Segmentation-free estimation of length distributions using sieves and RIA morphology

Cris L. Luengo Hendriks and Lucas J. van Vliet

Pattern Recognition Group, Delft University of Technology,
Lorentzweg 1, 2628 CJ Delft, The Netherlands

{cris,lucas}@ph.tn.tudelft.nl

Keywords: rotated line segment, object length, orientation space, scale space.

Abstract

Length distributions can be estimated using a class of morphological sieves constructed with a so-called Rotation-Invariant, Anisotropic (RIA) morphology. The RIA morphology can only be computed from an (intermediate) morphological orientation space, which is produced by a morphological operation with rotated versions of an anisotropic structuring element. This structuring element is defined as an isotropic region in a subspace of the image space (i.e. it has fewer dimensions than the image). A closing or opening in this framework discriminates on various object lengths, such as the longest or shortest internal diameter. Applied in a sieve, they produce a length distribution. This distribution is obtained from grey-value images, avoiding the need for segmentation. We apply it to images of rice kernels. The distributions thus obtained are compared with measurements on binarized objects in the same images.

1. Introduction

The fraction of broken rice kernels in a batch is used to determine its quality. The milling process used to extract the kernels from their husk breaks a certain amount of them. Broken rice causes the consumer's perception of quality to decrease, and so does the price. This makes it economically important to determine the fraction of broken kernels.

Because manual counting is both expensive and subjective (different people apparently produce different results!), an automated system is required. A flatbed scanner is an ideal instrument to image rice, but it takes a lot of time to distribute the rice kernels on it in such a way that segmentation is possible. Therefore, we have applied a segmentation-free measurement technique to estimate the length distribution of kernels in an image, which can be used to derive the fraction of broken ones. It involves morphological filtering (RIA morphology) at different scales, from which a particle length distribution is obtained. The length of a kernel can be used to determine if it is broken or not. This multi-scale morphological filtering is called sieving.

A sieve is a technique that builds a scale-space using a single morphological operation with a scale parameter. This operation has to be chosen carefully. The morphological operations that are allowed to be used in a sieve must satisfy three properties: increasingness, extensivity and absorption [6]. In this scale-space, image features are separated into different levels according to some size measure. The chosen morphological operation determines to what level each feature is assigned. Section 2 gives a description of sieves and the derivation of size distributions.

Since our application requires the measurement of object length, we need a morphological operation that discriminates features based on their length. To this end, we have developed a morphological framework based on a structuring element that is isotropic in a subspace of the image, and thus anisotropic in the image space itself. Since the structuring element has full rotational freedom, this framework is rotation-invariant. It also possesses most of the properties of regular morphology. We name it Rotation-Invariant Anisotropic (RIA) morphology. An RIA opening removes an object in the image if it cannot encompass the structuring element under any orientation. This allows the RIA opening to discriminate objects on their characteristic lengths (supposing convex objects). In the case of an ellipsoid, these would be the principal axes. On an *N*-dimensional (hyper-)ellipsoid, a 1-dimensional structuring element finds the longest axis, a 2-dimensional one the second longest, etc. An *N*-dimensional structuring element is isotropic in the image space, and therefore has no rotational free-

in: R.L. Lagendijk, J.W.J. Heijnsdijk, A.D. Pimentel, M.H.F. Wilkinson (eds.), Proc. 7th Annual Conference of the Advanced School for Computing and Imaging (Heijen, the Netherlands, May 30-June 1), ASCI, Delft, 2001, 129-136.

dom; its usage reverts to regular isotropic morphology.

The theory and implementation details of RIA morphology are discussed in Section 3. Finally, Section 4 contains the results of applying these methods to the images of rice kernels.

2. The Sieve

Morphological sieves were first proposed by Matheron [6]. They have been extensively used with both binary and grey-value morphology to measure particle-size distributions. Since 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 (see for example Alvarez and Morel [1], or Park and Lee [8]).

