

Written in Blood: Applying Shape Grammars to Retinal Vasculatures

Ryan Y. Yeh¹, Ken K. Nischal², Philip LeDuc¹, and Jonathan Cagan¹

¹ Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA

² Division of Pediatric Ophthalmology, Strabismus and Adult Motility, University of Pittsburgh Medical Center Children's Hospital, Pittsburgh, PA, USA

Correspondence: Jonathan Cagan, Department of Mechanical Engineering, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA. e-mail: cagan@cmu.edu

Philip LeDuc, Carnegie Mellon University, Department of Mechanical Engineering, 5000 Forbes Ave, Pittsburgh, PA 15213, USA. e-mail: prl@andrew.cmu.edu

Received: May 20, 2020

Accepted: July 9, 2020

Published: August 24, 2020

Keywords: Shape grammar; synthetic vasculature; retinal image analysis; image processing

Citation: Yeh RY, Nischal KK, LeDuc P, Cagan J. Written in blood: Applying shape grammars to retinal vasculatures. *Trans Vis Sci Tech.* 2020;9(9):36. <https://doi.org/10.1167/tvst.9.9.36>

Purpose: Blood vessel networks within the retina are crucial for maintaining tissue perfusion and therefore good vision. Their complexity and unique patterns often require a steep learning curve for humans to identify trends and changes in the shape and topology of the networks, even though there exists much information important to identifying disease within them.

Methods: Through image processing, the vasculature is isolated from other features of the fundus images, forcing the viewer to focus on the complex vascular feature. This article explores an approach using a grammar based on shape to describe retinal vasculature and to generate realistic and increasingly unrealistic artificial vascular networks that are then reviewed by ophthalmologists via digital survey. The ophthalmologists are asked whether these artificial vascular networks appeared realistic or unrealistic.

Results: With only three rules (initiate, branch, and curve), the grammar accomplishes these goals. Networks are generated by adding noise to rule parameters present in existing networks. Via the survey of synthetic networks generated with different noise parameters, a correlation between noise in the branch rule and realistic association is revealed.

Conclusions: By creating a language to describe retinal vasculature, this article allows for the potential of new insight into such an important but less understood feature of the retina, which in the future may play a role in diagnosing or helping to predict types of ocular disease.

Translational Relevance: Applying shape grammar to describe retinal vasculature permits new understanding, which in turn provides the potential for new diagnostic tools.

Introduction

Like most vascular systems in the body, the retinal vasculature develops to metabolically sustain cells. By providing nutrients to the inner part of the retina, the retinal vasculature helps to maintain cell viability that allows humans to see.^{1,2} As health issues arise, the shape of the vasculature may change, which can inhibit the delivery of nutrients to different regions of the eye. Fortunately, a quantitative approach for measuring shape has the opportunity to provide significant advances.³ One study with respect to vasculature, has quantified the increase in retinal blood flow within patients with diabetes mellitus and found a

significant increase in blood flow within the retina, though flow is not measured in typical retinal fundus images. These changes in blood flow preceded any rupturing of blood vessels or significant biomarkers. This indicates that there is the potential for vasculature to change due to disease before rupturing.⁴ One vascular related disease is diabetic retinopathy (DR). Doctors perform regular screenings to catch the disease before it permanently affects vision, which could lead to blindness if it goes untreated. Standard screening programs capture retinal fundus images, and ophthalmologist experts visually inspect images for leaking blood vessels, microaneurysms, retinal swelling, hemorrhages, cotton wool spots, exudates, retinal ischemia, or neovascularization before diagnosing a patient with




Figure 1. Retinal fundus image of DR patient from STARE dataset with unhealthy features highlighted.¹

DR.⁵ Several of these features are shown in Figure 1. Developing a better way to describe other subtle changes in vasculature could help with training for identification of important vascular cues of DR and allow for diseases to be detected *before* irreversible damage occurs to the retina.

As machine learning techniques have become more powerful, more researchers have attempted to use these tools to automate disease diagnosis in the retina.^{6,7} With the large number of images that have been recorded to train these algorithms, correlations between the images and other health factors have been explored. One study attempted to predict the systolic blood pressure from the fundus images and used the machine learning tool saliency maps⁸ that highlight pixels of the image that contributed most to the prediction. Although 98% of doctors agreed that these maps highlighted the blood vessels, the qualities of the blood vessels that conveyed the information was not known.⁹ This further supports the concept that vascular changes occur, but current methods could be significantly improved with language to observe and describe these differences.

Isolating vasculature from the other features within the retinal fundus image forces viewers to focus on the vasculature and consider the features within it. Through generating realistic but synthetic images of retinal vasculature, this work seeks to provide more data to train ophthalmologists and machine learning algorithms to process retinas with a purely vascular approach. Although separating the components of tissue where symptoms could exist may seem counter-productive, it is precisely this separation that may allow a fresh look at an old problem. Currently, machine learning methods are trained on the entire fundus, which can include vascular abnormalities and retinal abnormalities. The causal relationship between these abnormalities is not always clear, and with

certain diseases, removing retinal information may hurt diagnosis. With other diseases, this simplification may prove helpful. Isolating retinal vasculature could not only lead to disease diagnosis at a subclinical level but also be developed into a new training model for clinicians.

