

RESEARCH METHODS PAPERS

INFORMATION ON GRAIN SIZES IN GRAVEL-BED RIVERS BY AUTOMATED IMAGE ANALYSIS

L.C. SIME* AND R.I. FERGUSON

Department of Geography, University of Sheffield, Sheffield S10 2TN, U.K.

email: l.sime@uea.ac.uk

ABSTRACT: The time required in the field to characterize textural variation over gravel surfaces can be reduced by taking vertical photographs for subsequent image analysis. We present modified edge-detection algorithms which combine edge seeding with an image porosity concept and partial watershed segmentation. The methods allow quick, reliable, and operator-independent size analysis from a wide range of vertical bed-surface images. They are tested using 24 naturally lit images of an exposed river bed with mixed lithologies and partial burial of gravel by sand. Grain-size percentiles derived by automated image analysis correlate closely with those from manual image analysis, with only small and consistent degrees of bias. They also correlate well with percentiles from field measurements with substantial bias, which, however, is consistent so that it can be corrected for, leaving a residual scatter of $\sim 0.25 \psi$ (where $\psi = \log_2 \text{mm} = -\phi$) over a wide range of bed conditions. The bias depends somewhat on sand cover, and the biggest residual discrepancies are for tail percentiles.

image analysis) and errors associated with estimating the sizes of three-dimensional grains from a two-dimensional image. Our computer algorithms are freely available and may help others seeking rapid quantification of spatial variations in gravel size distributions.

FIELD SITE AND DATA COLLECTION

As part of an investigation of within-channel patchiness in gravel GSDs and its implications for bedload transport, digital photographs of the exposed bed were obtained at about 500 sites evenly distributed across two mildly braided reaches of Vedder River in British Columbia, Canada. This river flows across an alluvial fan from the edge of the Cascade Mountains into the Fraser River trench, and has been the subject of several investigations of channel change, bedload transport, and downstream fining (e.g., Martin and Church 1995, Ferguson et al. 2001). The bed material comprises a wide range of heavily metamorphosed sedimentary and volcanic rocks of Paleozoic and later age, and some granitic intrusives, so there is a wide range of texture and color. The GSDs at the two reaches are mainly in the cobble and gravel ranges but with substantial amounts of surficial sand in many places. The photos were taken in winter so sun angle was fairly low, the surface had recently been rain-washed, and organic debris was present at some sites. These complications make AIA of these images more challenging than in well-lit laboratory conditions.

INTRODUCTION

Surface grain-size distributions (GSDs) in gravel-bed rivers are of interest to sedimentologists and others because of their role in regulating bedload flux and its GSD, their value in paleoenvironmental reconstruction, and their ecological significance for both invertebrates and fish spawning. Gravel GSDs are spatially variable at several scales: broad down-basin or down-river trends, facies differences along and across channel bars, and seemingly random local patchiness. Thorough characterization of gravel river beds therefore requires determining GSDs at large numbers of points. This is extremely time-consuming by traditional manual field methods of bulk sieving or pebble counting, so there has long been an interest in reducing the fieldwork time required by taking low-level vertical photographs of the bed and determining the GSD afterwards by manual measurements on the image (e.g., Adams 1979; Ibbeken and Schleyer 1984).

Recently other forms of image acquisition have been investigated (e.g., laser altimetry in a flume: McEwan et al. 2000; scanned photos: Butler et al. 2001), and digital cameras and developments in computer software have opened up the possibility of simple and rapid automated image analysis (AIA hereafter). McEwan et al. (2000), Butler et al. (2001), and Lane (2001) have shown that AIA offers great promise as a tool for rapid, non-destructive characterization of textural variation over space or time. However, these authors have noted some methodological problems, particularly with the limited number of bed-surface datasets previously used.

