



HAL
open science

Efficient 3D Reconstruction of H&E Whole Slide Images in Melanoma

Janan Arslan, Mehdi Ounissi, Haocheng Luo, Matthieu Lacroix, Pierrick Dupré, Pawan Kumar, Arran Hodgkinson, Sarah Dandou, Romain M Larive, Christine Pignodel, et al.

► **To cite this version:**

Janan Arslan, Mehdi Ounissi, Haocheng Luo, Matthieu Lacroix, Pierrick Dupré, et al.. Efficient 3D Reconstruction of H&E Whole Slide Images in Melanoma. SPIE Medical Imaging 2023, Feb 2023, San Diego, California, United States. hal-03834014v2

HAL Id: hal-03834014

<https://hal.science/hal-03834014v2>

Submitted on 9 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Efficient 3D reconstruction of Whole Slide Images in Melanoma

J. Arslan^a, M. Ounissi^a, H. Luo^a, M. Lacroix^b, P. Dupré^b, P. Kumar^c, A. Hodgkinson^d, S. Dandou^c, R. Larive^b, C. Pignodel^b, L. Le Cam^b, O. Radulescu^c, and D. Racoceanu^a

^aSorbonne Université, Institut du Cerveau - Paris Brain Institute - ICM, CNRS, Inria, Inserm, AP-HP, Hôpital de la Pitié Salpêtrière, Paris, France

^bInstitut de Recherche en Cancérologie de Montpellier (IRCM), INSERM, Université de Montpellier, Institut régional du Cancer de Montpellier, Montpellier, France

^cLaboratory of Pathogen Host Interactions, Université de Montpellier, CNRS, Montpellier, France

^dQuantitative Biology & Medicine, Living Systems Institute, University of Exeter, UK

ABSTRACT

Cutaneous melanoma is an invasive cancer with a worldwide annual death toll of 57,000 (Arnold *et al.*, JAMA Dermatol 2022). In a metastatic state, surgical interventions are not curative and must be coupled with targeted therapy, or immunotherapy. However, resistance appears almost systematically and late-stage prognosis can remain poor. The complexity to eradicate melanoma stems from its plasticity; these cancer cells continually adapt to the tumor microenvironment, which leads to treatment resistance. Our primary assumption is that therapeutic resistance relies in part on a series of non-genetic transitions including changes in the metabolic states of these cancer cells. The 3D spatial distribution of blood vessels that are sources of nutrition and oxygen that drive this metabolic status is an important variable for understanding zoning aspects of this adaptation process. Using Whole Slide Images (WSI) of melanoma tumors from Patient-Derived Xenograft (PDX) mouse models, we build 3D vascular models to help predict and understand the metabolic states of cancer cells within the tumor. Our 3D reconstruction pipeline was based on PDX tumor samples sectioned over 2mm depth and stained with Hematoxylin and Eosin (H&E). The pipeline involves three primary steps, including 2D vessel segmentation using Deep Learning, intensity- and affine-based image registration, and 3D reconstruction using interpolation and 3D rendering (allowing for better interaction with biologists, pathologists, and clinicians). The originality of our computer-assisted pipeline is its capability to (a) deal with sparse data (i.e., not all tissue sections were readily available), and (b) adapt to a multitude of WSI-related challenges (e.g., epistemic uncertainty, extended processing times due to WSI scale, etc.). We posit both our 3D reconstruction pipeline, quantitative results of the major stages of the process, and a detailed illustration of the challenges faced, presenting resolutions to improve the pipeline's efficiency.

Keywords: cutaneous melanoma, whole slide images, 3D reconstruction, vascular reconstruction, personalized medicine

1. INTRODUCTION

Historically a rare cancer, cutaneous melanoma has rapidly evolved to become one of the most fatal forms of cancer today, accounting for approximately 57,000 deaths worldwide annually as of 2020.^{1,2} Early stages of the disease have increased survival rates with surgical interventions. However, treatment of late-stage metastatic melanoma is much more perplexing, relying on targeted therapy, or immunotherapy, for an improved line of defense.³ Despite these advanced therapeutic strategies, the clinical outcome of metastatic-affected patients remains poor because of major resistance that appears nearly systematically.⁴

Further author information: (Send correspondence to Janan Arslan)

J.A.: E-mail: janan.arslan@icm-institute.org, Telephone: +33 (0)6 75 22 39 53

The infiltration of resident host tissues by the cancer cells can stimulate molecular, cellular, and physical changes within the host tissue, creating a microenvironment that is conducive to the survival and proliferation of melanoma.⁵ Furthermore, melanoma cancer cells are plastic cells that have the ability to rewire their metabolism to adapt to changing environmental conditions (e.g., poorly vascularized and hypoxic tumor microenvironment, or as a result of the drug exposure), and thus thrive within their complex microenvironment.⁶

In this study, our primary assumption is that resistance to treatment stems at least in part from a series of non-genetic transitions and changed metabolic states. In a previous mathematical modeling work, we showed that the spatial heterogeneity of the blood vessels, which are sources of nutrition and drug compounds for the tumor, may generate zoning with resistant cells distributed preferentially in highly vascularized regions.⁷ We propose here a pipeline to reconstruct the 3D architecture of blood vessel distribution in a melanoma tumor using artificial intelligence (AI) algorithms as applied to histopathological samples. The AI pipeline includes pre-processing to remove artifacts, inpainting of extraneous regions that do not provide value to the final 3D vascular model, a registration process, an optimization component to expedite processing time, patch-level 2D segmentation, and finally a 3D reconstruction based on rendering and interpolation.

