

A peer-reviewed version of this preprint was published in PeerJ on 15 February 2018.

[View the peer-reviewed version](https://peerj.com/articles/4312) (peerj.com/articles/4312), which is the preferred citable publication unless you specifically need to cite this preprint.

Nunez-Iglesias J, Blanch AJ, Looker O, Dixon MW, Tilley L. 2018. A new Python library to analyse skeleton images confirms malaria parasite remodelling of the red blood cell membrane skeleton. PeerJ 6:e4312 <https://doi.org/10.7717/peerj.4312>

A new Python library to analyse skeleton images confirms malaria parasite remodelling of the red blood cell membrane skeleton

Juan Nunez-Iglesias ^{Corresp., 1}, Adam J Blanch ², Oliver Looker ², Matthew W Dixon ², Leann Tilley ²

¹ Melbourne Bioinformatics, The University of Melbourne, Australia

² Department of Biochemistry and Molecular Biology, Bio21 Institute, University of Melbourne, Melbourne, Australia

Corresponding Author: Juan Nunez-Iglesias

Email address: juan.n@unimelb.edu.au

We present Skan (Skeleton analysis), a Python library for the analysis of the skeleton structures of objects. It was inspired by the “analyse skeletons” plugin for the Fiji image analysis software, but its extensive Application Programming Interface (API) allows users to examine and manipulate any intermediate data structures produced during the analysis. Further, its use of common Python data structures such as SciPy sparse matrices and pandas data frames opens the results to analysis within the extensive ecosystem of scientific libraries available in Python. We demonstrate the validity of Skan’s measurements by comparing its output to the established Analyze Skeletons Fiji plugin, and, with a new scanning electron microscopy (SEM)-based method, we confirm that the malaria parasite *Plasmodium falciparum* remodels the host red blood cell cytoskeleton, increasing the average distance between spectrin-actin junctions.

1 A new Python library to analyse skeleton 2 images confirms malaria parasite 3 remodelling of the red blood cell membrane 4 skeleton

5 Juan Nunez-Iglesias^{1,2}, Adam J. Blanch^{1,3}, Oliver Looker³, Matthew W.
6 Dixon³, and Leann Tilley³

7 ¹Co-first author

8 ²Melbourne Bioinformatics, The University of Melbourne, Australia

9 ³Department of Biochemistry and Molecular Biology, Bio21 Institute, The University of
10 Melbourne, Australia

11 Corresponding author:

12 Juan Nunez-Iglesias²

13 Email address: juan.n@unimelb.edu.au

14 ABSTRACT

15 We present Skan (Skeleton analysis), a Python library for the analysis of the skeleton structures of objects.
16 It was inspired by the “analyse skeletons” plugin for the Fiji image analysis software, but its extensive
17 Application Programming Interface (API) allows users to examine and manipulate any intermediate data
18 structures produced during the analysis. Further, its use of common Python data structures such as SciPy
19 sparse matrices and pandas data frames opens the results to analysis within the extensive ecosystem of
20 scientific libraries available in Python. We demonstrate the validity of Skan’s measurements by comparing
21 its output to the established Analyze Skeletons Fiji plugin, and, with a new scanning electron microscopy
22 (SEM)-based method, we confirm that the malaria parasite *Plasmodium falciparum* remodels the host
23 red blood cell cytoskeleton, increasing the average distance between spectrin-actin junctions.

24 INTRODUCTION

25 Skeletons are single-pixel thick representations of networks within an image, and have wide application
26 to understanding the structural properties of objects. For example, skeletons have been used to model
27 human poses, neuronal morphology, nanofibre structure, road networks, kidney development, and vascular
28 networks, among others (Yim et al., 2000; Sundar et al., 2003; Bas and Erdogmus, 2011; Yuan et al.,
29 2009; Morales-Navarrete et al., 2015; Sambaer et al., 2011). These applications include both 2D and 3D
30 images, and often 3D images collected over time, underscoring the need for skeleton analysis software to
31 support multiple imaging modalities and dimensionality.

