

The Gradient Product Transform for Symmetry Detection and Blood Vessel Extraction

Christoph Dalitz¹, Regina Pohle-Fröhlich¹, Fabian Schmitt¹ and Manuel Jeltsch¹

¹*Institute for Pattern Recognition, Niederrhein University of Applied Sciences, Reinarzstr. 49, 47805 Krefeld, Germany*
{christoph.dalitz, regina.pohle}@hsnr.de

Keywords: symmetry detection, image filtering, blood vessels, vesselness filter, angiography

Abstract: The “gradient product transform” is a recently proposed image filter for assigning each image point a symmetry score based on scalar products of gradients. In this article, we show that the originally suggested method for finding the radius of the symmetry region is unreliable, and a more robust method is presented. Moreover, we extend the symmetry transform to rectangular symmetry regions so that it is more robust with respect to skew, and the transform is generalised to also work with three dimensional image data. We apply the transform to two different problems: detection of objects with rotational symmetry, and blood vessel extraction from medical images. In an experimental comparison with other solutions for these problems, the gradient product transform performs comparable to the best known algorithm for rotational symmetry detection, and better than the vesselness filter for blood vessel extraction.

1 Introduction

A *symmetry* is a geometric transformation under which a subset of an image is invariant. Formally, when $f(\vec{x})$ is the pixel value at position \vec{x} , then a region Ω is symmetric under the transformation σ if

$$f(\sigma(\vec{x})) = f(\vec{x}) \quad \text{for all } \vec{x} \in \Omega \quad (1)$$

The class of the transformation σ defines the symmetry type: *reflection*, *rotation*, or *translation*. In Figure 1, the symmetry region Ω is the church window and it is invariant under reflections at the six yellow axes, and also under rotations with angles $\alpha = k \cdot 2\pi/6$ around the intersection point of the axes. Rotations by multiples of $2\pi/m$ form the *symmetry group* C_m , with the special case C_∞ denoting rotation invariance under arbitrary angles.

A survey of symmetry detection algorithms can be found in (Liu et al., 2010). Generally, these algorithms assign each image point a “symmetry score” that measures how suitable the point is as a symmetry centre. The image of symmetry score values thus represents a “symmetry transform”, and symmetry points correspond to maxima in the symmetry transform image. The computation of the symmetry score depends on the type of symmetry and is specific to each algorithm.

For detecting C_{2m} (point reflection) symmetries in greyscale images, Dalitz et al. have recently proposed

an algorithm that is based on scalar products of grey value gradients at corresponding positions. Let us call the resulting symmetry transform the *gradient product transform*. In (Dalitz et al., 2013), this transform was defined based on a square shaped mask and methods were given for determining the size of the symmetry region and for discriminating between rotational symmetries and axial symmetries. The latter can be useful for finding actual rotation symmetries, because reflection symmetries can also show up as C_2 symmetries (see Figure 2).

The gradient product transform was reported in (Dalitz et al., 2013) to perform comparable to or even



Figure 1: A symmetric object. The church window shows reflection symmetries according to the yellow axes and a C_6 rotation symmetry.




Figure 2: A C_2 symmetry that also is a reflection symmetry.

better than the method by Loy & Eklundh (Loy and Eklundh, 2006), an algorithm that performed best in several comparative studies (Rauschert et al., 2011; Park et al., 2008; Liu et al., 2013). Nevertheless, there are still two shortcomings of the gradient product transform as defined by Dalitz et al., which we address in the present article. The first is that its definition is based upon square shaped regions, which is not appropriate for non circular point symmetric objects and which also makes the symmetry transform very sensible to skew. We therefore generalise the transform to work with rectangular regions and give a recursion formula that allows for an efficient computation of the transform.

The other shortcoming is that the original method for determining the size of the symmetry region is not very robust. In (Dalitz et al., 2013), it is defined as the radius with the highest occurring symmetry score, which worked in the experiments because the radii used for testing were not greater than the largest symmetric object in each image. In this article we demonstrate however that the symmetry score tends to increase with the size of the trial region, even when the size goes beyond the actual symmetry region. In other words, the detected symmetry size for the original method is generally simply the largest radius tried out. We therefore propose a normalisation that remedies this problem.

