Automated finding, segmenting, and measuring of grains in images of fluvial sediments – the potential of transfer learning in deep neural networks

David Mair[1], Guillaume Witz[2], Ariel Henrique Do Prado[1], Philippos Garefalakis[1], Fritz Schlunegger[1]

[1] Institute of Geological Sciences, University of Bern, Baltzerstrasse 1+3, 3012 Bern, Switzerland
[2] Data Science Lab, University of Bern, 3012 Bern, Switzerland

Correspondence to: David Mair (david.mair@unibe.ch)

Abstract. The size of sedimentary particles in gravel-bed rivers allows for inferring information on sediment entrainment or deposition mechanisms and on the hydraulic conditions controlling them. However, collecting data on these coarse-grained sediments is costly and time-consuming, both in the field and remotely from images. Therefore, recent attention has turned to machine learning models to improve such measurements. Despite their success, current methods need large quantities of data and yield results limited to a few percentile values of grain size datasets, often affected by systematic bias and low accuracy. 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 our re-trained models outperform existing methods so that the performance in segmentation tasks can be directly transferred to images of fluvial sediments. With our approach, these results are achievable with only 10-20% of the data that have been previously needed for training other machine learning models. Moreover, we find that primarily traits in our data control the segmentation performance, enabling data-driven approaches to improve future segmentation models. Additionally, comparing our automatically measured grains with the results retrieved from various image and field-based surveys confirms that these improvements in segmentation are directly leading to more precise and more accurate grain size data across different image settings. Finally, we release a software package, the trained models, and the used data. The goal is to offer a tool to efficiently segment and measure grains in images of sediments in an automated way, which can be adapted to different settings.

Keywords: Grain size, image segmentation, machine learning, neural network, fluvial gravel, sediment analysis

1 Introduction

Data on the size of the sediment transported by rivers, both in modern and ancient systems, is of crucial importance to understand the mechanisms, hydraulic conditions, and grain-to-grain interactions during sediment transport in fluvial systems (e.g., Piégay et al., 2020; Dunne and Jerolmack, 2018; Attal et al., 2015; Whittaker et al., 2010). In addition, grain size is critical to deciphering climate, tectonic, and supply signals preserved in the stratigraphic record of fluvial sediment routing systems (e.g., Tofelde et al., 2021; Allen et al., 2017; Schlunegger & Norton, 2015; Castelltort and Van Den Driessche, 2003). Standard field methods have been developed in the past years to measure the grain size of sediment transported in active rivers (e.g., Bunte and 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., Carbonneau et al., 2004; Butler et al., 2001). 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., 2015), semi-automated segmentation (e.g., Purinton and Bookhagen, 2019; Detert and Weitbrecht, 2012) to texture-based percentile predictions (e.g., Buscombe, 2013). While these methods allowed for a faster and remote measurement of larger numbers of grains, they still need 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; 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., 2022a; Mörtl et al., 2022; Soloy et al., 2020; Chen et al., 2023) or to directly predict percentile values of a grain size distribution (e.g., Lang et al., 2021; Buscombe, 2020). The main aims of all these works were to automate the measurements, improve reproducibility and scalability, and increase the number of observations. While deep learning did improve segmentation for some data and settings, challenges have remained (e.g., Chen et al., 2022a; Chen et al., 2023). In addition, current methods have limited ability to adapt to new, previously unseen data. They are restricted in predicting grain size data beyond a few and individual percentile values (e.g., D_{50}, D_{84}). 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., Fig. 1) pose challenges to current deep learning-based models. Such challenges might even prevent these neural network-based models from producing meaningful information in images for some data not seen during training (e.g., Szegedy et al., 2014; Zech et al., 2018). Thus, at least additional annotated data is 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 et al., 2022a,b and Lang et al., 2021, respectively) imposes an often prohibitive cost on re-training.

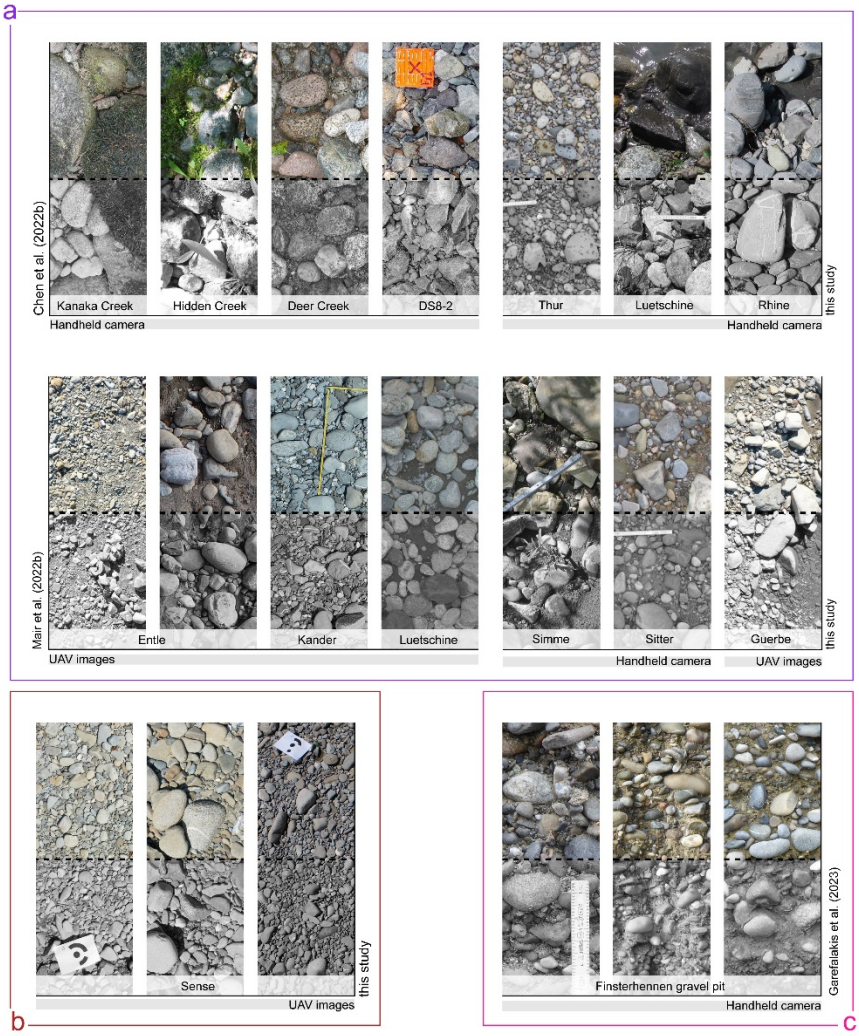

Figure 1: Examples of images taken from fluvial gravel bars used for this study. **a)** Images from Canadian and Swiss rivers (Chen et al., 2021; Mair et al., 2022b; respectively) that are used for the generalized APF ('all pebbles fluvial) data set. **b)** Uncrewed aerial vehicle (UAV) images from one 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, while the lower half depicts the single-channel greyscale image used for training and inference.

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), while others attempted to infer a size distribution from the roughness of a point cloud (e.g., Woodget et al., 2017). While 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 two main reasons. First, the related acquisition of data would be more complex, and the subsequent data processing time-consuming and thus expensive, and second such methods are not well suited to fit the geometry of angular sedimentary particles. 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 image segmentation, 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 new 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 designed to detect and segment cells and nuclei in biomedical images. We adapt this model to detect 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 fine-tuned based on models trained for nuclei segmentation or trained from scratch, vastly outperform existing models proposed for the segmentation of fluvial pebbles in all datasets we tested in this study, despite an order of magnitude smaller dataset size required for training. Furthermore, the models' flexibility and accessibility, as well as the ability to rapidly re-train them, allow us to generate task- or image-type-specific models, as we did here with different data sets. Dedicated models yield substantially better results for these specialized datasets than more generalizing models trained on larger datasets. In line with the approach of Pachitariu and Stringer (2022), we propose an interactive workflow to create high-performance models in a short time and from relatively small data sets. To facilitate access to these powerful segmentation models, we built an open-source software library, ImageGrains, which allows for (i) easy use of the Cellpose models we trained, (ii) straightforward training of custom models, as well as (iii) streamlined grain size measurements.

2 Methods

We employed an existing deep neural network model to segment and measure sedimentary particles (Fig. 2). Accordingly, our ImageGrains library (see Supplement S1 for a description) is built in Python v.3 and around the Cellpose library (Stringer et al., 2021; Pachitariu and Stringer, 2022; v.2.1.1), which we used as a backbone segmentation model. The Cellpose code itself is implemented in Python v.3, and it uses the pytorch, scipy, numpy, numba, OpenCV, PyQt, and pyqtgraph packages (Paszke et al., 2019; Virtanen et al., 2020; Van Der Walt et al., 2011; Lam et al., 2015; Bradski, 2000). In addition to these packages, ImageGrains uses matplotlib, pandas, scikit-image, and jupyter (Hunter, 2007; McKinney, 2010; van der Walt et al., 2014; Kluyver et al., 2016). In the following sections, we provide background information on the datasets and briefly describe the deep learning model used for segmentation and its training. Next, we evaluate the model's performance before explaining how the particle size was measured. We finally compare the resulting grain size data with the outcome of field surveys.

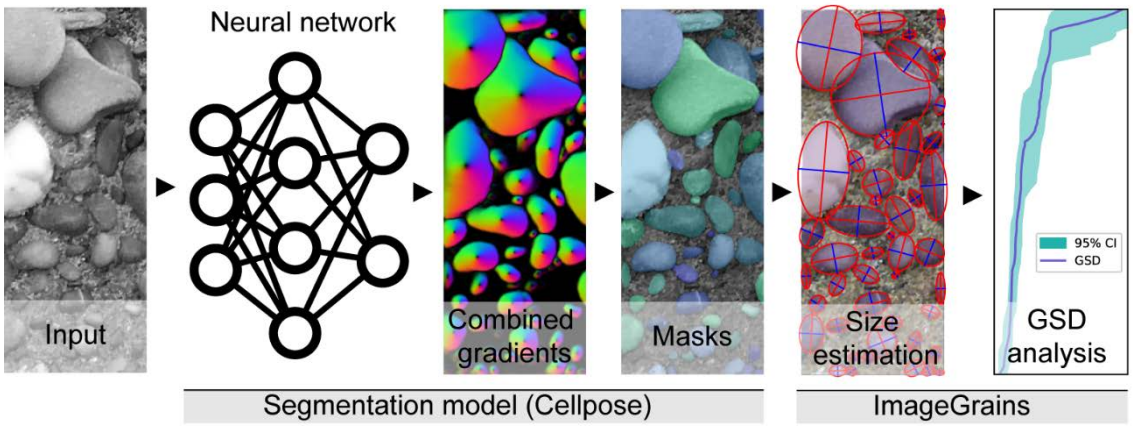

Figure 2: Overview of our workflow. We used the deep neural network model Cellpose (Stringer et al., 2021) as backbone model to segment sedimentary particles in images of fluvial sediments. On the resulting masks, we measured the sizes of individual grains, from which we compiled grain size distributions (GSD) with percentile-based uncertainty estimations (see section 2.4 and Appendix A for details).

2.1 Image data

We compiled a diverse dataset of fluvial sediment images with coarse-grained (> 2mm) fluvial sediment, so-called pebbles, into three basic datasets (Table S2), which are described below (see Fig. 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 the model training. 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 tiles while still capturing the visual complexity of the image. This led to a varying number of tiles for each image (Table S2).

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; Fig 1a) from several sites along Canadian (n_{images} = 4; Chen et al., 2021; Brayshaw, 2012) and Swiss (n_{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 et al. (2022a) to annotate grain boundaries more precisely. We complemented the APF dataset with seven additional photos from six sites along Swiss rivers (Fig. 1a), exhibiting challenging conditions (e.g., raindrops or wet gravel). These additional images were taken with a handheld camera (Litty and 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, while 9 tiles were kept as a test set. The second set, S1 (Fig. 1b), comprises five nadir images, all acquired with an UAV at one site in the Sense River (Switzerland) under homogenous light conditions. Set S1 contains 18 annotated tiles (14 for training and 4 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 (Fig. 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 (< 2mm grain size) matrix. This material fills the interstices and sometimes covers parts of the gravel grains. We annotated 7 tiles, one tile per image, for set FH, which we split by 6 to 1 for training and testing. From these annotated image tiles (see Supplementary Fig. S4), we trained a collection of segmentation models on different data splits (Fig. 3). We did so to assess the segmentation performance of the models as we vary the grade of specialization (Stringer et al., 2021) and to test how the training strategy influences segmentation performance.

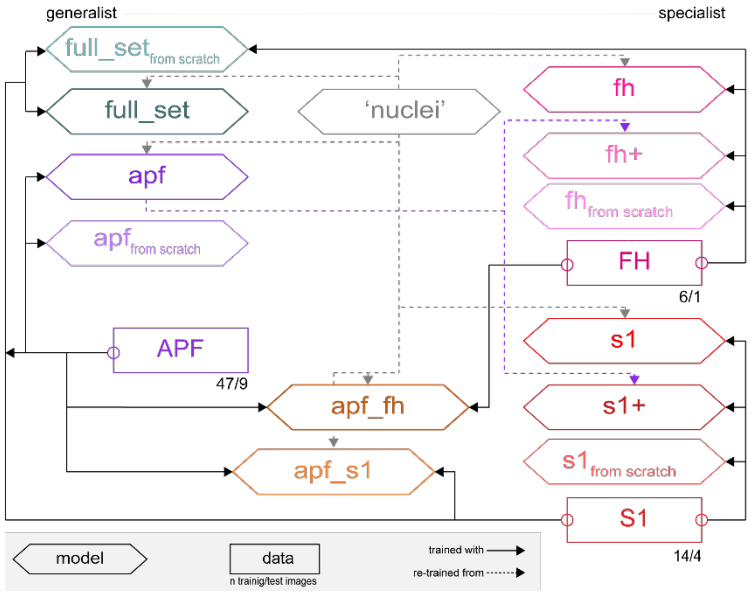

Figure 3: Overview of dataset splits and starting models used for training the segmentation models in this study. We refer to the main text for data description (Section 2.1) and model training (Section 2.2). The *nuclei* model refers to the pre-trained Cellpose model on images of cell nuclei of Stringer et al. (2021). We refer to section 3.1 in the main text for an overview of the model performances.

