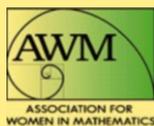


Association for Women in Mathematics Series

Ami Radunskaya · Rebecca Segal
Blerta Shtylla *Editors*

Understanding Complex Biological Systems with Mathematics



 Springer

Placental Vessel Extraction with Shearlets, Laplacian Eigenmaps, and a Conditional Generative Adversarial Network



Catalina Anghel, Kellie Archer, Jen-Mei Chang, Amy Cochran, Anca Radulescu, Carolyn M. Salafia, Rebecca Turner, Karamatou Yacoubou Djima, and Lan Zhong

Abstract The placenta is the key organ of maternal–fetal interactions, where nutrient, oxygen, and waste transfer take place. Differences in the morphology of the placental chorionic surface vascular network (PCSVN) have been associated with

C. Anghel (✉)
University of California Davis, Davis, CA, USA
e-mail: catalina.anghel@alum.utoronto.ca

K. Archer
The Ohio State University, Columbus, OH, USA
e-mail: archer.43@osu.edu

J.-M. Chang
California State University Long Beach, Long Beach, CA, USA
e-mail: jen-mei.chang@csulb.edu

A. Cochran
University of Wisconsin, Madison, WI, USA
e-mail: cochran4@wisc.edu

A. Radulescu
SUNY New Paltz, New Paltz, NY, USA
e-mail: radulesa@newpaltz.edu

C. M. Salafia
Placental Analytics, LLC., New Rochelle, NY, USA

R. Turner
The University of Auckland, Auckland, New Zealand
e-mail: rm.turner@auckland.ac.nz

K. Y. Djima
Amherst College, Amherst, MA, USA
e-mail: kyacouboudjima@amherst.edu

L. Zhong
University of Delaware, Newark, DE, USA
e-mail: lanzhong@udel.edu

© The Author(s) and the Association for Women in Mathematics 2018
A. Radunskaya et al. (eds.), *Understanding Complex Biological Systems with Mathematics*, Association for Women in Mathematics Series 14,
https://doi.org/10.1007/978-3-319-98083-6_8

171

developmental disorders such as autism, hinting that the PCSVN could potentially serve as a biomarker for early diagnosis and treatment of autism. Studying PCSVN features in large cohorts requires a reliable and automated mechanism to extract the vascular networks. This paper presents two distinct methods for PCSVN enhancement and extraction. Our first algorithm, which builds upon a directional multiscale mathematical framework based on a combination of shearlets and Laplacian eigenmaps, is able to intensify the appearance of vessels with high success in high-contrast images such as those produced in CT scans. Our second algorithm, which applies a conditional generative adversarial neural network (cGAN), was trained to simulate a human-traced PCSVN given a digital photograph of the placental chorionic surface. This method surpasses any existing automated PCSVN extraction methods reported on digital photographs of placentas. We hypothesize that a suitable combination of the two methods could further improve PCSVN extraction results and should be studied in the future.

Keywords Placenta · Autism · Vascular networks · Shearlets · Wavelets · Laplacian eigenmaps · Neural networks · Deep learning · cGAN · Generative models

1 Introduction

The placenta has been the subject of increased medical research attention, as the variability in its structure has been shown to indicate pregnancy complications [12, 18, 22] and developmental delays such as autism [6, 32, 34]. For example, placental chorionic surface vascular networks (PCSVNs) of a high-autism risk cohort consisting of infants with a biological sibling with autism [29] generally exhibit fewer branching points than their low-autism risk counterparts [6]. Although it is known that autism is highly heritable and more prevalent in males, a host of genetic and environmental factors have been associated with the disorder, including advanced maternal and paternal age, gestational diabetes, maternal infection and nutritional deficiencies, exposure to pharmaceutical drugs or pesticides and hypoxic damage [27, 28, 30]. No conclusive single factor or cause is known.

Studies linking PCSVN features and autism necessarily relied on completely connected PCSVNs to compute networks' geometric characteristics such as vessel curvature, branching angles, lengths, and thickness. Currently, the only reliable way to extract PCSVNs is through a laborious tracing process that is done manually by a trained researcher on a computer [33]. This process typically takes 4–8 h with a few rounds of validations, prohibiting any large-scale studies with PCSVNs.

In 2013, Chang et al. established an algorithm to extract the placental chorionic surface vascular networks with a multiscale framework [5]. This method worked relatively well in identifying small vessels; however, it was unable to produce a completely connected network due to the glare caused by the placenta's irregular texture and shape and the significant interference caused by the villus trees near

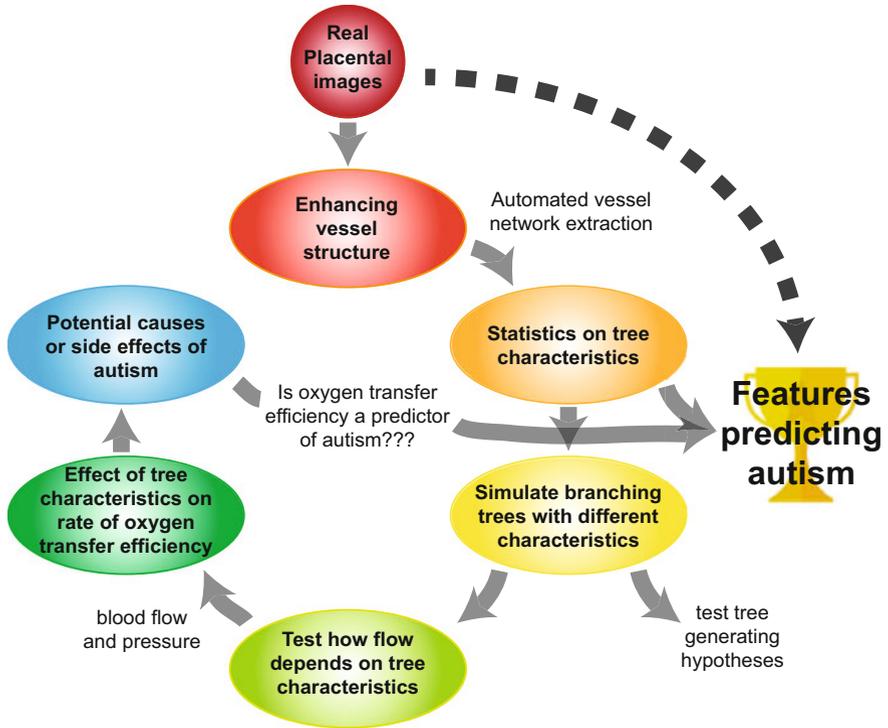


Fig. 1 Ongoing research goals

the edges of the placental chorionic surface. Even when photographs are taken with a polarized filter, the vessels are difficult to distinguish as the rough and irregular surface of the placental tissue is incorrectly identified as part of the vasculature. Creating high-contrast images of placenta ex-vivo to allow for an automated PCSVN extraction has been extremely challenging. In this paper, we present two distinct algorithms toward automating PCSVN enhancement and extraction. One algorithm combines a tool from harmonic analysis and signal processing with a technique from machine learning; the second algorithm is based on a particular type of neural network called a conditional generative adversarial network (cGAN).

This research is an essential part of a large research agenda to discover potential causes of autism through the structural and functional features exhibited in the PCSVNs, as shown in Fig. 1. Ultimately, we wish to decide whether an intervention for autism is needed upon delivery simply by taking a digital photograph of the PCSVN, as indicated by the dashed connection in Fig. 1. In order to accomplish this, we need to precisely identify which features from the PCSVNs are capable of differentiating placentas that are associated with autism from those in the general population. Extracting a detailed description of the vascular networks from digital photographs of the placenta, the purpose of this research project, serves as a crucial step in this ultimate research goal.

Due to the nature of our work, we chose to present the two methods in a parallel fashion where each section contains a complete treatment of the mathematical background, data sets used, parameter selections, and results. In particular, Sect. 2 discusses the work on the multiscale framework with shearlets and Laplacian eigenmaps, denoted by Shearlets-LE henceforth. Section 3 discusses the work with the conditional generative adversarial network. Section 4 concludes our work with a summary and future directions.

