SediNet: A configurable deep learning model for mixed qualitative and quantitative optical granulometry

Daniel Buscombe

Geosciences Division, School of Earth and Sustainability, Northern Arizona University

Abstract

I describe a configurable machine-learning framework to estimate a suite of continuous and categorical sedimentological properties from photographic imagery of sediment, and to exemplify how machine learning can be a powerful and flexible tool for automated quantitative and qualitative measurements from remotely sensed imagery. The model is tested on a large dataset consisting of 400 images and associated detailed label data. The data are from a much wider sedimentological spectrum than previous optical granulometry studies, consisting of both well- and poorly sorted sediment, terrigenous, carbonate, and volcaniclastic sands and gravels and their mixtures, and grain sizes spanning over two orders of magnitude. I demonstrate the model framework by configuring it in several ways, to estimate two categories (describing grain shape and population, respectively) and nine numeric grain-size percentiles in pixels from a single input image. Grain size is then recovered using the physical size of a pixel. Finally, I demonstrate that the model can be configured and trained to estimate equivalent sieve diameters directly from image features, without the need for area-to-mass conversion formulas and without even knowing the scale of one pixel. Thus, it is the only optical granulometry method proposed to date that does not necessarily require image scaling. The flexibility of the model framework should facilitate numerous application in the spatio-temporal monitoring of the grain size distribution, shape, mineralogy and other quantities of interest, of sedimentary deposits as they evolve as well as other texture-based proxies extracted from remotely sensed imagery.

1. Introduction

Sediment grain size fundamentally influences the physics of flows of water, wind, ice and sediment that continually shape landforms. Large sedimentological datasets have led to important discoveries in dynamic environments such as contemporary river beds, sea beds and aeolian sediment surfaces that are constantly changing under fluid power, for example in sediment transport (e.g. *Masteller & Finnegan*, 2017; *Rubin et al*., 2019), channel bed mobility (e.g. *Montgomery et al*., 1999), channel geometry (e.g. *Pfeiffer et al*., 2017), sediment provenance (e.g. *Paterson & Heslop*, 2015), sediment abrasion (e.g. *Novak-Szabo et al*., 2018), hydraulic resistance (e.g. *Rickenmann & Recking*, 2011), particle settling (e.g. *Sternberg et al*., 1999) and dispersal at coasts (e.g. *Wheatcroft & Borgeld*, 2000), and beach dynamics (e.g. *Bergillos et al.*, 2016). Traditionally, the means of acquiring large grain size (or shape, or any other metric) data sets has been laborious and time-consuming through laboratory analyses of samples taken in the field. Optical granulometry is the measurement of sediment from statistical analysis of image intensity and texture, and has been driven by instrumental (e.g. *Buscombe* et al., 2014; *Carbonneau et al*., 2018; *Rubin et al*., 2007; *Woodget et al.*, 2018) and analytical (e.g. *Black* et al., 2014; *Buscombe* et al., 2010; *Buscombe and Rubin*, 2012b; *Buscombe*, 2013; *Cheng and Liu*, 2015; *Carbonneau et al*., 2005a, 2005b; *Carbonneau et al*., 2004; *Cuttler et al*., 2017; *Dugdale* et al., 2010; *Legleiter* et al., 2016; *Rubin*, 2004; *Woodget et al*., 2017) developments over the past 15 years. Another set of deterministic methods known as `photosieving' (e.g. *Adams*, 1979) or object-based image analysis or OBIA (*Carbonneau et al*., 2018) have been developed (e.g. *Detert and Weitbrecht*, 2012; *Graham et al*., 2005) that aim to identify each individual grain and cannot therefore be used on grains smaller than one pixel (subpixel) which is not a theoretical limitation of optical granulometry techniques that statistically quantify image texture

(*Carbonneau et al*., 2004). One major goal of this corpus of work is to develop a reliable suite of techniques for spatio-temporal monitoring of the grain size of sedimentary deposits as they evolve, remotely and automatically. This has the potential to significantly alter the way geomorphological research is carried out (e.g. *Viles*, 2016) and may hopefully lead to significant discoveries in the two-way feedbacks between evolving sedimentary landform morphologies and the spatio-temporal dynamics of grain size, or 'morpho-sedimentary dynamics' (cf. *Buscombe and Masselink*, 2006), at large field scales. This will require measuring grain size at the same spatial (e.g. *Rubin et al.,* 2019) and temporal (e.g. *Buscombe et al*., 2014) coverage as is now possible with topographic measurements that can capture the spatio-temporal evolution of small-scale morphologies (e.g. *Austin et al*., 2007*; Nield et al*., 2011; *Turner et al.,* 2008*; Williams et al*., 2014).

The present study is motivated by five observations. First, the wavelet-based optical granulometry method of *Buscombe* (2013), while accurate for relatively well-sorted sediment (e.g. *Masteller and Finnegan*, 2017; *Michaelides et al*., 2018; *Prodger et al*., 2016; *Smith et al*., 2018), can be inaccurate for images of grains that are poorly sorted such as sand and gravel mixtures, or where there are relatively few individual grains in the image (hundreds to thousands of grains are typically required). For this study, I have collated a dataset of more than 100 images of sediment that mostly fall under these two categories, to augment the 300-image dataset used by *Buscombe* (2013) that contained a greater proportion of relatively well-sorted sediment, in order to develop a more generally applicable method. Some images contain as few as 10 individual grains, whereas others depict millions of individual grains.

Second, optical granulometry methods quantify the size of apparent axes of grains in the image plane, where many grains may be overlapping. If a bulk (i.e. by mass or by volume)

sample size distribution is the information required, the *Buscombe* (2013) or similar method can provide comparable grain size distributions to those derived using sieves or similar methods usually only if the appropriate conversion of area- to mass-by-size is made, which takes the form (*Diplas and Sutherland*, 1988; *Kellerhals and Bray*, 1971):

$$p(V-W)_i = \frac{p(A)_i D_i^x}{\sum p(A)_i D_i^x} \quad (1)$$

where $p(V-W)_i$ is the volume by weight proportion of the *i*th size fraction, $p(A)_i$is the image-derived areal proportion of the *i*th size fraction, D_i is the grain size of the *i*th size fraction and *x* is a conversion constant. See also *Graham et al*., (2012) for field applications of this conversion. *Diplas and Fripp* (1992) suggest that it is necessary to use different values for exponent *x* depending on grain size, but *Diplas et al*. (2008) suggest a pragmatic approach is to use an average value for *x*, which is determined empirically for each population of grains imaged. *Cuttler et al*. (2017) confirmed that *x* must be determined empirically for bioclastic carbonate sediment to avoid over-predicting sieve sizes and sediment settling velocities from parametric formulas, even though the *Buscombe* (2013) method worked well to estimate the apparent axes of grains from the imagery. Here, I demonstrate that machine learning can be used to map image features to sieve sizes directly, without the need for conversion formulas and without even knowing the scale of a pixel.

The third motivation for this study is provided by *Shojiet et al*. (2018) who demonstrated the utility of deep learning techniques to classify volcanic ash particles by shape, and specifically that a well-designed deep convolutional neural network (CNN) can automatically extract the relevant features from imagery of particles to estimate a categorical quantity. Here, that work is extended by demonstrating that the same CNN architecture can be used for both discrete (classification) and continuously varying quantities (regression) from a single image, by estimating categorical particle shape and population, and numerical percentiles of the grain

size distribution. CNNs are a type of artificial neural network (ANN) and part of a class of machine learning techniques called deep learning (*Goodfellow et al*., 2016) that have recently been shown to perform well for both classification and regression tasks equally, including in numerous geosciences applications where relevant image features are extracted automatically (e.g. *Buscombe and Carini*, 2019; *Buscombe et al*., 2019; *Buscombe and Ritchie*, 2018; *Linville et al*. 2019; *Luo et al*., 2018; *Jiang et al*., 2018; *Reichstein et al*., 2019). The basic premise of applications such as these, compared to those of other machine learning subcategories, is that it circumvents the need (and the effort required) to make decisions about what extracted image features are important to a specific task, which tends to make the models both more subjective and more powerful.

The fourth motivation is that predictive modeling techniques for both categorical and numerical output quantities in the geosciences is somewhat rare. Categorical variables are those that are ascribed an integer, but where the values themselves do not have a physical meaning as they simply enumerate the possible realizations of a phenomenon. As such, they are limited by our ability to identify and ascribe meaning to the phenomenon, and also as intra-categorical variation approaches inter-categorical variation. However, for trivial, well-known or unambiguously defined quantities, they are an essential part of the geosciences, but whereas some techniques are designed for handling continuous estimates, others are better for handling categorical or discrete variables. This typically requires the development of transforms that convert continuous to categorical (using discretization, dummy variables, etc.), which can be subjective if thresholds or discrete bins need to be defined. Here I describe a single empirical framework that can be trained to predict both categorical and continuous quantities, as needed, which might be useful in other geophysical contexts. Within the framework of an ANN, this is relatively straightforward: essentially, multinomial logistic

regression is used for image features that have been distilled by a CNN to estimate discrete variables (such as categorical grain shape), and linear regression for continuous variables (such as grain size). For the latter, the key to the framework is to provide the image features that scale linearly with the response variable (e.g. grain size) being estimated. Highlighting this relatively simple principle through demonstration is worthwhile if it motivates similar progress in other geophysical contexts.

The final and perhaps foremost motivation to developing yet another optical granulometry technique is the observation that the data-hungry nature of machine learning allows for collaborative tool development for extracting scientific information from images of sediment. Recognizing the variety of both sediment imagery, due to the inherent variability of natural sediment, and potential SediNet applications, the motivating idea behind the creation of the SediNet model and software (*SediNet online software*, 2019) is to foster the creation of such a community. Users can contribute imagery, models, and retrain existing models, as well as using existing SediNet models contained in the repository.