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 photogrammetry. Most 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, while others, e.g., 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 grain size uncertainty modeling (see section 2.4 below and Table S5), therein following Mair et al. (2022a). Ultimately, we used two orthoimage mosaics generated from UAV imagery on bars along the Swiss 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 (for all details, we refer to Stringer et al., 2021). It was trained with greyscale images and annotated ground truth of the objects to segment, i.e., cells or cell nuclei, and rock pebbles for our study (Fig. S4). The model itself (Fig. 2) is a deep neural network in the style of U-net (Ronneberger et al., 2015), which uses residual blocks (He et al., 2016). Similar to the original U-net, the model consists of a 4-level downsampling and upsampling pass (Fig. 2; Stringer et al., 2021), where the convolutions at each level resemble 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 et al., 2016; Karras et al., 2019). This style vector is used as input in the upsampling pass, influencing the segmentation results. It can thus be used to cluster data according to their respective segmentation style (Pachitariu and Stringer, 2022). The neural network itself is employed to predict gradients of a flow-vector field, which is simulated with an equation modeled on heat diffusion. These flow vectors are tailored to find the center 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 pixel routing to such a center can be identified, thereby allowing for the

segmentation of the object in the ROI and 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, therefore, trained mainly on various microscopy images of cells and a few images of repeated objects, e.g., fish scales, vegetables, or rocks. However, the model, which we use as a base for our re-training (see Fig. 3 for an overview), was exclusively trained on annotated data of cell nuclei in 1139 images of various sources (Stringer et al., 2021; Caicedo et al., 2019; Kumar et al., 2019; Coelho et al., 2009). All our models, presented in the following sections, were trained with stochastic gradient descent, a learning rate of 0.2, and a weight decay of 10^{-5}. Following the default schedule, the learning rate was annealed from zero to 0.2 over the first ten epochs. Similarly, for the last 100 epochs, the learning rate was reduced by a factor of 2 every ten epochs. The training occurred on batches of eight single-channel greyscale images with default image augmentation, which included random rotations, scaling, and translations. 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, e.g., 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 Supplement 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 and 64 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), to match the model predictions to the most similar ground-truth annotation. This is done by calculating the intersection over union (IoU) metric for each predicted object. Next, the 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. All reported mAP values in this study 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 (e.g., Lin et al., 2014; Rezatofigh et al., 2019), while the AP itself is calculated in a way that is typical for studies on biomedical images (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 et al. (2022a) and the PebbleCounts tool (Purinton and Bookhagen, 2019), which performs 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 et al., 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 and 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 emphasize here 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 segment all grains in an image in an automated way.

Finally, we inferred the style representation of the neural network (see section 2.2 above) for each of our 81 image tiles with our most generalist model (*full_set*). We used these 256-dimensional vectors for style clustering in a data-driven effort to find classes of similar image types to train style-specific models with improved segmentation capability. We followed here the general method of Pachitariu and Stringer (2022) by using the Leiden algorithm with 66 neighbors and a resolution of 0.8 for clustering (Traag et al., 2019) and t-distributed stochastic neighbor embedding (t-SNE) for visualization (Poličar et al., 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

We obtained delineated grains for each dataset using the segmentation model trained on the indicated data split (see Fig. 3) for inference. All predicted masks presented in the following sections are generated with default settings, including a minimum object diameter of 15 pixels ('min_size'). We rescaled each image for inference 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. In this study, we used for each grain 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). Any length measurements are converted from pixels to length units through the image-specific pixel resolution (Table S2). For all grains, the results are returned by default 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 modeled the grain size uncertainties for individual percentile values of GSDs (Eaton et al., 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 used this approach 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, i.e., the counting statistics. To account for further uncertainty introduced by measurements in images, we combined this bootstrapping with a one-dimensional error modeling. We accomplished this with randomization of each resampled axis (ax_i) with a length (ε_{length}) and scale (ε_{scale}) error component through:

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

The length error represents the measurement error along the axis length. It is by default implemented in the randomization as a normal distribution centered 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, i.e., through estimating the image resolution of an image. It is implemented by default in the randomization as a normal distribution centered 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 S3.

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).

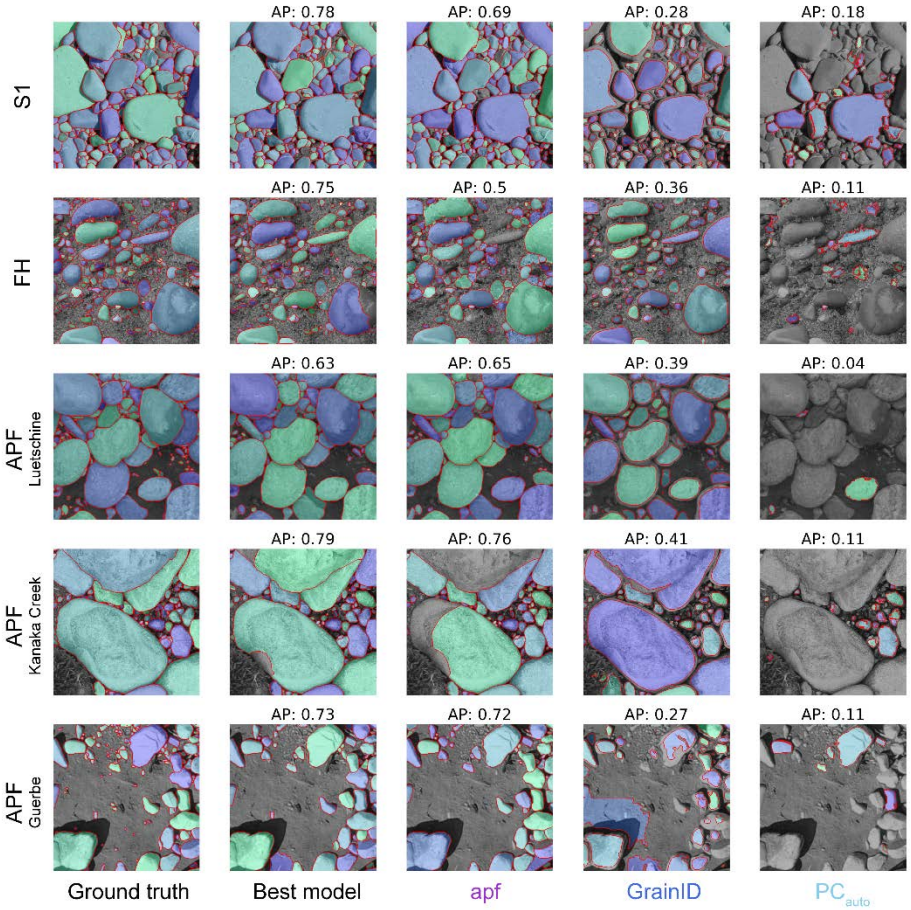

Figure 4: Examples of segmentations that resulted from our best performing models (S1: apf_s1, FH: fh+, and APF: full_set, respectively) compared to ground truth annotations, and results of a generalist model (*apf*), and segmentations of the benchmarks methods (GrainID; Chen et al., 2022a; and PCauto; Purinton et al., 2019) for selected test image tiles. AP = average precision at intersection over union threshold of 0.5 for the corresponding tile.

3.1.1 Overall performance

Our trained segmentation models are generally able to segment coarse sedimentary particles at high precision for all our datasets (for examples, see Fig. 4 and for all segmentation of the best performing models, we refer to Supplementary Fig. S6). While 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 Supplementary Figs. S7-9 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 ΔAP values range from 0.419 for our best model to 0.142 for our overall worst model when compared to the respective best benchmark model (Table 1). This means that our best model correctly segmented 42% more grains with an IoU of 0.5 or higher, while 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 (Fig. 5) and in each image tile (Fig. S6) than the other models. On a dataset level, our models performed better on specialized datasets, with average AP values of 0.753 (*full_set*) and 0.75 (*fh+*) in contrast to 0.633 (*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 perform similarly for the training and test sets (Supplementary Fig. S10). Only for *s1*, *s1+*, and *fh*, the models performed better for the training set, potentially exhibiting some overfitting to the training data.

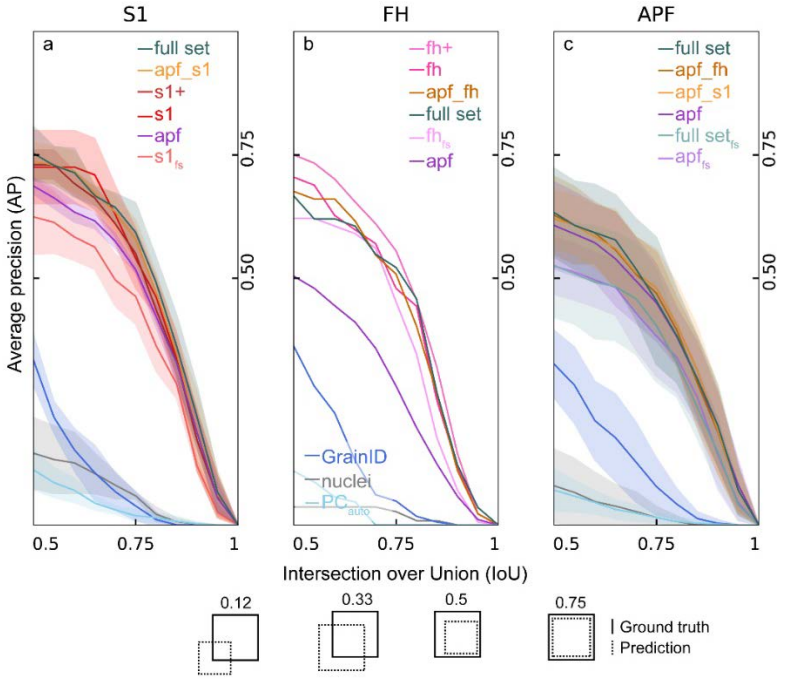

Figure 5: Segmentation performance for selected models applied on the image tiles that we used as test sets for the datasets S1 (a), FH (b), and APF (c). The lines represent the average precision that is averaged over the test tiles, while the shaded area represents the standard deviation (1 sigma). For the performance on individual train and test tiles and the average of the entire dataset, we refer to Supplementary Fig. S10. The segmentation performance of the benchmark methods (GrainID; Chen et al., 2022a; and PCauto; Purinton et al., 2019) are compared.

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

Table 1: Overall segmentation performance of our models and benchmarks for our test image tiles for all datasets. AP = average precision at the intersection over union (IoU) threshold of 0.5, averaged over the dataset; mAP = mean average precision over IoU threshold 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.

3.1.2 Systematic trends

Upon closer inspection, the segmentation results reveal systematic effects on the model performance from our data. First, using transfer learning, i.e., re-training the models from the *nuclei* model, improves the segmentation performance in all our datasets

compared to those we trained from scratch with the same data (denoted with the suffix '$_{fs}$'). 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 (Fig. 5; Table 1; see also Supplementary Fig. S9).

Second, the composition of the training data size and the content of the training data in combination with a re-training strategy (see Fig. 3 for an overview) had significant effects on the model's performance. Starting with the heterogeneous APF set, it is noteworthy that 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 training on the full dataset of 67 tiles (*full_set;* Fig. 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*; Fig. 5a), thereby even slightly outperforming more specialized models trained only on the S1 set (e.g., *s1*, *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, S1). In line with this, for the contrastingly different FH set, models that were trained only on the 6 tiles from the gravel pit outcrops (*fh*, *fh+*) performed better than 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% (Fig. 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, while for S1, such a two-fold training strategy yielded results that were inferior to 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 different 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 on the S1 training tiles (Supplementary Fig. S10).

3.1.3 Image style classes

Our results so far have shown that the composition of the training dataset systematically influences the segmentation performance in our data. By using the style vector used by the neural network and through clustering these, we obtained three distinct classes of image types (Fig. 6a), which were considered different by the model for segmentation. We found that these image classes consisted of (i) images with pebbles under sunny conditions with distinct shadows along granular interstices ("sunny pebbles"; SP), (ii) images 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 the segmentation models for which image tiles and ground truth masks were re-combined into datasets according to their image class. We evaluated the models (*ig*$_{mrg}$, *ig*$_{cv}$, and *ig*$_{sp}$) along with our generalist model on the respective datasets (Fig. 6b). We found that the style-specific models show either higher segmentation performance for two datasets (MRG and SP) or the same in one dataset (CV) when compared to our generalist *full_set* model. Moreover, this finding discloses that the segmentation performance of our models is lower for images with 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 (Fig. 6a). Therefore, we used the best-performing model on a dataset level for segmenting. We then measured the grains on these images in the following sections.

