

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

3 Daniel Buscombe

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

5

6 Abstract

7 I describe a configurable machine-learning framework to estimate a suite of continuous and  
8 categorical sedimentological properties from photographic imagery of sediment, and to  
9 exemplify how machine learning can be a powerful and flexible tool for automated quantitative  
10 and qualitative measurements from remotely sensed imagery. The model is tested on a  
11 dataset consisting of 409 images and associated detailed label data. The data are from a  
12 much wider sedimentological spectrum than previous optical granulometry studies,  
13 consisting of both well- and poorly sorted sediment, terrigenous, carbonate, and  
14 volcanoclastic sands and gravels and their mixtures, and grain sizes spanning over two orders  
15 of magnitude. I demonstrate the model framework by configuring it in several ways, to  
16 estimate two categories (describing grain shape and population, respectively) and nine  
17 numeric grain-size percentiles in pixels from a single input image. Grain size is then  
18 recovered using the physical size of a pixel. Finally, I demonstrate that the model can be  
19 configured and trained to estimate equivalent sieve diameters directly from image features,  
20 without the need for area-to-mass conversion formulas and without even knowing the scale  
21 of one pixel. Thus, it is the only optical granulometry method proposed to date that does not  
22 necessarily require image scaling. The flexibility of the model framework should facilitate  
23 numerous application in the spatio-temporal monitoring of the grain size distribution, shape,  
24 mineralogy and other quantities of interest, of sedimentary deposits as they evolve as well  
25 as other texture-based proxies extracted from remotely sensed imagery.

26

## 27 1. Introduction

28 Sediment grain size fundamentally influences the physics of flows of water, wind, ice and  
29 sediment that continually shape landforms. Large sedimentological datasets have led to  
30 important discoveries in dynamic environments such as contemporary river beds, sea beds  
31 and aeolian sediment surfaces that are constantly changing under fluid power, for example  
32 in sediment transport (e.g. *Masteller & Finnegan, 2017; Rubin et al., 2019*), channel bed  
33 mobility (e.g. *Montgomery et al., 1999*), channel geometry (e.g. *Pfeiffer et al., 2017*), sediment  
34 provenance (e.g. *Paterson & Heslop, 2015*), sediment abrasion (e.g. *Novak-Szabo et al.,*  
35 *2018*), hydraulic resistance (e.g. *Rickenmann & Recking, 2011*), particle settling (e.g.  
36 *Sternberg et al., 1999*) and dispersal at coasts (e.g. *Wheatcroft & Borgeld, 2000*), and beach  
37 dynamics (e.g. *Bergillos et al., 2016*). Traditionally, the means of acquiring large grain size  
38 (or shape, or any other metric) data sets has been laborious and time-consuming through  
39 laboratory analyses of samples taken in the field. Optical granulometry is the measurement  
40 of sediment from statistical analysis of image intensity and texture, and has been driven by  
41 instrumental (e.g. *Buscombe et al., 2014; Carbonneau et al., 2018; Rubin et al., 2007;*  
42 *Woodget et al., 2018*) and analytical (e.g. *Black et al., 2014; Buscombe et al., 2010;*  
43 *Buscombe and Rubin, 2012b; Buscombe, 2013; Cheng and Liu, 2015; Carbonneau et al.,*  
44 *2005a, 2005b; Carbonneau et al., 2004; Cuttler et al., 2017; Dugdale et al., 2010; Legleiter*  
45 *et al., 2016; Rubin, 2004; Woodget et al., 2017*) developments over the past 15 years.  
46 Another set of deterministic methods known as 'photosieving' (e.g. *Adams, 1979*) or object-  
47 based image analysis or OBIA (*Carbonneau et al., 2018*) have been developed (e.g. *Detert*  
48 *and Weitbrecht, 2012; Graham et al., 2005*) that aim to identify each individual grain and  
49 cannot therefore be used on grains smaller than one pixel (subpixel) which is not a theoretical  
50 limitation of optical granulometry techniques that statistically quantify image texture

51 (*Carbonneau et al.*, 2004). One major goal of this corpus of work is to develop a reliable suite  
52 of techniques for spatio-temporal monitoring of the grain size of sedimentary deposits as they  
53 evolve, remotely and automatically. This has the potential to significantly alter the way  
54 geomorphological research is carried out (e.g. *Viles*, 2016) and may hopefully lead to  
55 significant discoveries in the two-way feedbacks between evolving sedimentary landform  
56 morphologies and the spatio-temporal dynamics of grain size, or 'morpho-sedimentary  
57 dynamics' (cf. *Buscombe and Masselink*, 2006), at large field scales. This will require  
58 measuring grain size at the same spatial (e.g. *Rubin et al.*, 2019) and temporal (e.g.  
59 *Buscombe et al.*, 2014) coverage as is now possible with topographic measurements that  
60 can capture the spatio-temporal evolution of small-scale morphologies (e.g. *Austin et al.*,  
61 2007; *Nield et al.*, 2011; *Turner et al.*, 2008; *Williams et al.*, 2014).

62

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

73

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

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

$$80 \quad p(V - W)_i = \frac{p(A)_i D_i^x}{\sum p(A)_i D_i^x} \quad (1)$$

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

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

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

110

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

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

131

132 The final and perhaps foremost motivation to developing yet another optical granulometry  
133 technique is the observation that the data-hungry nature of machine learning allows for  
134 collaborative tool development for extracting scientific information from images of sediment.  
135 Recognizing the variety of both sediment imagery, due to the inherent variability of natural  
136 sediment, and potential SediNet applications, the motivating idea behind the creation of the  
137 SediNet model and software (*SediNet online software*, 2019) is to foster the creation of such  
138 a community. As explained further in Section 5.4, users can contribute imagery, models, and  
139 retrain existing models, as well as using existing SediNet models contained in the repository.