There are some prior works in which 3D reconstruction has been applied to histopathology in non-melanoma cases, with a predominant focus on the image registration component of the 3D pipeline. One of the earliest works includes a publication by Song *et al.* in which they evaluated an unsupervised content classification-based registration on singularly stained and multi-stained images and its impact on 3D reconstruction.⁸ Another study by Lotz *et al.* also focused on image registration for multi-stained tissue sections. Their tile-based method consisted of elastic deformation and the use of the normalized gradient field.⁹ Liang *et al.* developed a framework for 3D vessel reconstruction on liver tissue sections.¹⁰ Their pipeline included image registration (using rigid and non-rigid displacement), vessel segmentation (with color deconvolution), vessel cross-section association/interpolation (based on shape descriptors, spatial similarity, and trajectory smoothness), and volumetric rendering. The study by Paknezhad *et al.* was another registration-focused publication.¹¹ Tissue sections were first cleaned before undergoing a two-part registration process: a rough alignment created using a transformation matrix and a second, multi-scale and SIFT-based alignment focused on the region of interest selected by the user. Finally, the publication by Kiemen *et al.* illustrated a complete 3D reconstruction pipeline for histopathology images, consisting of: a coarse alignment using whole-field rigid body registration and elastic registration for local warping; high-throughput cell detection based on color deconvolution; a DeepLab semantic segmentation and a pre-trained ResNet50 network to segment tissue components; and 3D rendering.¹²

These prior works have described successful histopathology-based and reconstruction-focused pipelines. However, a common challenge encountered within these frameworks was processing time. For example, Song *et al.* noted that, in the absence of parallelization, their computation time was quite extensive (87-936 seconds). Furthermore, their registration performance was impacted when the gap between the tissue sections increased, suggesting their methods might be a challenge to use on sparse datasets.⁸ The proposed works of Lotz *et al.* improved registration performance at an average computational time of 40 seconds for every paired image.⁹ However, at this stipulated performance time, it would be assumed that registration on a stack of histopathological tissue sections could take approximately 1.5-2.5 hours to complete. Kiemen *et al.* also noted a computational time of 1-3 days for their pipeline.¹² Comparatively, the advantages of our proposed pipeline include (i) improving 3D reconstruction performance through the use of pre-processing steps which yield more manageable histopathological images, (ii) the optimization of computation speed using parallel processing (with our registration component only taking 10 minutes to complete), and (iii) the ability to handle sparse data through the use of appropriate rendering and interpolation techniques.

The proposed pipeline has a multi-function purpose: in addition to understanding the characteristics of the vascular network through reconstruction, we also automate previously human-run and manual pathology grading systems that are time-consuming and arduous. The pipeline was designed for Whole Slide Images (WSIs) of Hematoxylin and Eosin (H&E)-stained pathology sections from melanoma tumors obtained from Patient-Derived Xenograft (PDX) mouse models. In this study, we elucidate the steps of our pipeline, provide qualitative and quantitative assessment, and discuss WSI-related challenges (notorious and plenty owing to the size and depth of information in WSIs), and the solutions we have utilized to improve the efficiency and adaptability of our 3D reconstruction. This 3D vascular reconstruction can be used to explicate how characteristics such as vascular

shape, size, and bifurcations differ in alternate metabolic states and thus can be used in predicting treatment efficacy.

2. METHODOLOGY

2.1 3D Reconstruction Pipeline

2.1.1 Overall Pipeline

Fig. 1 provides an overall outline of the 3D reconstruction pipeline. Images first underwent preliminary pre-processing steps, including being (A) exported in Tagged Image File (TIF) Format, (B) split into top and bottom tissues, and (C) crop-centered, inpainted, and cleaned via the removal of the noisy background. Once images were pre-processed they were (D) registered using an optimized intensity- and affine-based registration method. This was followed by patch-level segmentation (E & F) to produce complete, segmented WSIs (G). All segmented WSIs were rendered and interpolated to produce a final 3D vascular model (H & I).