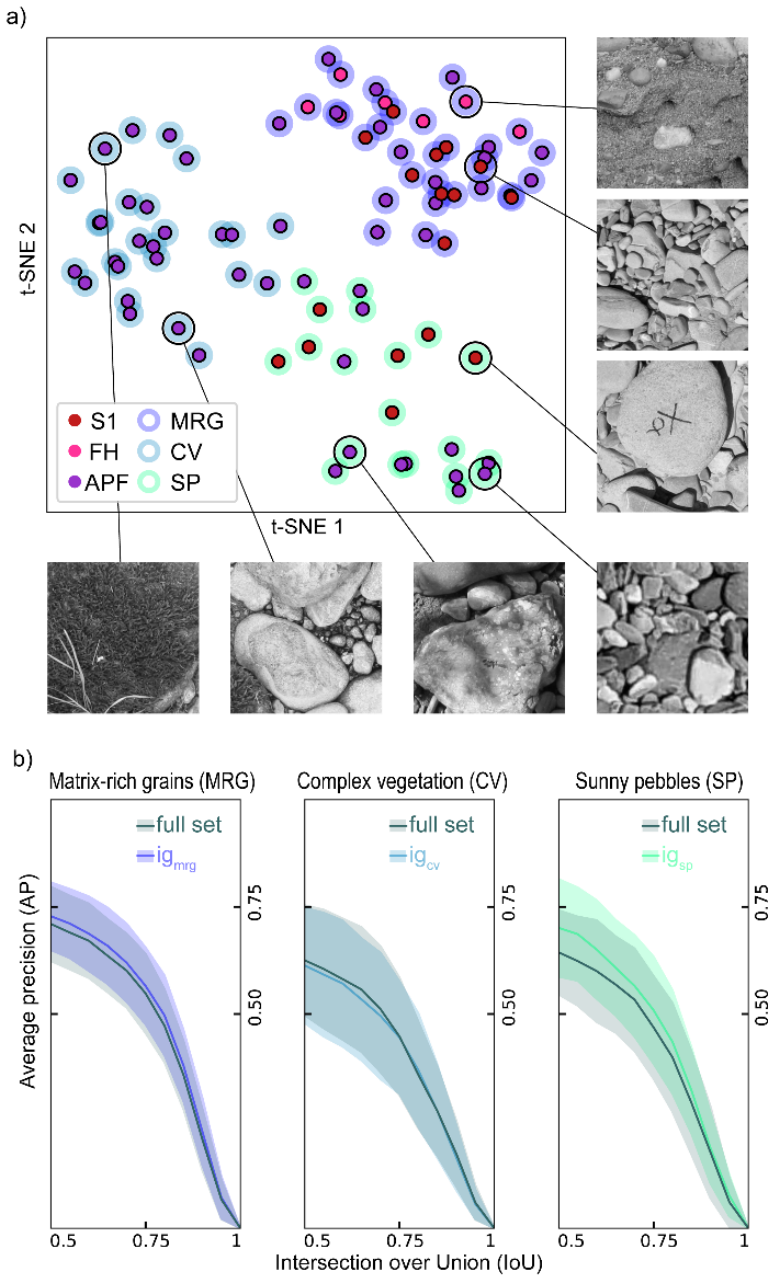

Figure 6: Classes of image types for segmentation inferred from the style vectors used by the neural network. (a) The clustering was visualized on the style vectors through t-distributed stochastic neighbor embedding (t-SNE) with image examples for each class (MRG, CV, and SP). Inner colors indicate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

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 images before measurements. Second, we report results after scaling, and we compared the data with grain sizes measured with different methods, including field measurements, in section 3.2.2. Finally, we present size distribution maps for 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, i.e., both ellipses and convex hulls yield similar size distributions (Figs. S12-15). In addition, when comparing to the ground truth data (Fig. 6), 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 78.6% (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 (Fig. 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 Figs. S12-15). However, overall high accuracy and precision in S1 and FH and across all percentile values are also evident when comparing individual key percentile values (Fig. S11). 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 (Fig. 7; Table 2).

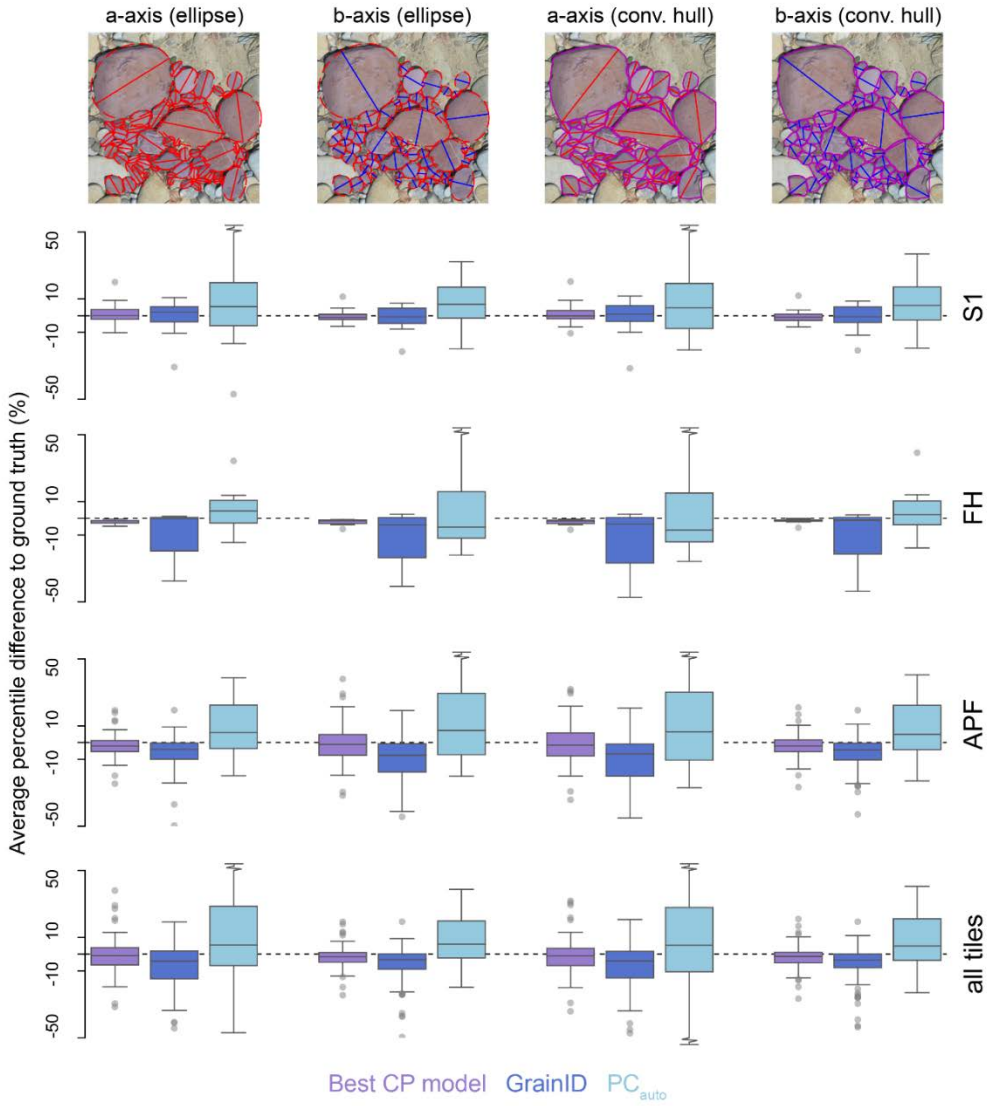

Figure 7: Overall quality of grain size data collected with different methods in image tiles. The quality is quantified by the closeness of predictions to ground truth for different grain size proxies for all tiles and the respective dataset splits (S1: 18, FH: 7, and APF: 56). We report the average 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 Supplementary Fig. S11. 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.

Grain size metric	Data	Percentage of tiles (%) identical with GT			Average percentile difference (px) to GT		
		Best model	GrainID	PC auto	Best CP model	GrainID	PebbleCounts auto
	S1	94.4	38.9	33.3	1.3 ± 9.7	-0.4 ± 15.5	6.8 ± 32.5

a-axis (ellipse)	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

Table 2: Statistical summary of the closeness between predictions and ground truth (GT) for different grain size proxies across the image tiles datasets. The table also 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, i.e., only grains with b-axes ≥ 12px and with a center-point situated within the central 64% of the image tile are taken into account (for visual reference, see Fig. 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 report it along with the associated 1σ standard deviation. We refer to Supplementary Figs. S12-S15 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 (Figs. 8a, b), which we compare to the data collected independently, in the field (K1, S1) and manually in images (FH; Figs. 8c, d). We use only the ellipse approximation since 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 results. All size distributions for our grain sizes 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 in the images are much closer to the reference data than others. Specifically, for the Kander site, the lengths of the measured b-axis represent the field data perfectly where the sizes differ by < 10% for all percentiles (Fig. 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 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 more significant differences that are evened out (Fig. 8d). Nevertheless, the differences between the percentile values of data collected from the prediction masks and the percentiles from the reference dataset never exceed ± 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 (Fig. 8d). However, the lengths of the a-axis are over-estimated 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 was collected.

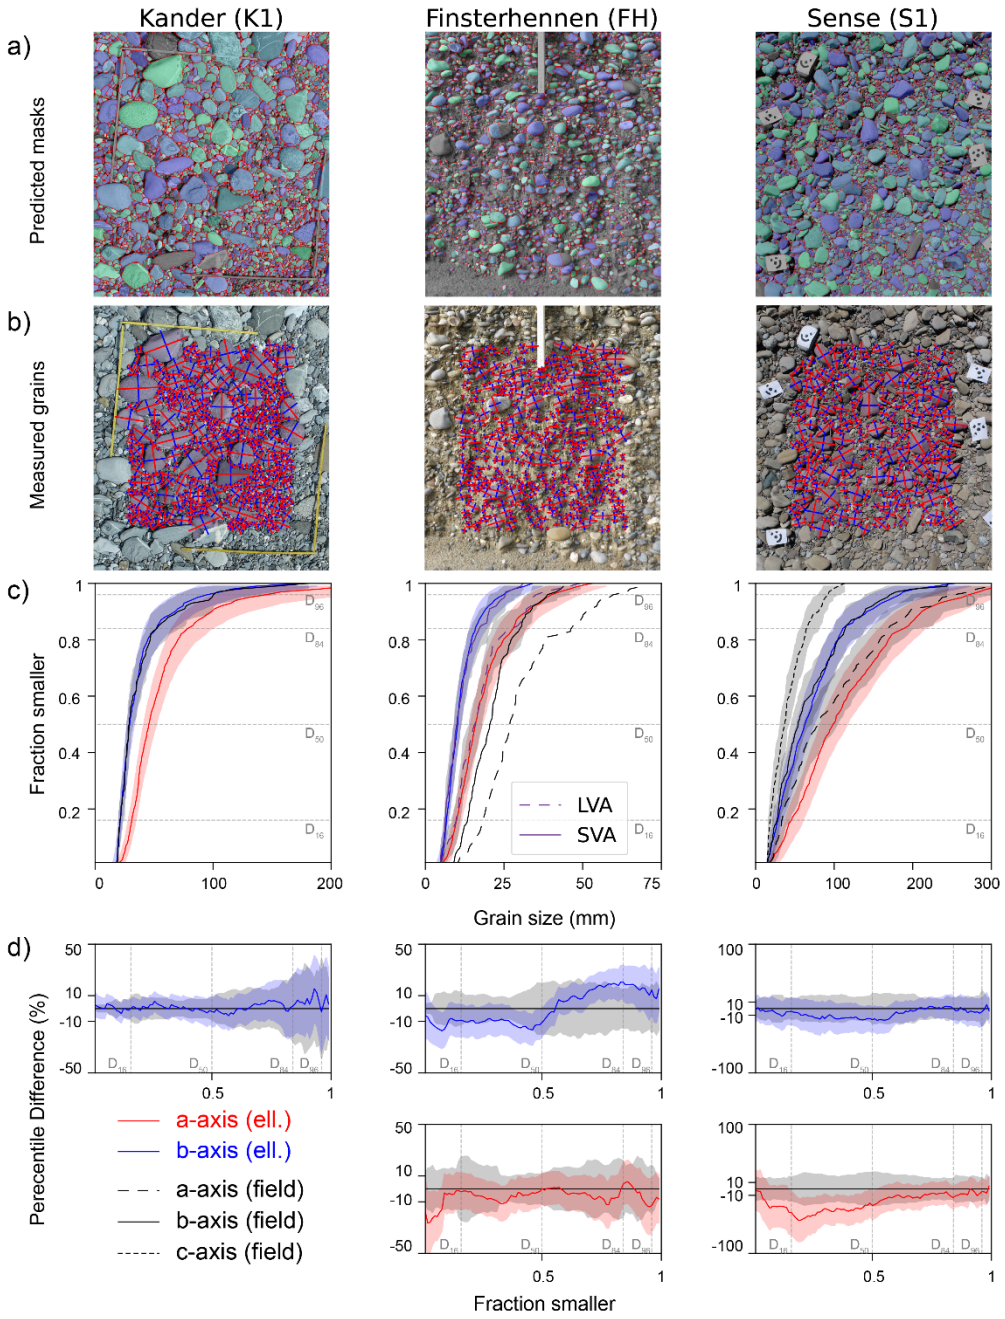

Figure 8: Comparison between grain sizes for 3 regions where data was collected on images and in the field. The predicted grain masks (a) of the best performing models (*full_set* for K1 and S1, *fh_boosted* for FH;) are filtered to represent the same area measured independently. Additionally, results 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).

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

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

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 automatically investigated in two orthoimage mosaics (Fig. 9), which were generated from close-range UAV surveys. This approach automatically delineated and measured > 268000 individual grains for site S1 and >143000 for site K1. The results disclose a high variability of grain sizes across the bars, as exemplified by the local variation in the D_{50} (Fig. 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 Fig 9a). These variations in the grain sizes are generally more significant than the uncertainty within the local images (see, e.g., Fig. 9b).