32 In this paper, we report Skan, a Python library that produces graphs and branch statistics from skeleton
33 images. Skan is written in Python using the Numba just-in-time (JIT) compiler (Lam et al., 2015) for
34 performance-critical code, including graph building and graph statistics computation. The source code
35 is available at <https://github.com/jni/skan> (under a BSD 3-clause license), and we encourage readers to
36 contribute code or raise GitHub issues where they require additional functionality to meet their needs.
37 Skan can be installed using standard tools from the two leading Python repositories, the Python Package
38 Index (PyPI) and conda-forge. Installation and usage instructions are available at <https://jni.github.io/skan>.

39 Skan works transparently with images of any dimensionality, allowing the analysis of 2D and 3D
40 skeletons. Out of the box, Skan provides functions to compute the pixel skeleton graph, compute statistics
41 about the branches of the skeleton, and draw skeletons and statistical overlays for 2D images.

42 The pixel skeleton graph maps which pixel is connected to which others in the skeleton image, as
43 well as the distances between them. This graph is provided in the standard `scipy.sparse.csr_matrix` sparse

44 matrix format, enabling further analysis using common tools for graph and array manipulation in the
45 scientific Python ecosystem.

46 From this graph, we can compute statistics about the branches of the skeleton, defined as junction-
47 junction and junction-endpoint paths in the pixel skeleton graph. These statistics include average branch
48 length, branch type, branch curvature, branch endpoints, branch euclidean length, and average image
49 intensity along the branch. We return these statistics as a pandas DataFrame, the de-facto standard format
50 for data tables in Python. The table includes the pixel IDs of the branch endpoints, allowing further
51 analysis of the junction-junction graph. Indeed, increasing the breadth of statistics computed by the
52 software was the primary motivation for Skan's development.

53 Skan further provides a rudimentary GUI to analyse batches of input images. The output of the
54 GUI interface is an Excel file, which contains all the above-mentioned statistics, as well as all analysis
55 parameters, to aid future reproducibility.

56 Because Skan uses common scientific Python data structures internally, it is easy to extend with new
57 statistics and analyses. The DataFrame of branch statistics follows the "tidy data" paradigm (Wickham,
58 2014), with each row representing one branch of a skeleton, facilitating downstream analysis, such as
59 computing summary statistics for each disjoint skeleton in an image.

60 To demonstrate Skan's 3D capabilities, we first compared its output to that of Fiji's Analyze Skeletons
61 (Arganda-Carreras et al., 2010), applied to a publicly available dataset of neuron skeleton traces. Then, we
62 used Skan to measure the spectrin cytoskeleton on the cytoplasmic side of the plasma membrane of red
63 blood cells (RBCs) infected with the malaria parasite *Plasmodium falciparum*, using a new SEM-based
64 protocol, and confirmed the remodelling of the RBC membrane skeleton by the parasite.

65 METHODS

66 Analysis of skeleton model from DIADEM challenge

67 We downloaded the olfactory projection neuron 1 (OP-1) model as a SWC file from DIADEM's website
68 at http://diademchallenge.org/data_set_downloads.html, along with its corresponding 3D TIFF image
69 stack. We then rasterised the model (i.e. converted it from a network of vertex coordinates to a set
70 of active pixels) by using the Simple Neurite Tracer (Longair et al., 2011) plugin for Fiji, function
71 "Analysis > Render/Analyze Skeletonized Paths." This produces a 6-connected skeleton path, which
72 we needed to convert to a (thinner) 26-connected path, so we further skeletonized the raster with the
73 `morphology.skeletonize3d` function from scikit-image (van der Walt et al., 2014), and saved it
74 as a compressed TIFF file.

75 Then, we imported this raster image into either Fiji or Python (using Christoph Gohlke's TIFFfile). In
76 both cases, we manually set the scale to $9.100602 \times 3.033534 \times 3.033534 \mu\text{m}$ per voxel, as documented
77 on the DIADEM website. In Fiji, we used "Analyze Skeletons" with the "Show detailed info" option
78 ticked, saved the results to csv, and loaded them into a pandas DataFrame in Python. In Python, we
79 used `skan.csr.summarise` to produce a corresponding pandas DataFrame for Skan's analysis.
80 Finally, we used `numpy.histogram` and `matplotlib.pyplot.hist` (Hunter, 2007) to produce
81 the histogram in Figure 1.

