



Automated detecting, segmenting and measuring of grains in images of fluvial sediments: The potential for large and precise data from specialist deep learning models and transfer learning

David Mair¹  | Guillaume Witz² | Ariel Henrique Do Prado¹ |
 Philippos Garefalakis¹  | Fritz Schlunegger¹

¹Institute of Geological Sciences, University of Bern, Bern, Switzerland

²Data Science Lab, University of Bern, Bern, Switzerland

Correspondence

David Mair, Institute of Geological Sciences, University of Bern, Baltzerstrasse 1+3, 3012 Bern, Switzerland.

Email: david.mair@unibe.ch

Abstract

The size of sedimentary particles in rivers bears information on the sediment entrainment or deposition mechanisms and the hydraulic conditions controlling them. However, collecting such data from coarse-grained sediments is work intensive, both in the field and remotely. Therefore, attention has turned to machine learning models to improve the data acquisition. Despite their success, current methods need large quantities of data and yield results limited to a few percentile values of grain size datasets, often additionally affected by a systematic bias. In most cases, the root of these limitations is the challenge of accurately segmenting grains. Here, we present a new approach to improve the segmentation of individual grains based on the capacity of transfer learning in convolutional neural networks. Specifically, we re-train a state-of-the-art model for cell segmentation in biomedical images to find and segment coarse-grained particles in images of fluvial sediments. Our results show that the performance in the segmentation tasks can be directly transferred to images of fluvial sediments and that our re-trained models outperform existing methods. We document that our results are achievable with only 10%–20% of the data needed for training other deep learning models designed to measure the size of fluvial sediments. Moreover, we find that traits in our data control the segmentation performance. This enables data-driven approaches to create specialist segmentation models. Additionally, comparing our automatically obtained datasets with the results retrieved from image and field-based surveys confirms that improvements in segmentation are directly leading to more precise and more accurate grain size data even if data collection occurs in images taken at different conditions. Finally, we release a software package, the trained models and our data. The goal is to offer a tool to efficiently segment and measure grains in sediment images in an automated way, which can be adapted to different settings.

KEYWORDS

deep learning, fluvial gravel, grain size, image segmentation, neural network, sediment analysis

1 | INTRODUCTION

Data on the size of the sediment transported by rivers, both in modern and ancient systems, are of crucial importance to understand the

mechanisms, hydraulic conditions and grain-to-grain interactions during sediment transport in fluvial systems (e.g., Attal et al., 2015; Dunne & Jerolmack, 2018; Piégay et al., 2020; Whittaker, Attal, & Allen, 2010). In addition, grain size is critical to deciphering climate,

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Earth Surface Processes and Landforms* published by John Wiley & Sons Ltd.

tectonic and supply signals preserved in the stratigraphic record of fluvial sediment routing systems (e.g., Allen et al., 2017; Castelltort & Van Den Driessche, 2003; Schlunegger & Norton, 2015; Tofelde et al., 2021). Standard field methods have been developed in the past years to measure the grain size of sediment transported in active rivers (e.g., Bunte & Abt, 2001). However, these methods are costly and yield limited or potentially biased data, which led to the development of approaches for measuring grain sizes in images (e.g., Butler, Lane, & Chandler, 2001; Carbonneau, Lane, & Bergeron, 2004). These image-based methods have been constantly improved over the last years, yielding a variety of approaches, ranging from manual annotation (e.g., Sulaiman et al., 2014), semi-automated segmentation (e.g., Detert & Weitbrecht, 2012; Purinton & Bookhagen, 2019) to texture-based percentile predictions (e.g., Buscombe, 2013). Although these methods allowed for a faster and remote measurement of a larger number of grains, they still need a calibration in the field for texture-based approaches, or they require a manual correction of individual grains for segmentation-based methods. In addition, both tend to systematically over-/underestimate the sizes of grains (e.g., Chardon et al., 2020; Chardon, Piasny, & Schmitt, 2022; Mair et al., 2022a). Therefore, and most recently, attention turned to deep neural networks, with the aim either to improve the segmentation in images (e.g., Chen et al., 2023; Chen, Hassan, & Fu, 2022; Mörtl et al., 2022; Soloy et al., 2020) or to directly predict percentile values of a grain size distribution (e.g., Buscombe, 2020; Lang et al., 2021). The main aims of all these works were to automate the measurements, improve the reproducibility and scalability, and to increase the number of observations. Although deep learning did improve the segmentation for some data and settings, challenges have remained (e.g., Chen et al., 2023; Chen, Hassan, & Fu, 2022). In addition, current methods have limited ability to adapt to new, previously unseen data. This, in turn, limits the uncertainty estimation and thus reduces the interpretability of the results (e.g., Reichstein et al., 2019). In particular, the large variety of information in images, different camera properties and the visual complexity of natural photographs (e.g., Figure 1) pose challenges to current deep learning-based models. Such challenges might even prevent these neural network-based models from producing meaningful information for some data not seen during training (e.g., Szegedy et al., 2014; Zech et al., 2018). Thus, additional annotated data may be required to apply these methods in new settings (e.g., Sun et al., 2017). However, for current models used for grain size measurements, the large amount of data required for training (e.g., >125 000 or >180 000 individually annotated grains for Chen, Hassan, & Fu, 2022, and Lang et al., 2021, respectively) imposes an often prohibitive cost on re-training.

In recent years, tools to estimate the grain size of fluvial pebbles from 3D point clouds have been developed in parallel with ongoing improvements in image-based methods. For example, some authors tried to segment individual grains by ellipsoidal fitting (e.g., Steer et al., 2022), whereas others attempted to infer a size distribution from the roughness of a point cloud (e.g., Woodget & Austrums, 2017). Although segmentation-based methods applied to point clouds can yield valuable additional information, such as the 3D orientation or the 3D shape, the current generation of such methods cannot be readily applied to every setting for several reasons. First, the related acquisition of data can be more complex, and the subsequent data processing is time-consuming and thus expensive. This is mainly due to the technical requirements to conduct a LIDAR (Light

Detection and Ranging) or RTK (real-time kinematic positioning) survey and to post-process the data. Second, such methods are not well suited to fit the geometry of angular sedimentary particles, because of inherent geometric restrictions that are related to the corresponding methods. Accordingly, despite many recent advances in measuring individual grain properties of fluvial pebbles, related surveys based on 3D point clouds and 2D images have remained a challenge. As such, there is an increasing need for a more accurate segmentation across different data.

Here, we present a new approach for improving instance segmentation of individual grains in images, employing the capability of transfer learning in deep neural networks (e.g., Lu et al., 2015; Yosinski et al., 2014). This allows us to adapt existing models for tasks similar to their original purpose. Specifically, we use the Python-based open-source tool Cellpose (Stringer et al., 2021), a state-of-the-art deep learning model for detecting and segmenting cells and nuclei in biomedical images. We adapt this model to find and delineate coarse-grained pebbles in images of fluvial gravel. Our underlying rationale is that sedimentary particles such as fluvial pebbles are geometrically similar to cell nuclei. Our results indicate that these models can indeed be retrained for segmenting sedimentary particles in images.

The resulting models, either re-trained from models that themselves are trained for nuclei segmentation or trained from scratch, vastly outperform existing models that have been proposed for the segmentation of fluvial pebbles in all datasets we tested in this study. Notably, we will document that although our approach requires an order of magnitude smaller dataset size for training than other methods (e.g., Chen, Hassan, & Fu, 2022; Lang et al., 2021), it yields better results. Furthermore, the models' flexibility and accessibility, as well as the ability to re-train them rapidly, allow us to generate task- or image-type-specific models, as we did here with different datasets. Such dedicated specialist models can yield substantially better results for these specialized datasets than more generalist models trained on larger datasets, which generally aim to segment many different types of objects (e.g., He et al., 2020; Kirillov et al., 2023). In line with the approach of Pachitariu and Stringer (2022), we additionally propose an interactive workflow to create specialized and high-performing models in a short time and from relatively small datasets. To facilitate access to these powerful instance segmentation models, we built an open-source software library, ImageGrains, which allows for (i) an easy use of the Cellpose models we trained, (ii) a straightforward training of custom models and (iii) streamlined grain size measurements.

2 | METHODS

We employed an existing deep neural network architecture, Cellpose (Stringer et al., 2021; Pachitariu & Stringer, 2022; v.2.1.1), as a backbone segmentation model within our ImageGrains software (see Data S1). This is realized with the goal to segment and measure sedimentary particles in curated image datasets (Figure 2). In the following sections, we provide background information on the datasets (Section 2.1) and briefly describe the deep learning architecture used for segmentation and its training (Section 2.2). Next, we explain how we evaluate the model's segmentation performance (Section 2.3) before describing how the particle sizes were measured (Section 2.4).



FIGURE 1 Examples of images taken from fluvial gravel bars used for this study. (a) Images from Canadian and Swiss rivers (Chen, Hassan, & Fu, 2021; Mair et al., 2022b; respectively) are used for the generalized APF ('all pebbles fluvial') dataset. (b) Uncrewed aerial vehicle (UAV) images from one site at the Swiss Sense River with homogenous image content and conditions are used in the specialized S1 dataset. (c) Images of a vertically oriented outcrop in the Finsterhennen gravel pit (Garefalakis et al., 2023) show gravel with fines and sand matrix between coarser grains. In all panels, the upper half shows the RGB image, whereas the lower half depicts the single-channel greyscale image used for training and inference.

2.1 | Image data

We compiled a diverse dataset of fluvial sediment images displaying coarse-grained (>2 mm) fluvial sediment, so-called pebbles, into three basic datasets (Table S1), which are described below (see Figure 1 for an overview). For each image, the same operator manually annotated square subset tiles with the size of 512 pixels to generate the ground truth for training the model. This was done using the LABKIT plugin (Arzt et al., 2022) for FIJI (Schindelin et al., 2012). Image sub-setting was accomplished semi-randomly to avoid overlaps between the tiles while still capturing the visual complexity of the image. This led to a varying number of tiles for each image (Table S1).

