

Sequential Graph-based Extraction of Curvilinear Structures

Shuaa S. Alharbi · Chris G. Willcocks · Philip T. G. Jackson · Haifa F. Alhasson · Boguslaw Obara

Received: date / Accepted: date

Abstract In this paper, a new approach is proposed to extract an ordered sequence of curvilinear structures in images, capturing the largest and most influential paths first and then progressively extracting smaller paths until a prespecified size is reached. The results are demonstrated both quantitatively and qualitatively using synthetic and real world images. The method is shown to outperform state-of-the-art methods for certain cases of noise, object class, and scale, while remaining fundamentally easier to use due to its low parameter requirement.

Keywords Curvilinear Structures · Centerline Enhancement · Graph-based Method · Object Detection.

1 Introduction

Extracting curvilinear structures, or sets of line segments, in digital images is an important low-level problem with many applications in computer vision [1]. There have been many attempts to redefine the ideal centreline properties and to develop bespoke algorithms in a variety of domains, from medicine to engineering [1–3].

Existing methods introduce a wide range of parameters that are often difficult to tune and/or unintuitive. Furthermore, these parameters often require to be tuned by the user for each image, resulting in faults or robustness issues outside the chosen values [4]. Table 1 compares the number of required parameters of some of the most popular centreline extraction methods. Other approaches [5] harness advances in machine learning to produce excellent results. The key problem with these approaches is that they require large training datasets, which are very expensive in terms of data acquisition and model training.

In this paper, we introduce a novel approach for sequential centreline extraction in two stages: (1) enhancement of the centreline based on prominent ridge detection and (2) extraction of the centreline based on an iterative graph-based optimisation. Graph-based

centreline extraction approaches [6, 7] are known to be robust to image noise, but they typically require user-specified start and end points. This approach searches through a subset of paths that fall along prominent ridges, sequentially returning the path with the largest intensity sum along its profile, and whose union with previous paths in the sequence maximally contributes to the output skeleton. Therefore, a user-specified seed is not required.

The main contribution of the proposed method is that it only requires a single tunable parameter ε and two optimisation parameters t_1 and t_2 without initialization regions or training data. The parameter ε corresponds to the length (in pixels) of the shortest lines to be captured. If the level of imaging noise is known *a priori*, ε is set to a slightly larger value than this level. The two optimisation parameters t_1 and t_2 are for the enhancement and extraction optimisation processes, respectively, and do not require tuning on a per-image basis. Therefore, this proposed algorithm works remarkably well as an out-of-the-box solution. The main stages of our pipeline are shown in Figure 1.

2 Related Work

2.1 Centreline Extraction

Many centreline extraction methods for 2D and 3D images have been proposed in the literature with applications in many different fields [1]. A summary is shown in Table 1.

2.1.1 Edge-based approaches

Simple gradient-based approaches are known to converge on local solutions, such as medial sheets in 3D. Gradient Vector Flow (GVF) addresses this problem by using a diffusion of the gradient information [2] in order to produce smooth centrelines. Smistad [3] combined the GVF, Hessian-based enhancement, and ridge-traversal. Leng et al. [9] propose enhancing the results of a simple Canny edge detector by using a voting field that enhances smooth curves and removes artificial lines caused by noise. Smistad [3] combined GVF, Hessian-based enhancement, and ridge-traversal. In order to




Fig. 1: Workflow of the proposed approach: (a) t_1 intensity profiles, defined by random lines parametrized with random positions $\mathbf{x}$ directions $\hat{n}$ and lengths l , are considered and cast votes at the locations of the maximum peaks into an accumulative space (b). This is followed by defining a graph $G(N, E)$ used to build a set of t_2 shortest paths between t_2 random start and end nodes N_i and N_j (c). Next, the most prominent shortest path in the set is found (d). The processes (c) and (d) are repeated until the length of the most prominent shortest path is no longer than ϵ , see the final extraction (e).

enhance the centreline, a searching algorithm is applied [10] that finds peaks in the intensity information along each axis using a sliding-window approach.