Generating medical images can also be useful for training doctors, validating image analysis techniques, and producing massive amounts of data needed for machine learning methods. Simulating the generation of blood vessels has been used for surgical training with simulation incorporated so doctors can more accurately experience what would happen when taking certain actions.¹⁰ Additionally, with the increasing demand for large amounts of data from machine learning methods that are not always available, synthetically produced data can be created to help train networks to diagnose patients for various diseases.¹¹ Although Costa et al.¹¹ used adversarial neural networks to produce realistic and new retinal fundus images from noise, their approach does not explicitly track how features vary across each of the images, which is important in determining disease diagnosis and training. These features are all implicitly learned by the neural networks.

Generating retinal images using shape grammars, which are described in the next paragraph, offers the potential to create new images that can be used for medical training, where important nuances in image visualization can alter the diagnosis via control over what types of images that are generated. One way this can be accomplished is through adding noise to rules from a shape grammar of existing images through parametric variation. This could produce unique images of a similar style as the original.

Shape grammars are a field originally introduced in architecture by Stiny and Gips,¹² where shapes are identified and then modified on the basis of the application of rules, through successive iteration, to change and build up an overall shape design.^{13,14} A shape grammar consists of a set of shapes S and a set of rules R , where the rules can act on the shapes S to generate new shapes S' , where S' is also known as a language:

$$r : s \rightarrow s',$$

where

$$r \in R, s \in S, s' \in S'.$$

The set of rules can continue to be applied to the shapes, until there are no more valid rules or a termination rule is applied. Shapes can also be parametric, resulting in a succinct grammar representing a potentially infinite shape space, such as the one presented in this article.

Such shape grammars have been shown to be able to capture particular styles, such as villas in the style of Palladio¹⁵ and the prairie houses of Frank Lloyd Wright.¹⁶ In these examples, the shape grammars were used to recreate existing and generate new designs that are of the same style as captured in the grammar language. Cagan and colleagues introduced shape grammars to the engineering design community, and they demonstrated the ability of shape grammars to capture and generate products representative of different brands, including coffee makers,¹⁷ motorcycle designs like Harley-Davidson brands,¹⁸ or cars in the style of Buick,¹⁹ among others.

The goals of applying shape grammars to biologic design are similar to the goals of grammars in architecture and mechanical design. Biologic systems contain patterns and features with some commonalities and some discrepancies. All blood vessel networks deliver nutrients to tissues within the body, but the specific paths of blood vessels in each human are unique. Retinal vasculature also contains common features as well, such as blood vessels coming from the optic nerve and distributing nutrients, but the patterns and shapes are unique for each person and change with time.

As with shape grammars for coffee makers and many other products, which leverage parameters to model functional, as well as varied form, designs, a retinal vascular grammar must also capture common functional qualities, like branching and curvature, but also enable varied generation of the shape of the vascular system. Furthermore, most design capture with shape grammar begins with an initial shape from which the overall design is built, such as the fireplace for the Frank Lloyd Wright prairie houses or the wheelbase of motorcycles. So, too, must a retinal vascular grammar begin with the optic nerve from which all vessels flow.

This article explores how a shape grammar is applied to retinal vasculature by considering only a three-rule grammar. Describing the vasculature and generating new vasculature show the beginning of the possibilities of applying shape grammars for medical imagery. Retinal diseases like diabetic retinopathy and retinopathy of prematurity cause symptoms in the retina due to retinal vascular changes.² Identifying changes in the retinal vasculature before retinal disease is clinically apparent affords an avenue of diagnostic capability that could disrupt clinical ophthalmology. They offer an alternative approach more explainable than existing machine learning diagnostic techniques, which implicitly learn only patterns from their training data.²⁰

The goals of this article are to create a framework based on shape grammars to better understand blood vessel networks and explain patterns that are

present. This is done through deriving a simple three-rule grammar and using it to break down existing vascular networks and generate new networks through reapplying these rules with noise. Shape grammars offer the potential to accurately describe features of the vasculature and allow comparison of these features between patients. With image processing techniques, the shape grammar approach could be used to identify issues within the vasculature. These findings may be useful for the detection of ocular disease caused by vasculature and the effects of primary retinal disease on vasculature.

Methods

Image Processing

The images used in the project are from the Structured Analysis of the REtina (STARE) dataset.¹ The STARE dataset was created to help with automated diagnosis of diseases in the human eye. The dataset consists of 402 raw images along with diagnoses for each image. Of the 402 raw images, 20 have blood vessel segmentation processing done by hand. Of the 20 images, only the images labeled as healthy ($n = 9$) or diagnosed with DR ($n = 2$) are further processed. This dataset is chosen because of the availability of expertly segmented images, the disease labeled data, and because it is publicly available.

