

Automated Sizing of Coarse-Grained Sediments: Image-Processing Procedures¹

David J. Graham,² Ian Reid,² and Stephen P. Rice²

This is the first in a pair of papers in which we present image-processing-based procedures for the measurement of fluvial gravels. The spatial and temporal resolution of surface grain-size characterization is constrained by the time-consuming and costly nature of traditional measurement techniques. Several groups have developed image-processing-based procedures, but none have demonstrated the transferability of these techniques between sites with different lithological, clast form and textural characteristics. Here we focus on image-processing procedures for identifying and measuring image objects (i.e. grains); the second paper examines the application of such procedures to the measurement of fluvially deposited gravels. Four image-segmentation procedures are developed, each having several internal parameters, giving a total of 416 permutations. These are executed on 39 images from three field sites at which the clasts have contrasting physical properties. The performance of each procedure is evaluated against a sample of manually digitized grains in the same images, by comparing three derived statistics. The results demonstrate that it is relatively straightforward to develop procedures that satisfactorily identify objects in any single image or a set of images with similar sedimentary characteristics. However, the optimal procedure is that which gives consistently good results across sites with dissimilar sedimentary characteristics. We show that neighborhood-based operations are the most powerful, and a morphological bottom-hat transform with a double threshold is optimal. It is demonstrated that its performance approaches that of the procedures giving the best results for individual sites. Overall, it out-performs previously published, or improvements to previously published, methods.

KEY WORDS: grain-size analysis, mathematical morphology, segmentation, bottom-hat transform, rivers, fluvial.

INTRODUCTION

Recent years have seen several groups attempt to derive automatically from digital images the surface grain-size distribution of fluvial sediment (McEwan and others, 2000; Butler, Lane, and Chandler, 2001; Reid, Rice, and Garcia, 2001; Sime and Ferguson, 2003). These workers have demonstrated that encouraging results can be

¹Received 17 November 2003; accepted 12 July 2004.

²Department of Geography, Loughborough University, Loughborough, Leicestershire LE11 3TU, United Kingdom; e-mail: d.j.graham@lboro.ac.uk; ian.reid@lboro.ac.uk; s.rice@lboro.ac.uk

achieved for small numbers of images collected under controlled conditions. The challenge now is to design a transferable procedure that performs demonstrably well under a range of sedimentary and sampling conditions (lithotype, fluvial texture, grain-arrangements, lighting) so that image-processing-based methods of grain-size measurement may be widely adopted. This is the first in a pair of papers in which we outline such a procedure. This paper is concerned with the identification and measurement of regions that correspond to grains in digital images (the shaded boxes in Fig. 1). The second paper examines how such procedures may be applied to derive reliable surface grain-size distributions for waterlain gravels.

The spatial variability of grain size at a variety of scales makes the characterization of fluvial sediment notoriously difficult (Church, McLean, and Wolcott, 1987; Wolcott and Church, 1991). Large sample sizes are necessary to ensure adequate representation of the clast population and sampling is therefore time-consuming, laborious and costly. Conventional sampling techniques leave potentially important textural variations unresolved in hydraulic, geomorphological and ecological studies of river channel behavior. So, for example, in grid-by-number sampling (Wolman, 1954), the establishment of a grid with adequate cell size may mean incorporating a number of micro- or meso-scale patches in which grain size differs significantly. Although representative of a broad area, such a sample may provide grain-size parameters that overestimate the entrainment threshold, thus displacing a chosen bedload function and leading to a serious underestimation of sediment flux.

The development of sampling techniques that achieve satisfactory characterization of grain size whilst simultaneously reducing the time spent in both the field and laboratory is highly desirable. Several researchers have used emulsion-based photographic data capture to reduce field time (Adams, 1979; Church, McLean, and Wolcott, 1987; Rice and Church, 1998), but laboratory analysis of the resulting photographs can prove extremely time-consuming. Even the most advanced of such procedures, photo-sieving (Ibekken and Schleyer, 1986; Diepenbroek, Martholoma, and Ibbeken, 1992; Diepenbroek and De Jong, 1994; Ibbeken, Warnke, and Diepenbroek, 1998), relies on manual identification and digitization of individual particle boundaries.

These limitations have led some workers to seek automated methods of extracting information from images. Such methods have been used extensively in biomedical applications, and these have driven much of the research on image segmentation and measurement (Bankman, 2000). Within the Earth sciences, image-processing techniques have been used successfully in studies of soil structure (Glasbey, Horgan, and Darbyshire, 1991; Horgan, 1998) and crystal shape (Bernard-Michel and others, 1997), following on from earlier manual methods of grain-size classification of fine-grained sediments from microscope thin-sections (Friedman, 1958; Kellerhals, Shaw, and Arora, 1975). In some respects, river beds

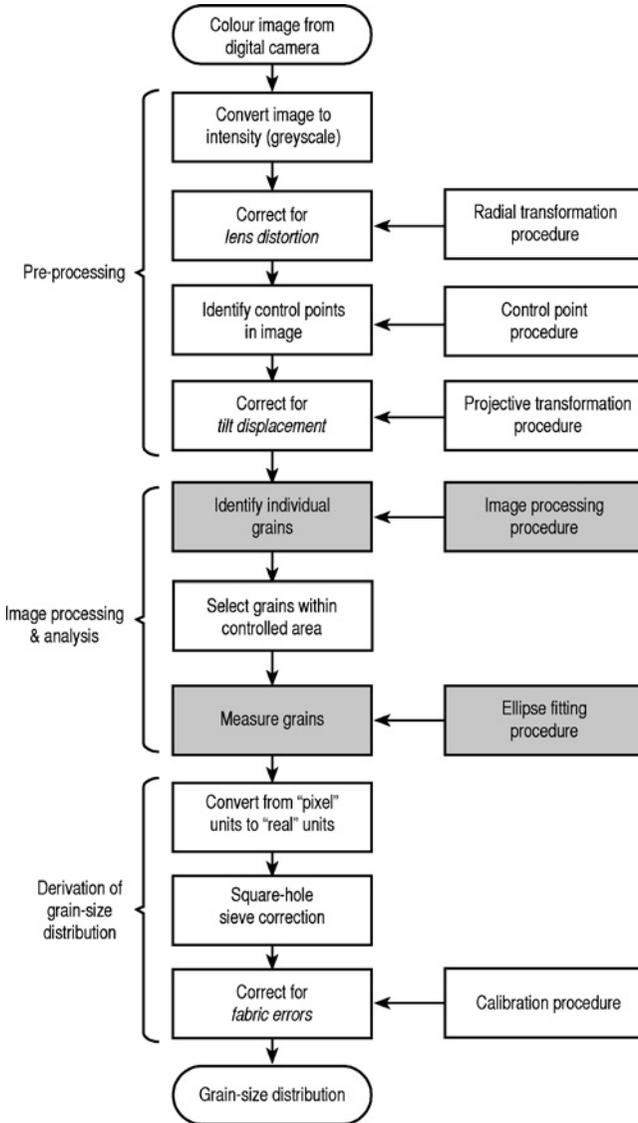


Figure 1. Illustration of the stages required to extract grain-size data from a digital image. The shaded gray boxes represent the image-processing and analysis procedures described here. A companion paper describes the remainder of the stages.

present a more complex problem in that grains may be partially hidden or inclined relative to the plane of the image, there may be significant heterogeneity in hue and grain-surface texture between and within individual grains, and the surface has elevation variations which may result in uneven lighting and shading across individual grains and across the image. Nevertheless, improvements in computer power and image-processing (i.e. enhancement and modification) and analysis (i.e. extraction of information) techniques have facilitated recent advances in automated grain identification and measurement within digital images (McEwan and others, 2000; Butler, Lane, and Chandler, 2001; Reid, Rice, and Garcia, 2001; Sime and Ferguson, 2003). However, although these results are encouraging, the test datasets used were small and the transferability of such approaches to a wide range of lithotypes, grain shapes, packing configurations and sizes is, as yet, unproven.

To be considered successful, automated techniques should: (i) be simple and rapid to employ without specialist equipment or technical knowledge; (ii) produce grain-size distributions with a precision comparable to, or better than, traditional manual approaches; and (iii) be applicable across a diverse range of sediment types. Keeping in mind the first of these objectives, we have focused exclusively on measurements from single images rather than stereo pairs, thus avoiding the use of specialist photogrammetric software and relatively complex field procedures. Butler, Lane, and Chandler (2001) found that the use of single images produces better results than photogrammetric or hybrid approaches.