140

## 141 2. Data

142 The model is trained and tested on a large data set consisting of 409 labeled images of  
143 sediment (Figures 1 and 2), with a large variation in the spatial footprint (field-of-view) of each  
144 image, the spatial resolution (physical size of a pixel), and variation in camera sensor. The  
145 data are from a wide sedimentological spectrum of well and poorly sorted sediment,  
146 consisting of terrigenous (derived by erosion of crystalline, volcanic, and sedimentary rocks),  
147 carbonate (skeletal grains, oolites, and some locally derived detrital carbonate), and  
148 volcanoclastic (lapilli, glass, and pyroclastic bombs) sands and gravels and their mixtures,  
149 and grain sizes in pixels spanning over two orders of magnitude. Out of the 409 images, 300

150 were compiled and used by *Buscombe* (2013) to develop a wavelet-based algorithm for  
151 estimating grain size from imagery (sets A and C in that paper). The remaining 109 samples  
152 were compiled for this study, from various fieldwork activities over more than 10 years in  
153 various coastal and riverine environments on several continents. The additional 100 samples  
154 were chosen specifically to better represent within the dataset both poorly sorted mixed sand  
155 and gravel sediment and (usually microscopic) imagery with relatively few numbers of grains.

156

### 157 2.1. Grain Size

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

166

### 167 2.2. Grain shape and population


168 The expanded dataset of 409 images contain a number of sediment populations (Figure 1)  
169 that I manually grouped into six categories: 1) well-sorted gravel; 2) well-sorted sand and  
170 shell hash from underwater camera (described in *Buscombe et al.*, 2014); 3) relatively poorly  
171 sorted gravel and sand-gravel mixtures (including imagery from *Warrick et al.*, 2009); 4) well-  
172 sorted sand; 5) miscellaneous terrigenous and volcanoclastic grains; and 6) miscellaneous  
173 bioclastic (carbonate) grains. Additionally, each of the 409 images were classified into four

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

176

177 Discrete classification schemes are subjective and the purpose is to train a machine to  
178 replicate or simulate expert judgement, as is the case here. I classified the shape and  
179 population of each image by eye. This was a relatively straightforward and objective for grain  
180 population, since populations 5 (terrigenous/volcaniclastic) and 6 (bioclastic) were the two  
181 largest groups when different clastic groupings (sand versus gravel) could be made.  
182 Populations 5 and 6 are obviously distinct because the latter are invariably white or near  
183 white, and almost all with regular rather than irregular shapes. Of the various sand and gravel  
184 groupings (two of each), the sand populations were naturally split in terms of  
185 subaerial/subaqueous. The only visually difficult distinction was then between 'well' and  
186 'poorly' gravels, which was determined using known grain-size distributions. Since this study  
187 is fully reproducible using software described in section 5.4, the interested reader is  
188 encouraged to explore different subjective groupings of the provided 409 sediment images  
189 and of their own. The process of visually classifying grain shape was more subjective,  
190 especially in the distinction between well-rounded and angular grains in marginal cases. The  
191 process was also iterative; originally, I did not make distinction between large (mean grain  
192 size > 100 pixels) and small (< 100 pixels) grains but the overall classification accuracy was  
193 not as high.

194



195


196

197

198

199

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 volcanoclastic grains; and 6) miscellaneous bioclastic (carbonate) grains.



200

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

203

204 **3. The SediNet model framework**

205 SediNet is a new deep learning model framework that uses ‘end-to-end’ training to extract  
 206 relevant features from imagery for a specific optical granulometry task. The framework refers  
 207 to the concepts and architecture of the model. Each trained model is a particular instance or  
 208 implementation of SediNet that has a specific purpose, which might be measuring a specific  
 209 grain size metric, or estimate a categorical variable. Completing those two tasks would

210 require two different SediNet models, but using the same basic model architecture  
211 repurposed for each task. The example implementations used to exemplify the SediNet  
212 framework, described below, each use the concept of ‘end-to-end’ learning (Goodfellow et  
213 al., 2016) whereby the framework is trained from scratch by optimizing the accuracy of the  
214 output by optimizing the assessment of the relevancy of each extracted image feature to the  
215 specific task. Therefore, the features extracted for grain shape estimation would be different  
216 from those extracted for grain size estimation, for example. However once the model is  
217 trained, it should generalize well to new, unseen imagery not that included in the training set.

218

219 ‘End-to-end’ training is different to the concept of ‘transfer learning’, which is the practice of  
220 using model layers to extract features from imagery, then using that ‘feature extractor’ model  
221 to predict an arbitrary set of classes (Buscombe and Ritchie, 2018) or continuous variable.  
222 This approach requires training data consisting of thousands to millions of example images  
223 and labels. I therefore deemed 409 images to be too small a data set for its successful  
224 application, but might be worth exploring in subsequent work with larger training data sets  
225 available.

226

### 227 3.1. Example SediNet Implementations

228 SediNet (Figure 3) is a supervised deep neural network model framework that can be used  
229 as presented in this paper, or alternatively configured for custom purposes, by training on  
230 any number of input images for any number of numeric or categorical outputs. For the  
231 purposes of demonstrating the model in this paper, several SediNet models were made:

232

- 233 1. To estimate nine percentiles of the cumulative grain size distribution in pixels, trained  
234 on 204 images and tested on 205 images, both drawn randomly. The train and test

235 sets consist of images of several populations of grains from a wide sedimentological  
236 spectrum

237 2. To estimate nine percentiles of the cumulative grain size distribution in pixels, trained  
238 on 15 images and tested on 16 images of one population (beach sands)


239 3. To estimate sieve size in microns directly, without first estimating the pixel size, trained  
240 on the same 15 images and tested on 16 images as above

241 4. To estimate six categorical populations of grains, trained on 204 images and tested  
242 on 205 images, both drawn randomly

243 5. To estimate four categorical grain shape/size classes, trained on 204 images and  
244 tested on 205 images, both drawn randomly

245

246



247

248 *Figure 3. Schematic of the SediNet architecture, as applied to estimating the grain size distribution, and categorical population and*  
 249 *shape/size. An input image is passed to the feature extractor consisting of a series of convolutional blocks. The last set of feature*  
 250 *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*  
 251 *of estimating grain size percentiles, sediment population, and grain shape.*

252

253 **3.2. SediNet architecture**

254 Deep learning models have multiple processing layers (called convolutional layers or blocks)  
 255 and nonlinear transformations (that include batch normalization, activation, and dropout,  
 256 which are explained below), with the outputs from each layer passed as inputs to the next.