2 Vessel Enhancement Using Shearlets and Laplacian Eigenmaps

Wavelets are used to obtain a multiscale representation for square-integrable functions and signals via decomposition into elements that give both location and scale information at increasingly precise resolution [8, 10]. Since they came to interest in the 1980s, wavelets have generated a plethora of fascinating mathematical results and numerical applications such as detecting singularities and denoising. In numerical applications, using wavelets offers two great advantages: the existence of fast transforms as well as fast pyramidal schemes based on the introduction of the multiresolution analysis (MRA) by Y. Meyer and S. Mallat in 1989 [3, 26]. Another benefit of wavelet-representation is the flexibility in the choice of the function which generates the system. For example, the standard Haar wavelet, a basic step function, has advantages such as simplicity, orthogonality of the induced system, and compact support; however, it presents poor differentiability properties, which can cause severe errors in certain approximations. To compensate for this, other types of wavelets include characteristics such as exponential decay, smoothness, or directionality; the last property caters specifically to our present needs.

Placenta chorionic surface vascular network images, like most natural images and signals, exhibit curves or line-like edges, i.e. discontinuities along smooth regions, some of which exhibit the same structure at different scales. In Fig. 2, via visual inspection, one can distinguish vessels of different sizes (or scales) with the same elongated and at times tortuous shapes from other regions of the placenta. The sharper differences in color between these regions, or, equivalently, the high pixel intensity gradient along the boundary of these distinct regions, constitute discontinuities in the image. To clarify further, we use the term “discontinuity” not to refer to the vessels being disconnected, but to the discontinuities in the pixel intensity values.

In a previous work, Chang et al. [5] established an algorithm based on a multiscale Frangi filter [14], a method that is based on images’ second-order characteristics.

Directional methods such as curvelets, contourlets, shearlets, etc., also target exactly this type of applications [4, 11, 17, 23]. These wavelet methods optimally, i.e., sparsely, represent functions with certain geometric features using basis

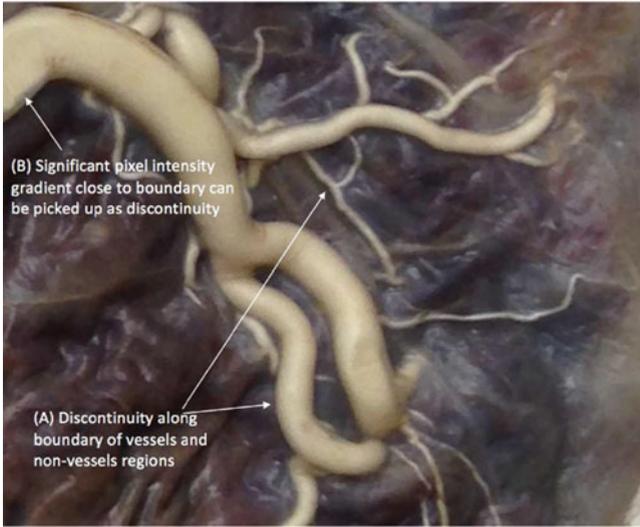


Fig. 2 Types of discontinuities in pixel intensities of the placenta image. Our method exploits strong differences in pixel values, i.e., the discontinuity along the boundary of vessel and non-vessel regions (a). However, a significant pixel intensity gradient close to the boundary can also be picked up as a discontinuity (b). That is, the boundaries of thicker vessels may be mistaken for vessels themselves

functions that include orientation in addition to location and scale information. In particular, as opposed to traditional wavelets, directional representations contain anisotropic elements occurring at all scales and locations and exploit the geometric regularity of edges. We thus expect directional shearlets to be an appropriate representation system for PCSVN images. In addition, shearlets-based algorithms are faster than Frangi-based filtering schemes, making them a superior choice for improved computational efficiency.

To further highlight the vessel structure after applying shearlets, we introduce a kernel method for manifold recovery called Laplacian Eigenmaps (LE) [2, 19, 25]. One reason for introducing this additional enhancement tool is that in shearlets and other similar methods, directionality is better captured when the images are smooth away from singularities along smooth curves, where the curves have bounded curvature. In the placenta images, we observed that certain non-vessel parts exhibit strong curve-like features or ripples due to glare and discoloration. Therefore, we cannot simply rely on directional methods to enhance the appearance of vessels. In addition, based on the work in [9, 35], we expect that different structures such as the difference between the interior or contours of the vessels will be enhanced by different vectors in the LE representation. These vectors of features will be referred to as eigenimages. Next, we give some background and algorithmic details on shearlets and LE.

2.1 Mathematical Background

2.1.1 Shearlets

Shearlet representations are obtained by applying parabolic dilations D_a , shearing, and translations to a wavelet function, ψ [17, 23]. The matrix D_a is a parabolic scaling, i.e., if $f(x, y) = \chi_{\{y \geq x^2\}}$,

$$f_a(x, y) = D_a f(x, y) = f(ax, a^2y),$$

as opposed to traditional scaling $f_a(x, y) = f(ax, ay)$ for classical wavelets. The parabolic scaling (for which the principal length² \sim width holds) allows for highly anisotropic elements at fine scales. Directionality is captured by a shearing operator S_m :

$$S_m = \begin{pmatrix} 1 & m \\ 0 & 1 \end{pmatrix}.$$

The variable m parametrizes the orientations with slopes rather than angles, as is the case for curvelets and contourlets [4, 11]. Hence, if m is an integer, the integer lattice is left invariant, a useful property in computations.

Now, suppose that $\psi \in L^2(\mathbb{R}^2)$ and $x \in \mathbb{R}^2$. A (regular) discrete shearlet system associated with a function ψ is the collection of dyadic parabolic dilations, and integer-shearing and -translations of the function ψ written as

$$\{\psi_{j,k,l}\} = \{2^{3j/4} \psi(S_k D_{2^j} x - l)\}_{j,k \in \mathbb{Z}, l \in \mathbb{Z}^2}, \quad (1)$$

where the coefficient $2^{3/4}$ is only needed for normalization purposes. Informally, a square integrable function f can be represented as a series of shearlets

$$f = \sum_{j,l,k} \langle f, \psi_{j,l,k} \rangle \psi_{j,l,k}, \quad (2)$$

where the mother wavelet, ψ , meets some admissibility conditions that we will not discuss here. In particular, ψ is chosen such that the collection $\{\psi_{j,l,k}\}$ forms a ‘‘tight frame’’ for $L^2(\mathbb{R}^2)$, i.e., if $f \in L^2(\mathbb{R}^2)$, then the equation

$$\|f\|_{L^2}^2 = \sum_{j,l,k} |\langle f, \psi_{j,l,k} \rangle|^2$$

holds. Hence, one can view a tight frame as an orthonormal basis, with the caveat that a tight frame may contain redundant elements.

Near edges, shearlets perform much better than a classical Fourier basis or a wavelet basis. While the k -term approximation error for these bases are, respec-

tively, $\|f - f_k\|_{L^2}^2 = O(k^{-1/2})$ and $\|f - f_k\|_{L^2}^2 = O(k^{-1})$ near edges [4], with shearlets, we obtain

$$\|f - f_k\|_{L^2}^2 = O\left(k^{-2} (\log k)^3\right)$$

within log factors of optimal rate k^{-2} [17, 23]. One major theoretical advantage of shearlets compared to other directional representation systems is that shearlets provide sparse approximation of anisotropic features while providing a unified treatment of the continuum and digital realm in the sense of allowing faithful implementation.

2.1.2 Laplacian Eigenmaps

Suppose that we are given the set of vectors $X = \{x_1, \dots, x_N\} \subset \mathbb{R}^D$, where D is large. The Laplacian Eigenmaps (LE) algorithm is a nonlinear, locality-preserving, dimensionality reduction algorithm that finds vectors $Y = \{y_1, \dots, y_N\} \subset \mathbb{R}^d$, where $d \ll D$, such that the important information, in particular, a certain notion of similarity between points in the original data X , is retained in Y . We assume that the data set X is sampled from a manifold embedded in \mathbb{R}^D . The steps of the LE algorithm are as follows [2].

Adjacency Matrix Construction

Given N vectors of dimension D sampled from a data set, we construct an adjacency graph G that represents the data with nodes (or vertices) formed by the N vectors, and with edges that represent the distances between the nodes that are defined to be “close.” In our work, the data vectors are each pixel of an image in three (RGB) dimensions; for example, for a 512×512 image, $N = 262,144$ and $D = 3$. There are two ways of establishing the notion of closeness for this neighborhood construction. In both cases, we start by computing the Euclidean distances between all pairs of points x_i and x_j , for $i, j = 1, \dots, N$. Then, one option is to define two nodes as connected by an edge if the Euclidean distance between them is less than some pre-defined, fixed ϵ . Although this option is geometrically intuitive, it can be difficult to choose an appropriate ϵ , and it often yields disconnected graphs. In our work, we use the second option: the k -nearest neighbors method. Here, the nodes are connected if x_i is among the k nearest neighbors of x_j or vice versa. The main advantage of using the k nearest neighbors scheme is that we have more control over the degree of connectivity of our graph. All of the information for the graph G is stored in an adjacency matrix \mathbf{A} :

$$\mathbf{A}_{ij} = \begin{cases} 1, & \text{if } x_j \text{ is in the } \epsilon\text{-neighborhood of } x_i, \\ 0, & \text{otherwise.} \end{cases}$$

where A is symmetric.