First, for the APF data ('all pebbles fluvial'), we selected 20 images taken from fluvial gravel with handheld cameras and with uncrewed

aerial vehicles (UAVs; Figure 1a). The images were taken from several sites along Canadian ($n_{\text{images}} = 4$; Chen, Hassan, & Fu, 2021; Brayshaw, 2012) and Swiss ($n_{\text{images}} = 9$; Mair et al., 2022b) rivers. Here, the goal was to generate a dataset of fluvial pebbles that is as general as possible. Note that we created new labels for images used by Chen, Hassan and Fu (2022) to annotate grain boundaries more precisely. We complemented the APF dataset with seven additional photos from six sites along Swiss rivers (Figure 1a), exhibiting challenging conditions (e.g., raindrops or wet gravel). These additional images were taken with a handheld camera (Litty & Schlunegger, 2017) and UAVs. We included these seven images to broaden the model's applicability to more general cases and increase the robustness of the model prediction. We annotated 56 tiles for the APF dataset, of which 47 were used for training, and nine tiles were kept as a test set. The second set, S1

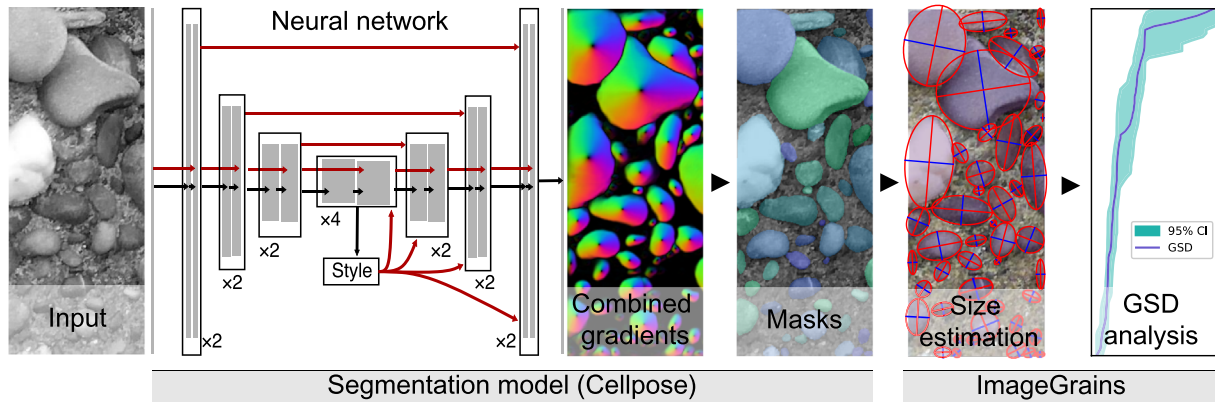


FIGURE 2 Overview of our workflow. We used the deep neural network model Cellpose (Stringer et al., 2021) as backbone model to segment sedimentary particles in images of fluvial sediments. On the resulting masks, we measured the sizes of individual grains, from which we compiled grain size distributions (GSD) with percentile-based uncertainty estimations (see Section 2.4 and Data S1 for details). The model architecture of Cellpose is reproduced from Stringer et al. (2021). Please note that the numbers in the architecture representation indicate the number of residual blocks in the layer, where each block is composed of two convolutions of the size 3×3 . Thus, they do not state the convolution size.

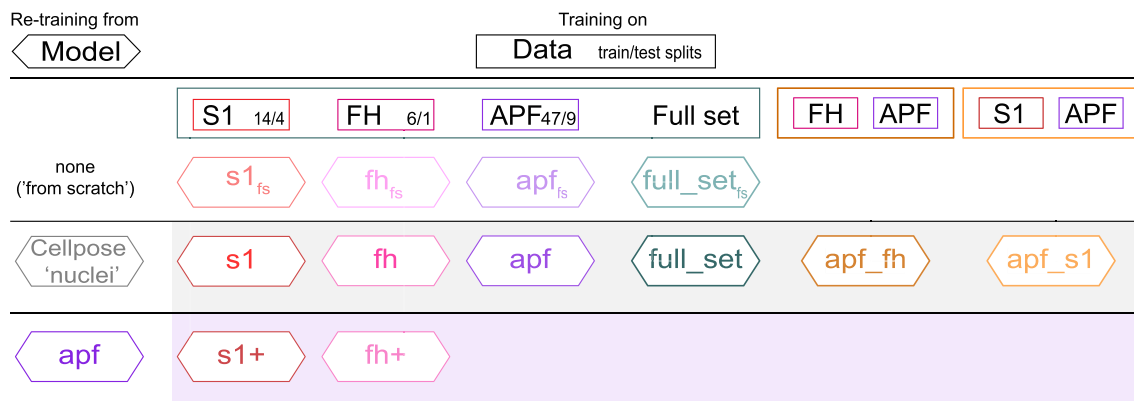


FIGURE 3 Overview of dataset splits and starting model weights used for training the segmentation models in this study. We refer to the main text for data description (Section 2.1) and model training (Section 2.2). The *nuclei* model refers to the pre-trained Cellpose model on images of cell nuclei of Stringer et al. (2021).

(Figure 1b), comprises five nadir images, all acquired with a UAV at one site in the Sense River (Switzerland) under homogenous light conditions. Set S1 contains 18 annotated tiles (14 for training and four as a test). The third dataset, referred to as FH ('Finsterhennen'), consists of seven vertically orientated images taken with a handheld camera from fluvial sediment in the Finsterhennen gravel pit in Switzerland (Figure 1c; Garefalakis et al., 2023). These images differ from the APF and S1 sets in their orientation and depositional nature showing gravels embedded in a fine-grained (<2 mm grain size) matrix. This material fills the interstices and sometimes covers parts of the clasts. We annotated seven tiles, one tile per image, for set FH, which we split by six to one for training and testing. From these annotated image tiles (see Figure S2), we trained a collection of segmentation models on different data splits (Figure 3). We did so to (i) assess the segmentation performance of the models as we varied the grade of specialization on specific data (Stringer et al., 2021) and to (ii) test how the training strategy influences the segmentation performance.

All UAV-derived images were aligned, undistorted and scaled using the structure from motion approach with Agisoft Metashape (v1.6 Pro), a standard software for SfM/MVS (Structure from Motion/Multi-View Stereo) photogrammetry. Most of such processed photos are undistorted nadir images, except one image from the Guerbe River, for

which we used an orthophoto mosaic. For some photos, the UAV image acquisition was accomplished in the raw (DNG) format, whereas others, for example, all S1 images, were acquired in a pre-processed JPEG format (for details, see Mair et al., 2022a). However, after the photogrammetric alignment, all images were converted to the JPEG format. Referencing was accomplished through ground control points for all UAV-derived imagery, measured with a Leica Zeno GG04 plus GNSS antenna and the real-time online Swipos-GIS/GEO RTK correction. This was accomplished at a precision of 2 cm (horizontally) and 4 cm (vertically; Swisstopo, 2022). We report key uncertainties of the photogrammetric models, which we used for the modelling of the grain size uncertainties (see Section 2.4 below and Table S2), therein following Mair et al. (2022a). Ultimately, we used two orthoimage mosaics generated from UAV imagery of bars along the Kander and Sense River (sites for which images are included in the APF and S1 datasets, respectively) for fully automated grain segmentation and size measurement.

2.2 | Segmentation model and training

Cellpose is a deep learning model designed for efficient and accurate segmentation of cells or cell nuclei in biomedical images, which at the

time of this study in early 2023 was one of the best-performing models for that specific task (for all details, we refer to Stringer et al., 2021; Pachitariu & Stringer, 2022). For our study, it was trained with single-channel, greyscale images and annotated ground truth of the objects to segment, that is, cells or cell nuclei, and rock pebbles (Figure S2). The model itself (Figure 2) is a deep neural network in the style of U-net (Ronneberger, Fischer, & Brox, 2015), which uses residual blocks (He et al., 2016). Similar to the original U-net, the model consists of a four-level downsampling and upsampling pass (Figure 2; Stringer et al., 2021), where the convolutions at each level can be related to different spatial scales. Between the two passes, the model computes a 256-dimensional vector that represents the style of each input image, for which a global average pooling of the convolutional maps at the smallest scale is applied (Gatys, Ecker, & Bethge, 2016; Karras et al., 2021). This style vector is used as input in the upsampling pass, influencing the segmentation results. Therefore, it can be understood as a numerical summary of an image (see Gatys, Ecker, & Bethge, 2016). This image style representation allows the clustering of data according to their respective segmentation style (Pachitariu & Stringer, 2022). The neural network itself is employed to predict the gradients of a flow-vector field, which is simulated with a heat diffusion equation. These flow vectors are tailored to find the centre of each object. In detail, the output consists of predictions of horizontal and vertical flow gradients and the probability of a pixel being inside or outside of a region of interest (ROI). Through gradient tracking (Li et al., 2008), any routing of a pixel to such a centre can be identified, thereby allowing for the segmentation of the object in the ROI and for mapping its precise outline. For our study, we used the standard architecture and default settings of Cellpose. These include a flow threshold of 0.4, a mask threshold of 0, a mean object diameter of 17 pixels during training for re-trained models and a scale-dependent resampling.

The original Cellpose models were created to segment cellular images. They were trained mainly on various microscopy images of cells and a few images of repeated objects, for example, fish scales, vegetables, or rocks. However, the model, which we use as a base for our re-training (see Figure 3 for an overview), was exclusively trained on annotated data of cell nuclei in 1139 images of various sources (Caicedo et al., 2019; Coelho, Shariff, & Murphy, 2009; Kumar et al., 2020; Stringer et al., 2021). All models presented in the following sections were re-trained from the weights of this *nuclei* model unless otherwise indicated with a suffix *_fs*, which means they were trained from scratch. Without exceptions, all our models were trained on the respective training splits (Table S1). For this, we employed stochastic gradient descent, a learning rate of 0.2 and a weight decay of 10^{-5} , and we validated every 10th epoch with test splits. Following the default schedule, the learning rate was annealed from zero to 0.2 over the first 10 epochs. Similarly, for the last 100 epochs, the learning rate was reduced by a factor of 2 every 10 epochs. The training occurred on batches of eight single-channel greyscale images with default image augmentation, which included random rotations, scaling and translations. The default configuration for training uses the L2 loss function. This training schedule was applied to all models, irrespective of the data subset or whether the models were re-trained or trained from scratch. When re-training from scratch, we set the mean diameter of the object to 17 pixels (i.e., the same as for the *nuclei* model used for re-training), which is used by the algorithm to re-scale every image during training. We assessed the effects of applying changes in

the training configuration where, for example, the number of epochs, the learning rate and the scale range are modified and where images were re-scaled during training. We also explored the effect of where the images are re-scaled and no minimum object sizes are applied (see Data S5). We found that the default setting produced the overall best segmentation results for our datasets and models. An exception is the number of epochs, which we increased accordingly to 1000 (from 500). All our training was accomplished on a stand-alone desktop PC with an Nvidia RTX 3070 GPU with 8 GB RAM.

2.3 | Quantification and optimization of segmentation performance

For assessing the segmentation performance, we employ the built-in method of Stringer et al. (2021), which is partly based on the approach of Schmidt et al. (2018), which matches the model predictions to the most similar ground-truth annotation. This is done by calculating the intersection over union (IoU) metric, which is the intersection of the ground truth and predicted mask, or bounding box, indicated as the area over the union of the same areas. We calculated an IoU for each predicted grain. Next, all predictions are evaluated for different IoU thresholds to calculate the number of valid matches (true positives; TP), the number of predictions without ground truth masks (false positives; FP) and the number of ground truth objects with no valid matches (false negatives; FN). Here, at higher IoU threshold values, predictions must resemble the corresponding ground truth object more closely than at lower IoUs. Thus, for all predictions and for each IoU threshold, the average precision (AP) can be calculated through

$$AP = \frac{TP}{TP + FP + FN} \quad (1)$$

Similarly, by taking the average of image AP values, we calculated the average precision at a specific IoU threshold for an entire dataset. The mean average precision (mAP) is then the corresponding average over several IoU values, either for an image or across a set of images. Here, all reported mAP values were calculated for IoUs from 0.5 to 0.9, evaluated at steps of 0.05. These metrics are standard for evaluating results where objects are detected in images in general (for more details see, e.g., Lin et al., 2014; Rezatofighi et al., 2019; Padilla, Netto, & da Silva, 2020; and references therein), or in biomedical images in particular (e.g., Caicedo et al., 2019; Schmidt et al., 2018). We calculated all AP values at all IoUs on masks that were already filtered to be larger than a cut-off value of 12 pixels (see Section 2.4 below).

