

Retinal Vessel Extraction with the Image Ray Transform

Alastair H. Cummings and Mark S. Nixon
{ahc08r,msn}@ecs.soton.ac.uk

University of Southampton, Southampton, SO17 1BJ, UK

Abstract. Extraction of blood vessels within the retina is an important task that can help in detecting a number of diseases, including diabetic retinopathy. Current techniques achieve good, but not perfect performance and this suggests that improved preprocessing may be needed. The image ray transform is a method to highlight tubular features (such as blood vessels) based upon an analogy to light rays. The transform has been employed to enhance retinal images from the DRIVE database, and a simple classification technique has been used to show the potential of the transform as a preprocessor for other supervised learning techniques. Results also suggest potential for using the ray transform to detect other features in the fundus images, such as the fovea and optic disc.

1 Introduction

Automated detection of blood vessels within eye fundus images is an important step in aiding diagnosis of a number of diseases, diabetic retinopathy in particular. The possible application of vision techniques to this problem was first identified by Chaudhuri et al. [1] through use of matched filters. A morphological approach was taken by Zana and Klein [2] whilst Jiang and Mojon [3] investigated different approaches to thresholding. Staal et al. [4] both assembled the DRIVE database upon which our work is tested and described a supervised learning method of classification of pixels based upon ridges within the image. Soares et al.[5] developed a supervised technique based upon Morlet wavelets and Niemeijer et al.[6] compared previous unsupervised techniques with a simple classifier on the DRIVE database. More recent work has built upon matched filters through optimisation of parameters [7] and use of extra methods to decide on thresholds [8]. The best supervised techniques are capable of achieving an accuracy on the DRIVE database of around 0.95, and we propose that superior preprocessing methods are necessary to improve performance further. We describe experiments with a preprocessor based upon a physical analogy to light rays; the image ray transform.

Physical analogies are an exciting paradigm in computer vision enabling the creation of novel techniques that approach the problems of feature extraction from entirely different angles [9]. These analogy based techniques have the advantage of being based on physical properties of natural phenomena such as water, heat or force fields and so are more easily understood by those using

them. In addition to the intuitive nature of the algorithms, the parameters used have meanings that are clear and have real world analogues. Although analogy operators are heavily based upon a well defined physical concepts, the analogies can be adapted outside this definition to increase their effectiveness and flexibility whilst maintaining the strengths provided by the analogy. These properties are a clear advantage over many standard techniques for which the mechanics can be hard to grasp and parameter selection is not clear.

Heat flow has been used as an analogy due to its smoothing properties. Anisotropic diffusion [10] is an edge-aware smoothing technique that allows heat to flow across areas of low but not high edge strength allowing Gaussian noise to be eliminated whilst maintaining important edge features. This is one of the earliest examples of a principled vision technique based upon an analogy. Water flow has also been used [11] as an analogy for image segmentation, and has been tested for medical applications including retinal images.

Other techniques have not used light in such a strongly analogical sense; The Eikonal equation describes the time a ray takes to travel from the boundary of an anisotropic medium to a point within it and has been used in a number of applications [12]. However none of these fully take advantage of the possible analogical formulation and is most often used as a distance metric. We use the image ray transform, a transform based upon an analogy to light rays, to detect retinal vasculature. Previously it has been used to enhance circle detection [13], but here we take advantage of its ability to highlight tubular features as shown by its previous use enabling enrolment for ear biometrics [14].

This paper describes a retinal blood vessel extraction technique using the image ray transform which detailed in section 2. Section 3 describes the application of the transform for retinal blood vessel extraction, showing how it can highlight vascular features well. Finally section 4 draws conclusions and describes future work; expanding on vessel extraction and expansion into other related areas.

2 The Image Ray Transform

The image ray transform is a novel technique for extracting tubular and circular features that are not often found by other methods. It uses a pixel based ray tracing technique and a subset of the laws of optics to trace rays through an image. These then react to certain structural features, emphasising them. Whilst the transform is based on the principles of optics, the details of the technique can be adjusted to suit successful feature extraction rather than accurate simulation of the natural phenomenon. The implementation capitalises only on the basis of the analogy; we do not simulate intricate details of light propagation.