Heat Kernel as Weights

The adjacency matrix \mathbf{A} can be modified by assigning weights to the edges of the graphs. Given $\sigma > 0$, we add weights to the edges of the graphs using the heat kernel as follows:

$$w_{i,j} = \begin{cases} e^{-\frac{\|x_i - x_j\|_2^2}{\sigma}} & \text{if } i \text{ and } j \text{ are connected,} \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Now, note that by using this type of weights, we would need to determine the appropriate σ . Although there are simpler alternatives (e.g., set $w_{i,j} = 1$ if x_i and x_j are connected by an edge and $w_{i,j} = 0$ otherwise) that avoid choosing σ , the heat kernel is a better option from a geometrical point of view, as it preserves all information and better encodes relative closeness between points. The Laplace operator that we define next, based on these weights, is analogous to the Laplace Beltrami operator on manifolds, whose eigenfunctions have properties that are desirable for embedding. Interested readers are referred to the original paper [2] for a complete discussion on the geometric implications of LE as well as the relation between heat flow and the Laplace Beltrami operator on manifolds.

Minimization Problem

Given the weight matrix $\mathbf{W} = [w_{i,j}]$, where $w_{i,j}$'s are the heat kernel weights defined above, we can set up a minimization problem that will allow us to find a collection of embedding vectors \mathbf{Y} that can be used to better represent the original observations. Consider the $d \times N$ matrix $\mathbf{Y}^T = [y_1, y_2, \dots, y_N]$, where y_i is a column vector that gives a d -dimensional representation of the i^{th} observation (node). Define an $N \times N$ diagonal matrix \mathbf{D} with components $d_{i,i}$ such that $d_{i,i} = \sum_j w_{i,j}$ and 0 otherwise. Each component $d_{i,i}$ that is associated with the i^{th} node gives a measure of how supported this node is, i.e., a large $d_{i,i}$ reflects the fact that the i^{th} node is strongly connected to other nodes in the graph. This could be because that there are many non-negligible weights associated with this node or fewer, but large weights. Now, to find \mathbf{Y} , we solve the minimization problem

$$\underset{\mathbf{Y}^T \mathbf{D} \mathbf{Y} = \mathbf{I}}{\operatorname{argmin}} \frac{1}{2} \sum_{i,j} \|y_i - y_j\|_2^2 w_{i,j} = \underset{\mathbf{Y}^T \mathbf{D} \mathbf{Y} = \mathbf{I}}{\operatorname{argmin}} \operatorname{trace}(\mathbf{Y}^T \mathbf{L} \mathbf{Y}), \quad (4)$$

where $\mathbf{L} = \mathbf{D} - \mathbf{W}$ is an $N \times N$ Laplace operator and \mathbf{I} is a $d \times d$ identity matrix.

Now, assume that the graph G is connected. This is a safe assumption as using the k nearest neighbors algorithm would ensure that. The Laplace operator is a symmetric, positive semidefinite matrix that can be thought of as an operator on

functions defined on vertices of G . Then, the constrained minimization problem is solved as follows:

- First, find the eigenvalues and eigenvectors solutions of the generalized eigenvalue problem

$$\mathbf{L}\xi_l = \lambda_l \mathbf{D}\xi_l, \quad l = 0, 1, \dots, N - 1. \quad (5)$$

- Order the eigenvalues from the lowest to the highest, $0 = \lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_{N-1}$.
- Ignore the eigenvectors corresponding to the zero eigenvalue and use the next d eigenvectors for embedding in the d -dimensional Euclidean space by setting the representation for x_i to

$$y_i = [\xi_1(i), \xi_2(i), \dots, \xi_d(i)]. \quad (6)$$

The justification for eliminating the zero eigenvalue is that since the graph is connected, the vector of all ones is the only eigenvector associated with this eigenvalue. For a formal justification of the above steps, please refer to [7].

2.2 Implementation

2.2.1 Data Sets

There were two types of image data available for our implementations. The first is a collection (201 images) of digital photographs of raw, formalin-fixed placental fetal surface images from the National Children's Study (used in Sects. 3.2.1 and 4.1). Secondly, we had access to a smaller collection (5 images) of digital photographs of PCSVN where the placental fetal vasculature was heparin flushed and perfused with a 1% agarose 30% barium sulfate solution. Within the barium-perfused data set, we have images of raw placenta ex-vivo without any alteration as well as images of formalin-fixed placentas. The purpose of injecting PCSVN with barium was to enhance the image contrast, thereby mitigating some of the imaging challenges that were native to the digital photographs of raw, formalin-fixed placentas. We focus our work on the Shearlets-LE algorithm to the barium-perfused PCSVN images in this section.

Figure 3 shows raw and formalin-fixed images from the same placenta. The rationale behind examining both raw and fixed images is that when we look for nodal connections along vessel paths each type of image presents some advantages and disadvantages. For example, fixed images do not exhibit bright speckles but the structural connection from one vessel to another may disappear; furthermore, individual vessels even appear fractured. In fact, they appear to have two distinct color profiles, i.e., non-negligible differences in the RGB values, between the



Fig. 3 Left: a patch of the barium-perfused *raw* PCSVN image. Right: a patch of the barium-perfused *formalin-fixed* PCSVN image. Both images belong to the same placenta

interior of the vascular region and its boundary—in the topological sense—away from the nonvascular regions. The significant pixel intensity gradient at those locations causes our algorithm to treat the boundary as a separate structure and separate vessels’ interior from their boundary. Thus, instead of being part of the vascular paths, some boundaries seem to form separate, thin, neighborhood vessels in some eigenimages. On the other hand, although the disconnection issue is not as severe in the raw image, fresh blood speckles caused by glare can manifest as tubular structures that lead to corrupted enhancement results.

2.2.2 Parameter Selection

Shearlets

- *Number of scales.* The parameter s gives us the depth of the decomposition, i.e., the number of scales (from largest to finest) used to decompose the image. Therefore, when combined with shearing, this increases the diversity of basis elements used. Since PCSVN images show significant curvilinear structure away from the vessels, we need to choose s carefully, as a large value may result in capturing unwanted details. In this work, the ideal number of scales was chosen to be 4, i.e., $s = 4$.
- *Thresholding factor.* We use a thresholding routine for selecting the shearlet coefficients; this essentially translates to a smoothing process since the coefficients of the shearlets are picking up noise when they are below a certain value. Therefore, thresholding along with varying scales allow us to control the amount of desired versus undesired features. Although an optimal threshold value is a function of the images, after a careful study of our images, we found that picking a threshold that keeps wavelet coefficients of magnitudes within the top 5–10% works well in general.

Laplacian Eigenmaps

- *Number of nearest neighbors, k .* Currently, there is not yet a systematic way to make an appropriate choice [2]. If N , the number of observations is large, it is important to pick the right k : (1) in general, k may increase with N , but aspects such as the sparsity of the data should come to play; (2) since the computational costs are higher as the number of neighbors increases (it affects the sparsity of W), our goal is to find k small enough such that the quality of the representation is good. In our work, we used several values of k and observed the resulting eigenimages. We found that, if $k \geq 5$, increasing the number of nearest neighbors does not affect the quality of the vessel enhancement. Since the cost of computing eigenvectors is high, we picked $k = 5$.
- *Heat kernel parameter.* There is also no principled way to find this parameter. In [2], the authors suggested that choosing a smaller σ tends to improve the quality of the representation for bigger but still relatively small k ; for small values of k , the results do not seem to depend significantly on σ . In our work, a value of σ around 1 gives good results based on visual inspection. Below 0.5, we obtain poor quality while increasing σ by orders of magnitude (up to 100) did not affect the quality of the vessel enhancement.

2.2.3 Algorithm

1. **Pre-processing.** Make subsections of the images with size 512×512 or 256×256 , as images of dyadic size are needed for the shearlet algorithm. For the barium-perfused images, we analyze subsections or patches of size 512×512 . After an initial crop to remove irrelevant background and undesired foreground objects such as scissors, the five placenta images in our data set yielded anywhere between 6 and 20 subsections with an average of 12.6 subsections.
2. **Shearlets.** Efficiently highlight curved, slanted vessel structures using a basis composed of directional elements at various scales.
3. **LE.** Emphasize vessel structures at different scales by finding points in neighborhoods with similar structures and enhancing the similarity between these points.
4. **Thresholding.** Treat the remaining highlighted, nonvascular areas as noise. Study the differences between the noise and the vessels, and throw away the noise by setting its intensity value to zero.

