Using Line Segments as Structuring Elements for Sampling-Invariant Measurements

C.L. Luengo Hendriks, *Member*, *IEEE*, and
L.J. van Vliet, *Member*, *IEEE*

Abstract—When performing measurements in digitized images, the pixel pitch does not necessarily limit the attainable accuracy. Proper sampling of a band-limited continuous-domain image preserves all information present in the image prior to digitization. It is therefore (theoretically) possible to obtain measurements from the digitized image that are identical to measurements made in the continuous domain. Such measurements are sampling invariant, since they are independent of the chosen sampling grid. It is impossible to attain strict sampling invariance for filters in mathematical morphology due to their nonlinearity, but it is possible to approximate sampling invariance with arbitrary accuracy at the expense of additional computational cost. In this paper, we study morphological filters with line segments as structuring elements. We present a comparison of three known and three new methods to implement these filters. The method that yields a good compromise between accuracy and computational cost employs a (subpixel) skew to the image, followed by filtering along the grid axes using a discrete line segment, followed by an inverse skew. The staircase approximations to line segments under random orientations can be modeled by skewing a horizontal or vertical line segment. Rather than skewing the binary line segment we skew the image data, which substantially reduces quantization error. We proceed to determine the optimal number of orientations to use when measuring the length of line segments with unknown orientation.

Index Terms—Mathematical morphology, granulometry, rotation invariance, translation invariance.

1 Introduction

THE basic operations in mathematical morphology, dilation and erosion, use a structuring element as filtering kernel. Flat structuring elements (i.e., sets) specify a neighborhood over which the operation has to take place (maximum or minimum filtering).[1] In the continuous domain, the set usually forms a compact shape such as a disk. In the discrete domain, the set is composed of grid points (pixels), often chosen to approximate a continuous-domain shape. Because of this, the operations in the two domains do not produce the same results. Note that this is different from continuous-time (or space) linear filters where the convolution integral between a band-limited signal and a band-limited impulse response can be replaced by sampling, discrete convolution, and reconstruction [1, Section 8.4]. We refer to this as sampling invariance. Most authors discuss the theory of morphological operations in the continuous domain, and present implementations and applications in the discrete domain. We combine these two domains, trying to obtain a discrete operation that, when applied to a properly sampled, band-limited image, yields the same result as the continuous-domain morphological filter. Unfortunately, this can only be approximated [2], [3]. This approximation can be significantly improved by interpolation [2]: First, we interpolate by a factor n (to reconstruct the continuous-domain input image from its samples [4]) followed by the morphological filter using an n-times larger structuring element (this reduces the influence of discretization effects at the borders of the structuring element). Any difference between the desired continuous result and the discrete result obtained in practice we call "discretization errors." In this paper, we aim to minimize the discretization errors for a very difficult class of structuring elements: line segments. Line segments as structuring elements typically produce a large discretization error because of their unfavorable surface to perimeter ratio; the discretization errors that occur when sampling a binary shape are always located in a 1-pixel wide strip along the boundary of the shape. This paper extends our earlier work [5], improving upon the performance evaluation, providing insight in computational complexity, and offering access to the implementation of the presented methods. Note that, unless explicitly stated, "translation invariance" is used to indicate invariance with respect to shifts in the continuous domain, that is, to (subpixel) shifts of the continuous image before sampling. "Compatibility under rotation" and "rotation invariance" are used in an analogous way.[2]

Using line segments as structuring elements we can construct rotation-invariant morphological operators, i.e., operators that adapt themselves to the local structure [6], [2]. It is for these operators that we study the implementation of discrete line segments here. Examples of other applications of line structuring elements are road detection in airborne images [7], [8], finding grid patterns on stamped metal sheets [9], and estimation of structure orientation [10], [11].

In Section 2, we summarize three known and three new implementations of the dilation or erosion with line structuring elements. All of these approaches are compared in Section 3, where they are tested for translation invariance (in the continuous-domain sense). Section 4 shows a typical application that demonstrates this translation invariance. Section 5 addresses the question of how many orientations of the line structuring element should be used when the orientation of the lines in the input image is not known.

2 Methods for Implementing Line Structuring Elements

The most simple implementation of the dilation with a line segment as structuring element uses a Bresenham line segment [12], for which an efficient, recursive algorithm by van Herk [13], [14] exists. The drawback of this operation is that it is not even translation invariant in the discrete sense (i.e., invariant over integer pixel shifts). The reason is that there are multiple Bresenham lines possible for a given orientation, which differ in where along the line one starts to count. Using periodic lines [15], it is possible to construct recursive dilations that are translation invariant in the discrete sense. This only yields a speed advantage over the basic implementation using Bresenham lines for certain orientations and lengths. Another solution, proposed by Soille and Talbot [11], combines multiple recursive openings or closings to obtain *discrete* translation invariance. This operation is closer to continuous-domain translation invariance than the basic implementation. However, it is limited to orientations for which the periodic line [15] has a low periodicity, due to the large number of openings or closings that would otherwise be required. This means that in practice the orientation and length will be rounded to more convenient values.