2.1.2 Graph-based approaches

Graph-based approaches find the shortest path or separating cut between intensity or color information. In particular, Jin et al. [7] find paths in objects with a tree-structure from a user-specified point, which is used to find the furthest point and the minimal connecting path. Similarly, Rouchdy and Cohen [6] consider the accumulative overlap of shortest paths from the image boundary towards a user-specified seed, which is shown to be robust to image noise.

Table 1: Comparison of centreline detection methods.

Representative Method	Approach	Dimensions	Notes
Zhang Suen et al. (Thinning) [8]	Thinning Process	2D	Requires binary segmentation
Leng et al. [9]	Canny edge and tensor voting	2D	4 tunable parameters
Smistad et al. [3]	Gradient vector and Hessian matrix	3D	9 tunable parameters
Lam et al. [2]	Normalised gradient vector field	2D	4 tunable parameters
Shao et al. [10]	Searching algorithm	2D	3 tunable parameters
Hassona et al. [11]	Level set	2D/3D	Requires binary segmentation
Shen et al. [12]	Deep learning	2D/3D	Requires large datasets and a long training process
Steger (UDCS) [4]	Ridge detecting	2D	3 tunable parameters
Sharma et al. (Voronoi) [13]	Voronoi diagram	2D	Sensitivity to noise
Lopez-Molina et al. (AGK) [14]	Anisotropic Gaussian kernels	2D	4 tunable parameters
Our method	Sequential graph-based extraction	2D/3D	1 tunable parameter and 2 optimization parameters

2.1.3 Learning-based approaches

Recent learning-based skeleton extraction methods are more suitable to deal with the scene complexity problem in natural images [5, 12, 15, 16]. Sironi et al. [5] method, learn huge numbers of parameters from large training datasets based on ground truths. As a result, their method is well-suited to ill-posed problems, such as centreline extraction. Recent work by Shen et al. [12, 15] proposed a multi-scaled learning framework that fuses the final output together. The object skeleton ground truth is used for certain object scales. The outputs from their method is a binary image denoting the detected skeletons, which are produced by thresholding a skeleton heatmap. Although the learning-based methods obtain good results, attaining ground truths can be expensive in terms of data acquisition and the training process itself.

2.1.4 Surface-skeleton approaches

Traditional thinning approaches [8] iteratively remove outer layers of boundary pixels according to local stopping criteria. However, these approaches are sensitive to surface noise. Similarly, geometric contraction approaches apply constrained iterative smoothing and/or merging, shrinking the surface into a thin centreline [17]. Alternatively, approaches grow or connect paths from locally centred regions; for instance, using a level set method propagating and connecting a wave front seeded from inside the deepest part of the object [11].

2.2 Ridge Detection

Ridges are defined as extrema of the image's largest surface curvature direction. Steger proposed a popular approach for extracting lines using Gaussian masks to estimate derivatives without bias in asymmetrical

lines [4]. Another approach by [14] handles crossings, junctions, and blobs by using anisotropic, multi-scale Gaussian kernels for second order image differentiation.

In conclusion, there are a large number of relevant approaches with varying input requirements, such as a reliance on training data, prior knowledge, or a number of unintuitive parameters. Graph-based approaches are shown to be robust to noise, which is often not the case with gradient-based detectors. Enhancement can be seen as an important pre-processing step, but it is difficult to parametrise the models, in particular choosing local orientation and scale.

3 Method

3.1 Overview

Enhancement is investigated based on the algorithmic definition of geographical prominence in 2D and 3D, which measures mountains from their summit to saddle points connecting to the next highest peak. Figure 2 shows how this property is desirable in images, whereby the most prominent peak (green arrow) captures the underlying object despite severe levels of multi-frequency inhomogeneous noise with higher local elevations, or summits, than the object itself.

The proposed method has two main conceptual stages: 1) centreline enhancement with geographical prominence and 2) sequential extraction of newly contributing bright paths. During the first step, an accumulative matrix is constructed using a voting procedure, whereby points in the matrix are at the index of prominent peaks of line (2D) or disk (3D) profiles. The second step searches through a subset of paths according to this matrix, sequentially extracting the most prominent paths that maximally contribute to the output skeleton according to a simple summation.