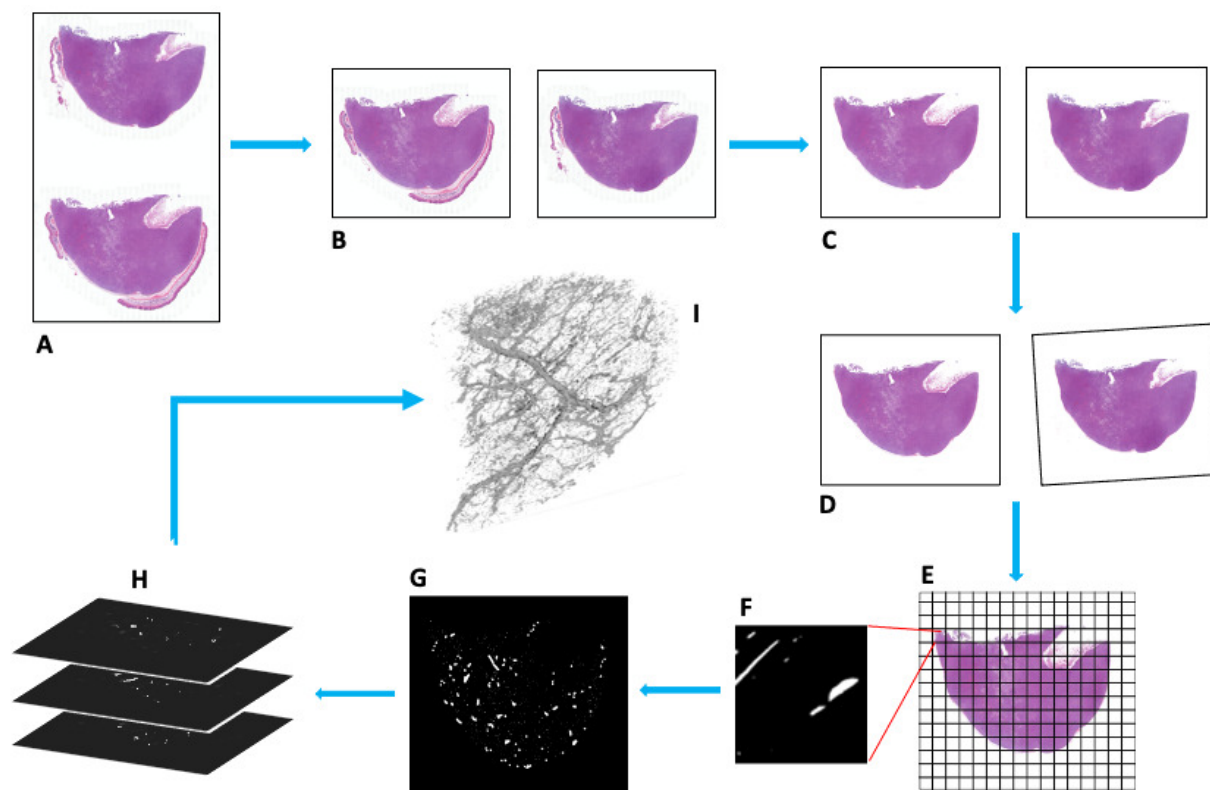


Figure 1: Overall pipeline for 3D vascular reconstruction from Hematoxylin and Eosin-stained Whole Slide Images.

2.1.2 Image Pre-processing

The pre-processing stage is designed to normalize the WSIs. In essence, the following steps are taken to create WSIs that are more uniform and almost mimic the standardized radiological images, such as those of computed tomography (CT) scans. To create this uniformity, several artifacts and variabilities in WSIs need to be accounted for (Fig. 2), with the most predominant concern being tissue shift. WSI acquisition is littered with epistemic uncertainty (i.e., uncertainty due to incomplete knowledge of a disease or process, or reducible errors, such as subjective, measurement, or human-related errors).¹³ This is mostly due to the process being heavily reliant on humans manually carrying out each task of WSI acquisition (i.e., the slicing of the tissue, its addition to a glass slide, etc). To minimize tissue shift, images were crop-centered. Crop-centering in this pipeline includes:

(i) establishing an image size with a width of 6000 and height of 5000 (choice of the size based on the size of the baseline and largest tissue section); (ii) the identification of the central point of the WSI; and (iii) from the identified center, cutting a perimeter based on the pre-set width and height. This process ensures all tissue sections were centrally located while retaining their original tissue size. The maintenance of the original tissue size takes priority. Cutting directly around the perimeter of each tissue would have resulted in small tissue sections appearing erroneously large that would not reflect their true size relative to the whole tissue. Thus, there would have been an inconsistency in the image size that could have impacted rendering and interpolation techniques later in the pipeline (refer to Section 2.1.6). An additional step was required prior to crop-centering. This included the addition of a border with a width of 2000 pixels around each tissue section following its split into top and bottom images. The reasoning for this addition was to account for the disparity in the location of some tissue sections. Some tissue samples were close to the edge of the slide as shown in Fig. 3. Thus, when crop-centering was attempted, the images would not be correctly centered, as there was insufficient space to cut a perimeter of 6000×5000 . The addition of a 2000-pixel border granted additional and sufficient space to ensure the correct and central placement of our tissue samples.

Noisy backgrounds were also removed (Fig. 4). Noisy backgrounds in WSIs consist of staining artifacts and pixelation, with the latter most likely the result of the scanning and digitization of the WSIs. Noisy background removal involved the tissue samples being isolated from their original background and placed on a clean background. This was achieved using a foreground extraction process. The original images were binarized using a simple Otsu thresholding technique, creating a foreground mask. To remove small and erroneous regions, such as staining artifacts surrounding the tissue, the largest contour (i.e., the tissue section) was selected, eliminating any small and unnecessary regions from the binary mask. Using the Bitwise Operation, the cleaned binarized masks were used to identify the region of interest within the original image that we wished to extract, leading to the extraction of foreground information only while eliminating the original noisy background. The final step was to simply add a white color to the background, resulting in a nice, cleaned tissue section.