2. Data

The model is trained and tested on a large data set consisting of 400 labeled images of sediment (Figures 1 and 2), with a large variation in the spatial footprint (field-of-view) of each image, the spatial resolution (physical size of a pixel), and variation in camera sensor. The data are from a wide sedimentological spectrum of well and poorly sorted sediment, consisting of terrigenous (derived by erosion of crystalline, volcanic, and sedimentary rocks), carbonate (skeletal grains, oolites, and some locally derived detrital carbonate), and volcaniclastic (lapilli, glass, and pyroclastic bombs) sands and gravels and their mixtures, and grain sizes in pixels spanning over two orders of magnitude. Out of the 400 images, 300

were compiled and used by *Buscombe* (2013) to develop a wavelet-based algorithm for estimating grain size from imagery (sets A and C in that paper). The remaining 100 samples were compiled for this study, from various fieldwork activities over more than 10 years in various coastal and riverine environments on several continents. The additional 100 samples were chosen specifically to better represent within the dataset both poorly sorted mixed sand and gravel sediment and (usually microscopic) imagery with relatively few numbers of grains.

2.1. Grain Size

The size distribution of intermediate axes of apparent (surface) grains was compiled for each image following the on-screen manual method of *Barnard et al*. (2007), which is the only way in which to reliably obtain a comparable grain-size distribution to that provided by image-based methods (*Baptista et al*., 2012; *Buscombe et al*., 2010; *Cuttler et al*., 2017). However, it is a time-consuming and meticulous process, usually taking a trained operator 30-60 minutes per image to measure the axes of up to 500 grains. Nine commonly utilized percentiles of the cumulative size distribution (namely 5, 10, 16, 25, 50, 75, 84, 90, and 95th percentiles) were calculated for each measured size distribution.

2.2. Grain shape and population

The expanded dataset of 400 images contain a number of sediment populations (Figure 1) that I manually grouped into six categories: 1) well-sorted gravel; 2) well-sorted sand and shell hash from underwater camera (described in *Buscombe et al*., 2014); 3) relatively poorly sorted gravel and sand-gravel mixtures (including imagery from *Warrick et al*., 2009); 4) well-sorted sand; 5) miscellaneous terrigenous and volcaniclastic grains; and 6) miscellaneous bioclastic (carbonate) grains. Additionally, each of the 400 images were classified into four

shape/size categories (Figure 2), namely 1) large well-rounded grains; 2) small well-rounded grains; 3) large angular grains; and 4) small angular grains.

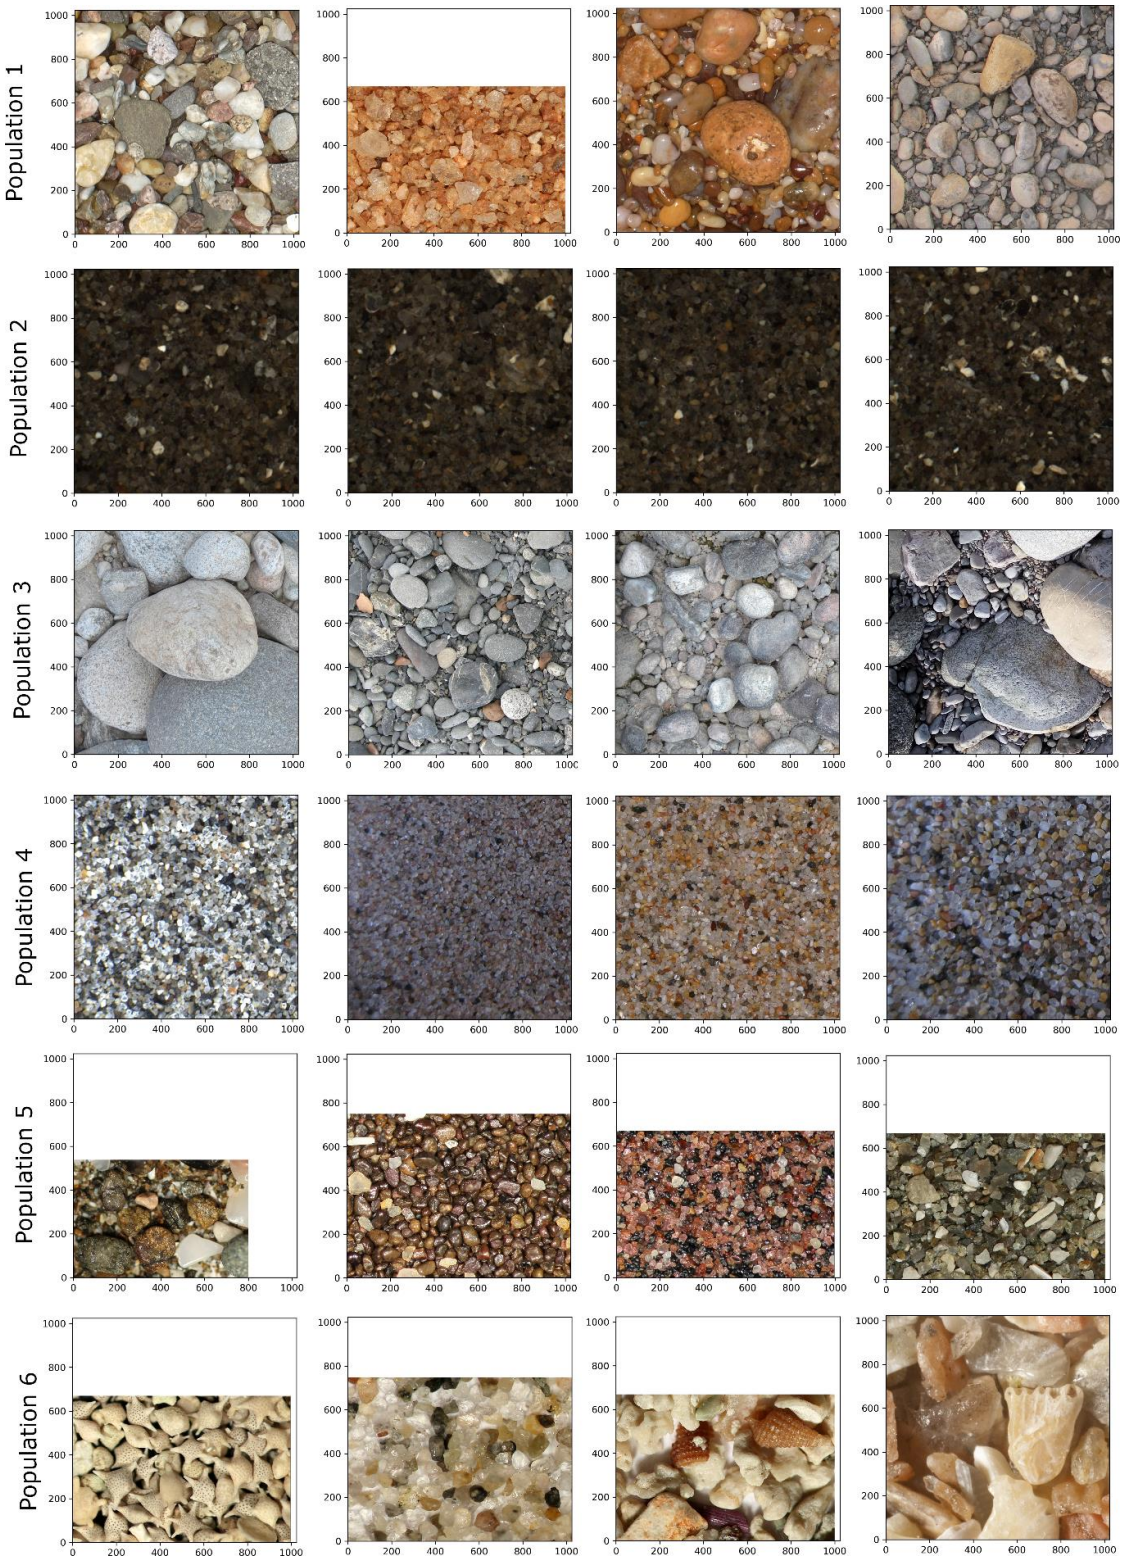

Figure 1: Four example 1024 x 1024 pixel subsets of images from each of six population categories. From top to bottom: 1) well-sorted gravel; 2) well-sorted sand and shell hash from underwater camera (described in Buscombe et al., 2014); 3) relatively poorly sorted gravel and sand-gravel mixtures (including imagery from Warrick et al., 2009); 4) well-sorted sand; 5) miscellaneous terrigenous and volcaniclastic grains; and 6) miscellaneous bioclastic (carbonate) grains.

Figure 2. Four example 1024 x 1024 pixel subsets of images from each of four shape categories. From top to bottom: 1) Large well rounded grains; 2) small well-rounded grains; 3) large angular grains; and 4) small angular grains.

3. SediNet model

Deep learning models have multiple processing layers (called convolutional layers or blocks) and nonlinear transformations (that include batch normalization, activation, and dropout, which are explained below), with the outputs from each layer passed as inputs to the next. SediNet (Figure 3) is a supervised deep neural network model framework that can be used as presented in this paper, or alternatively configured for custom purposes, by training on

any number of input images for any number of numeric or categorical outputs. For the purposes of demonstrating the model in this paper, several SediNet models were made:

1. To estimate nine percentiles of the cumulative grain size distribution in pixels, trained on 204 images and tested on 205 images, both drawn randomly. The train and test sets consist of images of several populations of grains from a wide sedimentological spectrum
2. To estimate nine percentiles of the cumulative grain size distribution in pixels, trained on 15 images and tested on 16 images of one population (beach sands)
3. To estimate sieve size in microns directly, without first estimating the pixel size, trained on the same 15 images and tested on 16 images as above
4. To estimate six categorical populations of grains, trained on 204 images and tested on 205 images, both drawn randomly
5. To estimate four categorical grain shape/size classes, trained on 204 images and tested on 205 images, both drawn randomly

Given the set of *n* images, let us denote one sample $X_\mu \in \mathbb{R}^p$ with $\mu = 1 \ldots n$, where *p* is the number of pixels. For each sample, X_μ there is label $y_\mu \in \mathbb{R}^q$ where *q* is the number of combined categorical and continuously distributed classes. Using the deep learning architecture described below, and the training data set $\{X_\mu , y_\mu\}$ consisting of 50 % of the total number of images, randomly selected, a function *f* is found such that $\hat{y} = f(X)$, where $\hat{y}$ is the predicted set of labels/metrics from sample image *X*. The remaining 50 % of the total data set was used as a test set to evaluate model performance.