To analyze the vasculature, the raw image must be processed. Figure 2 outlines the four steps described in this section: (1) segment blood vessels, (2) create edge list, (3) separate networks, and (4) analyze the networks using shape grammars. The image processing and grammar implementation throughout this project used Python with the NumPy, SciPy, and Matplotlib packages.^{21,22}

The first step of preparing the images for branching analysis involves identifying which pixels in the image contain blood vessels. In the STARE dataset, this step has already been performed by an expert, although there is the potential to automate this step by using combinations of filters and deep neural networks in the future.^{23–25} The segmented image is a binary image that contains pixels with a value of 1 if it contains any part of a blood vessel, and 0 otherwise.

After the segmentation, the binary image is broken down into edges. The segmented image is skeletonized to a 1-pixel-wide representation of each blood vessel by iteratively sweeping over an image and removing edge pixels while preserving connectivity. This process repeats until a stable skeleton is found.²⁶ From the skeletonized image, a 3-pixel-by-3-pixel mask is passed

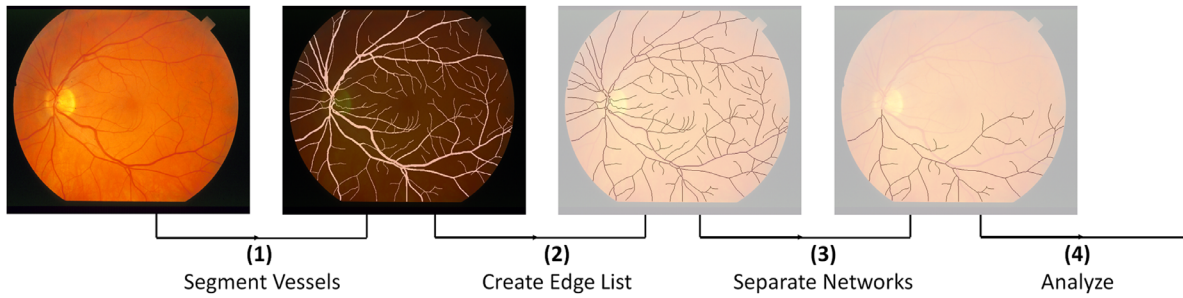


Figure 2. Image processing steps to prepare blood vessel networks for analysis. (1) First, blood vessels are segmented and separated from the rest of the image. (2) Then the segmented image is processed into a graph representation. (3) The vein and artery networks are separated. (4) Further processing is performed to apply the shape grammar.




Figure 3. Example of skeletonized network and sliding mask used to determine point types. This mask computes the sum of pixels within it to classify the points.

over the skeletonized image to identify end points and bifurcation points (Fig. 3).²⁷ From each bifurcation point, every connecting point is linked to a subsequent point using a recursive algorithm. Each list of points between two noninterior points becomes an edge.

With the edge list representation, the two primary networks (veins and arteries) often cross, creating loops in the graph representation. These connections do not exist in the vasculature but appear because the retina’s multiple layers are projected to a 2-D image. Due to the processing of the image, it is difficult to distinguish between the veins and arteries by looking at the skeletonized image. Even with the original retinal fundus image, it is challenging to identify whether smaller vessels are part of the vein or artery network, as illustrated in Figure 4. In this work, a human manually separates the networks from each other when they overlap, but this could be automated in the future as well.²⁸

With the edges of each network now identified, cubic b-splines with a smoothing condition are used to create fits for each edge, including all corresponding pixels as control points.²⁹ The curvature function describing the spline of an edge is stored independently




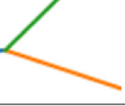




Figure 4. Veins and arteries can be difficult to identify before processing and after processing. *Left:* Raw image of artery crossing vein. *Right:* Same crossing after image processing.

from the start and end points of the edge. This allows a curvature function to be applied to any edge where it will be translated, rotated, and stretched to connect the points.

The processing results in data for each of the edges: parent, children, start point, end point, and path function. The parent is the edge that is hierarchically closer to the root, children are the edges that are hierarchically further away from root than the current edge, the start point and end point are coordinates for where the edge starts and ends, respectively, and the path function stores the parametric function for the spline describing the edge. With this information, the shape grammar is implemented.

Table 1. Description of the Three Rules Used in the Vascular Shape Grammar

Rule	Description	Before	After
Initiate	Creates a straight segment from the optic nerve, represented as a point.		
Branch	From a straight segment, creates two new child segments, using relative angles and proportion of original length for both children.		
Curve	Adds curvature to a straight segment by transforming a curve between the original endpoints.		