82 Tissue origin and ethics approval

83 This study made use of donated human red blood cells. All experiments were approved by The University
84 of Melbourne School of Biomedical Sciences, Human Ethics Advisory Group (HEAG), for project titled
85 "Characterising host cell interactions in the human malaria parasite, *Plasmodium falciparum*", and ethics
86 ID 1135799. Cells were obtained by a Material Supply Agreement with the Australian Red Cross Blood
87 Service, agreement number – 17-05VIC-23.

88 Sample preparation and SEM imaging

89 To prepare sheared membranes, infected and uninfected red blood cells were attached to 3-Aminopropyl-
90 triethoxysilane treated glass slides using the lectin erythro-agglutinating phytohemagglutinin (PHA-E)
91 and sheared in a hypotonic buffer according to a previously established procedure (Shi et al., 2013).
92 Sheared membranes were immediately fixed with 2.5% glutaraldehyde for 1 h before dehydration in a
93 series of ethanol:water mixtures of 20, 50, 70, 80, 90, 95 and (3x) 100% ethanol for 5 minutes each and
94 finally being allowed to dry in air.

95 Dried samples were gold coated on the rotating mount of a Dynavac SC100 sputter coating instrument
96 for 35 seconds using a 25 mA current, measuring 0.2 nm thickness on the quartz crystal microbalance.
97 The coating procedure was optimised to prevent under- or overcoating which presents problems with the
98 skeleton trace.

99 SEM images were recorded using the ETD detector (in Optiplan mode) of an FEI Teneo instrument
100 with a working distance of 5 mm, a beam current of 50 pA and a 2 kV accelerating voltage. Multiple
101 images at 200-250 k magnification were recorded per individual cell to cover a greater portion of the
102 membrane.

103 **Extraction of skeleton data from SEM images**

104 In our SEM images, the spectrin-actin network appears as bright (raised) patches over dark patches of
105 background (see Figure 1). We followed a simple approach to trace the midline of the spectrin branches:
106 smoothing the images, then thresholding them (Sauvola and Pietikäinen, 2000), and finally thinning them
107 (Zhang and Suen, 1984). The width of the Gaussian smoothing, the window size for thresholding, and the
108 offset for the thresholding are all parameters of our approach, and are recorded in the results output file of
109 a skeleton analysis (when using the graphical user interface).

110 **Data and code availability**

111 Our code is open source and available at <https://github.com/jni/skan>. Its documentation can be viewed at
112 <https://jni.github.io/skan> and includes all code necessary to reproduce Figure 2. Additional scripts used in
113 our analyses are available at <https://github.com/jni/skan-scripts>.

114 We have made the schizont SEM dataset available at the Open Science Framework (OSF), with DOI
115 10.17605/OSF.IO/SVPFU, together with an archive of the documentation at time of publication, and a
116 sample Excel file resulting from analysing the schizont dataset using the GUI.

117 **RESULTS**

118 **Comparison to Fiji's Analyze Skeletons plugin**

119 As a check that our software was producing results consistent with the existing literature on skeleton
120 analysis, we compared our software's results with that of Fiji's Analyze Skeletons plugin (Arganda-
121 Carreras et al., 2010). Although the original data from that paper is unavailable (Ignacio Arganda-
122 Carreras, pers. commun.), we compared the output of our software with that of Analyze Skeletons on a
123 neuron skeleton from the DIADEM Challenge (<http://diademchallenge.org>) (Figure 1A-B). Both software
124 packages found the same number of skeleton branches, with very close agreement between the two branch
125 length distributions (Figure 1C) and branch point locations (Figure 1D). Manual inspection confirmed that
126 the small differences observed result from the different treatment of branch junctions (see Supplementary
127 Information).

128 **Malaria parasites remodel the red blood cell inner membrane cytoskeleton**

129 Prior studies have shown that infection by *P. falciparum*, the most deadly malaria-causing parasite, results
130 in changes in the physical properties of the infected red blood cell (iRBC), and that these changes are
131 associated with an elongation of the spectrin skeleton branches in the inner RBC membrane skeleton
132 (Shi et al., 2013; Dearnley et al., 2016; Nans et al., 2011). A coarse-grained molecular model suggested
133 that this spectrin stretching could, in part, account for the deformability changes of the iRBC (Dearnley
134 et al., 2016), emphasizing the biological significance of the measurements. We sought to confirm these
135 observations using a novel scanning electron microscopy (SEM)-based protocol (Blanch A. et al., in
136 preparation). The method involves cross-linking the RBCs to a glass coverslip and shearing off the
137 upper membrane component, thus exposing the cytoplasmic/internal side of the cross-linked plasma
138 membrane (Shi et al., 2013). The membrane is chemically fixed, dehydrated and gold-coated before
139 imaging (Figure 2A). We automatically extracted spectrin skeletons (Figure 2B) from images produced
140 using both uninfected RBCs and RBCs infected with mature stage parasites (40-44h post infection). We
141 found that the average spectrin branch distance increased from 45.5nm to 49.2nm, an increase of 8%
142 (Figure 2C-D).