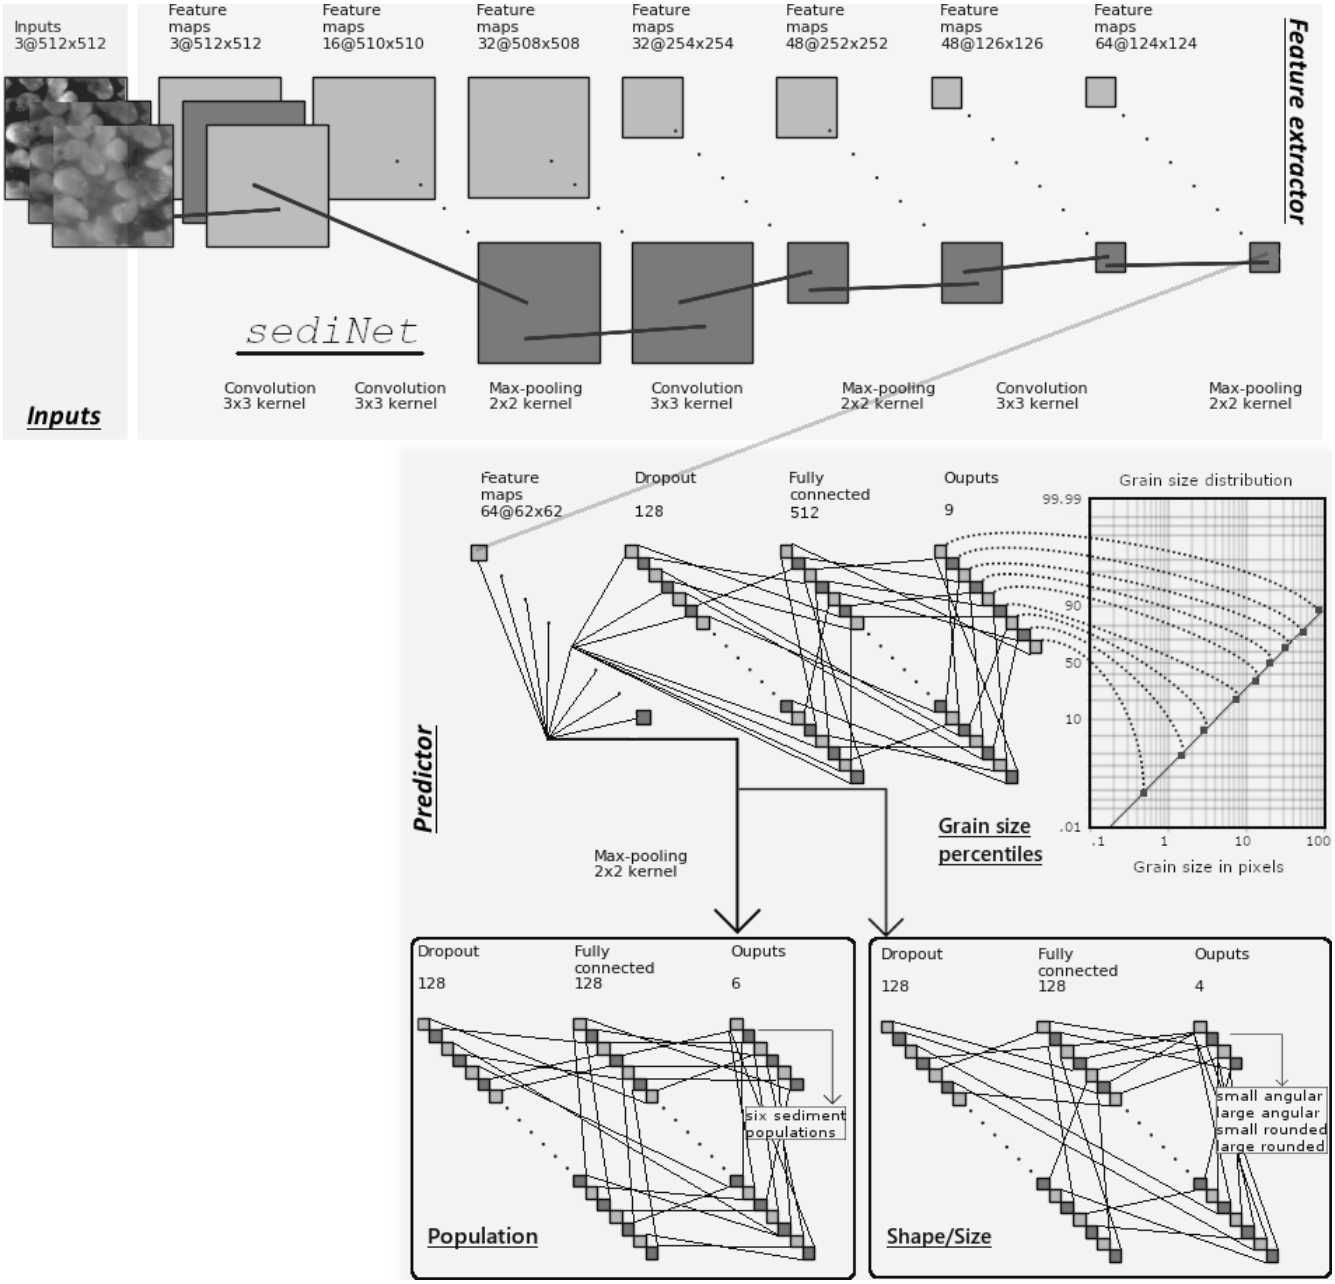

Figure 3. Schematic of the SediNet architecture, as applied to estimating the grain size distribution, and categorical population and shape/size. An input image is passed to the feature extractor consisting of a series of convolutional blocks. The last set of feature maps, which is the result of the last 2D max global pooling layer, is fed into one of three multi-layer perceptrons; one each for the task of estimating grain size percentiles, sediment population, and grain shape.

The image feature extractor consists of four convolutional blocks each consisting of a several two-dimensional convolutional filter layers, batch normalization layers, and two-dimensional max pooling layers (Figure 3). Batch normalization applies a transformation that maintains the mean neuron activation of zero and the activation standard deviation of one (*Ioffe and Szegedy*, 2015). Pooling layers are used to reduce the spatial dimensions of each of the

three-dimensional tensors associated with each pixel of the input image, from *h* x *w* x *d* to 1 · 1 x *d*, by averaging over *h* and *w*. This has the effect of reducing the total number of parameters in the model, thereby minimizing overfitting. The output of the last block is the input of the next. The number of filters increases for each of the four blocks, from 16 in the first block, 32 in the second, 48 in the third and finally 64 in the last block. After the last convolutional block, there is one more batch normalization and two-dimensional max pooling layer, and a dropout layer that randomly drops half the neurons (*Srivastava et al*., 2014). Batch normalization, max pooling, and dropout layers are techniques to prevent overfitting the model (i.e., memorizing the training data rather than learning a general trend). The extracted feature is fed into a series of multilayer perceptrons, one for each estimated quantity, that each culminates in a dense predicting layer with linear regression (known in machine learning literature as a linear activation function) for continuous variable prediction variables (such as grain size in pixels, or sieve size directly), or multinomial logistic regression (in machine learning parlance, a softmax activation function) for categorical variables such as grain shape and population.

The model was retrained 'end-to-end', which means it was initialized with random numbers for neuron weights $w \in \mathbb{R}^k$, then during training the value of those parameters was optimized by minimizing the discrepancy between known and estimated quantities by minimizing a loss function $L\left[f_w(X_\mu, y_\mu)\right]$ for each sample μ where f_w denotes weighted function. By doing so, the model simultaneously and automatically learns feature representations from imagery and a mapping from those features to the target values (e.g. grain size) or classes (e.g. grain shape). Models are trained over several epochs. One training epoch means that the learning algorithm has made one pass through the training dataset, where examples were separated into randomly selected batches of images. The number of training steps per epoch was

computed as the number of training images divided by the batch size. In this study, the batch size was set to eight and results were not sensitive to its value. Upon each step, the gradients of the network are updated and new weights assigned to each neuron. Stochastic gradient descent was used to iteratively adjusting the weights in the direction of the gradient of the average of the loss over the training set using $w^{t+1} = w^t - \lambda \nabla_w R(f_w)$, where t is iteration number (step within an epoch) and λ is the so-called 'learning rate', and where $R(f_w) = \sum L/n$ for the full training data is replaced by the contribution of just a few of the samples.

During model training, each *h* x *w* x 3 pixel input image was resized to 512 x 512 x 3 pixels for computational efficiency. With sufficient computing power, larger images and larger numbers of images could be used. That the image's aspect ratio is typically not preserved does not affect model performance (I revisit this point in the Discussion). The method was implemented in python 3.7 using the Tensorflow (*Abadi et al*., 2015) backend to the keras (*Chollet et al.*, 2015) module, on a GeForce RTX 2080Ti GPU with 11 GB of memory. The resolution of a given grain size estimate in pixels is approximately 2 pixels, determined as the range of that variable in the training data (in the present case, the largest grain size minus the smallest, which is approximately 1000 pixels) divided by the number of neurons in the final dense layer, which was set to 512 (Figure 3). Training utilized the popular Adam algorithm (*Kingma and Ba*, 2014) for stochastic optimization, with parameters β_1= 0.9 and β_1= 0.999 (*Buscombe et al*., 2019). During training, λ was automatically reduced when the loss function stabilized, i.e. when its value stopped decreasing, by a factor of 0.8 after 15 epochs had elapsed with no improvement (*Buscombe et al*., 2019). A lower bound on λ was set at 0.0001. The maximum number of training epochs was set to 100. Models stopped training early (i.e. before 100 epochs) if the validation loss failed to improve for 20 consecutive

epochs. Models typically trained for between 40 and 100 epochs before the criterion was met to stop training early.