257 The image feature extractor consists of four convolutional blocks each consisting of a several

258 two-dimensional convolutional filter layers, batch normalization layers, and two-dimensional  
259 max pooling layers (Figure 3). Batch normalization applies a transformation that maintains  
260 the mean neuron activation of zero and the activation standard deviation of one (*Ioffe and*  
261 *Szegedy, 2015*). Pooling layers are used to reduce the spatial dimensions of each of the  
262 three-dimensional tensors associated with each pixel of the input image, from  $h \times w \times d$  to  $1$   
263  $\cdot 1 \times d$ , by averaging over  $h$  and  $w$ . This has the effect of reducing the total number of  
264 parameters in the model, thereby minimizing overfitting. The output of the last block is the  
265 input of the next. The number of filters increases for each of the four blocks, from 16 in the  
266 first block, 32 in the second, 48 in the third and finally 64 in the last block. After the last  
267 convolutional block, there is one more batch normalization and two-dimensional max pooling  
268 layer, and a dropout layer that randomly drops half the neurons (*Srivastava et al., 2014*).  
269 Batch normalization, max pooling, and dropout layers are techniques to prevent overfitting  
270 the model (i.e., memorizing the training data rather than learning a general trend). The  
271 extracted feature is fed into a series of multilayer perceptrons, one for each estimated  
272 quantity, that each culminates in a dense predicting layer with linear regression (known in  
273 machine learning literature as a linear activation function) for continuous variable prediction  
274 variables (such as grain size in pixels, or sieve size directly), or multinomial logistic regression  
275 (in machine learning parlance, a softmax activation function) for categorical variables such  
276 as grain shape and population.

277

### 278 3.3. Training SediNet Implementations

279 Given the set of  $n$  images, let us denote one sample  $X_\mu \in \mathbb{R}^p$  with  $\mu = 1 \dots n$ , where  $p$  is the  
280 number of pixels. For each sample,  $X_\mu$  there is label  $y_\mu \in \mathbb{R}^q$  where  $q$  is the number of  
281 combined categorical and continuously distributed classes. Using the deep learning  
282 architecture described above, and the training data set  $\{X_\mu, y_\mu\}$  consisting of 50 % of the total

283 number of images, randomly selected, a function  $f$  is found such that  $\hat{y} = f(X)$ , where  $\hat{y}$  is  
284 the predicted set of labels/metrics from sample image  $X$ . The remaining 50 % of the total data  
285 set was used as a test set to evaluate model performance.


286

287 The model was retrained 'end-to-end', which means it was initialized with random numbers  
288 for neuron weights  $w \in \mathbb{R}^k$ , then during training the value of those parameters was optimized  
289 by minimizing the discrepancy between known and estimated quantities by minimizing a loss  
290 function  $L[f_w(X_\mu, y_\mu)]$  for each sample  $\mu$  where  $f_w$  denotes weighted function. By doing so,  
291 the model simultaneously and automatically learns feature representations from imagery and  
292 a mapping from those features to the target values (e.g. grain size) or classes (e.g. grain  
293 shape). Models are trained over several epochs. One training epoch means that the learning  
294 algorithm has made one pass through the training dataset, where examples were separated  
295 into randomly selected batches of images. The number of training steps per epoch was  
296 computed as the number of training images divided by the batch size. In this study, the batch  
297 size was set to eight and results were not sensitive to its value (I revisit this in the Discussion).  
298 Upon each step, the gradients of the network are updated and new weights assigned to each  
299 neuron. Stochastic gradient descent was used to iteratively adjusting the weights in the  
300 direction of the gradient of the average of the loss over the training set using  $w^{t+1} = w^t -$   
301  $\lambda \nabla_w R(f_w)$ , where  $t$  is iteration number (step within an epoch) and  $\lambda$  is the so-called 'learning  
302 rate', and where  $R(f_w) = \sum L/n$  for the full training data is replaced by the contribution of just  
303 a few of the samples.

304

305 During model training, each  $h \times w \times 3$  pixel input image was resized to  $512 \times 512 \times 3$  pixels  
306 for computational efficiency. With sufficient computing power, larger images and larger  
307 numbers of images could be used. That the image's aspect ratio is typically not preserved

308 does not affect model performance (I revisit this point in the Discussion). The method was  
309 implemented in python 3.7 using the Tensorflow (*Abadi et al.*, 2015) backend to the keras  
310 (*Chollet et al.*, 2015) module, on a GeForce RTX 2080Ti GPU with 11 GB of memory. The  
311 resolution of a given grain size estimate in pixels is approximately 2 pixels, determined as the  
312 range of that variable in the training data (in the present case, the largest grain size minus  
313 the smallest, which is approximately 1000 pixels) divided by the number of neurons in the  
314 final dense layer, which was set to 512 (Figure 3). Training utilized the popular Adam  
315 algorithm (*Kingma and Ba*, 2014) for stochastic optimization, with parameters  $\beta_1= 0.9$  and  
316  $\beta_2= 0.999$  (*Buscombe et al.*, 2019). During training,  $\lambda$  was automatically reduced when the  
317 loss function stabilized, i.e. when its value stopped decreasing, by a factor of 0.8 after 15  
318 epochs had elapsed with no improvement (*Buscombe et al.*, 2019). A lower bound on  $\lambda$  was  
319 set at 0.0001. The maximum number of training epochs was set to 100. Models stopped  
320 training early (i.e. before 100 epochs) if the validation loss failed to improve for 20 consecutive  
321 epochs. Models typically trained for between 40 and 100 epochs before the criterion was met  
322 to stop training early.



323

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

327

## 328 4. Results

### 329 4.1. Grain Size

330 The first implementation of SediNet estimated nine percentiles of the cumulative grain size


331 distribution in pixels, trained on 204 images with mean error between 24 and 52% depending

332 on percentile, and tested on 205 images with mean error between 24 and 45% again varying  
333 with percentile (Figures 4 and 5). Mean percent error for each percentile is computed as 100  
334 times the root-mean-squared error normalized by the mean grain size associated with that  
335 percentile. Overall, this SediNet model out-performed the wavelet technique of *Buscombe*  
336 (2013) and required fewer tunable parameters.

337

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


351



352

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

355



356

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

362

363

364 **4.2. Grain shape and population**


365 The fourth implementation of SediNet estimated six categorical populations of sediment,  
 366 trained on 200 images and tested on 200 images, both drawn randomly without replacement.