The errors associated with image-processing-based grain-size measurement may be separated into three distinct categories. (1) *Image-processing errors* are those associated with the ability of the image-processing and analysis procedures to identify accurately and measure individual regions in an image that correspond to actual grains. The procedure to minimize these errors is the subject of this paper. (2) *Spatial distortions* arise because the three-dimensional sediment surface is projected onto a two-dimensional image plane through an imperfect lens, giving rise to tilt distortion, relief distortion and lens distortion. (3) *Fabric errors* are also inherent in photographic approaches because the complex three-dimensional structure of an exposed sediment surface means that particles overlap, partially hiding one another (occlusion), and that grains are inclined relative to the plane of the image. The restitution of spatial distortions and the consequences of fabric errors are discussed in the companion paper.

In developing and testing appropriate image-processing procedures, a three-stage approach was adopted. First, seven procedures were developed based on published grain-sizing techniques and the image-processing literature. The performance of these procedures was evaluated subjectively to select the four most promising approaches. Second, the four selected procedures were rigorously tested against a control dataset based on 39 images representing a range of lithologies and textures in order to define the optimal procedure and set of internal parameters. Third, the performance of this optimal procedure and parameter set was evaluated

to specify the precision of the results. In order to emphasize that this paper is concerned with the evaluation of the image-processing procedures rather than the generation of calibrated grain-size distributions, the results are presented in pixels rather than metric units.

FIELD PROCEDURE

Fieldwork was undertaken at three sites chosen to represent a variety of lithologies and grain shapes. The sediment of Ettrick Water, in the Scottish Borders, is characterized generally by pale-colored grits and shales; individual clasts tend to be equant in shape. The Afon Ystwyth, Ceredigion, Wales, is dominated by clasts of fine-grained grit, darker in color than those of Ettrick Water and commonly platy in shape. In contrast, the clasts of the River Lune, Lancashire, England, consist primarily of limestone, occasionally of sandstone, and are predominantly equant in shape.

A total of 39 sediment patches were selected across the three sites to represent a variety of grain-size distributions. In order to characterize adequately the tails of a grain-size distribution, at least 300 grains need to be measured (Fripp and Diplas, 1993; Rice and Church, 1996). Assuming a simple geometric relation between median particle size D_{50} (mm) and area covered, the required sample area A (mm^2) varies as $A = nD_{50}^2$, where $n \geq 300$. Provided the median grain size (D_{50}) is less than 64 mm, an area of 1.2 m^2 thus yields an adequate sample. Experience suggests that where the median grain size is about 64 mm and where there are typical sorting values, maximum particle size can exceed 200 mm, so attention has not been unduly restricted to fine gravel beds. Each patch was rectangular, the aspect ratio reflecting that of the images recorded by the camera (4:3).

A wooden frame (1.27 by 0.95 m) was placed over each patch to ensure selection of a consistent area and the corners were marked using standard proprietary adhesive targets. The frame was then removed and the patch was photographed using a mid-range digital camera (Olympus C-3030 ZOOM, nominally 3.34 Megapixels, 2048 by 1536 maximum image resolution). Images were stored in color JPEG format using the minimum compression supported by the camera, resulting in image sizes of 1.5–2 MB. This compression resulted in no visually apparent loss in image quality when compared with an uncompressed image (TIFF format, 9 MB). Images were converted to 8-bit grayscale prior to the commencement of the image-processing procedures. Photographs were taken vertically from a platform height of approximately 1.5 m, giving a pixel resolution of approximately 0.7 mm on the ground. This height minimized barrel distortion associated with the use of a wide-angle lens whilst still making the camera easily accessible from the ground. The platform consisted of a horizontal bar suspended between two tripods and the camera was mounted on the bar using a tripod head.

Images were collected in a variety of natural and artificial lighting conditions (natural overcast and sunlit conditions as appropriate, artificially shaded, direct and bounced flash).

DEVELOPMENT OF THE IMAGE-PROCESSING PROCEDURES

The objective of the image-processing procedure is to generate a binary image in which each discrete region represents a grain. This is usually a two-stage process. The first stage produces a first partitioning of the image, but is not generally able to generate unique regions for every grain. Seven such procedures were developed and tested. Of these, four were selected for more detailed analysis on the basis of their promising performance (Fig. 2). The second stage, a *watershed segmentation*, tidies the binary image and attempts to split regions representing more than one grain into singular objects. Once the final segmentation has been achieved, an image-analysis procedure must be applied to extract useful information from the derived image, in this case the size of each object.

Image-processing procedures may operate either on the entire image, based on some global characteristic of the image (global operations), or on individual pixels, based on the characteristics of its neighbors (neighborhood operations). An example of a global operation is the application of an intensity threshold to a grayscale image in order to produce a binary image. Global operations have the advantage that they are very rapid to execute, but they are intolerant of non-uniform lighting and variations in grain hue. An example of a neighborhood operation is the application of a median filter, which replaces each pixel in the image with the median intensity of the pixels in an x by y pixel local neighborhood. Neighborhood operations are commonly used to enhance an image before the deployment of an intensity threshold to generate a binary image. A particularly powerful subset of neighborhood operations—which may be used for both image-processing and analysis—is grounded in the theory of mathematical morphology (Soille, 2003). Morphological operations are based on shape and modify the original (grayscale or binary) image by comparing it, using set theory, with a neighborhood, or structuring element, of a particular size and shape. Morphological operations differ from other neighborhood operations because they are based on a comparison between the original image and the structuring element itself rather than on being computed from the values of a set of neighboring pixels.

The four most promising image-processing procedures are discussed below.

Procedure 1: Simple Threshold Approach

This procedure (Fig. 2(A)) is based on that of Butler, Lane, and Chandler (2001) and Reid, Rice, and Garcia (2001) and utilizes a simple intensity threshold

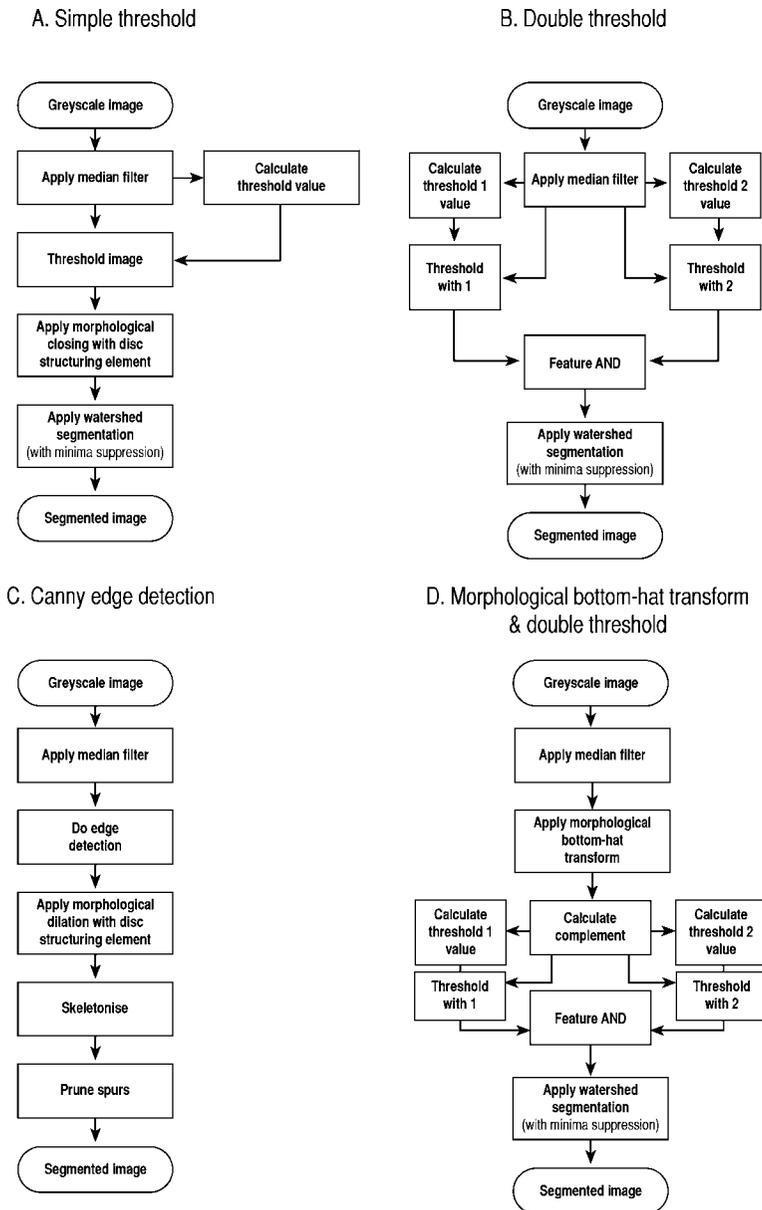


Figure 2. The generation of a segmented image using the four image-processing procedures described. A. The simple threshold approach. B. The double threshold approach. C. Application of the Canny edge detector. D. Application of a morphological bottom-hat transform and double threshold approach.