Comparative studies on symmetry detection usually discuss and evaluate symmetry transforms as a means for locating symmetric objects within an unsymmetrical environment (Park et al., 2008; Rauschert et al., 2011; Liu et al., 2013), thereby only taking the location of maxima of a symmetry transform into account. From a more general viewpoint, the symmetry transform can alternatively be considered as an *image filter* that transforms a greyscale image into an image of symmetry scores. To demonstrate the usefulness of such a filter, we apply the gradient product transform as a preprocessing operation to retina fundus photographs and show that blood vessels can be easily extracted from this filtered image. In the preliminary study (Pohle-Fröhlich and Stalder, 2014), Pohle-Fröhlich & Stalder have demonstrated that the gradient product transform can be useful for

blood vessel extraction in 3D medical images. In the present article, we do a systematic evaluation of the gradient product transform performance on the 2D dataset of retina images by Staal et al. (Staal et al., 2004) and compare it to the standard algorithm for blood vessel extraction, the *vesselness filter* (Drechsler and Laura, 2010).

2 The gradient product transform

The *gradient product transform* first computes for each image point $\vec{x} = (x, y)$ of a greyscale image and for a range of radii $\vec{r} = (r_x, r_y)$ a symmetry score $s(\vec{x}, \vec{r})$. From these values, an “optimal” radius $\vec{R}(\vec{x}) = (R_x, R_y)$ is determined for each image point, and the final symmetry transform is then defined as $S(\vec{x}) = s(\vec{x}, \vec{R}(\vec{x}))$. In (Dalitz et al., 2013), only the case of square shaped symmetry regions was considered, i.e. $r_x = r_y$. Here we give formulae for the general case.

Section 2.1 describes the computation of $s(\vec{x}, \vec{r})$, Section 2.2 defines a criterion for selecting $\vec{R}(\vec{x})$, and in Section 2.3 we give a criterion for discriminating between axial and actual rotational symmetries.

2.1 Symmetry score

Although homogeneous regions comply with the symmetry definition (1), they are of little interest for symmetry detection, because objects are better characterised by their edges. The starting point of the symmetry score computation is thus the gradient image $\vec{G}(\vec{x})$, which can be computed from a greyscale image, e.g., with a Sobel filter. The symmetry score formula is based on two observations:

- under point reflection symmetry, the mirrored gradient \vec{G}' is anti-parallel to the gradient \vec{G} (see Figure 3)
- the scalar product $\langle \vec{G}, \vec{G}' \rangle$ is minimal for $\vec{G}' = -\vec{G}$

The symmetry score for the rectangle with circumradius $\vec{r} = (r_x, r_y)$ around the centre point $\vec{x} = (x, y)$,




Figure 3: Effect of a point reflection at \vec{x} on the gradient. The gradient \vec{G} at $\vec{x} + \vec{d}$ is mapped onto the gradient $\vec{G}' = -\vec{G}$ at $\vec{x} - \vec{d}$.

i.e., the rectangle with upper left $\vec{x} - \vec{r}$ and lower right $\vec{x} + \vec{r}$, is therefore defined as

$$s(\vec{x}, \vec{r}) = - \sum_{d_x=1}^{r_x} \left\langle \vec{G}(d_x, 0), \vec{G}(-d_x, 0) \right\rangle - \sum_{d_y=1}^{r_y} \sum_{d_x=-r_x}^{r_x} \left\langle \vec{G}(d_x, d_y), \vec{G}(-d_x, -d_y) \right\rangle \quad (2)$$

where $\vec{G}(d_x, d_y)$ denotes the gradient at position $(x + d_x, y + d_y)$. The sum is split up so that each point pair only occurs once, and the minus sign is added for convenience so that a higher symmetry corresponds to a higher score value, not vice versa. As the size of the symmetry region is not known in advance, the score (2) must be computed for all radii $1 \leq r_x, r_y \leq r_{max}$. All scores $s(\vec{x}, \vec{r})$ can be computed in a single run over the symmetry region with the following recursion formula for $r_x, r_y > 1$:

$$s(\vec{x}, (r_x, r_y)) = s(\vec{x}, (r_x - 1, r_y)) + s(\vec{x}, (r_x, r_y - 1)) - s(\vec{x}, (r_x - 1, r_y - 1)) - \langle \vec{G}(r_x, r_y), \vec{G}(-r_x, -r_y) \rangle - \langle \vec{G}(r_x, -r_y), \vec{G}(-r_x, r_y) \rangle \quad (3)$$