367 Classification skill was evaluated using a ‘confusion matrix’ of normalised correspondences  
 368 between true and estimated labels (Figure 7, A - C). A perfect correspondence between true

369 and estimated labels is scored 1.0 along the diagonal elements of the matrix. Random  
 370 misclassifications are readily identified as off-diagonal elements with relatively small

371 magnitudes, and systematic misclassifications are recognized as off-diagonal elements with

372 relatively large magnitudes. The three confusion matrices for categorical sediment population  
 373 shown in Figure 7, A – C show skill for, respectively, training, testing and combined (i.e. all  
 374 409 images) data. The model overfits population 2 (underwater images of continental shelf  
 375 sand, Figure 1), evidenced by the large discrepancy between training skill (1.0) and test skill  
 376 (0.62; Figure 7A, B). However, overfitting is not evident for the other five classes, with test  
 377 scores being approximately equal to training scores. All classes are classified with accuracies  
 378 of > 70% for the combined model (Figure 7C).



379

380 *Figure 7. Confusion matrices for (A – C) categorical population and (D – F) categorical shape. Subplots A and B show training and*  
 381 *testing datasets. Subplot C shows classification accuracies for the combined train and test dataset.*

382

383 The fifth and final SediNet implementation reported here was configured to estimate four  
 384 categorical grain shape/size classes, trained on 200 images and tested on 200 images, both  
 385 drawn randomly. The three confusion matrices for categorical sediment shape shown in

386 Figure 7D – F show skill for, respectively, training, testing and combined (i.e. all 409 images)  
387 data. The similarity in train and test scores for all four classes demonstrates the model has  
388 not overfit the data. All classes are classified with accuracies of > 85% for the test, train and  
389 combined models (Figure 7D - F). Despite the subjective nature of manual image  
390 classification, the model performed excellently for grain shape. The same is true of population  
391 except that population classes 2 (well-sorted sand and shell hash from underwater camera)  
392 was often mistaken for class 4 (well sorted sand), which made physical sense because both  
393 samples are sand, therefore statistical explanations for the discrepancy were not sought. I  
394 conclude that either the model has not generalized well (i.e. that the 'sand' signal is more  
395 dominant than whether or not the imagery is dark/submerged) or that there are too few or too  
396 unequal numbers of images in each class. I revisit the potential effects of this so-called 'class  
397 imbalance' in section 5.4.

398

## 399 5. Discussion

### 400 5.1. Potential Applications


401 The task of quantifying and classifying natural objects and textures in images of sedimentary  
402 landforms is increasingly widespread in a wide variety of geomorphological research  
403 (*Franklin and Mulder, 2002; Mulder et al., 2011; Smith and Pain, 2009*), especially as imagery  
404 collection using UAVs becomes more prevalent (*Carbonneau et al., 2018; Gomez and*  
405 *Purdie, 2016; Turner et al., 2016*). The automated method to size and classify sediment  
406 described here could maximize speed and objectivity of sedimentary description at large  
407 scales, and might be applied to the analysis of datasets consisting of tens to millions of  
408 individual images. The model framework could enable spatio-temporal monitoring of grain  
409 size more efficiently, being configurable to estimate many custom-defined quantities and  
410 qualities for specific tasks. Given it is a data-driven approach, models trained for use in

411 specific environments will highly likely be as or more accurate than methods such as  
412 *Buscombe* (2013) and *Carbonneau et al.*, (2004) that are based on signal processing or  
413 random field theory, especially for poorly sorted sediment, small field-of-view, and large grain  
414 size compared to field-of-view (small numbers of individual grains). This is because those  
415 methods are not informed by data (i.e. only tested with data); therefore, the massive variation  
416 in natural sediment can only be a limitation in their application.

417

418 Convolutional neural networks have been particular useful for analysis of images because  
419 they implement invariance to translation and the convolution filters share weights spatially,  
420 which exploits stationarity in the image (*Buscombe and Carini*, 2019; *Goodfellow et al.*, 2016).  
421 There is typically a lot of stationarity (i.e. repeating spatial patterns) in images of sediment  
422 grains, because the location of grains of all sizes within the image is typically random. This  
423 is especially the case for relatively well-sorted sediment and or images of relatively large  
424 numbers of individual grains, because in those cases grains of all sizes are present in large  
425 numbers throughout the image. Training a deep neural network requires fitting a large  
426 number of parameters, which usually requires large training datasets. This paper has  
427 demonstrated that 409 images might be a sufficiently large data set to train a model that  
428 produces accurate predictions on unseen test images, but I would expect models only to  
429 improve by retraining and refining with more data. Data-driven models should also be highly  
430 accurate for smaller populations given large training data (Figure 6). Another approach to  
431 mitigating any reliance on large datasets is to use simulations to generate supplemental  
432 synthetic training data (e.g. *Buscombe*, 2013; *Buscombe and Rubin*, 2012a) or using data  
433 augmentation through random image synthesis (e.g. *Buscombe et al.*, 2019). Given recent  
434 progress in self-supervised deep learning models that do not require data labeling (e.g. *Oh*  
435 *et al.*, 2019), it might even soon be possible to estimate sedimentological quantities

436 accurately without manual image classification, manual axes measurements, or some other  
437 form of calibration.



438

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

441

442


## 443 5.2. Visualizing How a Model Works

444 It is useful to visualize which parts of a given image led the model to its final decision. Class  
445 Activation Map (CAM) visualization (*Selvaraju et al., 2017*) consists of computing 2D grids of

446 scores associated with a specific output value (such as a specific grain size), computed for  
447 every location in any input image, indicating how important each location is with respect to  
448 the output value. The “gradCAM” technique of *Selvaraju et al.*, (2017) computes the partial  
449 differentiation of the predicted output with respect to each channel in a previous layer (the  
450 layer for which we want visualize CAMs). The gradient of the resulting activations are scores  
451 of how important each channel is for the predicted output, which when multiplied by said  
452 channels acts to weigh each channel responsible for the predicted output. The weighted  
453 channel-wise mean is the CAM. I implemented this technique by computing the gradient of  
454 an image’s estimated grain size with regard to the output feature map of each of the four  
455 convolutional blocks in the SediNet grain-size model (Figure 3). Then I computed the product  
456 of 1) the mean of the gradient over each feature map channel and 2) each channel in the  
457 feature map. Finally, the channel-wise mean of the resulting feature map is our 2D heatmap  
458 of class activation scores. Figure 8 exemplifies this for one example image and the model-  
459 estimated grain size associated with the 5<sup>th</sup>, 50<sup>th</sup>, and 95<sup>th</sup> percentiles of the cumulative grain  
460 size distribution (rows in Figure), showing CAMs for all four convolution blocks in the SediNet  
461 grain-size model (Figure columns). One might interpret each of these 12 CAMs as a spatial  
462 map of how intensely the input image activates a specific grain size value, achieved by  
463 weighting a spatial map of how intensely the input image activates different channels in the  
464 convolutional block by another spatial map of how important each channel is with regard to  
465 the grain size value. The analysis demonstrates that each convolution block is weighted to  
466 activate different parts of the input image (Figure 8A). The first and second convolutional  
467 blocks tend to result in activations in grain interstices only, with generally stronger activations  
468 for larger percentiles (compare Figure 8B and 8J, and 8C and 8K). The third and fourth  
469 convolution block results in stronger activations for individual grains and grain outlines with

470 generally stronger activations for larger percentiles and for the largest grains (compare Figure  
471 8E and 8M).