We then compared the segmentation performance to existing methods. To do so, we compared our grain masks with the predictions from two other methods for all our datasets. The methods used are the neural network- and watershed-based model GrainID from Chen, Hassan and Fu (2022) and the PebbleCounts tool (Purinton & Bookhagen, 2019), which performs a classical edge detection in images. GrainID is a supervised machine learning model trained on over 125 000 annotated gravel instances in pictures taken from river sediments in Canada and China, as well as in images obtained from flume experiments (Chen, Hassan, & Fu, 2021). We used the model as

trained and published by the authors. For the second benchmark, we used `PebbleCountsAuto`, which is the automated version of `PebbleCounts`. We used the tool's default settings except for disabling the masking of fine-grained sediments. In addition, we employed this software without any manual adjustments (Purinton & Bookhagen, 2021). We pre-calculated the Otsu threshold for every image with an OpenCV routine, which we then used as input for `PebbleCounts` to improve the detection of grains. We acknowledge that `PebbleCounts` was designed to find grains in larger images and that the tool is intended to interactively find grains (which can be clearly segmented by edge detection) and not necessarily to automatically segment all grains in an image.

Finally, we inferred the representation of the image style for each of our 81 image tiles from the neural network (see Section 2.2 above). We did so using our most generalist model (`full_set`). We used these 256-dimensional vectors for image style clustering in a data-driven effort to find classes of similar image types. Here, the goal was to train style-specific models with potentially improved segmentation capability. For this, we followed the general method of Pachitariu and Stringer (2022) by using the Leiden algorithm with 66 neighbours. We employed a resolution of 0.8 for higher-dimension clustering (Traag, Waltman, & van Eck, 2019) and a t-distributed stochastic neighbour embedding (t-SNE) with dimensionality reduction for visualization (Poličar, Stražar, & Zupan, 2019; Van der Maaten & Hinton, 2008). We then trained models for each cluster, for which we used the same train and test tiles that were reorganized according to their respective style classes.

2.4 | Grain size and uncertainty

The segmentation model, which was trained on the indicated data split (see Figure 3) for inference, returned delineated grains for each dataset. All predicted masks presented in the following sections are generated with default settings, including a minimum object diameter of 15 pixels ('`min_size`'). For inference, we rescaled each image with the mean diameter of grains of the respective dataset, calculated with a circular approximation. We first performed a simple ellipsoidal approximation for each grain candidate to convert segmented ROI masks into grain size estimations. We then excluded grains, for which the minor axis was <12 pixels. The same filters were applied to both the predictions from the segmentation models and the ground truth masks. For each grain, we thus used (i) the minor and major axes of the approximated ellipses and (ii) the longest distance between points on the convex hull along with the largest distance perpendicular to it, as proxies for the a-axis (or longest visible axis) and the b-axis, respectively (or shortest visible axis). The results of all length measurements are converted from pixels to length units through the image-specific pixel resolution (Table S1). By default, the results are returned as one output file for every input image. This allows extracting the grain size distributions (GSDs) for sub-regions or a combination of multiple images.

We modelled the grain size uncertainties for individual percentile values of GSDs (Eaton, Moore, & MacKenzie, 2019; Mair et al., 2022a) with several approaches to account for the uncertainties of the varying input data types. Here, we followed the strategy of Mair et al. (2022a) and employed bootstrapping with replacement for

any axis value (ax_i) of a GSD. We did so to quantify the percentile uncertainty for grain sizes measured in the field and for grain sizes reported in image pixels before conversion to length units. However, any such percentile uncertainty only accounts for the variation introduced by the number of grains, that is, the counting statistics. To account for further uncertainty introduced by measurements in images, we combined this bootstrapping with a one-dimensional error modelling. We accomplished this by randomizing each resampled axis (ax_i) with two components, which consider the errors on the length (ϵ_{length}) and the scale (ϵ_{scale}):

$$ax_{sim} = (ax_i + \epsilon_{length}) * \epsilon_{scale}. \quad (2)$$

The length error represents the measurement error along the axis length. By default, it is implemented in the randomization as a normal distribution centred on zero, with a standard deviation set to 2 times the average length of a pixel's diameter. The scale error is a dimensionless factor that represents the uncertainty introduced by imperfect scaling, that is, through estimating the image resolution of an image. It is implemented, also by default, in the randomization as a normal distribution centred on 1. As standard deviation, it has the fractional uncertainty on the principal distance of the image; if no information is available on the uncertainty of the principal distance, an uncertainty of 10% (with a corresponding value of 0.1) is considered. We used this percentile uncertainty for all our data acquired with handheld cameras. For UAV images, we used the more complex parametrization of the error components described in Mair et al. (2022a; cf. section 2.4 therein), with the uncertainty quantities of the photogrammetric models provided in Table S2.

3 | RESULTS

3.1 | Grain segmentation

Here, we first report the performance of the models we trained, for which we compare their segmentation results with the results of other benchmark methods in Section 3.1.1. After that, in Section 3.1.2, we describe specific systematics that control our models' segmentation performance, which ultimately lead to image type-specific segmentation models (Section 3.1.3).

3.1.1 | Overall performance

Our trained segmentation models are generally able to segment coarse sedimentary particles at a high precision for all our datasets (see Figure 4, for example; for all segmentation of the best performing models, see Figure S3). Although the segmentation performance varied depending on the splits of the training data and the training or retraining strategy (see Sections 3.1.2), the Cellpose segmentation models generally vastly outperformed the other methods (Table 1; see Figures S4–S6 for segmentation results of GrainID, `PebbleCounts` and the `nuclei` model, respectively). This is evident when comparing average APs (at 0.5 IoU) for the test set, where Δ AP values range from 0.419 for our best model to 0.142 for our overall worst model when compared with the respective best benchmark model (Table 1). This

FIGURE 4 Examples of segmentations that resulted from our best performing models (S1: *full_set*, FH: *fh+*, and APF: *full_set*, respectively) compared with ground truth annotations, results of a generalist model (*apf*), and segmentations of the benchmarks methods (GrainID; Chen, Hassan, & Fu, 2022; and PCauto; Purinton & Bookhagen, 2019) for selected test image tiles. AP = average precision at intersection over union threshold of 0.5 for the corresponding tile.




TABLE 1 Overall segmentation performance of our models and benchmarks for our test image tiles for all datasets.

Dataset S1 (n = 347)				Dataset FH (n = 153)				Dataset APF (n = 768)			
Model	AP	mAP	n _{pred}	Model	AP	mAP	n _{pred}	Model	AP	mAP	n _{pred}
full_set	0.753	0.575	325	fh+	0.750	0.547	104	full_set	0.633	0.474	427
apf_s1	0.737	0.564	334	fh	0.705	0.498	113	apf_fh	0.628	0.470	438
s1+	0.730	0.546	334	apf_fh	0.676	0.494	91	apf_s1	0.620	0.475	448
s1	0.725	0.560	344	full_set	0.667	0.492	90	apf	0.607	0.458	437
apf	0.687	0.514	322	fh _{fs}	0.621	0.449	75	full _{fs}	0.526	0.399	321
s1 _{fs}	0.624	0.457	255	apf	0.504	0.318	97	apf _{fs}	0.525	0.397	320
GrainID	0.334	0.103	262	GrainID	0.362	0.126	79	GrainID	0.326	0.140	410
nuclei	0.145	0.075	118	PCauto	0.109	0.033	19	nuclei	0.079	0.031	131
PCauto	0.112	0.043	64	nuclei	0.036	0.025	22	PCauto	0.071	0.029	70

Note: AP, average precision at the intersection over union (IoU) threshold of 0.5, averaged over the dataset; mAP, mean average precision over IoU thresholds 0.5 to 0.9, again averaged over the dataset; n, number of grains in the ground truth; n_{pred}, number of predicted grains; see Section 2.3 of the main text for more details on the metrics.

means that our best model correctly segmented 42% more grains with an IoU of 0.5 or higher, whereas the worst still segmented 14% more grains correctly than the best benchmark model. Moreover, any of our trained models segmented grains with a higher average precision at any IoU threshold across all datasets (Figure 5) and in each image tile (Figure S3) than the other models. On a dataset level, our models performed better on specialized datasets, with average AP values of 0.753 (S1; *full_set*) and 0.75 (FH; *fh+*) in contrast to 0.633 (APF; *full_set*) on the respective test set, all calculated at 0.5 IoU for the best-performing models. Furthermore, the best models also performed better on the specialized test data (S1, FH) at higher IoU thresholds, thus achieving higher mAP scores (Table 1) than all models

evaluated on the APF set. Finally, most models performed similarly for the training and test sets (Figure S7). Only for *s1*, *s1+* and *fh*, the models performed better for the training set, potentially exhibiting some overfitting to the training data.

3.1.2 | Systematic trends

Upon closer inspection, the segmentation results reveal systematic effects on the model performance for our data. First, using transfer learning, that is, re-training the models from the *nuclei* model, yields models that have a better segmentation performance than those




FIGURE 5 Segmentation performance for selected models applied on the image tiles that we used as test sets for the datasets S1 (a), FH (b) and APF (c). Best-performing models are indicated. The lines represent the average precision that is averaged over the test tiles, and the shaded area represents the standard deviation (1 sigma). Higher AP values indicate a higher percentage of detected grains at a corresponding IoU threshold, whereas higher values on the IoU axis indicate stricter acceptance criteria for grains considered as detected. CP ‘nuclei’ refers to the performance of the original Cellpose model for cell nuclei segmentation without any re-training with our data. For the performance on individual train and test tiles and the average of the entire dataset, we refer to Figure S7. The segmentation performance of the benchmark methods (GrainID; Chen, Hassan, & Fu, 2022; and PCauto; Purinton & Bookhagen, 2019) are compared.

models we trained from scratch with the same data (denoted with the suffix ‘ $_{fs}$ ’). Notably, we made this observation for all datasets. The differences in average AP values (at an IoU of 0.5) amount to 11% for set APF (for both $full_set_{fs}$ and apf_{fs}) and 13% for sets FH (for fh_{fs}) and S1 ($S1_{fs}$). This indicates that the models still benefit from the learning that occurred on the much larger image dataset of cell nuclei (>1000 images), despite a very low predictive power of the *nuclei* model if evaluated on our images without the re-training (Figure 5 and Table 1; see also Figure S6).


