Available online at www.sciencedirect.com
ScienceDirect
Pattern Recognition Letters 28 (2007) 865–872
Pattern Recognition Letters
www.elsevier.com/locate/patrec

Improving the accuracy of isotropic granulometries

C.L. Luengo Hendriks [a,*], G.M.P. van Kempen [b], L.J. van Vliet [a]

[a] Quantitative Imaging Group, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands
[b] Unilever Research and Development Vlaardingen, Olivier van Noortlaan 120, 3133 AT Vlaardingen, The Netherlands

Received 16 August 2005; received in revised form 11 August 2006
Available online 25 January 2007

Communicated by M. Lindenbaum

Abstract

Morphological sieves are capable of classifying objects in images according to their size. They yield a granulometry, which describes the imaged structure. The discrete sieve has some disadvantages that its continuous-domain counterpart does not have: sampled disks (used as isotropic structuring elements) are rather anisotropic, especially at small scales, and their area, as a function of the size in the continuous domain, shows jumps at apparently arbitrary locations. These problems cause a severe bias and low precision of the derived size distribution. Therefore we propose a new digitization scheme for implementing continuous sieves. First we increase the sampling density of the structuring element and the image. This does not add new detail to the image, but yields a sampled structuring element that is a much better approximation to its continuous counterpart, and thereby substantially reduces the discretization error. The second innovation is to shift the structuring element with respect to the sampling grid; this makes the size increments smoother, and further reduces the discretization errors. These ideas are validated on synthetic images. We also show that the proposed improvements allow for a finer scale sampling.

Keywords: Sieve; Scale-space; Mathematical morphology; Image analysis; Size distribution

1. Introduction

Sieves and granulometries were first proposed by Matheron (1975). They have been used in both binary morphology and grey-value morphology to measure particle-size distributions (Tscheschel et al., 2000), as well as to characterize textures (Asano et al., 2003; Bangham et al., 1994) and shapes (Maragos, 1989). Because a sieve has an increasing scale parameter, it results in a scale-space. Many theoretical studies have been made, linking it with linear scale-space theory and other non-linear scale-spaces (Alvarez and Morel, 1994; Jackway and Deriche, 1996; Park and Lee, 1996; Bangham et al., 1996; Chen and Yan, 1989).

A comparison can be made between morphological sieves and sifting rocks in a gravel heap (Soille, 2003). The rocks are sifted through meshes of decreasing size, extracting rocks from the collection. Each mesh removes the set of rocks that fall through it, but did not fall through the ones before. Thus, each set contains rocks in a given size range. The morphological opening γ_B and closing ϕ_B operations perform a similar function to that of the mesh, removing from an image those features (i.e. local maxima or minima) that are smaller than the structuring element (SE). The weight of each of the sets of rocks provides a point in a size distribution, which gives information on the gravel heap. Likewise, the integral over the result of each opening or closing yields a point in the granulometry.

A sieve is defined in mathematical terms by a transformation Ψ_λ having a size parameter λ and satisfying

* Corresponding author. Present address: Life Science Division, Lawrence Berkeley National Laboratory, One Cyclotron Road, Mailstop 84R171, Berkeley, CA 94720, USA. Tel.: +1 510 4865359; fax: +1 510 4865730.
E-mail address: clluengo@lbl.gov (C.L. Luengo Hendriks).

0167-8655/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.patrec.2006.12.001

the three following axioms, enumerated by Matheron (1975)

- Extensivity or anti-extensivity: the rocks that remain in the sieve are a subset of the initial rocks: $\Psi_\lambda f \leqslant f$.
- Increasingness: adding rocks to the heap does not diminish the number of rocks that remain in the sieve: $f \leqslant g \Rightarrow \Psi_\lambda f \leqslant \Psi_\lambda g$.
- Absorption: sifting at two different sizes λ and v will give the same result regardless of the order of the sieves; the size of the largest sieve determines the result: $\Psi_\lambda \Psi_v = \Psi_v \Psi_\lambda = \Psi_\lambda$, $\forall \lambda \geqslant v$.

By definition, all closings and openings satisfy the first two properties (the closing is extensive, the opening anti-extensive). However, as shown by Matheron, not all closings or openings with SEs of increasing size satisfy the absorption property. For example, a Euclidean disk satisfies this property, but a sampled Euclidean disk does not because of the discretization errors (Vogt, 1988).

The accuracy of the size distribution obtained from a granulometry depends on the characteristics of the morphological operations used to compute this granulometry. Two possibilities can be found in the literature: using approximations to isotropic SEs of varying accuracy, such as rectangles, octagons and the more refined approximation as defined by Jones and Soille (1996); and alternative implementations of the opening or closing using curve evolution (Brockett and Maragos, 1994; Sapiro et al., 1993; van den Boomgaard and Smeulders, 1994). Using rectangles or octagons does not provide the rotation-invariance that would be expected from a size distribution measurement. Jones and Soille's method uses periodic line segments to construct a series of approximations to disks of increasing size that satisfy the absorption property. The SEs built this way are fairly isotropic, and the morphological operations performed with them have a low computational cost since they can be decomposed into simple operations. The drawback is that the increase in size of these SEs is erratic and the method yields a size distribution with low scale sampling density. Finally, methods using curve evolution are notoriously slow due to the small step size required to obtain a stable solution to the differential equation, and are therefore rarely used in practice. They provide perfect isotropy, but can only produce an approximation to a morphological operation with a flat SE on a grey-value image.