We present further developments in AIA for rapid grain-size determination in gravel-bed rivers, and test AIA on digital photos of a wide variety of natural gravel and gravel/sand surfaces from a Canadian river. Key features are (1) technical innovations to help process "difficult" images; (2) demonstration that our algorithms are robust, giving consistent results even for naturally lit images of gravels with complications such as mixed lithology, sand drapes, and moist patches; and (3) separate consideration of errors associated with AIA algorithms per se (as compared with manual

The camera had a 37 mm lens and a resolution of 1152×768 pixels. It was attached to a metal frame, pointing vertically down onto the bed, from a height of about 2 m. A 1-m-square frame was placed on the bed, and analysis was restricted to grains within this frame. The linear image scale was about 1.5 mm per pixel, which is quite coarse compared to the imagery used by McEwan et al. (2000) and Butler et al. (2001). Grain identification becomes unreliable with grains smaller than about 5×5 pixels, so grains with b axis < 8 mm were not counted separately but treated as fines. This particular cutoff is also widely used in manual pebble counting, so our results are comparable to standard procedures. Images were compressed to JPG format so that a few hundred could be stored on the camera flash card, but the compression was kept minimal to ensure low loss of picture information. Butler et al. (2001) show that the amount of image distortion produced by lenses similar to the one used here is not sufficient to warrant image rectification.

Test data were obtained in two ways: manual image analysis of 1 m^2 areas at 24 sites, and traditional field grain-size analysis at 12 of them. The manual image analysis was on the lines of a pebble count. Each image was overlaid with a 100-point grid and the visible b axis of each particle under a grid point was manually digitized. The median fines content at these sites was 6%, but four sites had $> 30\%$ fines, with a maximum of $> 60\%$. Figure 1 reproduces four of the photos to indicate the range of surface texture and condition. The field GSD determination involved spraying the surface with paint, removing all painted clasts from within the quadrat, sieving at 0.5ψ intervals down to 8 mm, and counting the number in each size class. These area-by-number GSDs were converted to a grid-by-number basis using a D^2 weighting in accordance with the literature (e.g., Church et al. 1987).

*Present address: School of Environmental Sciences, University of East Anglia, U.K.




FIG. 1.—Selected images of gravel and gravel/sand bed surfaces in Vedder River. They show approximately 1 m² of bed surface at a resolution of about 650 × 650 pixels.

AUTOMATED IMAGE ANALYSIS

Any kind of AIA method for deriving a GSD from a photographic or other image of a gravel surface involves three stages: (1) detect possible edges between grains, from their distinctively dark appearance or the rapidity of tonal variation across them; (2) clean up the initial edge map to minimize errors and ensure that edges are connected to define discrete grains; and then (3) determine properties of interest for the separate grains. We discuss these stages in turn, then evaluate the results of using different combinations of the choices involved at each stage.

Edge Detection and “Image Porosity”

Texts such as Gonzalez and Woods (1992) identify two main approaches to identifying particle boundaries from grayscale values (or other kinds of digital image intensity): thresholding either the image itself, or a first- or second-derivative map obtained by filtering the image. Image thresholding works on the assumption that dark areas are gaps between grains; derivative methods work on the assumption that the zones of most rapid tonal change are between grains, or at edges between grains and sand or shadow. Both approaches have been tried in previous work on AIA of gravels. McEwan et al. (2000) used a second-derivative mapping known as the Canny method, which computes the local maximum of the gradient of image intensity. This has the advantage of defining edges only a single pixel wide, so it may be the best option in situations with strong edge definition and low fines content, but it is sensitive to noise, which can be difficult to avoid

except with carefully controlled vertical lighting (Butler et al. 2001) or elevation imagery (McEwan et al. 2000). We test this method on our images, but also present results from a modified Sobel first-derivative and a threshold method.

One stage in the Canny method is edge seeding. We have developed algorithms which use it in two other edge-finding methods. Edge seeding involves using two threshold values to give “strong” and “weak” edge maps. The strong edges from the higher threshold are all accepted, but weak edges are included only if connected to strong ones. This filters out small false edges relating to intra-grain noise. We found that for naturally lit images this is an important modification. Butler et al. (2001) noted the problem of selecting thresholds for edge seeding. In the context of gravel surfaces we propose that it is helpful to consider what we term “image porosity”: the proportion of pixels in an image that are classified as edges. If these edges represent gaps between grains, the image porosity is a 2-D analogue of the volumetric porosity familiar to sedimentologists, and it seems reasonable to expect it to have a similar value in the final edge map. Because edge-detection methods differ in how they respond to areas of sand or shadow, the target image-porosity value should vary according to the choice of edge-detection method. We developed an iterative procedure in which a strong edge map is produced by a constant high threshold, and the low threshold is varied until the end product converges on the target image porosity. This is not the same as using a simple image-density histogram to select the weak map threshold, because that would count intra-grain noise as grain porosity. However, if there is little or no intra-grain




FIG. 2.—Edge maps generated from image B in Figure 1 using different edge-detection methods: A) Sobel; B) Canny; C) dark threshold; D) aggregate of Sobel and dark threshold.

noise a less sophisticated threshold from a image density histogram will give similar results.