472



473

474 *Figure 9. Per-image percent error in three grain-size percentile estimates, as a function of A) the image's original aspect ratio and B)*  
475 *the change in aspect ratio due to image resizing. The lack of correlation suggests*

476

### 477 5.3. Image Resolution and Aspect Ratio

478 The use of SediNet models currently requires that all input imagery to be the same size as  
479 that used to train the model. Images were resized to 512 x 512 x 3 pixels, irrespective of  
480 original size that was typically much larger. However, there is no correlation between  
481 prediction error and an image's aspect ratio (Figure 9A), nor is there correlation between  
482 error and the change in aspect ratio as a result of resizing to 512 x 512 pixels (Figure 9B). In  
483 addition, there is not a consistent image size or aspect ratio per class; images in most classes  
484 have a wide range of aspect ratios. Therefore, the success of the SediNet approach reveals  
485 two interesting phenomena. First, an image's aspect ratio does not need necessarily to be  
486 preserved to provide an accurate grain size, shape or population estimate. Second, those

487 quantities can be estimated even with many subpixel grains, which is the case for relatively  
488 fine grains and/or images that have undergone a relatively large amount of downsizing. This  
489 is because the model apparently learns which textures are associated with each grain size,  
490 at the scale of imagery provided but regardless of the scale and distortion of pixels. Intuitively,  
491 the image texture should be sensitive to image distortion, as it will change the anisotropy of  
492 the grain axis. While aspect ratio preservation may improve model results, warranting further  
493 investigation in subsequent work, there is such a wide variety of image aspect ratios  
494 represented in the training data, from 46% to 139% (Figure 9), that the model training  
495 automatically picks up on image features that are less sensitive to distortion. For example,  
496 Figure 8 clearly demonstrates that the algorithm is largely agnostic to that distortion because  
497 it isn't activating the pixels associated with individual grains. Rather, it is scaling activations  
498 between grains and interstices to make predictions.

499

500 This observation bodes well for applications of this or similar technique on aerial or satellite  
501 imagery of sedimentary deposits where most grains exist at subpixel scales, but only where  
502 spatial resolution is sufficient to create images textures uniquely diagnostic of grain size.  
503 Optical granulometry methods similar to *Carbonneau et al. (2004)* operate under the same  
504 principles, except in those methods image features are extracted using prescribed filters (and  
505 their hyperparameters) such as entropy (and kernel size) rather than those features extracted  
506 through an iterative procedure that is optimized to minimize observation-estimate error. That  
507 said a consistent behavior observed in the three SediNet models for grain-size was over-  
508 estimation of the size of finer grains. Examination of these images reveals that image  
509 downsizing has degraded the spatial resolution to the point where the distinction between  
510 individual grains cannot be made; therefore, I hypothesize that the model can't preferentially  
511 activate interstices (cf. Figure 8) for these relatively small grains. Therefore, I conclude that

512 preservation of image resolution is more important than preservation of aspect ratio for the  
513 success of SediNet models.

514

#### 515 5.4. Using and contributing to SediNet software

516 This work is fully reproducible using freely available data and code hosted on a github  
517 repository (SediNet software, 2019), which also includes further examples of how to configure  
518 SediNet for different purposes. The motivating idea behind SediNet is community  
519 development of tools for better generic information extraction from images of sediment. You  
520 can use SediNet "off-the-shelf", or other people's models, or configure it for your own  
521 purposes. You can even choose to contribute imagery back to the project, so we can build  
522 bigger and better models collaboratively. Within this package, there are several examples of  
523 different ways it can be configured for estimating categorical variables and various numbers  
524 of continuous variables.

525

526 Instructions are provided for how to run the program locally on a machine, on a cloud  
527 computer, or in a browser through jupyter notebooks. Some notebooks on cloud-hosted  
528 jupyter notebook servers are provided. Users can interact with the software in a few different  
529 ways, to 1) replicate the results of this paper, 2) explore additional provided examples, 3) use  
530 the models built for this paper on their own data, or 4) to train models on their own data for  
531 their own purposes. The program reads a comma-separated (csv) file containing a list of  
532 image file names and the quantities of interest associated with each one. The program is  
533 interacted with using a configuration file that specifies where the training images are, where  
534 the corresponding csv file is, and values for model hyper-parameters.

535

536 I show here that SediNet models can achieve high accuracy for a wide range of sediment  
537 and metrics even on small datasets. However, there are several indications that SediNet  
538 models would improve and be more generally applicable if trained with more data, or better  
539 subsets of sediment populations gleaned from large data sets. For example, the relatively  
540 large classification errors between classes 2, 3 and 4 (Figure 7) may be due to small sample  
541 size. A data set of 400-labeled images, while relatively large by standards set by previous  
542 optical granulometry techniques papers, is very small for deep learning models. Small batch  
543 sizes dictated by small sample sizes may lead to erratic training behavior such as increasing  
544 loss or large fluctuations in loss upon successive epochs, which will produce non-optimal  
545 results. Such small batch sizes may have only worked for the model to estimate sieve sizes  
546 of sand because the sample images were relatively homogeneous in grain shape and size,  
547 and the image-to-image variability in scale, perspective, brightness and contrast was minimal.  
548 In the models trained for this paper, each step of each epoch would randomly select a batch  
549 of 8 training images out of the training set; for a 16-image data set that implies two steps per  
550 epoch and for a 200 image set, 25 steps. These are unusually small for a deep learning  
551 model, but SediNet is relatively small, with thousands to tens-of-thousands of tunable  
552 parameters rather than millions to hundreds of millions of parameters that have made generic  
553 societally oriented breakthroughs in strategy games, image recognition, self-driving cars, fake  
554 video, etc. Larger, more general models will likely require much larger data sets.