2.1 Laws of Optics

Rays are a method of modelling the propagation of waves, most often light. Specific regard is given for the direction of the wave as it interacts with its environment, and wave-like interactions such as diffraction are ignored. The path of a light ray will be altered when it crosses the boundary with a medium of

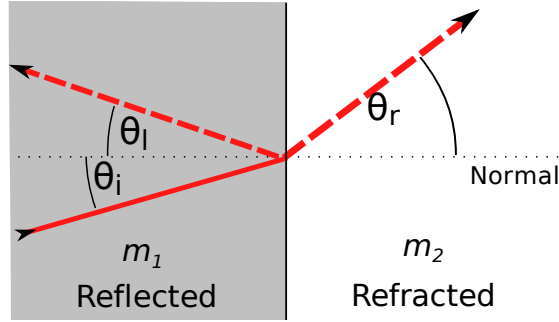


Fig. 1. Refraction and reflection of light at the boundary of two media.

different refractive index, refracting and/or reflecting (see figure 1). Light crosses the boundary between the media at an angle of θ_i to the normal of the boundary (the dotted line). If reflection occurs, then $\theta_r = \theta_i$. If refraction occurs, the light refracts at an angle of θ_r to the normal where θ_r is calculated from the refractive indices of n_1 and n_2 of the media m_1 and m_2 . Refractive index n is the ratio of the speed of light in a vacuum to the speed of light within the medium, in nature $1 \leq n \lesssim 4$. θ_r is found by Snell's Law:

$$\frac{\sin \theta_i}{\sin \theta_r} = \frac{n_2}{n_1} \quad (1)$$

If $n_1 < n_2$, light bends towards the normal, so $\theta_r < \theta_i$. If $n_1 > n_2$, light bends away from the normal, as long as it is below the critical angle θ_c . This is the angle for which θ_r would be equal to 90° and is calculated as:

$$\theta_c = \sin^{-1} \left(\frac{n_2}{n_1} \right) \quad (2)$$

Reflected angles above this are physically impossible. In this case, the light is totally internally reflected. In the natural world, the amount of refraction and reflection that occurs depends on the media and in most cases is a combination of the two, some part passing through, and some being reflected back. These rules form the basis of the image ray transform.

2.2 The Image Ray Transform

The image ray transform takes the laws of optics described previously and applies them to the problem of feature extraction. The transform operates by tracing a number of rays through an image. The paths of these rays are then used to generate a new image, with tubular and circular features emphasised. The image is analogised to a matrix of two dimensional glass blocks, each representing a pixel whose refractive index is related to the intensity of the pixel in the image. One method to assign a refractive index to a pixel with intensity i is shown in

equation 3 (another is described in section 2.3), where n_{\max} defines the maximum possible refractive index. The indices are then spaced evenly between 1 and n_{\max} :

$$n_i = 1 + \left(\frac{i}{255} \right) \cdot (n_{\max} - 1) \quad (3)$$

The splitting of rays into reflected and refracted parts is not considered as in order to reduce computational complexity.

In this matrix of blocks, a ray is created with the position (x and y) and the direction (ϕ) drawn from a uniform distribution. For an image of size $w \times h$:

$$x \in [0, w), y \in [0, h), \phi \in [0, 2\pi) \quad (4)$$

ϕ is converted into a unit vector, \mathbf{V} , representing the direction of the ray as this is more convenient for calculation. The position vector \mathbf{p} initialises the ray at:

$$\mathbf{p}^{<0>} = (x, y)^T \quad (5)$$

and at iteration i :

$$\mathbf{p}^{<i+1>} = \mathbf{p}^{<i>} + \mathbf{V} \quad (6)$$

At each pixel through which the ray passes, \mathbf{A} is updated to show that the ray has passed through: $\mathbf{A}(\mathbf{p})$ is only increased once per ray, so as to prevent small loops in the ray's path repeatedly increasing a single pixel and excessively emphasising noise or other undesirable features. When crossing a pixel boundary between media of differing refractive indices, a new path must be calculated using a vector formation of the laws described in section 2.1. Using the direction of the ray \mathbf{V} , the normal of the boundary \mathbf{N} , and the refractive indices of the first and second media n_1 and n_2 the new direction can be calculated. If $n_1 > n_2$ we have to test whether we should internally reflect or refract. The critical angle θ_c must be found from equation 2 and if $n_2 > n_1$ or $\theta_i < \theta_c$ the ray refracts in direction \mathbf{R}_r , otherwise it reflects in direction \mathbf{R}_l . The equations for calculating these from \mathbf{V} and \mathbf{N} can be found in [15]. The ray is traced in this manner until it has either undergone d reflections or refractions or the ray exits the image. This is repeated for N rays or until a stopping condition (section 2.3) is reached, and the normalised accumulator gives the transformed image.