This paper proposes two improvements in the definition of the discrete disk that greatly improve on the accuracy of the granulometry with respect to other approximations to isotropic SEs. We start by introducing the sieve (Section 2) and deriving a size distribution from it. This is examined in the continuous domain. Then we go into the implementation: Section 3 discusses the discretization of the continuous-domain granulometry, Section 4 contains the two improvements to the sampled Euclidean SE, and Section 5 deals with the implementation details. The proposed improvements are evaluated in Section 6 by comparing the various approximations to the discrete disk.

2. Sieves, granulometries and size distributions

To illustrate the principles of sieving, we use the structural closing (i.e. closing with a SE) as the sieving operation. However, by substituting it for the opening, an equivalent sieving operation is obtained. The difference is that from the closing sieve one can derive a size distribution of dark objects, whereas from the opening sieve a size distribution of the light objects would be obtained.

We will use an isotropic SE (e.g. a disk in two dimensions), expressed as $D(x,r) = \{x \mid \|x\| \leqslant r\}$. In this section, it is assumed that the image $f(x)$ is continuous and does not have a boundary ($f : \mathbb{R}^n \to \mathbb{R}$). We construct a scale-space $F(x,r)$ by closing (ϕ) the image $f(x)$ at all scales $r \in (0,\infty)$

$$F(x,r) = \phi_{D(x,r)} f(x), \tag{1}$$

$F(x,r)$ is an image with one more dimension than $f(x)$. Each image $F(x,r_0)$ contains only dark objects larger than r_0. We define $F(x,0) = f(x)$. Note that the continuous sieve is a continuous operation that produces a continuous scale-space when applied to a continuous image.

The granulometry $P(r) = \int F(x,r)\,\mathrm{d}x$ is a function of the scale r. Its scale derivative $p(r) = \frac{\mathrm{d}}{\mathrm{d}r} P(r)$ is referred to as pattern spectrum (Maragos, 1989; Soille, 2003) The granulometry is an increasing function (or decreasing in the case of the opening-sieve), with apparent transitions at the scales where image features disappear (as in Fig. 1). The pattern spectrum shows peaks at these locations. Normalization of the pattern spectrum yields a probability density function $h(r)$ that is referred to as the size distribution,

$$h(r) = \frac{p(r)}{\int_{r=0}^{\infty} F(x,r)\,\mathrm{d}x} = \frac{p(r)}{P(\infty) - P(0)}. \tag{2}$$

2.1. Invariances of the (cumulative) size distribution

Based on the closing scale-space, it is relatively easy to construct a size distribution of the dark objects. As mentioned in the introduction, the weight of each of the sets of rocks provides a point of the size distribution of these rocks. Similarly, the integral of each of the images in the closing scale-space can be used to construct a cumulative distribution. This distribution is rotation and translation invariant, because the closing is (Soille, 2003). By normalizing the cumulative distribution such that it ranges from 0 to 1, we make it invariant to the image size, offset and contrast, as well as the area fraction of the image covered by the objects. The cumulative distribution is thus defined as

$$H(r) = \frac{\int F(x,r)\,\mathrm{d}x - \int F(x,0)\,\mathrm{d}x}{\int F(x,\infty)\,\mathrm{d}x - \int F(x,0)\,\mathrm{d}x}, \tag{3}$$

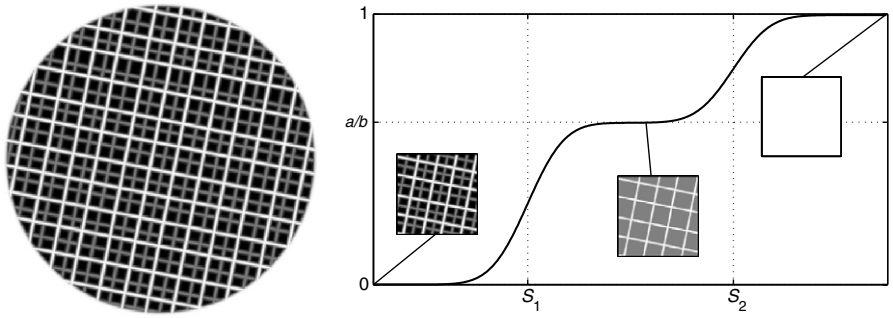

Fig. 1. Demonstration of the granulometry. Left: Synthetic test image with two structures at different scales superimposed. The image is formed by lines of grey-value a and b. Right: Abstracted granulometric curve that shows two jumps at the two scales present in the image.

where $F(x,\infty)$ is the original image closed with an infinite SE, and is thus equal to the image filled with its maximal grey-value. A size distribution is obtained by taking the scale derivative of $H(r)$,

$$h(r) = \frac{\mathrm{d}}{\mathrm{d}r} H(r). \tag{4}$$

3. Discrete granulometries

A discrete sieve is a discrete operation (with discrete SEs), applied to a digitized image, and resulting in a discrete scale-space. This requires