Fig. 2: An intensity profile along a random line. (a) The area of interest in the image has inhomogeneous backgrounds, highlighted by the red dashed square and the random line in blue, and (b) an illustration of the profile that shows prominent peaks compared with the most prominent peak (green arrow) that captures the underlying object structure.

3.2 Enhancement Process

In geography, the prominence of a peak in a mountain range is defined as the difference in elevation between the summit of the peak and the highest saddle that connects that summit to any higher area [18]. This algorithm requires no parameters to be tuned.

The proposed enhancement algorithm is formulated as a voting procedure based on prominence. Let us consider an image $I(\mathbf{x}) \in \mathbb{R}^n$. If $\mathbf{x}$ is a random point in the image, $\hat{n}$ is a random unit vector, and l is a random length between 0 and the size of the smallest dimension of I , then a random intensity profile along a line from $\mathbf{x}$ to $\mathbf{x} + l\hat{n}$ is defined as:

$$f = I(\mathbf{x} + i\hat{n}), \quad \forall i \in [0, l], \quad (1)$$

The intensity profile f is sampled and the index of the highest prominence peak is found, as in [18], as follows:

$$j = \text{prom}(f), \quad (2)$$

and its pixel location:

$$\mathbf{x}' = \mathbf{x} + j\hat{n}. \quad (3)$$

Then, a vote is cast into an accumulative map A :

$$A(\mathbf{x}') = A(\mathbf{x}') + 1, \quad (4)$$

as shown in Figure 3a. This process is repeated t_1 times, gradually accumulating more votes along ridge pixels. The accumulative map A is then normalized to be in the range $[0, 1]$.




Fig. 3: An illustration of finding the most prominent peaks in 2D and 3D images. (a) The proposed method uses a random line in 2D to find the most prominent peak $\mathbf{x}'$ and accumulates it into A . (b) The most prominent peak on a random plane in 3D.

3.3 Extraction Process

In order to extract the centreline, an undirected, weighted graph $G(N, E)$ is constructed, where each pixel position $\mathbf{x}$ corresponds to a node in N . Graph edges E connect each node N at pixel $\mathbf{x}$ to its 8-connected neighbourhood of pixels. Graph node costs are defined by $1/A$, such that the prominent ridges are set to small values and background regions tend to infinity.

From this graph definition G , a set of t_2 random shortest paths is produced iteratively:

$$P = \{\mathbf{p}_i\}, \forall i \in [1, t_2], \quad (5)$$

where each shortest path $\mathbf{p}_i$ is calculated between two random nodes using Dijkstra's algorithm [19].

The most prominent path $\tilde{\mathbf{p}}$ is found and stored in $\tilde{P} = \{\tilde{\mathbf{p}}\}$ at each iteration, where $\tilde{\mathbf{p}}$ is defined:

$$\tilde{\mathbf{p}} = \arg \max_{\mathbf{p}} \left(\sum_{\mathbf{x} \in \mathbf{p} \setminus \tilde{P}} A(\mathbf{x}) \right), \mathbf{p} \in P \setminus \tilde{P}, \quad (6)$$

which is the longest segment of the path, from the set of paths in P , with the highest cost across the accumulative map A , that does not intersect paths found in previous iterations $\tilde{P}$. All paths in the set $\tilde{P}$ must be disjointed (not sharing any pixels with other paths).

Equation 6 is iterated (Figure 4) until the length of the most prominent path $\tilde{\mathbf{p}}$ is shorter than ε .




Fig. 4: The first three iterations of the sequential extraction process. After these three iterations, there is a large drop in the length of the next path, which is captured easily by ε .

3.4 Centreline in 3D Images

The method is easily extensible to 3D images. In the enhancement procedure, the profile of a random disk is considered instead of a line. In particular, a random plane is chosen, parametrized by a point $\mathbf{x}$ and normal $\hat{n}$ and the single most prominent peak over the set of lines (of length l) that sweep through 360° on this plane is chosen. Only the most prominent peak from the set of lines on the plane will cast a vote in the accumulator array, see Figure 3b. In the second step, the graph edges are defined as the full connectivity between the current pixel and its 26 neighbours instead of 8 neighbours in the 2D version. The final skeleton is extracted as described in section (3.2, 3.3).