A sieve can be built with any closing or opening operation Ψ that satisfies these three axioms [6]:

- Extensivity ($\Psi(f) \geq f$) or anti-extensivity ($\Psi(f) \leq f$),
- Increasingness (if $f \leq g$, then $\Psi(f) \leq \Psi(g)$), and
- Absorption (if $\lambda > \nu$ then $\Psi_\lambda(\Psi_\nu(f)) = \Psi_\nu(\Psi_\lambda(f)) = \Psi_\lambda(f)$).

By definition, all openings and closings satisfy the first two axioms, but many do not satisfy the third one [6]. In the next section we will introduce a closing that we use in the application in Section 4, and which does satisfy all three axioms.

In this section we illustrate the notion of sieving with a generic isotropic closing.

2.1. The Closing Scale-Space

We construct a (continuous) scale-space (see Figure 1) by closing (ϕ) the image at all scales $r \in (0,\infty)$,

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

Each image $F(x,r_0)$ contains only dark features larger than r_0. This is the closing scale-space. We define $F(x,0) = f(x)$. Sampling at discrete scales $r = s[i]$ denoted by $i \in \mathbb{N}$ produces a sampled scale-space,

$$F[x,i] = \phi_{D(s[i])} f(x) \quad . \tag{2}$$

For uniform scale sampling, $s[i] = i + 1$. However, if the relative error should be kept constant, logarithmic sampling suffices. In this case, $s[i] = b^i$ with $b = 2^{1/n}$, in which n denotes the number of samples per octave. The difference between subsequent scales contains features with sizes between $s[i]$ and $s[i+1]$, and is called a *granule image*:

$$g[x,i] = F[x,i+1] - F[x,i] \quad . \tag{3}$$

$g[x,i]$ forms the granule scale-space. Note that, as stated in Section 1, the measure for the size of a feature is determined by the chosen morphological operation.

Some structures contain different scales. Think about a telephone cable, composed of many bundles, each of which is made out of hundreds of thin wires. The wires are part of two structures at different scales. The morphological scale-space as described in this section is capable of finding both scales. Take as an example the middle pore in Figure 1, which is formed by walls of grey-value *a*, inside a region bounded by a higher grey-value *b*. A closing at some scale will cover the pore with grey-value *a*. A closing at a larger scale will cover the whole containing region with grey-value *b*. This results in a single pore being represented at two levels of the scale-space: at the one level with grey-value *a*, and at the other with grey-value *b*-*a*.

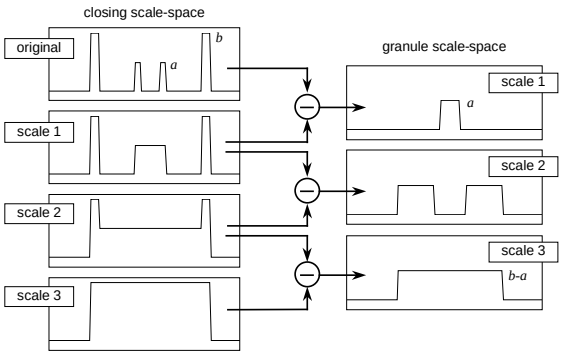

Figure 1. Sieve with flat closings.

2.2. Size Distributions

Based on either the granule scale-space or the closing scale-space, it is relatively easy to construct a size distribution. The grey-value sum of each of the granule images provides information on the amount of features of each size. Similarly, the sum of each of the images in the closing scale-space can be used to construct a cumulative distribution. Both distributions are rotation and translation invariant, since the closing is too [9]. The difference is that the cumulative distribution is independent of the chosen scales, whereas the granule images are not. By normalizing the cumulative distribution, it is made independent from the image size, contrast, and the fraction of objects. The cumulative distribution is thus defined as

$$H[i] = \frac{\sum_x F[x,i] - \sum_x F[x,0]}{\sum_x F[x,\infty] - \sum_x F[x,0]} \quad , \tag{4}$$

where $F[x,\infty]$ is the original image closed with an infinitely large structuring element, and is thus equal to an image filled with its maximum grey-value. A size distribution is the derivative of the cumulative distribution, which can be approximated by

$$h[i] = \frac{H[i+1] - H[i]}{s[i+1] - s[i]} \quad . \tag{5}$$

This is actually the same as what results after normalizing the distribution obtained using the sum of the granule images, i.e.

