diagnosis and accurate management of the patient.




Fig. 1. Our pipeline for generating WSI-quality images from video and the corresponding scanned WSI for leishmaniasis cases.

Curated dataset. Skin biopsies from 12 patients diagnosed with cutaneous leishmaniasis (spanning the disease differentiation spectrum) were acquired from Agha Khan University with Internal Review Board approval # 2024-9038-28175. Six of these biopsies, stained with hematoxylin & eosin (H&E), were scanned using Huron’s TissueScope digital scanner at 40X. A cheap-yet-clinical grade

Table 1. Model training parameters/configurations.

Model	Upsampling Layer	Padding Type	Generator Loss	LR	Padding
#1	transposed conv	zero	vanilla + L1	0.0002	zero
#2	resize conv	zero	vanilla + L1	0.0008	zero
#3	transposed conv	reflect	vanilla + L1	0.0008	reflect
#4	resize conv	reflect	vanilla + L1	0.0008	reflect
#5	transposed conv	reflect	vanilla + L1 + VGG	0.0004	reflect
#6	resize conv	reflect	vanilla + L1 + VGG	0.0004	reflect
#7	resize conv	reflect	LSGAN + L1 + VGG	0.0004	reflect

(\$200) microscope (BS-2020MD Digital Microscope from BestScope International Limited, China) with a built-in digital 1.3 megapixel CMOS USB camera and 1W 5-LED illumination with Halogen Lamp 6V/20W was then used to capture videos for all 12 skin biopsies at 10X resolution. To show generalizability of our whole-slide image creation workflow for other common pathologies beyond cutaneous leishmaniasis, we also acquired videos for core biopsies of breast, liver, duodenum, stomach, and lymph node.

4 New Slide Digitization Workflow

4.1 Video creation guidelines and stitching

Pathologists were instructed to record videos using a methodical and consistent movement over the slide, such as starting at one corner and sweeping complete rows, moving down a step and doing the sweep in the opposite direction, all the while pausing for a second between frames with at least 20–30% overlap between adjacent frames. Following these instructions, the videos were recorded at 10X resolution for all the core biopsies, spanning 2–8 mins; we focus on core biopsies since these dominate the bulk of diagnosis work across diseases and these are fastest to sweep with video capture.


For stitching, we automatically extract the static (or “pause”) frames using optical flow. A de-duplication step is run to discard the same frames. With the prior knowledge that these frames are contiguous with at least 20% overlap, we use the open-source Python OpenStitching library (<https://github.com/OpenStitching/stitching>) to then stitch the extracted frames recursively in batches of 30–50 (stitched frames merged into super-stitched frames till a single WSI is generated). The recursive stitching brought down the time for WSI creation to a few minutes (in the order of 5–10 mins) from few hours with the naive stitching approach, which performs image registration between all frame pairs.

4.2 Image-to-image translation

Once the images were stitched, we performed affine co-registration between the 6 corresponding video-stitched WSIs and the scanned WSIs. Two of these co-registered datasets were used to train a conditional generative adversarial network (pix2pix) with the remaining 4 used for testing; the 6 cases without scanned

Table 2. Performance metrics on the test set.

Model	SSIM	PSNR
#1	0.361 ± 0.033	11.078 ± 1.079
#2	0.363 ± 0.034	11.144 ± 1.546
#3	0.362 ± 0.039	11.090 ± 1.445
#4	0.376 ± 0.044	11.619 ± 1.604
#5	0.317 ± 0.054	10.733 ± 1.221
#6	0.319 ± 0.064	10.587 ± 1.431
#7	0.424 ± 0.041	11.167 ± 1.450

**Fig. 2.** Visual comparison of our seven trained models and the original video data. Model #5 gives the best visual results and is used for our final results.

WSIs were used for additional validation and qualitative assessment by an expert pathologist. In total, 1408 image pairs were used for training and 453 for validation. For the generator, we used the standard 9-block ResNet module. We opted for the standard transposed convolution layers after trying checkerboard-artifact-mitigating alternatives such as resize-convolution, which in our use case introduced new artifacts such as gray spots and larger color variation. Reflect padding was used to reduce artifacts at the borders of the generated images and the VGG loss in the generator below dramatically improved the visual quality of the generated tiles. The discriminator was the 4-layer PatchGAN recommended in the original pix2pix paper [5]. When feeding images into the discriminator, low-quality video patches were fed and high-quality patches were the generated condition to the ground truth high-quality scanned patch.

In our application, the vanilla discriminator performed better than LSGAN discriminator for more realistic image generation. For the generator, a combination of mean-squared adversarial loss, VGG loss, and L1 loss was used with all given weights = 1. In line with Deblur-GAN-v2 [7], we also observed that the




Fig. 3. Comparison with other image restoration methods.

best generator was not the one with the best peak signal/noise ratio and structural similarity index measure scores; in fact a different model (Model #5, using reflect padding + VGG, in Table 2) generated the best visual image quality. Ablation study and details of different training configurations are given in Figure 2 and Tables 1 and 2. We implemented the model using PyTorch 1.10.2 framework. Adam optimizer was used to update gradients. The model was trained for 400 epochs with the learning rate set to 0.004 for the first 200 epochs and decays linearly. With batch size 16, the training took 84 hours on 2 A100 (40 GB) GPUs. The average runtime for training was 60 secs per epoch and for inference, it was 0.26 per tile. For memory footprint, training used 32.5 GB GPU memory while inference used 1.9 GB GPU memory.