- Discretizing the input image: $f(x) \to \hat{f}[k]$, where $\hat{f} : \mathbb{Z}^n \to \mathbb{R}$, $k \in \mathbb{Z}^n$, $x = k\Delta$ and Δ the sample spacing.
- Discretizing the scales: $r \to i$. We sample the scale-axis at $r = s[i]$, where $i \in \mathbb{N}$ and $s : \mathbb{N} \to \mathbb{R}$ a scale-generating function. For logarithmic sampling, $s[i] = 2^{i/p}$, with p the number of scale samples per octave.
- Discretizing the SEs: $D(x, s[i]) \to \widehat{D}[k,i]$.

This results in a discrete scale-space computed by

$$\widehat{F}[k,i] = \phi_{\widehat{D}[k,i]} \hat{f}[k]. \tag{5}$$

Discretizing the input image poses no challenge. As long as the input image $f(x)$ is band-limited, it can be sampled without loss of information (Nyquist, 1928; Shannon, 1949). Such a sampled image exactly represents the continuous image. Sampling the SE, on the other hand, causes a large discretization error, the error that originates after reconstructing it from its samples. Due to this and the lack of sampling-invariance of the closing operation itself, the computed $\widehat{F}[k,i]$ is not equal to the sampled version of $F(x,r)$ at $x = k\Delta$, $r = s[i]$. The difference between $\widehat{F}[k,i]$ and $F(k\Delta, s[i])$ is caused by discretization errors. Also, because $F(x,r)$ is not band-limited, it is not possible to reconstruct it from $F(k\Delta, s[i])$. This is important especially along the scale-axis, where it might be interesting to find the exact location of peaks in $h(r)$.

We need to stress here that pattern spectra as in e.g. Maragos (1989) are discrete implementations of the sieve with a series of SEs that do not necessarily increase uniformly in size. These are defined to satisfy the absorption property, and the shape may vary to accomplish this. In this paper we define a discrete granulometry that approximates the continuous one, but only satisfies the absorption property by approximation. That is, the shape and size of the SEs must approximate those of the continuous granulometry as defined above. The more accurately this is accomplished, the closer it gets to satisfying the absorption property.

3.1. Sampling the flat structuring element

Sampling a binary function causes severe quantization errors. This has two problematic consequences:

- The original function cannot be reconstructed given the samples. The discretization errors occur in a strip along the boundary. The width of this strip is equal to the sample spacing Δ, the relative error is thus proportional to Δ.
- The size increments of a finely sampled scale-space become very erratic due to the irregular increase in size of the sampled disk as a function of the radius.

An obvious solution is not to use flat but grey-valued SEs, which can be sampled more accurately. A prominent example of a grey-valued SE is the parabola (Dorst and van den Boomgaard, 1994). However, a parabolic closing does not produce the desired result, because a sieve that uses parabolic closings will split a single object over a whole range of scales. This makes the transformation of the granulometry into a size distribution difficult, and its results would be even more difficult to interpret. Section 4 discusses how to implement a binary disk minimizing the discretization error.

3.2. Sampling the scale-axis

How to sample the scale axis is the second problem in discretizing the granulometry. There is relatively little literature on this topic. In most articles, one-pixel increments are used as a default solution. We suggest to use logarithmic sampling, so as to keep the relative error constant. One

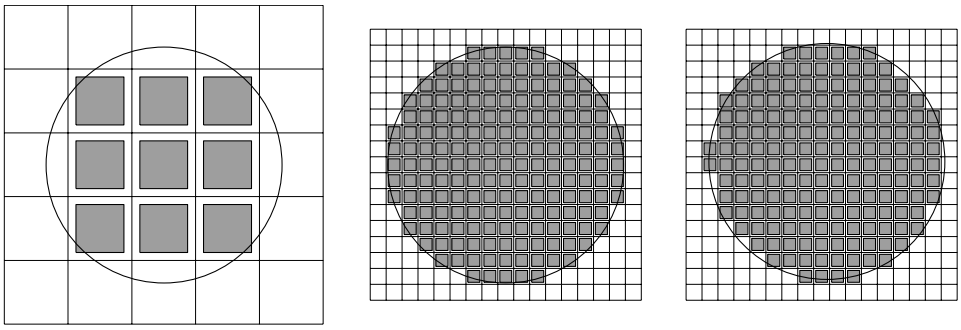

(a) Sampled disk (b) Denser grid (c) Shifted

Fig. 2. (a) Disk of radius 1.85 pixels, sampled (expected area = 10.7521, sampled area = 9). (b) Same disk, sampled on a grid four times as dense (area = 11.0625). This is the SE one would use after interpolating the image four times. (c) Same disk, centered at (0.19, 0.31) (area = 10.7500).

might want to distinguish between 3-pixel objects and 4-pixel ones, but not between 100-pixel objects and 101-pixels ones. Additionally, large objects are sparse in an image, causing a linearly sampled size distribution to be inaccurate at large scales. Particle or pore-size distributions often are log-normal (Aitchison and Brown, 1957), meaning that, when logarithmically sampled, they resemble a normal distribution.