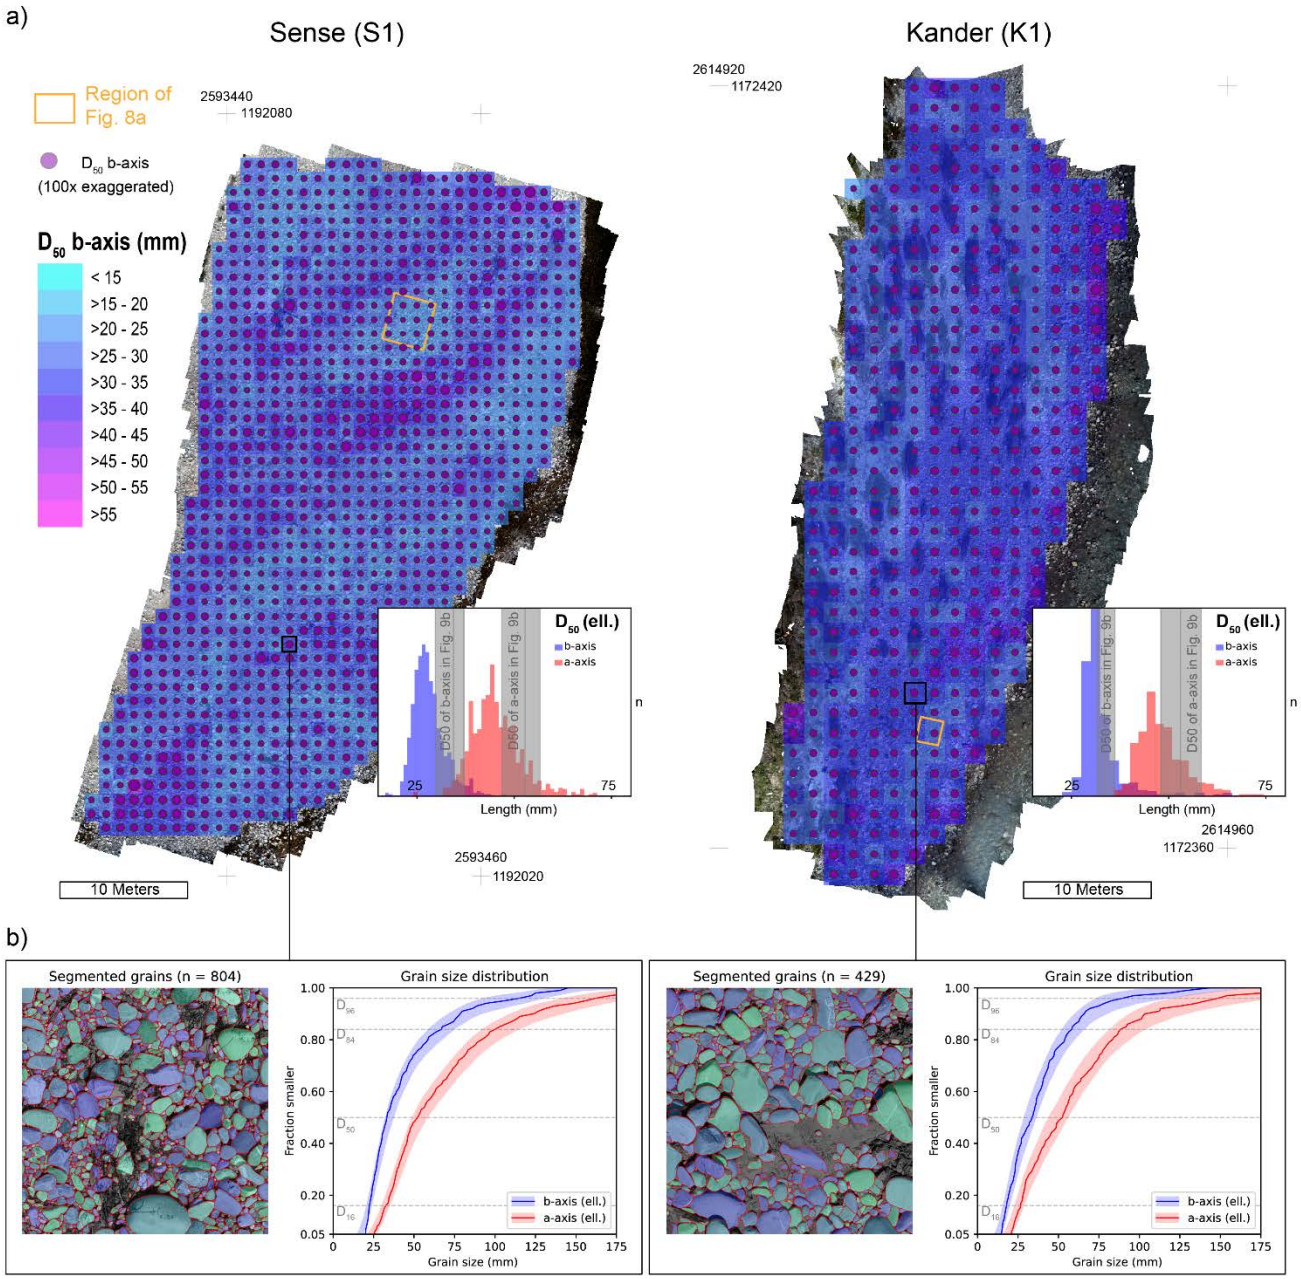

Figure 9: Gain sizes of the surface layer of entire gravel bars measured with ImageGrains for two sites at rivers in Switzerland. (a) Maps of the local D_{50} for the b-axis were obtained for subsets of the orthoimage mosaic with a size of 1000x1000 pixels. The frequency histograms (insert) show high variability in the local D_{50}. (b) Example of segmentation mask 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.

4 Discussion

Our results show that transfer learning allows training 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 higher number of observations than what can be achieved with the other benchmark methods. Furthermore, we can achieve these excellent results with relatively small datasets compared to other approaches (e.g., Chen et al., 2022a). For example, we used less than 1000 objects from only seven tiles for the training (Table S2) 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 likewise 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 datasets, thereby documenting the potential of a data-driven approach for this type of analysis (section 4.2). Furthermore, precise and automated segmentations allow for 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 and their limits. First, with our ImageGrains software library, our workflow can be directly applied to segment and measure coarse sedimentary particles in a wide 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 et al., 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 mask of our models might also form the base for other machine-learning applications, e.g., the training of a classifier to identify the particles' petrography. All these applications can be realized on a regular stand-alone computer without specialized knowledge of coding with machine learning libraries. In particular, the segmentation model architecture allows fast training and inference on a desktop PC with a consumer-grade GPU. Our grain size analysis with ImageGrains allows for measurements with a similar speed where thousands of grains can be analyzed within minutes. The Cellpose model architecture itself is rigorously tested (Stringer et al., 2021; Pachitariu and Stringer, 2022), 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 Supplement S5). 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., 2022; Sun et al., 2017). For our case, overfitting can be prevented by using appropriate ground truth data for testing and choosing a suitable training strategy (see Supplement S5 and section 4.2 below). To ensure that the model has learned to segment grains on images correctly, the results should be compared to 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 segmenting grains. We emphasize this because, despite the overall excellent results of our models, they failed to predict grains for some challenging image tiles (Fig. S6). In summary, we find that on a technical level, primarily the data used for training conditions, the capability of the model to segment grains. Therefore, 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, e.g., of 4.7 to 18.1 mm for the images

in Fig. 8a. While lower cut-off values are possible, we opted for the more conservative value. The reasons for our 12 px cut-off are that i) we found it hard to delineate grains smaller than those of the ground truth visually, and ii) the model rescales images during inference and training in the configuration we use (see S5 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 be why other approaches were based on similar cut-off values of 20 pixels (Chen et al., 2022a; Purinton and Bookhagen, 2019). Accordingly, measuring small grains has remained challenging for fluvial settings (e.g., Carbonneau et al., 2018; Marchetti et al., 2021; 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 (Structure from Motion / Multi-View Stereo) methods (e.g., James et al., 2020; James et al., 2019). 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 and 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., 2017). 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., 3.1.3, and Fig. 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 model performance (Pachitariu and Stringer; 2022). For example, we notice that while using more data for training did improve the segmentation for the S1 tiles, it did not improve the segmentation performance for the FH tiles (Fig. 5). A direct consequence is for selecting datasets, it is, therefore, essential to pay attention to dataset balance and composition, concerning image content and visual complexity, e.g., 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. Suppose the inferred image style differs from the data we used for training. In that case, a custom model will likely exhibit a better segmentation performance than existing models. This can also point to the kind of data split, which might be most promising for training, i.e., our full set or a specific class of image style. Related to the data composition, we find that the most effective re-training strategy also 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 (Fig. 6a). Thus, datasets with similar embedded distributions might benefit from the same training strategy 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 data set. 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 fluvial grain size measurements

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. Due to 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 and 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 (Fig. 8), which was accomplished by a grid-by-number approach. Such method-specific traits become even more critical when analyzing 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 (Fig. 9). Similar trends of locally high variability in grain size distributions have also been observed in other field surveys (e.g., Gómez et al., 2022; Chardon et al., 2020); Rice and Church, 1998). In addition, spatial differences in sedimentary patterns, e.g., vertical and lateral sorting and/or armoring (see also Bunte and 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 Figs. 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 base 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 uncertainty estimation methods for various image types, including UAV-derived imagery, nadir images, orthophoto mosaics, and photographs from handheld cameras.

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

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.

Data availability

The images 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.

Author contribution

DM conceptualized the research and developed the code together with GW. DM, AHDP, and PG did data collection in the field. DM was responsible for the data curation, which includes creating the ground truth annotations. DM interpreted the results with scientific input from GW and FS, and prepared the paper and figures with contributions from all co-authors.

Competing interests

The authors declare that they have no conflict of interest.

Acknowledgments

We thank 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.

References

Allen, P. A., Michael, N. A., D'Arcy, M., Roda-Boluda, D. C., Whittaker, A. C., Duller, R. A., and Armitage, J. J.: Fractionation of grain size in terrestrial sediment routing systems, Basin Res., 29, 180–202, https://doi.org/10.1111/bre.12172, 2017.

Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M. A., Al-Amidie, M., and Farhan, L.: Review of deep learning: concepts, CNN architectures, challenges, applications, future directions, J. Big Data, 8, 53, https://doi.org/10.1186/s40537-021-00444-8, 2021.

Arzt, M., Deschamps, J., Schmied, C., Pietzsch, T., Schmidt, D., Tomancak, P., Haase, R., and Jug, F.: LABKIT: Labeling and Segmentation Toolkit for Big Image Data, Front. Comput. Sci., 4, 1–12, https://doi.org/10.3389/fcomp.2022.777728, 2022.

Attal, M., Mudd, S. M., Hurst, M. D., Weinman, B., Yoo, K., and Naylor, M.: 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 Surf. Dyn., 3, 201–222, https://doi.org/10.5194/esurf-3-201-2015, 2015.

Bradski, G.: The openCV library. Dr. Dobb's Journal: Software Tools for the Professional Programmer, 25(11), 120-123, 2000.

Brayshaw, D. D.: Bankfull and effective discharge in small mountain streams of British Columbia, https://doi.org/10.14288/1.0072555, 2012.

Bunte, K. and Abt, S. R.: Sampling Surface and Subsurface Particle-Size Distributions in Wadable Gravel- and Cobble-Bed Streams for Analyses in Sediment Transport , Hydraulics , and Streambed Monitoring, 450, 2001.

Buscombe, D.: Transferable wavelet method for grain-size distribution from images of sediment surfaces and thin sections, and other natural granular patterns, Sedimentology, 60, 1709–1732, https://doi.org/10.1111/sed.12049, 2013.

Buscombe, D.: SediNet: a configurable deep learning model for mixed qualitative and quantitative optical granulometry, Earth Surf. Process. Landforms, 45, 638–651, https://doi.org/10.1002/esp.4760, 2020.

Butler, J. B., Lane, S. N., and Chandler, J. H.: Automated extraction of grain-size data from gravel surfaces using digital image processing, J. Hydraul. Res., 39, 519–529, https://doi.org/10.1080/00221686.2001.9628276, 2001.

Caicedo, J. C., Goodman, A., Karhohs, K. W., Cimini, B. A., Ackerman, J., Haghighi, M., Heng, C. K., Becker, T., Doan, M., McQuin, C., Rohban, M., Singh, S., and Carpenter, A. E.: Nucleus segmentation across imaging experiments: the 2018 Data Science Bowl, Nat. Methods, 16, 1247–1253, https://doi.org/10.1038/s41592-019-0612-7, 2019.

Carbonneau, P. E. and Dietrich, J. T.: Cost-effective non-metric photogrammetry from consumer-grade sUAS: implications for direct georeferencing of structure from motion photogrammetry, Earth Surf. Process. Landforms, 42, 473–486, https://doi.org/10.1002/esp.4012, 2017.

Carbonneau, P. E., Lane, S. N., and Bergeron, N. E.: Catchment-scale mapping of surface grain size in gravel bed rivers using airborne digital imagery, Water Resour. Res., 40, 1–11, https://doi.org/10.1029/2003WR002759, 2004.

Carbonneau, P. E., Bizzi, S., and Marchetti, G.: Robotic photosieving from low-cost multirotor sUAS: a proof-of-concept, Earth Surf. Process. Landforms, 43, 1160–1166, https://doi.org/10.1002/esp.4298, 2018.Castelltort, S. and Van Den Driessche, J.: How plausible are high-frequency sediment supply-driven cycles in the stratigraphic record?, Sediment. Geol., 157, 3–13, https://doi.org/10.1016/S0037-0738(03)00066-6, 2003.

Chardon, V., Schmitt, L., Piégay, H., and Lague, D.: Use of terrestrial photosieving and airborne topographic LiDAR to assess bed grain size in large rivers: a study on the Rhine River, Earth Surf. Process. Landforms, 45, 2314–2330, https://doi.org/10.1002/esp.4882, 2020.