$$h[i] = \frac{\sum_x g[x,i]}{\left(s[i+1] - s[i]\right)\left(\sum_x F[x,\infty] - \sum_x F[x,0]\right)} \quad . \tag{6}$$

and is often called a *size histogram* [5].

2.3. Implementation Aspects

When looking at the description of a sieve, it is obvious that image features composed of grey-value ramps will be separated into many scales, just like the multi-scale pore of Figure 1. This can be dealt with by an appropriate pre-processing step (e.g. high pass-filtering, line or edge detection).

Another important question is how to sample the scale-space. There is relatively little literature on this topic, and in most articles, one-pixel increments are used as a default solution. However, we believe it makes sense to use logarithmic sampling, since we might want to distinguish between 3-pixel features and 4-pixel ones, but not between 100-pixel features and 101-pixels ones. As stated before, this causes the relative error to remain constant across the scales. We will be using four samples per octave for the current application, which means that $s[i] = 2^{i/4}$.

3. RIA Morphology

As stated in the introduction, this morphological framework is based on structuring elements that are isotropic in a subspace of the image, and thus anisotropic in the image space itself. By allowing these structuring elements to rotate, we can create rotation-invariant operators. The operators in this framework that are comparable to the dilation and erosion are actually not a dilation and erosion in the strict morphological sense, since they don't distribute with the intersection and union, respectively. Therefore we will name them *sedimentation* and *wear*, two words with a similar meaning, but without the morphological connotations. The other two operations defined in this framework are the closing and the opening.

In this section, we use δ as the symbol for dilation, ε for erosion, and, as in the previous section, ϕ for closing. As subscripts to these, we provide its structuring element. Translation is also denoted with a subscript: $f_x(t) = f(t-x)$.

3.1. Sedimentation and Wear

By decomposing the dilation with an isotropic structuring element D with radius r into a union of dilations with rotated one-dimensional isotropic elements L_φ with radius r and orientation φ, we get

$$\delta_D f = f \oplus D = f \oplus \bigcup_\varphi L_\varphi = \bigvee_\varphi \delta_{L_\varphi} f \quad . \tag{7}$$

Note that here φ is taken as a multi-dimensional orientation, or *orientation vector*. If, instead of taking the maximum over the dilations, we take the minimum, we get a new operator, which we will call RIA sedimentation,

$$\delta_L^{\measuredangle} f = \bigwedge_\varphi \delta_{L_\varphi} f \quad . \tag{8}$$

Here L can be any isotropic support with less dimensions than the image itself, and thus does not need to be a line. This operator takes the maximum of the image over the structuring element, rotated in such a way as to minimize this maximum. Figure 2 gives an example of the effect that this operator has on an object boundary. Note that a convex object boundary is not changed, but a concave one is.

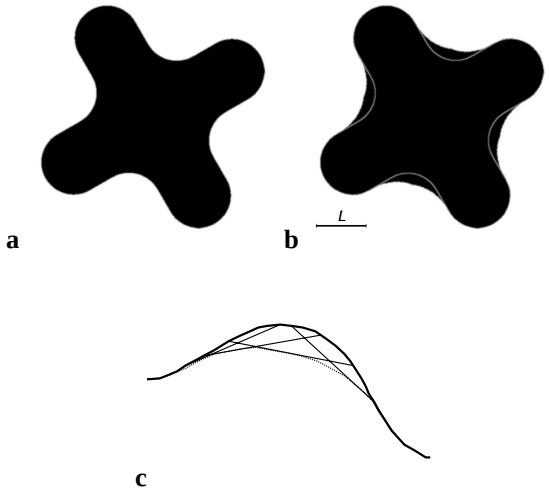

Figure 2. Effect of the RIA dilation on an object boundary. The white line in b represents the original object boundary. In c, its construction.