A sampled $F[k,i]$ also produces a sampled $H[i] = H(r)|_{r=s[i]}$. In this case, the derivative in (4) can only be obtained by approximation,

$$h[i] = \frac{H[i+1] - H[i]}{s[i+1] - s[i]}. \tag{6}$$

An upper bound to the number of scale samples that are useful is given by the discretization of the SE. If too many samples are taken, then the differences between these SEs are very small, and the discretization error has a large influence on the results. That is, the averaging effect introduced by the limited number of samples along the scale-axis hides some of the errors made in the spatial domain. This notion links the scale sampling to the sampling of the SE, and is illustrated in Section 6.

Even for band-limited images, the scale-space $F(x,r)$ is not band-limited along the scale-axis. Therefore, $H(r)$ is also not band-limited. This makes it impossible to obtain all information on $H(r)$ using a predefined set of scales $s[i]$. Using adaptive sampling, however, it might be possible to find the location of large jumps in $H(r)$ (i.e. the peaks of $h(r)$). Note that $H(r)$ can be sampled at random locations; it is not necessary to do this in a fixed order or with fixed steps. Such a procedure is similar to finding zero-crossings of a function numerically.

4. Sampling the binary structuring element

Section 3 discussed problems encountered when discretizing the granulometry. One of the issues is how to discretize the SE, which should be flat (and thus binary) if one wishes to obtain a size distribution. This section proposes two innovations to the discrete disk that greatly diminish the discretization error: interpolation of the (grey-value) image data while increasing the sampling density of the disk by the same factor, and shifting the SE with respect to the sampling grid. This is not a heuristic method, but inspired by sampling theory. The relative discretization error of sampled randomly positioned disk (using a regular grid) is proportional to $R^{-1.5}$ (van Vliet et al., 2004). Here we skip the random positioning and use the optimal fixed position with respect to the sampling grid.

4.1. Increasing accuracy with interpolation

We stated in the previous section that the discretization error occurs in a strip along the boundary of the disk, whose width is equal to the pitch of the sampling grid. By decreasing this pitch the strip becomes thinner and the error diminishes (see Fig. 2b).

The scaling property of the closing

$$\{\phi_{D(\alpha^{-1}x,r)} f(\alpha^{-1}x)\}(\alpha x) = \phi_{D(x,r)} f(x) \tag{7}$$

teaches us that we can up-scale the input image and the SE, and down-scale the result. Replacing the left-hand side by a discrete closing, the equality is only true for $\alpha \to \infty$. Increasing α (i.e. up-sampling the image), the discrete closing becomes a better approximation to the continuous closing. Optionally, one could down-sample the input image for the larger SEs, so that the discretization error is approximately equal for all levels of the scale-space.[1]

4.2. Increasing accuracy with a shift

Due to the symmetry of a sampled disk on a square grid, all discretization effects are enhanced: when increasing the size of a discrete disk, a multiple of four or eight pixels are added to it. Placing the center of the disk away from the origin, this symmetry is broken (Fig. 2c), allowing the disk to be sampled more accurately. That is, we create a discrete disk by

[1] This also greatly diminishes the computational cost of the operation at large scales.

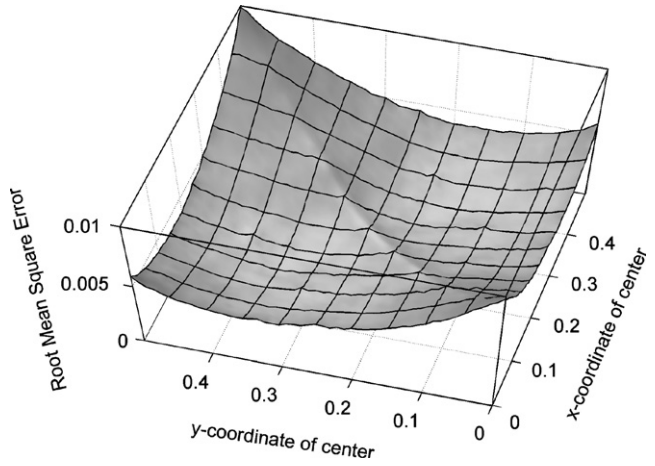

Fig. 3. Mean square relative error made when discretizing Euclidean disks (with random radii), for different positions of the disk's center with respect to the sampling grid. There is a minimum at coordinates (0.19, 0.31), but the exact position is not very important, because quite a large region around this minimum produces small errors.

$$\check{D}[n,i] = \{n \mid \|n\Delta + \delta\| \leqslant s[i]\}, \tag{8}$$

where δ is the shift. The optimal δ can be determined experimentally by computing the mean square relative error of the area for disks of many different sizes, centered at each location. This results in the graph of Fig. 3. There is a clear minimum at (0.19, 0.31), and, due to symmetry, another one at (0.31, 0.19). Fig. 4 shows the relative error made when centering the disk at this location in comparison to centering it at the origin. Remember that the continuous closing is invariant to translation of the SE. For higher-dimensional structures, the same experiment can be performed to determine the optimal location of the origin. For the three-dimensional sphere the optimal δ is (0.16, 0.24, 0.34).