Finally, the epidermis was inpainted using a whole-image level trained U-Net model (Fig. 5). Initial registration attempts showed the epidermis hindered the performance of the registration and its presence provided little value for the final vascular reconstruction. The epidermis U-Net model had mostly similar hyperparameters to that of our patch-level blood vessel segmentation model (refer to Section 2.1.4) and included: batch size=32, the ADAM optimizer, learning rate= 3×10^{-5} , epoch=300, and the Binary Cross Entropy Dice loss function (BCE DICE). Given the (almost) consistent nature of the epidermis (i.e., shape, location, patterns, etc.), a smaller epoch size proved to be more than adequate for inpainting. The BCE DICE was chosen based on its ability to better discern segmentation boundaries as compared to other loss functions.¹⁴

2.1.3 Image Registration

Histopathological image registration poses several challenges, with feature-based image registration methods limited by their ability to identify distinct points throughout time-series images, while intensity-based methods alone result in a high number of local optima when there is noisy data. In this study, for the registration of the WSI stack, we propose an improved version of the symmetric, intensity-based affine registration framework introduced by Öfverstedt *et al.*¹⁵ More specifically, the authors combined intensity-based registration with either fuzzy point-to-set bidirectional or fuzzy point-to-set inwards distances into an asymmetric average minimal distance. The limitation of this method is that, when compared to traditional similarity measures (i.e., Sum of Squared Differences, Pearson Correlation Coefficient, and Mutual Information), this method requires substantial memory to store auxiliary data structures (e.g., a single, paired-image registration may require up to 4GB of working memory). This results in extensive computation time. For example, the initial version of this algorithm took approximately one week to finalize the registration of our entire WSI dataset, even when executed on a high-performance cluster (HPC).

In order to optimize the execution time, we profiled this approach and deduced the following bottlenecks/flaws:

- Sequential image registration takes longer to finish registering all slides for a given tumor because a given $slide_{s+1}$ requires the registered $slide_s$ (which is the registration results of $slide_{s-1}$ and $slide_s$) to initiate the subsequent registration. Thus, we can not take advantage of the parallelization offered by HPC. This dependency limits the scalability of the algorithm.




Figure 2: Sample image artifacts impacting pipeline flow. Due to the mostly manual acquisition of tissue sections in whole slide image processing, these imaging modalities are highly prone to several artifacts, such as erroneous staining and tissue shifts. There are also pixelations (shown as horizontal bars) that appear around the perimeter of tissue sections.




Figure 3: Demonstration of proper and improper crop-centering. In some cases, tissue sections would appear at the edge of a slide, resulting in a lack of space to cut the pre-set perimeter of 6000×5000 . This often resulted in images that were not centered. To correct this anomaly, an additional step was taken to first add a border size of 2000 pixels before the crop-centering.




Figure 4: Noisy background removal through foreground extraction. This involved the creation of a binary mask using the Otsu thresholding, selecting the largest contour (i.e., the centralized tissue sections, eliminating all smaller artifacts that may be surrounding the tissue), then using this binary mask as a reference to extract the regions of interest from the original image. Finally, a new white background was re-added, leading to a cleaned tissue section.

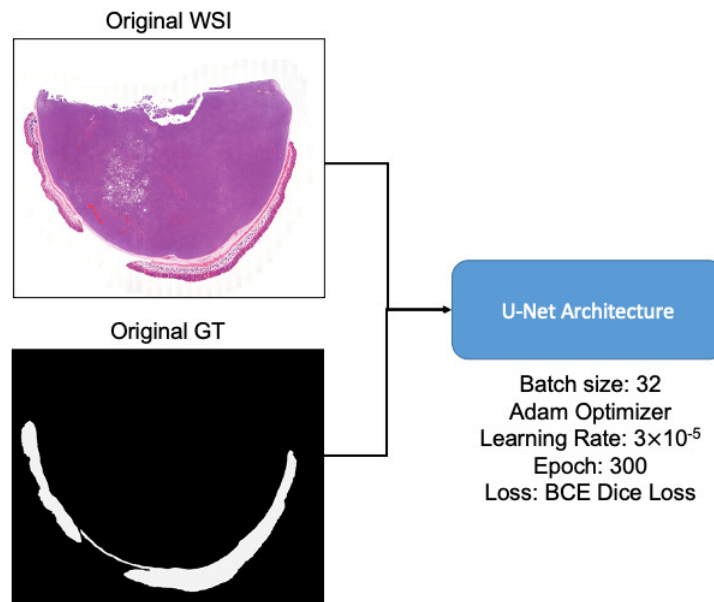


Figure 5: Whole-image level epidermis model trained on the U-Net architecture. To remove the epidermis, which provided little value in terms of vascular reconstruction but caused a hindrance in the registration process, a U-Net model was trained to detect and inpaint the epidermis.

- Registration using the native resolution is unjustifiably costly for such a registration task. It uses a global approach that estimates one affine transformation applicable to all slide pixels. For instance, given a pair of RGB slides with 6000×5000 pixels each, when re-scaling to 90% of the original size, the registration takes 1961.59 sec while yielding a mean absolute error (MAE) of 45.72. However, re-scaling to 10% of the original scale takes only 16.30 sec with a 45.53 error.
- The tumors analyzed in this study are ball-like. Thus, the mid-section slides contain more tissue when compared to the top/bottom slides. This registration approach does not account for this specification since it stretches the slides tissue (affine transformation) to minimize the global error.

We propose a generic optimization scheme for WSI registration (Fig. 6): (i) Choose the down-scaled percentage that does not influence the registration performance, (ii) Register each two consecutive WSI in a separate CPU (job array in case of HPC), and (iii) Apply a correction ratio based on the stretched area and original tissue. In order to quantify the added value from the optimization scheme, we registered 220 consecutive H&E WSI. It takes less than 10 minutes compared to one week previously (without the optimization). Note that the same HPC resources were used on both tests: 1 core/WSI-pair Intel® Xeon® Processor E5-2680 v4 35M Cache, at 2.40 GHz clock rate with 1GB RAM and for the previous approach 16GB RAM was required (bigger resolution). Code for the optimization scheme can be found in Section 7.