In the 2D case, in which L is a line, we can compare this sedimentation operator with a train running along a track. The train wagons (which are joined at both ends to the track) require some extra space at the inside of the curves. This sedimentation, applied to a train track, and using a structuring element with the length of the wagons, reproduces the area required by them. Note that this analogy is only true if the length of the structuring element is small com-

pared to the curvature of the boundary. This is always true for a train track, but not necessarily so for a grey-value image.

By duality, one can define the RIA wear as the maximum of a set of erosions with rotated line segments.

3.2. Opening and Closing

The closing is usually defined as a dilation followed by an erosion. However, it is easier to understand (and modify) if we see it as the maximum of the image over the support of the structuring element D, after shifting it in such a way that it minimizes this maximum, but still hits the point t at which the operation is being evaluated (see Figure 3a). Or, in other words, the 'lowest' position we can give D by shifting it over the 'landscape' defined by the function f,

$$\phi_D f = \bigwedge_{x \in D} \bigvee_{y \in D_x} f_y$$

$$\left[= \bigwedge_{x \in D} \left(\bigvee_{y \in D} f_y \right)_x = \varepsilon_D \delta_D f \right] \quad . \tag{9}$$

In accordance to this, we define a new morphological operation, RIA closing, as the 'lowest' position we can give the linear structuring element L, by shifting and rotating it over the 'landscape' f, such that it still hits the point x being evaluated (see Figure 3b). It is defined by

$$\phi_L^{\curlywedge} f = \bigwedge_{\varphi} \bigwedge_{x \in L_\varphi} \bigvee_{y \in L_{\varphi,x}} f_y \quad . \tag{10}$$

This turns out to be the same as the minimum of the closings, at all orientations, with the structuring element L (but not equal to an RIA sedimentation followed by an RIA wear),

$$\phi_L^{\curlywedge} f = \bigwedge_{\varphi} \bigwedge_{x \in L_\varphi} \left(\bigvee_{y \in L_\varphi} f_y \right)_x =$$

$$= \bigwedge_{\varphi} \varepsilon_{L_\varphi} \delta_{L_\varphi} f = \bigwedge_{\varphi} \phi_{L_\varphi} f \quad . \tag{11}$$

According to Matheron, this operation is an algebraic closing since it is an intersection of morphological closings [6]. This implies that extensivity, increasingness and absorption are satisfied, and they can be used in a sieve. The two-dimensional case is an intersection of closings with rotated lines, which have been used before (see for example Soille [9]), and we will use in the next section.

By duality, one can define the RIA opening as the maximum of openings with rotated line segments.

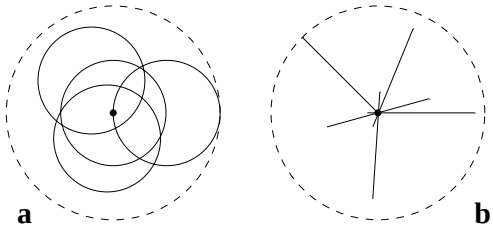

Figure 3. Construction of the 2D isotropic closing (a), and that of the 2D RIA closing (b).

3.3. Morphological Orientation Space and RIA Morphology

The morphological operation ψ by rotated versions of an anisotropic structuring element L can be used to construct a morphological orientation space

$$f_{\Psi,L}(x, y, \varphi) = \Psi_{L_\varphi} f(x, y) \quad . \tag{12}$$

The RIA sedimentation and closing now result from a maximum projection along the orientation axis,

$$\Psi_L^{\curlywedge} f(x, y) = \bigwedge_{\varphi} f_{\Psi,L}(x, y, \varphi) \quad . \tag{13}$$

The RIA wear and opening result from a minimum projection.