A target was not set for the Canny method because dark intergrain areas are not counted as edges, making the porosity analogy inappropriate. In the Sobel method the target image porosity was set at 0.35 and the ratio of low and high thresholds was maintained at 2 as both were varied. In the image thresholding method, a local mean grayscale value computed over a sliding neighborhood was used for images with sparse fines. This deals with variable lighting, moisture, or grain grouping across the image: the threshold changes in accordance. For images with significant sand content a single global mean was used, to avoid sand patches being misidentified as single large grains. The global or local mean was then used with a high and low threshold to obtain a seeded edge map. Target image porosity was set at $0.35 + 0.65F$, where F is a visual estimate of the areal fraction of fines. This variation in image porosity allows accurate setting of the upper threshold to identify most sand as edge area, even where portions of the sand are quite dry. This is our recommended method for use where substantial areas (say 20% or more) of the image are taken up by fines.

It is possible to combine different edge maps using a logical OR operation; for example Butler et al. (2001) aggregated edges detected from radiometric and photogrammetric data. We have tried aggregating edges detected in different ways from radiometric data alone. The image-threshold method finds areas of sand more efficiently than other approaches, whilst the local-derivative methods find other edges more efficiently, so the two methods combined ought to be more effective than either on its

own. We combined a dark-threshold edge map with a skeletonized Sobel edge map (preferred to the Canny map because of its lower sensitivity to noise). The skeletonization minimizes excessive edge finding (e.g., McEwan et al. 2000).

Figure 2 shows examples of edge maps derived from the images in Figure 1B using these implementations of the dark-threshold, Sobel, Canny, and aggregate methods. Edges detected by the Sobel and Canny methods tend to occur a little inside the true grain edge. This tends to reduce grain areas and diameters slightly. Edges detected by dark thresholding tend to correspond better with true edges. Where the semi-global, sliding-neighborhood local mean value is used, even grain boundaries formed by grain shadows tend to be located at the true grain boundary rather than just inside.

Edge Tidying and Watershed Segmentation

Edges identified by any of these methods may not be completely linked (grains may not be completely segmented). Also, despite edge seeding and spatially varied thresholds, some false intra-grain edge noise remains. Secondary processing is therefore required to ensure complete linking of true edges and minimal survival of false edges which might interfere with subsequent edge linking.

Unconnected edges smaller than 0.0001 of the image area were removed. They generally relate to intra-grain dark spots such as grain chips, raindrop marks, or visible portions of vein not attached to the main area of edge. The filtered size could be varied with better knowledge of the likely size

of defects. Filtering was not applied to Canny maps because it removes too much true edge.

The next step is to ensure that edges define the outlines of discrete grains. A sophisticated way to do this is watershed segmentation, in which the edge map is expanded through successive episodes of dilation (addition of pixels adjacent to those already identified as edges) until no grain space is left. The map is then eroded (the dilations are reversed) but with the proviso that no merging of separate grains is allowed. This should mean that grains which were incompletely separated by edge spurs are now fully separated. In principle this is an elegant approach, but for natural river gravels it tends to produce oversegmented results (Butler et al. 2001). The larger the grain, the more chance there is that false edge spurs will occur somewhere on the edge, leading to frequent false segmentation of larger grains. This biases the number of grains detected, though with negligible effect on the derived GSD according to Butler et al. (2001). We found that segmentation errors were minimized by using only a small number of episodes of edge dilation before regrowing the grains. This modification has the effect of setting some maximum distance across which edges can be connected, so that larger grains are not split. Results are presented below from watershed erosion of edges dilated up to eight times. This means unconnected pixels separated by gaps up to 16 pixels (~ 23 mm) wide can be connected, but grains larger than 23 mm are not falsely split. Results deteriorated with more than eight dilations, confirming that complete watershed erosion is not a good option. The possibility of elongated grains splitting preferentially across their a axis remains; it has relatively little effect on GSD estimation from b axis measurements but could be a problem if grain orientation information is required.