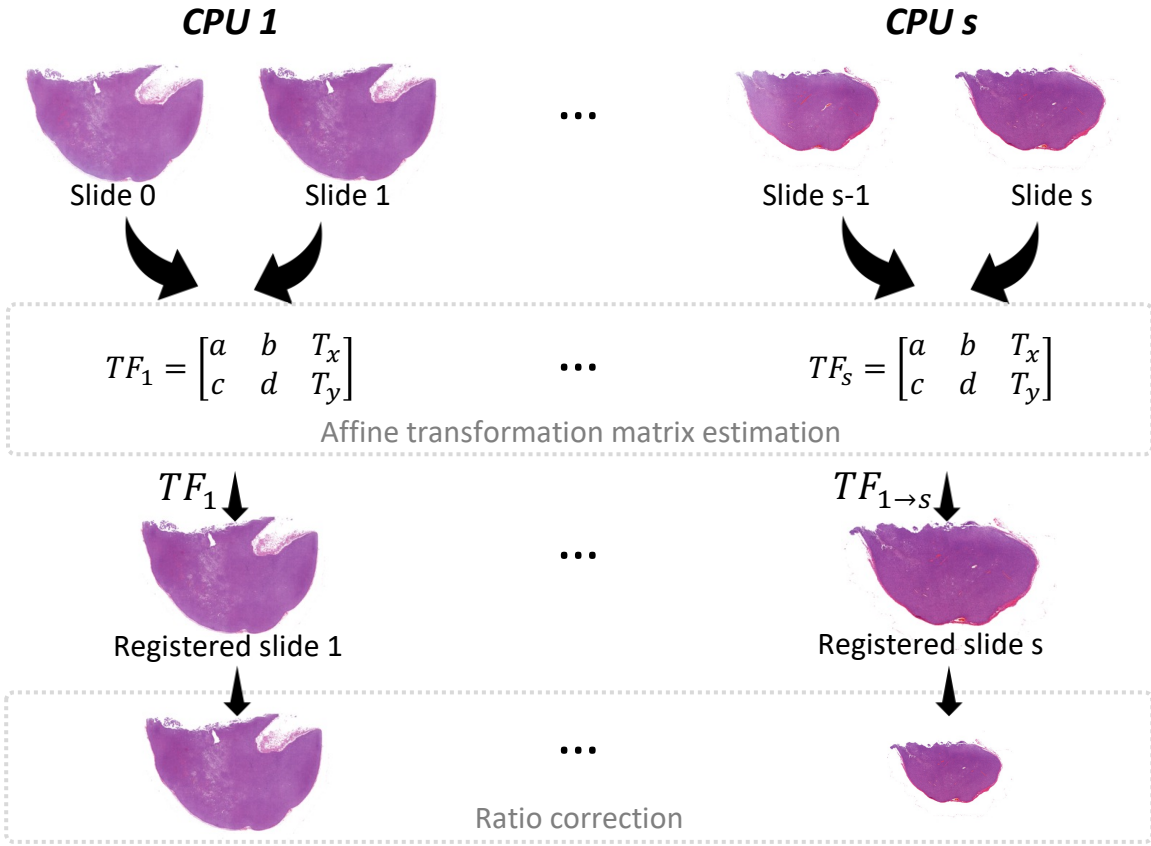


Figure 6: Optimization scheme for Whole Slide Image registration.

2.1.4 2D Blood Vessel Segmentation

Similar to the epidermis inpainting discussed in Section 2.1.2, our blood vessel segmentation model was trained using the U-Net architecture, with the primary differences being (i) the blood vessel model was trained at a patch-level, with patch sizes of 512×512 , and (ii) due to class imbalances and the variety of shapes and locations (even within patches) of blood vessels, the training epoch was extended to 1000 (Fig. 7). Furthermore, data were trained and tested using a 5-fold cross-validation approach, with each fold split into 80% training and 20% test. The final hyperparameters included: batch size=32, the ADAM optimizer, learning rate= 3×10^{-5} , epoch=1000, and the BCE DICE. As stated earlier, the BCE DICE was used for its ability to better discern boundaries. This is particularly imperative in the case of vascular network segmentation. Our hypothesis is based on the evaluation of non-genetic transitions that may impact chemotherapeutic response. Thus, capturing the branching and size of blood vessel networks can impact our understanding of nutrient and oxygen diffusion, so we need the most accurate representation of the vascular network. Further information regarding the ground truth (GT) labels used in the training process can be found in Section 3.1.1.




Figure 7: Patch-level blood vessel segmentation model trained on U-Net architecture. Whole Slide Images and equivalent Ground Truth labels were divided into patches of size 512×512 and fed into the U-Net architecture.