In a sieve, this orientation space would be extended with a scale dimension.

3.4. Implementation

Morphological operations that use line segments as structuring elements are often implemented using approximations of line segments, such as periodic lines [4]. These might be a very good solution for binary images, but for grey-value images we can do much better.

Since grey-value images are an exact representation of a continuous function (if sampled correctly), they can be resampled at arbitrary locations. This makes it possible to rotate the image instead of the line segment. Using a horizontal or vertical line produces the smallest discretization error for the structuring element. For each two orientations, the image must be rotated, the one-dimensional filter applied in both horizontal and vertical directions, the two results combined, and the result rotated back to be combined with the results for other orientations. Of course, this is only possible because the results are combined using a maximum or minimum operation.

The region of influence for the dilation and erosion is a circle with radius r (being the diameter of the structuring element in the subspace). For the opening and closing it has a diameter of $2r$. To take all the pixels in these neighborhoods into account, respectively $r\pi/2$ and $r\pi$ orientations are needed.

This means that $r\pi$ rotations of the image are required for an opening operation. To save on some computation, we skew the image instead of rotating it. However, each skew only allows for the computation of a single orientation, which means that for each orientation two skews are required instead of three (a rotation can be accomplished with three skews [3, 7, 10]).

The quantification error is produced by the approximated line segment length in the skewed image, which is not the same for all orientations.

4. Length Measurement of the Rice Kernels

Figure 4 shows two images of rice kernels obtained by placing the rice on a flatbed scanner. The image on the left has all kernels manually separated before acquisition, which takes about 15 minutes. The one on the right contains the same kernels randomly scattered on the scanning surface. As stated before, it is not trivial to correctly segment such an image. Thus, the classical measuring paradigm (threshold, label, measure the segmented objects) is not easily applied.

In total we have 10 images of the same sample, 20% of which consists of broken kernels:

- two images with only the broken kernels (one touching, one separated),
- two images with only the intact kernels (again, one touching, one separated), and
- six images with all kernels (two touching, four separated).