The computation of all scores up to a radius r_{max} thus only requires $O(r_{max}^2)$ additions, not $O(r_{max}^4)$ that would be required by a naïve application of Eq. (2). To generalise Eq. (2) to more than two dimensions, it is more convenient to notate the formula in a compact vector notation:

$$s(\vec{x}, \vec{r}) = -\frac{1}{2} \sum_{\substack{\vec{d} \in \Omega(-\vec{r}, \vec{r}) \\ \vec{d} \neq \vec{0}}} \left\langle \vec{G}(\vec{d}), \vec{G}(-\vec{d}) \right\rangle \quad (4)$$

where $\vec{G}(\vec{d})$ denotes the greyscale gradient at position $\vec{x} + \vec{d}$, and $\Omega(-\vec{r}, \vec{r})$ denotes the region of the rectangle or cuboid with corners $-\vec{r}$ and \vec{r} . The factor $1/2$ compensates for the fact that each pair of corresponding mirror points occurs twice in the sum. Eq. (4) holds for any dimension.

2.2 Symmetry size

When the symmetry score (2) is computed for a range of circumradii \vec{r} , the question arises which of these radii represents the size of the symmetry region. In (Dalitz et al., 2013), Dalitz et al. suggested to simply take the radius with the highest score, which would be for our generalised case of rectangular regions

$$\vec{R} = \arg_{\vec{r}} \max \{s(\vec{x}, \vec{r}) \mid \max(r_x, r_y) \leq r_{max}\} \quad (5)$$




Figure 4: Symmetry score as a function of the radius. On average, the symmetry score $s(r)$ increases with r even when r becomes larger than the actual size r_0 of the symmetric object. Division by a factor r^α can compensate for this effect.

The problem with this definition is that it tends to pick a radius quite close to r_{max} , even when the actual symmetric object is much smaller. To understand this phenomenon, we have computed the symmetry score as a function of the radius r with square shaped regions ($r_x = r_y = r$) for all ground truth symmetry points in the data set (Dalitz et al., 2013). To make the measurements comparable, we have normalised all radii with the ground truth radius r_0 , and all scores with the score $S(r_0)$. As can be seen in Figure 4 (curve for $\alpha = 0$), the score increases steeply for $r \leq r_0$, as expected, but still increases as r becomes larger than r_0 , albeit not as steep. This is simply due to the fact that the number of summands in (2) grows with r so that random fluctuations can increase the sum.

To remedy this problem, the symmetry score needs to be normalised by a factor that depends on \vec{r} . Even though the number of summands in (2) is proportional to the area of the symmetry region, which is $r_x \cdot r_y$, only scalar products with large gradients will contribute significantly to the sum. As the gradient is large for edges, and the number of edge points is proportional to the surface of the symmetry region, it seems natural to normalise with $r_x + r_y$. As can be seen from the curve for $\alpha = 1.0$ in Figure 4, this indeed leads on average to a maximum score at the actual symmetry size, but it also has the side effects that the score values of small symmetry regions are increased.

To control the trade off between these two effects, we normalise with $(r_x + r_y)^\alpha$, where $\alpha = 0.5$ is a compromise between both effects. The effect of the choice of α on the symmetry radius detection can be seen in Figure 5, where we have set r_{max} to the size of the image. We therefore define the optimum circumradius $\vec{R} = (R_x, R_y)$ and the resulting symmetry transform




Figure 5: Impact of α on the symmetry radius determination. Symmetry circumradius around the highest symmetry score in the image with $\alpha = 0$ (yellow), $\alpha = 0.5$ (green), and $\alpha = 1.0$ (red).

value S at position \vec{x} as

$$\vec{R}(\vec{x}) = \arg_{\vec{r}} \max \left\{ \frac{s(\vec{x}, \vec{r})}{(r_x + r_y)^\alpha} \mid \max(r_x, r_y) \leq r_{max} \right\}$$

$$\text{and } S(\vec{x}) = \frac{s(\vec{x}, \vec{R}(\vec{x}))}{(R_x(\vec{x}) + R_y(\vec{x}))^\alpha} \quad (6)$$

For more than two dimensions, the normalisation denominator $(r_x + r_y)^\alpha$ in Eq. (6) must be replaced with $(\sum_{i=1}^n |r_i|)^\alpha$ for the circumradius $\vec{r} = (r_1, \dots, r_n)$.