In the next three sections, we will propose alternative operations that improve on the ones mentioned above in terms of translation invariance and compatibility under rotation. They accomplish this by interpolating in the input image (Sections 2.1 and 2.2) or by using gray-value structuring elements (Section 2.3).

1. We do not make a distinction in this paper between the maximum and the supremum over a set since the sets we consider are closed and the functions are smooth.

2. Compatibility under rotation means $\Psi R = R\Psi$, whereas rotation invariance means $\Psi R = \Psi$, using R as the rotation operator.

- *C.L. Luengo Hendriks is with the Life Sciences Division, Lawrence Berkeley National Laboratory, One Cyclotron Road, MS 84R171, Berkeley, CA 94720. E-mail: clluengo@lbl.gov.*
- *L.J. van Vliet is with the Quantitative Imaging Group, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands. E-mail: L.J.vanVliet@ph.tn.tudelft.nl.*

Manuscript received 1 July 2004; revised 11 Feb. 2005; accepted 1 Apr. 2005; published online 14 Sept. 2005.
Recommended for acceptance by J. Goutsias.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0332-0704.

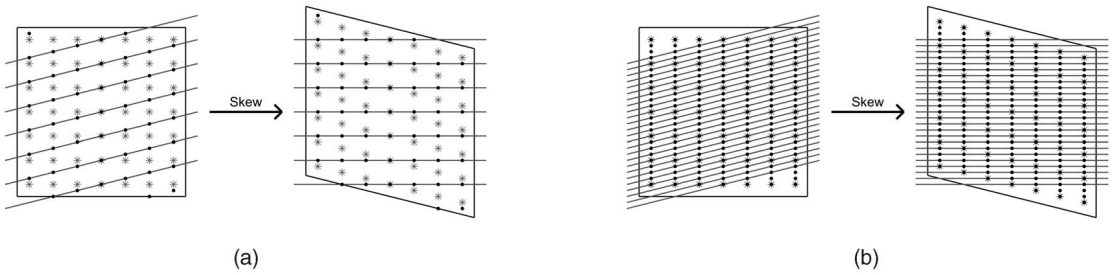

Fig. 1. After skewing the image, horizontal lines correspond to lines under a certain orientation with respect to the image data. In (a), (Interpolated line segments, Section 2.1), some of the original image samples ($*$) fall exactly on these lines, but most samples used ($\cdot$) lie in between original grid points. The value at these pixels is obtained by interpolation. In (b), (True interpolated line segments, Section 2.2), each of the output pixels is directly computed to avoid interpolation of the output of the morphological operation. This does increase the computational cost of the operation.

2.1 Interpolated Line Segments by Skewing the Image

Operations along a Bresenham line can be implemented by skewing the image, applying the operation along a column (or row), and skewing the image back. Here, we consider this operation using image skews with interpolation (that is, the rows or columns of the image are shifted by subpixel quantities); see Fig. 1a.

The interpolation method used is an important factor in the correctness of the output. The better the method is, the smaller the error will be. We used cubic convolution to implement the skews. This method is a good compromise between accuracy, computational cost and window size [16].

Note that—in contrast to the Bresenham lines—all samples lie exactly on the line, which yields a better approximation to translation invariance. The drawbacks are that the result needs to be skewed back, which again requires interpolation, and that the sample spacing along the line depends on the orientation. The former is serious because the result of the discrete operation does not have a direct relation to the result of the continuous operation. Interpolation in this image can potentially introduce undesired image features.

2.2 True Interpolated Line Segments

The skewing method presented above yields the best accuracy on only those columns (or rows) of the output image with an integer shift (thereby avoiding interpolation). This typically occurs for only a few columns (the distance between these columns is given by the periodicity P_φ of the Bresenham line of the given angle [15]). It is possible to obtain the same accuracy for all output pixels by skewing the image multiple times, or equivalently, increasing the sampling density of each column, in such a way that the operation gives an output value at each of the output locations. This is represented in Fig. 1b.

We implemented this method by skewing the image once for each column, changing the offset of the skew, and using only the one output column that corresponds to that offset. A more efficient implementation would use all those columns that have an integer shift, so that only P_φ skews are required; this would reduce the number of operations substantially for a selected set of angles (those that yield rational slopes).