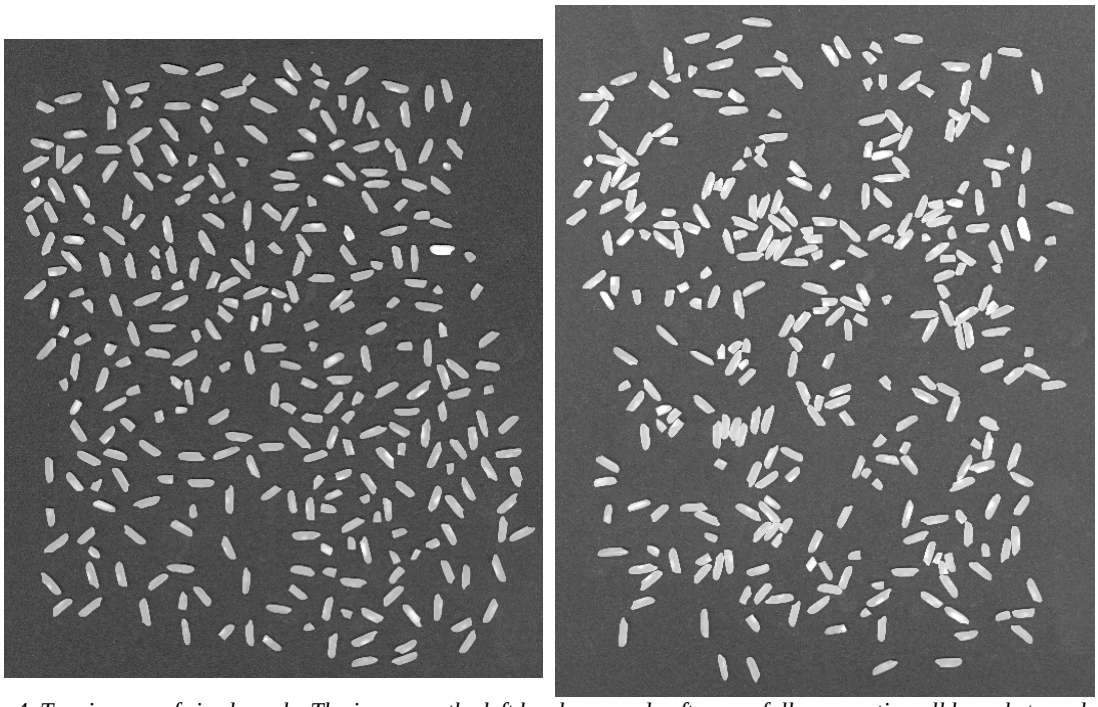

Figure 4. Two images of rice kernels. The image on the left has been made after carefully separating all kernels to make segmentation easy.

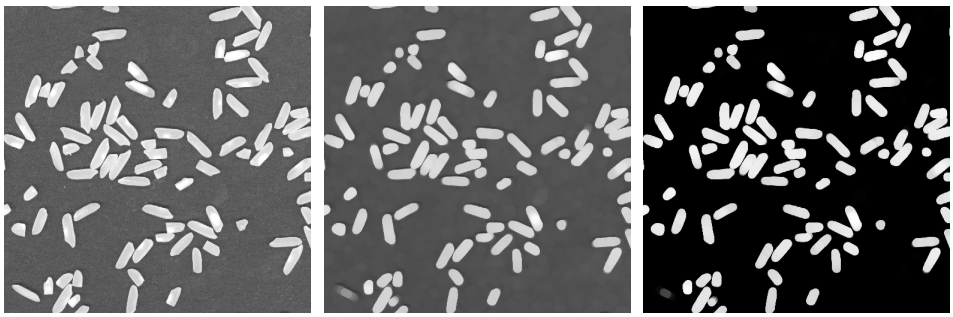

Figure 5. Preprocessing of the images: first, an opening removes thin elements, which are not counted in de length distributions (middle). Then, an error-function clip is applied (right).

Accurate and reliable measurements can be obtained by applying the sieve with an RIA opening to an image with all kernels separated (as the image on the left of Figure 4). However, that is the easy problem. An image with touching kernels will produce an over-estimation of the particle sizes, since groups of kernels can accommodate larger line segments than single ones. The solution would be to use 'thick lines' (ellipses). If these are thick enough not to fit through the union point between touching rice kernels (which is usually thinner than the kernels themselves), the measurements produce the same results as on the first image.

4.1. Preprocessing

To increase the accuracy of the measurements we do some preprocessing on the images (see Figure 5 for the results of these steps). The goals are two-fold:

1. remove imaging artifacts, and
2. remove kernels that are thinner than the structuring element to avoid an underestimation of the lengths.

Because we use thick line segments as structuring elements, all kernels and portions of kernels that are thinner than these will be put into the smallest scale of the granulometry. To overcome this we remove these features by an opening with a disk of diameter equal to the width of the line segments. Since very few rice kernels are too thin, removing them introduces only a very small imprecision in the measurements. The thinner portions of the kernels that are also removed cause these to be somewhat shorter. This yields a systematic error, an average underestimation of the lengths of four pixels (result obtained experimentally). This causes a shift to the left of the cumulative length distribution. This error would, however, also be produced by the introduction of thick line segments (in addition to an overestimation of the smallest scale).

The second operation that is applied to the images is an error-function clip. This is a clipping that introduces less aliasing than hard clipping [11]. Its need is two-fold: removing noise in the background, and equalizing the grey-value over the kernels. Some of these contain a chalky portion, caused by an unbalanced growing process. This chalky portion is imaged whiter than the rest of the kernel, and would influence the length distribution by adding weight to the smaller scales. Figure 6 shows the result of this error-function clip on a one-dimensional portion of the image through a rice kernel.