The one-dimensional isotropic SE makes this concept easier to explain. If a segment is centered at position 0, two pixels are added (one at each end) at the same time when increasing its size (i.e. only lengths with an odd number of pixels are possible). Centering it at position 0.5 the same is true, but now it is always even in size. Optimal asymmetry is obtained by centering it at position 0.25. In this case, one pixel at a time is added to the SE. That is, more distinct lengths are possible.

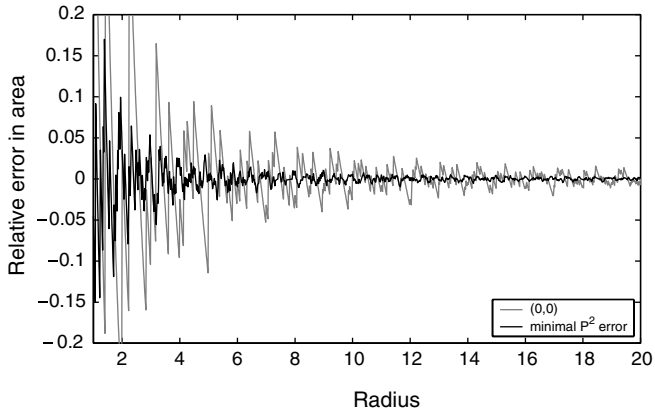

Fig. 4. Relative error in area for a disk centered at (0,0), and a disk centered at (0.19, 0.31), the location of the minimum in Fig. 3.

5. Implementation aspects

Implementing the granulometry as proposed here can be accomplished with the following steps:

1. Determine the required interpolation factors α_i given the required filter sizes s_i and a few parameters: maximum and minimum filter size, and maximum and minimum allowed interpolation factor.
2. Determine the offset u and scaling v used for normalization, as in (3): $u = \sum f$ and $v = n(\bigvee f) - u$, where n is the number of pixels in f.
3. For each filter size s_i:
   (a) Resample the image if needed: $f'[k] = f[k/\alpha_i]$.
   (b) Apply the closing (or opening) with a SE of size $\alpha_i s_i$: $g_i = \phi_{D(\alpha_i s_i)} f'$.
   (c) Sum the pixel values, divide by the interpolation factor and normalize: $H[i] = \left(\sum g_i/\alpha_i^N - u\right)/v$ (with N the image dimensionality).

We use values for α that are as small as possible, within the preset constraints, but are always a power of 2. Taking small values for α reduces the computation time: the closing with a SE that cannot be trivially decomposed (as ours are) is an operation with a computational complexity of $\mathbb{O}(n\sqrt{m})$, with n the number of pixels in the image and m the number of pixels in the SE (van Droogenbroeck and Talbot, 1996). Reducing the number of pixels by a factor α reduces the computation time by a factor $\alpha^{3/2}$. We used cubic convolution (Keys, 1981) to up-sample the images when $\alpha > 1$. This method is a good compromise between accuracy, computational cost and window size. In the experiments presented in the next section we have not down-sampled the image for larger SEs.

MATLAB sources for the algorithm are available for download at http://clluengo.lbl.gov/granulometry.html, together with the scripts we used to run the experiments described below.

6. Method evaluation

To illustrate the ideas on the sampling of the sieve presented above, the closing-sieve is applied to 13 identical but rotated versions of a synthetic, band-limited image (Fig. 1, left). It is a superposition of two line patterns, the finer one forming squares half the size, and with a grey-value 50% lower, than the coarser one. The position of the lines that compose the image have been distorted by white noise (standard deviation of 1 pixel; this distortion is identical in each instance).