Measuring Grain Sizes

The final stage of image processing is feature extraction: measurement of grain size, shape, or orientation from a fully connected edge map. We concentrate on what is probably the most widely used end product, the GSD based on b axis diameters of grains. These b axes are those visible on the planimetric image, and may differ from the true diameters which would be determined after removing grains from the bed. The possible bias introduced is considered later. Butler et al. (2001) showed that manual image measurements of apparent b axis diameters were matched well in an AIA approach which estimated the b axis from the a axis and area of each identified grain, on the assumption of an elliptical shape. Our work confirms this so we used ellipse matching to size each grain, assigned it to a 0.5ψ size class, then summed the total area of grains in each class. This gives results equivalent to the numerical grid Wolman count of McEwan et al. (2000) but taken to the highest possible pixel resolution: each pixel of each grain size class is measured. Experiments using lower-resolution numerical sampling grids indicate the pixel-count method is best.

TESTING THE ALGORITHMS

Grain measurement errors are of two types, which we discuss in turn: those introduced by automatic image processing (errors in grain identification and size measurement), then those generated by bed structuring (errors caused by the use of a photograph as a substitute for manual sieving).

A visual estimate of the fraction area of fines for each image is required for various stages of processing. This estimate and a pixel calibration were taken for each image prior to processing. Any operator error in these preliminaries is unquantified but is expected to be insignificant.

Automatic Versus Manual Image Analysis

The data for this comparison relate to the 24 sites for which images were manually digitized in the manner of a 100-pebble Wolman count truncated at 3ψ (8 mm). The GSDs so obtained are shown in Figure 3 to emphasize how wide a range of conditions was included. Maximum grain sizes vary




Fig. 3.—Cumulative grain-size distributions obtained by manual image analysis for 24 images used to test automated procedures.

between 5.5 and 7.5 ψ (45–180 mm), median diameters between 3.9 and 5.9 ψ (15–60 mm), standard deviations between about 0.5 and 1.0 ψ , and sand area between 0 and 70%.

Rather than compare each field GSD directly with the corresponding GSDs obtained by alternative AIA procedures, which would require a large number of graphs, we compare values of seven percentiles of each GSD. This enables the performance of each procedure for the set of 24 sites to be depicted in one plot, whilst still allowing investigation of any size dependency in the accuracy of the automated procedures. The percentiles chosen are those most commonly used in sedimentological description: 5, 16, 25, 50, 75, 84, and 95. Because of the approximately lognormal shape of the GSDs we analyze ψ values (i.e., \log_2 mm) rather than diameters in millimeters, and refer to them where necessary as P_{50} (etc.) rather than D_{50} .

Edge Connectivity and Edge-Detection Method.—The 24 images were subjected to each of the four edge detection methods outlined above, followed by edge tidying. Each of the 96 incomplete edge maps was then processed eight times, using 1, 2, ..., 8 dilations before watershed segmentation and GSD determination. This allows investigation of the optimum number of dilations (abbreviated to ND henceforth) for each procedure.

The results are summarized in Figure 4. We define errors in the following way, always in ψ units. For each edge-detection method and ND, the mean error b and root-mean-square error r are computed as $b = \sum(P_m - P_a)/n$ and $r^2 = \sum(P_m - P_a)^2/n$, respectively, where the sums are over $n = 168$ differences (7 percentiles \times 24 images) between manual (m) and automated (a) estimates. The irreducible random error e of the estimates is then calculated as $e^2 = r^2 - b^2$. In terms of a plot of P_m against P_a (Figure 4A), b represents the mean offset of points leftwards from the line of equality and e represents the scatter around a line offset by this amount. Thus b quantifies bias and e quantifies precision. We regard e as more useful than r because e shows the irreducible imprecision even after correcting for any bias, whereas r conflates two sources of error which may be conditioned by different factors.