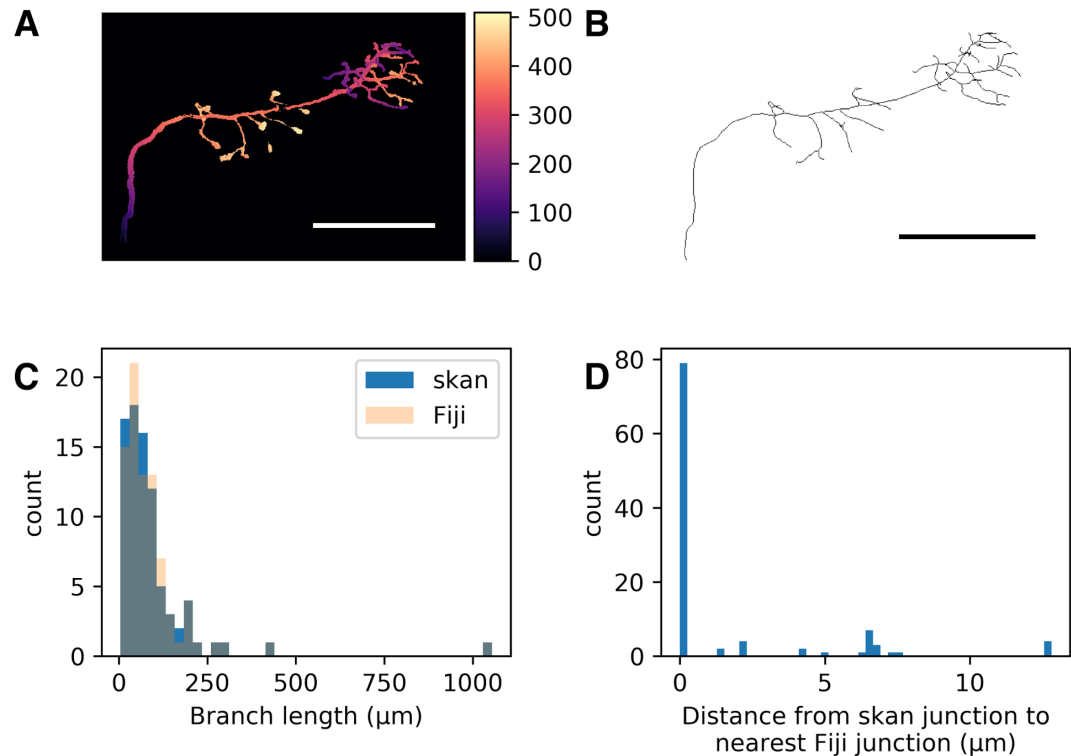


Figure 1. Comparison of skan and Fiji analysis results of the neuronal skeleton from olfactory projection neuron 1 (OP-1) from the DIADEM challenge. (A) Depth projection of the neuron. Scale bar: 500 μm . Colour map: height in μm . (B) Skeleton of the neuron. (C) Distribution of 82 branch lengths between 103 branch points measured by Skan and Fiji in the neuronal skeleton. (D) Distance from 103 skan junction points to the nearest Fiji junction point. Note that the voxel spacing is approximately $9 \times 3 \times 3 \mu\text{m}$, so almost all of these distances are less than one pixel apart.

143 DISCUSSION AND CONCLUSIONS

144 Spectrin remodelling by *P. falciparum*

145 The remarkable deformability and durability of the RBC membrane derives from its membrane skeleton
 146 (Zhang et al., 2015). The skeleton is composed of a regular hexagonal array of “spring-like” proteins
 147 forming a meshwork at the cytoplasmic surface of the RBC. Spectrin heterodimers constitute the cross-
 148 beams of the molecular architecture and are connected to integral membrane proteins in the plasma
 149 membrane. Previous studies using atomic force microscopy (AFM) and transmission electron microscopy
 150 (TEM), followed by manual selection and measurement of skeleton branches, revealed reorganization and
 151 expansion of the spectrin network of the host cell membrane (Shi et al., 2013; Millholland et al., 2011;
 152 Cyrklaff et al., 2011).