2.3 Results

Our discussion on the performance of the Shearlets-LE algorithm will be strictly qualitative as the ground truth traces to the barium-perfused images are not available at this time. Basically, the algorithm takes an RGB image as an input and outputs a grayscale image in which the vessels are enhanced. In the experiments performed on

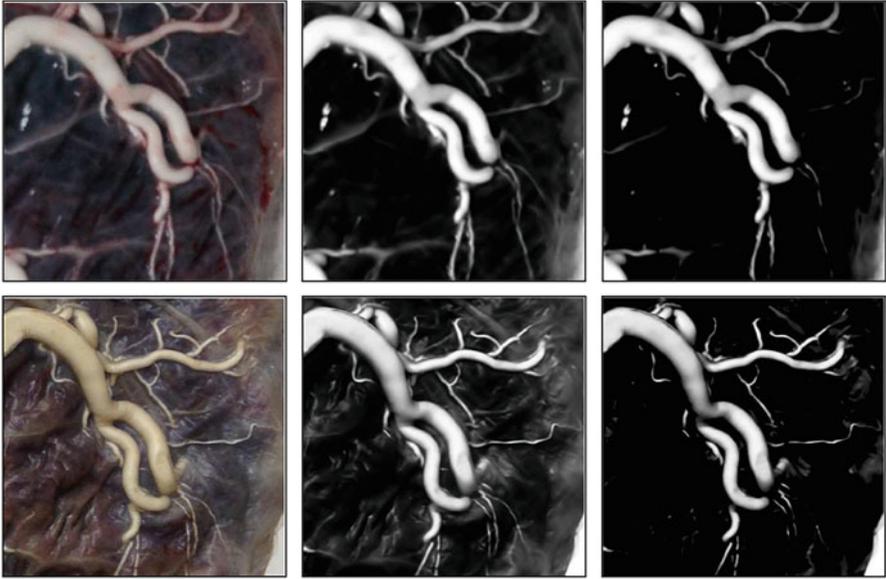


Fig. 4 Top row: Shearlets-LE enhancement result on a selected barium-perfused *raw* placenta image. Bottom row: Shearlets-LE enhancement result on a selected barium-perfused *fixed* placenta image. Left: Original PCSVN image; Middle: Enhanced image after applying the Shearlets-LE algorithm to the left image; Right: Enhanced image after applying a thresholding routine to the middle image. The results here clearly show that the vascular structure is significantly highlighted using the Shearlets-LE algorithm

raw and Barium-perfused images, where hard-thresholding was possible, pixels in regions detected as non-vessels have value zero while regions identified as vessels have non-zero values. This allows us to have a notion of vessel extraction in the same sense as the cGAN output in the next section. On the other hand, the algorithm did not produce substantial intensity value differences between vessel-like and non-vessel-like regions on the National Children’s Study (NCS) images using hard thresholding.

The preliminary results show that the vessel structure of the PCSVN is effectively captured by the Shearlets-LE algorithm on fixed and raw placentas. Figure 4 gives a representative result from running the algorithm on five placentas (about 60 images after cropping). Note that we only present the first eigenimages, but we observed that for many subsections of the placenta images, the next three eigenimages give a particularly precise delineation of the vessels. At this point, we do not have a technique to interpolate those images, but we believe that those images may be very useful when we combine Shearlets-LE and neural networks in future work. Visually, the algorithm gives a satisfying enhancement of arteries and veins of various sizes; the algorithm effectively learned directional and anisotropic variability exhibited through the tortuous vasculature network. With the appropriate parameters described previously, the shearlets were able to capture directional and

curved structures in the original images while smoothing structures away from the vessels. They do so by treating variations in nonvascular regions as noise. LE emphasizes vessel structures remaining after applying shearlets at different scales by finding points in neighborhoods with similar structures and enhancing the similarity among these points. In addition, intensity hard-thresholding further enhanced the image contrast.

To do the hard-thresholding, we studied the color scale of a few eigenimages using MATLAB's `colorbar` command. We noted that many pixels in the lower grayscale could be classified as background and have their values set to zero. For some eigenimages where the difference was particularly obvious, we simply use the `imbinarize` command in MATLAB. This was more effective in the fixed images, where noisy parts were removed while the vessels were maintained. Noticeably, superior smoothing was observed on fixed images in non-vessel-like region due to the lack of speckles induced by glare and fresh blood. It is worth noting that every step of the algorithm was essential; using only shearlets or applying LE directly to images produced significantly inferior results.

The Shearlets-LE algorithm successfully amplifies the appearance of vascular networks in barium sulfate perfused placentas in the absence of glare and discoloration. Perfusing placental vascular networks with a barium solution helps significantly with the automation of our proposed research agenda; however, a major limitation to generalize this approach is the lack of data available, as it is extremely difficult to obtain this type of data in an uncontrolled clinical setting. Furthermore, our algorithm was only tested on a small number of images. Future work on expanding our results to a larger collection of images is much needed.

3 Vessel Extraction Using a Conditional Generative Adversarial Network

Taking advantage of the recent research in deep learning and the improvement in computational resources, we propose to tackle the PCSVN extraction problem with a neural network in our second approach. Previously, a classic neural network algorithm was applied to a set of 16 PCSVN images from the University of North Carolina's Pregnancy, Infection, and Nutrition Study (UNC-PIN) [1]. Pixel-wise features such as the magnitude and the direction of the intensity gradient were fed into the network. The target output was a binary value representing whether or not the pixel represented a blood vessel. Consequently, the networks were small, consisting of layers of 8–15 neurons (see Sect. 3.1). Because of the memory limitations, only a number of random pixels were chosen from each image and passed into the neural network in order to limit the size of the training set.

In contrast, our work here uses a deep neural network with a complex architecture and millions of parameters. Each input is a cropped RGB photo of the placenta itself, rather than a feature vector; and the target output is a corresponding traced PCSVN

image, obtained by a trained expert [33]. Isola et al. [21] recently developed a very flexible conditional generative adversarial network (cGAN) for image-to-image translation, called `pix2pix`. We use their implementation in this work. Briefly, the purpose of the cGAN is to learn the relationship between pairs of corresponding images, such that given one member of the pair it can generate the other. Examples of pairs include aerial photographs and maps, night and day photos of the same scene and, of special interest to us, photo and line drawings of the same object.

3.1 Machine Learning Background

3.1.1 Convolutional Neural Networks

Neural Network Basics

The neural networks discussed here are networks used for supervised learning. The neural network learns the mapping between the input and the desired output when it is provided with a sufficient number of matched training examples. The validation set is used to select the best model hyperparameters while the test set is used to evaluate the performance of the final selected model. The test sets are not used in the training nor the validation; thus, the performance on this set gives the best indication of whether the neural network could be used to automate the process of vessel extraction.

A neural network is a machine-learning method to model a function f^* from inputs \mathbf{x} to desired outputs \mathbf{y} . If $\mathbf{y} = f^*(\mathbf{x})$, a feedforward network defines a mapping f to approximate f^* ; that is, $\mathbf{y} \approx f(\mathbf{x}; \boldsymbol{\theta})$, in terms of the parameters $\boldsymbol{\theta}$ which are learned during training [16]. Note that *hyperparameters* are values chosen in the design of the neural network, such as the loss function, number of layers, initial learning rate, etc.; *parameters* refer to values learned from the data by the neural network, not set manually.

The function f is a composition of simpler functions. A neural network is composed of layers of *neurons*. Each neuron applies a very simple function to its input. For instance, given an input \mathbf{x} , an ReLU neuron outputs $\max(0, \mathbf{x})$; it is often used in conjunction with an affine function, returning $\max(0, \mathbf{w}^T \mathbf{x} + b)$, where \mathbf{w} and b are network parameters (Fig. 5).

The networks used here, feedforward neural networks, can be represented by an acyclic graph which shows how the functions are composed together. The neurons in one layer receive inputs from multiple neurons in the previous layer, and send output to multiple neurons in the succeeding layer (Fig. 5).