4 Results and Discussion

This section provides quantitative and qualitative validations against synthetic and real-world data.

4.1 Synthetic Validation

4.1.1 Noise Sensitivity

The robustness of the proposed method to Gaussian, salt & pepper and speckle noise types is measured, as

well as its robustness to a structured multi-frequency noise that simulates intensity inhomogeneity. In particular, a 2D synthetic image of a curvilinear object with a known ground-truth is created and the image is corrupted with noise, as shown in Table 3. The synthetic 2D curve is generated with a thickness radius of 3 and the resulting 2D image (128×128) is corrupted with different levels and types of noise. The Hausdorff distance is calculated between the extracted curve and its analytical ground truth (see Figure 5). The results show that our method remains stable under severe quantities of noise, heavily corrupting the object to a PSNR of about 11.




Fig. 5: The Hausdorff distance between the extracted curve and the ground truth as shown in Table 3. These results use the same parameters: $t_1 = 10000$, $t_2 = 2000$, and $\varepsilon = 2$. The lines show the mean across 100 experiments and the error envelopes (transparent shaded regions) show the standard deviation.

4.1.2 Parameter Space

In order to understand the effect of the optimization parameters t_1 and t_2 in the two stages (Section 3.2, 3.3), the two parameters are varied in 2D space: $t_1 = [5, 10000]$ and $t_2 = [1, 2000]$ and the Hausdorff distance between the ground-truth and extracted curve is plotted, shown in Table 2 (middle). As expected, there is a large rectangular flat region and increases in t_1 and t_2 do not produce a significant increase in quality. In particular, the Hausdorff distance = 2 at the base of the hill (at the point $t_1 = 2750$ $t_2 = 1000$) implies that t_1 and t_2 require very little tuning in practice. Generally, these parameters can be set as high as possible within the processing requirements.

4.1.3 Comparison with The Existing Methods

The proposed method is also validated against a wide range of centreline extraction methods using synthetic 2D data. In particular, we show (1) a modified tradi-

Table 2: Illustration of the Hausdorff distance (middle) and method runtime (right) with varying optimization parameters $t_1 = [5, 10000]$ and $t_2 = [1, 2000]$. We set the tunable parameter $\varepsilon = 2$. The Hausdorff distance is computed between the analytical ground-truth (blue curve, left-upper) and the extracted curve (red curve, left-lower). The dashed red lines show the point where the method completes at 5.7s with a Hausdorff distance of 2 pixels. The flat region (middle) shows the stability of these parameters.




Table 3: Comparison between the extraction centerline obtained from the proposed method and the-state-of-art methods for different noise types. The parameters for each method are tuned across all 6 scenarios and tuned manually per-image.

	Ground truth	Parameters tuned per-image			Parameters tuned across all images & held constant				
		Thinning [8]	UDCS [4]	AGK [14]	Thinning [8]	UDCS [4]	Voronoi [13]	AGK [14]	Our method (no tuning)
a									
b									
c									
d									
e									
f									

Table 4: The Hausdorff distance measures corresponding to the images are shown in Table 3. The last two rows show the means and standard deviations of the Hausdorff distance across state-of-the-art methods.

Image	Parameters tuned per-image			Parameters tuned across all images & held constant				
	Thinning [8]	UDCS [4]	AGK [14]	Thinning [8]	UDCS [4]	Voronoi [13]	AGK [14]	Our method (no tuning)
a	4.24	3.16	9.49	75.5	71.02	29.99	9.49	3.67
b	4.78	5.10	6.93	50.21	5.05	46.75	8.06	3.60
c	43.68	4.24	34.83	84.9	26.17	36.98	40.4	4.48
d	5	3.49	8.94	28.03	4.74	75.35	15.11	7.07
e	1.41	1.9	12	78	74.73	60.8	12	3
f	48.08	49.51	34.55	76.22	69.38	68.97	51.95	58.01
Mean (pixel)	17.86	11.23	17.79	65.48	41.85	56.37	22.83	13.30
Std (pixel)	21.78	18.78	13.19	21.84	33.67	18.21	18.60	21.95