The Retinal Vascular Shape Grammar

There are many possible shape grammars that can describe a retinal network. A better grammar is one that can be used to describe many networks within the same class, but not those outside of the class, with a succinct set of rules. Each of the venous and arterial networks within the retina is processed separately. The grammar used in this work consists of three simple rules described in [Table 1](#): initiate, branch, and curve. Each rule is applied on a feature and then returns a new feature to replace the previous. The initiate rule creates the first edge of the network from a given point, representing where the blood vessel exits the optic nerve. Branch replaces a single edge with a shorter edge and adds two new child edges. Limiting the branching to only bifurcations more accurately mimics the physiology of retinal networks. Curve uses a curvature function, which contains the spline information independent of the endpoints. The edge is modified by applying the curvature function between the two points. Additionally, the branch rule and curve rule can only be applied to edges with no curvature. This implies no subsequent rules can be applied to a curved edge and removes the need for a separate terminate rule in the grammar. When applying these rules to generate retinal networks, typically a sequence is followed. First initiate is applied on a seed point. Then the branch rule is applied on successive edges. Once sufficient branching is complete, the curve rule is applied to each edge, thus ending potential rule application. Of note the curve rule could be applied at any time during generation; however, if applied to a leaf branch, no additional branching could take place. [Figure 5](#) demonstrates these rules can be used to create a retinal network.

Because only certain combinations of parameters used with these rules can generate retinal networks, existing retinal networks are first deconstructed by applying the rules in reverse ([Fig. 6](#)). Through the reverse application of the rules, parameters of each rule's application are recorded. After all the rules have been applied in reverse to a blood vessel network, no edges remain.

To use the shape grammar to generate new retinal blood vessel networks, the rules are run in the forward direction with noise added to each of the parameters to generate similar, but varied, networks ([Fig. 6](#)). See Supplementary Movie S1 for an animated example of how the rules can be applied in reverse and reapplied. Noise can be added to the proportion of the edge that is used in the branching, the branching angles of both children, and the curvature function. In our application the topology of the original network is maintained, while allowing for a diverse set of new networks that can be generated. Adding noise to branching results in changing the angle of the children edges from a branch and the subsequent children that follow. Adding length noise can change the length of a branch, making it shorter or longer between each junction. When adding noise, the mean of the distribution is centered at zero, so the average distance between junctions is approximately unchanged. Similarly, when adding noise to angles, the noise is not biased to wider angles or narrower angles.

Using this technique, many networks can be generated from a single realistic network. These networks can be combined with other networks from the same image to produce the vasculature for an entire eye.

Survey of Ophthalmologists

To analyze the relationship between the retinal shape grammar with respect to changing the vascular pattern and its pattern relative to ocular disease, an online survey was distributed to ophthalmologists within the World Society of Paediatric Ophthalmology and Strabismus who have experience analyzing retinal fundus images. The survey began with an explanation and example of how the vasculature is extracted from the retinal fundus image. The survey consisted of a yes/no question, "Does this image appear realistic?" repeated 15 times each with different images. The questions were ordered randomly. The 15 images consisted of three unaltered original images and four perturbations of these images. The perturbations had varying amounts of angle and length noise as shown in [Table 2](#).