555

556 Users are therefore strongly encouraged to contribute data. This is best achieved by  
557 submitting a 'pull request' to a special repository designed for collaboration (SediNet  
558 collaborative software, 2019). Initially, this is a copy of the main github repository (SediNet  
559 software, 2019) until the first user-contributed dataset. Users contribute data by forking this  
560 repository, adding their files to their forked version, optionally adding their name and contact

561 details to the list of contributors, then make a pull request. The main repository moderator  
562 reviews and merges the changes into the main repository. Code improvements may also be  
563 suggested in this way. Models can then be periodically retrained using the new data. Over  
564 time, if enough new imagery is amassed, model architecture may need to change by adding  
565 or changing convolution layers in order to uncover and exploit additional useful features  
566 extracted from the new data.

567

568 Future changes to model architecture should handle the effects of class imbalance, which is  
569 where different classes have many different numbers of examples. For example, the unequal  
570 numbers of images in population classes 2, 3 and 4 (Figure 7A, B) may be behind the  
571 misclassifications. If the training set consists of many more images of one class than another,  
572 the model may tend to classify the class better represented in the set. This might be overcome  
573 by weighting the cost function used to train the model by the relative abundance of classes  
574 in the data set. Weighted cross-entropy is a popular choice in the deep learning literature.

575

## 576 Conclusions

577 I have described a configurable machine-learning framework called SediNet for estimating  
578 either (or both) continuous and categorical variables from a photographic image of clastic  
579 sediment. To demonstrate the framework, five separate models were configured and trained,  
580 three of which for estimating various grain size metrics on both mixed and single populations  
581 of sediment, and two for classifying aspects of grain shape and population. Perhaps of most  
582 significance is that SediNet can be configured and trained to estimate equivalent sieve  
583 diameters directly from image features, without the need for area-to-mass conversion  
584 formulas and without even knowing the scale of one pixel. As such, it is the only optical  
585 granulometry method proposed to date that does not necessarily require image scaling.

586 SediNet will allow for reliable estimation of several sedimentological variables from arbitrary  
587 imagery of sediment, where grains may be either supra- or sub-pixel in scale, and where  
588 conversions between grain size measurements on different physical or statistical scales  
589 might be learnt directly from the data. The model framework should therefore find numerous  
590 application in the spatio-temporal monitoring of the grain size distribution, shape, mineralogy  
591 and other quantities of interest, of sedimentary deposits as they evolve. This study has also  
592 served to exemplify how machine learning can be a powerful tool for automated and  
593 simultaneous quantitative and qualitative measurements from the same remotely sensed  
594 imagery.

595

#### 596 Acknowledgments

597 Thanks to Evan Goldstein and Cedric John for their review comments that collectively  
598 improved the manuscript.

#### 599 References

- 600 1. Abadi M.; Agarwal A.; Barham P.; Brevdo E.; Chen Z.; and 35 others. 2015.  
601 *TensorFlow: Large-scale machine learning on heterogeneous systems*. Software  
602 available online: <https://www.tensorflow.org> (accessed on 1 July, 2019).
- 603 2. Adams J. 1979. Gravel size analysis from photographs. *Journal of the Hydraulics*  
604 *Division. American Society of Civil Engineers*: 1247–1255.
- 605 3. Austin, M.J., Masselink, G., O'Hare, T.J. and Russell, P.E. 2007. Relaxation time  
606 effects of wave ripples on tidal beaches. *Geophysical Research Letters*, 34(16).
- 607 4. Baptista, P., Cunha, T., Gama, C., and Bernardes, C. 2012. A new and practical  
608 method to obtain grain size measurements in sandy shores based on digital image  
609 acquisition and processing. *Sedimentary Geology*, 282, 294–306.

- 610 5. Barnard, P., Rubin, D., Harney, J., and Mustain, N. 2007. Field test comparison of an  
611 autocorrelation technique for determining grain size using a digital beachball camera  
612 versus traditional methods. *Sedimentary Geology*, 201 (1-2), 180–195.
- 613 6. Bergillos, R.J., Ortega-Sánchez, M., Masselink, G. and Losada, M.A. 2016. Morpho-  
614 sedimentary dynamics of a micro-tidal mixed sand and gravel beach, Playa Granada,  
615 southern Spain. *Marine Geology*, 379, 28-38.
- 616 7. Black, M., Carbonneau, P., Church, M., and Warburton, J. 2014. Mapping sub-pixel  
617 fluvial grain sizes with hyperspatial imagery. *Sedimentology*, 61 (3), 691–711.
- 618 8. Buscombe, D. 2013. Transferable wavelet method for grain-size distribution from  
619 images of sediment surfaces and thin sections, and other natural granular patterns.  
620 *Sedimentology*, 60 (7), 1709–1732.
- 621 9. Buscombe, D., and Carini, R. J. 2019. A data-driven approach to classifying wave  
622 breaking in infrared imagery. *Remote Sensing*, 11 (7), 859.
- 623 10. Buscombe, D., Carini, R., Harrison, S., Chickadel, C., and Warrick, J. 2019. Optical  
624 wave gauging with deep neural networks. *Coastal Engineering*, in review July 2019
- 625 11. Buscombe, D. and Masselink, G. 2006. Concepts in gravel beach dynamics. *Earth-  
626 Science Reviews*, 79 (1-2), 33-52.
- 627 12. Buscombe, D. and Ritchie, A. 2018. Landscape classification with deep neural  
628 networks. *Geosciences*, 8 (7), 244.
- 629 13. Buscombe, D., and Rubin, D. M. 2012a. Advances in the simulation and automated  
630 measurement of well-sorted granular material: 1. Simulation. *Journal of Geophysical  
631 Research: Earth Surface*, 117 (F2).
- 632 14. Buscombe, D., and Rubin, D. M. 2012b. Advances in the simulation and automated  
633 measurement of well-sorted granular material: 2. Direct measures of particle  
634 properties. *Journal of Geophysical Research: Earth Surface*, 117 (F2).