153 In this work we have applied a novel SEM-based method to image the RBC membrane skeleton, and a
 154 fully automated method for selection and measurement of the spectrin branch distances. We observed an
 155 8% increase in the length of the spectrin cross-members, in reasonable agreement with previous studies.
 156 Our data are consistent with the Cyrklaff et al. (2011) cryo-electron tomography study that provided
 157 evidence that the RBC membrane skeleton is reorganised as a result of mining of the actin junction points
 158 to generate actin filaments that connect parasite-derived organelles known as Maurer’s clefts to the knobs.

159 Numba and performance

160 An interesting aspect of Skan’s development is its use of Numba, a just-in-time compiler for Python code.
 161 Skan is one of the first scientific packages to make extensive use of Numba to speed up its operations.

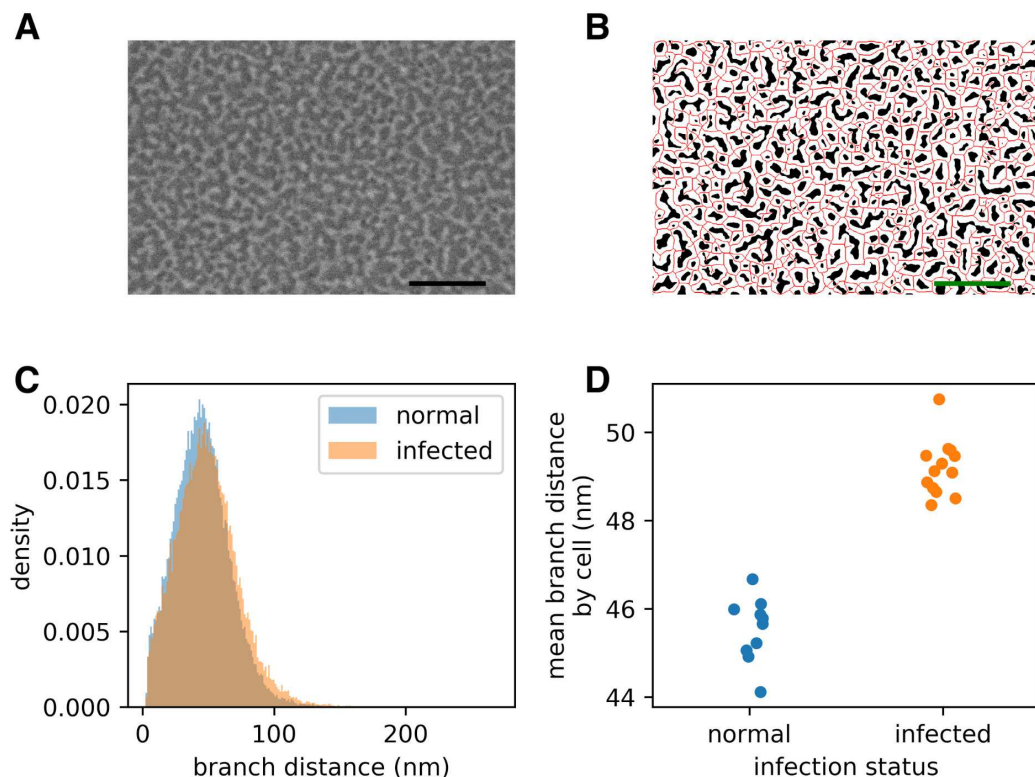


Figure 2. Infection by the malaria parasite remodels the spectrin skeleton of the host red blood cell in the asexual developmental stage. (A) Example image produced by our protocol. Scale bar: 300 nm. (B) Thresholding (white) and skeletonisation (red) of the image in (A). (C) Complete distribution of measured spectrin branch lengths for normal and infected RBCs. (D) Mean spectrin branch length by cell ($n_{\text{norm}} = 10$, $n_{\text{inf}} = 13$).

162 In our hands, Numba has been able to dramatically speed up our code, in some cases approaching the
 163 theoretical maximum performance of our CPUs.