The values of the normals (\mathbf{N}) used to calculate the new directions of rays are always set to the normal of the edge direction found by the Sobel operator at that point. Currently, we only perform refraction when moving from higher to lower refractive indices as this tends to improve the quality of the transform result.

2.3 Refinements

In addition to the basic transform, a number of additional parameters and processes can be performed with the transform to give a wider range of results and to extend the range of circumstances in which the transform is useful.

Intensity and Refractive Index The transform will only extract features that have larger refractive indices than the surrounding area because total internal reflection only prevents rays from passing from materials with higher to lower refractive indices, not vice versa. As the refractive index is proportional to intensity, this has the effect of extracting features that are lighter than their surrounding area. Depending on what is known about the properties of the features to be extracted different measures can be taken to improve the result of the transform.

If the approximate intensity of the desired features is known then the image can be easily transformed to make the target intensity have the largest value. This can be done by finding the difference from the target intensity t to the original intensity i_o for each pixel, as in equation 7:

$$i_t = |i_o - t| \quad (7)$$

It may also be the case that the difference in intensity between features and the surrounding area is not significant enough to extract it with the linear refractive indices calculated by equation 3. In such a case an alternative version can be used (equation 8) that assigns refractive indices exponentially, to ensure greater difference, and more refraction and reflection:

$$n_i = e^{\frac{i}{k}} \quad (8)$$

In this case it is k rather than n_{\max} that controls the scale of the refractive indices.

Automatic Stopping Condition Finding the optimal number of rays that should be traced in order to ensure the result of the transform is of sufficient quality is a challenge that must be met. Rather than use fixed values of N , an improved method for automatically deciding when the transform should cease is to monitor the resultant image and stop when it no longer changes significantly between iterations.

The best method that has been found to do this is to measure the difference in the normalised accumulator image between iterations. We use the root mean squared (RMS) difference between the intensities:

$$D^{<t>}(\mathbf{I}'^{<t>}, T) = \sqrt{\frac{1}{|\mathbf{I}'|} \sum_{i \in \mathbf{I}'} \left(\mathbf{I}'^{<t>}(i) - \mathbf{I}'^{<t-T>}(i) \right)^2} \quad (9)$$

where T is the number of iterations between each comparison. Rather than using successive images, the operation of a transform is sufficiently fast that large intervals should be used, that is, $T \gg 1$. The results produced by this method are consistent with the observed resultant images, with the size of D accurately reflecting the change that can be observed in the successive images.

The RMS difference measure is also cheap to calculate when T is set high enough, as expensive accumulator normalisations are calculated infrequently. It can be used as a stopping condition by setting a minimum value (D_S) to stop the process. From experimentation it has been found that a value for D_S of 1 when $T = 1000$ allows termination with a stable result.

3 Segmentation of Blood Vessels in Retinal Images

The ability of the image ray transform to highlight tubular curvilinear structures by turning them into optic fibres was shown in [14] and detection of blood vessels within eye fundus images is an appropriate application to further demonstrate this. As we do not use any advanced techniques for classifying the images that result from the transform, we do not expect accuracy superior to other techniques (although we try to establish that this is possible with superior classification techniques). We primarily intend to show that the transformed image is a suitable preprocessor for many retinal vessel extraction techniques, more appropriate than either intensity or edge detected images. Current techniques do not achieve accuracy of much more than 95%, and through use of an appropriate preprocessing technique such as the image ray transform we suggest that these results may be improved further.

3.1 Extraction Technique

We use the DRIVE database [4] to test our technique and compare to other techniques. The database consists of 20 test and 20 training fundus images, of which we only use the test images. The green channel of the retinal image is extracted for use as it provides the greatest contrast for blood vessels (figure 2(a)). The image is effectively inverted by the use of the target parameter (section 2.3), so that the lighter vascular structures are highlighted most strongly by the transform. Use of the transform also introduces additional problems. The edge of the fundus image acts as a circle and is highlighted by the transform; we expand the masks included in the database by 3 pixels in order to remove these unwanted features. The exponential transform also highlights the fovea strongly, which is undesired and so we used an automatic method (template matching with a Gaussian template) to find and remove them.