Fig. 6: A selection of images in the Ghent University Fungal Images (GUF1-1) dataset [14, 20] alongside the ground truth centrelines and extracted curves from multiple methods. (a) Original image, (b) Ground truth, (c) Thinning [8], (d) UDCS [4], (e) Voronoi [13], (f) AGK [14], and (g) our method

tional thinning approach [8], which has been improved with both median and Gaussian filtering, and then thresholded to give a fair chance of success, (2) ridge detecting using Steger’s popular approach of unbiased detector of curvilinear structures (UDCS) [4], (3) a Voronoi skeleton approach [13], and (4) the recent anisotropic multiscale Gaussian kernels (AGK) [14]). The results are shown in Table 3. Furthermore, the methods with their tuned parameters are compared across all 6 scenarios and are tuned manually on a per-image basis. In addition, the Hausdorff distance between our extracted curve and the ground truth is calculated for these 6 scenarios. The results are presented in Table 4.

Our proposed approach succeeds under severe levels of noise where other approaches fail. The extracted centreline is robust, well-centred, and homotopic to the original object. However, it is not as smooth as with methods that use a Gaussian approximation of the derivatives. In other approaches, the tuning of parameters requires significantly more effort. In particular, the AGK method utilizes 11 different configurations to determine the set of parameter values. The UDCS method [4] gives excellent results when the parameters are carefully tuned on a per-scenario basis, but the method fails when the parameters are held constant. However, even simple approaches, such as thinning, outperform other methods when tuned on a per-image basis using different filtering techniques. In contrast, our method does not require such tuning.

4.2 Real-World Validation

Our method was evaluated with the Ghent University Fungal Images dataset 1 (GUF1-1) [14, 20], which is a popular and varied biological dataset with images of fungi growing in vitro. Each image has a resolution of 300×300 pixels and is accompanied with manually labelled ground truths for use in evaluating centreline extraction methods. These images contain a variety of ridges, different degrees of contamination, frequent overlaps and junctions, and large regions without edges and the dataset is therefore considered to

be rather challenging. The ground truth centrelines are not always thin (with 1 pixel width) and, therefore, the ground truth was manually altered to ensure the required thinness. The ground truth has been established independently by two experts and the Hausdorff distance was then calculated between the two sets in order to ensure the validity of proposed ground truth. The mean Hausdorff distance is 0.86 pixels, where the standard deviation is 0.54 pixels. The ground truths were revised by two other experts. The proposed method was evaluated alongside a variety of other methods, calculating the Hausdorff distance between the extracted skeletons and the ground truth. The results are shown in Figure 6 and Table 5. The results show that our method comes closest to the ground truth across the dataset, however certain features are missing which may be better captured by other methods [4, 8, 14]. However, unlike other methods, our approach tends to be cleaner and results in fewer artifacts, such as spurious branches. Therefore, it is inferred that the proposed method is beneficial in the task of collecting reliable metrics across a large dataset, but, for some specific analyses of individual images, it is recommended to consider more sensitive methods, such as [4].

Table 5: Comparing the proposed method with the state-of-the-art methods in term of the Hausdorff distance using the GUF1-1 dataset.

Method	Mean (pixel)	Std (pixel)
Thinning [8]	77.95	51.98
UDCS [4]	81.99	50.83
Voronoi [13]	87.13	56.72
AGK [14]	66.87	45.18
Our method	56.74	56.67

The proposed method is also validated qualitatively against several 2D and 3D real world images and the results are shown in Tables 6 and 7, respectively. Images showing dark objects on bright backgrounds are inverted before processing. In the 3D validation, the extracted curve is compared with the thinning method [21], which has a publicly available 3D implementation.

Table 6: Results for different 2D real-world images: (a) wires, (b) spirillum [22], (c) straight hair [23], (d) retina networks [24], (e-f) treponema [25, 26], (g) telomeric DNA from HeLa cell clone [27], and (h) electron microscopic image of isolated mouse circular mtDNA [28].