164 As just one example, in the context of implementing Sauvola image thresholding, we developed a
 165 function for the cross-correlation of an n-dimensional image with a sparse kernel. Sauvola thresholding
 166 requires computing the local mean and standard deviation for every pixel in an image. This can be
 167 optimally achieved by computing the integral of both the original image and the image of squared
 168 intensity values, and then convolving each with a kernel consisting of the outer product of the vector
 169 $(-1, 0, 0, \dots, 0, 1)$ with itself, where the number of zeros is equal to the width of the neighbourhood minus
 170 one. This definition results in an extremely sparse kernel, which is not efficiently used by conventional
 171 convolution functions available in NumPy (v1.12) and SciPy (v0.19) (Walt et al., 2011; Oliphant, 2007).

172 The function is implemented as `correlate_sparse`, which handles boundary conditions and
 173 formatting of the kernel, and then calls the Numba-jitted function `_correlate_sparse_offsets`,
 174 which iterates through the array, performing the cross-correlation.

175 The result is striking. For a 2048 by 2048 pixel image and a 31 by 31 kernel size, `correlate_sparse`
 176 takes 130ms, somewhat slower than SciPy's `ndimage.correlate`, which takes 57ms. For a much
 177 bigger 301 by 301 kernel, however, `correlate_sparse` takes a similar amount of time — 135ms —
 178 while SciPy takes 17s. Furthermore, if we analyse just the inner loop of the computation, the part handled
 179 by Numba, we measured a time of 1.8ns per loop in our 1.3GHz (i.e. 0.77ns per cycle) CPU. Each loop
 180 performs two additions and a multiplication, in addition to array access, suggesting that Numba is close to
 181 achieving optimal performance for our problem and CPU. This example illustrates the power of Numba
 182 to speed up numerical Python code.

183 We also take this opportunity to note the loop order in the code of `_correlate_sparse_offsets`.

184 For every non-zero element of the kernel, we make a full pass over the input image. When picturing a
185 convolution, this is slightly counter-intuitive: most people would instead consider, for each pixel position,
186 correlating all the non-zero elements of the kernel (thus examining each pixel only once).

187 However, that order of operations is poorly optimised for modern processor architectures, which
188 fetch RAM contents by chunks into the processor cache. Once a chunk has been loaded, access-
189 ing elements of that chunk is 20-200 times faster than fetching more data from RAM (Jonas Bonér,
190 <https://gist.github.com/jboner/2841832>). One consequence is that algorithms that access data in the order
191 in which it is stored in RAM end up being much faster, by virtue of using processor cache to the maximum
192 extent possible.

193 In our case, this translated to a 10-fold speedup when changing the order from (for pixel in
194 image: for elem in kernel) to (for elem in kernel: for pixel in image),
195 even though these two expressions are mathematically equivalent.

196 ACKNOWLEDGEMENTS

197 We thank Ignacio Arganda-Carreras for suggesting the dataset to compare Skan to Analyze Skeletons,
198 and Alan Rubin for help with Skan's GUI code.