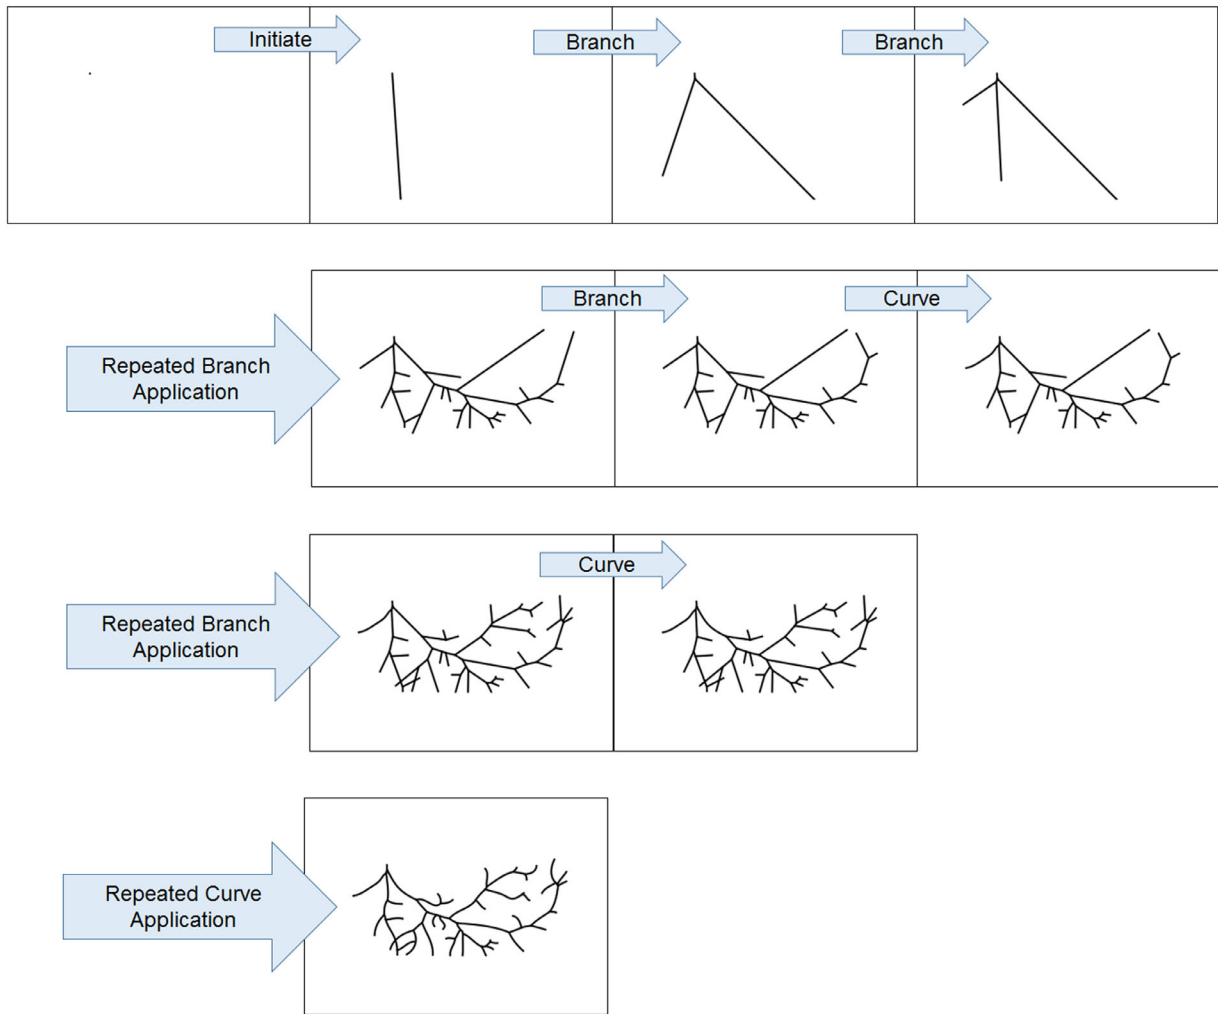


Figure 5. Example of forward rule application to generate a retinal network.

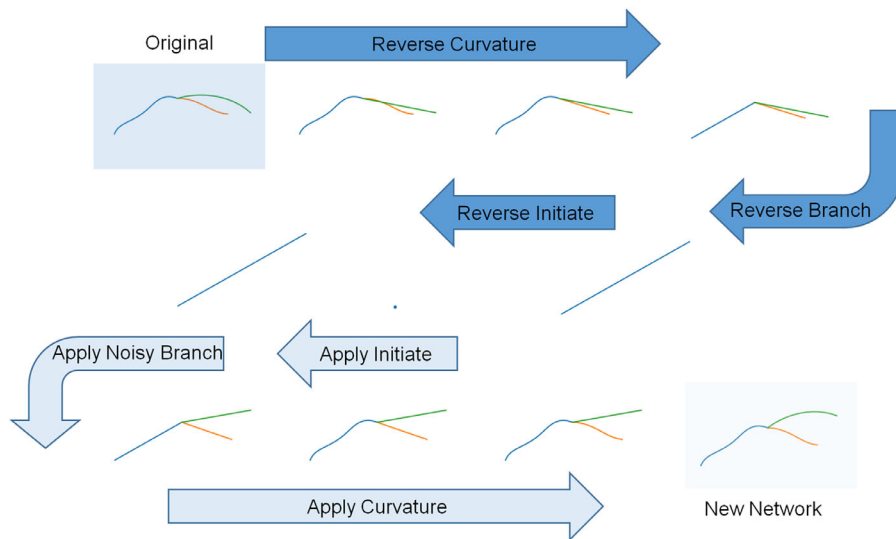


Figure 6. Example of how rules can be reversed and then applied to generate networks. First the curve rule is applied in reverse, then the branch rule, followed by initiate. The opposite sequence is followed in the forward direction.

Table 2. Image Names and Corresponding Noise Levels That Were Used in the Survey

Root Image Name	Noise Levels (Length Noise Proportion, Angle Noise Degrees)				
8c	(0, 0)	(0.05, 5)	(0.02, 8)	(0.08, 2)	(0.15, 15)
8d	(0, 0)	(0.05, 5)	(0.02, 8)	(0.08, 2)	(0.15, 15)
8e	(0, 0)	(0.05, 5)	(0.02, 8)	(0.08, 2)	(0.15, 15)

Each noise combination listed here corresponds to an image shown in the survey, a total of 15 images.

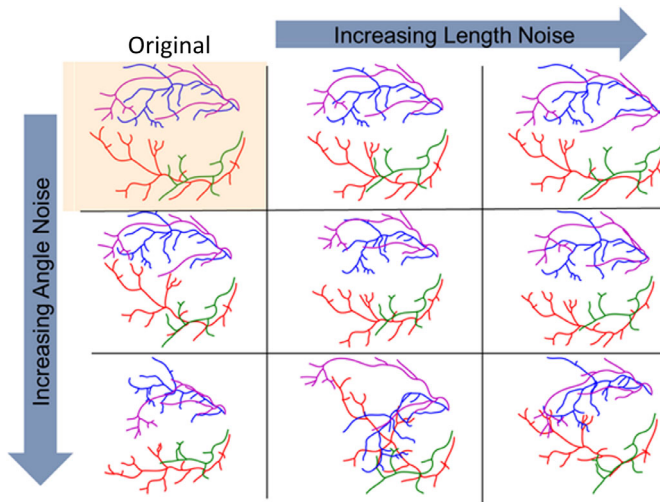


Figure 7. Length and angle noise are applied to the original vasculature of the retina through the reapplication of rules to create a variety of new retinal vasculature.