Table 7: Comparison between our method and topological thinning [21] using 3D images of: neocortical axon trees (a) [29], olfactory projection neuron tree (b) [29], and visual cortical trees (c) [29].



5 Future work

In our implementation¹, the performance of the algorithm may be significantly improved by removing low-weighted edges in the graph according to a thresholding parameter with very little impact on quality. Indeed, this is especially useful for segmenting large 3D images. In this paper, we have explored the performance of the proposed algorithm and the optimisation of our method is left as an area of future work.

Although the proposed method gives good 2D/3D results for biomedical and non-biomedical images, there is room for improvement. In particular, the enhancement process should be removed and the graph cost function should be extended in order to ensure centeredness in flat intensity distributions, as well as the better handling of junctions. This extension could also consider a distance metric that gives preference to network growth, penalising the creation of disconnected branches in distant regions.

6 Conclusion

A novel approach is described which searches through a subset of paths that fall along prominent ridges, sequentially returning the most prominent path that maximally contributes to the output centreline. The method is evaluated against real-world images and is shown to be comparable to the state-of-the-art with extracted centrelines that are close to the ground truth. In particular, the method has different strengths and weaknesses, such as its ability to bridge inhomogeneous gaps (Table 7b) and to handle complex/busy backgrounds. Furthermore, the method requires little tuning and works well as an out-of-the-box solution in medical (e.g., blood-vessel extraction) and biological (e.g., plant roots and neural networks) imaging applications.

Acknowledgment

Shuaa Alharbi is supported by the Saudi Arabian Ministry of Higher Education Doctoral Scholarship and Qasim University in Saudi Arabia.

References

1. C. Kirbas and F. Quek, "A review of vessel extraction techniques and algorithms," *ACM Computing Surveys*, vol. 36, no. 2, pp. 81–121, 2004.
2. B. S. Y. Lam and H. Yan, "A novel vessel segmentation algorithm for pathological retina images based on the divergence of vector fields," *IEEE Transactions on Medical Imaging*, vol. 27, no. 2, pp. 237–246, 2008.
3. E. Smistad, A. C. Elster, and F. Lindseth, "GPU-based airway segmentation and centerline extraction for image guided bronchoscopy," *Norsk Informatikkonferanse*, vol. 2012, no. 1, pp. 129–140, 2012.
4. C. Steger, "An unbiased detector of curvilinear structures," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 2, pp. 113–125, 1998.
5. A. Sironi, E. Türetken, V. Lepetit, and P. Fua, "Multi-scale centerline detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 7, pp. 1327–1341, 2016.
6. Y. Rouchdy and L. D. Cohen, "Geodesic voting for the automatic extraction of tree structures. methods and applications," *Computer Vision and Image Understanding*, vol. 117, no. 10, pp. 1453–1467, 2013.
7. D. Jin, K. S. Iyer, C. Chen, E. A. Hoffman, and P. K. Saha, "A robust and efficient curve skeletonization algorithm for tree-like objects using minimum cost paths," *Pattern Recognition Letters*, vol. 76, no. 1, pp. 32–40, 2016.
8. Y.-S. Chen and W.-H. Hsu, "A modified fast parallel algorithm for thinning digital patterns," *Pattern Recognition Letters*, vol. 7, no. 2, pp. 99–106, 1988.
9. Z. Leng, J. R. Korenberg, B. Roysam, and T. Tasdizen, "A rapid 2D centerline extraction method based on tensor voting," in *IEEE International Symposium on Biomedical Imaging: From Macro to Nano*, (Chicago, IL), pp. 1000–1003, Mar 2011.
10. Y. Shao, B. Guo, X. Hu, and L. Di, "Application of a fast linear feature detector to road extraction from remotely sensed imagery," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 4, no. 3, pp. 626–631, 2011.
11. M. S. Hassouna and A. A. Farag, "Variational curve skeletons using gradient vector flow," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 12, pp. 2257–2274, 2009.
12. W. Shen, K. Zhao, Y. Jiang, Y. Wang, X. Bai, and A. Yuille, "Deepskeleton: Learning multi-task scale-associated deep side outputs for object skeleton extraction in natural images," *IEEE Transactions on Image Processing*, vol. 26, no. 11, pp. 5298–5311, 2017.
13. O. Sharma, D. Mioc, and F. Anton, "Voronoi diagram based automated skeleton extraction from colour scanned maps," in *International Symposium on Voronoi Diagrams in Science and Engineering*, (Banff, Canada), pp. 186–195, Jul 2006.
14. C. Lopez-Molina, G. V.-D. de Ulzurrun, J. M. Baetens, J. Van den Bulcke, and B. De Baets, "Unsupervised ridge detection using second order anisotropic Gaussian kernels," *Signal Processing*, vol. 116, no. 1, pp. 55–67, 2015.
15. W. Shen, X. Bai, Z. Hu, and Z. Zhang, "Multiple instance subspace learning via partial random projection tree for local reflection symmetry in natural images," *Pattern Recognition*, vol. 52, no. 1, pp. 306–316, 2016.
16. W. Shen, K. Zhao, Y. Jiang, Y. Wang, Z. Zhang, and X. Bai, "Object skeleton extraction in natural images by fusing scale-associated deep side outputs," in *IEEE Conference on Computer Vision and Pattern Recognition*, (Las Vegas, Nevada), pp. 222–230, Jun 2016.
17. C. G. Willcocks and F. W. Li, "Feature-varying skeletonization," *The Visual Computer*, vol. 28, no. 6–8, pp. 775–785, 2012.
18. A. Helman, *The finest peaks: Prominence and other mountain measures*. Victoria, Canada: Trafford Publishing, 2005.
19. E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
20. "Fungal images," 2014. <http://www.kermit.ugent.be/software.php?navigatieId=135&categorieId=43>.