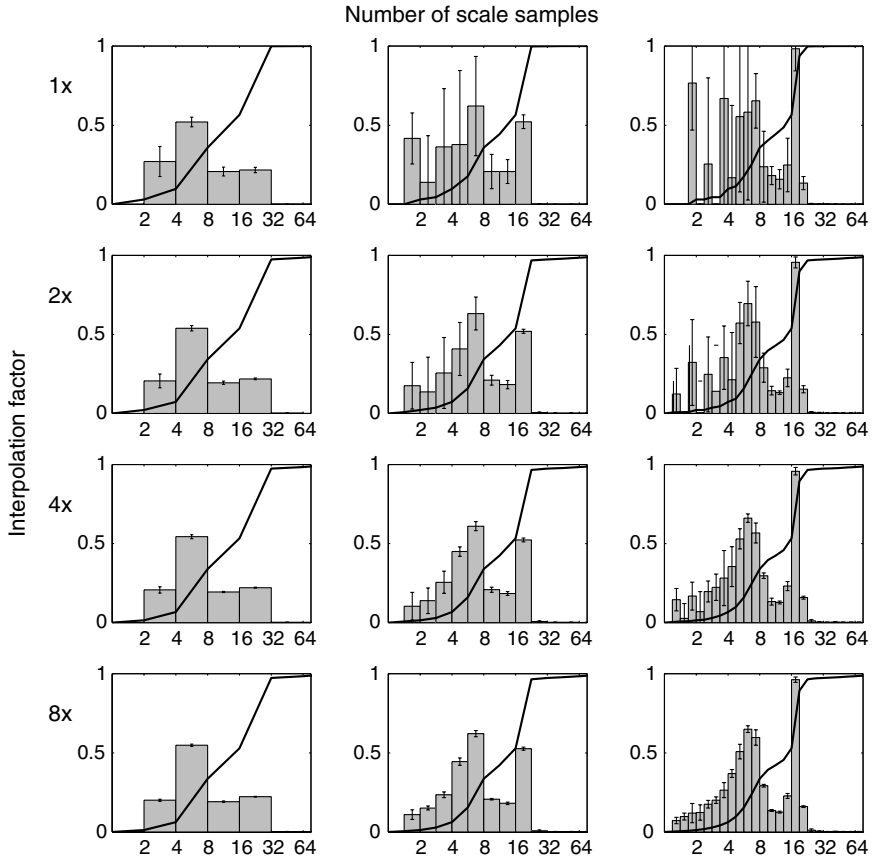

Fig. 5. Size distribution measured using the sieve with different interpolation factors and scale sampling densities (1, 2 and 4 samples per octave), and averaged over results on 13 rotated versions of the test image in Fig. 1. The SE used is a disk centered on the origin of the grid. The size of the detected dark objects is plotted on the x-axis, whereas the fraction of detected objects is plotted on the y-axis. The continuous line represents the average measured cumulative distribution. The error bars indicate the standard deviation.

Fig. 5 shows size distributions for this image, as calculated using the closing-sieve without shifting the SE. The various distributions result from choosing different interpolation factors (1, 2, 4 or 8 times) and scale-sampling densities (1, 2 or 4 samples per octave). Obviously, the best precision is obtained when the image is interpolated eight times.

The discretization error is accentuated by a scale-sampling rate that is too high. That is, the finer the scale sampling, the more influence the spatial discretization error has on the distribution. The reason is that the difference between SEs of subsequent scales is smaller for a finer scale sampling; this difference should always be larger than the discretization error. This gives a relation between the accuracy of the sampling of the SE, and the number of scale samples that can be computed.

We repeated this experiment using the shifted disk as SE. The optimal shift of $(0.19, 0.31)$ found through Fig. 3 was used. The results are shown in Fig. 6. It can clearly be seen that this greatly improves the precision of the method, especially at small scales. Using such a disk, an interpolation factor of 2 might be enough to obtain the precision at small scales otherwise only obtained with an interpolation factor of 8. The improvement is insignificant for the very large scales.

To demonstrate the increased sensitivity of this method we compare it to other methods used in the literature: using octagonal SEs, using sampled disk SEs, using Jones and Soille's approximation to the isotropic SEs (Jones and Soille, 1996), and using a curve evolution approach with the PDE as described by Maragos et al. (1994) (with a step size of 0.01, which yields an algorithm orders of magnitude slower than the proposed method). We applied the various granulometries to a collection of band-limited disks with diameters of 6 and 10 pixels (50 of each), and random sub-pixel offset. The resulting size distributions are shown in Fig. 7. Jones and Soille's method suffers from uneven and large increases in the size parameter[2]: both disk cohorts fall into the same bin. The octagonal SEs also do not provide enough resolution to distinguish the two diameters, since they require two-pixel increments. The curve evolution approach provides a very high level of rotational and translational invariance (data not shown), but the dilation itself is a smoothed approximation to the isotropic dilation. This smoothing effect also yields a smoothed gran-

[2] The diameter of the SE was determined by $2\sqrt{n/\pi}$, with n the surface area of the generated disk.