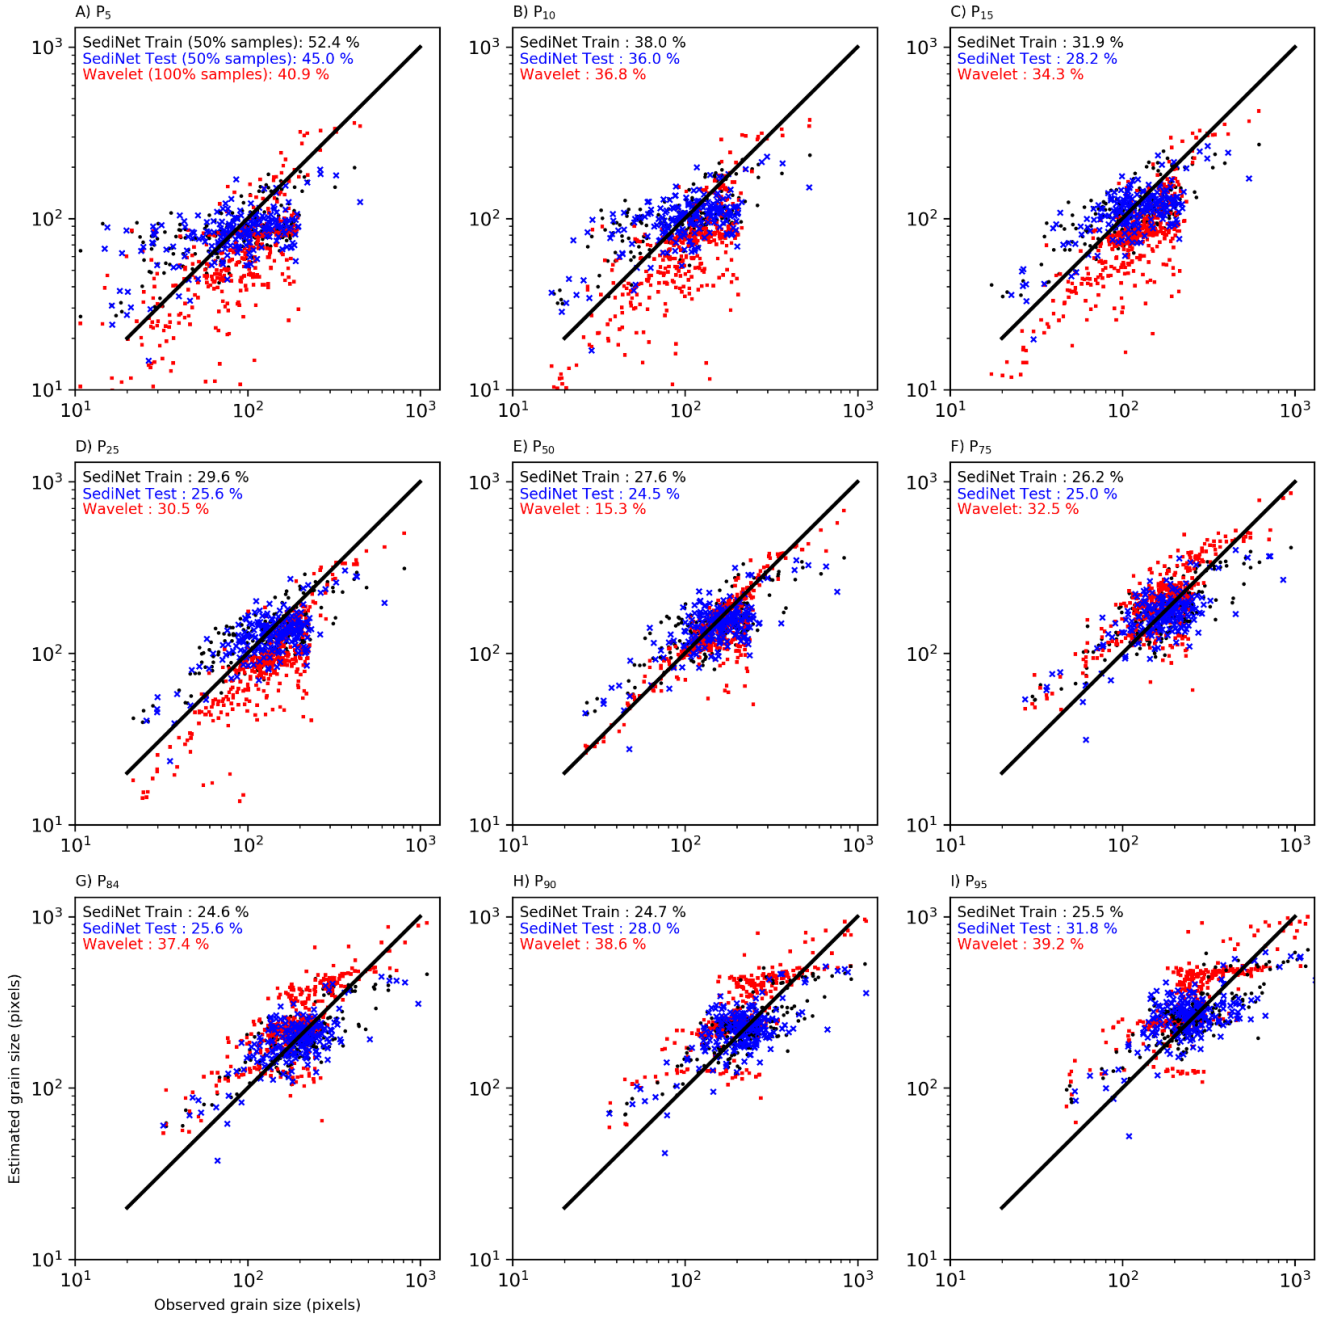

Figure 4. Observed versus estimated grain size percentiles in pixels, for all 409 images. Black dots are the estimate from the training image set (204 samples). Blue crosses are the estimates from the remaining 205 test images. Red dots are all 409 samples analyzed using the wavelet method of Buscombe (2013).

4. Results

4.1. Grain Size

The first implementation of SediNet estimated nine percentiles of the cumulative grain size distribution in pixels, trained on 204 images with mean error between 24 and 52% depending on percentile, and tested on 205 images with mean error between 24 and 45% again varying with percentile (Figures 4 and 5). Mean percent error for each percentile is computed as 100 times the root-mean-squared error normalized by the mean grain size associated with that percentile. Overall, this SediNet model out-performed the wavelet technique of *Buscombe* (2013) and required fewer tunable parameters.

The second implementation of SediNet was for estimating nine percentiles of the cumulative grain size distribution in pixels for a smaller population of sediment images from a given environment (Figure 6). I chose a set of 31 images of sieved beach sand, separated into 16 test and 15 training images. Mean error on the training set was between 7 and 29%, and between 16 and 29% for the test set (Figure 6, A – I). The third SediNet implementation estimated sieve size directly from the same imagery without first estimating the grain size in pixels. Therefore, it implicitly learned the actual size of an image pixel. This model tended to slightly underestimate grain size, with train and test mean errors of 29 and 22%, respectively. The slight bias in the prediction might be corrected empirically, such as by means of parameter x in equation (1), or through further refinement of the model architecture or training procedure. In all three SediNet grain size models, the mean errors for test and train datasets were similar, strongly indicating that the model has generalized well to the data and has not overfit the training data.