Second, the composition and the content of the training data in combination with the re-training strategy (see Figure 3 for an overview) had significant effects on the segmentation performance. Starting with the heterogeneous APF set, it is noteworthy that a training on 53 (apf_fh) and 61 (apf_s1) tiles yields a similar segmentation performance (within ~1% difference on the average AP at 0.5 IoU score in the test set) as a training on the full dataset of 67 tiles ($full_set$; Figure 5c). This is different for the homogeneous and specialized set S1, where the use of all tiles improves the performance drastically by ~7% ($full_set$ vs apf ; Figure 5a), thereby even slightly outperforming more specialized models trained only on the S1 set (e.g., $s1$ and $s1+$). This shows that adding the data from the gravel pit (FH) only marginally increases the predictive power for purely fluvial settings (i.e., APF and S1). In line with this, for the contrastingly different FH set, models that were trained only on the six tiles from the gravel pit outcrops (fh , $fh+$) performed better than the models trained on larger sets (e.g., apf_fh). As a result, while showing the highest

score in both S1 and APF test sets, the $full_set$ model falls behind the best segmenting model ($fh+$) in the FH test data by 8% (Figure 5b). This systematic influence of how the training data is composed further plays a role in which a re-training strategy leads to the best segmentation performance. For FH, a model ($fh+$) that was trained twice, i.e., that was trained from apf that itself had been re-trained from the *nuclei* model, performed best, whereas for S1, such a two-fold training strategy yielded results that were worse than those where training occurred on all training tiles. Hence, for the FH data, it is beneficial to start from the generalist weights and to train the model only on the dataset that differs from the generalized and more homogenous data, thereby allowing the model to learn a specific representation. For the S1 data, such an approach (i.e., $s1+$), along with training only on the S1 tiles ($s1$, $s1_{fs}$), is potentially hampered by overfitting the S1 training tiles (Figure S7).

3.1.3 | Image style classes

Our results so far have shown that the composition of the training dataset systematically influences the segmentation performance. Through the higher-dimensional clustering of the image style representation by the neural network, we obtained three distinct classes of image style classes (Figure 6a). We found that these classes consisted of (i) images with pebbles under sunny conditions with distinct shadows along granular interstices (‘sunny pebbles’; SP), (ii) images

FIGURE 6 Classes of image types for segmentation inferred from the style vectors used by the neural network. (a) The clustering of the style vectors is visualized through a dimensionality reduction by a t-distributed stochastic neighbor embedding (t-SNE) with image examples for each class (MRG, CV and SP). Inner colors indicate the original allocation in our dataset (S1, FH or APF). For more details, including the results of a principal components analysis of the style vectors, we refer to Figure S8. (b) Segmentation performance for models trained on the respective data split compared with the performance of the generalist model *full_set* on the same set.



that featured coarse particles within a sandy matrix ('matrix-rich' gravels; MRG) and (iii) images with higher visual complexity by vegetation, its shadow and/or water ('complex vegetation'; CV). We trained segmentation models (ig_{mrg} , ig_{cv} and ig_{sp}) for which image tiles and ground truth masks were re-combined into datasets according to their image class. We evaluated these models along with our generalist model on the respective datasets (Figure 6b). We found that the style-specific models show either a higher segmentation performance for two datasets (MRG and SP) or the same for one dataset (CV) when compared with our generalist *full_set* model. Moreover, this finding discloses that the segmentation performance of our models is lower for images with a higher visual complexity (i.e., AP scores are higher in datasets MRG and SP). We note that the class boundaries do not overlap with dataset boundaries and hence that none of our datasets consists exclusively of one image type (Figure 6a). Therefore, we used

the best-performing model on a dataset level for segmenting, when we measured the grains on images in the following sections.

3.2 | Grain sizes

Here, we report the results of our grain size measurements. We do so first on an image tile basis and for unscaled data (Section 3.2.1) to compare the grain size data resulting from the different methods to those where grains were manually annotated in the ground truth images before measurements. Second, we report the results after scaling, and we compare the data with grain sizes measured with different methods, including field measurements, in Section 3.2.2. Finally, we present maps displaying the distribution of grain sizes on entire gravel bars in Section 3.2.3.

3.2.1 | Measurement quality

We measured the size of grains whose shapes were approximated by either an ellipse or a convex hull. Interestingly, the approximation method has no significant influence on the grain sizes for our models, that is, both ellipses and convex hulls yield similar size distributions (Figures 7 and S9). In addition, when comparing to the ground truth data (Figure 7), our best segmenting models (*full_set* for S1 and APF, and *fh+* for FH) are returning values that are overall accurate for all tiles and for each dataset, independent of the approximation method. The observation that at least 88.9% (S1), 85.7% (FH) and 60.7% (APF) of the grain size results cannot statistically be distinguished from the size distribution of the ground truth data within 95% confidence ($p > 0.05$,

two-sample Kolmogorov–Smirnov test; Table 2) confirms this. In addition, the results of the segmentation models are very precise for S1 and FH. This is inferred from the relative differences of <10% on average for any percentile values in almost all tiles (Figure 7), and the absolute average difference that is <10 pixels (Table 2). For the APF, the precision is slightly lower. Furthermore, despite the overall good performance, the results do not match the ground truth data for a few tiles (see Figures S10–S13). However, the overall high accuracy and precision in S1 and FH and across all percentile values are also evident when comparing individual key percentile values (Figure S9). Finally, comparing our models' results with the benchmark methods' results reveals that our models deliver precise and accurate results for any tested grain approximation across all datasets (Figure 7 and Table 2).

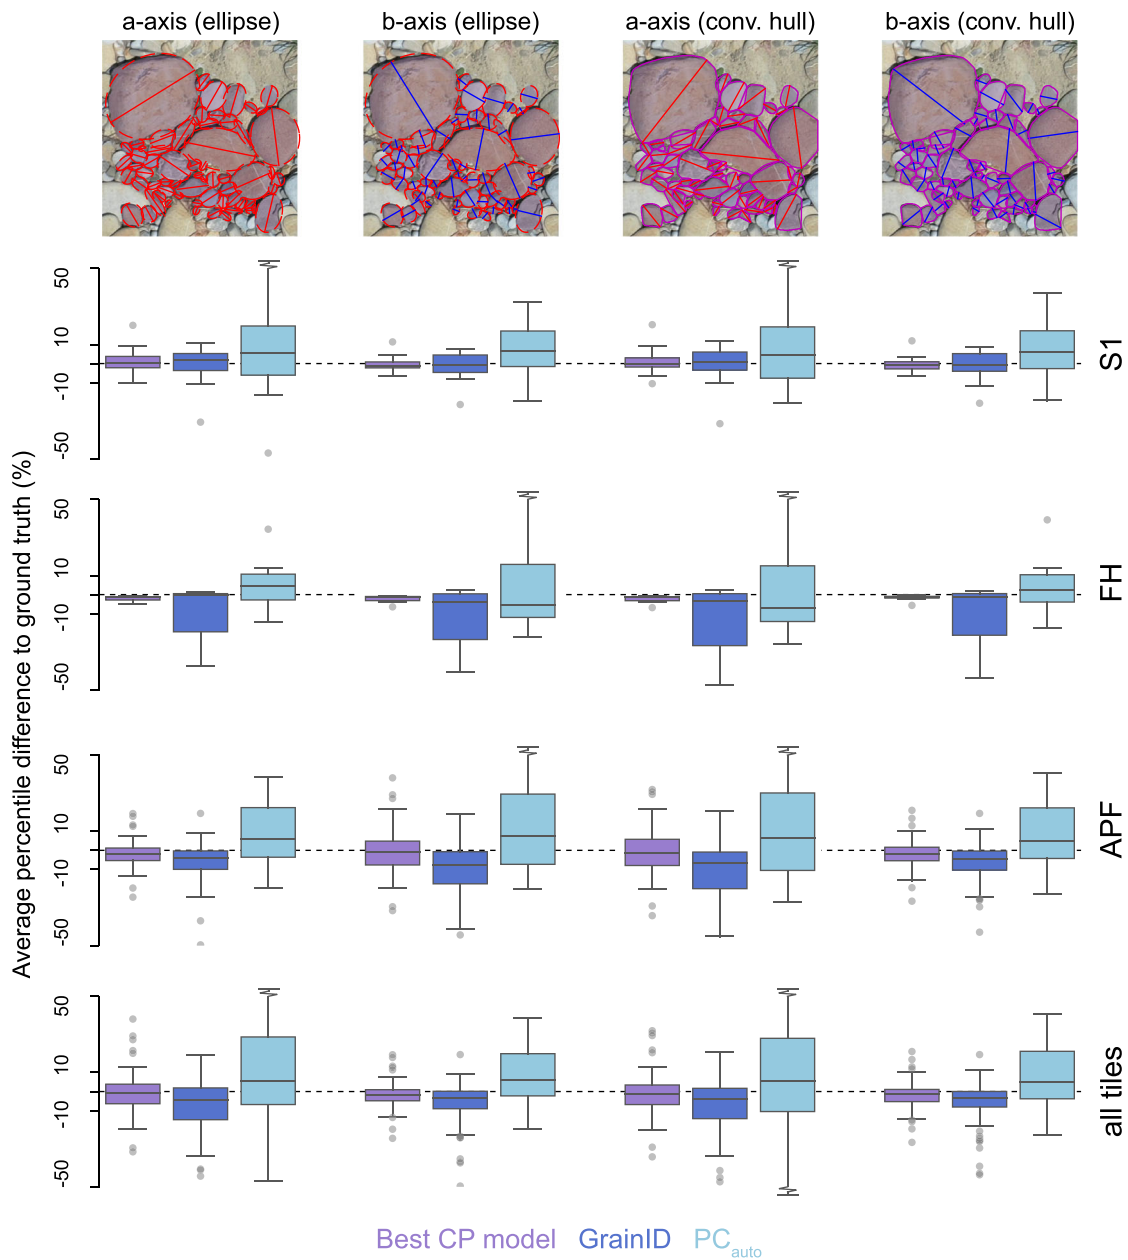


FIGURE 7 Overall quality of grain size data collected with different methods in image tiles. The quality is quantified by the closeness of predictions to ground truth for different grain size proxies for all tiles and the respective dataset splits (S1: 18, FH: 7, and APF: 56). We report the average difference of all percentiles as the relative difference between each percentile of the prediction set and the respective ground truth data. For information on the average percentile difference of key percentiles (i.e., D_{16} , D_{50} , D_{84} and D_{96}), we refer to Figure S9. The best Cellpose models (CP) refer to the models with the highest average AP score (0.5 IoU; Section 3.1.1) and are *full_set* for S1 and APF and *fh+* for FH, respectively.

TABLE 2 Statistical summary of the closeness between predictions and ground truth (GT) for different grain size proxies across the image tiles datasets.