Figures 4B and 4C show how b and e alter with the number of dilations used after each edge-detection procedure. As ND increases the bias switches from positive (Canny and dark threshold) or near-zero (Sobel) to negative (Fig. 4B). This is because the more dilations are used, the lower the chance of missing a true grain boundary during AIA and the greater the likelihood of falsely subdividing a grain. The aggregate method has a nearly constant negative bias from the outset, i.e., a tendency to oversegment, because it combines edges identified in two different ways. The amount of bias is nearly constant for ND from 2 or 3 to 6 in each method, suggesting the possibility of using simple empirical corrections. The random error (Fig. 4C) is also near-constant over this range of ND. In general bias is least

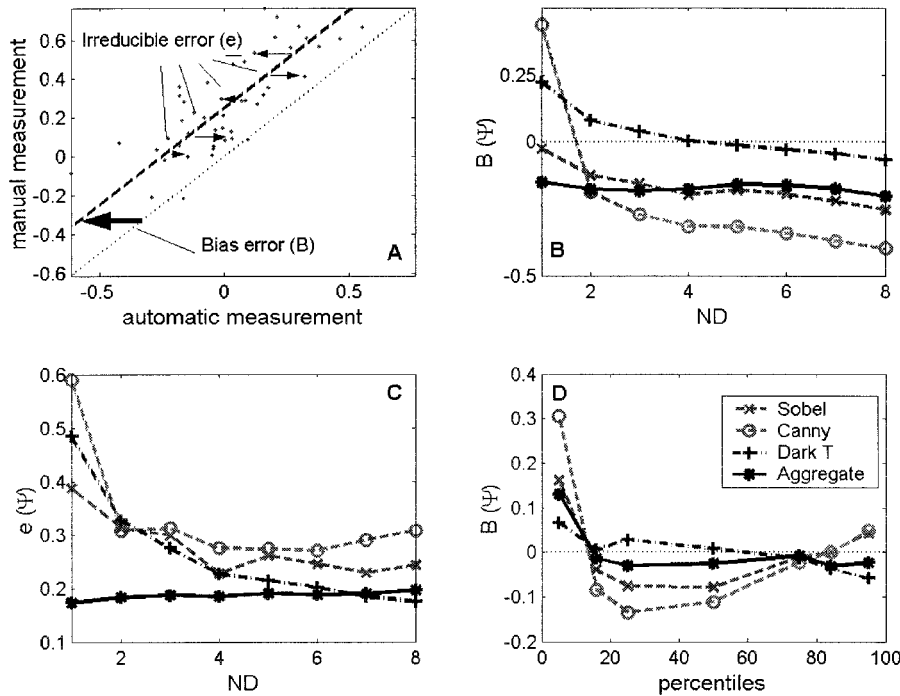


FIG. 4.—Comparison of automated with manual image analysis of grain-size percentiles (ψ) for 24 images. **A)** Schematic plot of manual against automated percentiles to show definitions of bias (b) and irreducible random error (e). **B)** Dependence of bias upon method used for edge detection and number of dilations done before watershed segmentation. **C)** Dependence of random error upon method used for edge detection and number of dilations done before watershed segmentation. **D)** Variation between percentiles in mean bias for different methods of edge detection each with the optimum number of dilations. Where shown, the light dotted line is the line of equality.

with the dark-threshold method, and random error least with this or (at low ND) the aggregate method. The success of the dark-threshold method must depend on the sensitivity of the thresholding to surface condition and lighting, but it appears our extensions to the algorithm increase its robustness.

Figure 5 plots P_m against P_a at the ND giving maximum precision (minimum e). Bias is essentially independent of absolute grain size: the clouds of points are offset only a little from the line of equality and are parallel to it. Both of our preferred methods have bias and random uncertainty of the order of 0.1–0.2 ψ . This is comparable with the uncertainty which Rice and Church (1996) found in manual grid-by-number sampling: their boot-

strap precision values for 100 pebbles at two sites range from ~ 0.1 to 0.3 ψ according to percentile. Calculating mean errors separately for each percentile of interest, averaged across the 24 images, reveals a slight dependence of bias on the *relative* grain size within a GSD (Figure 4D). P_{16} and P_{75} are estimated well but P_5 is always overestimated, and P_{25} and P_{50} are slightly underestimated by each method except dark thresholding. This pattern of error is not sensitive to ND. We have two explanations for it. AIA grain edges are generally somewhat inside the true edges, making typical-sized grains a little too small. Second, and more importantly, AIA tends to find too many fine grains as a result of oversegmentation. The size depen-