to generate the first partitioning of the image. The procedure works by assuming that high intensity values are associated with grains and low intensity values are associated with shading of the interstices between grains. The quality of the segmentation obtained is highly dependent on the selection of an appropriate threshold value. There are many automated procedures for selecting threshold values based on properties of the intensity frequency histogram (see review by Sahoo and others, 1988). Some of these procedures may be applied to the unimodal, sharply peaked, intensity frequency distributions that are characteristic of images of gravel surfaces, but there is no theoretical reason to believe that any are optimal for this application. Butler, Lane, and Chandler (2001) found that automated procedures performed poorly for such images and so they set the threshold subjectively, assessing the effects of different values on the resulting grain-size distribution. Since this would be a time-consuming process where many images collected under different conditions must be processed, an alternative method of defining the threshold value was sought.

The method selected defines the threshold in terms of a given percentile of the intensity frequency histogram rather than as an absolute intensity value. Such an approach minimizes the problems associated with variations in, for example, lighting conditions between images. The assumption implicit in this approach is that the general form of the intensity frequency histogram will be similar for all images of gravel river beds and that an empirically determined percentile value used for thresholding will, therefore, be generally applicable.

A problem with the simple threshold approach is that surface texture within individual grains can lead to fragmentation of the objects in the partitioned image. However, setting a threshold level that minimizes this fragmentation may result in poor definition of the grain boundaries. Applying a median filter before thresholding can reduce the effects of this by smoothing out noise within grains whilst preserving edges (Russ, 1999).

Procedure 2: Double Threshold Approach

Even after the application of a median filter, the selection of a single threshold level that adequately defines the grain boundaries whilst minimizing fragmentation may be impossible. A refinement of the simple threshold approach (Fig. 2(B)) utilizes two threshold values in an attempt to overcome this problem (Russ, 1999). The first threshold value is selected to identify the grain boundaries, accepting that there will be noise within some of the grains. The second threshold is selected to identify the darkest parts of the image—those that definitely represent interstices. A *feature-AND* operation is then undertaken, selecting all features in the first thresholded image that are connected to features in the second thresholded image. This approach results in a significant improvement in grain delimitation compared

to a simple threshold approach. As in the simple threshold approach, the two threshold values are defined in terms of percentiles of the intensity frequency distribution.

Procedure 3: Canny Edge Detection

Both simple and double threshold approaches are based on global operations that utilise constant threshold values across each image. Where the images contain grains of mixed lithologies, where the shading varies across individual grains, or where the lighting is not uniform across the image, constant threshold values may result in poor segmentation of parts of the image. Segmentation based on neighborhood operations may overcome these limitations. One class of neighborhood operations includes edge-detection algorithms, which identify steep intensity gradients in images. The Canny (1986) algorithm is particularly powerful as it is less likely to be confused by noise than other approaches, whilst still detecting weak edges. The output of the algorithm is a binary edge map. The edges in this are poorly connected, so, they are first dilated (thickened) to ensure good connectivity and then skeletonized (thinned) in preparation for object measurement (Fig. 2(C)). Dilation and skeletonization are types of morphological operation (Soille, 2003). McEwan and others (2000) applied this approach to high-resolution laser altimeter data of artificial mixes of fine-grained sediment ($D_{50} \approx 5$ mm), with satisfactory results. They argued that the technique could be applied to digital photographs, given further development.

Because it is based on a neighborhood operation, edge detection is able to cope well with variations in grain lithology and uneven illumination. The procedure produces good grain boundaries for well-sorted sediments having clasts with rounded edges and few surface features. As with the previous approaches, the initial application of a median filter may help to reduce fragmentation associated with angular grains and grain-surface texture. A particular limitation of this technique is that small grains are lost during edge dilation.

Procedure 4: Morphological Bottom-Hat Transformation and Double Threshold

An alternative neighborhood operation (Fig. 2(D)) utilizes a bottom-hat (also called: black top-hat, top-hat by closing) transformation to identify small, dark parts of an image. Several different operations, with similar functions but differing implementation, are described in the literature (Jackway, 2000). The operation used here is the morphological operation introduced by Meyer (1979). This is implemented as the difference between an image and its morphological

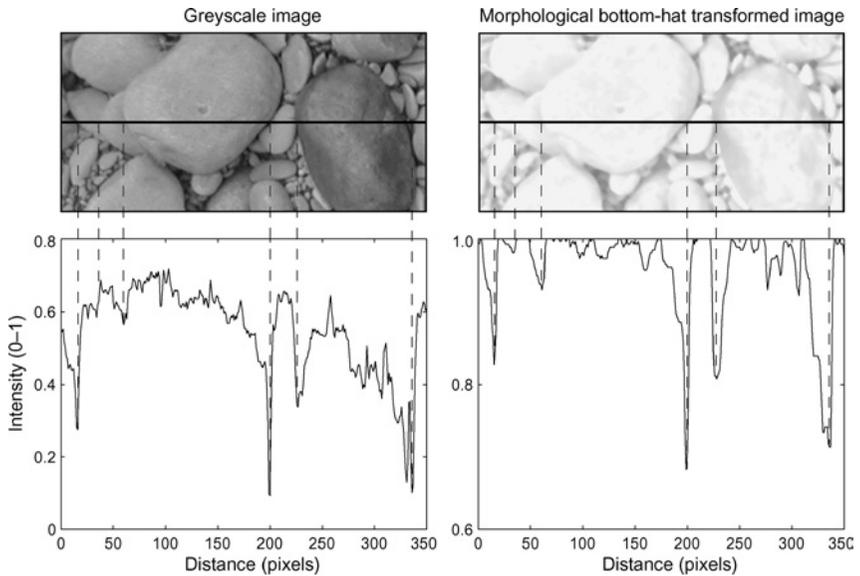


Figure 3. Illustration of the effect of the application of the morphological bottom-hat transform. The top panels show the original and transformed images. The bottom panels show an intensity profile (scaled from 0 to 1) across the images. Vertical (dashed) lines show the locations of the grain boundaries along the intensity profiles. Differences in intensity between and across grains in the grayscale image complicate the selection of appropriate threshold values. These differences are largely removed by the application of the transform and most of the intensity variations remaining are associated with grain boundaries.

grayscale closing (Soille, 2003). When applied to images of river beds, this approach enhances the interstices between grains relative to the grains themselves, simplifying the selection of appropriate threshold values (Fig. 3). The initial application of a median filter reduces the response of the bottom-hat transformation to intra-grain noise. Applying a double threshold approach (based on percentiles of the intensity frequency distribution in the bottom-hat transformed image) to the transformed image successfully partitions the image.

Watershed Segmentation

Following the simple and double threshold and bottom-hat transformation approaches, regions representing separate grains in the resulting binary image are frequently connected and additional processing is required to separate them into individual objects. The dilation/skeletonization procedure performs this role for the Canny edge detection approach. With other approaches, a watershed segmentation

is commonly applied to overcome such problems (Vincent and Soille, 1991). This is derived from the Euclidean Distance Map (EDM) of the binary image, in which each pixel that forms part of a “grain” is given a value inversely proportional to its distance from the nearest ‘non grain’ pixel. The resulting grayscale image may be envisaged as a topographic surface in which ‘grains’ are represented by depressions. The ‘watersheds’ between these depressions may be used to segment the binary image. Each depression is gradually ‘flooded’ until ‘water’ from one depression overflows into its neighbor. The line along which this occurs is then marked as a watershed and the flooding and marking continues until the image is entirely submerged. The watersheds thus defined are then used to segment the binary image. Several watershed algorithms have been proposed. That which is used here is the slightly modified version of the Vincent and Soille (1991) algorithm implemented in MATLAB.[®]

The watershed-segmented image has one particle for each regional minimum in the EDM. Unfortunately, the watershed procedure commonly results in significant over-segmentation of the binary image (Bleau and Leon, 2000). A number of approaches have been developed to overcome this problem. Marker-based watershed segmentation imposes a single regional minimum on the EDM for each particle, based on some *a priori* knowledge of the locations of individual grains (e.g. Vincent and Dougherty, 1994). Several methods have been proposed for identifying such markers (e.g. Beucher, 1992; Talbot and Vincent, 1992), but this may be the most difficult stage of the image-processing procedure and it has not proved possible to develop a robust method of uniquely marking each grain automatically. Development of region-merging approaches for reducing the over-segmentation are a promising area of current research (Bleau and Leon, 2000; Vanderstockt and Whyte, 2002), but they are not currently implemented in commercially available software.