Grain size metric	Data	Percentage of tiles (%) identical with GT			Average percentile difference (px) to GT		
		Best CP model	GrainID	PCauto	Best CP model	GrainID	PCauto
a-axis (ellipse)	S1	94.4	38.9	33.3	1.3 ± 9.7	-0.4 ± 15.5	6.8 ± 32.5
	FH	100.0	71.4	0.0	-2.4 ± 4.0	-17.5 ± 29.5	5.4 ± 21.1
	APF	78.6	55.4	16.1	-0.8 ± 14.5	-11.7 ± 18.6	14.6 ± 37.6
b-axis (ellipse)	S1	88.9	44.4	27.8	-0.5 ± 6.7	-1.1 ± 10.7	8.0 ± 22.5
	FH	85.7	71.4	14.3	-2.1 ± 2.3	-10.7 ± 18.9	5.7 ± 13.8
	APF	67.9	46.4	21.4	-2.0 ± 9.7	-6.8 ± 12.5	10.5 ± 26.1
a-axis (convex hull)	S1	100.0	38.9	22.2	1.0 ± 9.7	-0.3 ± 15.9	5.8 ± 33.9
	FH	100.0	71.4	0.0	-2.5 ± 4.0	-19.0 ± 32.5	3.9 ± 21.8
	APF	78.6	53.6	12.5	-1.1 ± 15.5	-11.8 ± 19.3	13.6 ± 39.1
b-axis (convex hull)	S1	88.9	38.9	16.7	-0.7 ± 7.4	-0.7 ± 11.9	7.4 ± 23.9
	FH	85.7	71.4	14.3	-1.8 ± 2.9	-12.1 ± 22.8	5.3 ± 14.5
	APF	60.7	37.5	23.2	-2.1 ± 10.4	-7.1 ± 14.0	10.2 ± 27.9

Note: The table shows the percentage of tiles for which the size distribution of grains is not statistically distinguishable from those in the ground truth dataset within 95% confidence (i.e., $p \geq 0.05$ for a two-sample Kolmogorov–Smirnov test). All grain sizes are measured on filtered masks, that is, only grains with b-axes ≥ 12 px and with a centre point situated within the central 64% of the image tile are taken into account (for visual reference, see Figure 7). We calculated the average percentile difference as the mean of the difference between the percentiles of the respective prediction set and the ground truth, and reported it along with the associated 1σ standard deviation. We refer to Figures S10–S13 for all tile-by-tile results. The best Cellpose models (CP) refer to the models with the highest average AP score (0.5 IoU; Section 3.1.1) and are *full_set* for S1 and APF and *fh+* for FH, respectively.

3.2.2 | Size accuracy

Here, we present the results where grains were measured on images after scaling (Figure 8a,b), which we compare to the data collected independently, in the field (K1, S1) and manually in images (FH; Figure 8c,d). We used only the ellipse approximation because this method yielded similar results as the convex hull (see above). All our independent measurements were conducted with grid sampling, either in the field or digitally. Therefore, we resampled all our grains with a similar digital grid to allow a direct comparison of the results. All grain size distributions statistically represent the respective reference measurement ($p > 0.05$; two-sample Kolmogorov–Smirnov test; see Table 3). However, a more detailed inspection reveals that some axes values in the images are much closer to the reference data than others. Specifically, for the Kander site (K1), the lengths of the b-axis represent the field data perfectly where the sizes differ by $<10\%$ for all percentiles (Figure 8d) and where the average of the difference between the percentile values of the reference data and the data collected with our approach is -0.4 ± 3.3 mm (Table 3). Similarly, for the Finsterhennen (FH) example, the average difference between the percentile values is generally small both for the b-axis (0.1 ± 0.9 mm) and the a-axis (-0.9 ± 1.2 mm). Yet, a look at the whole grain size distribution discloses much larger differences that are evened out (Figure 8d). Nevertheless, the differences between the percentile values of data collected from the prediction masks and the percentiles from the reference dataset never exceed $\pm 20\%$. Furthermore, they agree within uncertainties with each other. As another example, the b-axis values from the Sense (S1) images are also in overall good agreement with the reference data and constantly within the uncertainty of the data collected from the field (Figure 8d). However, the lengths of the a-axis are overestimated for the percentile values D_5 to D_{50} . Unfortunately, we cannot resolve

whether this is an effect of the field sampling, the grain occlusion in the images, or whether this can be explained by a potential offset between the location of the field survey and the area on the image where data were collected.

3.2.3 | Size maps on the bar-scale

We tested our workflow where grains in the surface layer of gravel bars along Swiss rivers are investigated in two orthoimage mosaics (Figure 9), which were generated from close-range UAV surveys. This approach automatically delineated and measured $>268\,000$ individual grains for site S1 and $>143\,000$ for site K1. The results show a high variability of grain sizes across the bars, as exemplified by the local variation in the D_{50} (Figure 9a). This large variability in particle size in the surface layer (the D_{50} ranges from 20.6 to 46.5 mm for S1 and from 23.2 to 46.5 mm for K1) allows disclosing areas of coarse- and fine-grained gravel (e.g., S1 in Figure 9a). These variations in the grain sizes are generally more significant than the uncertainty within the local images (see, e.g., Figure 9b).

4 | DISCUSSION

Our results show that using specialist deep learning models and transfer learning allows the training of state-of-the-art segmentation models for images taken from coarse fluvial sediments. Moreover, our re-trained models delineate pebbles with high accuracy and precision in a fully automated way. These improvements in segmentation directly translate into results where grain sizes are determined more precisely, more accurately and with a larger number of observations than what can be achieved with the other benchmark methods.

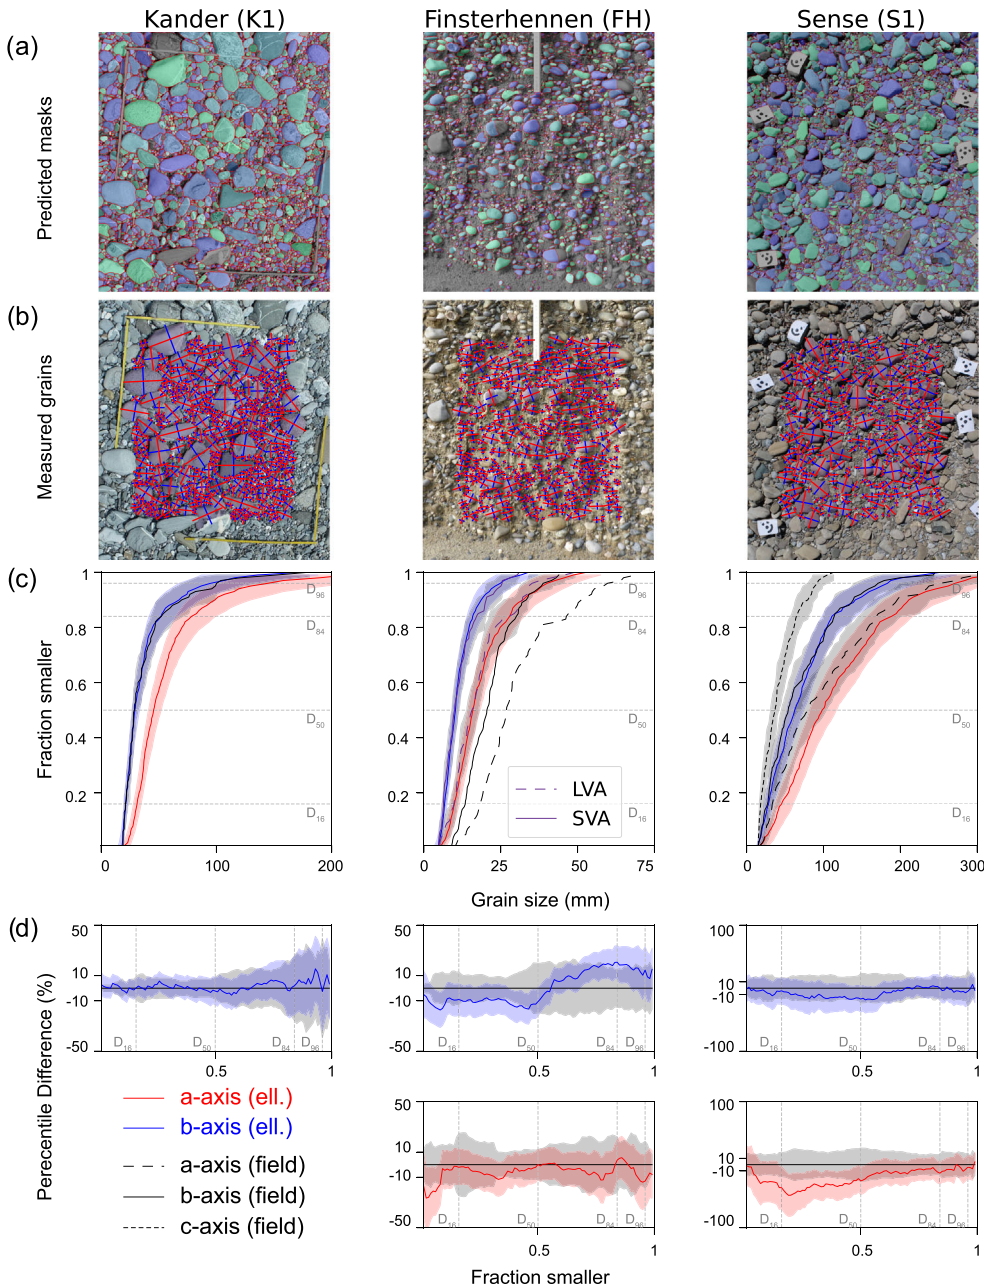


FIGURE 8 Comparison between grain sizes for three regions where data were collected on images and in the field. The predicted grain masks (a) of the best-performing models (*full_set* for K1 and S1, *fh+* for FH) are filtered to represent the same area measured independently. Additionally, results were re-sampled along a digital image grid to compare the different datasets (b). The resulting grain size distributions (c) are compared with the independently measured data on a percentile basis (d). Uncertainties are displayed as shaded areas and correspond to each percentile's 95% confidence interval (see Section 2.4 for details on the estimation). LVA, longest visible axis; SVA, shortest visible axis. We note here that for FH, we compared the predictions where grains are measured on undistorted images with a grid sampling approach. We do so because we expect a significant underestimation of the axes' lengths due to the occlusion of grains by the sandy matrix (for further details, see Garefalakis et al., 2023).

TABLE 3 Difference between values for key percentiles of grain size distributions collected from images and percentile values in the reference results.

Site	Ellipse axis	ΔD_{16} (mm)	ΔD_{50} (mm)	ΔD_{84} (mm)	ΔD_{96} (mm)	Avg. ΔD (mm)	p
K1	b	0.3	0.7	-1.1	2.8	-0.4 ± 3.3	> 0.99
FH	b	0.3	0.3	-1.4	0.5	0.1 ± 0.9	0.91
FH	a	0.1	0	-1.1	4.8	-0.9 ± 1.2	0.91
S1	b	1.4	9.4	0	8.6	-3.7 ± 4.3	0.82
S1	a	12.1	20.7	23.2	11.2	-14.2 ± 6.3	0.11

Note: The uncertainty refers to the 1σ standard deviation of the average percentile difference. All grain size distributions are statistically not different from the results of the reference measurements ($p > 0.05$; two-sample Kolmogorov-Smirnov test).