- 635 15. Buscombe, D., Rubin, D. M., Lacy, J. R., Storlazzi, C. D., Hatcher, G., Chezar, H., and  
636 Sherwood, C. R. 2014. Autonomous bed-sediment imaging-systems for revealing  
637 temporal variability of grain size. *Limnology and Oceanography: Methods*, 12 (6), 390–  
638 406.
- 639 16. Buscombe, D., Rubin, D., and Warrick, J. 2010. A universal approximation of grain  
640 size from images of noncohesive sediment. *Journal of Geophysical Research: Earth  
641 Surface*, 115 (F2).
- 642 17. Carbonneau, P., Bizzi, S., and Marchetti, G. 2018. Robotic photosieving from low-cost  
643 multicopter sUAS: A proof-of-concept. *Earth Surface Processes and Landforms*, 43 (5),  
644 1160–1166.
- 645 18. Carbonneau PE, Bergeron N, Lane SN. 2005a. Automated grain size measurements  
646 from airborne remote sensing for long profile measurements of fluvial grain sizes.  
647 *Water Resources Research*: 41.
- 648 19. Carbonneau PE, Bergeron NE, Lane SN. 2005b. Texture-based image segmentation  
649 applied to the quantification of superficial sand in salmonid river gravels. *Earth Surface  
650 Processes and Landforms* 30: 121–127.
- 651 20. Carbonneau, P.E., Lane, S.N. and Bergeron, N.E. 2004. Catchment-scale mapping of  
652 surface grain size in gravel bed rivers using airborne digital imagery. *Water Resources  
653 Research*, 40 (7).
- 654 21. Chollet, F. 2015. Keras. Software available online: <https://keras.io> (accessed on 1 July,  
655 2019).
- 656 22. Cheng, Z. and Liu, H. 2015. Digital grain-size analysis based on autocorrelation  
657 algorithm. *Sedimentary Geology*, 327, 21-31.

- 658 23. Cuttler, M. V., Lowe, R. J., Falter, J. L., and Buscombe, D. 2017. Estimating the settling  
659 velocity of bioclastic sediment using common grain-size analysis techniques.  
660 *Sedimentology*, 64 (4), 987–1004.
- 661 24. Detert M, Weitbrecht V. 2012. Automatic object detection to analyze the geometry of  
662 gravel grains – a free stand-alone tool. In River Flow, 2012, Munoz RM (ed). CRC  
663 Press: London; 595–600.
- 664 25. Diplas, P. and Fripp, J.B. 1992. Properties of various sediment sampling procedures.  
665 *Journal of Hydraulic Engineering*, 118, 955–970.
- 666 26. Diplas, P., Kuhnle, R.A., Gray, J.R., Glysson, G.D. and Edwards, T.E. 2008.  
667 Sediment transport measurements. In: Sedimentation Engineering: Processes,  
668 Measurements, Modeling, and Practice (Ed. M.H. Garcia), 110, 307–353. American  
669 Society of Civil Engineering Manuals and Reports on Engineering Practice, Reston,  
670 VA
- 671 27. Diplas, P. and Sutherland, A.J. 1988. Sampling techniques for gravel sized sediments.  
672 *Journal of Hydraulic Engineering*, 114, 484–501.
- 673 28. Dugdale, S. J., Carbonneau, P. E., and Campbell, D. 2010. Aerial photosieving of  
674 exposed gravel bars for the rapid calibration of airborne grain size maps. *Earth Surface  
675 Processes and Landforms*, 35 (6), 627–639.
- 676 29. Franklin, S., and Wulder, M. 2002. Remote sensing methods in medium spatial  
677 resolution satellite data land cover classification of large areas. *Progress in Physical  
678 Geography*, 26 (2), 173–205.
- 679 30. Gomez, C., and Purdie, H. 2016. UAV-based photogrammetry and geocomputing for  
680 hazards and disaster risk monitoring—A review. *Geoenvironmental Disasters*, 3 (1), 23.
- 681 31. Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. 2016. *Deep learning* (Vol. 1).  
682 MIT press Cambridge.
- 683 32. Graham DJ, Reid I, Rice SP. 2005. Automated sizing of coarse-grained sediments:  
684 image-processing procedures. *Mathematical Geology* 37: 1–28.

- 685 33. Graham, D., Rollet, A., Rice, S., and Piegay, H. 2012. Conversions of surface grain-  
686 size samples collected and recorded using different procedures. *Journal of Hydraulic*  
687 *Engineering*, 138 (10), 839–849.
- 688 34. Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network  
689 training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- 690 35. Jiang, G.-Q., Xu, J., and Wei, J. 2018. A deep learning algorithm of neural network for  
691 the parameterization of typhoon-ocean feedback in typhoon forecast models.  
692 *Geophysical Research Letters*, 45 (8), 3706–3716.
- 693 36. Kellerhals, R., and Bray, D. I. 1971. Sampling procedures for coarse fluvial sediments.  
694 *Journal of the Hydraulics Division*, 97 (8), 1165–1180.
- 695 37. Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization.  
696 *arXiv preprint arXiv:1412.6980*.
- 697 38. Legleiter, C.J., Stegman, T.K. and Overstreet, B.T. 2016. Spectrally based mapping  
698 of riverbed composition. *Geomorphology*, 264, 61-79.
- 699 39. Linville, L., Pankow, K., and Draelos, T. 2019. Deep learning models augment analyst  
700 decisions for event discrimination. *Geophysical Research Letters*, 46 (7), 3643–3651.
- 701 40. Luo, J. Y., Irisson, J.-O., Graham, B., Guigand, C., Sarafraz, A., Mader, C., and  
702 Cowen, R. K. 2018. Automated plankton image analysis using convolutional neural  
703 networks. *Limnology and Oceanography: Methods*, 16 (12), 814–827.
- 704 41. Masteller, C.C. and Finnegan, N.J. 2017. Interplay between grain protrusion and  
705 sediment entrainment in an experimental flume. *Journal of Geophysical Research:*  
706 *Earth Surface*, 122 (1), 274-289.
- 707 42. Michaelides, K., Hollings, R., Singer, M.B., Nichols, M.H. and Nearing, M.A. 2018.  
708 Spatial and temporal analysis of hillslope–channel coupling and implications for the