2.3 Symmetry type

As already pointed out in the introduction (see Figure 2 above), some axial symmetries can also be C_2 symmetries. A practical example can be seen in Figure 6: the highest local maximum in the symmetry transform corresponds to the region highlighted in red, which has both a C_2 and an axial symmetry, while the second highest local maximum in the symmetry transform corresponds to the C_∞ symmetry highlighted in green. In (Dalitz et al., 2013), a simple threshold based discrimination between these two types of symmetry was proposed, based upon the edge directedness in the neighbourhood of the symmetry centre (the threshold was set to 0.27).

In this section, we investigate this discrimination in more detail with a quadratic discriminant analysis. As can be seen from the symmetry transform image in Figure 6, for axial symmetries the entire symmetry axis tends to obtain high symmetry scores, while for pure rotational symmetries the symmetry score decays with the distance to the symmetry centre independent from the direction. We therefore compute the following features for a candidate symmetry centre \vec{x} , i.e., for a local maximum of the symmetry score $S(\vec{x})$:




Figure 6: Symmetry transform of a sample image. The two highest symmetry scores and the corresponding circumradii according to Eq. (6) with $\alpha = 0.5$.

Edge directedness. This feature is based on the angle histogram of the gradients of the symmetry transform in a $k \times k$ window centred at \vec{x} . We compute the gradients by applying a Sobel filter to the symmetry transform and build a weighted histogram with 16 bins of the gradient angles. Each angle is weighted in the histogram with the corresponding absolute value of the gradient. The edge directedness is then defined as the highest occurring frequency in this histogram, and it should be higher for axial symmetries.

Skeleton size. Starting from \vec{x} , we follow the skeleton until the score falls below a percentage p of $S(\vec{x})$. The feature is then the ratio between the skeleton length and the size $\sqrt{R_x R_y}$ of the symmetry region, where \vec{r} is the circumradius of the symmetry region. It should be higher for axial symmetries.

Antiparallel directions. This feature is based on the gradient of the original greyscale image. We compute the direction histogram with eight bins of all gradients in a window with the radius $\min\{R_x, R_y\}$ according to Eq. (6). Only those gradients are taken into account for which the mirrored gradient is “antiparallel”, i.e. the cosine of the angles between the gradients is less than -0.975. The feature is the highest relative frequency in the direction histogram. The value for “antiparallel direc-

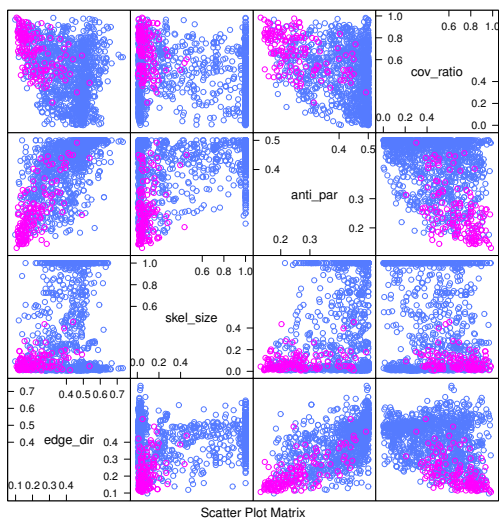


Figure 7: Scatterplots for the four features for symmetry type discrimination. The blue points belong to axial symmetries, the magenta points to rotational symmetries.

tions” should be lower for rotational symmetries.

Covariance eigenratio. For the points in a $k \times k$ window around (x, y) in the symmetry transform image, we compute the covariance matrix \mathbf{K} as

$$\mathbf{K} = \frac{1}{N} \sum_{dx=-k/2}^{k/2} \sum_{dy=-k/2}^{k/2} S(x+dx, y+dy) \times \begin{pmatrix} dx dx & dx dy \\ dy dx & dy dy \end{pmatrix} \quad (7)$$

where $S(x, y)$ is the symmetry transform value at (x, y) , and the normalisation factor N is the sum over all symmetry values in the window. The eigenvalues of \mathbf{K} indicate how strongly the values spread in the direction of the corresponding eigenvector. Consequently, the ratio between the smaller and the larger eigenvalue should be higher for rotational symmetry.