The image ray transform was applied to the green channel images (figure 2(b)) using both linear (referred to as RT-n) and exponential refractive indices (RT-k), as both gave results that differed significantly. Parameters used were $d = 256$, $t = 0$ and $D_s = 1$. For the linear refractive indices $n_{\max} = 40$ and for the exponential $k = 7$. We also tested on aggregated version of linear and exponential indices (RT-nk), so as to exploit the advantages of both techniques.

The results of the transform then had their histograms equalised (figure 2(c)) as this allows thresholds to be selected that are appropriate across all transformed images. Finally hysteresis thresholding (figure 2(d)) was performed to segment the image into vessel and background pixels. The upper threshold for hysteresis thresholding was set to 253 for all ray transform techniques, whilst the lower threshold was selected to give the highest accuracy across the database (RT-n: 234, RT-k: 235 and RT-nk: 230).

We compare pixels in the thresholded images with the ground truth, and calculate the performance of the technique. Two common metrics were calculated that are often used to test the strength of retinal blood vessel extraction: maximum average accuracy (MAA) and the area under the curve of a ROC graph (AUC).




Fig. 2. Steps taken to extract retinal blood vessels.

3.2 Results

This technique is capable of classifying pixels with accuracy comparable to a selection of contemporaneous unsupervised techniques. Figure 3 is a ROC graph showing the ability of each variant of our technique to classify vascular pixels correctly. Whilst all variants are significantly better than the untransformed images, the aggregation of the two different types produces a marked increase in discriminatory ability, suggesting that different vascular structures are emphasised by each. The transforms themselves do not have prohibitive computational cost: the linear transform took 4.65s on average whilst the exponential transform took only 1.91s on a 2.53GHz processor. The transform with exponential refractive indices was significantly faster as rays adhered to the vessels more strongly and so the resultant image had less noise and converged more quickly. These times are, however, many times quicker than the rest of the classification process took.

In figure 4 a number of transformed images and their thresholded versions are displayed. The ray transform with linear refractive indices (figures 4(a) and (d)) is adept at highlighting larger features and those near the optic disc but introduces some noise across the image. In contrast, using exponential refractive indices (figures 4(b) and (e)) forces rays to adhere to small vascular features more strongly, introducing less noise across the whole image. However it highlights the fovea in the centre of the image and fails around the optic disc. The highlighting of the fovea was removed before classification. These complementary results allow the aggregated images (figures 4(c) and (f)) to have the strengths of both and negates some of the weaknesses. Most vessels highlighted by either version are present in the thresholded combined version, and it has been improved by automated removal of the highlighted fovea. Results for different lower thresholds retain more vascular pixels, but misclassified more noisy background pixels as well, a classification technique less vulnerable to noise would improve results considerably.

Table 1 shows the results for our technique as well as a range of others, also on the DRIVE database. Using our simple classification technique, the ray transformed images achieve superior MAA and AUC values than the original




Fig. 3. ROC curve of the discriminatory ability of variants of our technique with the image ray transform.




Fig. 4. Selecting retinal blood vessels by variants of the ray transform and hysteresis thresholding (HT).

Technique	MAA	AUC	Notes
Ray Transform-n	0.9082	0.8100	Linear Refractive Indices
Ray Transform-k	0.9118	0.8139	Exponential Refractive Indices
Ray Transform-nk	0.9237	0.8644	Aggregated RT-k and RT-n
Original Image	0.8863	0.7151	Original, our classification technique
All Background	0.8727	-	Most likely class
Chauhuri et al. [1, 6]	0.8773	0.7878	Unsupervised
Zana et al.[2, 6]	0.9377	0.8984	Unsupervised
Jiang et al.[3, 6]	0.9212	0.9114	Unsupervised
Staal et al. [4]	0.9441	0.9520	Supervised
Soares et al. [5]	0.9467	0.9598	Supervised

Table 1. Maximum average accuracy (MAA) and area under the ROC curve (AUC) for our technique and others

intensity images, implying that the ray transform has emphasised the vascular features to a greater extent than they had been before. Our method is comparable to other unsupervised techniques, despite the simplicity of our classifier. As expected, supervised techniques produce superior results, but through use of the image ray transform as a preprocessor to a supervised learning method it should be possible to improve results further.