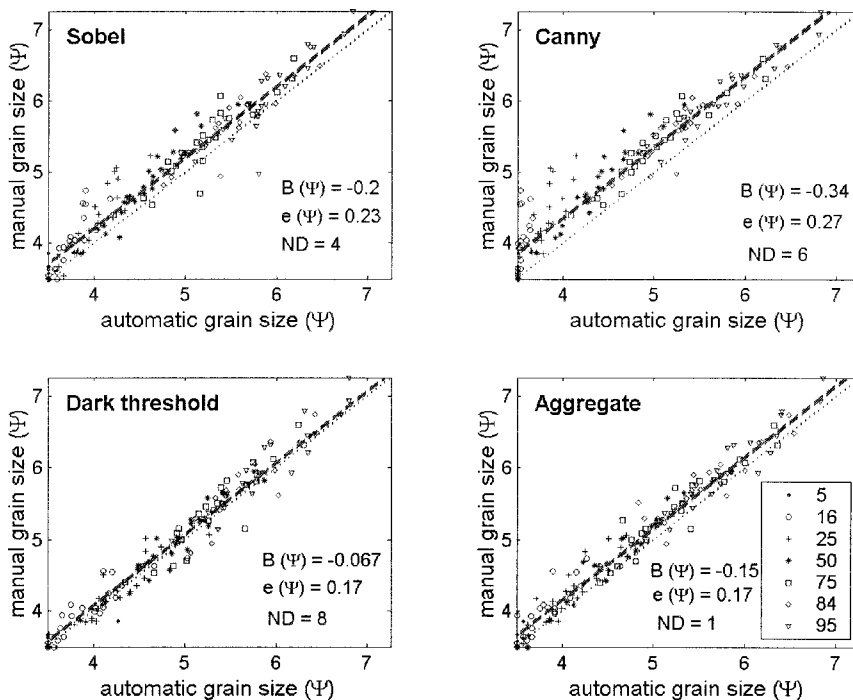


FIG. 5.—Comparison between grain-size percentiles (seven for each of 24 images) estimated by manual (axis digitization) and four methods of automated image analysis. Edge detection was followed by the number of dilations (ND) giving maximum precision (minimum random error e). Light dotted line is line of equality, heavy dashed line is after correcting for overall mean bias (b) in each method.