5 Results and Discussion

We have compared our approach against three image restoration approaches: DeblurGAN-v2 [7], SRGAN [8], and Real-ESRGAN [15]. DeblurGAN corrects for Gaussian and motion blur (similar to Restore-GAN) where as SRGAN and Real-ESRGAN upsample images to 4X similar to what we are doing in this work (10X to 40X). As shown in the comparisons below in Figure 3, our model achieved better performance against these methods when run on real microscope video imagery. In particular, DeblurGAN and SRGAN show little to no visual improvement to the input images. Real-ESRGAN sharpens the image where there is sufficient contrast, but does not restore or enhance areas of low contrast.

The key difference between these other methods and our method is that we train on image data captured from real low quality videos and corresponding high quality scanned images. These videos contain natural variations in focus and illumination, which cause real world blur and brightness/contrast changes across different frames. Other models are trained on synthetic Gaussian and motion blur, which are not representative of the more complex distortions seen in real world data. Their training data also starts as high quality data, while we




Fig. 4. Annotated granulomas identified in our stitched image from video (top row) and our generated WSI-quality (bottom row) images.




Fig. 5. Results on non-leishmaniasis core biopsies.

train with true lower quality image data that contains all defects found in real world video captured from microscopes.

We demonstrate our WSI creation workflow on Leishmaniasis cases, as shown in Figures 1 and 4. An expert pathologist reviewed the cases and annotated granulomas as well as LD bodies on the stitched images and our generated images in Figure 4. To show generalizability of our approach, we also show results on core biopsies of breast, duodenum, stomach, liver, and lymph nodes (Figure 5).

Limitations. Our current model results in significant artifacts at the boundary of the tissue. We will address these and the square artifacts at the borders more rigorously using bi-direction feature fusion GANs [13]. In this work, we used just one microscope which was the cheapest-yet-clinical grade microscope available on market (in developing regions). In the future, we will compare our stitching results on microscopes from different vendors with different resolution cameras. We will also explore using smartphone cameras to acquire videos rather than microscope cameras. We currently require pathologists to pause between frames (ideally with $> 50\%$ overlap) to get the best-quality results. In the future, we will explore ways to handle stitching even from faster videos with minimal overlaps. We will also extend our stitching platform to immunohistochemistry-stained images where manual cell counting is tedious and more advanced deep learning approaches can play a critical role in more accurate, objective, reproducible, and faster protein expression quantification. Finally, we will develop deep learning algorithms for detecting granulomas and LD bodies specific to Leishmaniasis.

Data use declaration and acknowledgment: This study is not Human Subjects Research because it was a secondary analysis of results from biological specimens that were not collected for the purpose of the current study and for which the samples were fully anonymized. This work was supported by MSK Cancer Center Support Grant/Core Grant (P30 CA008748).

Acknowledgments. This research was funded in part through the NIH/NCI Cancer Center Support Grant P30 CA008748.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ahmad, S., Obaid, M.K., Taimur, M., Shaheen, H., Khan, S.N., Niaz, S., Ali, R., Haleem, S.: Knowledge, attitude, and practices towards cutaneous leishmaniasis in referral cases with cutaneous lesions: A cross-sectional survey in remote districts of southern khyber pakhtunkhwa, pakistan. *Plos one* **17**(5), e0268801 (2022)
2. Bilgic-Temel, A., Murrell, D.F., Uzun, S.: Cutaneous leishmaniasis: a neglected disfiguring disease for women. *International journal of women's dermatology* **5**(3), 158–165 (2019)
3. Farahani, N., Parwani, A.V., Pantanowitz, L.: Whole slide imaging in pathology: advantages, limitations, and emerging perspectives. *Pathology and Laboratory Medicine International* pp. 23–33 (2015)
4. Handler, M.Z., Patel, P.A., Kapila, R., Al-Qubati, Y., Schwartz, R.A.: Cutaneous and mucocutaneous leishmaniasis: Differential diagnosis, diagnosis, histopathology, and management. *Journal of the American Academy of Dermatology* **73**(6), 911–926 (2015)
5. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1125–1134 (2017)
6. Krishna, V., Joshi, A., Vrabac, D., Bulterys, P., Yang, E., Fernandez-Pol, S., Ng, A.Y., Rajpurkar, P.: Gloflow: Whole slide image stitching from video using optical flow and global image alignment. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VIII* 24. pp. 519–528. Springer (2021)
7. Kupyn, O., Martyniuk, T., Wu, J., Wang, Z.: Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 8878–8887 (2019)
8. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4681–4690 (2017)
9. Organization, W.H., et al.: Report of the first meeting of the who diagnostic technical advisory group for neglected tropical diseases: Geneva, switzerland, 30–31 october 2019 (2020)
10. Pellikka, M., Lahtinen, V.: A robust method for image stitching. *Pattern Analysis and Applications* **24**(4), 1847–1858 (2021)

11. Rong, R., Wang, S., Zhang, X., Wen, Z., Cheng, X., Jia, L., Yang, D.M., Xie, Y., Zhan, X., Xiao, G.: Enhanced pathology image quality with restore-generative adversarial network. *The American Journal of Pathology* **193**(4), 404–416 (2023)
12. Sasidharan, S., Saudagar, P.: Leishmaniasis: where are we and where are we heading? *Parasitology research* **120**, 1541–1554 (2021)
13. Sun, K., Chen, Z., Wang, G., Liu, J., Ye, X., Jiang, Y.G.: Bi-directional feature fusion generative adversarial network for ultra-high resolution pathological image virtual re-staining. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3904–3913 (2023)
14. de Vries, H.J., Schallig, H.D.: Cutaneous leishmaniasis: a 2022 updated narrative review into diagnosis and management developments. *American Journal of Clinical Dermatology* **23**(6), 823–840 (2022)
15. Wang, X., Xie, L., Dong, C., Shan, Y.: Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 1905–1914 (2021)