An alternative and simple method to reduce the over-segmentation is to modify the EDM to remove shallow regional minima that lead to over-segmentation but probably do not uniquely mark a grain (Eddins, 2002). This may be achieved using the *h-minima transform*, which removes all minima in the EDM whose depth is less than a given threshold value (Soille, 2003). The application of watershed segmentation to the modified EDM results in a reduction in the degree of over-segmentation, but at the cost of under-segmenting some of the grains (Fig. 4). Nevertheless, this procedure currently provides the best results.

Object Measurement

Once the image has been satisfactorily segmented, the final procedure is to identify uniquely each object (grain) in the segmented image and measure it. Identifying and labeling individual objects is achieved using standard software

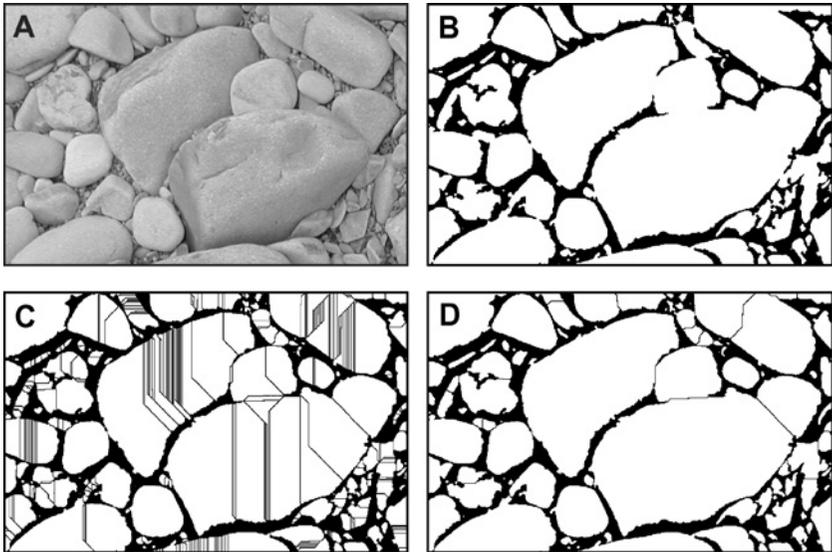


Figure 4. Example of the application of the watershed segmentation algorithm. A. The original grayscale image. B. The image after the application of the morphological bottom-hat filter and double threshold. Some adjacent grains have not been completely separated. C. The result of applying the watershed segmentation algorithm to B. There is significant over-segmentation of many of the grains. D. The result of applying the watershed segmentation algorithm after the application of the *h-minima* transform with a threshold value of 1. The segmentation is much improved, with many formerly joined grains being separated and over-segmentation removed.

tools. A number of measures of object size may be envisaged, including area or perimeter length (by counting pixels) or the diameter of the circle with the same area (Ghalib and Hryciw, 1999). By convention, the *b*-axis of a grain is used as an index of size in studies of river sediments. This is defined as the maximum length orthogonal to the longest dimension, or *a*-axis, of the grain. In a photograph, the *a*-*b* plane of the grain may not be coincident with the plane of the image and only the projected axes of the grain may therefore be measured from the image. Nevertheless, the projected minor axis of the image object provides the best estimate of the true *b*-axis of the grain obtainable from a photograph, and is therefore the desired property to measure. The effects of grain orientation (and hiding by other grains) are discussed in the companion paper.

Measurement of the minor axis of each object in the labeled image may be based on some mathematical property of each object or on direct measurements. For 25 objects in a single image, Butler, Lane, and Chandler (2001) argued that measuring an ellipse with a second moment that is the same as that of the object provides an adequate estimate of the actual size of the object as measured with a

ruler. In contrast, McEwan and others (2000) described an algorithm for directly measuring the minor axes in an image. This algorithm performed well for the majority of objects, but the results are indeterminate for objects with large perimeter concavities (such as may occur, for example, where sediment grains overlap).

The use of the ellipse-fitting method has the significant advantages that it is simple to implement (it is a standard feature in image-processing software) and considerably faster than direct measurements. However, its use can only be justified if it provides an unbiased estimate of the true minor axis of each object, a result that is yet to be adequately demonstrated. To this end, we developed an alternative algorithm to that of McEwan and others (2000), which gives a unique and correct solution for each object, regardless of shape. The major-axis length is calculated as the Euclidean distance between the two furthest pixels in the perimeter of the object and the minor axis is defined as the maximum Euclidean distance across the object along a line perpendicular to the major axis. This is implemented by drawing a line across the object perpendicular to the major axis and calculating the distance between the intersections of this line and the region perimeter. Since the line may intersect the perimeter in more than two locations (owing to concavities in the grain boundary), the maximum Euclidean distance between intersecting points is the distance of interest. This calculation is repeated for lines passing through each pixel along the major axis in turn and the minor axis is the maximum of the distances calculated.

The output of this algorithm was compared with ellipse-based measurements for 900 manually digitized grains. These grain boundaries were generated from two images from each of the three field sites. The results are excellent for objects larger than approximately 12 pixels across (about 8 mm) (Fig. 5). As the object size approaches the pixel size the proportional errors increase because an ellipse is not able to satisfactorily represent the true object shape. For objects with a minor axis larger than 12 pixels, the regression coefficient is 0.95 with an R^2 value of 0.99 and a RMS error of 3.1 pixels. The error in the median object size (P_{50}) is 0.05 pixels (0.1% of P_{50}). The errors at the tails of the size distribution are larger, but still acceptable, with a P_5 error of -0.6 pixels (-4%) and a P_{95} error of 3.9 pixels (3.6%). These results indicate that the use of ellipse-based measurements of object size is appropriate for this application.

OPTIMAL IMAGE-PROCESSING PROCEDURE SELECTION

Control Dataset Generation

Since the aim of the image-processing procedures is to remove the time-consuming requirement to digitize the grain boundaries in each image manually, a fair test of the image-processing procedures is their ability to identify the grains

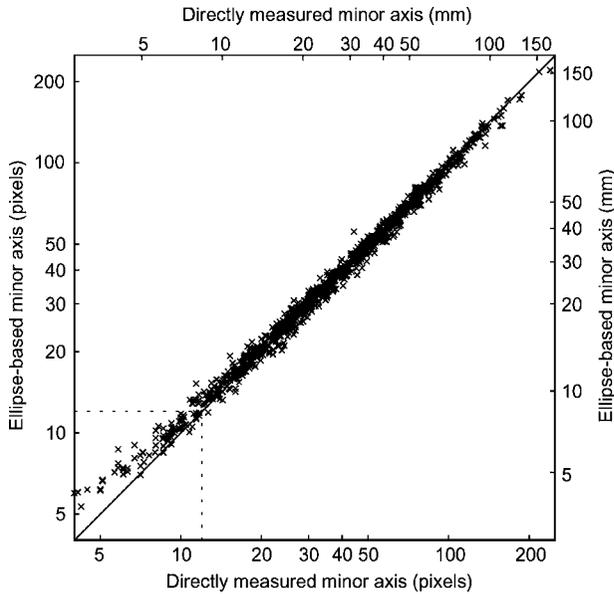


Figure 5. Performance of the ellipse-fitting procedure relative to direct measurement of the minor axes of 900 manually digitized objects. The line represents a 1:1 relation. The second set of axes shows the object sizes in mm, each pixel representing approximately 0.7 mm.

that are recognizable by a human operator. For each of the 39 field sediment patches, the digital photograph with the clearest definition of particle boundaries was selected for manual digitization. This was found generally to be an image collected under diffuse natural light with an overhead, camera-mounted flash. A regular grid was laid over the image on the screen in a vector drawing package and the boundaries of the grains beneath each grid intersection were digitized on-screen using a mouse. The grid spacing was chosen to maximize the number of grains selected whilst minimizing the number of grid intersections that fell on individual or adjacent grains. The digitized image was saved as a binary bitmap at the same resolution as the original image. Where adjacent grains had been digitized, the conversion to a bitmap sometimes resulted in them being merged. This was the primary reason for using the grid-based approach instead of digitizing every grain in the image. The problem was overcome by tidying each digitized image using a drawing package in order to separate touching grains.