For all discrete line segments mentioned up to now, the number of samples used in the computation of the morphological operation depends not only on the length of the segment, but also on the orientation. Line segments along the grid are sampled with the highest density; diagonal segments are sampled with the lowest density. Thus, for some orientations there is a higher probability to miss a local maximum (i.e., the maximum falls in between samples) than for others. This makes the approximation to continuous-domain translation invariance better for horizontal and vertical lines than for diagonal lines. This difference also has repercussions for the compatibility under rotation. Ideally, one would like to sample all of these lines with the same density. Rotating the image instead of skewing also alleviates this problem. However, when rotating, only a limited set of samples falls exactly on output samples and, in the worst case, this happens only for the sample in the origin of the rotation. This means that an even larger number of operations is required to compute the result of the operation at all output pixels.

2.3 Band-Limited Line Segments

Another option to overcome sampling problems with discrete line segments is to use smooth (approximately bandlimited) gray-value structuring elements. Such a segment is sampling invariant and does not have a limited set of available lengths. Even though the morphological operation itself is still not sampling invariant (meaning $\Psi(f)\cdot \mathrm{III} \neq \Psi(f\cdot \mathrm{III})$, where III is the pulse train used for sampling), it is expected that such a structuring element reduces the discretization errors.

A band-limited shape has a smooth edge, such that the continuous-domain shape can be reconstructed from the samples. This implies that the position of the boundary is known with subpixel accuracy. We obtain smooth edges by convolving the ideal, continuous-domain shape with a Gaussian function, which is approximately band-limited if the sample spacing is smaller or equal to σ [17]. Starting from an infinitely thin line segment, we obtain

$$L_{(\ell,\sigma)}(x,y) = A\cdot\frac{1}{2}\left\{1-\mathrm{erf}\left(\frac{\ell-2|x|}{2\sigma}\right)\right\}\cdot\exp\left(\frac{-y^2}{2\sigma^2}\right), \qquad (1)$$

where ℓ is the length of the line segment, x is the coordinate-axis in the direction of the segment and y is the coordinate-axis perpendicular to it. Fig. 2a shows an example of such a band-limited line segment, and Fig. 2b gives some details on its construction. Note that the gray-value of the segment is 0, and the background has a value of $-A$. A is the scaling of the image $L_{(\ell,\sigma)}$, and depends on the gray-value range in the image to be processed. A must be larger than the gray-value range of the image, so that the background of the structuring element never interacts with the image. Even though the image $L_{(\ell,\sigma)}$ is approximately band-limited for any A, its slopes are

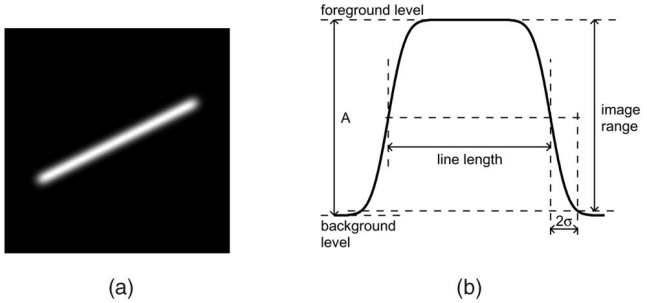

Fig. 2. An approximately band-limited line segment constructed with (1), to be used as gray-value structuring element. The difficulty lies in the normalization (with respect to the image range) and defining the length of the line. (a) Gray-value structuring element. (b) Construction.

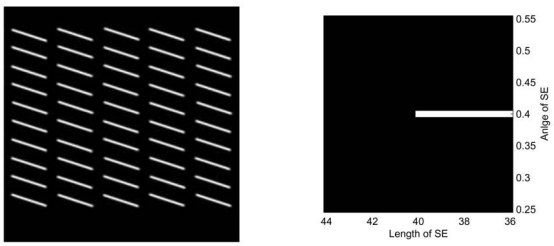

Fig. 3. (a) Input image and (b) ideal response for graphs shown in Fig. 4. A value of 1 is expected for the openings in which the orientation of the structuring element matches that of the segments in the input image (0.4 rad), and the length ℓ is smaller or equal to the length of these segments (40 pixels). A value of 0 is expected for any other parameter of the structuring element. Because the input image is band-limited, a smooth transition is actually obtained.

not invariant to gray-value scaling. Since morphological operations can be written as an interaction between slopes [18], it follows that this scaling has an influence on the result of the operation. By relating the value of A to the range of gray-values in the image, the operation is invariant to gray-value scaling of the image, but not invariant to, e.g., impulse noise (which increases the gray-value range) or gray-value scaling of individual objects in the image. We will be using a value of 1.0233 times the gray-value range of the input. This value is chosen to make the cut-off point of the structuring element image to be $|x| \leq \frac{1}{2}\ell + 2\sigma$ (this gives the size of the image needed to generate the structuring element).