The data set (Dalitz et al., 2013) contains lists of ground truth points that are labelled as either axial or rotational symmetry centres, which provides a test set for this classification tasks. The scatterplots in Figure 7 actually reveal the tendencies in the above descriptions, but there is considerable overlap between the classes. Nevertheless, the scatterplots indicate that it is better to use more than one feature, and we therefore applied a quadratic discriminant analysis as implemented in the MASS package of the R statistical language¹, which decides for the class with the largest discriminant function

$$g_i(\vec{x}) = -w_i - \|(\vec{x} - \vec{\mu}_i) \cdot W_i\|^2 \quad (8)$$

¹<http://www.r-project.org>

where the index $i \in \{a, r\}$ denotes the class (axial or rotational) and \vec{x} is the row vector of the four features in the same order as in the list above. The parameters $\vec{\mu}_i$, w_i and W_i were evaluated by the R function *qda* as

$$\begin{aligned} w_a &= -15.03338 & w_r &= -21.49041 & (9) \\ \vec{\mu}_a &= (0.4334, 0.4975, 0.4417, 0.4352) \\ \vec{\mu}_r &= (0.1859, 0.0654, 0.2302, 0.7136) \\ W_a &= \begin{pmatrix} -9.3921 & 3.1451 & 7.8727 & 0.8677 \\ 0 & -2.4364 & 0.3522 & 0.1780 \\ 0 & 0 & -18.5275 & 3.1907 \\ 0 & 0 & 0 & 4.3365 \end{pmatrix} \\ W_r &= \begin{pmatrix} 14.3116 & -6.1912 & 15.5453 & -6.7814 \\ 0 & 20.3183 & 0.3754 & -0.9648 \\ 0 & 0 & -20.5954 & -3.8730 \\ 0 & 0 & 0 & -7.7489 \end{pmatrix} \end{aligned}$$

In our experiments in Section 3.1, this quadratic discrimination rule improved the recognition rates for rotational symmetries by about 8% compared to the simple decision rule based on the edge directedness alone.

3 Experimental results

We have evaluated the usefulness of the gradient product transform and compared it to other state-of-the-art algorithms for two use cases: detection of rotational objects in real world images, and extracting the blood vessel skeleton from two-dimensional retina fundus photographs. For the first problem, we have chosen the algorithm by Loy & Eklundh (Loy and Eklundh, 2006) as a reference, because it performed best in different comparison studies (Rauschert et al., 2011; Park et al., 2008; Liu et al., 2013). For the second problem, we have chosen the vesselness filter as a reference algorithm, because it is a widely utilised general purpose algorithm for blood vessel extraction. In the comparative study (Drechsler and Laura, 2010), Drechsler et al. demonstrated that each of the three vesselness filter variants (Frangi et al., 1998; Sato et al., 1998; Erdt et al., 2008) has its strengths and weaknesses. In our experiments, we have implemented the filter by Frangi et al., because the retina images were two-dimensional, and the filter by Frangi et al. is the only vesselness filter that can also be applied to 2D images in a straightforward way.

3.1 Rotational symmetry detection

To evaluate the detection of rotational symmetric objects, we have used the dataset that was used in the study (Dalitz et al., 2013), and that is freely available

	Count	square		rectangular		Loy & Eklundh
		$\alpha = 0$	$\alpha = 0.5$	$\alpha = 0$	$\alpha = 0.5$	
Detail	83	0.58	0.60	0.59	0.58	0.37
Context	76	0.16	0.30	0.22	0.32	0.37
Front	43	0.51	0.63	0.56	0.53	0.70
Light skew	57	0.39	0.51	0.46	0.53	0.39
Strong skew	59	0.27	0.29	0.27	0.32	0.12
Total	159	0.38	0.46	0.42	0.45	0.37
Area precision		0.39	0.70	0.37	0.80	

Table 1: Effect of the normalisation with α and the extension to nonsquare symmetry regions on the recognition rates for symmetric objects in the test set (Dalitz et al., 2013). For comparison, the results for the reference algorithm by Loy & Eklundh (Loy and Eklundh, 2006) are given. The last row gives the average area precision p_A according to Eq. (10) for the correctly detected objects.

on the Internet². This dataset consists of 159 images of size 600×400 , showing 27 different subjects. Each subject is shown in different perspectives and both in context and in detail. This allowed us to investigate how an algorithm performs under skew and contextual noise.