For this test stage, the images were not rectified to correct spatial distortions because such rectification requires the interpolation of pixel values onto a new spatial coordinate system and the derived image is one step removed from the

original data. Without rectification, the real area represented by a pixel varies slightly between images, so it is not possible to convert the measured sizes of the object in the image into spatial units precisely. Since, for individual images, the scale of the manually digitized control grains and equivalent image-processing-derived objects is the same, this has no bearing on the assessment of the ability of the image-processing procedures to identify grains.

Test Methodology

For each of the four image-processing procedures, tests were carried out to identify the range of values of the internal parameters that gave acceptable results. The values of the parameters eventually selected are presented in Table 1 along with the total number of permutations required to give every possible combination of parameters for each procedure (a total of 416). The same digital photographs of the 39 sediment patches used for the generation of the control images were processed using each image-processing procedure with all permutations of the internal parameters, producing a total of 16,224 segmented images.

The centroid of each of the grains in the manually digitized control images was determined and used to select the equivalent object in each of the image-processing-derived segmented images. The length of the minor axis of the control grains and of the image-processed equivalents were then measured using the ellipse-fitting procedure outlined above to determine their size in pixels. A database was created containing the paired (control and image-processing) minor-axis lengths, grain by grain. Where the control grain had no equivalent object in the image-processing-based image, a place-holder was used in the database. For the purposes of selecting the optimal image-processing procedure, the data were aggregated for each field site, giving between 1000 and 2000 control grains per site. With three field sites and 416 permutations of the image-processing procedures, this gave 1248 sets of data for evaluation, each containing between 1000 and 2000 pairs of values.

At this stage, the data were truncated to remove the smallest objects. There are two justifications for this. First, as the object size approaches pixel size, the percentage error associated with measuring the minor axis increases. Second, there is some error associated with the manual digitization of the boundaries of the control grains. This takes two forms. Positioning the grain boundaries with a precision greater than a few pixels is difficult in practice, representing a significant proportional error on smaller grains. In addition, for the smallest particles it is not possible to identify clearly the boundaries between adjacent grains. In practice, it was found that a truncation at 23 pixels (approximately 16 mm at the scale of imagery used) removed the worst of these effects and this value was used throughout the subsequent evaluation of the image-processing procedures. A full investigation of the effect of the truncation point is discussed below.

The performance of each image-processing procedure was assessed using three derived statistics. The first was designed to assess how precisely the procedures measured the grains they identified. RMS error was selected because it provides a measure of the overall precision of the image-processing-derived minor-axis measurements. However, those grains in the control dataset that have no pair in the image-processing dataset (i.e. those where the image-processing procedure had failed to identify the grain) are excluded from this statistic. Procedures that produce good estimates of the sizes of those grains they identify may still miss grains.

To overcome this problem, two additional statistics were used to assess the overall quality of the derived object sizes. The two-sample Kolmogorov-Smirnov statistic assesses whether two samples have the same distribution. The associated *p-value* gives the probability of observing the given sample distribution if the source population has the same distribution. A larger value gives us greater confidence in the performance of the image-processing procedure. The advantage of this statistic is that the *p-value* takes into account the sample size and allows the direct comparison of samples of different sizes (such as those from the three field sites). However, the statistic operates by measuring the greatest difference between the cumulative size-distribution curves of two samples and thus gives no indication of the overall fit of the two distributions. The third statistic was designed to overcome this limitation and is the sum of the magnitudes of the differences between seven predefined percentiles across the entire object-size distribution ($P_{5,16,25,50,75,84,95}$). Therefore, it provides an approximation of the area between the two cumulative size distribution curves. Good performance of an image-processing procedure is then defined as that which minimizes the RMS error and sum of the magnitudes of the percentile errors whilst maximizing the Kolmogorov-Smirnov test *p-value*.

Results of the ‘Optimal’ Image-Processing Procedure

Given the large number of datasets evaluated (1248), it is not possible to present all of the results here (they are available at the website described below). Figure 6 gives, for one of the performance measures, an example of the results achieved by each of the image-processing procedures for a small number of the possible permutations of internal parameters (bold/italic values in Table 1).

The analysis demonstrates that, for particular sets of internal parameters, good results are achievable at individual sites with all four of the image-processing procedures. Furthermore, for each of the sites, the procedure that achieves the best results is different (Table 2). However, the objective here is to define the procedure and the set of internal parameters where the performance is *consistently* good for all three sites and, therefore, may be considered ‘optimal’ for a range of lithotypes and grain morphology. As might be expected on theoretical grounds, those procedures

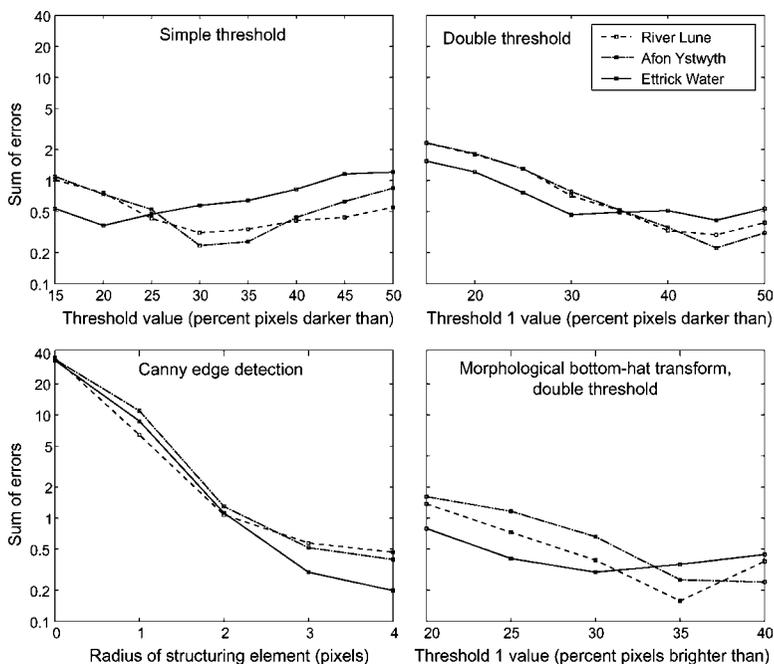


Figure 6. Example of the performance of each of the image-processing procedures with specified internal parameters. The other internal parameters are bold/italic values in Table 1. Data for each river are aggregated. The performance indicator is the sum of the errors at seven percentiles in the size distribution; low values represent better performance (note log scale). All of the examples were filtered with a 5×5 median filter.

that are based on neighborhood operations (the Canny and bottom-hat approaches) perform most consistently between sites. For these procedures, it is possible to define values of the internal parameters that perform almost as well at all three sites as the best performing procedure at each site. In practice, the differences between some of the procedures and internal parameter sets are marginal, but overall the bottom-hat approach produces the best results and the appropriate internal parameters are specified in Table 2.

PERFORMANCE EVALUATION OF THE OPTIMAL PROCEDURE

At this point, it is important to remember that, here, we are concerned with accurate sizing of objects in an image rather than the grain sizes of the photographed sediment. The object-size frequency distributions and percentiles considered below are not fluvial sediment grain-size distributions per se. They

Table 2. Definition of the Optimal Image-Processing Procedure and the Procedure that Performs Best at Each Site

	Median filter	Threshold 1 (%)	Threshold 2 (%)	Structuring element radius
Overall optimal procedure				
Morphological bottom-hat transform and double threshold	5 × 5	35	1	N/A
Best procedure at individual sites				
River Lune				
Morphological bottom-hat transform and double threshold	5 × 5	35	4	N/A
Afon Ystwyth				
Double threshold	None	35	3	N/A
Ettrick Water				
Canny edge detection	5 × 5	N/A	N/A	3

Note. N/A means not applicable.

require conversion to grain-size distributions and percentiles using the corrections discussed in the companion paper.

Precision for Aggregated Data

The first test of the performance of the optimal image-processing procedure defined above is how well it is able to reproduce the object-size distributions for each of the three sites using the aggregated data of each site. The results are extremely encouraging, with the cumulative-frequency plots of the image-processing data showing excellent correspondence with the manually derived control data at all three sites (Fig. 7(A)) and the percentiles of the control and image-processing frequency distributions showing only small deviations from a 1:1 relation (Fig. 7(B)). The precision at each of five percentiles is given in Table 3. Many values are sub-pixel in magnitude. In general the largest errors (in both absolute and percentage terms) are at the coarse end of the distributions as would be expected. Overall, the technique performs best on the Afon Ystwyth alluvium.