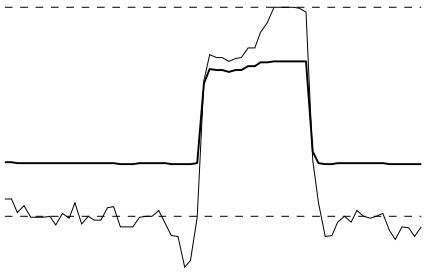

Figure 6. Error-function clip on a 1D portion of the image through a rice kernel. The bottom dotted line is the median over the whole image; the top one is the maximum value. The chosen range of the clipping is half the distance between these values, so that the influences of the background noise and the dense portion of the rice kernel are reduced.

4.2. The Classical Measuring Paradigm

To compare our results with those obtained with an existing algorithm, we measured the length distribution using the Feret length measure [2] on the thresholded and segmented image. This works well on the images where the kernels have been manually separated before acquisition, but produces poor results on the images with touching kernels. The algorithm we used to determine the Feret length uses a chain-code representation of the object boundary, which can be easily rotated. The longest projection of the boundary is used as the object length.

4.3. Results

The length distributions of the two images in Figure 4, obtained by the proposed sieve as well as the Feret length, are plotted in Figure 7. The results for both images using the sieve are almost identical and only slightly different from the measurement obtained using the classical method applied to the image with separated kernels. However, the classical method applied to the image with touching kernels produces a very large over-estimation of the sizes.

Figure 8 shows the results obtained by the sieve on the ten images. In all cases, the sieve applied to the images with touching kernels produces only a minimal overestimation of the kernel length.

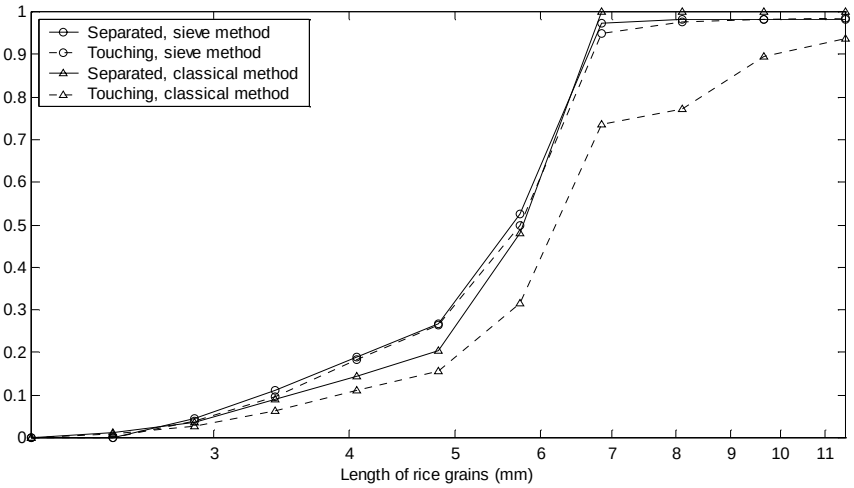

Figure 7. Comparison of the classical segment and measure method, and the sieve with the RIA opening. For the latter, touching rice kernels do not influence the measurement very much. Note the logarithmic scaling of the horizontal axis.