Furthermore, we can achieve such good results with relatively small datasets compared with other approaches (e.g., Chen, Hassan, & Fu, 2022). For example, we used less than 1000 objects from only seven tiles for the training (Table S1) of model *fh+*, which achieves a high segmentation performance (Table 1). Our trained models can directly be employed for segmenting grains, and our approach and

data can be used to train custom segmentation models. In addition, our software can readily be used for measuring the size of segmented grains (Section 4.1). Our results, particularly the custom models, underscore the importance of the composition of the datasets, thereby documenting the potential of a data-driven approach for this type of analysis (Section 4.2). Furthermore, precise and automated




FIGURE 9 Grain sizes of the surface layer of entire gravel bars measured with ImageGrains for two sites at rivers in Switzerland. (a) Maps of the local D_{50} for the b-axis were obtained for subsets of the orthoimage mosaic with a size of 1000×1000 pixels each. The frequency histograms (insert) show high variability in the local D_{50} . (b) Example of segmentation masks and resulting grain size distributions for local image subsets, which are used to generate the maps in (a). All measurements of axes lengths were obtained through an ellipse approximation. Local Swiss coordinates (CH1903+) are provided for reference.

segmentations yield more extensive and spatially resolved grain size information for fluvial settings (Section 4.3).

4.1 | Applicability and limitations

We see several applications of our work, outlined in this section together with their limits. First, with our ImageGrains software library, our workflow can be directly applied to segment and measure coarse sedimentary particles in a large range of images. Second, the segmentation models we trained can now be used for any segmentation-based workflow that intends to measure the sizes of grains on images similar to ours (e.g., Carbonneau, Bizzi, &

Marchetti, 2018). Third, our annotated dataset can be used to train custom segmentation models for other image data types. Finally, by obtaining precisely segmented grain masks, crucial data on particle sphericity, roundness and orientation (e.g., Steer et al., 2022) can be obtained. They can inform novel data-driven study designs (e.g., Chen et al., 2023). Furthermore, grains from the prediction masks of our models might also form the basis for other machine learning applications, for example, the training of a classifier to identify the particles' petrography. All these applications can be realized with a regular stand-alone computer without specialized knowledge of coding with machine learning libraries. In particular, the segmentation model architecture allows a fast training and inference on a desktop PC with a consumer-grade GPU. For example, we could train our *full_set* model

for 1000 epochs in less than 30 minutes using an NVIDIA GeForce RTX 3070 GPU with 8GB RAM. On the same GPU, inference for 1000×1000 pixel images, for example, as used in Figure 9, took less than two minutes per image. Our grain size analysis with ImageGrains allows for measurements with a similar rate, where thousands of grains can be analysed within minutes. The Cellpose model architecture itself was rigorously tested (Pachitariu & Stringer, 2022; Stringer et al., 2021), and for our case, we found no evidence for a need to change the default configuration as this would not improve the segmentation performance (see Data S3). Thus, on a technical level, our approach is limited by considering the general characteristics of such machine learning-based methods.

Deep learning models, such as convolutional neural networks, tend to overfit the training data, and their trained state is difficult to interpret (e.g., Alzubaidi et al., 2021; Sun et al., 2017). For our case, overfitting can be detected by using appropriate ground truth data for testing and choosing a suitable training strategy (see Data S3 and S7 and Section 4.2 below). To ensure that the model has learned to segment grains on images correctly, the results should be compared with ground truth data by evaluating the segmentation performance with suitable benchmarks, which are the AP or mAP scores with defined IoU thresholds for object detection tasks. Thus, when applying pre-trained models to images with no ground truth, a simple visual inspection of the segmentation results is recommended to ensure that the model is properly segmenting grains. We emphasize this because, despite the overall good results of our models, they failed to predict grains for some challenging image tiles (Table 2; see Figure S3 for details). In summary, we find that primarily the data used for training control the capability of the model to segment grains. Therefore, mostly inherent image characteristics influence the applicability of our approach.

The nature of images itself imposes some limits on our workflow. First, all data on objects extracted from images have a minimum size controlled by the image resolution. In our case, we used a rigorous cut-off of 12 pixels length upon inference and measuring the size of grains, which leads to image-specific minimum grain sizes, for example, of 4.7 to 18.1 mm for the images in Figure 8a. Although lower cut-off values are possible, we opted for the more conservative value. The reasons for our 12 pixel cut-off are that (i) we found it hard to delineate grains smaller than this for the ground truth visually, and (ii) the model rescales images during inference and training in the configuration we use (see Data S2 and Table S3 for more details on the effect of this). Thus, predictions of smaller grains might yield unstable results. We note that such challenges are typical for this type of imagery, which could explain why other approaches were based on similar cut-off values of 20 pixels (Chen, Hassan, & Fu, 2022; Purinton & Bookhagen, 2019). Accordingly, measuring small grains has remained a challenge for fluvial settings (e.g., Carbonneau, Bizzi, & Marchetti, 2018; Marchetti et al., 2022; Steer et al., 2022). Second, image data need to be scaled and pre-processed accordingly, which might include a rectification and a photogrammetric alignment through SfM/MVS methods (e.g., James et al., 2019, 2020). Especially for data acquired with UAVs, image distortion and systematic errors stemming from the photogrammetric alignment can have a significant impact on the results' quality (e.g., Carbonneau & Dietrich, 2017; Woodget et al., 2018; Mair et al., 2022a). Third, our approach and other models (e.g., Weigert et al., 2020), which are based on

microscopy images, are not well suited for a 3D segmentation of sedimentary particles, despite a dedicated 3D segmentation functionality. Such models infer 3D shapes from a stack of images of the same objects. This is achieved either through slicing the objects of interest or through applications of non-destructive imaging methods, which is impossible for topographic point clouds. Thus, segmentation methods that use ellipsoidal fitting in a semi-automated fashion (Steer et al., 2022) might be more suitable for such data. Nevertheless, deep learning might advance the segmentation of topographic point clouds in the future, possibly by improving and/or modifying existing methods with neural networks (e.g., Qi et al., 2018). Once such models and datasets for segmenting sedimentary clasts would be available, we would expect similar systematics to govern the segmentation performance, as is currently the case for our image-based method.

4.2 | Custom segmentation models

Our results show that aside from the training strategy, the main control on segmentation performance is distinct differences in image content (see Sections 3.1.2 and 3.1.3 and Figure 6). Consequently, dataset balance and composition are more critical than dataset size for our models, despite the almost universal agreement in computer vision literature that more data improve the model performance. In line with the finding of Pachitariu and Stringer (2022), we hypothesize that the reason for this is that specialist models like ours are trained to find only objects of a few classes or a single class in a much narrower range of images, whereas more generalist models, for example, the Segment Anything Model (SAM; Kirillov et al., 2023), are required to detect objects from many classes in a much larger variety of images. For example, we notice that while using more data for training a specialized model for the S1 tiles, it did not improve the segmentation performance when compared with segmentation performance of the best model for the FH tiles (Figure 5). A direct consequence is the selection of datasets. It is, therefore, essential to pay attention to dataset balance and composition, concerning image content and visual complexity, for example, shadows, vegetation, water, particle size, pebble shape and so on. Our classes of image styles (Section 3.1.3) can help when facing the question of which model to select and which data to use for a custom model. Specifically, before annotating any images, one can infer style vectors for new images and embed them together with the style vectors of our data. In case where the inferred image style differs from the data used for training, a custom model will likely exhibit a better segmentation performance than existing models. This can also point to the kind of data split that might be most promising for training, that is, our full set or a specific class of image style. Related to the data composition, we find that the most effective re-training strategy is influenced by the dataset composition (see Section 3.1.2). Hence, the best strategy for custom models might depend on the kind of data. Again, a style vector clustering might inform the decision on the optimal strategy. In particular, the style vectors for the FH images are located closely together (Figure 6a). Thus, datasets with similar embedded distributions might benefit from the same training strategy that we used for *fh+*. Furthermore, our models can predict an initial approximation of masks at the annotating step, facilitating ground truth generation. The masks can

then be manually corrected before being used as ground truth for model training. This can be further sped up using the human-in-the-loop approach (Pachitariu & Stringer, 2022), for which a custom model is trained after a newly annotated image is added to the dataset. In such an iterative process, the updated model would be used to create increasingly precise masks, thereby reducing the effort of manual corrections.

4.3 | Implications for measurements of the sizes of fluvial grains

Our segmentation and workflow for measuring grain size with ImageGrains allow for a near-complete delineation and measurement of grains in images of coarse fluvial particles. Because of the high segmentation accuracy, the resulting grain size dataset could be considered as if it had been collected by an area-by-number sampling approach (Bunte & Abt, 2001; and references therein). However, our workflow also allows for a grid (or random) resampling of grains, thus requiring an explicit choice of data type. For example, a grid resampling is needed to compare the results of our image analysis to the reference data collected in the field (Figure 8), which was accomplished by a grid-by-number approach. Such method-specific traits become even more critical when analysing partially occluded or partially buried particles. This is the case for FH, where grain size measurements on images yielded different results than field-based surveys where grains are manually measured (Garefalakis et al., 2023). Furthermore, our approach allows us to measure grains almost continuously and identify spatial variations within grain size data. For our K1 and S2 examples, the grain size patterns change significantly with respect to the sampling location within the same gravel bar (Figure 9). Similar trends of locally high variability in grain size distributions have also been observed in other field surveys (e.g., Chardon et al., 2020; Díaz Gómez et al., 2022; Rice & Church, 1998). In addition, spatial differences in sedimentary patterns, for example, vertical and lateral sorting and/or armouring (see also Bunte & Abt, 2001), can cause a change in the obtained results of grain size patterns. Such local variations in grain sizes might be distinctly different for different types of rivers (e.g., Guerit et al., 2018). Our examples K1 and S1 from Figures 8 and 9 are from small and alpine streams with high sediment throughput, where such a spatial variability in grain sizes is typical. Consequently, with a strategy where grain sizes would be measured only in isolated patches or with a different binning or resampling approach, much of the data variability might not be captured, potentially introducing a bias, particularly upon interpreting the data. Therefore, an eye has to be kept on such scale-related effects.

5 | CONCLUSIONS

Our workflow efficiently finds, segments and measures the size of individual coarse sedimentary particles in a broad range of images. We achieved this through transfer learning, which enabled us to train segmentation models for such sedimentary particles with a neural network-based model designed to segment cells in biomedical images. With this approach, we can improve the segmentation of sedimentary particles significantly compared to existing methods. This improvement

in segmentation allows us to overcome one of the major roadblocks for automated grain size measurements in images in the past. To make our workflow available for the larger community, we released the open-source ImageGrains software along with our annotated data and segmentation models.