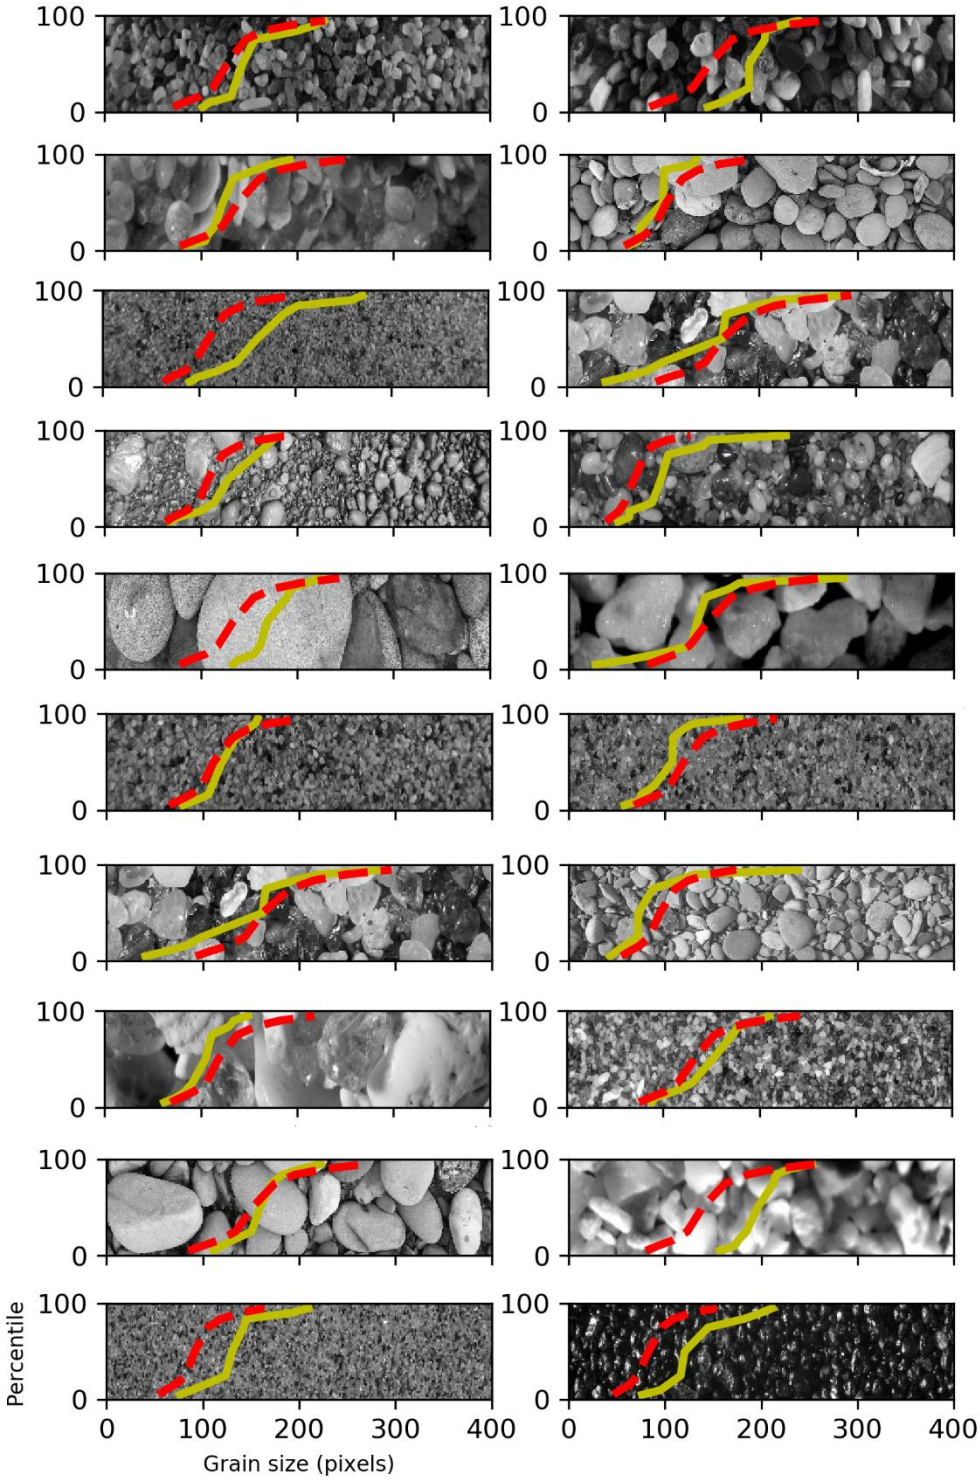

306

307 *Figure 5. Example true (solid yellow line) and estimated (dashed red line) cumulative distributions for 20 randomly selected images,*
308 *small subsets of which are shown in the background of each subplot.*

309

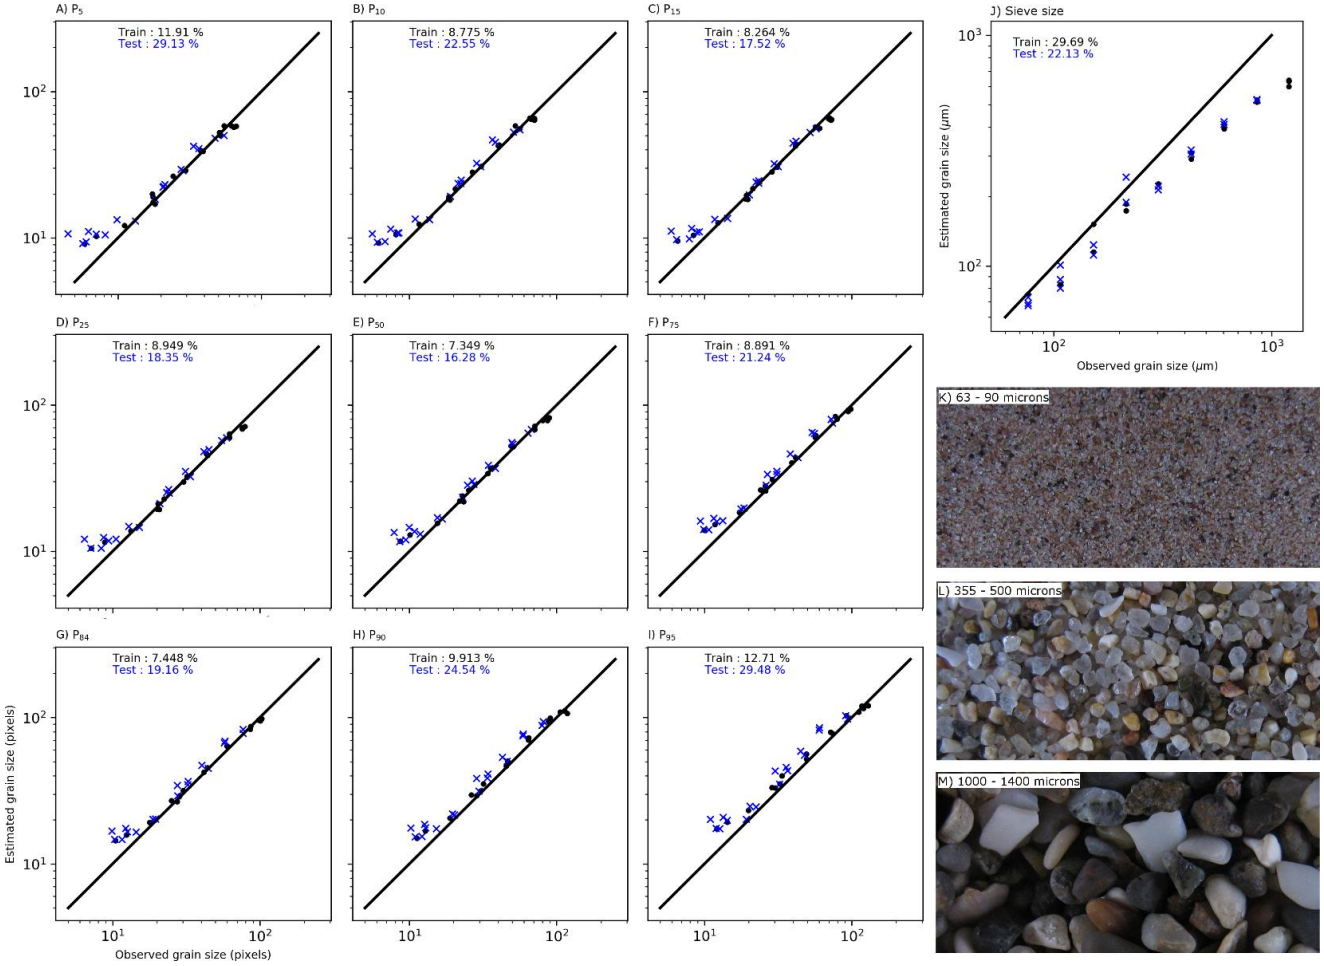

Figure 6. Analysis of one sediment population, consisting of 31 images of sieved beach sands from samples taken at Pescadero in California (images courtesy of David Rubin). A – I) Observed versus estimated grain size percentiles in pixels where black dots are the estimate from the training image set (15 samples) and blue crosses are the estimates from the remaining 16 test images.; J) observed versus estimated mid-sieve size, obtained directly from the image without knowledge of the pixel size; and K – M) example images of three sieve fractions.

4.2. Grain shape and population

The fourth implementation of SediNet estimated six categorical populations of sediment, trained on 204 images and tested on 205 images, both drawn randomly. Classification skill was evaluated using a 'confusion matrix' of normalised correspondences between true and estimated labels (Figure 7, A - C). A perfect correspondence between true and estimated labels is scored 1.0 along the diagonal elements of the matrix. Random misclassifications are readily identified as off-diagonal elements with relatively small magnitudes, and systematic misclassifications are recognized as off-diagonal elements with relatively large magnitudes.

The three confusion matrices for categorical sediment population shown in Figure 7, A – C show skill for, respectively, training, testing and combined (i.e. all 400 images) data. The model overfits population 2 (underwater images of continental shelf sand, Figure 1), evidenced by the large discrepancy between training skill (1.0) and test skill (0.62; Figure 7A, B). However, overfitting is not evident for the other five classes, with test scores being approximately equal to training scores. All classes are classified with accuracies of > 70% for the combined model (Figure 7C).