2.1.5 Pipeline Evaluation

The evaluation metrics used to assess the 2D segmentation of the blood vessels included: Accuracy (ACC; $\frac{TP+TN}{TP+TN+FP+FN}$), Dice coefficient (DICE; $\frac{2TP}{2TP+FP+FN}$), Dice Loss (i.e., 1 - DICE; DICE LOSS), BCE DICE (BCE = $y\log(\hat{y})+(1-y)\log(1-\hat{y})$), MAE ($\frac{1}{n}\sum_{i=1}^n |y_i - x_i|$), Precision (PREC; $\frac{TP}{TP+FP}$), recall (REC; $\frac{TP}{TP+FN}$), specificity (SPEC; $\frac{TN}{TN+FP}$), the F_1 -score ($2 \times \frac{PREC \times REC}{PREC+REC}$), and Matthews Correlation Coefficient (MCOR; $\frac{(TP \times TN) - (FP \times FN)}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$). Image registration performance was validated using the Structural Similarity Index Measure (SSIM; $\frac{(2\mu_x\mu_y+c_1)(2\sigma_{xy}+c_2)}{(\mu_x^2+\mu_y^2+c_1)(\sigma_x^2+\sigma_y^2+c_2)}$).

2.1.6 3D Rendering and Interpolation

For a complete 3D reconstruction, a hybrid solution was created using rendering and interpolation, where rendering stacks existing data to create a preliminary vascular volume while the interpolation imputes the gaps between the H&E rendered tissues sections. Rendering was achieved with marching cubes. The shape- and distance-based method proposed by Schenk *et al.* was used for interpolation.¹⁶ Their proposed interpolation involves the following:

- First, binary scenes for the i and $i+1$ images are established. While in their publication these are user-defined contours, in our application the binary scenes involve the segmented and sequentially stacked WSIs.
- Gray-level distance maps are generated for the i and $i+1$ binary scenes, with distance being relative to the binary scene boundaries. Distances within the boundaries are given positive values, while those outside of the boundaries are negative.
- Gray-level distance maps are interpolated using conventional gray-scale interpolation techniques, such as linear interpolation.
- The interpolated gray-level distance maps are converted back into interpolated binary scenes.

2D-level blood vessel model training was based on the tiles extracted at the highest resolution. For the purposes of 3D rendering and interpolation, the whole slides at their highest resolution would be too large to export (requiring inordinate amounts of memory) and execute in a timely manner. Therefore, while the segmentation model was trained on patches at the highest resolution, the 3D reconstruction pipeline included whole tissue sections that were exported with a downsampling factor of 8. The highest resolution has a magnification of $\times 40$, which translates to 0.25 μm per pixel. At downsampling of 8, we have a magnification level of $\times 5$ and a resolution of 2 μm per pixel. Given the training of the U-Net (described in Section 2.1.4) was based on patch sizes of 512×512 at a resolution of 0.25 μm per pixel, the equivalent patch size and resolution at a downsampling of 8 would be 64×64 pixels, and thus this is the patch size we see in part (E) of Fig. 1. As illustrated in Fig. 8, 512×512 patches with an $\times 40$ magnification (A) and 64×64 patches with an $\times 5$ magnification (C) capture the same contextual and spatial information within WSIs. If we were, hypothetically, to run patch-level segmentation at 512×512 patches with an $\times 5$ magnification (B), we would capture a larger spatial zone, which would lead to inefficient blood vessel segmentation, as our model was trained to see blood vessels 'close up' within these patches.

3. DATA

3.1 Data preparation

3.1.1 Raw data

The PDX model is illustrated in Fig. 9a. Cutaneous melanoma tumors were extracted from humans during biopsy, then implanted onto immunodeficient mice for amplification. PDX samples underwent serial sectioning at every 12 μm over a depth of 2mm and were stained with H&E. Each tissue section had a thickness of 4 μm . This sectioning protocol resulted in gaps between each tissue section, which were imputed using interpolation




Figure 8: Patch sizes, magnifications, and resolutions of Whole Slide Images. A patch size of 512×512 at a magnification of $\times 40$ captures the same spatial and contextual information as a patch size of 64×64 at a magnification of $\times 5$.

methods (refer to Section 2.1.6). Slides were digitized into MIRAX format (multi-file with very complicated proprietary metadata and indexes). Each slide contained two tissue sections (i.e., top and bottom) as shown in Fig. 9b; this method was used to expedite data acquisition time and reduce production costs. A total of 120 slides (240 tissues sections) were available. Some of these sections were removed entirely due to severe artifacts that could not be corrected with the pre-processing steps outlined in Section 2.1.2. This led to the removal of 20 tissue sections, leaving a final sample size of 220 tissues. Blood vessels were annotated by an expert pathologist (co-author *CP*) using the open-source software QuPath.¹⁷

3.1.2 Sampling

In preparation for training using deep learning methods, WSIs along with their annotated counterparts were exported as tiles/patches from QuPath. Patches were sized 512×512 (with size selection based on preliminary experimentation with various patch sizes) and were exported at their highest resolution ($0.25 \mu\text{m}$ per pixel). A total of 1000 patches with positive (i.e., presence of blood vessels) and negative (i.e., background/other molecular markers) cases were captured.

4. RESULTS

2D segmentation performs generally well as demonstrated in Table 1, with $\text{DICE} > 80\%$ in training and test results. The 3D reconstruction model illustrated in part (I) of Fig. 1 further validates this as we can qualitatively see the majority of the blood vessels (particularly the larger blood vessels, which are suggested to play an important role in understanding chemotherapeutic response). Registration performance was qualitatively and quantitatively assessed as shown in Table 2 and Fig. 10. We noticed that larger tissue sections had lower SSIMs as compared to smaller sections, despite the alignment being near-perfect visually (Fig. 10). We assume the lower SSIM is a reflection of the greater heterogeneity of blood vessels in the larger sections (i.e., many small blood vessels that do not appear across all large tissue sections). Smaller tissues, however, have larger blood vessels that appear iteratively, and thus demonstrate higher performance given the similarity in features ($\text{SSIM} > 80\%$). The qualitative and quantitative results from this study demonstrate an overall good performance. The final 3D reconstructed model (Fig. 1) demonstrates a clear vascular network consisting of small capillaries and large blood vessels. Quantifying these vascular networks based on shape and spatial distribution can mark a stepping stone in understanding how non-genetic metabolic states could impact chemotherapeutic response in melanoma patients.