199 REFERENCES

- 200 Arganda-Carreras, I., Fernández-González, R., Muñoz-Barrutia, A., and Ortiz-De-Solorzano, C. (2010).
201 3D reconstruction of histological sections: Application to mammary gland tissue. *Microsc. Res. Tech.*,
202 73(11):1019–1029.
- 203 Bas, E. and Erdogmus, D. (2011). Principal curves as skeletons of tubular objects. *Neuroinformatics*.
- 204 Cyrklaff, M., Sanchez, C. P., Kilian, N., Bisseye, C., Simporé, J., Frischknecht, F., and Lanzer, M. (2011).
205 Hemoglobins S and C interfere with actin remodeling in plasmodium falciparum-infected erythrocytes.
206 *Science*, 334(6060):1283–1286.
- 207 Dearnley, M., Chu, T., Zhang, Y., Looker, O., Huang, C., Klonis, N., Yeoman, J., Kenny, S., Arora, M.,
208 Osborne, J. M., Chandramohanadas, R., Zhang, S., Dixon, M. W. A., and Tilley, L. (2016). Reversible
209 host cell remodeling underpins deformability changes in malaria parasite sexual blood stages. *Proc.*
210 *Natl. Acad. Sci. U. S. A.*, 113(17):4800–4805.
- 211 Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing In Science & Engineering*,
212 9(3):90–95.
- 213 Lam, S. K., Pitrou, A., and Seibert, S. (2015). Numba: A LLVM-based python JIT compiler. In
214 *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM '15*, pages
215 7:1–7:6, New York, NY, USA. ACM.
- 216 Longair, M. H., Baker, D. A., and Armstrong, J. D. (2011). Simple neurite tracer: open source software
217 for reconstruction, visualization and analysis of neuronal processes. *Bioinformatics*, 27(17):2453–2454.
- 218 Millholland, M. G., Chandramohanadas, R., Pizarro, A., Wehr, A., Shi, H., Darling, C., Lim, C. T.,
219 and Greenbaum, D. C. (2011). The malaria parasite progressively dismantles the host erythrocyte
220 cytoskeleton for efficient egress. *Mol. Cell. Proteomics*, 10(12):M111.010678.
- 221 Morales-Navarrete, H., Segovia-Miranda, F., Klukowski, P., Meyer, K., Nonaka, H., Marsico, G.,
222 Chernykh, M., Kalaidzidis, A., Zerial, M., and Kalaidzidis, Y. (2015). A versatile pipeline for
223 the multi-scale digital reconstruction and quantitative analysis of 3D tissue architecture. *Elife*, 4.
- 224 Nans, A., Mohandas, N., and Stokes, D. L. (2011). Native ultrastructure of the red cell cytoskeleton by
225 cryo-electron tomography. *Biophys. J.*, 101(10):2341–2350.
- 226 Oliphant, T. E. (2007). SciPy: Open source scientific tools for python. *Computing in Science and*
227 *Engineering*, 9:10–20.
- 228 Sambaer, W., Zatloukal, M., and Kimmer, D. (2011). 3D modeling of filtration process via polyurethane
229 nanofiber based nonwoven filters prepared by electrospinning process. *Chem. Eng. Sci.*, 66(4):613–623.
- 230 Sauvola, J. and Pietikäinen, M. (2000). Adaptive document image binarization. *Pattern Recognit.*,
231 33(2):225–236.
- 232 Shi, H., Liu, Z., Li, A., Yin, J., Chong, A. G. L., Tan, K. S. W., Zhang, Y., and Lim, C. T. (2013). Life
233 cycle-dependent cytoskeletal modifications in plasmodium falciparum infected erythrocytes. *PLoS*
234 *One*, 8(4):e61170.

- 235 Sundar, H., Silver, D., Gagvani, N., and Dickinson, S. (2003). Skeleton based shape matching and
236 retrieval. In *Shape Modeling International, 2003*, pages 130–139.
- 237 van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart,
238 E., Yu, T., and scikit-image contributors (2014). scikit-image: image processing in python. *PeerJ*,
239 2:e453.
- 240 Walt, S. v. d., Colbert, S. C., and Varoquaux, G. (2011). The NumPy array: A structure for efficient
241 numerical computation. *Comput. Sci. Eng.*, 13(2):22–30.
- 242 Wickham, H. (2014). Tidy data. *J. Stat. Softw.*, 59(10).
- 243 Yim, P. J., Choyke, P. L., and Summers, R. M. (2000). Gray-scale skeletonization of small vessels in
244 magnetic resonance angiography. *IEEE Trans. Med. Imaging*, 19(6):568–576.
- 245 Yuan, X., Trachtenberg, J. T., Potter, S. M., and Roysam, B. (2009). MDL constrained 3-D grayscale
246 skeletonization algorithm for automated extraction of dendrites and spines from fluorescence confocal
247 images. *Neuroinformatics*, 7(4):213–232.
- 248 Zhang, T. Y. and Suen, C. Y. (1984). A fast parallel algorithm for thinning digital patterns. *Commun.*
249 *ACM*, 27(3):236–239.
- 250 Zhang, Y., Huang, C., Kim, S., Golkaram, M., Dixon, M. W. A., Tilley, L., Li, J., Zhang, S., and Suresh,
251 S. (2015). Multiple stiffening effects of nanoscale knobs on human red blood cells infected with
252 plasmodium falciparum malaria parasite. *Proc. Natl. Acad. Sci. U. S. A.*, 112(19):6068–6073.