- 709 longitudinal profile in a dryland basin. *Earth Surface Processes and Landforms*, 43 (8),  
710 1608-1621.
- 711 43. Montgomery, D. R., Panfil, M. S., and Hayes, S. K. 1999. Channel-bed mobility  
712 response to extreme sediment loading at Mount Pinatubo. *Geology*, 27 (3), 271–274.
- 713 44. Mulder, V., De Bruin, S., Schaepman, M., and Mayr, T. 2011. The use of remote  
714 sensing in soil and terrain mapping: A review. *Geoderma*, 162 (1-2), 1–19.
- 715 45. Nield, J.M., Wiggs, G.F. and Squirrell, R.S. 2011. Aeolian sand strip mobility and  
716 protodune development on a drying beach: examining surface moisture and surface  
717 roughness patterns measured by terrestrial laser scanning. *Earth Surface Processes  
718 and Landforms*, 36 (4), 513-522.
- 719 46. Novak-Szabo, T., Sipos, A.A., Shaw, S., Bertoni, D., Pozzebon, A., Grottoli, E.,  
720 Jerolmack, D. J. 2018. Universal characteristics of particle shape evolution by bed-  
721 load chipping. *Science Advances*, 4 (3).
- 722 47. Oh, C., Ham, B., Kim, H., Hilton, A. and Sohn, K. 2019. OCEAN: Object-centric  
723 arranging network for self-supervised visual representations learning. *Expert Systems  
724 with Applications*, 125, 281-292.
- 725 48. Paterson, G. A., and Heslop, D. 2015. New methods for unmixing sediment grain size  
726 data. *Geochemistry, Geophysics, Geosystems*, 16 (12), 4494–4506.
- 727 49. Pfeiffer, A. M., Finnegan, N. J., and Willenbring, J. K. 2017. Sediment supply controls  
728 equilibrium channel geometry in gravel rivers. *Proceedings of the National Academy  
729 of Sciences*, 114 (13), 3346–3351.
- 730 50. Prodger, S., Russell, P., Davidson, M., Miles, J. and Scott, T. 2016. Understanding  
731 and predicting the temporal variability of sediment grain size characteristics on high-  
732 energy beaches. *Marine Geology*, 376, 109-117.

- 733 51. Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N.  
734 2019. Deep learning and process understanding for data-driven Earth system science.  
735 *Nature*, 566 (7743), 195.
- 736 52. Rickenmann, D., and Recking, A. 2011. Evaluation of flow resistance in gravel-bed  
737 rivers through a large field data set. *Water Resources Research*, 47 (7).
- 738 53. Rubin, D. 2004. A simple autocorrelation algorithm for determining grain size from  
739 digital images of sediment. *Journal of Sedimentary Research*, 74 (1), 160–165.
- 740 54. Rubin, D., Buscombe, D., Wright, S., Topping, D., Grams, P., and Schmidt, J. 2019.  
741 What grain size reveals about suspended-sand transport in the Colorado River in  
742 Grand Canyon. *Journal of Geophysical Research - Earth Surface*, in review.
- 743 55. Rubin, D., Chezar, H., Harney, J., Topping, D., Melis, T., and Sherwood, C. 2007.  
744 Underwater microscope for measuring spatial and temporal changes in bed-sediment  
745 grain size. *Sedimentary Geology*, 202 (3), 402–408
- 746 56. Sedinet v.1.0 online and data repository software. 2019.  
747 <https://github.com/MARDAScience/SediNet>. <http://doi.org/10.5281/zenodo.3466045>
- 748 57. Sedinet collaborative software and data repository. 2019.  
749 <https://github.com/MARDAScience/SediNet-Contrib>.
- 750 58. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D. and Batra, D. 2017.  
751 Grad-cam: Visual explanations from deep networks via gradient-based localization. In:  
752 *Proceedings of the IEEE International Conference on Computer Vision*, 618-626.
- 753 59. Shoji, D., Noguchi, R., Otsuki, S. and Hino, H. 2018. Classification of volcanic ash  
754 particles using a convolutional neural network and probability. *Scientific Reports*, 8 (1),  
755 8111.
- 756 60. Smith, M., and Pain, C. 2009. Applications of remote sensing in geomorphology.  
757 *Progress in Physical Geography*, 33 (4), 568–582.

- 758 61. Smith, M.E., Werner, S.H., Buscombe, D., Finnegan, N.J., Sumner, E.J. and Mueller,  
759 E.R. 2018. Seeking the shore: Evidence for active submarine canyon head incision  
760 due to coarse sediment supply and focusing of wave energy. *Geophysical Research*  
761 *Letters*, 45 (22), 12-403.
- 762 62. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. 2014.  
763 Dropout: a simple way to prevent neural networks from overfitting. *The Journal of*  
764 *Machine Learning Research*, 15 (1), 1929–1958.
- 765 63. Sternberg, R.W., Berhane, I. and Ogston, A.S. 1999. Measurement of size and settling  
766 velocity of suspended aggregates on the northern California continental shelf. *Marine*  
767 *Geology*, 154 (1-4), 43-53.
- 768 64. Turner, I. L., Harley, M. D., and Drummond, C. D. 2016. UAVs for coastal surveying.  
769 *Coastal Engineering*, 114, 19–24.
- 770 65. Turner, I.L., Russell, P.E. and Butt, T. 2008. Measurement of wave-by-wave bed-  
771 levels in the swash zone. *Coastal Engineering*, 55 (12), 1237-1242.
- 772 66. Viles, H. 2016. Technology and geomorphology: Are improvements in data collection  
773 techniques transforming geomorphic science? *Geomorphology*, 270, 121–133.
- 774 67. Warrick, J.A., Rubin, D.M., Ruggiero, P., Harney, J.N., Draut, A.E. and Buscombe, D.  
775 2009. Cobble Cam: Grain-size measurements of sand to boulder from digital  
776 photographs and autocorrelation analyses. *Earth Surface Processes and Landforms*,  
777 34 (13), 1811-1821.
- 778 68. Wheatcroft, R.A. and Borgeld, J.C. 2000. Oceanic flood deposits on the northern  
779 California shelf: large-scale distribution and small-scale physical properties.  
780 *Continental Shelf Research*, 20 (16), 2163-2190.

- 781 69. Williams, R.D., Brasington, J., Vericat, D. and Hicks, D.M. 2014. Hyperscale terrain  
782 modelling of braided rivers: fusing mobile terrestrial laser scanning and optical  
783 bathymetric mapping. *Earth Surface Processes and Landforms*, 39 (2), 167-183.
- 784 70. Woodget, A., and Austrums, R. 2017. Subaerial gravel size measurement using topo-  
785 graphic data derived from a UAV-SfM approach. *Earth Surface Processes and*  
786 *Landforms*, 42 (9), 1434–1443.
- 787 71. Woodget, A., Fyffe, C., and Carbonneau, P. 2018. From manned to unmanned aircraft:  
788 Adapting airborne particle size mapping methodologies to the characteristics of sUAS  
789 and SfM. *Earth Surface Processes and Landforms*, 43 (4), 857–870.