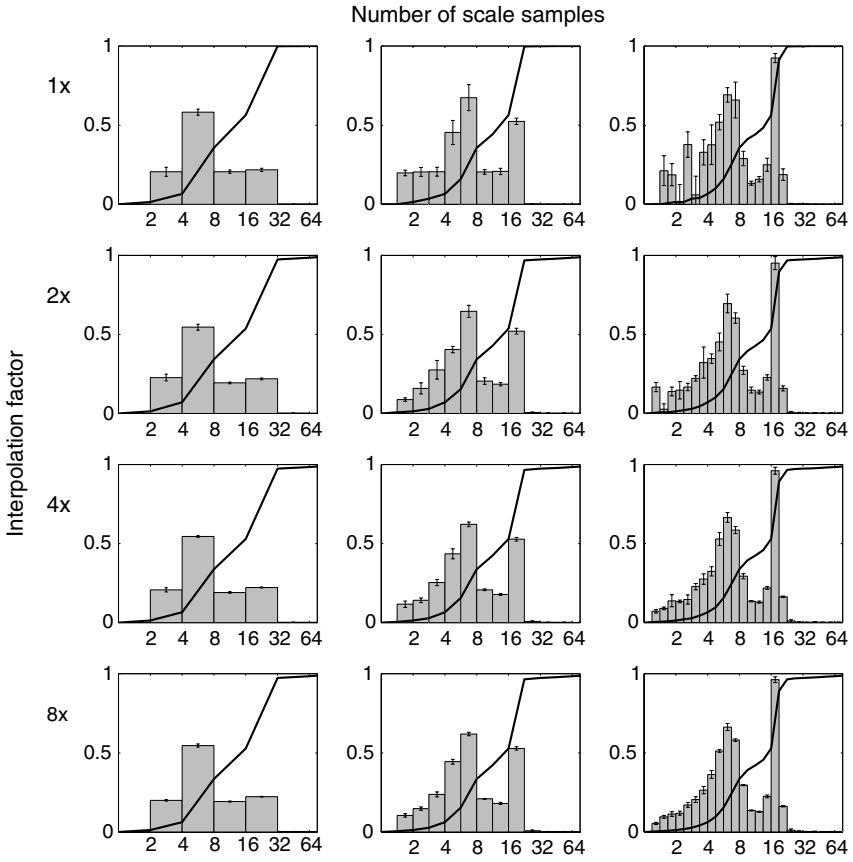

Fig. 6. Measured size distribution, as in Fig. 5, using a SE optimally placed with respect to the sampling grid.

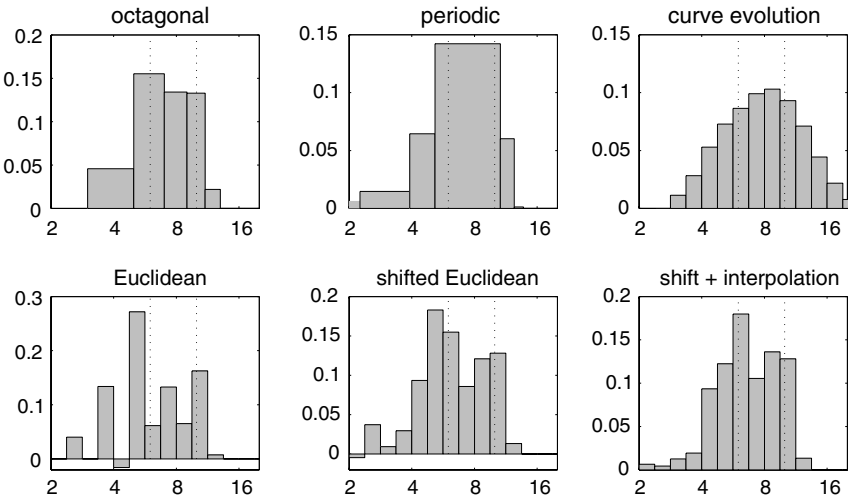

Fig. 7. Size distribution of a collection of 100 disks, half of which has a diameter of 6 pixels, the other half a diameter of 10 pixels, as measured with a granulometry using the various options for an approximation to an isotropic SE. The 'octagonal' method uses octagons. The 'periodic' method uses SEs constructed according to Jones and Soille (1996). The 'curve evolution' method uses the PDE described by Maragos et al. (1994) (step size 0.01). The 'Euclidean' method uses sampled disks. The 'shifted Euclidean' method uses sampled disks that were optimally shifted prior to sampling, as proposed in this paper. The 'shift + interpolation' method uses both the optimal shift and interpolation as proposed (interpolation factor up to 4, for a minimum diameter of 8 pixels).

ulometry, reducing its resolving power. The sampled Euclidean SEs yield a poor size distribution because they do not satisfy the absorption property required by the granulometry. This causes a very large error that reduces the resolving power of the method. Optimally shifting the disk before sampling greatly increases the accuracy, and yields a series of SEs that better (but not fully) satisfies the absorption property. This improvement by itself

enables the granulometry to distinguish two peaks in the size distribution. Adding the interpolation further increases the accuracy, enabling a correct measurement at even smaller scales. With interpolation, both peaks appear at the expected locations.

7. Conclusions

To overcome the problems induced by the discretization of binary disks used in discrete morphological sieves, we have proposed two improvements of the morphological operations:

- interpolate the input image for small scales, which allows for a denser sampling of the discrete disk; and
- place the origin of the disk away from the center (coordinates $(0.19, 0.31)$ were found to be optimal for two-dimensional disks), which suppresses the irregular size increments of the discrete disk.

These modifications are necessary to obtain an accurate and precise size distribution. Even though the resulting operation is strictly speaking not a sieve (because the absorption property is not exactly satisfied), it produces results that are closer to the results expected of a continuous-domain sieve on the continuous image.

We show that the granulometry with these modifications has a higher resolving power, especially at small scales, when compared to the implementation using discretized disk SEs, octagonal SEs, Jones and Soille's method using cascades of periodic line segments, or a curve evolution approach. Nonetheless, the method proposed here has a higher computational cost than any of the other methods (except for curve evolution): the SE is not separable, and the interpolation greatly increases the cost (doubling the sampling density in each dimension increases the computation time by a factor 8 for a two-dimensional image).

We have also shown that the discretization error imposes an upper bound to the number of samples that can be computed for the cumulative size distribution. Reducing the number of samples hides the discretization errors.