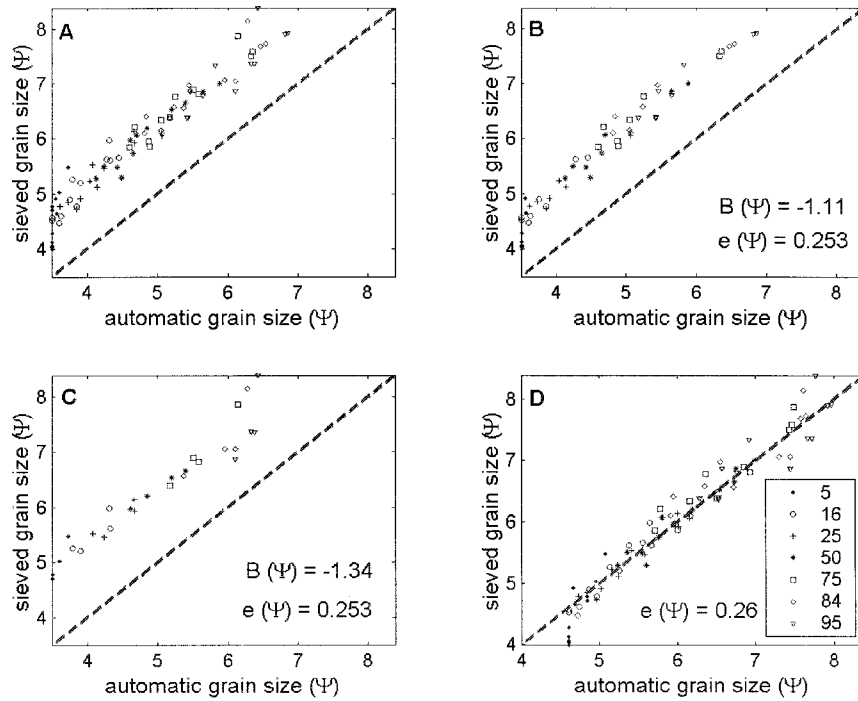


FIG. 6.—Comparison between grain-size percentiles (seven for each of 12 images) obtained from field measurements and by automated image analysis using the aggregate edge-detection method. **A)** All sites and percentiles; **B)** openwork gravels only; **C)** closed-framework sites; **D)** fit after applying relevant bias correction. Line of equality is shown in each case.

dence is particularly pronounced with the Canny method, which again performs less well than the other methods on the relatively noisy images we have studied.

Taking the results in Figures 4 and 5 as a whole, it appears that the dark-threshold method of edge detection, alone or combined with the first-derivative (Sobel) method, gives the most reliable estimates of GSD percentiles for gravel beds. Results are slightly biased compared to manual digitization, but by an amount that is small compared to the error (discussed below) inherent in estimating a volumetric GSD from a 2-D image. A small constant bias correction could be applied using the values shown in Figure 5. If dark thresholding is used on its own, it should be followed by partial watershed segmentation using 6–8 dilations to reduce bias and improve precision, but an aggregate edge map only needs one dilation.

Automatic Image Analysis Versus Field Sieving

It has long been realized (e.g., Adams 1979) that even an accurately measured image GSD will not generally match the GSD obtained by traditional field methods. Image analysis estimates the GSD entirely from the planimetric appearance of the surface whereas field methods involve the third, vertical dimension to greater or lesser extent. In bulk sieving the GSD relates to a 3-D volume, not a 2-D area, and the mass of a grain depends on its 3-D shape rather than planimetric area. A pebble-count GSD does relate, like an image-analysis GSD, to a surface not a volume, but with the difference that grains are removed from the bed for sizing; this gives higher grain diameters than image analysis to the extent that grains overlap each other and/or are tilted. Ibbeken and Schleyer’s (1986) manual analysis of photographs of water-worked gravels showed that effects of non-equiaxial grains and grain positioning cancel out to some extent, but only in the middle of the GSD; systematic bias can be expected in the coarsest percentiles. Also, their work was restricted to open-framework gravels. Partial grain burial will clearly be more likely in the bed surfaces examined here because a large amount of fine material is generally present. GSDs obtained by AIA are therefore likely to be too fine by an unknown value (partly dependent on the percentage of fines present) but to have approximately the right shape apart from the extreme tails.

Our field test data consist of area-by-number GSDs for 12 sites, truncated

at 8 mm and converted to be equivalent to grid-by-number distributions. The counts averaged 429 grains per image, with a minimum of 72 for a surface which was 45% covered by fines. Percentiles of the corrected GSDs are compared in Figure 6A with percentiles obtained by AIA using the aggregate method with one dilation. AIA is seen to underestimate true grain size by rather over 1 ψ on average, but the data plot parallel to the line of equality; the bias is therefore fairly consistent across grain sizes between and within sites. Part of the scatter relates to the distinction between open-framework gravels (Figure 6B) and those with significant areas partly or fully covered by fines (Figure 6C): the mean biases are $b = 1.11$ and 1.34ψ , respectively. When these values are used as correction factors, the data plot close to the line of equality with irreducible scatter $e = 0.25 \psi$ (Figure 6D). The biggest discrepancies are for tail percentiles, as expected.