Results

Sample Images Generated

Figure 7 demonstrates the results of applying the reconstruction and deconstruction technique to generate new networks. As noise increases, the similarity to the original decreases. In Figure 7, the horizontal axis from left to right is increased in the length parameter for branching. The vertical axis from top to bottom shows the noise is increased for the angle parameter in branching.

A total of 95 people from the pediatric ophthalmology community completed the survey. The results are shown in Figure 8. See Supplementary Document S2 for the images used in the survey.

Discussion

Figure 8 shows how adding noise overall decreased the number of respondents who saw an image as being realistic. Realism dropped off as noise increased in both angle and length dimensions as observed

across all three original images. Despite this overall trend, there are differences in how this trend is established. Figure 8c has a strong realistic association with the original image with zero noise. There is a clear curve of the primary vessels around the area where the fovea is expected to be. The fovea contains many photoreceptor cells and much of the vasculature feeds it, though it is clearly not visible on the extracted vasculature patterns. In Figure 8c, many of the smaller vessels reach toward this foveal point. The others had less of a realistic association with the original, but with all base images there is a clear drop off in realistic association as noise increases. Figure 8d in the original form was ranked less realistic than Figure 8e in the highest noise case. This could be because Figure 8d contains fewer common patterns. The main blood vessels do not appear to curve around the putative fovea, and the large number of smaller vessels accompanied by frequent branching make it appear less typical. It is also important to note that the noise applied in each image has an element of randomness associated with it.

Creating a new image, using the same noise parameters used in the survey, can produce a significantly different image. Figures 8d and 8e had no change in realistic association from the no-noise original to the images with (0.05, 5) noise. This could indicate that noise levels below a certain threshold produce images that still appear realistic. With the noise levels of (0.02, 8) and (0.08, 2), the angle noise and the length noise, respectively, cross over this threshold, and the proportions of doctors that see the image as realistic decreases.

The image with weakest realistic association is Figure 8d with (0.02, 8) noise (Fig. 9). In this image there is no clear curvature around the putative fovea, the green network reaches through where the putative fovea should be and contains excessive branching. In Figure 8e with (0.15, 15) noise (Fig. 10), 30 of 95 doctors responded stating the image appears realistic. This image contains some curvature around the putative fovea with some smaller branches extended toward the putative fovea, but the red network also extends away from the center. Together these features may lead to the overall weak realistic association.