Acknowledgement

This research was partially supported by the Dutch Ministry of Economic Affairs under project number IOP-IBV98006.

References

Aitchison, J., Brown, J.A.C., 1957. The Lognormal Distribution. Cambridge University Press, London.

Alvarez, L., Morel, J.M., 1994. Morphological approach to multiscale analysis: From principles to equations. In: ter Haar Romeny, B.M. (Ed.), Geometry-Driven Diffusion in Computer Vision. Kluwer Academic Publishers, Dordrecht, pp. 229–254.

Asano, A., Ohkubo, T., Muneyasu, M., Hinamoto, T., 2003. Primitive and point configuration texture model and primitive estimation using mathematical morphology. In: Bigun, J., Gustavsson, T. (Eds.), SCIA 2003, 13th Scandinavian Conference on Image Analysis, Lecture Notes in Computer Science, vol. 2749. Springer-Verlag, Berlin, pp. 178–185.

Bangham, J.A., Campbell, T.G., Aldridge, R.V., 1994. Multiscale median and morphological filters for 2D pattern recognition. Signal Process. 38 (3), 387–415.

Bangham, J.A., Chardaire, P., Pye, C.J., Ling, P.D., 1996. Multiscale nonlinear decomposition: The sieve decomposition theorem. IEEE Trans. Pattern Anal. Machine Intell. 18 (5), 529–539.

Brockett, R.W., Maragos, P., 1994. Evolution equations for continuous-scale morphological filtering. IEEE Trans. Signal Process. 42 (12), 3377–3386.

Chen, M.H., Yan, P.F., 1989. A multiscaling approach based on morphological filtering. IEEE Trans. Pattern Anal. Machine Intell. 11 (7), 694–700.

Dorst, L., van den Boomgaard, R., 1994. Morphological signal processing and the slope transform. Signal Process. 38 (1), 79–98.

Jackway, P.T., Deriche, M., 1996. Scale-space properties of the multiscale morphological dilation-erosion. IEEE Trans. Pattern Anal. Machine Intell. 18 (1), 38–51.

Jones, R., Soille, P., 1996. Periodic lines: Definition, cascades, and application to granulometries. Pattern Recognition Lett. 17 (10), 1057–1063.

Keys, R.G., 1981. Cubic convolution interpolation for digital image processing. IEEE Trans. Acoust. Speech Signal Process. 29 (6), 1153–1160.

Maragos, P., 1989. Pattern spectrum and multiscale shape representation. IEEE Trans. Pattern Anal. Machine Intell. 11 (7), 701–716.

Maragos P., 1994. Differential morphology: Multiscale image dynamics, max–min difference equations, and slope transforms. In: Proc. ICIP-94, IEEE Internat. Conf. on Image Process., vol. 2, pp. 545–549.

Matheron, G., 1975. Random Sets and Integral Geometry. Wiley, New York.

Nyquist, H., 1928. Certain topics in telegraph transmission theory. Trans. AIEE, 617–644 (Reprinted in: Proceedings of the IEEE, vol. 90, no. 2, pp. 280–305, February 2002).

Park, K.R., Lee, C.N., 1996. Scale-space using mathematical morphology. IEEE Trans. Pattern Anal. Machine Intell. 18 (11), 1121–1126.

Sapiro, G., Kimmel, R., Shaked, D., Kimia, B., Bruckstein, A.M., 1993. Implementing continuous-scale morphology via curve evolution. Pattern Recognition 26 (9), 1363–1372.

Shannon, C.E., 1949. Communication in the presence of noise. Proc. IRE 37 (1), 10–21 (Reprinted in: Proceedings of the IEEE, vol. 86, no. 2, pp. 447–457, February 1998).

Soille, P., 2003. Morphological Image Analysis: Principles and Applications, second ed. Springer, Berlin.

Tscheschel, A., Stoyan, D., Hilfer, R., 2000. Erosion-dilation analysis for experimental and synthetic microstructures of sedimentary rock. Physica A 284, 46–58.

van den Boomgaard, R., Smeulders, A., 1994. The morphological structure of images: The differential equations of morphological scale-space. IEEE Trans. Pattern Anal. Machine Intell. 16 (11), 1101–1113.

van Droogenbroeck, M., Talbot, H., 1996. Fast computation of morphological operations with arbitrary structuring elements. Pattern Recognition Lett. 17 (14), 1451–1460.

van Vliet, L.J., Verbeek, P.W., Young, I.T., 2004. Quantitative imaging: How to measure size features in digitized images. In: Denney, T. (Ed.), ISBI'04: Proc. IEEE Int. Symp. Biomedical Imaging: From Nano to Macro (Arlington, VA, April 15–18). IEEE Press, Piscataway, pp. 1227–1230.

Vogt, R.C., 1988. Morphological operator distributions based on monotonicity and the problem posed by digital disk-shaped structuring elements. In: Juday, R.D. (Ed.), Digital and Optical Shape Representation and Pattern Recognition, SPIE Proceedings, vol. 938. SPIE, Bellingham, WA, pp. 1019–1022.