Our contribution allows anyone to use our segmentation models for images of fluvial gravel directly. Additionally, our models and data can form the basis of custom segmentation models for other types of images, for which we provide training guidance. In our release, we included software tools to apply commonly used methods, such as ellipse fits and convex hull outline approximations, for obtaining the sizes of delineated grains. Furthermore, our approach includes quantitative methods to estimate the uncertainties for various image types, including UAV-derived imagery, nadir images, orthophoto mosaics and photographs from handheld cameras.

More generally, a precise segmentation of grains enables spatially resolved and accurate grain size measurements with high precision. The results of our analyses allow us to disclose distinct grain size patterns on a river bar scale. Furthermore, precisely segmented particles can be further investigated, for example, for shape and orientation or petrography.

AUTHOR CONTRIBUTIONS

David Mair conceptualized the research and developed the code together with Guillaume Witz. David Mair, Ariel Henrique Do Prado and Philippos Garefalakis did data collection in the field. David Mair was responsible for the data curation, which included creating the ground truth annotations. David Mair interpreted the results with scientific input from Guillaume Witz and Fritz Schlunegger and prepared the paper and figures with contributions from all co-authors.

ACKNOWLEDGEMENTS

We are indebted to Mykhailo Vladymyrov and Sigve Haug from the Data Science Lab of the University of Bern for their invaluable input and discussion during the conception of this study. We thank Fanny Ville, University of Lleida, for feedback on the first version of the manuscript. We acknowledge valuable comments from two reviewers and the editors that helped to improve this article.

DATA AND CODE AVAILABILITY

All code used in this study, which includes the trained models, is released as part of the ImageGrains library at <https://github.com/dmair1989/imagegrains>.

The images and grain size data used in this study, along with all our annotations are available under <https://doi.org/10.5281/zenodo.8005771>. The Cellpose nuclei dataset is available at <https://www.cellpose.org/dataset>.

ORCID

David Mair  <https://orcid.org/0000-0002-7018-6416>

Philippos Garefalakis  <https://orcid.org/0000-0002-6171-1258>

REFERENCES

Allen, P.A., Michael, N.A., D'Arcy, M., Roda-Boluda, D.C., Whittaker, A.C., Duller, R.A., et al. (2017) Fractionation of grain size in terrestrial sediment routing systems. *Basin Research*, 29(2), 180–202. Available from: <https://doi.org/10.1111/bre.12172>

- Alzubaidi, L., Zhang, J., Humaidi, A.J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., et al. (2021) Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. *Journal of Big Data*, 8(1), 53. Available from: <https://doi.org/10.1186/s40537-021-00444-8>
- Arzt, M., Deschamps, J., Schmied, C., Pietzsch, T., Schmidt, D., Tomancak, P., et al. (2022) LABKIT: labeling and segmentation toolkit for big image data. *Frontiers of Computer Science*, 4, 1–12. Available from: <https://doi.org/10.3389/fcomp.2022.777728>
- Attal, M., Mudd, S.M., Hurst, M.D., Weinman, B., Yoo, K. & Naylor, M. (2015) Impact of change in erosion rate and landscape steepness on hillslope and fluvial sediments grain size in the Feather River basin (Sierra Nevada, California). *Earth Surface Dynamics*, 3(1), 201–222. Available from: <https://doi.org/10.5194/esurf-3-201-2015>
- Brayshaw, D. D. (2012). Bankfull and effective discharge in small mountain streams of British Columbia. <https://doi.org/10.14288/1.0072555>
- Bunte, K. & Abt, S.R. (2001) Sampling surface and subsurface particle-size distributions in wadable gravel-and cobble-bed streams for analyses in sediment transport, hydraulics, and streambed monitoring. U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station. Available from: <https://doi.org/10.2737/rmrs-gtr-74>
- 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. Available from: <https://doi.org/10.1111/sed.12049>
- Buscombe, D. (2020) SediNet: a configurable deep learning model for mixed qualitative and quantitative optical granulometry. *Earth Surface Processes and Landforms*, 45(3), 638–651. Available from: <https://doi.org/10.1002/esp.4760>
- Butler, J.B., Lane, S.N. & Chandler, J.H. (2001) Automated extraction of grain-size data from gravel surfaces using digital image processing. *Journal of Hydraulic Research*, 39(5), 519–529. Available from: <https://doi.org/10.1080/00221686.2001.9628276>
- Caicedo, J.C., Goodman, A., Karhohs, K.W., Cimini, B.A., Ackerman, J., Haghighi, M., et al. (2019) Nucleus segmentation across imaging experiments: the 2018 data science bowl. *Nature Methods*, 16(12), 1247–1253. Available from: <https://doi.org/10.1038/s41592-019-0612-7>
- Carbonneau, P.E., Bizzi, S. & Marchetti, G. (2018) Robotic photosieving from low-cost multirotor sUAS: a proof-of-concept. *Earth Surface Processes and Landforms*, 43, 1160–1166. Available from: <https://doi.org/10.1002/esp.4298>
- Carbonneau, P.E. & Dietrich, J.T. (2017) Cost-effective non-metric photogrammetry from consumer-grade sUAS: implications for direct georeferencing of structure from motion photogrammetry. *Earth Surface Processes and Landforms*, 42(3), 473–486. Available from: <https://doi.org/10.1002/esp.4012>
- Carbonneau, P.E., Lane, S.N. & Bergeron, N.E. (2004) Catchment-scale mapping of surface grain size in gravel bed rivers using airborne digital imagery. *Water Resources Research*, 40(7), 1–11. Available from: <https://doi.org/10.1029/2003WR002759>
- Castelltort, S. & Van Den Driessche, J. (2003) How plausible are high-frequency sediment supply-driven cycles in the stratigraphic record? *Sedimentary Geology*, 157, 3–13. Available from: [https://doi.org/10.1016/S0037-0738\(03\)00066-6](https://doi.org/10.1016/S0037-0738(03)00066-6)
- Chardon, V., Piasny, G. & Schmitt, L. (2022) Comparison of software accuracy to estimate the bed grain size distribution from digital images: a test performed along the Rhine River. *River Research and Applications*, 38(2), 358–367. Available from: <https://doi.org/10.1002/rra.3910>
- Chardon, V., Schmitt, L., Piégay, H. & Lague, D. (2020) Use of terrestrial photosieving and airborne topographic LiDAR to assess bed grain size in large rivers: a study on the Rhine River. *Earth Surface Processes and Landforms*, 45(10), 2314–2330. Available from: <https://doi.org/10.1002/esp.4882>
- Chen, X., Hassan, M. A., and Fu, X.: CNN for image-based sediment detection applied to a large terrestrial and airborne dataset, Zenodo. <https://doi.org/10.5281/zenodo.5240906>, 2021.
- Chen, X., Hassan, M.A. & Fu, X. (2022) Convolutional neural networks for image-based sediment detection applied to a large terrestrial and airborne dataset. *Earth Surface Dynamics*, 10(2), 349–366. Available from: <https://doi.org/10.5194/esurf-10-349-2022>
- Chen, Z., Scott, C., Keating, D., Clarke, A., Das, J. & Arrowsmith, R. (2023) Quantifying and analysing rock trait distributions of rocky fault scarps using deep learning. *Earth Surface Processes and Landforms*, 1–17(6), 1234–1250. Available from: <https://doi.org/10.1002/esp.5545>
- Coelho, L.P., Shariff, A. & Murphy, R.F. (2009) Nuclear segmentation in microscope cell images: a hand-segmented dataset and comparison of algorithms, proceedings. In: 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 518–521. Available from: <https://doi.org/10.1109/ISBI.2009.5193098>
- Detert, M. & Weitbrecht, V. (2012) Automatic object detection to analyze the geometry of gravel grains—a free stand-alone tool, River Flow 2012 - Proc. International Conference on Fluvial Hydraulics, 1, 595–600.
- Díaz Gómez, R., Pasternack, G.B., Guillon, H., Byrne, C.F., Schwindt, S., Larrieu, K.G. et al. (2022) Mapping subaerial sand-gravel-cobble fluvial sediment facies using airborne lidar and machine learning. *Geomorphology*, 401, 108106. Available from: <https://doi.org/10.1016/j.geomorph.2021.108106>
- Dunne, K.B.J. & Jerolmack, D.J. (2018) Evidence of, and a proposed explanation for, bimodal transport states in alluvial rivers. *Earth Surface Dynamics*, 6(3), 583–594. Available from: <https://doi.org/10.5194/esurf-6-583-2018>
- Eaton, B.C., Moore, R.D. & MacKenzie, L.G. (2019) Percentile-based grain size distribution analysis tools (GSDtools)—estimating confidence limits and hypothesis tests for comparing two samples. *Earth Surface Dynamics*, 7(3), 789–806. Available from: <https://doi.org/10.5194/esurf-7-789-2019>
- Garefalakis, P., do Prado, A., Mair, D., Douillet, G.A., Nyffenegger, F. & Schlunegger, F. (2023) Comparison of three grain size measuring methods applied to coarse-grained gravel deposits. *Sedimentary Geology*, 446, 106340. Available from: <https://doi.org/10.1016/j.sedgeo.2023.106340>
- Gatys, L.A., Ecker, A.S. & Bethge, M. (2016) Image style transfer using convolutional neural networks. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2414–2423. Available from: <https://doi.org/10.1109/CVPR.2016.265>
- Guerit, L., Barrier, L., Liu, Y., Narteau, C., Lajeunesse, E., Gayer, E., et al. (2018) Uniform grain-size distribution in the active layer of a shallow, gravel-bedded, braided river (the Urumqi River, China) and implications for paleo-hydrology. *Earth Surface Dynamics*, 6(4), 1011–1021. Available from: <https://doi.org/10.5194/esurf-6-1011-2018>
- He, K., Gkioxari, G., Dollar, P. & Girshick, R. (2020) Mask R-CNN. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2), 386–397. Available from: <https://doi.org/10.1109/TPAMI.2018.2844175>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778, <https://doi.org/10.1109/CVPR.2016.90>
- James, M.R., Chandler, J.H., Eltner, A., Fraser, C., Miller, P.E., Mills, J.P. et al. (2019) Guidelines on the use of structure-from-motion photogrammetry in geomorphic research. *Earth Surface Processes and Landforms*, 44(10), 2081–2084. Portico. Available from: <https://doi.org/10.1002/esp.4637>
- James, M.R., Antoniazza, G., Robson, S. & Lane, S.N. (2020) Mitigating systematic error in topographic models for geomorphic change detection: accuracy, precision and considerations beyond off-nadir imagery. *Earth Surface Processes and Landforms*, 45(10), 2251–2271. Portico. Available from: <https://doi.org/10.1002/esp.4878>
- Karras, T., Laine, S. & Aila, T. (2021) A style-based generator architecture for generative adversarial networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(12), 4217–4228. Available from: <https://doi.org/10.1109/TPAMI.2020.2970919>
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Berg, A., Lo, W.-Y., Doller, P., & Girshick, R. (2023). Segment Anything, Meta AI Research.
- Kumar, N., Verma, R., Anand, D., Zhou, Y., onder, O.F., Tsougenis, E., et al. (2020) A multi-organ nucleus segmentation challenge. *IEEE*