3 Comparison of Implementations for Line Structuring Elements

We have implemented the following versions of the opening with a line segment structuring element:[3]

- Method 1: Bresenham line segment.
- Method 2: recursive algorithm for the Bresenham line segment [14].
- Method 3: composite discrete line segment (union of openings) [11].
- Method 4: interpolated line segment by skewing the image (Section 2.1).
- Method 5: true interpolated line segment (Section 2.2).
- Method 6: band-limited line segment (Section 2.3).

To compare the different methods, an image was generated that contains many line segments of fixed length $\ell_0 = 40$ pixels and orientation $\varphi_0 = 0.4$ rad (not a rational slope), but random subpixel position (see Fig. 3a). These line segments were drawn using (1). Openings with each of the implemented methods were applied to this image, varying both the length and orientation of the structuring element. The result of each operation is integrated, normalized, and plotted in a graph. The result for these operations should be

$$\frac{\int [L_{(\ell,\varphi)}(f)](x)dx}{\int f(x)dx} = \begin{cases} 1 & (\varphi = \varphi_0) \wedge (\ell \leq \ell_0), \\ 0 & \text{elsewhere}, \end{cases} \quad (2)$$

as in Fig. 3b. Since the image contains band-limited lines, we expect the result to obtain a smooth transition from one state to the other. The more the result approximates the ideal situation, the better the specificity of the operator.

These results are plotted in Fig. 4. Table 1 shows the time needed to compute the data in Fig. 4 for each algorithm, as well as the computational complexity of these algorithms, as discussed in Section 2. We have timed the algorithms for different values of φ, L, and N, and verified these complexities. A number of performance-related observations can be obtained from these graphs and table:

- The simple discrete, noninterpolated implementations (methods 1 and 2) never reach values approximating 1. The composed, interpolated, and gray-value methods (methods 3, 4, 5, and 6) reach higher values, close to the ideal value of 1.
- All methods except method 6 show a staircase-like dependency on the length. This is because of the inherent discretized length of these segments. Note that the actual length of the structuring element depends on the orientation. In methods 1 and 3, the angle is discrete as well because it is limited by the possible orientations of a discrete line of finite length; this results in plateaus in which the structuring element does not change.
- There are very few differences between methods 4 and 5, which employ interpolation.
- The result of the gray-value method is very smooth, but shows some "ringing." This can be explained by the sampling of the structuring element and the image. Results of morphological filtering depend on the position of local maxima or minima in relation to the sampling grid. Small changes in the orientation of the line cause that a different set of grid points will sit close to maxima or minima (i.e., the ridge of the line).
- The computational cost of methods 3, 5, and 6 is one to two orders of magnitude higher than the cost of the other three methods. For method 3, this is on average over all the angles used in this experiment. This method has a running time that changes wildly with the chosen angle and length. For some specific line segments, method 3 yields a highly efficient algorithm that is comparable in execution time to method 2.

Taking these observations into account, it can be said that methods 3, 4, 5, and 6 produce results that come much closer to the expectations than the simple discrete methods. Method 3 has the disadvantage of discrete angles and a strongly angle-dependent running time. Method 6 is computationally more expensive and it depends on the application if its superior results warrant this expense. Finally, it does not appear to be necessary to use method 5, since it produces a result almost identical to method 4. Method 4 is much simpler to implement and computationally cheaper.

If the implementation of the opening satisfies the absorption criterion, the graph should be decreasing for increasing length and fixed angle. By looking at the difference between subsequent points in each graph, it appears that only methods 2 and 5 satisfy the absorption criterion, as expected: For a given angle, increasing the length of the structuring element only adds pixels to it, unlike in method 1; for method 5 this also implies that the set of skews used do not change with the length of the structuring element, meaning that the interpolation errors are identical as well. Method 4 only violates the absorption criterion in a few isolated cases, due to the interpolation in the result of the morphological operation. Method 6, even though it produces a very smooth result, only satisfies this criterion by approximation. Finally, our implementation of methods 1 and 3 only satisfies the absorption criterion for lines along the cardinal directions of the grid. Increasing the length of an arbitrary line segment causes a slight change of the apparent angle (we round both the length and the angle to the nearest possible discrete line segment), which causes a violation of the absorption property. It is possible to implement method 3 to round only the length, keeping the angle fixed. In this implementation, the absorption property is satisfied. Most angles yield an irrational slope, giving the Bresenham line an infinite periodicity. This means that the number of openings over which the union is taken is given by the image size. When rounding the angle, the slope is made rational, which results in a finite periodicity and a much smaller number of openings. We must stress that none of these alternative implementations have a significant effect on the results shown in Section 4.