Because the objective of this paper is to produce a technique that is of utility at sites with a wide range of sedimentary characteristics, the selected (optimal) technique is a compromise and not tailored specifically to any one field sites. It is therefore of interest to examine the extent to which the technique that performs best at each individual site outperforms that which performs best overall. These estimates are presented in the final column of Table 3. Although there is an increase in precision associated with what is locally the best technique, the improvement is generally very small, engendering confidence in the prospect

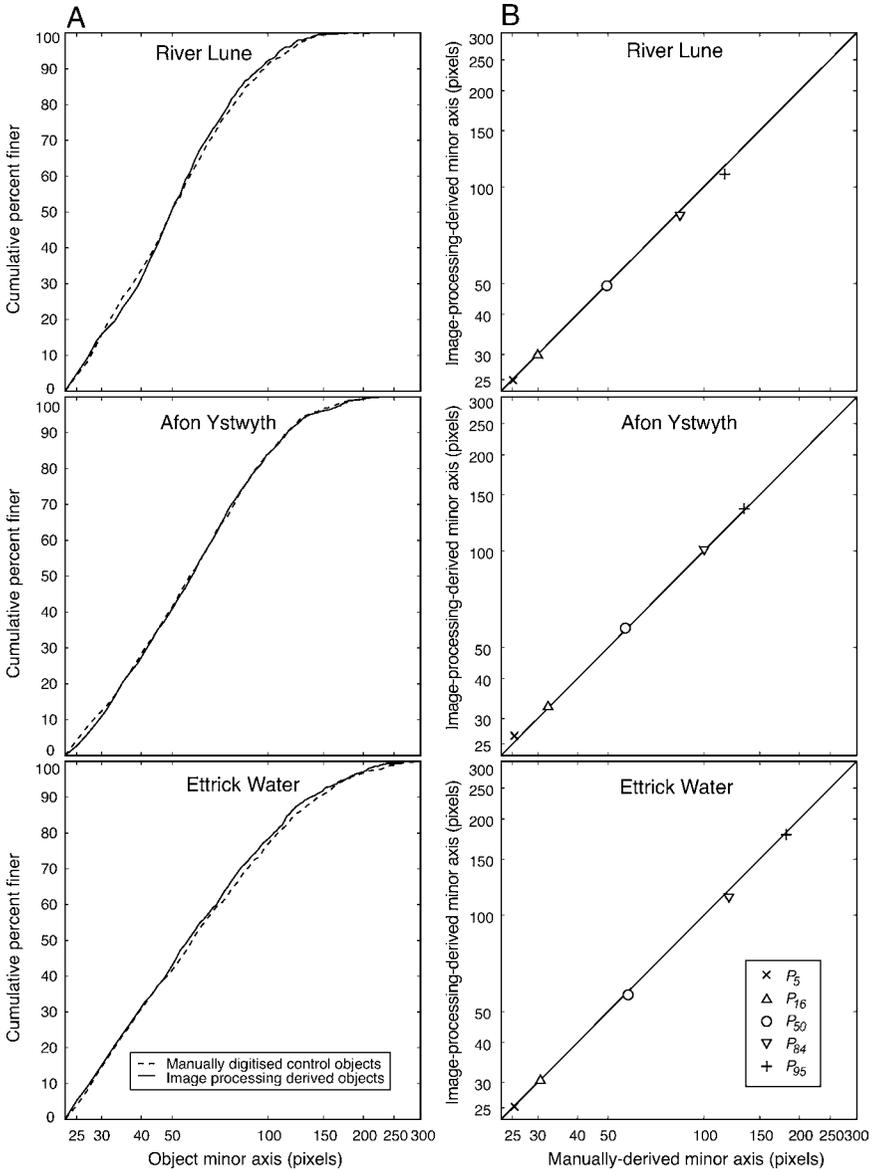


Figure 7. Performance of the optimal image-processing procedure at each of the three field sites. Data for each river are aggregated. A. Cumulative object-size distributions of the control (manually digitized) and image-processing procedure. B. Object sizes at specific percentiles of the size distribution as defined by image-processing and manual measurement.

Table 3. The Performance of the Optimal Procedure and the Best Procedure at Each Site

Percentile	Manually digitized control		Image-processing: optimal procedure		Image-processing: best procedure for site	
	Minor axis length (pixels)	Minor axis length (pixels)	Minor axis length (pixels)	Error (pixels [%])	Minor axis length (pixels)	Error (pixels [%])
River Lune						
P_5	25.2	25.1	25.2	-0.1 [0.4]	25.2	0.0 [0.1]
P_{16}	30.2	29.9	30.0	-0.2 [0.7]	30.0	-0.1 [0.2]
P_{50}	49.8	49.3	49.5	-0.5 [1.0]	49.5	-0.3 [0.6]
P_{84}	84.2	80.8	82.5	-3.3 [3.9]	82.5	-1.6 [1.9]
P_{95}	116.3	110.0	112.0	-6.3 [5.4]	112.0	-4.3 [3.7]
Afon Ystwyth						
P_5	25.5	26.3	26.1	0.8 [3.3]	26.1	0.6 [2.3]
P_{16}	32.4	32.8	32.9	0.4 [1.2]	32.9	0.4 [1.4]
P_{50}	56.9	57.5	57.5	0.6 [1.1]	57.5	0.6 [1.0]
P_{84}	100.1	100.6	100.8	0.5 [0.5]	100.8	0.7 [0.7]
P_{95}	134.3	135.7	136.1	1.4 [1.1]	136.1	1.8 [1.4]
Eitrick Water						
P_5	25.5	24.9	24.9	-0.6 [2.3]	24.9	-0.6 [2.4]
P_{16}	30.8	30.6	30.3	-0.3 [0.9]	30.3	-0.5 [1.6]
P_{50}	58.0	56.3	58.0	-1.6 [2.8]	58.0	0.0 [0.0]
P_{84}	120.4	113.8	120.9	-6.7 [5.6]	120.9	0.5 [0.4]
P_{95}	181.5	179.8	183.7	-1.8 [1.0]	183.7	2.2 [1.2]

Note. A pixel represents approximately 0.7 mm.

Table 4. The Mean Errors in Estimates of Various Percentiles of the Object-Size Distributions at the Three Field Sites and Their 95% Confidence Intervals

	Mean error in percentile estimate and 95% confidence interval (pixels)				
	P_5	P_{16}	P_{50}	P_{84}	P_{95}
River Lune	-0.38 ± 1.22	0.97 ± 2.50	-1.06 ± 1.79	-4.79 ± 2.78	-5.34 ± 2.92
Afon Ystwyth	0.88 ± 1.40	0.75 ± 1.97	1.06 ± 3.47	-2.43 ± 2.73	4.62 ± 7.47
Ettrick Water	-2.02 ± 3.38	-3.24 ± 2.69	-4.09 ± 4.25	-6.47 ± 4.51	-9.62 ± 8.82

Note. A pixel represents approximately 0.7 mm.

of deploying the morphological bottom-hat transformation and double threshold approach universally.

Precision for Individual Photographs

The second test of the optimal-image-processing procedure is how well it is able to reproduce the object-size distribution of each individual image. Although there is inevitably a little more scatter than for the aggregated data, once again the results are very encouraging. Space prohibits the presentation of cumulative-size distribution curves for each of the 39 images used, but comparison of the percentiles of the control and image-processed distributions shows a satisfactorily small deviation from a 1:1 relation (Fig. 8). Table 4 summarizes the mean error in various percentile estimates site by site along with the 95% confidence interval for each. However, it should be noted that a single image (from the Afon Ystwyth) has been excluded from this analysis owing to its exceptionally poor performance. Examination of the segmented and original images revealed that segmentation had failed in the corners of the image where patches of very fine sediment coincided with significant chromatic aberration and slight optical vignetting. This serves as a warning that even a procedure that performs generally well may fail in specific circumstances. It is recommended, therefore, that the segmentation of each image be briefly examined to check for gross errors before relying on the data.