The parameters of every neuron are updated using *backpropagation* [36], a method consisting of the following two stages applied repeatedly. During the forward pass, the values from inputs to outputs are computed using the composition of functions from each layer, keeping the parameters constant. During the backward pass, the error between the output $f(\mathbf{x}; \boldsymbol{\theta})$ and the target \mathbf{y} is computed and

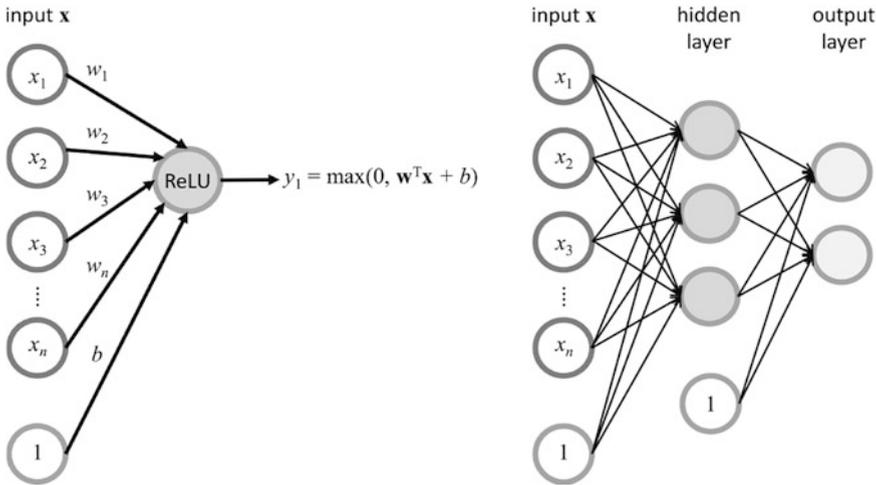


Fig. 5 Left: A schematic diagram of one ReLU neuron. Right: An example of a neural network composed of layers. The input layer consists of n input nodes and one bias term. In this case, the network has a single hidden layer with three neurons. The output layer contains two neurons. All neurons within one layer apply the same activation function

propagated back through the network layers. The parameters θ of the neurons are updated by gradient descent such that the error, or the loss, is minimized.

Convolutional Layers

The types of neural networks used in image processing have convolutional layers which learn the filters to best extract features from the image, with no prior knowledge. Unlike the fully connected layers described previously, filters take a local region of the input at a time, but sweep across the entire image.

We can think of the input \mathbf{x} of a convolutional layer as a volume of size $W \times H \times D$, where W is the width of the image, H is the height and $D = 3$ if we restrict to RGB images of three channels. A convolutional filter \mathbf{k} is of size $F \times F \times D$, e.g. $4 \times 4 \times 3$. The entries of the filter are learned parameters of the network. The output at a position along an interior point of the image is the dot product of the filter and the values from an $F \times F$ patch of the image [16]:

$$S(i, j) = (\mathbf{k} * \mathbf{x})(i, j) = \sum_{l=1}^3 \sum_m \sum_n \mathbf{x}(i + m, j + n, l) \mathbf{k}(m, n, l). \quad (7)$$

Each filter produces a two-dimensional activation map when it is passed along, in “hops” of stride S , across the image. If D_2 different filters are used for that convolutional layer, the layer’s output will have depth D_2 . Output width and height

Fig. 6 The weights of the 4×4 filters (R channel) of the first convolutional layer in the generative network applied to placenta images



depend on F and S , in the simplest case. This output volume is then processed by the next layer of the network. A nonlinear function (such as ReLU) can be applied after a convolutional layer.

After training, the filters of the first convolutional layer learn low-level features. Figure 6 shows the 64 filters, with $F = 4$, for the first convolutional layer of the trained network in this project. Most filters do not converge to random weights, but show a gradient of intensity, or an area of greater intensity, across the small square. Empirically, after training, the filters learned by the first layer of a convolutional neural network often resemble Gabor filters [24] used in texture analysis; these filters activate when they encounter edges along a certain orientation or a patch of a certain color [16]. Deeper layers of the network capture more complex features.

3.1.2 Conditional Generative Adversarial Network

As mentioned above, a convolutional neural network learns the mapping from the input image to the output image. A conditional generative adversarial neural network (cGAN) additionally learns the loss function. Instead of minimizing a metric of the discrepancy between the output and the target images, cGAN uses a coupling of two neural networks to create output images which are indistinguishable from the targets. Given one member in the pair of the training images (e.g., a digital photograph of the placenta), a generator network generates the other (e.g., the traced PCSVN), competing against an adversary who tries to distinguish it against the ground truth (e.g., manually traced PCSVN). A detailed exposition of the method can be found in [21]; here, we present a brief overview.

Generative adversarial networks (GANs) and their optimization are introduced in [15]. In a GAN, a generator network G produces samples $G(\mathbf{z})$ from a noise

prior $p_z(\mathbf{z})$, in order to approximate the target images \mathbf{y} . The discriminator D is the (adversarial) neural network with input either \mathbf{y} or $G(\mathbf{z})$; it outputs the probability that the input came from the data rather than being generated. In other words, the network D can be represented as a function with a scalar output in the range $[0,1]$.

For a GAN, the optimization function can be written as follows (adapted from equations (1) of [15] and (20.81) of [16]):

$$\mathcal{L}(G, D) = \mathbb{E}_{\mathbf{y} \sim p_{\text{data}}} [\log D(\mathbf{y})] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [\log (1 - D(G(\mathbf{z})))] . \quad (8)$$

The adversarial networks work in competition with each other: G to minimize the objective and D to maximize it, so that

$$G^* = \arg \min_G \max_D \mathcal{L}(G, D) . \quad (9)$$

A *conditional* GAN (i.e., cGAN) uses not only \mathbf{z} , but also an input image \mathbf{x} (in our case, the photographs) to simulate the outputs \mathbf{y} (traces). More formally, let \mathbf{x} be an input image, \mathbf{y} an output image, and \mathbf{z} a noise vector. The generator, G , is trained to learn a mapping $G(\mathbf{x}, \mathbf{z}) \approx \mathbf{y}$. The discriminator, D , is trained to distinguish between $G(\mathbf{x}, \mathbf{z})$ and \mathbf{y} . From equation (1) of [21], the objective function is now:

$$\mathcal{L}(G, D) = \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p_{\text{data}}(\mathbf{x}, \mathbf{y})} [\log D(\mathbf{x}, \mathbf{y})] + \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x}), \mathbf{z} \sim p_z(\mathbf{z})} [\log (1 - D(\mathbf{x}, G(\mathbf{x}, \mathbf{z})))] . \quad (10)$$

Additionally, as in regular convolutional neural networks, we would like the output image, $G(\mathbf{x}, \mathbf{z})$, to be close to the target \mathbf{y} using an appropriate norm. In [21], this is achieved by adding an L_1 penalty to the objective function. The final objective is then:

$$G^* = \arg \min_G \max_D \mathcal{L}(G, D) + \lambda \mathbb{E}_{\mathbf{x}, \mathbf{y}, \mathbf{z}} [|\mathbf{y} - G(\mathbf{x}, \mathbf{z})|_1] . \quad (11)$$

The L_1 norm was chosen as it produces less blurry outputs compared to L_2 .

The structure of the G network is based on the U-Net architecture [31], an established, frequently used network for image segmentation. The D network classifies small patches of the image as real (from \mathbf{y}) or fake (from $G(\mathbf{x}, \mathbf{z})$) and averages these decisions across all responses for one image. Therefore, the discriminator can work on arbitrarily large images since it works on a local scale. Both G and D networks are formed from modules of three layers: a convolutional layer followed by a layer which performs batch normalization [20], followed by a layer of ReLU neurons. (Normalizing data by subtracting the mean and dividing by the standard deviation is a common preprocessing step in machine learning. Batch normalization is a variant of this transformation applied before non-linear function layers within the network itself. Without it, small changes in the input of a deep neural network can propagate into large changes across the network, making training difficult.)

3.2 Application

3.2.1 Data Sets

Training a deep neural network, i.e., one with many layers and many weights, requires a sufficiently large number of images. To this end, we used 201 digital photographs of formalin-fixed placentas from the National Children’s Study (NCS). The photographs were taken at delivery or following pathology evaluation using a polarized filter. The images were made available to us through Placental Analytics, LLC. The NCS data set has the benefit of having high-quality, manually traced PCSVNs, serving as the ground truth in the supervised training of the cGAN. The tracings were done with a validated protocol [13, 33]; colors were used to differentiate vessel diameters which range from 3 to 19 pixels. A random sample of 10% of the tracings were traced by a second tracer, to check the tracing for consistency [6].