For each image, we have computed all local maxima of the gradient product transform ($r_{max} = 300$, i.e. the maximum possible radius) and, among these, searched for the highest value that was classified as “rotational” by the quadratic discriminant analysis described in Section 2.3. The algorithm took in total about 3.6 sec per image on an Intel i7-4770 CPU 3.40GHz. When the position of this maximum had a distance less than 10 pixels from a ground truth symmetry point, we considered a symmetry to be correctly detected, otherwise it was an error. The resulting recognition rates are shown in Table 1. Compared to the results in (Dalitz et al., 2013), which only used no normalisation ($\alpha = 0$) and square shaped regions ($r_x = r_y$), the recognition rates in the corresponding column are lower. This is because in the experiments of (Dalitz et al., 2013), r_{max} was set to the largest ground truth symmetry radius for each image, thereby introducing a priori bias, especially as for $\alpha = 0$ the symmetry radius tends to be selected as r_{max} .

The recognition rates in Table 1 show that the normalisation of the symmetry score with $(r_x + r_y)^{0.5}$ (see the columns for $\alpha = 0.5$) significantly improves the symmetry detection for the images in the “context” category, but has little effect for the “detail” category. This is because for the context images, the radii of the symmetry regions are typically much smaller than r_{max} , and are thus less likely to be found without normalisation ($\alpha = 0$). The detail images on the other hand are mostly filled out by the symmetric object so that the correct symmetry radius is close to r_{max} .

A comparison of the recognition rates for square

²<http://informatik.hsnr.de/~dalitz/data/visappl3/>

shaped regions versus rectangular regions shows that the extension to rectangular regions indeed improves symmetry detection for skewed symmetries, independent of the choice for α . For unskewed symmetries in the category “front” which actually have a circular symmetry region, the recognition rate for rectangular regions was however 10% smaller for $\alpha = 0.5$, due to the larger search space for the optimum circumradius (r_x, r_y) . Depending on the application, it is thus recommendable to restrict the ratio r_x/r_y to a maximum amount of permitted skew.

The recognition rates only measure whether the location of a symmetry centre was correctly found, but not whether the entire symmetry region has been correctly detected. We have therefore additionally measured the area precision of the detected circumradius for the correctly detected symmetries. Let G be the smallest rectangle enclosing the ground truth symmetry region, and T be the detected symmetry region with corners $\vec{x} \pm \vec{R}(\vec{x})$. Then we defined the *area precision* p_A as

$$p_A = \frac{|T \cap G|}{|T \cup G|} \quad \text{with } |X| = \text{area of } X \quad (10)$$

As can be concluded from the last row in Table 1, both the normalisation and the generalisation to rectangular regions clearly improve the accuracy of the shapes of the detected symmetry regions.

Compared to the algorithm by Loy & Eklundh, the gradient product transform is unaffected by light skew, but is less robust with respect to contextual noise. For the algorithm by Loy & Eklundh, it is just the other way round.

3.2 Blood vessel extraction

The problem of blood vessel extraction occurs for many types of medical images, e.g. for angiographs, fundus photographs, or MRT-TOF datasets. For blood vessel extraction from two-dimensional retina fundus photographs, Staal et al. have created an evaluation

dataset known as the DRIVE database (“Digital Retinal Images for Vessel Extraction”) (Staal et al., 2004). It consists of 40 colour images, divided into 20 training and 20 test images, with ground truth data manually extracted by medical experts (see Figure 8).