Effect of Varying the Size Truncation

Throughout the preceding analysis the data have been truncated at 23 pixels (approximately 16 mm). Understanding the relation between the truncation point and the quality of the resulting object-size distribution is important because, for a given minimum resolvable object size, the areal coverage and the number of grains in the image are inversely proportional to the square of the truncation size. Thus,

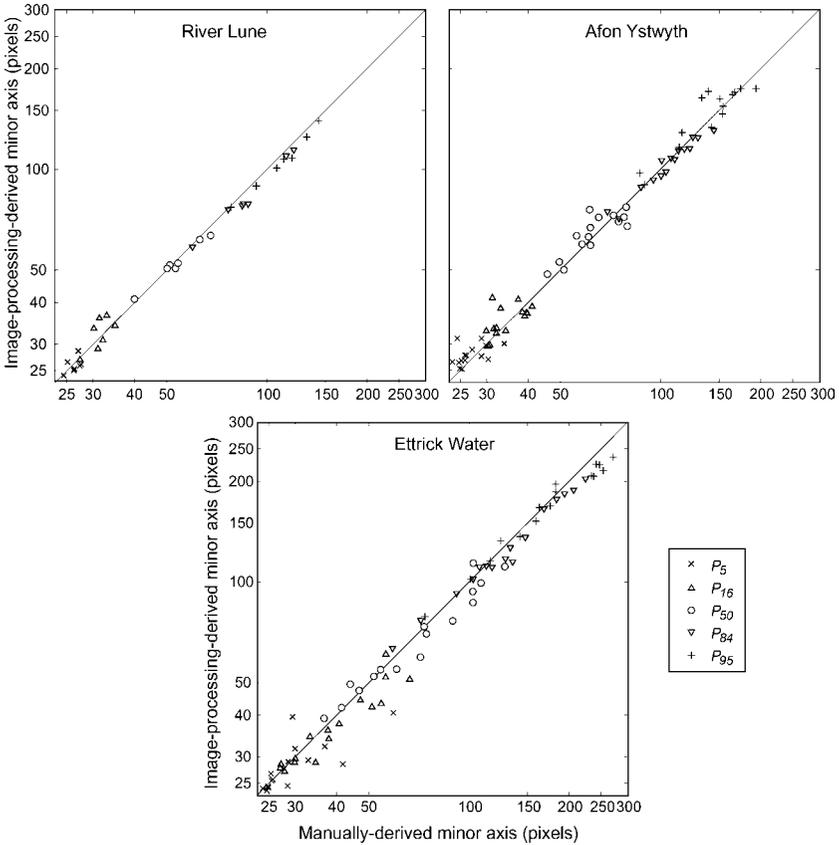


Figure 8. Object size at specified percentiles of the size distribution as defined by the optimal image-processing procedure and manual measurement for each of 39 patches grouped by field site.

it is desirable to use the smallest truncation size that gives satisfactory results in order that areal coverage and sample size are maximized.

Using the aggregated data for each of the field sites, an assessment was made of how different levels of truncation affect the quality of the resulting object-size distributions. Figure 9 gives the results of this analysis. Performance is represented by the error (image-processing minus control) in each of two percentiles (P_5 and P_{50}) expressed as a percentage of the appropriate control-percentile value for truncations of between 0 (no truncation) and 100 pixels. Although there is a lot of noise, the lower percentile does seem to reach a reasonably consistent level at just above 20 pixels. The higher percentile is not particularly sensitive to the truncation point except in the case of the Etrick Water.

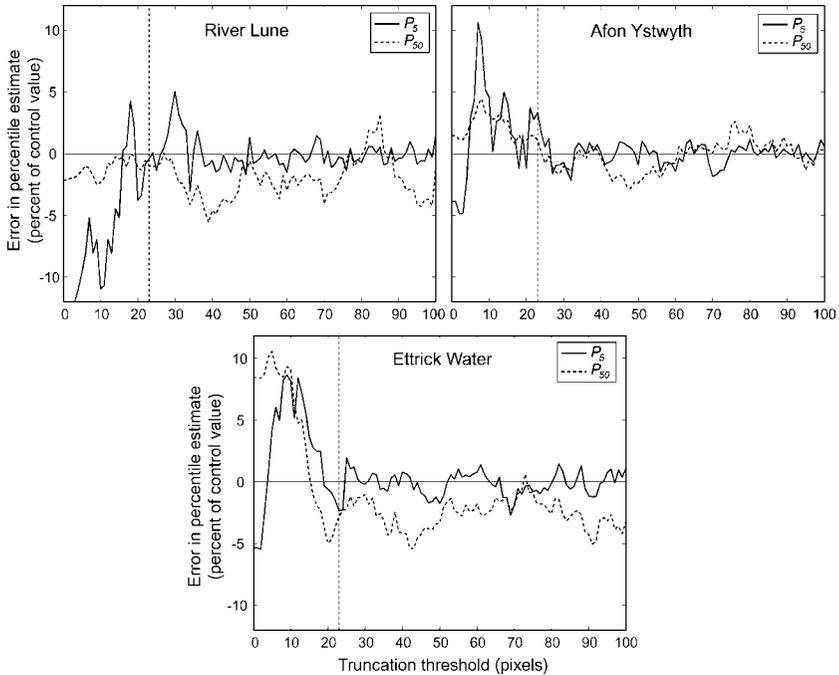


Figure 9. The effect of varying the size truncation on the error (image processed minus manually determined) in the P_5 and P_{50} at each of the three field sites. Values normalized by expressing the errors in the percentile estimates as a percentage of the manually digitized percentile value. The vertical (pecked) line is the truncation applied in this paper.

Based on this analysis, a truncation point of 23 pixels seems appropriate. At the scale of imagery used in this paper, this translates into a truncation at about 16 mm. It seems reasonable to assume that this truncation will apply at all scales of photography. In order to be able to truncate at 8 mm, the conventional point used by many geomorphologists when ‘Wolman sampling’, the areal coverage would have to be reduced from 1.2 to 0.4 m². For a given focal length, ground coverage is proportional to camera height. However, because the focal length can be varied using the camera zoom (and, on compact camera models, is often not known), specifying the minimum resolvable object size as a function of areal coverage rather than camera height may be more useful. Therefore, for a range of common camera resolutions, Figure 10 gives the relation between minimum resolvable grain size (assuming a truncation at 23 pixels) and areal coverage. It should be noted, however, that other factors, such as the quality of the camera optics, the size and type of recording array, and the on-camera image-processing are also likely to affect the minimum resolvable object size.

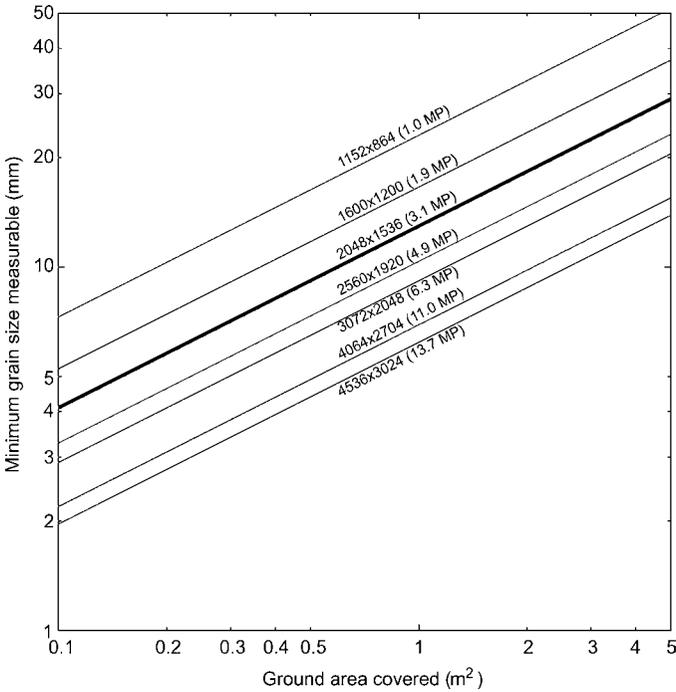


Figure 10. The relation between minimum resolvable grain size and ground area coverage for various camera resolutions. The bold line represents the camera used in this study (2048 × 1536 pixels). MP = megapixels.

CONCLUSIONS

- A. It is relatively easy to develop image-processing procedures that satisfactorily segment small numbers of images collected at a single site, but much less easy to develop a technique that is transferable between sites.
- B. As expected from theory, operations based on pixel neighborhoods are the most successful because they are tolerant of variations in lighting conditions and grain hue. This paper presents a procedure, based on a morphological transformation that has been shown to perform well on 39 images collected at three sites with very different sediment characteristics. This procedure outperforms previously published, or improvements on previously published, procedures.
- C. The optimal procedure utilizes a morphological bottom-hat transform and double threshold approach.