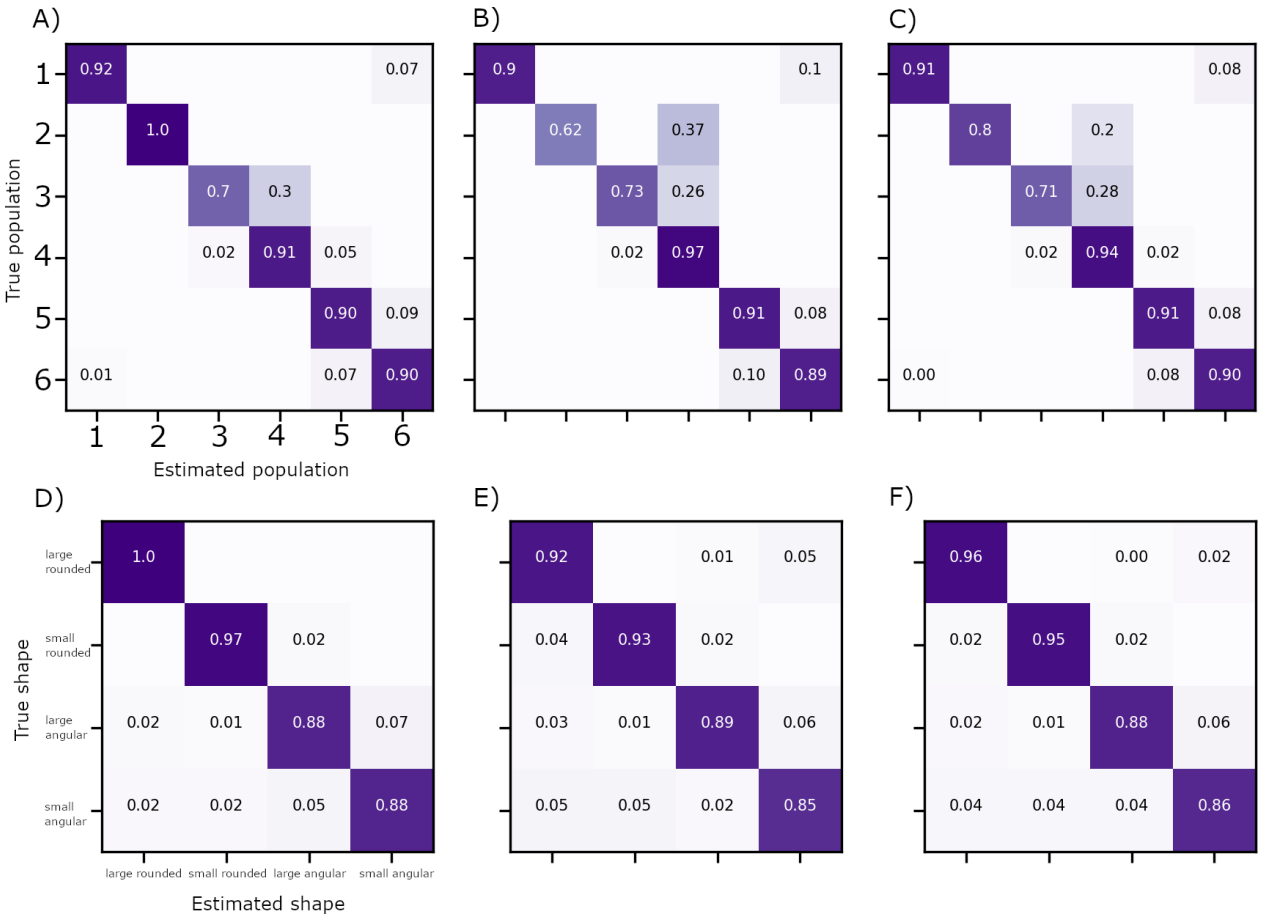

Figure 7. Confusion matrices for (A – C) categorical population and (D – F) categorical shape. Subplots A and B show training and testing datasets

The fifth and final SediNet implementation reported here was configured to estimate four categorical grain shape/size classes, trained on 204 images and tested on 205 images, both drawn randomly. The three confusion matrices for categorical sediment shape shown in

Figure 7D – F show skill for, respectively, training, testing and combined (i.e. all 400 images) data. The similarity in train and test scores for all four classes demonstrates the model has not overfit the data. All classes are classified with accuracies of > 85% for the test, train and combined models (Figure 7D - F).

5. Discussion

The task of quantifying and classifying natural objects and textures in images of sedimentary landforms is increasingly widespread in a wide variety of geomorphological research (*Franklin and Mulder*, 2002; *Mulder et al*., 2011; *Smith and Pain*, 2009), especially as imagery collection using UAVs becomes more prevalent (*Carbonneau et al*., 2018; *Gomez and Purdie*, 2016; *Turner et al*., 2016). The automated method to size and classify sediment described here could maximize speed and objectivity of sedimentary description at large scales, and might be applied to the analysis of datasets consisting of tens to millions of individual images. The model framework could enable spatio-temporal monitoring of grain size more efficiently, being configurable to estimate many custom-defined quantities and qualities for specific tasks. Given it is a data-driven approach, models trained for use in specific environments will highly likely be as or more accurate than methods such as *Buscombe* (2013) and *Carbonneau et al*., (2004) that are based on signal processing or random field theory, especially for poorly sorted sediment, small field-of-view, and large grain size compared to field-of-view (small numbers of individual grains). This is because those methods are not informed by data (i.e. only tested with data); therefore, the massive variation in natural sediment can only be a limitation in their application.

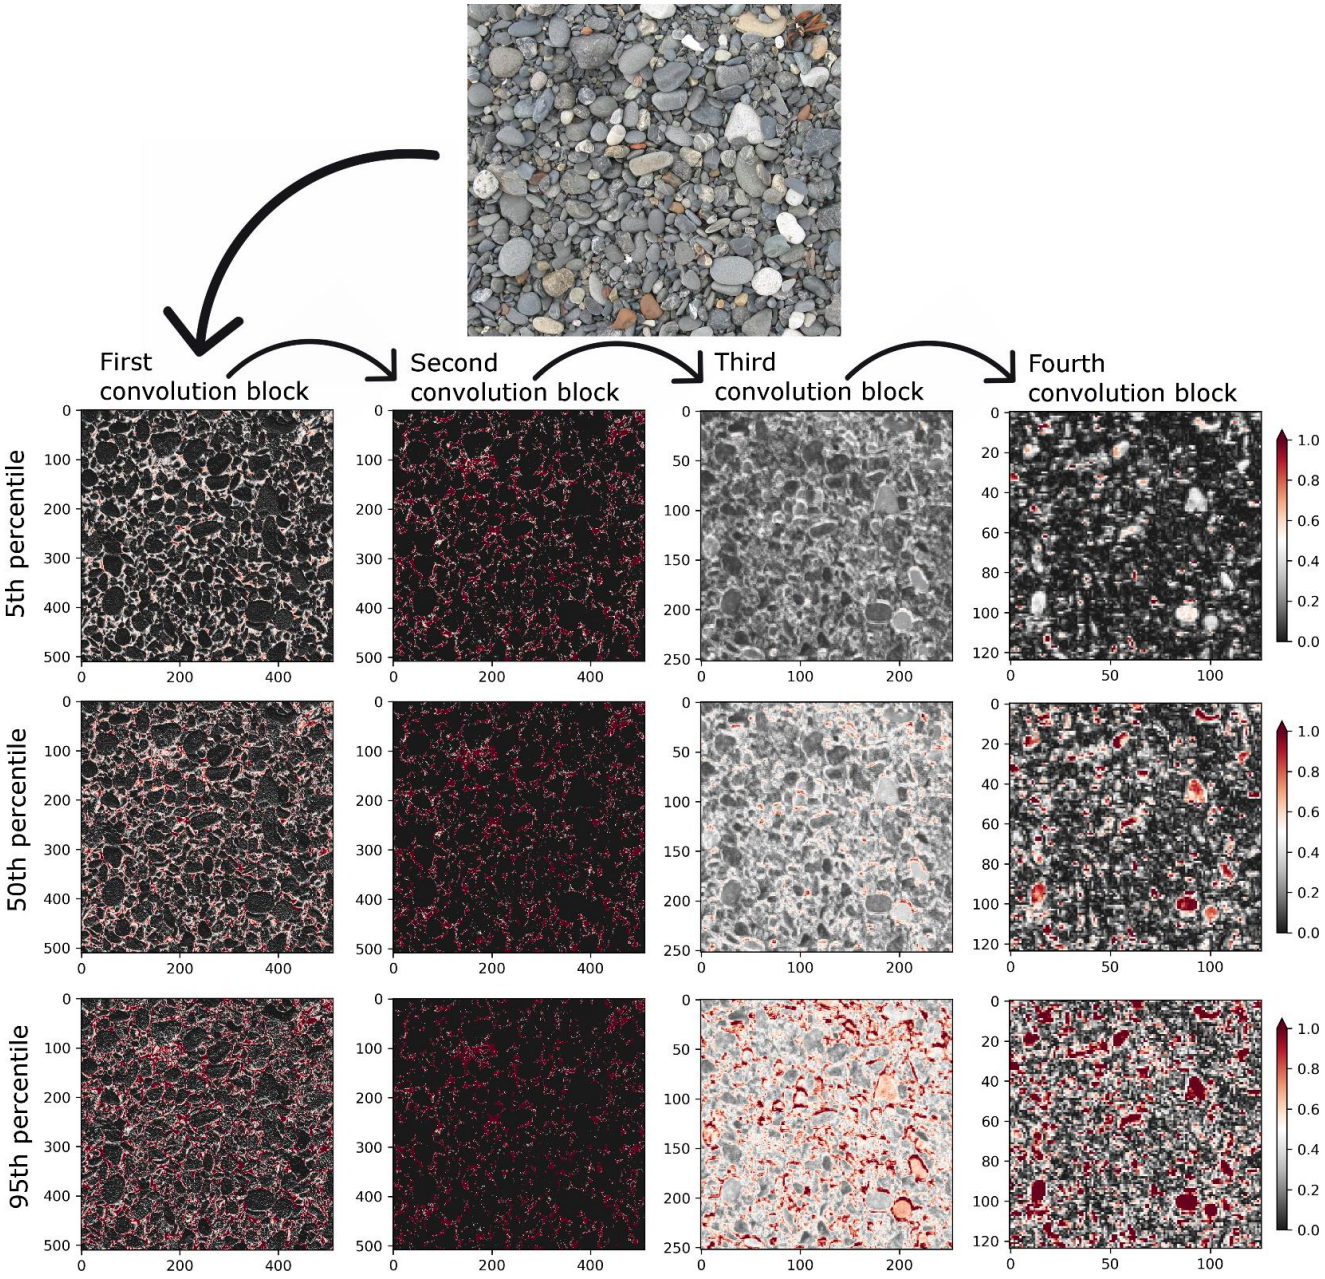

Figure 8. Activation map outputs from each of the four convolutional blocks (columns) in the SediNet model, for three grain size percentiles (rows) for an example image of gravels. Red areas indicate relatively high activation values.