Chardon, V., Piasny, G., and Schmitt, L.: Comparison of software accuracy to estimate the bed grain size distribution from digital images: A test performed along the Rhine River, River Res. Appl., 38, 358–367, https://doi.org/10.1002/rra.3910, 2022.

Chen, X., Hassan, M. A., and Fu, X.: Convolutional neural networks for image-based sediment detection applied to a large terrestrial and airborne dataset, Earth Surf. Dyn., 10, 349–366, https://doi.org/10.5194/esurf-10-349-2022, 2022a.

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, Z., Scott, C., Keating, D., Clarke, A., Das, J., and Arrowsmith, R.: Quantifying and analysing rock trait distributions of rocky fault scarps using deep learning, Earth Surf. Process. Landforms, 1–17, https://doi.org/10.1002/esp.5545, 2023.

Coelho, L. P., Shariff, A., and Murphy, R. F.: NUCLEAR SEGMENTATION IN MICROSCOPE CELL IMAGES: A HAND-SEGMENTED DATASET AND COMPARISON OF ALGORITHMS., Proceedings. IEEE Int. Symp. Biomed. Imaging, 5193098, 518–521, https://doi.org/10.1109/ISBI.2009.5193098, 2009.

Detert, M. and Weitbrecht, V.: Automatic object detection to analyze the geometry of gravel grains - A free stand-alone tool, River Flow 2012 - Proc. Int. Conf. Fluv. Hydraul., 1, 595–600, 2012.

Dunne, K. B. J. and Jerolmack, D. J.: Evidence of, and a proposed explanation for, bimodal transport states in alluvial rivers, Earth Surf. Dyn., 6, 583–594, https://doi.org/10.5194/esurf-6-583-2018, 2018.

Eaton, B. C., Moore, R. D., and MacKenzie, L. G.: Percentile-based grain size distribution analysis tools (GSDtools) – estimating confidence limits and hypothesis tests for comparing two samples, Earth Surf. Dyn., 7, 789–806, https://doi.org/10.5194/esurf-7-789-2019, 2019.

Garefalakis, P., do Prado, A., Mair, D., Douillet, G. A., Nyffenegger, F., and Schlunegger, F.: Comparison of three grain size measuring methods applied to coarse-grained gravel deposits, Sediment. Geol., 106340, https://doi.org/10.1016/j.sedgeo.2023.106340, 2023.

Gatys, L. A., Ecker, A. S., and Bethge, M.: Image Style Transfer Using Convolutional Neural Networks, Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit., 2016-December, 2414–2423, https://doi.org/10.1109/CVPR.2016.265, 2016.

Guerit, L., Barrier, L., Liu, Y., Narteau, C., Lajeunesse, E., Gayer, E., and Métivier, F.: 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 Surf. Dyn., 6, 1011–1021, https://doi.org/10.5194/esurf-6-1011-2018, 2018.

Hunter, J. D.: Matplotlib: A 2D Graphics Environment, Comput. Sci. Eng., 9, 90–95, https://doi.org/10.1109/MCSE.2007.55, 2007.

Karras, T., Laine, S., and Aila, T.: A Style-Based Generator Architecture for Generative Adversarial Networks, IEEE Trans. Pattern Anal. Mach. Intell., 43, 4217–4228, https://doi.org/10.1109/TPAMI.2020.2970919, 2021.

Kluyver, T., et al.: Jupyter Notebooks—a publishing format for reproducible computational workflows, Position. Power Acad. Publ. Play. Agents Agendas - Proc. 20th Int. Conf. Electron. Publ. ELPUB 2016, 87–90, https://doi.org/10.3233/978-1-61499-649-1-87, 2016.

Kumar, N., et al.: A Multi-Organ Nucleus Segmentation Challenge, IEEE Trans. Med. Imaging, 39, 1380–1391, https://doi.org/10.1109/TMI.2019.2947628, 2020.

Lam, S. K., Pitrou, A., and Seibert, S.: Numba: A LLVM-based Python JIT Compiler, Proc. LLVM-HPC 2015 2nd Work. LLVM Compil. Infrastruct. HPC - Held conjunction with SC 2015 Int. Conf. High Perform. Comput. Networking, Storage Anal., 2015-January, https://doi.org/10.1145/2833157.2833162, 2015.

Lang, N., Irniger, A., Rozniak, A., Hunziker, R., Wegner, J. D., and Schindler, K.: GRAINet: mapping grain size distributions in river beds from UAV images with convolutional neural networks, Hydrol. Earth Syst. Sci., 25, 2567–2597, https://doi.org/10.5194/hess-25-2567-2021, 2021.

LI, G., LIU, T., NIE, J., GUO, L., CHEN, J., ZHU, J., XIA, W., MARA, A., HOLLEY, S., and WONG, S. T. C.: Segmentation of touching cell nuclei using gradient flow tracking, J. Microsc., 231, 47–58, https://doi.org/10.1111/j.1365-2818.2008.02016.x, 2008.

Li, Y., Onasch, C. M., and Guo, Y.: GIS-based detection of grain boundaries, J. Struct. Geol., 30, 431–443, https://doi.org/10.1016/j.jsg.2007.12.007, 2008.

Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L.: 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, 2014.

Litty, C. and Schlunegger, F.: Controls on Pebbles' Size and Shape in Streams of the Swiss Alps, J. Geol., 125, 101–112, https://doi.org/10.1086/689183, 2017.

Lu, J., Behbood, V., Hao, P., Zuo, H., Xue, S., and Zhang, G.: Transfer learning using computational intelligence: A survey, Knowledge-Based Syst., 80, 14–23, https://doi.org/10.1016/j.knosys.2015.01.010, 2015.

Mair, D., Do Prado, A. H., Garefalakis, P., Lechmann, A., Whittaker, A., and Schlunegger, F.: Grain size of fluvial gravel bars from close-range UAV imagery – uncertainty in segmentation-based data, Earth Surf. Dyn., 10, 953–973, https://doi.org/10.5194/esurf-10-953-2022, 2022a.

Mair, D., Henrique, A., Prado, D., Garefalakis, P., Lechmann, A., Whittaker, A., and Schlunegger, F.: 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, 2022b.

McKinney, W.: Data Structures for Statistical Computing in Python, Proc. 9th Python Sci. Conf., 1, 56–61, https://doi.org/10.25080/majora-92bf1922-00a, 2010.

Mörtl, C., Baratier, A., Berthet, J., Duvillard, P.-A., and De Cesare, G.: 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, 2022.

Pachitariu, M. and Stringer, C.: Cellpose 2.0: how to train your own model, Nat. Methods, 19, 1634–1641, https://doi.org/10.1038/s41592-022-01663-4, 2022.

Paszke, A.,et al.: PyTorch: An Imperative Style, High-Performance Deep Learning Library, Proceedings of the 33rd International Conference on Neural Information Processing Systems, 721, 8026–8037, 2019.

Piégay, H., Arnaud, F., Belletti, B., Bertrand, M., Bizzi, S., Carbonneau, P., Dufour, S., Liébault, F., Ruiz-Villanueva, V., and Slater, L.: Remotely sensed rivers in the Anthropocene: state of the art and prospects, Earth Surf. Process. Landforms, 45, 157–188, https://doi.org/10.1002/esp.4787, 2020.

Poličar, P. G., Stražar, M., and Zupan, B.: OpenTSNE: A modular Python library for t-SNE dimensionality reduction and embedding, 1–2, https://doi.org/10.1101/731877, 2019.

Purinton, B. and Bookhagen, B.: Introducing PebbleCounts: a grain-sizing tool for photo surveys of dynamic gravel-bed rivers, Earth Surf. Dyn., 7, 859–877, https://doi.org/10.5194/esurf-7-859-2019, 2019.

Purinton, B. and Bookhagen, B.: Tracking Downstream Variability in Large Grain-Size Distributions in the South-Central Andes, J. Geophys. Res. Earth Surf., 126, 1–29, https://doi.org/10.1029/2021JF006260, 2021.

Qi, C. R., Liu, W., Wu, C., Su, H., and Guibas, L. J.: Frustum PointNets for 3D Object Detection from RGB-D Data, Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit., 918–927, https://doi.org/10.1109/CVPR.2018.00102, 2018.

Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., and Prabhat: Deep learning and process understanding for data-driven Earth system science, Nature, 566, 195–204, https://doi.org/10.1038/s41586-019-0912-1, 2019.

Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., and Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression, Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit., 2019-June, 658–666, https://doi.org/10.1109/CVPR.2019.00075, 2019.

Ronneberger, O., Fischer, P., and Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation, IEEE Access, 9, 16591–16603, https://doi.org/10.48550/arXiv.1505.04597, 2015.

Schindelin, J., et al..: Fiji: An open-source platform for biological-image analysis, Nat. Methods, 9, 676–682, https://doi.org/10.1038/nmeth.2019, 2012.

Schlunegger, F. and Norton, K. P.: 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 Res., 27, 230–245, https://doi.org/10.1111/bre.12070, 2015.

Schmidt, U., Weigert, M., Broaddus, C., and Myers, G.: Cell detection with star-convex polygons, Springer International Publishing, 265–273 pp., https://doi.org/10.1007/978-3-030-00934-2_30, 2018.

Soloy, A., Turki, I., Fournier, M., Costa, S., Peuziat, B., and Lecoq, N.: A deep learning-based method for quantifying and mapping the grain size on pebble beaches, Remote Sens., 12, 1–23, https://doi.org/10.3390/rs12213659, 2020.

Steer, P., Guerit, L., Lague, D., Crave, A., and Gourdon, A.: Size, shape and orientation matter: fast and semi-automatic measurement of grain geometries from 3D point clouds, Earth Surf. Dyn., 10, 1211–1232, https://doi.org/10.5194/esurf-10-1211-2022, 2022.

Stringer, C., Wang, T., Michaelos, M., and Pachitariu, M.: Cellpose: a generalist algorithm for cellular segmentation, Nat. Methods, 18, 100–106, https://doi.org/10.1038/s41592-020-01018-x, 2021.

Sulaiman, M. S., Sinnakaudan, S. K., Ng, S. F., and Strom, K.: Application of automated grain sizing technique (AGS) for bed load samples at Rasil River: A case study for supply limited channel, 121, 330–343, https://doi.org/10.1016/j.catena.2014.05.013, 2014.

Sun, C., Shrivastava, A., Singh, S., and Gupta, A.: 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, 2017.

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R.: Intriguing properties of neural networks, 2nd Int. Conf. Learn. Represent. ICLR 2014 - Conf. Track Proc., 1–10, 2014.

Tofelde, S., Bernhardt, A., Guerit, L., and Romans, B. W.: Times Associated With Source-to-Sink Propagation of Environmental Signals During Landscape Transience, Front. Earth Sci., 9, 1–26, https://doi.org/10.3389/feart.2021.628315, 2021.

Traag, V. A., Waltman, L., and van Eck, N. J.: From Louvain to Leiden: guaranteeing well-connected communities, Sci. Rep., 9, 5233, https://doi.org/10.1038/s41598-019-41695-z, 2019.

Van der Maaten, L. and Hinton, G.: Visualizing Data using t-SNE, J. Mach. Learn. Res., 9, 2579–2605, 2008.

Van Der Walt, S., Colbert, S. C., and Varoquaux, G.: The NumPy array: A structure for efficient numerical computation, Comput. Sci. Eng., 13, 22–30, https://doi.org/10.1109/MCSE.2011.37, 2011.

Van Der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., and Yu, T.: Scikit-image: Image processing in python, PeerJ, 2014, 1–18, https://doi.org/10.7717/peerj.453, 2014.

Virtanen, P., et al.: SciPy 1.0: fundamental algorithms for scientific computing in Python, Nat. Methods, 17, 261–272, https://doi.org/10.1038/s41592-019-0686-2, 2020.

Whittaker, A. C., Attal, M., and Allen, P. A.: Characterising the origin, nature and fate of sediment exported from catchments perturbed by active tectonics, Basin Res., 22, 809–828, https://doi.org/10.1111/j.1365-2117.2009.00447.x, 2010.

Woodget, A. S. and Austrums, R.: Subaerial gravel size measurement using topographic data derived from a UAV-SfM approach, Earth Surf. Process. Landforms, 42, 1434–1443, https://doi.org/10.1002/esp.4139, 2017.

Yosinski, J., Clune, J., Bengio, Y., and Lipson, H.: How transferable are features in deep neural networks?, Adv. Neural Inf. Process. Syst., 4, 3320–3328, 2014.

Zech, J. R., Badgeley, M. A., Liu, M., Costa, A. B., Titano, J. J., and Oermann, E. K.: Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study, PLoS Med., 15, 1–17, https://doi.org/10.1371/journal.pmed.1002683, 2018.

Supplement S1: ImageGrains workflow

We built a software library in Python v.3, which we streamlined for automated grain measurements of coarse sedimentary particles ("pebbles"). We used segmentation masks of grain instances as a basis for this. The library consists of three main modules, roughly representative of the three main steps in the principal workflow (Fig. A1). Here, we used the Cellpose algorithm (Stringer et al., 2021) as backbone model for segmentation. At its core, the grain size module is built with functions of the scikit-image package (van der Walt et al., 2014), while our percentile-based uncertainty estimation is mainly built on the approach of Mair et al. (2022a), which we complemented with functionalities proposed by Eaton et al. (2019), and Purinton and Bookhagen (2021). ImageGrains can be executed locally via jupyter notebooks, enabling an online execution in a cloud environment (e.g., Google Colab). We also released a set of jupyter notebooks, allowing a step-by-step analysis with more custom options and instructions for image pre-processing and custom model training. In general, ImageGrains is designed to work with sets of images in provided file locations or with individual image files. The whole processing chain, or parts of it, can be readily included in custom scripts.