- D. The procedure requires no user intervention to set appropriate values of the internal parameters or correct poor segmentation beyond operator inspection to reject rare cases of error.
- E. The smallest grain that can be reliably resolved has an image size of about 23 pixels. It is thought that this is a general result that will apply at different scales of photography, but one that may be affected to some extent by the quality of the camera.
- F. The derivation of grain-size distributions requires several additional steps, such as the correction of spatial distortions and fabric errors. The application of the image-processing procedure developed here to the derivation of grain-size data is discussed in a companion paper.

SOFTWARE

All image-processing and analysis were undertaken with MATLAB[®] and the Image Processing toolbox (v3.2) because they are widely available within academia and industry. They provide a powerful development environment and programming language. They also enable the generation of stand-alone applications that can be run independently of MATLAB[®]. Further information is available via <http://www.lboro.ac.uk/research/phys-geog/>.

ACKNOWLEDGMENTS

The work was supported by the Leverhulme Trust and Loughborough University. The Countryside Council for Wales, the Tweed Foundation and several individual landowners are thanked for permission to access field sites. Assistance in the field was provided by Tom Buffin-Bélanger, Phillippa Noble and Georgina Endfield and the manual digitizing was performed by Natasha Todd-Burley. Celso Garcia undertook much of the initial “proof of concept” image-processing work.

REFERENCES

- Adams, J., 1979, Gravel size analysis from photographs: J. Hydraulics Division, ASCE, v. 105, no. 10, p. 1247–1255.
- Bankman, I. N., ed., 2000, Handbook of medical imaging: Processing and analysis: Academic Press, London, 901 pp.
- Bernard-Michel, B., Rohani, S., Pons, M. N., Vivier, H., and Hundal, H. S., 1997, Classification of crystal shape using fourier descriptors and mathematical morphology: Particle and Particle Syst. Characterization, v. 14, no. 4, p. 193–200.
- Beucher, S., 1992, The watershed transformation applied to image segmentation: Scanning Microscopy, v. Supplement 6, p. 299–314.

- Bleau, A., and Leon, L. J., 2000, Watershed-based segmentation and region merging: *Computer Vision Image Understanding*, v. 77, no. 3, p. 317–370.
- Butler, J. B., Lane, S. N., and Chandler, J. H., 2001, Automated extraction of grain-size data from gravel surfaces using digital image processing: *J. Hydraulic Res.*, v. 39, no. 4, p. 1–11.
- Canny, J., 1986, A computational approach to edge-detection: *IEEE Trans. Pattern Anal. Machine Intell.*, v. 8, no. 6, p. 679–698.
- Church, M., McLean, D., and Wolcott, J. F., 1987, River bed gravels: Sampling and analysis, *in* Thorne, C. R., Bathurst, J. C., and Hey, R. W., eds., *Sediment transport in gravel bed rivers*: Wiley, Chichester, p. 43–79.
- Diepenbroek, M., Martholoma, A., and Ibbeken, H., 1992, How round is round? A new approach to the topic “roundness” by Fourier grain shape analysis: *Sedimentology*, v. 39, no. 3, p. 411–422.
- Diepenbroek, M., and De Jong, C., 1994, Quantification of textural particle characteristics by image analysis of sediment surfaces—examples from active and paleo-surfaces in steep, coarse-grained mountain environments, *in* Ergenzinger, P., and Schmidt, K. H., eds., *Dynamics and geomorphology of mountain rivers*: Springer-Verlag, Berlin, Heidelberg, p. 301–314.
- Eddins, S., 2002, The watershed transform: Strategies for image segmentation: <http://www.mathworks.com/company/newsletter/win02/> [10 September 2003].
- Friedman, G. M., 1958, Determination of sieve-size distribution from thin section data for sedimentary petrological studies: *J. Geol.*, v. 66, p. 394–416.
- Fripp, J. B., and Diplas, P., 1993, Surface sampling in gravel streams: *J. Hydraulic Eng.*, v. 119, no. 4, p. 473–490.
- Ghalib, A. M., and Hryciw, R. D., 1999, Soil particle size distribution by mosaic imaging and watershed analysis: *J. Comput. Civil Eng.*, v. 13, no. 2, p. 80–87.
- Glasbey, C. A., Horgan, G. W., and Darbyshire, J. F., 1991, Image analysis and three-dimensional modeling of pores in soil aggregates: *J. Soil Sci.*, v. 42, no. 3, p. 479–486.
- Horgan, G. W., 1998, Mathematical morphology for analyzing soil structure from images: *Eur. J. Soil Sci.*, v. 49, no. 2, p. 161–173.
- Ibbeken, H., Warnke, D. A., and Diepenbroek, M., 1998, Granulometric study of the Hanaupah Fan, Death Valley, California: *Earth Surface Processes Landforms*, v. 23, no. 6, p. 481–492.
- Ibeken, H., and Schleyer, R., 1986, Photo-sieving: A method for grain-size analysis of coarse-grained, unconsolidated bedding surfaces: *Earth Surface Processes Landforms*, v. 11, no. 1, p. 59–77.
- Jackway, P. T., 2000, Improved morphological top-hat: *Electr. Lett.*, v. 36, no. 14, p. 1194–1195.
- Kellerhals, R., Shaw, J., and Arora, V. K., 1975, On grain size from thin sections: *J. Geol.*, v. 83, no. 1, p. 79–96.
- McEwan, I. K., Sheen, T. M., Cunningham, G. J., and Allen, A. R., 2000, Estimating the size composition of sediment surfaces through image analysis: *Proc. Inst. Civil Eng., Water Maritime Eng.*, v. 142, no. 4, p. 189–195.
- Meyer, F., 1979, *Cytologie Quantitative et Morphologie Mathématique*: Unpublished Doctoral Dissertation, Ecole des Mines, Paris.
- Reid, I., Rice, S., and Garcia, C., 2001, Discussion of “The measurement of gravel-bed river morphology”, *in* Mosley, M. P., ed., *Gravel-bed rivers V*: New Zealand Hydrological Society, Wellington, p. 325–327.
- Rice, S. P., and Church, M., 1996, Sampling fluvial gravels: Bootstrapping and the precision of size distribution percentile estimates: *J. Sediment. Res.*, v. 66, no. 3, p. 654–665.
- Rice, S. P. and Church, M., 1998, Grain size along two gravel-bed rivers: Statistical variation, spatial pattern and sedimentary links: *Earth Surface Processes Landforms*, v. 23, no. 4, p. 345–363.
- Russ, J. C., 1999, *The image processing handbook*, 3rd edn.: CRC Press, Boca Raton, Florida.
- Sahoo, P. K., Soltani, S., Wong, A. K. C., and Chen, Y. C., 1988, A survey of thresholding techniques: *Computer Vision, Graphics Image Process.*, v. 41, no. 2, p. 233–260.

- Sime, L. C., and Ferguson, R. I., 2003, Information on grain sizes in gravel-bed rivers by automated image analysis: *J. Sediment. Res.*, v. 73, no. 4, p. 630–636.
- Soille, P., 2003, *Morphological image analysis: Principles and applications*, 2nd edn.: Springer-Verlag, Berlin, 407 pp.
- Talbot, H., and Vincent, L., 1992, Euclidean skeletons and conditional bisectors, *in* Maragos, P., ed., *Proceedings of SPIE on Visual Communications and Image Processing'92*, Vol. 1818, p. 862–876.
- Vanderstockt, Y., and Whyte, R. N., 2002, Watershed transformation: Reducing the over-segmentation problem by applying a noise reducer and a region merger: *J. WSCG*, v. 10, no. 3: http://wscg.zcu.cz/wscg2002/wscg2002_program.htm [14 October 2003].
- Vincent, L., and Dougherty, E. R., 1994, Morphological segmentation for textures and particles, *in* Dougherty, E. R., ed., *Digital image processing: Fundamentals and applications*: Marcel-Dekker, New York, p. 43–102.
- Vincent, L., and Soille, P., 1991, Watersheds in digital spaces: An efficient algorithm based on immersion simulations: *IEEE Trans. Pattern Anal. Machine Intell.*, v. 13, no. 6, p. 583–598.
- Wolcott, J. F., and Church, M., 1991, Strategies for sampling spatially heterogeneous phenomena: The example of river gravels: *J. Sediment. Res.*, v. 61, no. 4, p. 534–543.
- Wolman, M. G., 1954, Method of sampling coarse river bed material: *Trans. Am. Geophys. Union*, v. 35, no. 6, p. 951–956.