3. Available online at http://clluengo.lbl.gov/line_se.html.

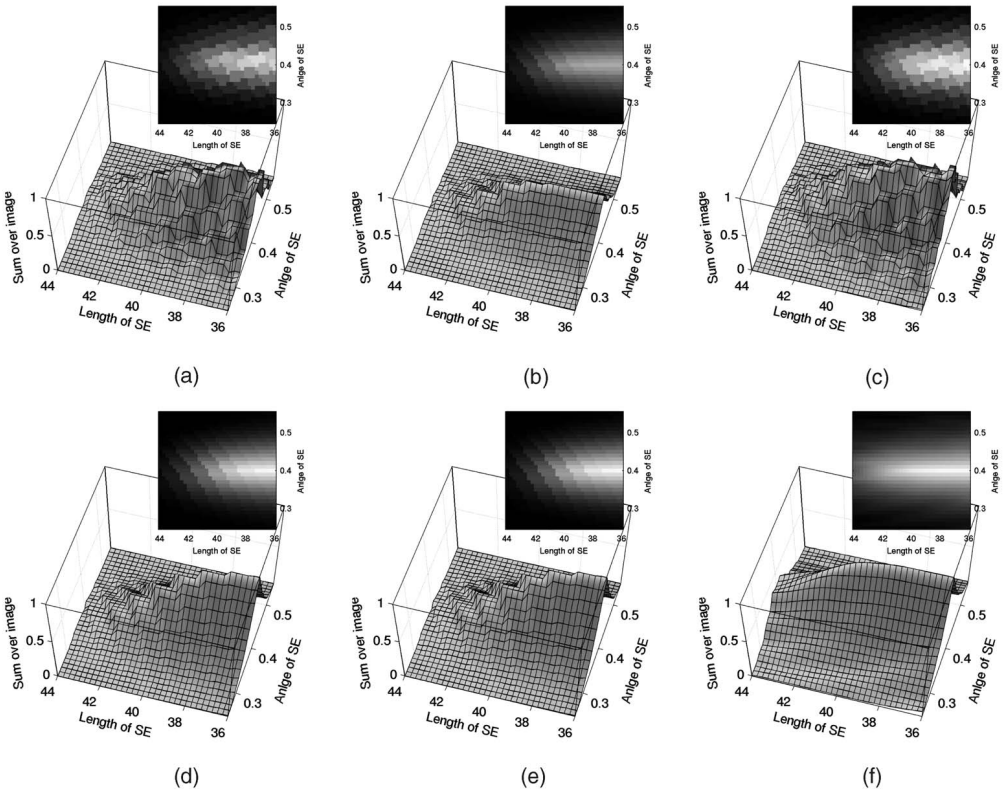

Fig. 4. Comparison of different implementations of the opening with a line segment structuring element. (a) Method 1. (b) Method 2. (c) Method 3. (d) Method 4. (e) Method 5. (f) Method 6. See text for details.

TABLE 1
Computational Complexity of the Tested Algorithms and the Time in Seconds required to Compute Each Graph of Fig. 4

method description	method	complexity	total time (sec)
Bresenham line segment	1	$\mathcal{O}(NL)$	40
recursive algorithm	2	$\mathcal{O}(N)$	20
composite discrete line segment	3	$\mathcal{O}(NP_\varphi)$	500
interpolated line segment by skewing the image	4	$\mathcal{O}(N)$	30
true interpolated line segment	5	$\mathcal{O}(NW)$	1900
band-limited line segment	6	$\mathcal{O}(NL\sin(\varphi)\cos(\varphi))$	800

N is the number of pixels in the image, L is the length of the structuring element and φ is its angle, P_φ is the periodicity of a Bresenham line under angle φ, and W is the width in pixels of the image if $|\varphi| < \pi/4$, its height otherwise.

4 Application

We have applied the closing with these methods to a series of nine images, recorded with a wide-field microscope (Zeiss Axioplan, with a Zeiss Fluar 10x lens and a Millix MicroImager camera). The stage (Ludl Electronics) was moved back and forth between recordings, to obtain subpixel shifts of the subject. Fig. 5a shows one of these images. Using image number 5 as the reference image, we determined the shift of the other eight images using the iterative, gradient-based shift estimation method as detailed in [19]. The closing with the six methods studied above was applied to all images, using a line segment of length 15 pixels, parallel to one of the sets of lines on the test slide, to obtain images similar to Fig. 5b (the angle was determined manually on the reference image by drawing a long discrete line; it had a slope of $-472/166$, corresponding to an angle of -0.3379). For each method, we shifted and compared the results with the result on the reference image, using as an error measure the mean square difference (computed only in a small region around the lines that are being detected); these numbers are shown in Table 2. The smaller these differences, the better the method is at approximating continuous-domain translation invariance, as discussed in the introduction. Note that the closing should decrease the difference because it removes noise. Methods 4 and 5 produce the smallest error, and there is not much difference between the two. Method 3 is somewhat better than method 1, but still increases rather than decreases the error. Method 2, lacking translation invariance in the discrete sense, dramatically increases its error for larger shifts between the images. Finally, method 6 does not perform as well as expected, both because its larger width causes it to be more sensitive to noise and because of the uneven illumination across the image, which makes the scaling factor A (chosen globally) to be too large for the features being filtered. This shows that, even though this method seems good in theory, it is not very useful in practice.