- Transactions on Medical Imaging*, 39(5), 1380–1391. Available from: <https://doi.org/10.1109/TMI.2019.2947628>
- Lang, N., Irmiger, A., Rozniak, A., Hunziker, R., Wegner, J.D. & Schindler, K. (2021) GRAINet: mapping grain size distributions in river beds from UAV images with convolutional neural networks. *Hydrology and Earth System Sciences*, 25(5), 2567–2597. Available from: <https://doi.org/10.5194/hess-25-2567-2021>
- Li, G., Liu, T., Nie, J., Guo, L., Chen, J., Zhu, J., et al. (2008) Segmentation of touching cell nuclei using gradient flow tracking. *Journal of Microscopy*, 231(1), 47–58. Available from: <https://doi.org/10.1111/j.1365-2818.2008.02016.x>
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., & Zitnick, C. L. (2014). Microsoft COCO: Common Objects in Context, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 8693 LNCS, 740–755, https://doi.org/10.1007/978-3-319-10602-1_48
- Litty, C. & Schlunegger, F. (2017) Controls on pebbles' size and shape in streams of the Swiss Alps. *Journal of Geology*, 125(1), 101–112. Available from: <https://doi.org/10.1086/689183>
- Lu, J., Behbood, V., Hao, P., Zuo, H., Xue, S. & Zhang, G. (2015) Transfer learning using computational intelligence: a survey. *Knowledge-Based Systems*, 80, 14–23. Available from: <https://doi.org/10.1016/j.knsys.2015.01.010>
- Mair, D., Do Prado, A.H., Garefalakis, P., Lechmann, A., Whittaker, A. & Schlunegger, F. (2022a) Grain size of fluvial gravel bars from close-range UAV imagery—uncertainty in segmentation-based data. *Earth Surface Dynamics*, 10(5), 953–973. Available from: <https://doi.org/10.5194/esurf-10-953-2022>
- Mair, D., Henrique, A., Prado, D., Garefalakis, P., Lechmann, A., Whittaker, A., & Schlunegger, F. (2022b). Data and code for: Grain size of fluvial gravel bars from close-range UAV imagery—uncertainty in segmentation-based data, Zenodo. <https://doi.org/10.5281/zenodo.6415047>
- Marchetti, G., Bizzi, S., Belletti, B., Lastoria, B., Comiti, F. & Carbonneau, P.E. (2022) Mapping riverbed sediment size from Sentinel-2 satellite data. *Earth Surface Processes and Landforms*, 47(10), 2544–2559. Portico. Available from: <https://doi.org/10.1002/esp.5394>
- Mörtl, C., Baratiar, A., Berthet, J., Duvillard, P.-A., and De Cesare, G. (2022). GALET: A deep learning image segmentation model for drone-based grain size analysis of gravel bars. In: Proceedings of the 39th IAHR World Congress, 5326–5335, <https://doi.org/10.3850/IAHR-39WC252171192022895>
- Pachitariu, M. & Stringer, C. (2022) Cellpose 2.0: how to train your own model. *Nature Methods*, 19(12), 1634–1641. Available from: <https://doi.org/10.1038/s41592-022-01663-4>
- Padilla, R., Netto, S. L., & da Silva, E. A. B. (2020). A Survey on Performance Metrics for Object-Detection Algorithms, in: 2020 International Conference on Systems, Signals and Image Processing (IWSSIP), 237–242, <https://doi.org/10.1109/IWSSIP48289.2020.9145130>
- Piégay, H., Arnaud, F., Belletti, B., Bertrand, M., Bizzi, S., Carbonneau, P., et al. (2020) Remotely sensed rivers in the Anthropocene: state of the art and prospects. *Earth Surface Processes and Landforms*, 45(1), 157–188. Available from: <https://doi.org/10.1002/esp.4787>
- Poličar, P. G., Stražar, M., & Zupan, B. (2019) OpenTSNE: A modular Python library for t-SNE dimensionality reduction and embedding, 1–2, <https://doi.org/10.1101/731877>
- Purinton, B. & Bookhagen, B. (2019) Introducing PebbleCounts: a grain-sizing tool for photo surveys of dynamic gravel-bed rivers. *Earth Surface Dynamics*, 7(3), 859–877. Available from: <https://doi.org/10.5194/esurf-7-859-2019>
- Purinton, B. & Bookhagen, B. (2021) Tracking downstream variability in large grain-size distributions in the South-Central Andes. *Journal of Geophysical Research - Earth Surface*, 126(8), 1–29. Available from: <https://doi.org/10.1029/2021JF006260>
- Qi, C.R., Liu, W., Wu, C., Su, H. & Guibas, L.J. (2018) Frustum PointNets for 3D object detection from RGB-D data. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 918–927. Available from: <https://doi.org/10.1109/CVPR.2018.00102>
- Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., et al. (2019) Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743), 195–204. Available from: <https://doi.org/10.1038/s41586-019-0912-1>
- Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I. & Savarese, S. (2019) Generalized intersection over union: a metric and a loss for bounding box regression. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 658–666. Available from: <https://doi.org/10.1109/CVPR.2019.000075>
- Rice, S. & Church, M. (1998). Grain size along two gravel-bed rivers: statistical variation, spatial pattern and sedimentary links. *Earth Surface Processes and Landforms*, 23(4), 345–363. Available from: [https://doi.org/10.1002/\(SICI\)1096-9837\(199804\)23:4%3C345::AID-ESP850%3E3.0.CO;2-B](https://doi.org/10.1002/(SICI)1096-9837(199804)23:4%3C345::AID-ESP850%3E3.0.CO;2-B)
- Ronneberger, O., Fischer, P. & Brox, T. (2015) U-net: convolutional networks for biomedical image segmentation. *IEEE Access*, 9, 16591–16603. Available from: <https://doi.org/10.48550/arXiv.1505.04597>
- Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., et al. (2012) Fiji: an open-source platform for biological-image analysis. *Nature Methods*, 9(7), 676–682. Available from: <https://doi.org/10.1038/nmeth.2019>
- Schmidt, U., Weigert, M., Broaddus, C. & Myers, G. (2018) Cell detection with star-convex polygons. *Lecture Notes in Computer Science*, 265–273. Available from: https://doi.org/10.1007/978-3-030-00934-2_30
- Schlunegger, F. & Norton, K.P. (2015) Climate vs. tectonics: the competing roles of Late Oligocene warming and Alpine orogenesis in constructing alluvial megafan sequences in the North Alpine foreland basin. *Basin Research*, 27(2), 230–245. Available from: <https://doi.org/10.1111/bre.12070>
- Soloy, A., Turki, I., Fournier, M., Costa, S., Peuziat, B. & Lecoq, N. (2020) A deep learning-based method for quantifying and mapping the grain size on pebble beaches. *Remote Sensing*, 12(21), 1–23. Available from: <https://doi.org/10.3390/rs12213659>
- Steer, P., Guerit, L., Lague, D., Crave, A. & Gourdon, A. (2022) Size, shape and orientation matter: fast and semi-automatic measurement of grain geometries from 3D point clouds. *Earth Surface Dynamics*, 10(6), 1211–1232. Available from: <https://doi.org/10.5194/esurf-10-1211-2022>
- Stringer, C., Wang, T., Michaelos, M. & Pachitariu, M. (2021) Cellpose: a generalist algorithm for cellular segmentation. *Nature Methods*, 18(1), 100–106. Available from: <https://doi.org/10.1038/s41592-020-01018-x>
- Sulaiman, M.S., Sinnakaudan, S.K., Ng, S.F. & Strom, K. (2014) Application of automated grain sizing technique (AGS) for bed load samples at Rasil River: a case study for supply limited channel. *Catena*, 121, 330–343. Available from: <https://doi.org/10.1016/j.catena.2014.05.013>
- Sun, C., Shrivastava, A., Singh, S., & Gupta, A. (2017). Revisiting Unreasonable Effectiveness of Data in Deep Learning Era, Proc. IEEE Int. Conf. Comput. Vis., 2017-Octob, 843–852, <https://doi.org/10.1109/ICCV.2017.97>
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2014). Intriguing properties of neural networks, 2nd Int. Conf. Learn. Represent. ICLR 2014 - Conf. Track Proc., 1–10.
- Tofelde, S., Bernhardt, A., Guerit, L. & Romans, B.W. (2021) Times associated with source-to-sink propagation of environmental signals during landscape transience. *Frontiers in Earth Science*, 9, 1–26. Available from: <https://doi.org/10.3389/feart.2021.628315>
- Traag, V.A., Waltman, L. & van Eck, N.J. (2019) From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports*, 9(1), 5233. Available from: <https://doi.org/10.1038/s41598-019-41695-z>
- Van der Maaten, L. & Hinton, G. (2008) Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9, 2579–2605.
- Weigert, M., Schmidt, U., Haase, R., Sugawara, K. & Myers, G. (2020) Star-convex Polyhedra for 3D Object Detection and Segmentation in Microscopy. In: 2020 IEEE Winter Conference on Applications of Computer Vision (WACV). Available from: <https://doi.org/10.1109/wacv45572.2020.9093435>

- Whittaker, A.C., Attal, M. & Allen, P.A. (2010) Characterising the origin, nature and fate of sediment exported from catchments perturbed by active tectonics. *Basin Research*, 22(6), 809–828. Available from: <https://doi.org/10.1111/j.1365-2117.2009.00447.x>
- Woodget, A.S. & Austrums, R. (2017) Subaerial gravel size measurement using topographic data derived from a UAV-SfM approach. *Earth Surface Processes and Landforms*, 42(9), 1434–1443. Available from: <https://doi.org/10.1002/esp.4139>
- Woodget, A.S., Fyffe, C. & Carbonneau, P.E. (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. Portico. Available from: <https://doi.org/10.1002/esp.4285>
- Yosinski, J., Clune, J., Bengio, Y. & Lipson, H. (2014) How transferable are features in deep neural networks? *Advances in Neural Information Processing Systems*, 4, 3320–3328.
- Zech, J.R., Badgeley, M.A., Liu, M., Costa, A.B., Titano, J.J. & Oermann, E.K. (2018) Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS Medicine*, 15(11), 1–17. Available from: <https://doi.org/10.1371/journal.pmed.1002683>

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Mair, D., Witz, G., Do Prado, A.H., Garefalakis, P. & Schlunegger, F. (2023) Automated detecting, segmenting and measuring of grains in images of fluvial sediments: The potential for large and precise data from specialist deep learning models and transfer learning. *Earth Surface Processes and Landforms*, 1–18. Available from: <https://doi.org/10.1002/esp.5755>