From the colour images, we have used the green channel as input to both our symmetry transform and the vesselness filter, because this channel had the highest contrast. It should be noted that Shahbeig recently suggested sophisticated preprocessing operations on these particular images which strongly enhance the visibility of the blood vessels (Shahbeig, 2013). As we were less concerned with the specific problem of retina photographs, but wanted to evaluate the vessel extraction on typical medical images, we restricted ourselves to a Gaussian blur with window size $k = 3$ and variance $\sigma^2 = 1.0$ as the only preprocessing operation on the green channel images; this preprocessing operation was only applied before the symmetry transform, but not before the vesselness filter because the latter does a Gaussian blurring anyway.

When applying the gradient product transform to blood vessel extraction, there are two issues:

- 1) Depending on the imaging technique, blood vessels are either darker or lighter than the surrounding tissue, but the gradient product transform invariably detects symmetries in both dark and light objects.
- 2) The gradient product transform results in a float image of symmetry scores and not in a skeleton image. The same holds for the vesselness filter.

To solve the first issue, we have modified the symmetry score from Eq. (4) to

$$s(\vec{x}, \vec{r}) = -\frac{1}{2} \sum_{\substack{\vec{d} \in \Omega(-\vec{r}, \vec{r}) \\ \vec{d} \neq \vec{0}}} \langle \vec{G}(\vec{d}), \vec{G}(-\vec{d}) \rangle \times H(\pm \langle \vec{d}, \vec{G}(\vec{d}) \rangle) \quad (11)$$

where H denotes Heavyside’s step function (zero for negative arguments, and one for positive arguments), and the sign in its arguments controls which type of




Figure 8: Sample image from the DRIVE dataset and its corresponding ground truth.




Figure 9: Impact of Heavyside’s step function on the gradient product transform. The left image is the result of the original gradient product transform, and the right image the result after the modification (11) (the float images have been converted to greyscale and the greyscale histogram was equalised).

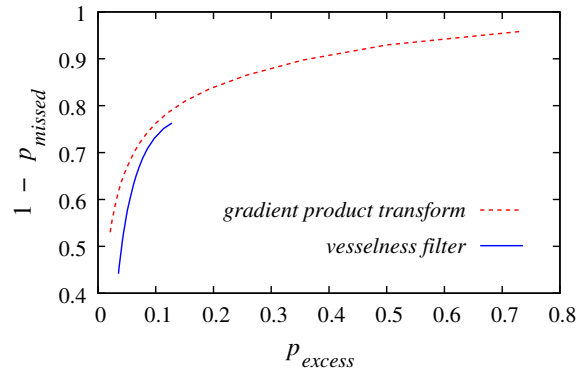


Figure 10: ROC-curves for blood vessel extraction on the DRIVE dataset. p_{missed} and p_{excess} have been averaged over all of the 40 images. Different points on the same curve correspond to different threshold values.

objects are detected. For dark vessels, the plus sign should be used because it has the effect that only gradients pointing outwards contribute to the sum. Figure 9 shows the impact of this modification.

Concerning the second issue, we have binarised the float image with a global threshold and extracted the medial axis transform skeleton from the binary image with the algorithm by Haralick & Shapiro (Haralick and Shapiro, 1992), as implemented in the Gamera framework³. To compare a resulting skeleton T with the skeleton S of the corresponding ground truth image, we computed the percentage of missed and excess pixels as

$$p_{missed} = \frac{|S \setminus (T \oplus H)|}{|S|} \quad (12)$$

$$p_{excess} = \frac{|T \setminus (S \oplus H)|}{|T|}$$

where “ $\oplus H$ ” denotes morphological dilation with a 3×3 structuring element H and $|X|$ denotes the number of pixels in X .

³<http://gamera.sf.net/>

There is a trade-off between the values for p_{missed} and p_{excess} , depending on the threshold value: for a lower threshold, the skeleton includes more pixels and p_{missed} is smaller, but p_{excess} is greater. To compare the performances of gradient product transform and vesselness filter, we therefore tried a range of thresholds and computed the receiver operating characteristics (ROC), which is shown in Figure 10. For the gradient product transform, we have chosen $r_{max} = 3$ and $\alpha = 0.5$. For the vesselness filter, we have tried all values $\sigma \in \{0.1, 0.2, \dots, 3.0\}$ and selected for each pixel the largest vesselness score.