3.2.2 Preprocessing

The preprocessing steps began with turning the traced PCSVN images into black and white. The images were then cropped into non-overlapping squares of 256×256 pixels for the neural network. Since the feature of our interest, the vasculature, does not have an orientation, we also rotated each cropped image by 90° , 180° , and 270° to augment the data set for the training. After removing 80% of the images under 10 KB in size, which were blank or non-informative, the resulting data set has 7210 training and 2357 validation images associated with 121 and 40 placentas, respectively. The remaining 40 placentas constitute the testing set.

3.2.3 Conditional Generative Adversarial Network

We used the `pix2pix` implementation from Isola et al. [21] with a few adjustments. A thorough parameter search is needed in future work; at present, the network ran for 25 iterations at a learning rate of 0.0002, followed by 25 iterations for a decaying rate. The L_1 regularization parameter λ in the code was set to 100. Complete details of the options can be found at <https://github.com/canghel/placenta/docs>.

3.2.4 Postprocessing

Given a cropped (256×256) image, `pix2pix` outputs the reconstructed estimate of the vasculature trace for that image. The cropped images were then patched together to form a fully reconstructed PCSVN. The outputs from the rotated patches were rotated and reassembled back into the respective traces for the same placenta, which were then averaged together. Doing this produced an undesirable tiling artifact since

the squares do not overlap. A simple way to improve the reconstructed result is to generate overlapping squares from the photographs, pass them through the cGAN, and average the resulting reconstructions. To this end, we translated the original tiling by 64, 128, and 192 pixels to generate additional reconstructions for the interior of the images. Note that the pre-and post-processing will likely change in future work. Full-size images can be input into the `pix2pix` network, and cropped automatically during training only.

3.3 Results

The Matthews Correlation Coefficient (MCC) is used to quantitatively evaluate the performance of the vessel extraction method. It is an aggregate measure of the confusion matrix when the accuracy is expected to be low and the classes are imbalanced, as is the case for placenta images. Additionally, it allows for a more consistent comparison to previous work in [1, 5].

$$\text{MCC} = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}. \quad (12)$$

The averaged reconstructions from `pix2pix` were thresholded to create binary, black and white images. In our case, “positives” represent the pixels identified as belonging to vessels in the manual, ground truth trace. Thus, TP represents the number of pixels identified as blood vessel in both the cGAN reconstruction and the trace. Similarly, TN is the number of true negatives, FP is the number of false positives, and FN is the number of false negatives.

Figure 7 gives an illustrative example of the worst, average, and best cGAN reconstructed result on the test set. Notice that the cGAN reconstructed traces captured large veins and arteries and filtered out noise from the irregular texture of the surrounding tissue. The region near the umbilical cord insertion was blurry and the thinner vessels were not captured well. Interestingly though, the variation in the thickness of the larger vessels often matches the photograph more closely than the hand-drawn trace. This is because the hand-traces are limited to lines of fixed pixel widths, while cGAN produces smoother transitions for vessels which vary in thickness.

The average MCC on the 40 test placentas is 0.76, ranging from 0.67 to 0.84, as shown in Fig. 8. In comparison, previous methods returned a maximum MCC value of 0.4 on 16 images of a similar data set (UNC-PIN) [1, 5].

As expected, reconstructions of the placentas used from training had higher MCC values than those for the validation and testing sets. Using the overlapping patches in the cGAN reconstructions helped to increase the MCC in the test set (Wilcoxon signed rank test, p -value = $1.82e-12$) and in the validation sets (Wilcoxon signed rank test, p -value = $3.46e-11$), but not for the training set, as illustrated in Fig. 9. As averaging may have the effect of blurring an image, this result may indicate that

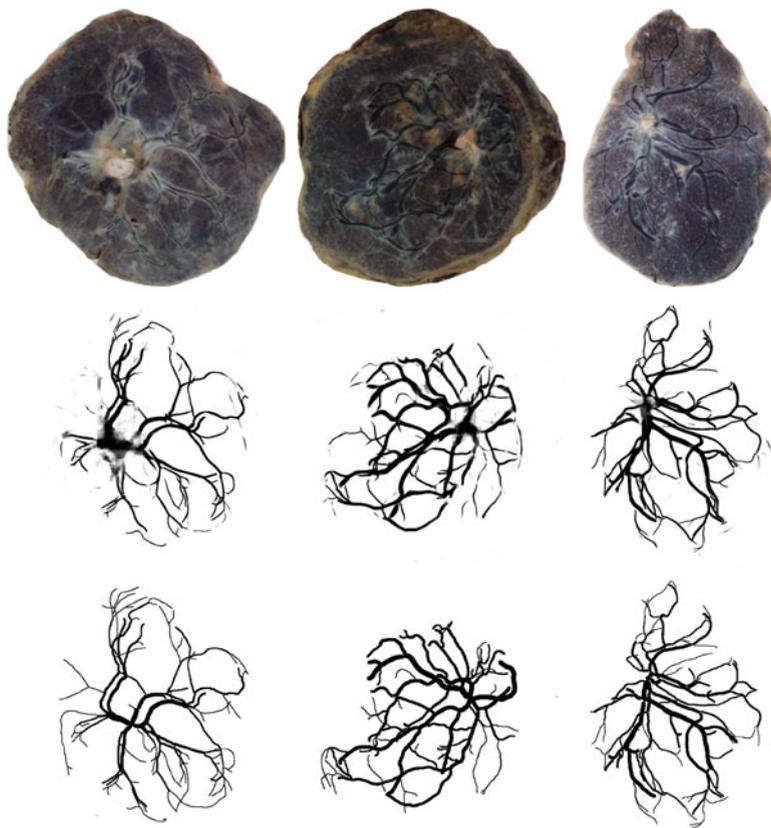


Fig. 7 Left to right: An illustrative example of the worst, average, and best cGAN reconstructed results on the test set in terms of the Matthews Correlation Coefficient (MCC) values. Top row: Test NCS photographs; Middle row: cGAN reconstructed PCSVN images; Bottom row: Manually traced ground truth images

cGAN was overfitting the training set, highlighting the importance of further work on a thorough parameter search.

4 Discussion

4.1 Comparison of Shearlets-LE and cGAN

We have developed and presented two parallel methods for the automation of PCSVN extraction from digital photographs of placentas. Here we apply both algorithms to cropped photographs (256×256) in the NCS data set for comparison.

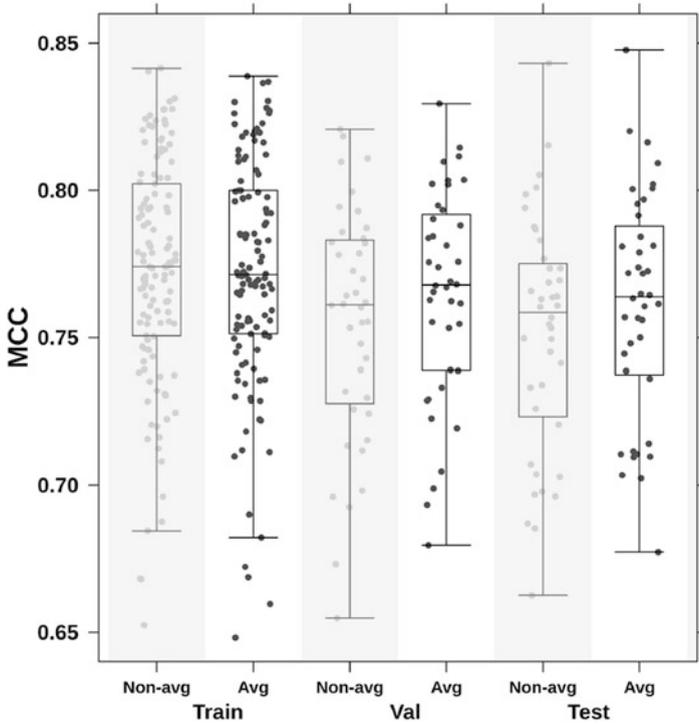


Fig. 8 MCC distributions for the training, validation, and test sets. For each data set, the *averaged* results consist of averaged reconstructions from both rotated and overlapping 256×256 pixel squares. The *non-averaged* results are obtained from four reconstructions each consisting of non-overlapping squares, one for each angle of rotation

By incorporating structures at various scales and orientations as well as neighborhood information, the combination of shearlets and LE is effective in picking up directional and curved structures in the image, and is particularly successful in highlighting small vessels as shown in the bottom row, second column of Fig. 10.