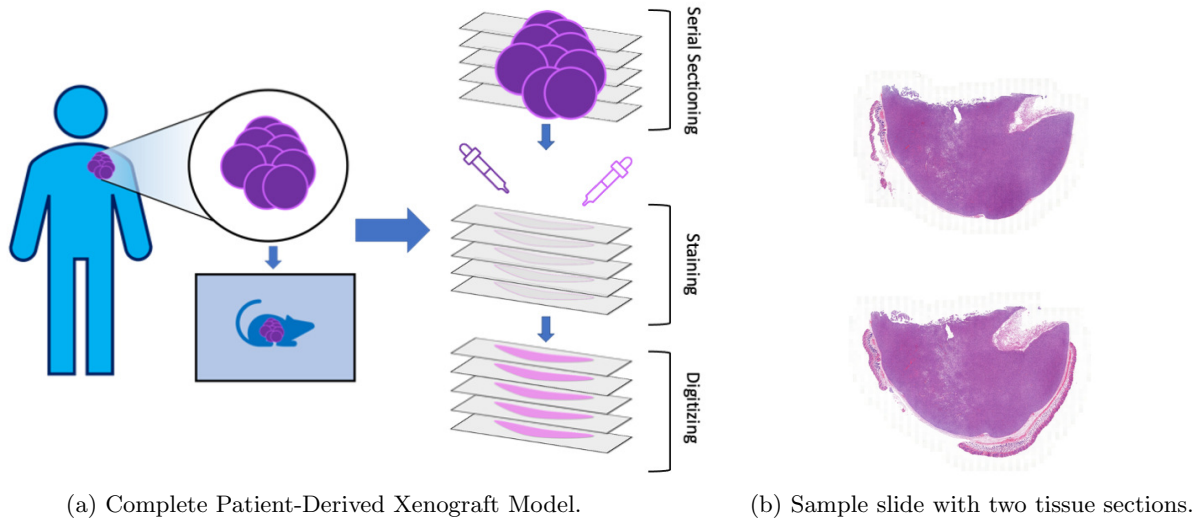


Figure 9: Human tumors were implanted onto immunodeficient mice and grown. Mice-harvested samples then underwent serial sectioning. Slides stained with Hematoxylin and Eosin were digitized into Whole Slide Images.

Table 1: Average training and test results for 2D-level blood vessel segmentation based on U-Net architecture across 5-fold cross-validation.

Metrics	Training		Testing	
	Mean	SD	Mean	SD
BCE DICE	0.081	0.019	0.183	0.083
MAE	0.009	0.001	0.027	0.017
ACC	0.992	0.001	0.973	0.016
DICE	0.886	0.037	0.813	0.086
DICE LOSS	0.114	0.037	0.188	0.085
REC	0.836	0.052	0.767	0.129
PREC	0.842	0.049	0.847	0.082
F_1	0.832	0.054	0.795	0.112
SPEC	0.996	0.001	0.985	0.015
MCOR	0.830	0.053	0.784	0.105

SD = Standard Deviation

Table 2: Quantitative assessment of registered images using the Structural Similarity Index Measure.


Mean	SE	Median	SD	Min.	Max.
0.864	0.0032	0.861	0.048	0.768	0.964

SE = Standard Error

SD = Standard Deviation

5. NOVELTY OF WORK

We present an end-to-end approach for reconstructing a vascular 3D model based on WSIs. An original feature of our pipeline is its ability to handle sparse data, given that not all tissue sections were made readily available. This is achieved through rendering and interpolation. Furthermore, the pipeline has been built to accommodate the multitude of challenges faced with processing WSIs. These include: conducting patch-level training and segmentation; creating uniformity in the available data; and improving registration by using inpainting, crop-centering, and artifact removal techniques to minimize registration-related errors. Other breakthrough components of our



(a) Overlap between large tissue sections following registration. (b) Overlap between small tissue sections following registration.

Figure 10: Qualitative assessment of registered images. While visually near-perfect, variability in the Structural Similarity Index Measure between larger and smaller tissues are attributable to greater heterogeneity of blood vessels in larger tissue sections.

project include the optimization of WSI registration via parallel processing. This optimization has considerably improved the time efficiency of our pipeline, reducing our computational processing time to mere minutes, thus improving the overall efficiency and speed of our reconstruction process.

6. CONCLUSIONS

In this study, we have proven that AI can be used to develop automated and clinically operational pipelines for WSIs. Using a combination of deep learning, image processing, registration, statistical and parallel processing techniques, we can develop 3D models that are adaptable to WSI challenges and produce results which elucidate vascular networks that could predict chemotherapy efficacy in melanoma patients. This body of work can be readily extended to other biomarkers and cancer types and is not limited in its scope or application.