The success of the SediNet approach using images resized to 512 x 512 x 3 pixels, irrespective or original size that was typically much larger, reveals two interesting phenomena. First, an image's aspect ratio does not need to be preserved to provide an accurate grain size, shape or population estimate. Second, those quantities can be estimated even with many subpixel grains, which is the case for relatively fine grains and/or images that have undergone a relatively large amount of downsizing. This is because the model learns

which textures are associated with each grain size, at the scale of imagery provided but regardless of the scale and distortion of pixels. Therefore, the use of the model requires all input imagery to be the same size as that used to train the model. This observation bodes well for applications of this or similar technique on aerial or satellite imagery of sedimentary deposits where most grains exist at subpixel scales, but where spatial resolution is sufficient to create images textures uniquely diagnostic of grain size. Optical granulometry methods similar to *Carbonneau* et al. (2004) operate under the same principles, except in those methods image features are extracted using prescribed filters (and their hyperparameters) such as entropy (and kernel size) rather than those features extracted through an iterative procedure that is optimized to minimize observation-estimate error.

It is useful to visualize which parts of a given image led the model to its final decision. Class Activation Map (CAM) visualization (*Selvaraju et al*., 2017) consists of computing 2D grids of scores associated with a specific output value (such as a specific grain size), computed for every location in any input image, indicating how important each location is with respect to the output value. The “gradCAM” technique of *Selvaraju et al*., (2017) computes the partial differentiation of the predicted output with respect to each channel in a previous layer (the layer for which we want visualize CAMs). The gradient of the resulting activations are scores of how important each channel is for the predicted output, which when multiplied by said channels acts to weigh each channel responsible for the predicted output. The weighted channel-wise mean is the CAM. I implemented this technique by computing the gradient of an image’s estimated grain size with regard to the output feature map of each of the four convolutional blocks in the SediNet grain-size model (Figure 3). Then I computed the product of 1) the mean of the gradient over each feature map channel and 2) each channel in the feature map. Finally, the channel-wise mean of the resulting feature map is our 2D heatmap

of class activation scores. Figure 8 exemplifies this for one example image and the model-estimated grain size associated with the 5^{th}, 50^{th}, and 95^{th} percentiles of the cumulative grain size distribution (rows in Figure), showing CAMs for all four convolution blocks in the SediNet grain-size model (Figure columns). One might interpret each of these 12 CAMs as a spatial map of how intensely the input image activates a specific grain size value, achieved by weighting a spatial map of how intensely the input image activates different channels in the convolutional block by another spatial map of how important each channel is with regard to the grain size value. The analysis demonstrates that each convolution block is weighted to activate different parts of the input image (Figure 8A). The first and second convolutional blocks tend to result in activations in grain interstices only, with generally stronger activations for larger percentiles (compare Figure 8B and 8J, and 8C and 8K). The third and fourth convolution block results in stronger activations for individual grains and grain outlines with generally stronger activations for larger percentiles and for the largest grains (compare Figure 8E and 8M).

Convolutional neural networks have been particular useful for analysis of images because they implement invariance to translation and the convolution filters share weights spatially, which exploits stationarity in the image (*Buscombe and Carini*, 2019; *Goodfellow et al*., 2016). There is typically a lot of stationarity (i.e. repeating spatial patterns) in images of sediment grains, because the location of grains of all sizes within the image is typically random. This is especially the case for relatively well-sorted sediment and or images of relatively large numbers of individual grains, because in those cases grains of all sizes are present in large numbers throughout the image. Training a deep neural network requires fitting a large number of parameters, which usually requires large training datasets. This paper has demonstrated that 400 images might be a sufficiently large data set to train a model that

produces accurate predictions on unseen test images, but I would expect models only to improve by retraining and refining with more data. Data-driven models should also be highly accurate for smaller populations given large training data (Figure 6). Another approach to mitigating any reliance on large datasets is to use simulations to generate supplemental synthetic training data (e.g. *Buscombe*, 2013; *Buscombe and Rubin*, 2012a) or using data augmentation through random image synthesis (e.g. *Buscombe et al*., 2019). Given recent progress in self-supervised deep learning models that do not require data labeling (e.g. *Oh et al*., 2019), it might even soon be possible to estimate sedimentological quantities accurately without manual image classification, manual axes measurements, or some other form of calibration.

Conclusions

I have described a configurable machine-learning framework called SediNet for estimating either (or both) continuous and categorical variables from a photographic image of clastic sediment. To demonstrate the framework, five separate models were configured and trained, three of which for estimating various grain size metrics on both mixed and single populations of sediment, and two for classifying aspects of grain shape and population. Perhaps of most significance is that SediNet can be configured and trained to estimate equivalent sieve diameters directly from image features, without the need for area-to-mass conversion formulas and without even knowing the scale of one pixel. As such, it is the only optical granulometry method proposed to date that does not necessarily require image scaling. SediNet will allow for reliable estimation of several sedimentological variables from arbitrary imagery of sediment, where grains may be either supra- or sub-pixel in scale, and where conversions between grain size measurements on different physical or statistical scales might be learnt directly from the data. The model framework should therefore find numerous

application in the spatio-temporal monitoring of the grain size distribution, shape, mineralogy and other quantities of interest, of sedimentary deposits as they evolve. This study has also served to exemplify how machine learning can be a powerful tool for automated and simultaneous quantitative and qualitative measurements from the same remotely sensed imagery.

Acknowledgments

This work is fully reproducible using data and code at https://github.com/MARDAScience/SediNet, which also includes further examples of how to configure SediNet for different purposes.

References

1. Abadi M.; Agarwal A.; Barham P.; Brevdo E.; Chen Z.; and 35 others. 2015. *TensorFlow: Large-scale machine learning on heterogeneous systems*. Software available online: https://www.tensorflow.org (accessed on 1 July, 2019).
2. Adams J. 1979. Gravel size analysis from photographs. *Journal of the Hydraulics Division. American Society of Civil Engineers*: 1247–1255.
3. Austin, M.J., Masselink, G., O'Hare, T.J. and Russell, P.E. 2007. Relaxation time effects of wave ripples on tidal beaches. *Geophysical Research Letters*, 34(16).
4. Baptista, P., Cunha, T., Gama, C., and Bernardes, C. 2012. A new and practical method to obtain grain size measurements in sandy shores based on digital image acquisition and processing. *Sedimentary Geology*, 282, 294–306.
5. Barnard, P., Rubin, D., Harney, J., and Mustain, N. 2007. Field test comparison of an autocorrelation technique for determining grain size using a digital beachball camera versus traditional methods. *Sedimentary Geology*, 201 (1-2), 180–195.

6. Bergillos, R.J., Ortega-Sánchez, M., Masselink, G. and Losada, M.A. 2016. Morpho-sedimentary dynamics of a micro-tidal mixed sand and gravel beach, Playa Granada, southern Spain. *Marine Geology*, 379, 28-38.
7. Black, M., Carbonneau, P., Church, M., and Warburton, J. 2014. Mapping sub-pixel fluvial grain sizes with hyperspatial imagery. *Sedimentology*, 61 (3), 691–711.
8. Buscombe, D. 2013. Transferable wavelet method for grain-size distribution from images of sediment surfaces and thin sections, and other natural granular patterns. *Sedimentology*, 60 (7), 1709–1732.
9. Buscombe, D., and Carini, R. J. 2019. A data-driven approach to classifying wave breaking in infrared imagery. *Remote Sensing*, 11 (7), 859.
10. Buscombe, D., Carini, R., Harrison, S., Chickadel, C., and Warrick, J. 2019. Optical wave gauging with deep neural networks. *Coastal Engineering*, in review July 2019
11. Buscombe, D. and Masselink, G. 2006. Concepts in gravel beach dynamics. *Earth-Science Reviews*, 79 (1-2), 33-52.
12. Buscombe, D. and Ritchie, A. 2018. Landscape classification with deep neural networks. *Geosciences*, 8 (7), 244.
13. Buscombe, D., and Rubin, D. M. 2012a. Advances in the simulation and automated measurement of well-sorted granular material: 1. Simulation. *Journal of Geophysical Research: Earth Surface*, 117 (F2).
14. Buscombe, D., and Rubin, D. M. 2012b. Advances in the simulation and automated measurement of well-sorted granular material: 2. Direct measures of particle properties. *Journal of Geophysical Research: Earth Surface*, 117 (F2).
15. Buscombe, D., Rubin, D. M., Lacy, J. R., Storlazzi, C. D., Hatcher, G., Chezar, H., and Sherwood, C. R. 2014. Autonomous bed-sediment imaging-systems for revealing

496 temporal variability of grain size. *Limnology and Oceanography: Methods*, 12 (6), 390–
497 406.
498 16. Buscombe, D., Rubin, D., and Warrick, J. 2010. A universal approximation of grain
499 size from images of noncohesive sediment. *Journal of Geophysical Research: Earth*
500 *Surface*, 115 (F2).
501 17. Carbonneau, P., Bizzi, S., and Marchetti, G. 2018. Robotic photosieving from low-cost
502 multirotor sUAS: A proof-of-concept. *Earth Surface Processes and Landforms*, 43 (5),
503 1160–1166.
504 18. Carbonneau PE, Bergeron N, Lane SN. 2005a. Automated grain size measurements
505 from airborne remote sensing for long profile measurements of fluvial grain sizes.
506 *Water Resources Research*: 41.
507 19. Carbonneau PE, Bergeron NE, Lane SN. 2005b. Texture-based image segmentation
508 applied to the quantification of superficial sand in salmonid river gravels. *Earth Surface*
509 *Processes and Landforms* 30: 121–127.
510 20. Carbonneau, P.E., Lane, S.N. and Bergeron, N.E. 2004. Catchment-scale mapping of
511 surface grain size in gravel bed rivers using airborne digital imagery. *Water Resources*
512 *Research*, 40 (7).
513 21. Chollet, F. 2015. Keras. Software available online: https://keras.io (accessed on 1 July,
514 2019).
515 22. Cheng, Z. and Liu, H. 2015. Digital grain-size analysis based on autocorrelation
516 algorithm. *Sedimentary Geology*, 327, 21-31.
517 23. Cuttler, M. V., Lowe, R. J., Falter, J. L., and Buscombe, D. 2017. Estimating the settling
518 velocity of bioclastic sediment using common grain-size analysis techniques.
519 *Sedimentology*, 64 (4), 987–1004.