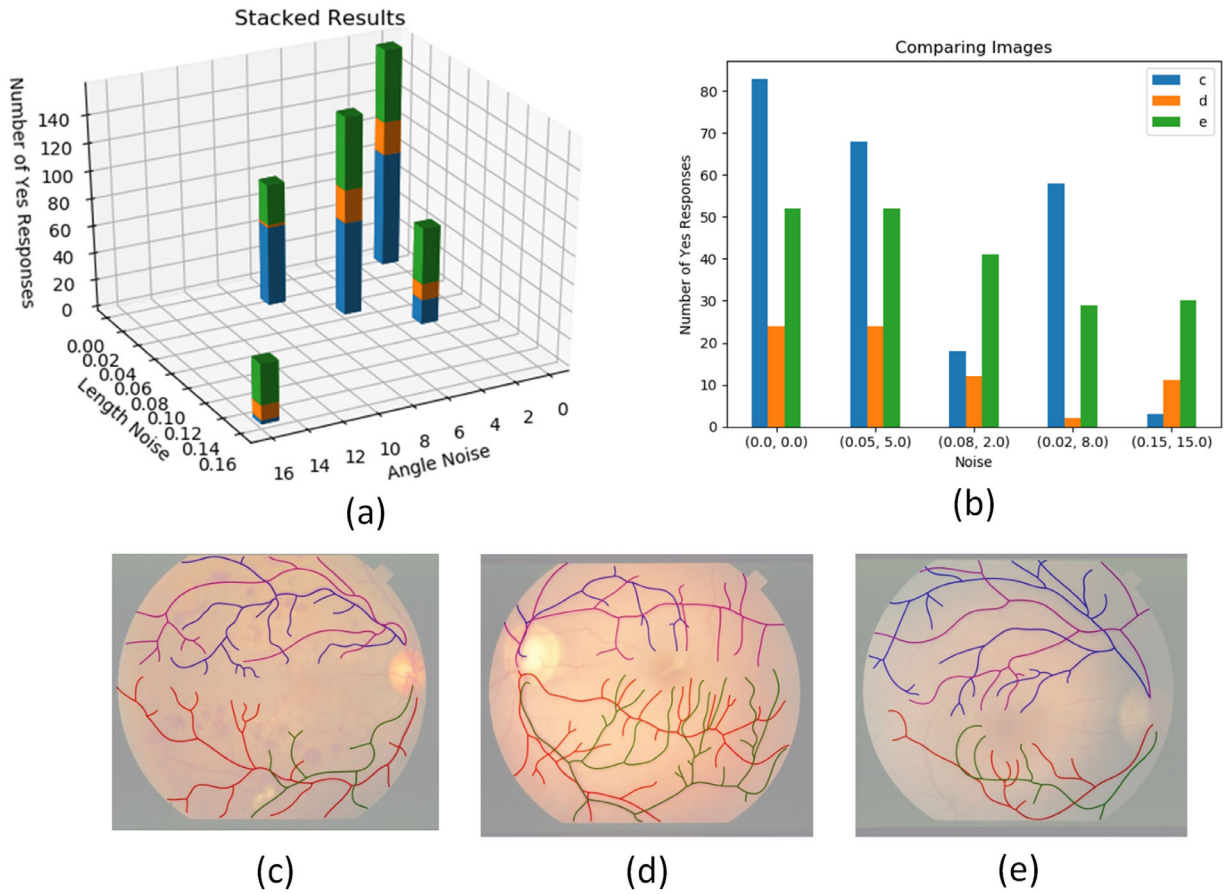


Figure 8. Survey results shown with noise as a 3D bar plot of stacked responses (a) and a 2D bar plot (b). The original images before adding noise are shown below (c), (d), (e).




Figure 9. Figure 8d with (0.02, 8) noise is the image with the weakest realistic association. There is no clear curvature around the putative fovea with a network crossing over where the fovea should be.




Figure 10. Figure 8e with (0.15, 15) noise. This image had 32% of respondents describe it as realistic. There appears to be some curvature around the putative fovea, but also a network extending away from the center.

When processing the images, there currently are limitations that lead to inaccuracies. Due to the methods of obtaining retinal fundus images some blood vessels extended beyond the image frame. In the grammar representation, these appeared like terminating blood vessels, while realistically they may continue to spread and branch many more levels. Additionally, because of the manual steps involved in processing the images, there is a lack of uniformity in how the images are processed. To automate this process in a consistent way, more work needs to be done on automatically identifying blood vessels and discriminating between arteries and veins.

Although there are issues that prevent the vasculature from being captured accurately in its entirety, our shape grammar approach accomplishes many unique functions. The grammar can accurately break down each of the networks from the retinal images and reassemble them by applying the rules in the forward direction to produce the same original image. This demonstrates the grammar is flexible enough to describe retinal vascular networks. Without limitations on the parameters, this grammar is not constrained to creating only retinal vasculatures and could serve as a grammar for any vascular system given an initial point. The current method to control the grammar to produce only retinal vasculature is to start with a real network and constrain the amount of noise applied to it. To create a more general grammar, existing relationships between parameters applied at each step would need to be determined. This could be done through using more complex rules or understanding constraints to which the vasculature is subjected. An example complex rule could be to turn an edge into a long edge with several terminated branches extending from it. An example of a constraint could be not allowing branching where blood vessels from the same network would overlap. Whiting et al.³⁰ have investigated automating rule induction to identify higher level rules in their work, and a similar approach could be applied here. With these relationships of more precisely determined parameters, the options within the grammar could be reduced to contain only networks that appear as viable retinal vasculature and not be limited to being based off networks in the dataset.

Through this shape grammar approach, realistic images of only vasculature have been produced through applying noise. These vascular images could then be used to train machine learning algorithms to use vasculature to identify diseases and train ophthalmologists to recognize retinal vascular anomalies better. By using only the vasculature for disease identification, there is also the poten-

tial to find subclinical or preclinical retinal disease because of the changes that occur in the blood vessel structure.

Conclusion

This article introduces a new method for synthetic generation of varied yet realistic vascular networks. Parametric shape grammars succinctly model vascular systems resulting in a generative method to create new and varied networks. The three-rule shape grammar described in this work can accurately describe retinal blood vessel networks. The flexibility and usefulness of shape grammars are demonstrated by generating similar but unique blood vessel networks through the application of noise through parameter variation, making new networks of the same style as the original network. Feedback from ophthalmologists through a survey illustrates that there is a limit where, when too much noise is added, the networks no longer have a realistic form. Through applying shape grammars to retinal vasculature, this work could provide a basis to uncover a fundamental theory to understand vascular systems in the context of ocular disease.

Acknowledgments

The authors thank the participation of the members of the World Society of Paediatric Ophthalmology and Strabismus (www.wspos.org) Regd. 1144806 Charity Commission UK.

Supported by the Office of Naval Research (N00014-17-1-2566), the Air Force Office of Scientific Research (FA9550-18-1-0262), the National Institute of Health (5-R01-AG061005), the National Science Foundation (CMMI-1946456), and the Pennsylvania Department of Health (AP4100077084).

Disclosure: **R.Y. Yeh**, None; **K.K. Nischal**, None; **P. LeDuc**, None; **J. Cagan**, None

References

1. Hoover A, Kouznetsova V, Goldbaum M. Locating blood vessels in retinal images by piece-wise threshold probing of a matched filter response Department of Ophthalmology University of California, San Diego. *IEEE Trans Med Imaging*. 2000;19:203–210.