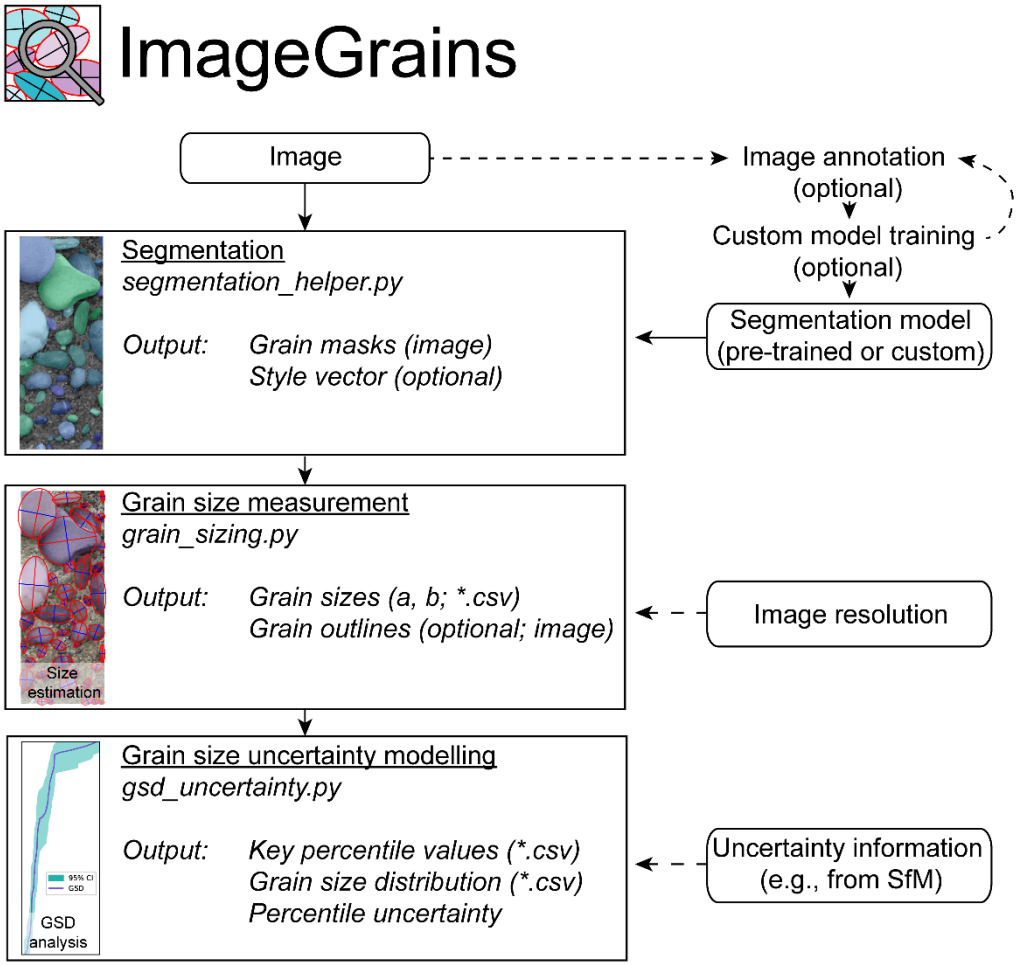

Figure S1: Principal of the ImageGrains workflow with main Python modules and their inputs and outputs.

Set	Site	Image	Resolution (mm px^{-1})	Tiles	Grains	Train	Test	Reference
S1	Sense	S1_a_003_2	0.53	6	441	401	40	this study
		S1_b_103_2	0.81	1	232	232	0	
		S1_b_289_3	0.75	1	260	260	0	
		S1_g_171_2	0.76	4	381	199	182	
		S1_g_208_2	0.80	6	794	642	152	
			S1 total	**18**	**2108**	**1734**	**374**	
FH	Finsterhennen	1_P1060330	0.29	1	125	125	0	Garefalakis et al. (2023)
		2_P1060338	0.21	1	47	47	0	
		3_P1060343	0.21	1	191	191	0	
		4_P1060348	0.22	1	153	0	153	
		5_P1060351	0.18	1	169	169	0	
		6_P1060355	0.17	1	137	137	0	
		7_P1060359	0.17	1	136	136	0	
			FH total	**7**	**958**	**805**	**153**	
APF	Deer Creek	deer_creek_004	~0.3	5	433	318	115	Chen et al. (2022), Brayshaw (2012)
	unknown	DS8-S2	-	5	348	232	116	
	Hidden Creek	hiddencreek_003	~0.3	5	396	358	38	
	Kanaka Creek	kanakacreek_020	~0.3	5	222	148	74	
	Kander	K1_2_C1_DJI_0385_clipped	1.50	7	279	207	72	Mair et al. (2022a)
		K1_2_C1_DJI_0385	1.51	1	320	320	0	
		K1_2_C1_DJI_0449	1.51	1	421	421	0	
	Luetschine	L2_1_C1_DJI_0352	1.47	3	846	846	0	
		L2_1_C1_DJI_0371	1.21	3	276	157	119	
	Entle	S9_5_C1_DJI_0043	1.49	1	579	579	0	
		S9_5_C1_DJI_0102	1.33	5	492	446	46	
		S9_6_C1_DJI_0677	1.47	3	284	284	0	
		S9_6_C1_DJI_0730	1.45	3	465	406	59	
	Guerbe	GU_Bar_6	1.00	1	161	161	0	this study, Litty & Schlunegger (2017)
		GU_Bar_7	1.00	2	772	643	129	
	Luetschine	Luetschine_1652	-	2	146	146	0	
	Rhine	Rhein_1679	-	2	91	91	0	
	Simme	Simme_1633	-	1	29	29	0	
	Thur	Thur_1781	-	1	1037	1037	0	
			APF total	**56**	**7597**	**6829**	**768**	
			Total	**81**	**10663**	**9368**	**1295**	

Table S2: Overview of the dataset of the images and image tiles that we used with the respective train and test splits as indicated.

Survey	K1_2_C1	S1
Flight height nominal (m)	6	5
Ground sampling distance (mm)	1.6	1.4
Number of images used	400	781
Image acquisition format	JPEG	JPEG
Number of used GCPs	16	43
Control Point RMSE (mm)		
X, Y	13.1	10.1
Z	95.3	16.7
Check Point RMSE (mm)		
X, Y	9.4	10
Z	106	19.3
Tie Point Observation Distance (m)		
Mean	4.92	4.07
Std	0.37	0.62
Sparse point cloud precision (mm)		
Mean Z	9.7	8.8
Std Z	6.9	6.8
Image error (px)	2	0.7
Reprojection error RMSE (px)	1	0.7
Doming amplitude (m)	0.46	-0.03
Orthophoto mosaic resolution (mm px^{-1})	1.59	1.11

Table S3: Summary of the UAV surveys and resulting SfM model uncertainties. RMSE = root mean square error, Std = standard deviation, px = pixel.

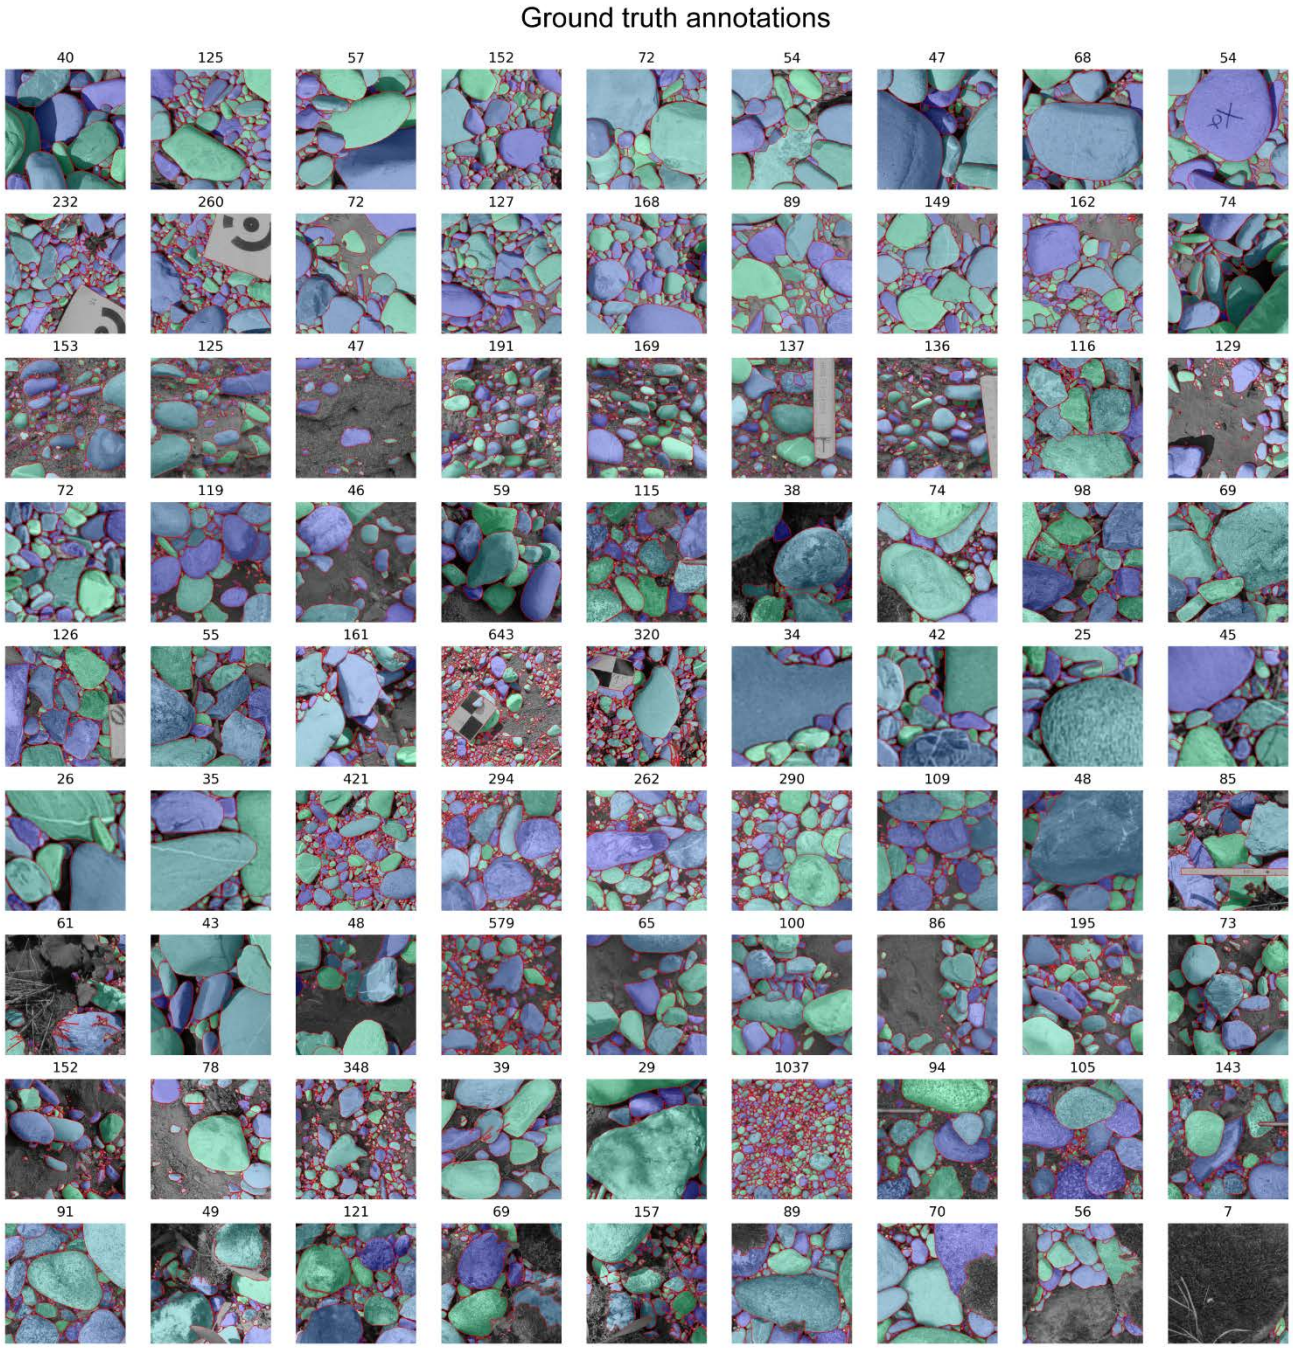

740 **Figure S4:** Ground truth annotations of coarse sedimentary grains (“pebbles”) used in this study. The caption represents the number of grains per image.

Supplement S5: The effect of different model configurations for training and inference

We tested several configurations for our models to assess some model aspects that might influence the performance of the Cellpose architecture upon application to sedimentary images. For an overview of how hyper-parameters or test-time modifications affect the architecture, we refer to Stringer et al. (2021). Specifically, we considered the object scales because Cellpose models re-scale images during training with the mean diameter of the objects in the ground truth. In retrained models, we used the mean diameter of the original dataset (in the case of the nuclei model: 17 pixels). Because grain size data often follows a lognormal distribution, we suspected that this might pose two challenges for our case: 1) Scaling data that has a lognormal distribution with its mean in combination with a minimum object size might lead to a loss of information, especially for the smallest grains. To test if this was the case, we trained models in a scale-invariant manner, i.e., by not rescaling images during training to the mean object diameter. 2) During the training, specifically at the image augmentation step, a default scale range effectively rescales images in a range of 0.75 to 1.25 of the input image size. This range might be too narrow for our data, which has a much more extensive size range than the original nuclei set, partly due to the underlying logarithmic distribution (see above). To test whether this is affecting the segmentation capability for our case, we modified the training so that rescaling was done with the median object size instead of the mean. We additionally increased the scale range during training to 0.25 to 1.75 of the input image size. Finally, we also tested whether an increase in the number of epochs would improve the segmentation performance. Consequently, we had to run inference for all models with and without a minimum object diameter (which is by default set to 15 pixels for circular approximation and not to confuse with our filter of 12 pixels for the b-axis of approximated ellipses) and with and without an automated size estimation for rescaling during prediction. For predictions without automated re-scaling, we used each dataset's respective mean object size diameter for rescaling.