The `pix2pix` cGAN is a completely general method, initially blind to the relationship between the images. It learns this relationship from the images themselves in two ways: by determining the parameterization of suitable filters and by using an adversarial loss function which extends the simple penalty of deviation in the L_1 norm from the ground truth. It is particularly suited to vessel detection in placentas where both the vasculature and the tissue itself vary greatly in shape, color, and texture. It performed much better than all previous approaches. The cGAN reconstructed images, as shown in the third column of Fig. 10, exhibit well-captured global structures, i.e., large vessels.

The two methods are complementary, and we expect a combination of the two can lead to improved results.

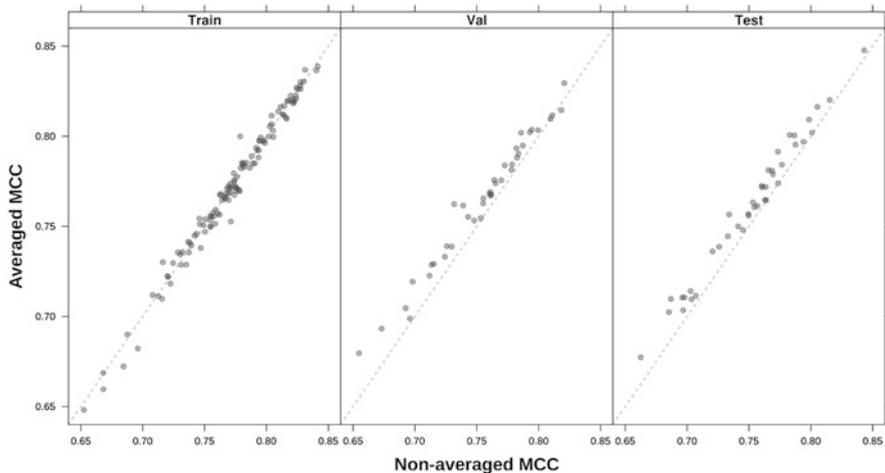


Fig. 9 Comparison of the MCC values in *averaged* and *non-averaged* cGAN reconstructions. The dashed line is the 45° line $y = x$. Using the overlapping patches in the cGAN reconstructions helped to increase the MCC in the test and validation sets

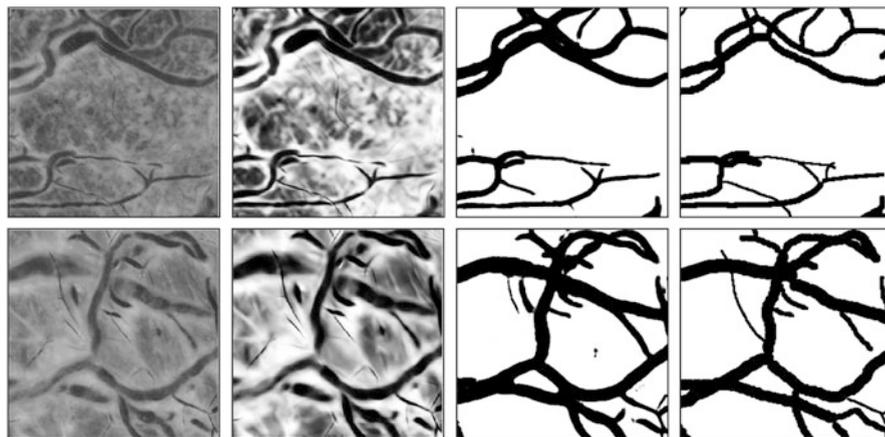


Fig. 10 Top and bottom: Two illustrative examples to compare the results of PCSVN extraction using the Shearlets-LE and cGAN algorithms. Left: 256×256 patch images in the NCS data set; Middle left: Shearlets-LE enhanced result; Middle right: cGAN reconstructed result; Right: Traced ground truth. Notice that the quality of the Shearlets-LE enhancement varies, especially in the background. On the other hand, the cGAN was able to capture large vessels very well but fails at the finer scales

4.2 Future Directions

As described above, we would like to combine the strengths of both methods in our future work. One simple way to do this is to input Shearlets-LE enhanced

images into the cGAN directly or as additional pseudo-color channels of the RGB images. In some way, the enhanced images may also incorporate prior knowledge of the vessel morphology into the cGAN. The Shearlets-LE transform highlights the curvilinearity of the vessels, a feature that is important in distinguishing vessels from non-vessels.

Another property of the vessels is their connectedness. In training the cGAN, the traces may also be replaced by their skeletonizations, i.e., thin lines marking the midpoints of the vessels, to emphasize the tree structure of the vasculature and to de-emphasize the thickness of the vessels. The outputs from the different cGANs trained on different training images (Shearlets-LE transformed inputs, skeletonization outputs, etc.) can be averaged to form the final estimate of a reconstructed trace. Such ensemble models frequently improve performance. The first step, however, is to optimize the performance of the `pix2pix` cGAN by performing a full parameter search for learning rate, etc., and to compare it to previous methods on the same image data set.

Further postprocessing of the full reconstructed traces from the cGANs may be an equally important step. The intuition is similar to that reported in [5], where curvilinear and vessel enhancement steps removed noise from the image obtained with only a multiscale filter. We only used cropped images from the placentas in training the cGAN; the full image of a trace contains additional structural information that was not exploited, such as the position of the umbilical chord and the tree structure of the vasculature.

The work presented here has focused on enhancing vessel structure and automated vessel network extraction (Fig. 1). We are optimistic that, with the additional extensions described, it will soon be possible to obtain black and white tracings of the vasculature. From these tracings, the skeletonization of the traces is straightforward, and various properties of the vessels such as mean tortuosity, mean thickness, etc., can be computed as in [6]. While identifying arteries and veins separately will be an additional challenge, using the entire PCSVN as a whole may be sufficient to identify interesting associations with ASD risk. The advantage of automation is that studies linking PCSVN features with ASD and other developmental delays can be easily scaled up to hundreds and thousands of images.

5 Computational Time and Software Specifications

For the Shearlets-LE method, the details of the implementation are as follows:

- The proposed method took roughly 12 min for a 512×512 patch image on a 2.9 Ghz Intel Core i7 CPU. This calculation was not computationally prohibitive.
- For shearlets, we used the shearlets package by www.shearlab.org. A useful companion can be found at http://www.math.tu-berlin.de/fileadmin/i26_fg-kutyniok/Kutyniok/Papers/ShearLab3D.pdf.

- For Laplacian Eigenmaps, we used Matlab’s Toolbox for Dimensionality Reduction by Laurens van der Maaten (<https://lvdmaaten.github.io/drtoolbox/>) with various additions/optimizations by A. Cloninger, T. Doster, A. Halevy, K. Yacoubou Djima.

The details for the cGAN work are as follows:

- Image preprocessing was done in Python 3.6.1. Figures were produced in the R programming language (v3.3.3) with the BoutrosLab.plotting.general (BPG) package (v5.3.4) <http://labs.oicr.on.ca/boutros-lab/software/bpg>. Codes for all processing can be found at <https://github.com/canghel/placenta/clean-code>.
- The cGAN implementation comes from the `pix2pix` model in [21], implemented in PyTorch <https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix> (Retrieved July 19, 2017).
- Training required 8 h on Intel Xeon Processor (10 M Cache, 3.50 Ghz) CentOS 6.5 64 bit with NVIDIA GeForce GTX 1080 GPU, but once trained the testing time was minimal (4 min and 15 s for all 7124 overlapping 256×256 cropped test images).

Acknowledgements The project was part of the MBI Women Advancing Mathematical Biology: Understanding Complex Biological Systems with Mathematics 2017 Workshop organized by the Association for Women in Mathematics. Funding for the workshop was provided by MBI, NSF ADVANCE “Career Advancement for Women Through Research-Focused Networks” (NSF-HRD 1500481), Society for Mathematical Biology, and Microsoft Research.

Over the course of the project, we received biology expertise and support from Drs. Ruchit Shah, George Merz, and Richard K. Miller.