2. Selvam S, Kumar T, Fruttiger M. Retinal vasculature development in health and disease. *Prog Retin Eye Res.* 2018;63:1–19.
3. Safi H, Safi S, Hafezi-Moghadam A, Ahmadi H. Early detection of diabetic retinopathy. *Surv Ophthalmol.* 2018;63:601–608.
4. Burgansky-Eliash Z, Barak A, Barash H, et al. Increased retinal blood flow velocity in patients with early diabetes mellitus. *Retina.* 2012;32:112–119.
5. Bhaskaranand M, Ramachandra C, Bhat S, et al. Automated diabetic retinopathy screening and monitoring using retinal fundus image analysis. *J Diabetes Sci Technol.* 2016;10:254–261.
6. Gardner GG, Keating D, Williamson TH, Elliott AT. Automatic detection of diabetic retinopathy using an artificial neural network: a screening tool. *Investig Ophthalmol Vis Sci.* 1996;37:940–944.
7. Gulshan V, Peng L, Coram M, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA.* 2016;316:2402–2410.
8. Simonyan K, Vedaldi A, Zisserman A. Deep inside convolutional networks: visualising image classification models and saliency maps. arXiv:1312.6034 [cs.CV] 2014, <http://arxiv.org/abs/1312.6034>.
9. Poplin R, Varadarajan AV, Blumer K, et al. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nat Biomed Eng.* 2018;2:158–164.
10. Liu X, Liu H, Hao A, Zhao Q. Simulation of blood vessels for surgery simulators. *2010 Int Conf Mach Vis Human-Machine Interface, MVHI 2010.* 2010:377–380.
11. Costa P, Galdran A, Meyer MI, et al. End-to-end adversarial retinal image synthesis. *IEEE Trans Med Imaging.* 2018;37:781–791.
12. Stiny G, Gips J. Shape grammars and the generative specification of painting and sculpture. *Proc Inf Process Congr.* 1972;71:125–135.
13. Stiny G. An introduction to shape and shape grammars. *Environ and Plan B Plan Des.* 1980;7:343–351.
14. Stiny G. *Shape: Talking about Seeing and Doing.* Cambridge, MA: MIT Press; 2006.
15. Stiny G, Mitchell WJ. The Palladian grammar. *Environ Plan B Plan Des.* 1978;5:5–18.
16. Koning H, Eizenberg J. The language of the prairie: Frank Lloyd Wright's prairie houses. *Environ Plan B Plan Des.* 1981;8:295–323.
17. Agarwal M, Cagan J. A blend of different tastes: the language of coffee makers. *Environ Plan B Plan Des.* 1998;25:205–226.
18. Pugliese MJ, Cagan J. Capturing a rebel: modeling the Harley-Davidson brand through a motorcycle shape grammar. *Res Eng Des.* 2002;13:139–156.
19. McCormack JP, Cagan J, Vogel CM. Speaking the Buick language: capturing, understanding, and exploring brand identity with shape grammars. *Des Stud.* 2004;25:1–29.
20. Wainberg M, Merico D, DeLong A, Frey BJ. Deep learning in biomedicine. *Nat Biotechnol.* 2018;36:829–838.
21. Oliphant T, Millma Jk. A guide to NumPy. New York: Trelgol Publishing; 2006.
22. Hunter JD. Matplotlib: A 2D Graphics Environment. *Comput Sci Eng.* 2007;9:90–95.
23. Fu H, Xu Y, Wong DWK, Liu J. Retinal vessel segmentation via deep learning network and fully-connected conditional random fields. *Proc - Int Symp Biomed Imaging.* 2016;698–701.
24. Ronneberger O, Fischer P, Brox T. U-net: convolutional networks for biomedical image segmentation. *Lect Notes Comput Sci.* 2015;9351:234–241.
25. Zhang B, Zhang L, Zhang L, Karray F. Retinal vessel extraction by matched filter with first-order derivative of Gaussian. *Comput Biol Med.* 2010;40:438–445.
26. Lee T-C, Kashyap RL, Chu C-N. Building skeleton models via 3-D medial surface/axis thinning algorithms. *CVGIP Graph Model Image Process.* 1994;56:462–478.
27. Hichem G, Chouchene F, Belmabrouk H. 3D model reconstruction of blood vessels in the retina with tubular structure. *Int J Electr Eng Informatics.* 2015;7:724–734.
28. Martinez-Perez ME, Hughes AD, Stanton AV, et al. Retinal vascular tree morphology: a semi-automatic quantification. *IEEE Trans Biomed Eng.* 2002;49:912–917.
29. Dierckx P. Algorithms for smoothing data with periodic and parametric splines. *Comput Graph Image Process.* 1982;20:171–184.
30. Whiting ME, Cagan J, Leduc P. Efficient probabilistic grammar induction for design. *Artif Intell Eng Des Anal Manuf AIEDAM.* 2018;32:177–188.