¹ The software has been implemented in MATLAB and is made available at: (Github repository will be included).

21. M. Kerschnitzki, P. Kollmannsberger, M. Burghammer, G. N. Duda, R. Weinkamer, W. Wagermaier, and P. Fratzl, "Architecture of the osteocyte network correlates with bone material quality," *Journal of Bone and Mineral Research*, vol. 28, no. 8, pp. 1837–1845, 2013.
22. "Spirillum oil bacterium," 2004. http://w3.marietta.edu/~spilatrs/biol202/labresults/spirillum_oil.jpg.
23. C. G. Willcocks, P. T. Jackson, C. J. Nelson, and B. Obara, "Extracting 3D parametric curves from 2D images of helical objects," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 9, pp. 1757–1769, 2017.
24. M. Niemeijer, J. Staal, B. van Ginneken, M. Loog, and M. D. Abramoff, "Comparative study of retinal vessel segmentation methods on a new publicly available database," in *Image Processing in Medical Imaging*, (San Diego, CA), pp. 648–657, May 2004.
25. "Treponema pallidum bacterium," 2016. https://phil.cdc.gov/phil/details_linked.asp?pid=2333.
26. "Treponema bacteria," 2004. <http://w3.marietta.edu/~spilatrs/biol202/labresults/treponema.html>.
27. J. D. Griffith, L. Comeau, S. Rosenfield, R. M. Stansel, A. Bianchi, H. Moss, and T. De Lange, "Mammalian telomeres end in a large duplex loop," *Cell*, vol. 97, no. 4, pp. 503–514, 1999.
28. J. H. Sinclair and B. J. Stevens, "Circular DNA filaments from mouse mitochondria," *Proceedings of The National Academy of Sciences*, vol. 56, no. 2, pp. 508–514, 1966.
29. K. M. Brown, G. Barrionuevo, A. J. Canty, V. De Paola, J. A. Hirsch, G. S. Jefferis, J. Lu, M. Snippe, I. Sugihara, and G. A. Ascoli, "The DIADEM data sets: representative light microscopy images of neuronal morphology to advance automation of digital reconstructions," *Neuroinformatics*, vol. 9, no. 2-3, pp. 143–157, 2011.