24. Detert M, Weitbrecht V. 2012. Automatic object detection to analyze the geometry of gravel grains – a free stand-alone tool. In River Flow, 2012, Munoz RM (ed). CRC Press: London; 595–600.

25. Diplas, P. and Fripp, J.B. 1992. Properties of various sediment sampling procedures. *Journal of Hydraulic Engineering*, 118, 955–970.

26. Diplas, P., Kuhnle, R.A., Gray, J.R., Glysson, G.D. and Edwards, T.E. 2008. Sediment transport measurements. In: Sedimentation Engineering: Processes, Measurements, Modeling, and Practice (Ed. M.H. Garcia), 110, 307–353. American Society of Civil Engineering Manuals and Reports on Engineering Practice, Reston, VA

27. Diplas, P. and Sutherland, A.J. 1988. Sampling techniques for gravel sized sediments. *Journal of Hydraulic Engineering*, 114, 484–501.

28. Dugdale, S. J., Carbonneau, P. E., and Campbell, D. 2010. Aerial photosieving of exposed gravel bars for the rapid calibration of airborne grain size maps. *Earth Surface Processes and Landforms*, 35 (6), 627–639.

29. Franklin, S., and Wulder, M. 2002. Remote sensing methods in medium spatial resolution satellite data land cover classification of large areas. *Progress in Physical Geography*, 26 (2), 173–205.

30. Gomez, C., and Purdie, H. 2016. UAV-based photogrammetry and geocomputing for hazards and disaster risk monitoring–A review. *Geoenvironmental Disasters*, 3 (1), 23.

31. Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. 2016. *Deep learning* (Vol. 1). MIT press Cambridge.

32. Graham DJ, Reid I, Rice SP. 2005. Automated sizing of coarse-grained sediments: image-processing procedures. *Mathematical Geology* 37: 1–28.

33. Graham, D., Rollet, A., Rice, S., and Piegay, H. 2012. Conversions of surface grain-size samples collected and recorded using different procedures. *Journal of Hydraulic Engineering*, 138 (10), 839–849.

547 34. Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network
548 training by reducing internal covariate shift. *arXiv preprint* arXiv:1502.03167.
549 35. Jiang, G.-Q., Xu, J., and Wei, J. 2018. A deep learning algorithm of neural network for
550 the parameterization of typhoon-ocean feedback in typhoon forecast models.
551 *Geophysical Research Lett*ers, 45 (8), 3706–3716.
552 36. Kellerhals, R., and Bray, D. I. 1971. Sampling procedures for coarse fluvial sediments.
553 *Journal of the Hydraulics Division*, 97 (8), 1165–1180.
554 37. Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization.
555 *arXivpreprint* arXiv:1412.6980.
556 38. Legleiter, C.J., Stegman, T.K. and Overstreet, B.T. 2016. Spectrally based mapping
557 of riverbed composition. *Geomorphology*, 264, 61-79.
558 39. Linville, L., Pankow, K., and Draelos, T. 2019. Deep learning models augment analyst
559 decisions for event discrimination. *Geophysical Research Letters*, 46 (7), 3643–3651.
560 40. Luo, J. Y., Irisson, J.-O., Graham, B., Guigand, C., Sarafraz, A., Mader, C., and
561 Cowen, R. K. 2018. Automated plankton image analysis using convolutional neural
562 networks. *Limnology and Oceanography: Methods*, 16 (12), 814–827.
563 41. Masteller, C.C. and Finnegan, N.J. 2017. Interplay between grain protrusion and
564 sediment entrainment in an experimental flume. *Journal of Geophysical Research:*
565 *Earth Surface*, 122 (1), 274-289.
566 42. Michaelides, K., Hollings, R., Singer, M.B., Nichols, M.H. and Nearing, M.A. 2018.
567 Spatial and temporal analysis of hillslope–channel coupling and implications for the
568 longitudinal profile in a dryland basin. *Earth Surface Processes and Landforms*, 43 (8),
569 1608-1621.
570 43. Montgomery, D. R., Panfil, M. S., and Hayes, S. K. 1999. Channel-bed mobility
571 response to extreme sediment loading at Mount Pinatubo. *Geology*, 27 (3), 271–274.

572 44. Mulder, V., De Bruin, S., Schaepman, M., and Mayr, T. 2011. The use of remote
573 sensing in soil and terrain mapping: A review. *Geoderma*, 162 (1-2), 1–19.
574 45. Nield, J.M., Wiggs, G.F. and Squirrell, R.S. 2011. Aeolian sand strip mobility and
575 protodune development on a drying beach: examining surface moisture and surface
576 roughness patterns measured by terrestrial laser scanning. *Earth Surface Processes*
577 *and Landforms*, 36 (4), 513-522.
578 46. Novak-Szabo, T., Sipos, A.A., Shaw, S., Bertoni, D., Pozzebon, A., Grottoli, E.,
579 Jerolmack, D. J. 2018. Universal characteristics of particle shape evolution by bed-
580 load chipping. *Science Advances*, 4 (3).
581 47. Oh, C., Ham, B., Kim, H., Hilton, A. and Sohn, K. 2019. OCEAN: Object-centric
582 arranging network for self-supervised visual representations learning. *Expert Systems*
583 *with Applications*, 125, 281-292.
584 48. Paterson, G. A., and Heslop, D. 2015. New methods for unmixing sediment grainsize
585 data. *Geochemistry, Geophysics, Geosystems*, 16 (12), 4494–4506.
586 49. Pfeiffer, A. M., Finnegan, N. J., and Willenbring, J. K. 2017. Sediment supply controls
587 equilibrium channel geometry in gravel rivers. *Proceedings of the National Academy*
588 *of Sciences*, 114 (13), 3346–3351.
589 50. Prodger, S., Russell, P., Davidson, M., Miles, J. and Scott, T. 2016. Understanding
590 and predicting the temporal variability of sediment grain size characteristics on high-
591 energy beaches. *Marine Geology*, 376, 109-117.
592 51. Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N.
593 2019. Deep learning and process understanding for data-driven Earth system science.
594 *Nature*, 566 (7743), 195.
595 52. Rickenmann, D., and Recking, A. 2011. Evaluation of flow resistance in gravel-bed
596 rivers through a large field data set. *Water Resources Research*, 47 (7).

53. Rubin, D. 2004. A simple autocorrelation algorithm for determining grain size from digital images of sediment. *Journal of Sedimentary Research*, 74 (1), 160–165.
54. Rubin, D., Buscombe, D., Wright, S., Topping, D., Grams, P., and Schmidt, J. 2019. What grain size reveals about suspended-sand transport in the Colorado River in Grand Canyon. *Journal of Geophysical Research - Earth Surface*, in review.
55. Rubin, D., Chezar, H., Harney, J., Topping, D., Melis, T., and Sherwood, C. 2007. Underwater microscope for measuring spatial and temporal changes in bed-sediment grain size. *Sedimentary Geology*, 202 (3), 402–408
56. Sedinet online software. 2019. https://github.com/MARDAScience/SediNet
57. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D. and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE International Conference on Computer Vision*, 618-626.
58. Shoji, D., Noguchi, R., Otsuki, S. and Hino, H. 2018. Classification of volcanic ash particles using a convolutional neural network and probability. *Scientific Reports*, 8 (1), 8111.
59. Smith, M., and Pain, C. 2009. Applications of remote sensing in geomorphology. *Progress in Physical Geography*, 33 (4), 568–582.
60. Smith, M.E., Werner, S.H., Buscombe, D., Finnegan, N.J., Sumner, E.J. and Mueller, E.R. 2018. Seeking the shore: Evidence for active submarine canyon head incision due to coarse sediment supply and focusing of wave energy. *Geophysical Research Letters*, 45 (22), 12-403.
61. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15 (1), 1929–1958.

62. Sternberg, R.W., Berhane, I. and Ogston, A.S. 1999. Measurement of size and settling velocity of suspended aggregates on the northern California continental shelf. *Marine Geology*, 154 (1-4), 43-53.
63. Turner, I. L., Harley, M. D., and Drummond, C. D. 2016. UAVs for coastal surveying. *Coastal Engineering*, 114, 19–24.
64. Turner, I.L., Russell, P.E. and Butt, T. 2008. Measurement of wave-by-wave bed-levels in the swash zone. *Coastal Engineering*, 55 (12), 1237-1242.
65. Viles, H. 2016. Technology and geomorphology: Are improvements in data collection techniques transforming geomorphic science? *Geomorphology*, 270, 121–133.
66. Warrick, J.A., Rubin, D.M., Ruggiero, P., Harney, J.N., Draut, A.E. and Buscombe, D. 2009. Cobble Cam: Grain-size measurements of sand to boulder from digital photographs and autocorrelation analyses. *Earth Surface Processes and Landforms*, 34 (13), 1811-1821.
67. Wheatcroft, R.A. and Borgeld, J.C. 2000. Oceanic flood deposits on the northern California shelf: large-scale distribution and small-scale physical properties. *Continental Shelf Research*, 20 (16), 2163-2190.
68. Williams, R.D., Brasington, J., Vericat, D. and Hicks, D.M. 2014. Hyperscale terrain modelling of braided rivers: fusing mobile terrestrial laser scanning and optical bathymetric mapping. *Earth Surface Processes and Landforms,* 39 (2), 167-183.
69. Woodget, A., and Austrums, R. 2017. Subaerial gravel size measurement using topo-graphic data derived from a UAV-SfM approach. *Earth Surface Processes and Landforms*, 42 (9), 1434–1443.
70. Woodget, A., Fyffe, C., and Carbonneau, P. 2018. From manned to unmanned aircraft: Adapting airborne particle size mapping methodologies to the characteristics of sUAS and SfM. *Earth Surface Processes and Landforms*, 43 (4), 857–870.