5 Angular Selectivity

5.1 Rotation Invariant Morphology Using Anisotropic Structuring Elements

In [6], [2], we introduced a rotation-invariant anisotropic (RIA) morphology. It is constructed by decomposing an isotropic structuring element D, i.e., a disk in two dimensions, into its

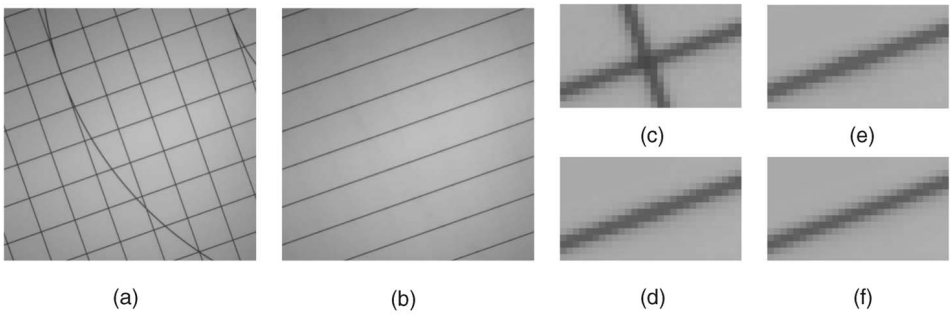

Fig. 5. (a) An image of a test slide (512 by 512 pixels). (b) Closing with a line structuring element $\ell = 15$. (c) Input image. (d) Result of method 3. (e) Result of method 1. (f) Result of method 4. The details show the differences between the various results.

diameters, L_φ. This is generalized as an infinite set of rotated versions of a lower-dimensional structuring element; for example, a ball in three dimensions can be decomposed into rotated versions of a disk or a line. In this paper, we limit ourselves to two-dimensional images, such that

$$D = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq r^2\}, \tag{3}$$

where r is the radius of the disk and

$$L_\varphi = \left\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq r^2 \wedge \tan(\varphi) = \frac{y}{x}\right\}. \tag{4}$$

The opening in this framework, which we denote by γ^*, corresponds to a supremum of openings, which is, as proven by Matheron [20], an algebraic opening:

$$\gamma^*_L(f) = \bigvee_\varphi \gamma_{L\varphi}(f). \tag{5}$$

Note that such an operator was already proposed in [20] and has been used by many authors afterward. Being an algebraic opening, the increasingness and antiextensivity properties are satisfied. We had proven earlier that the RIA opening (in the continuous domain) also satisfies the absorption property [2]. All of these three properties are required for an operator to be used in a granulometry [21], [22]. However, a discretized version of this operator will only be absorbing by approximation.

5.2 Discretizing the RIA Morphology

Implementation of the algebraic opening of (5) requires that we apply our line segments at a finite number of orientations. Here, we determine how many orientations are required. Since the morphological orientation-space $F(x, \varphi) = [\gamma_{L\varphi}(f)](x)$ is not bandlimited along the φ-axis we cannot apply the Nyquist condition for sampling and reconstruction. Instead, we will look for the number of orientations that will reduce the error below a certain level. This number is directly proportional to the circumference of the disk that is being probed by the rotated line segments.

We use method 4, interpolated line segments, to implement the openings. We experimentally verified that the conclusions are valid for the other methods as well. In Fig. 6, we plotted the full width at half the maximum (FWHM) of peaks in the orientation-space (as resulting from the input image in Fig. 3a). This is a good measure for the width of the response as a function of the line segment length. The figure shows an inverse relation between segment length and response width, which confirms the direct relation between the segment length ($\sim$ circumference) and the number of orientations that are needed.

To get an approximation to the number of orientations required, we will assume that the distance along the orientation-axis between samples is given by q/ℓ (in rad), ℓ being the segment length and q being some constant. This is equivalent to taking $\lfloor \pi\ell/q \rfloor$ samples between 0 and π rad. By varying the value of q and computing a granulometry, we find a suitable value for q such that decreasing it further has little effect on the measurement result. The granulometry was applied to an image similar to the one in Fig. 3a, but with the line segments of random orientation. We applied RIA openings with interpolated line segments (method 4). The results are shown in Fig. 7. For values of $q \leq 0.16$, the graph is quite consistent. This implies around 800 orientations should be used when computing a RIA opening with a line of 40 pixels in length.