From Figure 10, we conclude that the gradient image transform performs better than the vesselness filter for blood vessel extraction, because it yields more correctly detected vessels for the same rate of false positives. Note that the ROC curve for the vesselness filter is shorter, because the threshold cannot become less than zero. The same excess rate as with the vesselness filter for threshold zero was obtained with the gradient product transform for a threshold set to half the mean positive symmetry score.

4 Conclusions

The symmetry score normalisation for the gradient product transform, as suggested in this article, clearly improves the detection of the size of the symmetry region and removes the necessity of an approximate a priori guess of the maximum circumradius r_{max} . For the power α in the normalisation, we recommend $\alpha = 0.5$. The generalisation to rectangular regions makes the gradient product transform more robust with respect to light skew.

Moreover, we have demonstrated that the gradient product transform can be useful as an image filter beyond the narrow problem of symmetry detection. An additional factor in the symmetry transform that suppresses symmetry responses from light (or dark) regions makes the gradient product transform well suited for blood vessel extraction from medical images. To this end, it performs better than the vesselness filter, and can thus be a replacement for this standard algorithm for blood vessel extraction. As a global threshold for the symmetry score of blood vessels, 0.5 times the mean positive score worked well in our examples.

To enable others to easily build upon our results, we make the source code of our new generalised symmetry transform and for the blood vessel extraction freely available on our website⁴.

⁴<http://informatik.hsnr.de/~dalitz/data/visappl3/>

REFERENCES

- Dalitz, C., Pohle-Fröhlich, R., and Bolten, T. (2013). Detection of symmetry points in images. In *International Conference on Computer Vision Theory and Applications (VISAPP)*, pages 577–585.
- Drechsler, K. and Laura, C. (2010). Comparison of vesselness functions for multiscale analysis of the liver vasculature. In *Information Technology and Applications in Biomedicine (ITAB), 2010 10th IEEE International Conference on*, pages 1–5.
- Erdt, M., Raspe, M., and Sühling, M. (2008). Automatic hepatic vessel segmentation using graphics hardware. In *Medical Imaging and Augmented Reality, 4th International Workshop, MIAR 2008*, pages 403–412.
- Frangi, A., Niessen, W., Vincken, K., and Viergever, M. (1998). Multiscale vessel enhancement filtering. In *Medical Image Computing and Computer-Assisted Intervention — MICCAI’98*, pages 130–137.
- Haralick, R. M. and Shapiro, L. G. (1992). *Computer and Robot Vision*, volume 1. Addison-Wesley, Reading, MA.
- Liu, J., Slota, G., Zheng, G., Wu, Z., Park, M., Lee, S., Rauschert, I., and Liu, Y. (2013). Symmetry detection from real world images competition 2013: Summary and results. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Liu, Y., Hel-Or, H., Kaplan, C., and Gool, L. V. (2010). Computational symmetry in computer vision and computer graphics. *Foundations and Trends in Computer Graphics and Vision*, 5:1–195.
- Loy, G. and Eklundh, J. (2006). Detecting symmetry and symmetric constellations of features. In *European Conference on Computer Vision (ECCV)*, pages 508–521.
- Park, M., Lee, S., Chen, P., Kashyap, S., Butt, A., and Liu, Y. (2008). Performance evaluation of state-of-the-art discrete symmetry detection algorithms. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8.
- Pohle-Fröhlich, R. and Stalder, D. (2014). 3D-Symmetrietransformation zur Gefäßsegmentierung in MRT-TOF-Daten. In *Bildverarbeitung für die Medizin 2014*, pages 144–149.
- Rauschert, I., Brockelhurst, K., Liu, J., Kashyap, S., and Liu, Y. (2011). Workshop on symmetry detection from real world images - a summary. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sato, Y., Nakajima, S., Shiraga, N., Atsumi, H., Yoshida, S., Koller, T., Gerig, G., and Kikinis, R. (1998). Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical Image Analysis*, 2:143–168.
- Shahbeig, S. (2013). Automatic and quick blood vessels extraction algorithm in retinal images. *Image Processing, IET*, 7:392–400.
- Staal, J., Abramoff, M., Niemeijer, M., Viergever, M., and van Ginneken, B. (2004). Ridge-based vessel segmentation in color images of the retina. *Medical Imaging, IEEE Transactions on*, 23:501–509.