In Table S5, we report the key outcomes of these tests in the form of the overall segmentation performance for those setups that came closest and farthest to the one used in the study. We note that the default configurations yielded the best results for our data. An exception was the number of epochs, for which 1000 epochs of training are better than 500. However, scale-invariant and modified models showed only slightly worse segmentation performance in some cases. Sometimes these performed better in an individual set than the configuration we considered the overall best, and we thus used. Therefore, such modifications might prove helpful when the default settings are not producing good results for datasets with an even more extensive range of object sizes.

Dataset	Model	Performance (Test split)		Performance (Train split)		Configuration
		AP	mAP	AP	mAP	
S1	full_set	0.758	0.580	0.720	0.547	
	apf_s1	0.755	0.575	0.742	0.555	E
	full_set	0.753	0.575	0.764	0.582	G
	apf_s1	0.747	0.560	0.741	0.554	B
	apf_s1	0.745	0.562	0.764	0.582	G
	apf_s1	0.742	0.566	0.745	0.560	B
	full_set	0.741	0.571	0.735	0.552	F
	s1_fs	0.624	0.457	0.685	0.494	A
	s1_fs	0.622	0.454	0.642	0.466	G
	s1_fs	0.570	0.409	0.640	0.45	D
	s1_fs	0.570	0.409	0.654	0.472	A
	s1_fs	0.570	0.382	0.588	0.404	C
FH	fh+	0.755	0.575	0.684	0.504	E
	fh+	0.750	0.580	0.649	0.469	C
	fh+	0.750	0.547	0.684	0.503	D
	fh+	0.745	0.560	0.631	0.469	B
	fh+	0.740	0.575	0.683	0.504	A
	fh+	0.735	0.572	0.647	0.485	D
	fh+	0.721	0.551	0.717	0.554	C
	apf	0.504	0.318	0.491	0.333	G
	apf	0.500	0.314	0.457	0.331	B
	apf	0.480	0.305	0.499	0.334	E
	apf	0.462	0.270	0.480	0.328	B
	apf	0.435	0.299	0.449	0.326	A
APF	apf_s1	0.642	0.488	0.626	0.466	E
	full_set	0.633	0.474	0.639	0.478	G
	apf_s1	0.631	0.477	0.626	0.469	B
	full_set	0.631	0.466	0.604	0.451	F
	apf_fh	0.628	0.470	0.639	0.478	G
	apf_fh	0.626	0.469	0.627	0.472	B
	full_set	0.625	0.470	0.606	0.460	F
	apf_fs	0.484	0.361	0.521	0.381	F
	full_fs	0.481	0.356	0.518	0.380	C
	apf_fs	0.474	0.335	0.503	0.369	E
	apf_fs	0.473	0.351	0.489	0.349	D
	apf_fs	0.460	0.348	0.499	0.363	A

Configurations:
A: Cellpose default.
B: Cellpose default with 1000 epochs (configuration used in this study).
C: Cellpose default with size invariant training (rescale = False).
D: Rescale range for augmentation increased and scaled to median object size during training; size invariant (rescale = False); mean object diameter used for rescaling during inference.
E: Rescale range for augmentation increased and scaled to median object size during training.
F: Rescale range for augmentation increased and scaled to median object size during training; size invariant (rescale = False);1000 epochs; mean object diameter used for rescaling during inference.
G: Size invariant training (rescale = False); 1000 epochs.

Table S5: Overall segmentation performance of our models and benchmarks for our test image tiles for all datasets with different configurations. AP = average precision at the intersection over union threshold of 0.5, averaged over the dataset; mAP = mean average precision over IoU threshold 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.

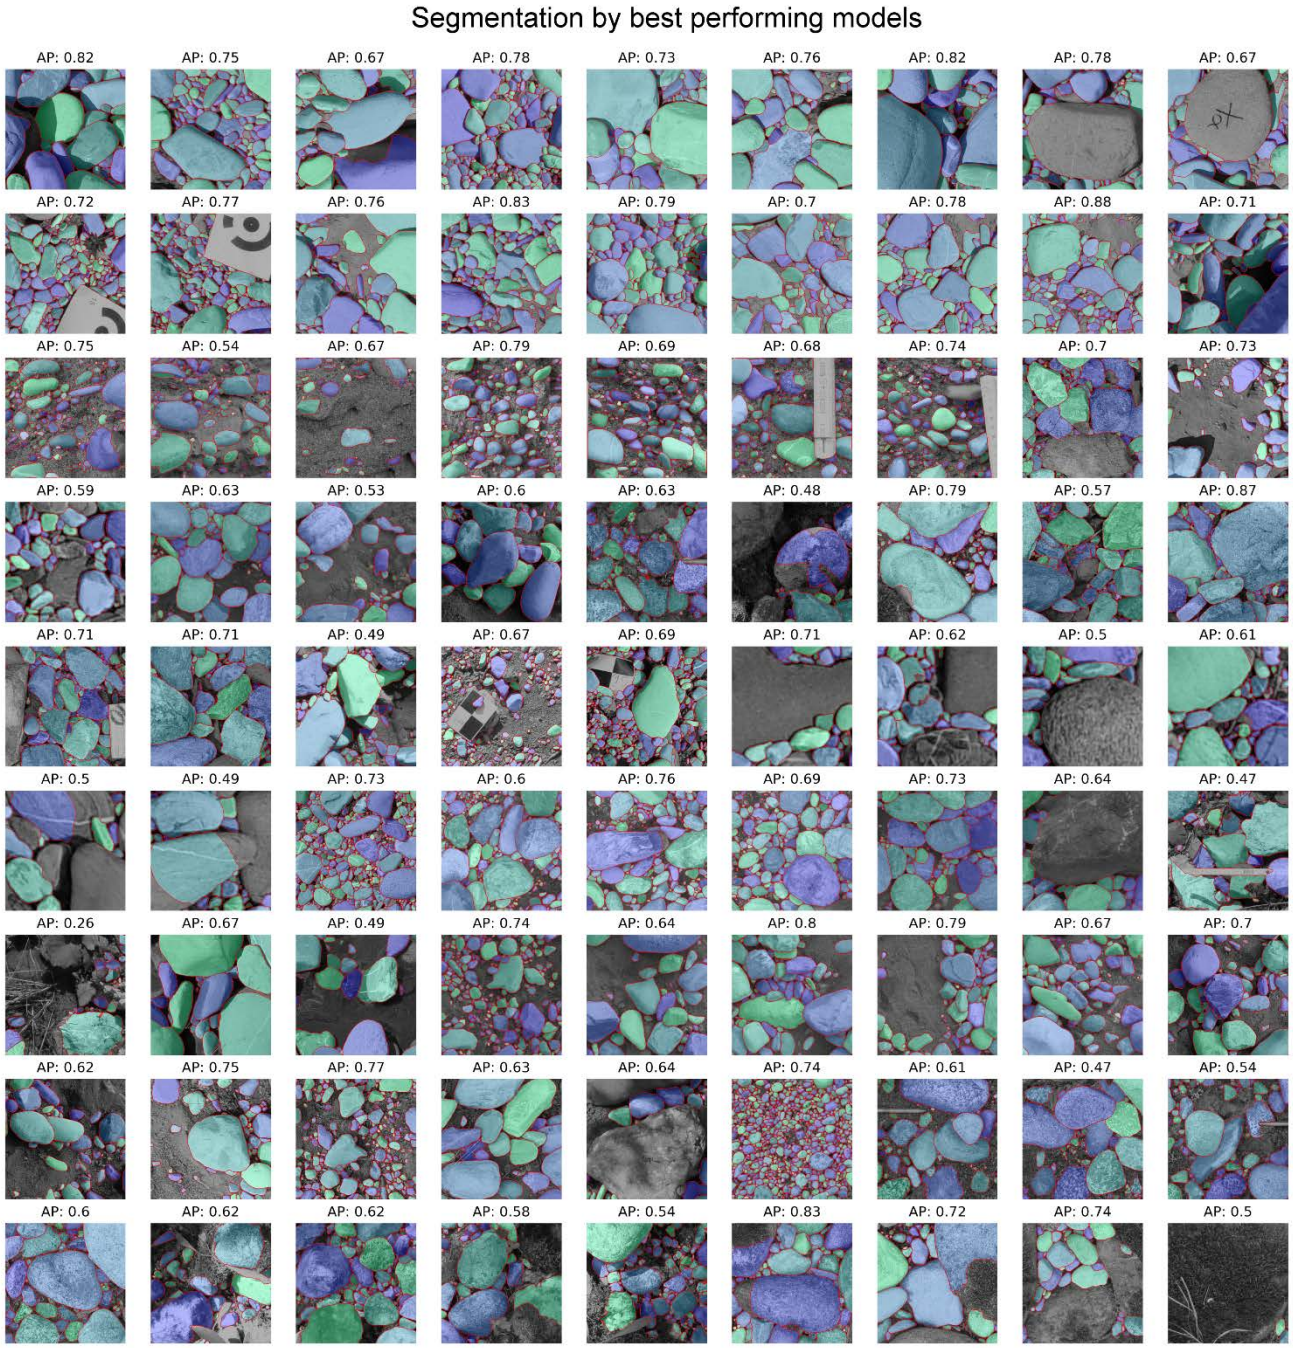

Figure S6: Segmentation results for all image tiles for the best performing model that was applied on the respective data splits, i.e., *full_set* (S1), *fh*+ (FH), and *full_set* (APF) again. AP = average precision at intersection over union threshold of 0.5 for the corresponding tile.

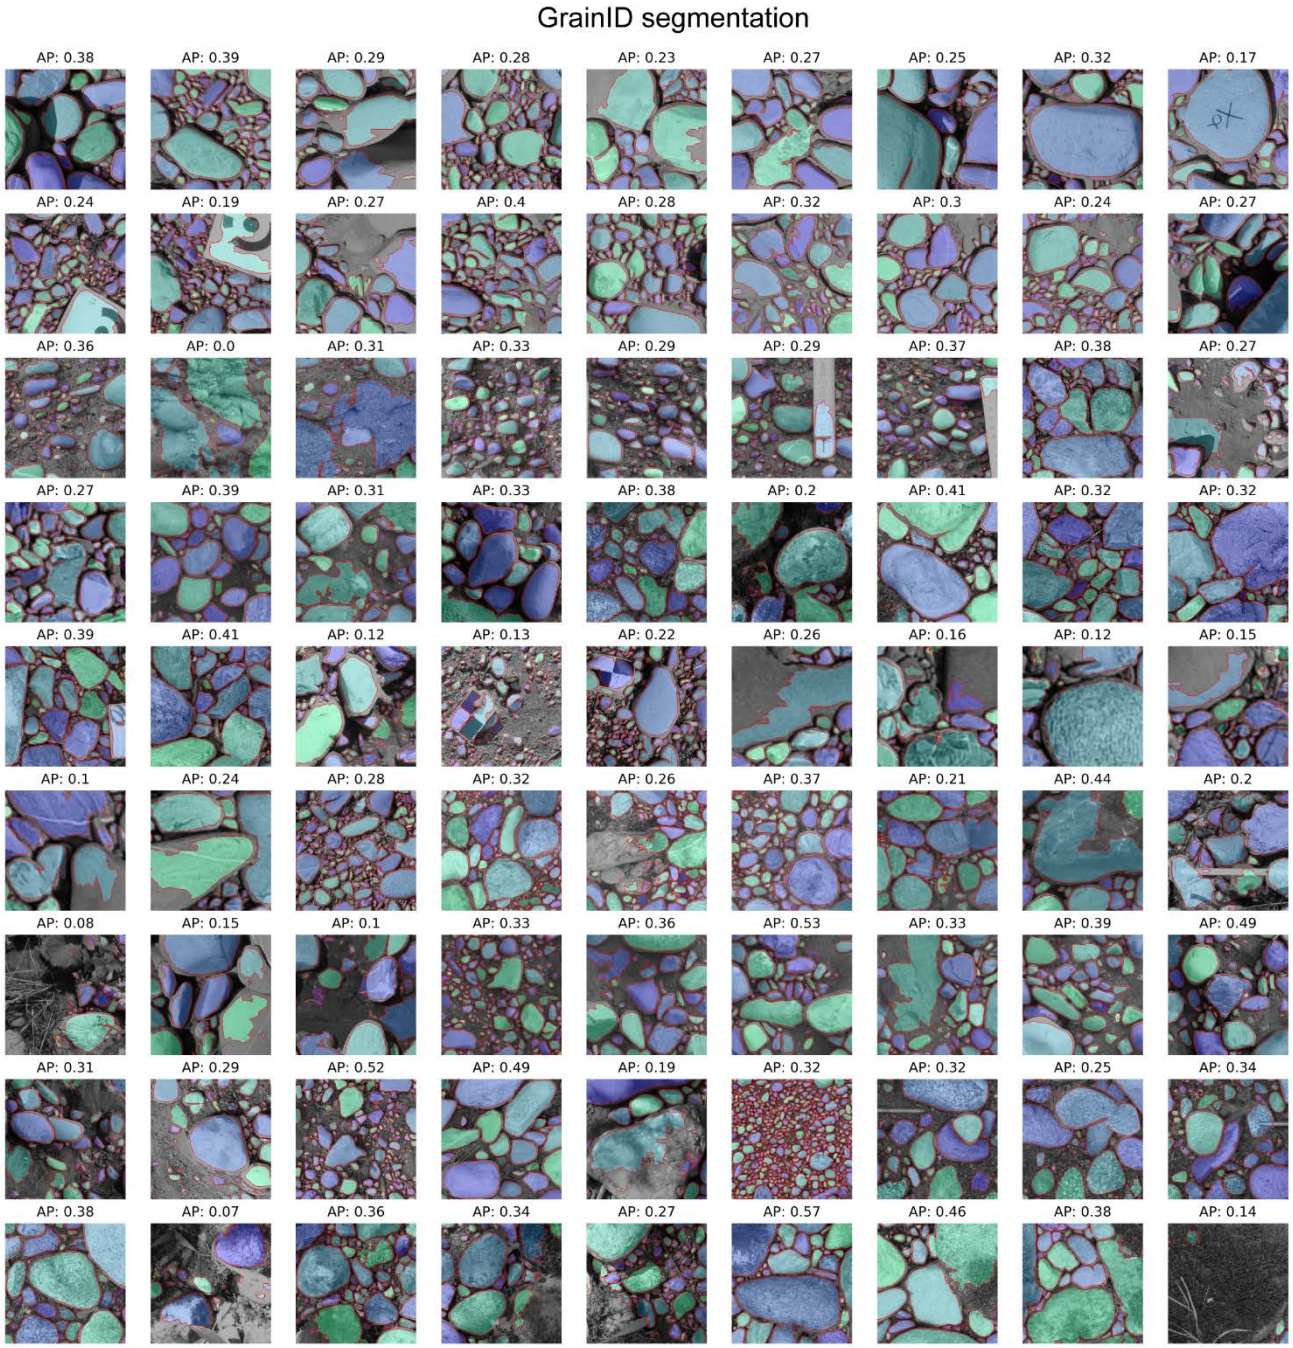

Figure S7: Segmentation results for all image tiles for the GrainID model of Chen et al. (2022a) used as a benchmark. AP = average precision at intersection over union threshold of 0.5 for the corresponding tile.