Fig. 7 shows a granulometric line (for small q) that is negative at smaller ℓ. As stated in Section 3, the method used here to implement the openings does not satisfy the absorption criterion because of the interpolation in the output of the morphological filter. This same error is the one causing these negative values, which are not produced by methods 2 or 5.

If this experiment is repeated for a method that has a discrete set of possible angles (method 1 or 3), a larger value for q would be obtained. Increasing the number of orientations would only cause

TABLE 2
Mean Square Error (MSE) between Images Before and After Applying the Closing with the Different Methods

image #	MSE	MSE after closing with method #						measured shift	
	input	1	2	3	4	5	6	x	y
1	5.80	15.78	32.78	10.46	0.99	0.71	4.47	-0.374	0.364
2	5.25	13.48	23.96	8.98	0.96	0.82	4.08	-0.310	-0.695
3	2.01	3.78	2.07	2.55	0.46	0.56	1.57	0.029	-0.114
4	5.43	15.33	8.93	10.22	0.60	0.74	4.53	0.062	-0.348
6	2.27	4.72	2.68	3.15	0.49	0.60	1.83	0.038	-0.136
7	3.91	10.50	5.61	6.88	0.53	0.68	3.28	0.019	-0.250
8	5.28	14.61	8.96	9.83	0.70	0.87	4.39	-0.070	0.343
9	4.35	11.95	7.08	7.84	0.55	0.67	3.64	-0.032	-0.282

All images are compared against image 5. The MSE is computed in a small area around the lines that are expected not to change.

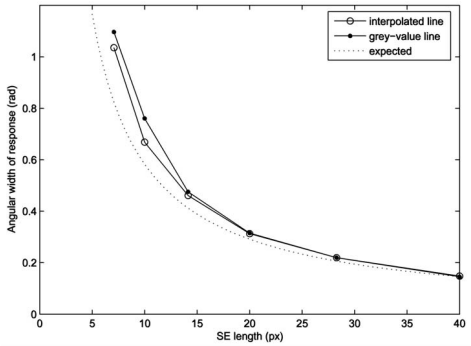

Fig. 6. The width (FWHM) of the response along the orientation-axis of the morphological orientation-space $F(x, \varphi)$, together with a line FWHM $\propto \ell^{-1}$. The disagreement for shorter line segments is due to the band-limit of the objects in the input image.

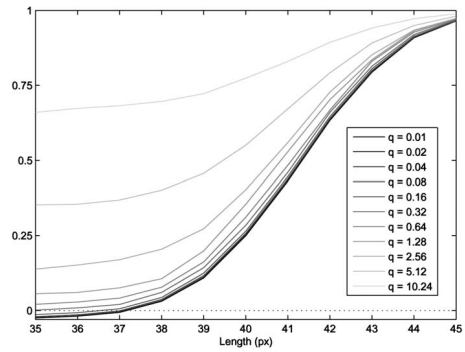

Fig. 7. Granulometries computed using RIA openings with different number of orientations sampled. The number of orientations is given by $\pi\ell/q$.

the same angles to used more often because they are rounded to the nearest possible values. The result of the operation is not changed.

6 Conclusions

We studied and compared different implementations of morphological operations with line structuring elements, under the assumption that the input image is band-limited and properly sampled. All differences between the result of the discrete operation and the result expected for the equivalent continuous-domain operation were considered discretization errors. We showed that the two methods that use interpolation produce notably smaller discretization errors than the methods that are commonly used. We suggest the use of interpolated structuring elements obtained by (subpixel) skewing of the input image (with interpolation by cubic convolution), applying the morphological operation along one of the grid axes, and then skewing the image back. This method provides a good compromise between accuracy of the result and computational cost. Interpolating the output image of the morphological operation is not correct since it is not a properly sampled version of a band-limited continuous image. However, for the test image we used this had only a small negative effect, as seen by the failure of the absorption criterion and the negative values in Fig. 7. We showed the validity of this method by applying it in a typical application.

To roughly determine the number of orientations that should be used when applying a linear structuring element under all orientations, we computed a granulometry with RIA morphology. We conclude that about $\lfloor 6.5\pi\ell \rfloor$ orientations should be used for a structuring element of length ℓ. This value is far too large for real-world applications, where a constrained is often put on the execution time, yielding a larger error.

Acknowledgments

The authors would like to thank Pierre Soille for the discussions and valued comments and Rodrigo Fernandez Gonzalez for technical assistance. This research was partially supported by the Dutch Ministry of Economic Affairs under project number IOP-IBV98006, and by the US Department of Energy under contract number DE-AC03-76SF00098.