253 SUPPLEMENTARY RESULTS

254 Pixel-level comparison with Analyze Skeletons

255 Although the results obtained by our software are broadly similar to those of Analyze Skeletons, we
256 investigated the behaviour of both tools on a tiny example image, to ensure that we understood the minor
257 discrepancies between the two.

258 The differences between the libraries occur in the handling of clusters of junction nodes. It is
259 impossible to guarantee that several branches of a skeleton will converge on a single “junction” pixel. For
260 example, Supplementary Figure 1A shows a fully-reduced skeleton of 10 pixels. Counting the number of
261 neighbours of each pixel, we classify pixels as end-points, paths, and junctions. However, we can see that
262 four contiguous pixels at the bottom-centre are all considered junctions (Supplementary Figure S1B-C).
263 We saw three possible ways to deal with such clusters when computing branch statistics:

- 264 1. Ignore them — the branch from pixel 1 to pixel 3 will be considered to have length 2, as will the other
265 branches; the junction clusters are thus considered to have some “spatial extent.” (Supplementary
266 Figure S1D)
- 267 2. Replace them by a single pixel, perhaps the pixel closest to the centroid of the pixels.
- 268 3. Replace them by their centroid. (Supplementary Figure S1E)

269 In Skan, we added a flag, `unique_junctions`, that selects between options 1 and 3. The Analyze
270 Skeletons plugin, in contrast, uses a variant of option 2, where the pixel selected to represent the cluster is
271 the “earliest,” in lexicographical order of (x, y, z) position (Supplementary Figure S1F). Although in
272 most situations these small subtleties would make little difference to the downstream results, we believe
273 our approach is closer to what a user would expect to get from their skeleton images.

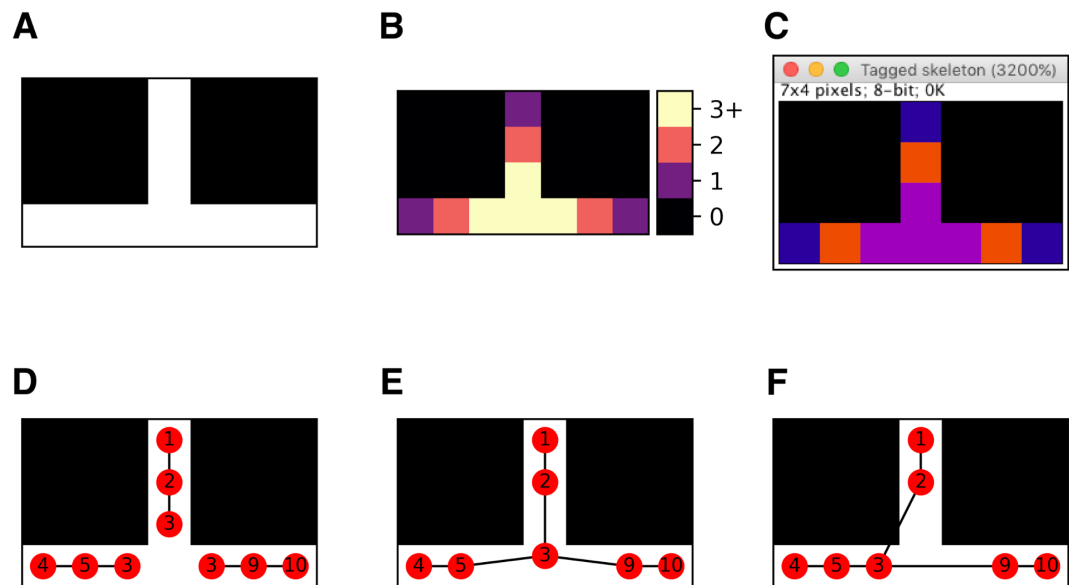


Figure S1. Strategies for resolving skeleton junctions. (A) A minimal skeleton. (B) Skan's classification of pixels into endpoints, paths, and junctions based on the number of neighbours (1, 2, and 3 or more, respectively). (C) Identical classification in Fiji's Analyze Skeletons. (D) Skeleton measurement when junctions are assigned an implicit "extent". (E) Skeleton measurement when all adjacent junction pixels are replaced by their centroid (our default strategy). (F) Skeleton measurement used in Fiji's Analyze skeletons (mid-2017 version).