CONCLUSIONS

Our investigations confirm suggestions by McEwan et al. (2000) and Butler et al. (2001) that automated image analysis is now a viable method for obtaining grain size descriptors of river-bed gravel surfaces. We have extended its application to moist and partially sand-draped gravels under a range of uncontrolled light conditions, from cloud or moderate rainfall to sunshine at various orientations. Our AIA methods, using dark thresholding alone or aggregated with a first-derivative map, appear sufficient to give accurate grain size results so long as edge finding is based on our edge-seeded image-porosity concept and edge detection is followed by incomplete watershed segmentation. These innovations, and the use of local rather than global thresholds, ensure that the AIA methods are robust when applied to “difficult” images. After initial calibration of each image for pixel size and fines content (~ 30 seconds per image), all our results were generated automatically. This opens up greater possibilities than previously for detailed investigation of spatially variable gravel surfaces, because large numbers of images can be processed for grain-size distributions (or other grain information) in the time to take the photographs plus about an hour of operator time per 100 images. The MATLAB code written by the first author to perform the analysis is documented and available on the second author’s web pages (http://www.shef.ac.uk/geography/staff/ferguson_rob.html).

Our results must be qualified in two respects. First, natural river gravels vary greatly in surface character (in such respects as nature of clast interlocking, amount and nature of partial burial by sand, and within-patch variation in grain size, shape, and color) so it is unlikely that the correction factors in Figure 6 would be optimal elsewhere. More research is needed to determine how surface and lighting conditions affect results. Experiments using simulated grain arrangements might also help isolate effects of orientation and structuring. Second, the applicability of our methods to underwater areas remains to be tested. Obtaining good results from underwater images may be problematic because of degradation of the textural information. High-resolution elevation imagery may prove preferable to radiometric sensing, particularly in the laboratory (McEwan et al. 2000). However, if appropriate submerged images can be obtained by whatever means, the nondestructive nature of automated image analysis makes it ideal for studying bed evolution in either flume or field.

ACKNOWLEDGMENTS

The algorithm development and application was done by the first author while supported by a Ph.D. studentship from the Department of Geography, Sheffield University. Many thanks also to Mike Church for providing field equipment, accommodation, and helpful advice; to Erica Ellis for field assistance; and to Lynne Frostick and Stuart Lane for helpful review comments.

REFERENCES

- ADAMS, J., 1979, Gravel size analysis from photographs: American Society of Civil Engineers, Proceedings, Journal of the Hydraulics Division, v. 104, p. 1247–1255.
- BUTLER J.B., LANE S.N., AND CHANDLER J.H., 2001, Automatic extraction of grain-size data from gravel surfaces using digital image processing: American Society of Civil Engineers, Journal of the Hydraulic Division, v. 39, p. 519–529.
- CHURCH M.A., McLEAN D.G., AND WOLCOTT, J.F., 1987, River bed gravels: sampling and analysis, in Thorne, C.R., Bathurst, J.C., and Hey, R.D., eds., *Sediment Transport in Gravel-Bed Rivers*: Chichester, U.K., Wiley, p. 43–79.
- FERGUSON, R.I., CHURCH, M., AND WEATHERLY, H., 2001, Fluvial aggradation in Vedder River, British Columbia: testing a one-dimensional sedimentation model: *Water Resources Research*, v. 37, p. 3331–3347.
- GONZALEZ, R.C., AND WOODS, R.E., 1992, *Digital Image Processing*: Reading, Massachusetts, Addison-Wesley, 716 p.
- IBBEKEN, H., AND SCHLEYER, R., 1984, Photo-sieving: a method for grain-size analysis of coarse-grained, unconsolidated bedding surfaces: *Earth Surface Processes and Landforms*, v. 11, p. 59–77.
- McEWAN, I.K., SHEEN T.M., CUNNINGHAM, G.J., AND ALLEN, A.R., 2000, Estimating the size composition of sediment surfaces through image analysis: *Proceedings of the Institute of Civil Engineers: Water, Maritime and Energy*, v. 142, p. 189–195.
- MARTIN, Y., AND CHURCH, M., 1995, Bed-material transport estimated from channel surveys: Vedder River, British Columbia: *Earth Surface Processes and Landforms*, v. 20, p. 347–361.
- LANE, S.N., 2001, The measurement of gravel-bed river morphology, in Mosley, M.P., ed., *Gravel-bed Rivers V*: Christchurch: Caxton Press, p. 291–338.
- RICE, S., AND CHURCH, M., 1996, Sampling surficial fluvial gravels: the precision of size distribution percentile estimates: *Journal of Sedimentary Research*, v. 66, p. 654–665.

Received 13 May 2002; accepted 21 November 2002.