The authors also wish to thank the following people who contributed to the collection of the placentas in the National Children’s Study Placenta Consortium: C.J. Stodgell, L. Salamone, L.I. Ruffolo, A. Penmetsa, P. Weidenborner (University of Rochester), J. Culhane, S. Wadlinger, M. Pacholski, M.A. Kent, L. Green (University of Pennsylvania), R. Wapner, C. Torres, J. Perou (Columbia University), P. Landrigan, J. Chen, L. Lambertini, L. Littman, P. Sheffield, A. Golden, J. Gilbert, C. Lendor, S. Allen, K. Mantilla, Y. Ma (Ichan School of Medicine), S. Leuthner, S. Szabo (Medical College of Wisconsin), J.L. Dalton, D. Misra (Placental Analytics), N. Thiex, K.Gutzman, A. Martin, B. Specker (South Dakota University), J. Swanson, C. Holliday, J. Butler (University of California at Irvine), A. Li, R.M.A.P.S. Dassanayake, J. Nanes, Y. Xia (University of Illinois at Chicago), J.C. Murray, T.D. Busch, J. Rigdon (University of Iowa), Kjersti Aagaard, A. Harris (Baylor College of Medicine), T.H. Darrah, E. Campbell (Boston University), N. Dole, J. Thorp, B. Eucker, C. Bell (University of North Carolina at Chapel Hill), E.B. Clark, M.W. Varner, E. Taggart, J. Billy, S. Stradling, J. Leavitt, W. Bell, S. Waterfall (University of Utah), B. O’Brien, M. Layton, D. Todd, K. Wilson, M.S. Durkin, M.-N. Sandoval (Westat, Inc).

Most importantly, we thank the study participants who donated their placentas.

References

1. N. Almoussa, B. Dutra, B. Lampe, P. Getreuer, T. Wittman, C. Salafia, L. Vese, Automated vasculature extraction from placenta images, in *Proceedings of SPIE Medical Imaging Conference*, vol. 7962, 2011
2. M. Belkin, P. Niyogi, Laplacian Eigenmaps for dimensionality reduction and data representation. *Neural Comput.* **15**, 1373–1396 (2002)

3. P. Burt, E. Adelson, The Laplacian pyramid as a compact image code. *IEEE Trans. Commun.* **31**, 532–540 (1983)
4. E. Candès, L. Demanet, D. Donoho, L. Ying, Fast discrete curvelet transforms. *Multiscale Model. Simul.* **5**(3), 861–899 (2006)
5. J.-M. Chang, N. Huynh, M. Vasquez, C. Salafia, Vessel enhancement with multi-scale and curvilinear filter matching for placenta images, in *Proceeding of the 2013 20th International Conference on Systems, Signals and Image Processing (IWSSIP)*, 2013, pp. 125–128
6. J.-M. Chang, H. Zeng, Y.-M. Chang, R. Shah, C. Salafia, C. Newschaffer, R. Miller, P. Katzman, M.-F. Moye, C. Walker, L. Croen, Autism risk classification using placental chorionic surface vascular network features. *BMC Med. Inform. Decis. Mak.* **17**(1), 162 (2017)
7. F.R.K. Chung, *Spectral Graph Theory* (American Mathematical Society, Providence, 1997)
8. R. Coifman, Y. Meyer, V. Wickerhauser, Adapted wave form analysis, wavelet-packets and applications, in *Proceedings of the second International Conference on Industrial and Applied Mathematics, ICIAM 91* (Society for Industrial and Applied Mathematics, Philadelphia, 1992), pp. 41–50
9. W. Czaja, M. Ehler, Schrödinger Eigenmaps for the analysis of bio-medical data. *CoRR*, abs/1102.4086, 2011
10. I. Daubechies, Orthonormal bases of compactly supported wavelets. *Commun. Pure Appl. Math.* **41**(7), 909–996 (1988)
11. M. Do, M. Vetterli, Contourlets: a directional multiresolution image representation, in *ICIP (1)*, 2002, pp. 357–360
12. M. Egbor, T. Ansari, N. Morris, C.J. Green, P.D. Sibbons, Morphometric placental villous and vascular abnormalities in early- and late-onset pre-eclampsia with and without fetal growth restriction. *J. Obstet. Gynecol.* **113**(5), 580–589 (2006)
13. E. Farnell, S. Farnell, J.-M. Chang, M. Hoffman, R. Belton, K. Keaty, S. Lederman, C. Salafia, A shape-context model for matching placental chorionic surface vascular networks. *Image Anal. Stenogr.* **37**(1), 55–62 (2018)
14. A.F. Frangi, W.J. Niessen, K.L. Vincken, M.A. Viergever, *Multiscale Vessel Enhancement Filtering* (Springer, Berlin, 1998)
15. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in *Advances in Neural Information Processing Systems*, vol. 27, ed. by Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, K.Q. Weinberger (Curran Associates, Inc., Red Hook, 2014), pp. 2672–2680
16. I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning* (MIT press, Cambridge, 2016)
17. K. Guo, D. Labate, Optimally sparse multidimensional representation using shearlets. *SIAM J. Math. Anal.* **39**(1), 298–318 (2007)
18. E. Haeussner, C. Schmitz, H.G. Frank, F.E. von Koch, Novel 3D light microscopic analysis of IUGR placentas points to a morphological correlate of compensated ischemic placental disease in humans. *Sci. Rep.* **6**, 24004 (2016)
19. J. Ham, D. Lee, S. Mika, B. Schölkopf, A kernel view of the dimensionality reduction of manifolds, in *Proceedings of the Twenty-first International Conference on Machine Learning, ICML '04* (ACM, New York, 2004), pp. 47–55
20. S. Ioffe, C. Szegedy, Batch normalization: accelerating deep network training by reducing internal covariate shift, in *International Conference on Machine Learning*, 2015, pp. 448–456
21. P. Isola, J.-Y. Zhu, T. Zhou, A.A. Efros, Image-to-image translation with conditional adversarial networks, in *CVPR*, 2017
22. P. Jarmuzek, M. Wielgos, D.A. Bomba-Opon, Placental pathologic changes in gestational diabetes mellitus. *Neuroendocrinol. Lett.* **36**(2), 101–105 (2015)
23. D. Labate, W.-Q. Lim, G. Kutyniok, G. Weiss, Sparse multidimensional representation using shearlets. *Opt. Photon.* **2005**, 59140U (2005)
24. T.S. Lee, Image representation using 2D Gabor wavelets. *IEEE Trans. Pattern Anal. Mach. Intell.* **18**(10), 959–971 (1996)
25. J. Lee, M. Verleysen, *Nonlinear Dimensionality Reduction*, 1st edn. (Springer, New York, 2007)

26. S. Mallat, A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.* **11**, 674–693 (1989)
27. L. Matelski, J. Van de Water, Risk factors in autism: thinking outside the brain. *J. Autoimmun.* **67**, 1–7 (2016)
28. A. Modabbernia, E. Velthorst, A. Reichenberg, Environmental risk factors for autism: an evidence-based review of systematic reviews and meta-analyses. *Mol. Autism* **8**, 13 (2017)
29. C.J. Newschaffer, L.A. Croen, M.D. Fallin, I. Hertz-Picciotto, D.V. Nguyen, N.L. Lee, C.A. Berry, H. Farzadegan, H.N. Hess, R.J. Landa, S.E. Levy, M.L. Massolo, S.C. Meyerer, S.M. Mohammed, M.C. Oliver, S. Ozonoff, J. Pandey, A. Schroeder, K.M. Shedd-Wise, Infant siblings and the investigation of autism risk factors. *J. Neurodev. Disord.* **4**, 1–16 (2012)
30. H.R. Park, J.M. Lee, H.E. Moon, D.S. Lee, B.N. Kim, J. Kim, D.G. Kim, S.H. Paek, A short review on the current understanding of autism spectrum disorders. *Exp. Neurobiol.* **25**(1), 1–13 (2016)
31. O. Ronneberger, P. Fischer, T. Brox, U-net: convolutional networks for biomedical image segmentation, in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Springer, Berlin, 2015), pp. 234–241
32. C.M. Salafia, C. Platt, T. Girardi, R. Shah, G. Merz, D.P. Misra, Placental structure in ASD: does the placenta mirror the diagnosis? in *2014 International Meeting for Autism Research*, page Abstract No. 17578, May 14–17 2014
33. R.G. Shah, C.M. Salafia, T. Girardi, L. Conrad, K. Keaty, A. Bartleotc, Shape matching algorithm to validate the tracing protocol of placental chorionic surface vessel networks. *Placenta* **36**(8), 944–946 (2015)
34. C.K. Walker, P. Krakowiak, A. Baker, R.L. Hansen, S. Ozonoff, I. Hertz-Picciotto, Preeclampsia, placental insufficiency, and autism spectrum disorder or developmental delay. *JAMA Pediatr.* **169**(2), 154–162 (2015)
35. K. Yacoubou Djima, L. Simonelli, D. Cunningham, W. Czaja, Detection of anomaly in human retina using Laplacian Eigenmaps and vectorized matched filtering, in *Proceedings of SPIE*, vol. 9413, Mar 2015, pp. 94132F-1–94132F-11
36. S.K. Zhou, H. Greenspan, D. Shen, *Deep Learning for Medical Image Analysis* (Elsevier Science, Amsterdam, 2017)