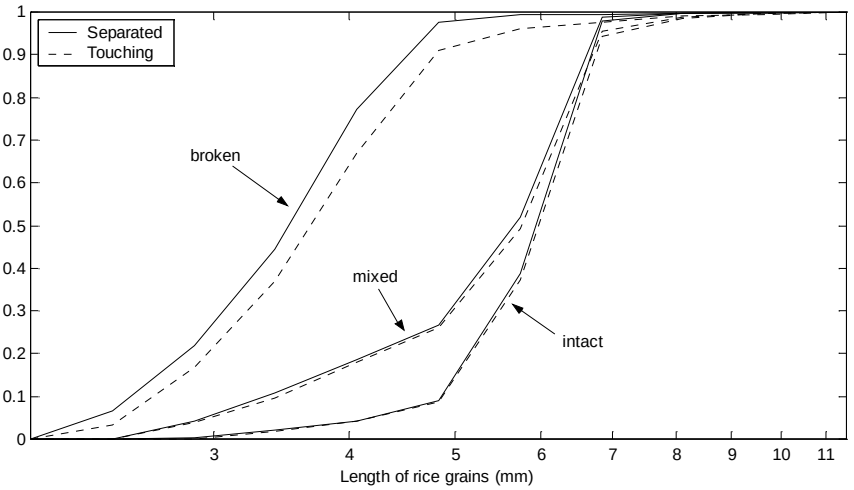

Figure 8. Cumulative distribution measured for the images. This figure shows that it is easy to measure the fraction of broken kernels in this way. The difference induced by the contact between rice kernels is very small.

5. Conclusion

For the application discussed in this article, as well as many other applications, segmentation is very difficult or not possible at all. Segmentation-free measurement techniques are therefore desirable. Sieves, a form of multi-scale morphological filtering, are very useful in this context. Sieves produce size distributions from grey-value images, and the measure for the size of the image features is determined by the chosen morphological operation. Since this application requires length measurement, RIA openings have been used in a sieve to obtain length distributions.

RIA (Rotation-Invariant Anisotropic) openings are the openings in a new morphological framework that results from decomposing an isotropic structuring element into rotated lower-dimensional isotropic structuring elements. An RIA opening only removes an image feature if the chosen structuring element does not fit under any orientation.

The proposed sieve is applied to measure the length distribution of rice kernels acquired with a flatbed scanner. These were scattered quickly onto the scanning surface, so that many are touching. To minimize the influence of the touching kernels, we have modified the RIA openings slightly, using line segments of certain width, instead of using one-pixel thin line segments. With this modification, the ob-

tained distributions are almost identical for the images with separated and touching kernels. In contrast, the classical measuring paradigm (which uses a threshold, segmentation of the objects, and measuring the length based on these binarized shapes) produces incorrect results for the image with the touching kernels.

6. References

[1] L. Alvarez and J.-M. Morel, “Morphological Approach to Multiscale Analysis: From Principles to Equations,” in Geometry-Driven Diffusion in computer Vision, M. A. Viergever, Ed. Dordrecht: Kluwer Academic Publishers, 1994.

[2] L. R. Feret, “La grosseur des grains,” Assoc. Intern. Essais Math. 2D, Zurich, 1931.

[3] B. Jähne, Practical handbook on image processing for scientific applications. Boca Raton (FL): CRC Press, 1997.

[4] R. Jones and P. Soille, “Periodic Lines: Definition, cascades, and application to granulometries,” Pattern Recognition Letters, vol. 17, pp. 1057-1063, 1996.

[5] P. Maragos, “Pattern Spectrum and Multiscale Shape Representation,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 11, pp. 701-716, 1989.

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

[7] A. W. Paeth, “A Fast Algorithm for General Raster Rotation,” presented at Graphics Interface '86, 1986.

[8] K.-R. Park and C.-N. Lee, “Scale-Space Using Mathematical Morphology,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 18, pp. 1121-1126, 1996.

[9] P. Soille, Morphological Image Analysis. Berlin: Springer-Verlag, 1999.

[10] A. Tanaka, M. Kameyama, S. Kazama, and O. Watanabe, “A Rotation Method for Raster Image Using Skew Transformation,” presented at CVPR'86, IEEE Conference on Computer Vision and Pattern Recognition, 1986.

[11] L. J. van Vliet, "Grey-Scale Measurements in Multi-Dimensional Digitized Images." Ph.D. Thesis, Pattern Recognition Group, Faculty of Applied Physics, Delft University of Technology, Delft, 1993.