References

[1] A.V. Oppenheim, A.S. Willsky, and I.T. Young, *Signals and Systems.* London: Prentice-Hall Int'l, 1983.

[2] C.L. Luengo Hendriks, "Structure Characterization Using Mathematical Morphology," PhD dissertation, Delft Univ. of Technology, Delft, The Netherlands, 2004, http://www.qi.tnw.tudelft.nl/Publications/phd_theses.html.

[3] C.L. Luengo Hendriks and L.J. van Vliet, "Basic Morphological Operations, Band-Limited Images and Sampling," *Proc. Scale Space Methods in Computer Vision,* pp. 313-324, 2003.

[4] H. Nyquist, "Certain Topics in Telegraph Transmission Theory," *Trans. AIEE,* pp. 617-644, 1928.

[5] C.L. Luengo Hendriks and L.J. van Vliet, "Discrete Morphology with Line Structuring Elements," *Proc. Computer Analysis of Images and Patterns,* pp. 722-729, 2003.

[6] C.L. Luengo Hendriks and L.J. van Vliet, "A Rotation-Invariant Morphology for Shape Analysis of Anisotropic Objects and Structures," *Proc. Fourth Int'l Workshop Visual Form,* pp. 378-387, 2001.

[7] J. Chanussot and P. Lambert, "An Application of Mathematical Morphology to Road Network Extractions on SAR Images," *Math. Morphology and Its Applications to Image and Signal Processing,* pp. 399-406, Dordrecht: Kluwer, 1998.

[8] A. Katartzis, V. Pizurica, and H. Sahli, "Applications of Mathematical Morphology and Markov Random Field Theory to the Automatic Extraction of Linear Features in Airborne Images," *Math. Morphology and Its Applications to Image and Signal Processing,* pp. 405-414, 2000.

[9] A. Tuzikov, P. Soille, D. Jeulin, H. Bruneel, and M. Vermeulen, "Extraction of Grid Patterns on Stamped Metal Sheets Using Mathematical Morphology," *Proc. 11th Int'l Conf. Pattern Recognition,* vol. 1, pp. 425-428, 1992.

[10] P. Soille and H. Talbot, "Image Structure Orientation Using Mathematical Morphology," *Proc. 14th Int'l Conf. Pattern Recognition,* vol. 1, pp. 1467-1469, 1998.

[11] P. Soille and H. Talbot, "Directional Morphological Filtering," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 23, no. 11, pp. 1313-1329, Nov. 2001.

[12] J.E. Bresenham, "Algorithm for Computer Control of a Digital Plotter," *IBM Systems J.,* vol. 4, no. 1, pp. 25-30, 1965.

[13] M. van Herk, "A Fast Algorithm for Local Minimum and Maximum Filters on Rectangular and Octagonal Kernels," *Pattern Recognition Letters,* vol. 13, pp. 517-521, 1992.

[14] P. Soille, E.J. Breen, and R. Jones, "Recursive Implementation of Erosions and Dilations Along Discrete Lines at Arbitrary Angles," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 18, no. 5, pp. 562-567, May 1996.

[15] R. Jones and P. Soille, "Periodic Lines: Definition, Cascades, and Application to Granulometries," *Pattern Recognition Letters,* vol. 17, no. 10, pp. 1057-1063, 1996.

[16] R.G. Keys, "Cubic Convolution Interpolation for Digital Image Processing," *IEEE Trans. Acoustics, Speech, and Signal Processing,* vol. 29, no. 6, pp. 1153-1160, 1981.

[17] L.J. van Vliet, "Grey-Scale Measurements in Multi-Dimensional Digitized Images," PhD dissertation, Delft Univ. of Technology, Delft, The Netherlands, 1993, http://www.qi.tnw.tudelft.nl/Publications/phd_theses.html.

[18] L. Dorst and R. van den Boomgaard, "Morphological Signal Processing and the Slope Transform," *Signal Processing,* vol. 38, no. 1, pp. 79-98, 1994.

[19] T.Q. Pham, M. Bezuijen, L.J. van Vliet, K. Schutte, and C.L. Luengo Hendriks, "Performance of Optimal Registration Estimators," *Visual Information Processing XIV (Proc. SPIE Defense and Security Symp. 2005),* vol. 5817, pp. 133-144, 2005.

[20] G. Matheron, *Random Sets and Integral Geometry.* New York: Wiley, 1975.

[21] J. Serra, *Image Analysis and Mathematical Morphology.* London: Academic Press, 1982.

[22] P. Soille, *Morphological Image Analysis,* second ed. Berlin: Springer-Verlag, 2003.