7. CODE AVAILABILITY

The optimized registration code detailed in Section 2.1.3 is available on the following GitHub page: <https://github.com/ounissimehdi/WSI-OPT-registration>

ACKNOWLEDGMENTS

This study was supported by The Cancer ITMO of the French National Alliance for Life and Health Sciences (AVIESAN): “MIC 2020” - project ref. C20051DS ”MALMO” (Mathematical Approaches to Modelling Metabolic Plasticity and Heterogeneity in Melanoma). We would like to acknowledge our discussions with Prof. Denis Wirtz, Prof. Pei-Hsun Wu, and their team in the Wirtz/Wu Lab of John Hopkins University and thank them for their valuable time.

REFERENCES

- [1] Arnold, M., Singh, D., Laversanne, M., Vignat, J., Vaccarella, S., Meheus, F., Cust, A. E., de Vries, E., Whiteman, D. C., and Bray, F., “Global burden of cutaneous melanoma in 2020 and projections to 2040,” *JAMA Dermatol.* **158**, 495–503 (May 2022).
- [2] Erdei, E. and Torres, S. M., “A new understanding in the epidemiology of melanoma,” *Expert Rev. Anti-cancer Ther.* **10**, 1811–1823 (Nov. 2010).

- [3] Davis, L. E., Shalin, S. C., and Tackett, A. J., “Current state of melanoma diagnosis and treatment,” *Cancer Biol. Ther.* **20**, 1366–1379 (Aug. 2019).
- [4] Garbe, C., Eigentler, T. K., Keilholz, U., Hauschild, A., and Kirkwood, J. M., “Systematic review of medical treatment in melanoma: current status and future prospects,” *Oncologist* **16**, 5–24 (Jan. 2011).
- [5] Anderson, N. M. and Simon, M. C., “The tumor microenvironment,” *Curr. Biol.* **30**, R921–R925 (Aug. 2020).
- [6] Ratnikov, B. I., Scott, D. A., Osterman, A. L., Smith, J. W., and Ronai, Z. A., “Metabolic rewiring in melanoma,” *Oncogene* **36**, 147–157 (Jan. 2017).
- [7] Hodgkinson, A., Trucu, D., Lacroix, M., Le Cam, L., and Radulescu, O., “Computational model of heterogeneity in melanoma: Designing therapies and predicting outcomes,” *Frontiers in Oncology* **12**, 857572 (2022).
- [8] Song, Y., Treanor, D., Bulpitt, A. J., and Magee, D. R., “3d reconstruction of multiple stained histology images,” *Journal of Pathology Informatics* **4**, 7 (Mar. 2013).
- [9] Lotz, J., Berger, J., Müller, B., Breuhahn, K., Grabe, N., Heldmann, S., Homeyer, A., Lahrmann, B., Laue, H., Olesch, J., Schwier, M., Sedlacek, O., and Warth, A., “Zooming in: high resolution 3d reconstruction of differently stained histological whole slide images,” in [*SPIE Proceedings*], Gurcan, M. N. and Madabhushi, A., eds., SPIE (Mar. 2014).
- [10] Liang, Y., Wang, F., Treanor, D., Magee, D., Teodoro, G., Zhu, Y., and Kong, J., “Liver whole slide image analysis for 3d vessel reconstruction,” in [*2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*], IEEE (Apr. 2015).
- [11] Paknezhad, M., Loh, S. Y. M., Choudhury, Y., Koh, V. K. C., Yong, T. T. K., Tan, H. S., Kanavar, R., Tan, P. H., Peng, J. Y. S., Yu, W., Tan, Y. B., Loy, Y. Z., Tan, M.-H., and Lee, H. K., “Regional registration of whole slide image stacks containing major histological artifacts,” *BMC Bioinformatics* **21** (Dec. 2020).
- [12] Kiemen, A. L., Braxton, A. M., Grahn, M. P., Han, K. S., Babu, J. M., Reichel, R., Jiang, A. C., Kim, B., Hsu, J., Amoa, F., Reddy, S., Hong, S.-M., Cornish, T. C., Thompson, E. D., Huang, P., Wood, L. D., Hruban, R. H., Wirtz, D., and Wu, P.-H., “CODA: quantitative 3d reconstruction of large tissues at cellular resolution,” *Nature Methods* **19**, 1490–1499 (Oct. 2022).
- [13] Arslan, J. and Benke, K. K., “Progression of geographic atrophy: Epistemic uncertainties affecting mathematical models and machine learning,” *Translational Vision Science & Technology* **10**, 3 (Nov. 2021).
- [14] Jadon, S., “A survey of loss functions for semantic segmentation,” in [*2020 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB)*], 1–7 (2020).
- [15] Öfverstedt, J., Lindblad, J., and Sladoje, N., “Fast and robust symmetric image registration based on distances combining intensity and spatial information,” **28**, 3584–3597 (2018).
- [16] Schenk, A., Prause, G., and Peitgen, H.-O., “Efficient semiautomatic segmentation of 3D objects in medical images,” in [*Medical Image Computing and Computer-Assisted Intervention – MICCAI 2000*], *Lecture notes in computer science*, 186–195, Springer Berlin Heidelberg, Berlin, Heidelberg (2000).
- [17] Bankhead, P., Loughrey, M. B., Fernández, J. A., Dombrowski, Y., McArt, D. G., Dunne, P. D., McQuaid, S., Gray, R. T., Murray, L. J., Coleman, H. G., James, J. A., Salto-Tellez, M., and Hamilton, P. W., “QuPath: Open source software for digital pathology image analysis,” *Sci. Rep.* **7**, 16878 (Dec. 2017).