4 Discussion and Future Work

This work has shown that the image ray transform has considerable potential for use in enhancing detection of retinal blood vessels, and that performance comparable with other techniques can be achieved with use of a very simple classification technique. The transform is appropriate and successful at this task due to its inherent ability to detect tubular features such as blood vessels. Results show that the ray transformed images highlighted vascular features to a greater extent than they are in the original image, suggesting that use of the transform as a preprocessor of better classification methods will increase their performance. The transform is not computationally expensive in comparison with such classification techniques, and so does not increase their execution time significantly.

Future work will begin by applying the image ray transform as a preprocessor for current state of the art supervised and unsupervised methods, as it should be able to increase the performance of these techniques. Combining the two versions of the transform clearly improves results, but a better fusion would be able to use both results while excluding their weaknesses. Whilst the aim of this work was to detect blood vessels, it is clear that the ray transform also has potential to detect the fovea and the optic disc. This was a challenge in extracting vascular features, but detection of the fovea and optic disc are problems for which computer vision has been applied [16] and in the future the ray transform should also be evaluated for its ability to detect these features.

Acknowledgements

We gratefully acknowledge Alastair Cummings' EPSRC CASE studentship funded by the National Physical Laboratory (NPL).

References

1. Chaudhuri, S., Chatterjee, S., Katz, N., Nelson, M., Goldbaum, M.: Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE Transactions on Medical Imaging* **8** (1989) 263–269
2. Zana, F., Klein, J.: Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation. *IEEE Transactions on Image Processing* **10** (2001) 1010–1019
3. Jiang, X., Mojon, D.: Adaptive local thresholding by verification-based multi-threshold probing with application to vessel detection in retinal images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **25** (2003) 131–137
4. Staal, J., Abramoff, M., Niemeijer, M., Viergever, M., van Ginneken, B.: Ridge based vessel segmentation in color images of the retina. *IEEE Transactions on Medical Imaging* **23** (2004) 501–509
5. Soares, J., Leandro, J., Cesar Jr, R., Jelinek, H., Cree, M.: Retinal vessel segmentation using the 2-D Morlet wavelet and supervised classification. *IEEE Transactions on Medical Imaging* **25** (2006) 1214–1222
6. Niemeijer, M., Staal, J., van Ginneken, B., Loog, M., Abramoff, M.: Comparative study of retinal vessel segmentation methods on a new publicly available database. In: *Proc. SPIE Medical Imaging 2004*. (2004) 648–656
7. Al-Rawi, M., Qutaishat, M., Arrar, M.: An improved matched filter for blood vessel detection of digital retinal images. *Computers in Biology and Medicine* **37** (2007) 262 – 267
8. Zhang, B., Zhang, L., Zhang, L., Karray, F.: Retinal vessel extraction by matched filter with first-order derivative of Gaussian. *Computers in Biology and Medicine* **40** (2010) 438–445
9. Nixon, M.S., Liu, X.U., Direkoglu, C., Hurley, D.J.: On using physical analogies for feature and shape extraction in computer vision. *The Computer Journal* (2009)
10. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12** (1990) 629–639
11. Liu, X.U., Nixon, M.: Medical image segmentation by water flow. In: *Proc. Medical Image Understanding and Analysis (MIUA) 2007*. (2007)
12. Maragos, P.: PDEs for morphological scale-spaces and eikonal applications. In Bovik, A.C., ed.: *The Image and Video Processing Handbook*. 2nd edn. Elsevier Academic Press (2005) 587–612
13. Cummings, A.H., Nixon, M.S., Carter, J.N.: Circle detection using the image ray transform. In: *Int'l Conf. Computer Vision Theory and Applications (VISAPP 2010)*. (2010)
14. Cummings, A.H., Nixon, M.S., Carter, J.N.: A novel ray analogy for enrolment of ear biometrics. In: *4th IEEE Int'l Conf. on Biometrics Theory, Applications Systems (BTAS 10)*. (2010)
15. Hill, F.: 12. In: *Computer graphics using OpenGL*. 3rd edn. Prentice Hall (2000) 678
16. Niemeijer, M., Abramoff, M., van Ginneken, B.: Fast detection of the optic disc and fovea in color fundus photographs. *Medical Image Analysis* **13** (2009) 859–870