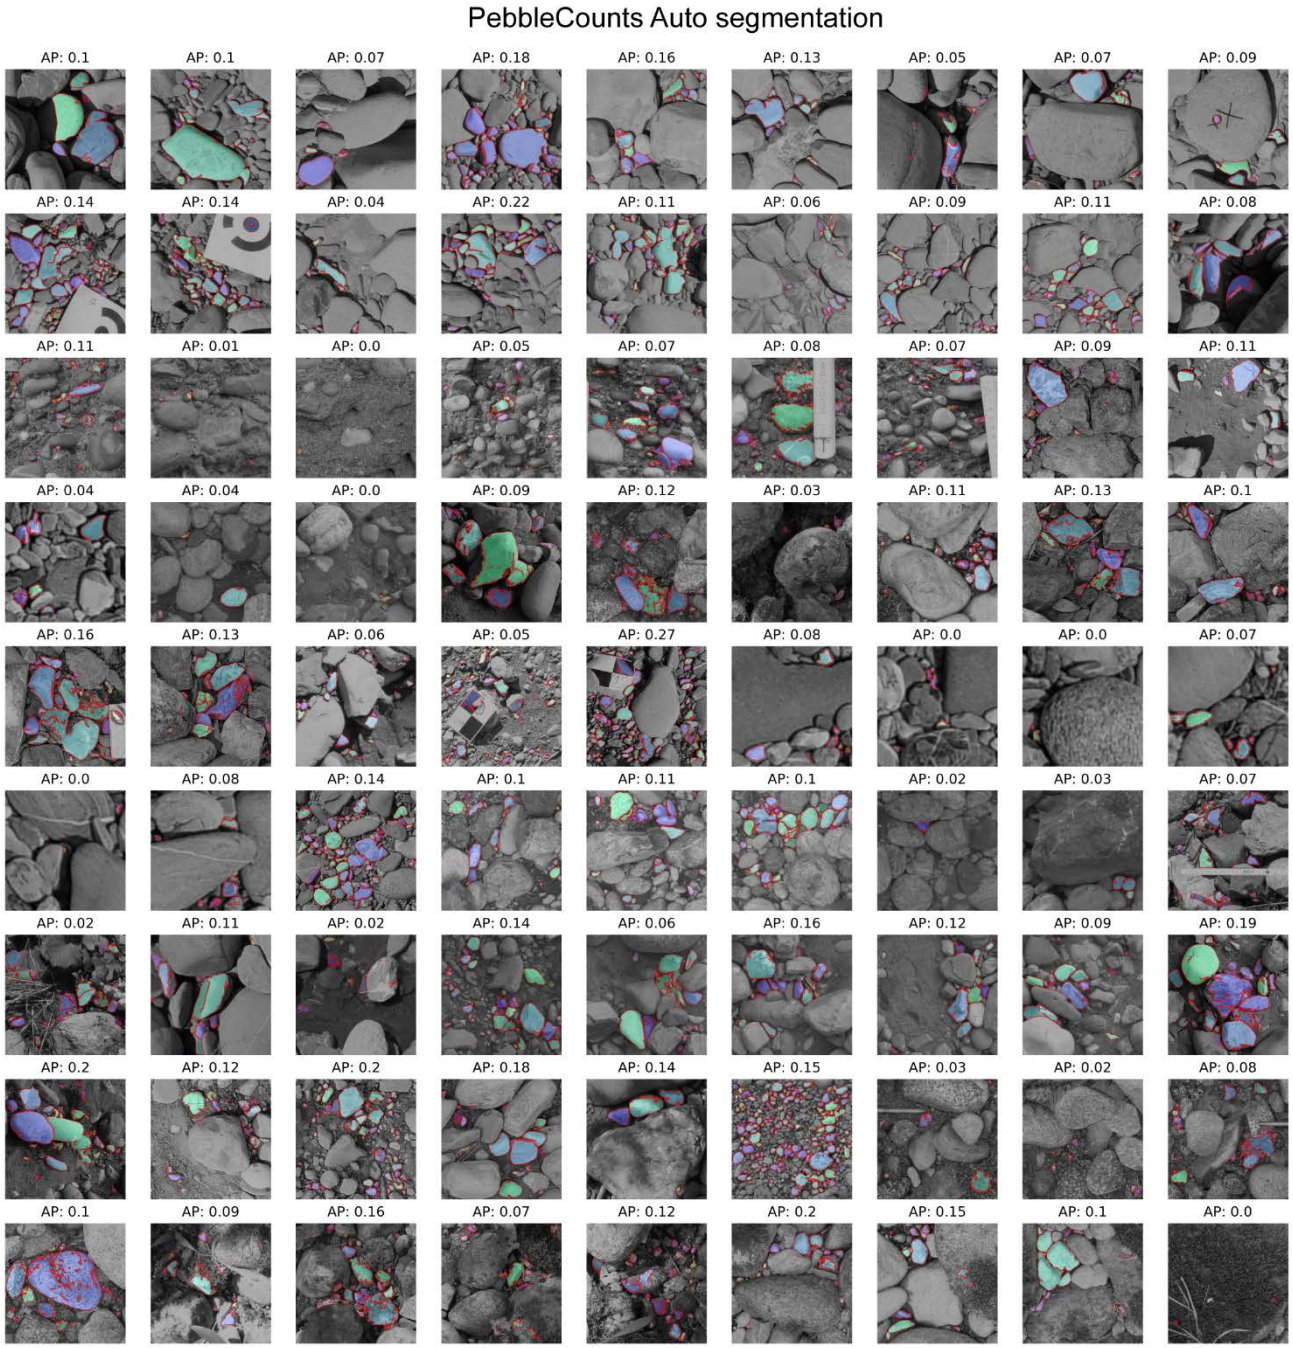

Figure S8: Segmentation results for all image tiles for the PebbleCounts Auto tool of Purinton et al. (2019) used as a benchmark. AP = average precision at intersection over union threshold of 0.5 for the corresponding tile.

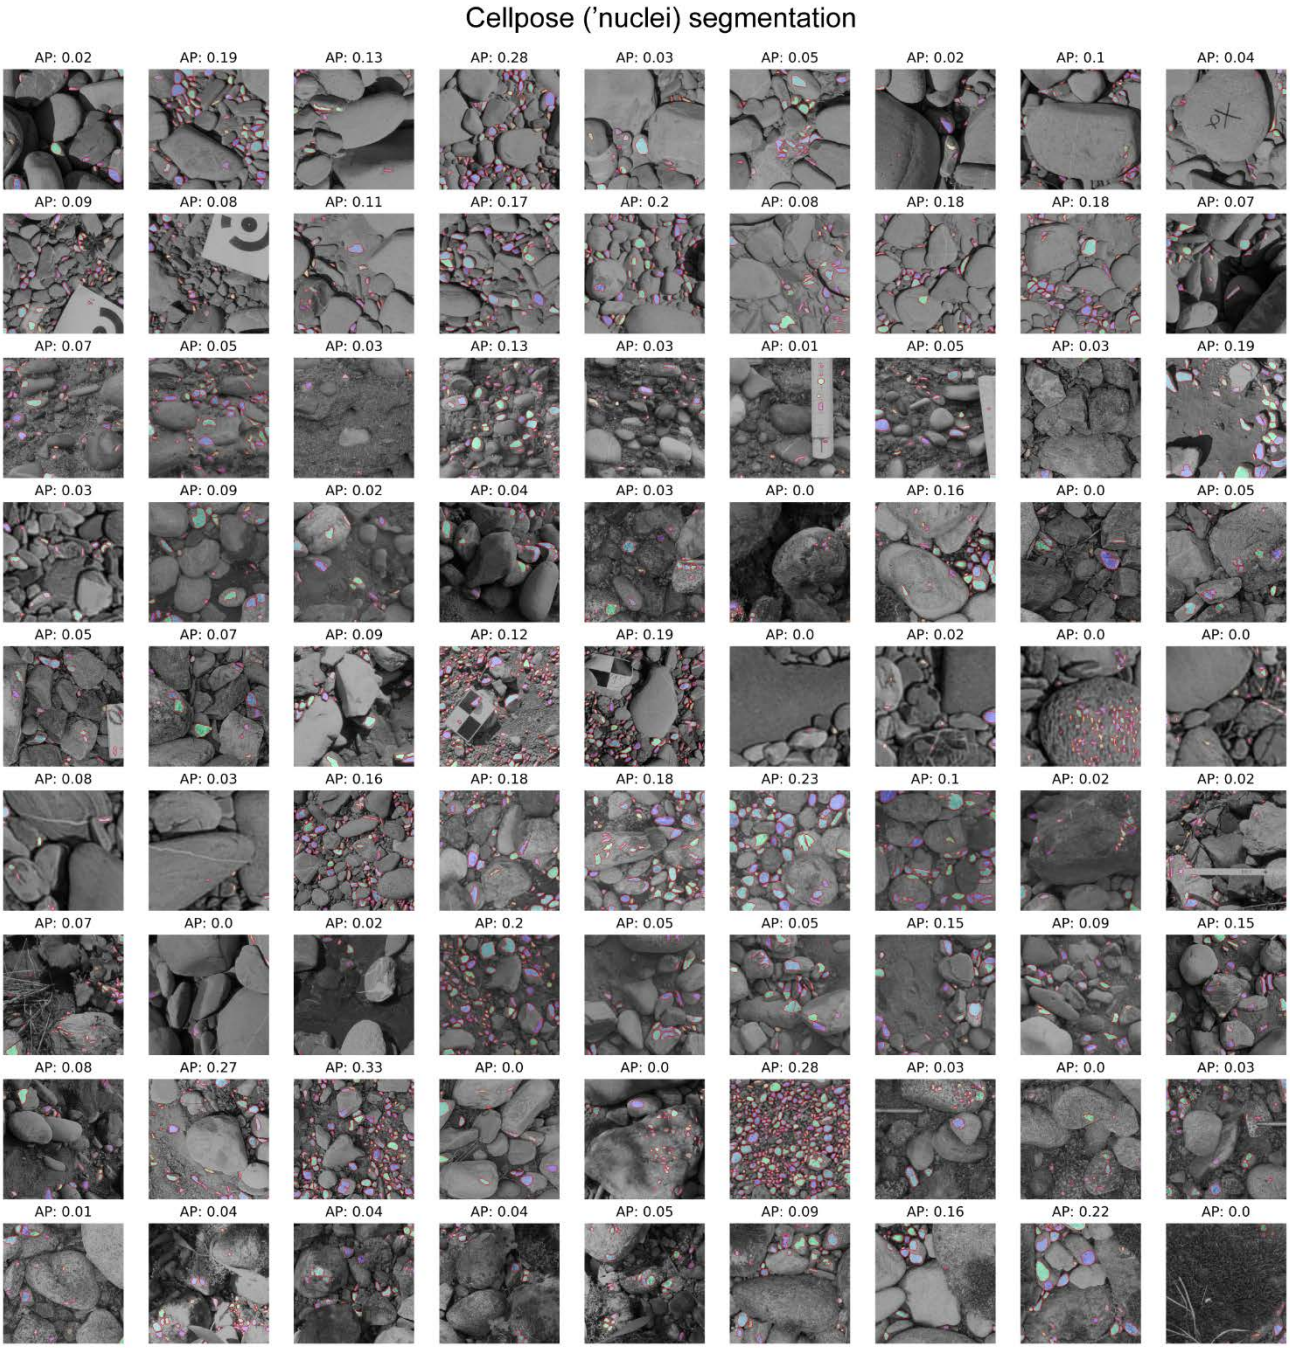

Figure S9: Segmentation results for the Cellpose ‘nuclei’ model of Stringer et al. (2021) without re-training that was used as a benchmark. AP = average precision at intersection over union threshold of 0.5 for the corresponding tile.

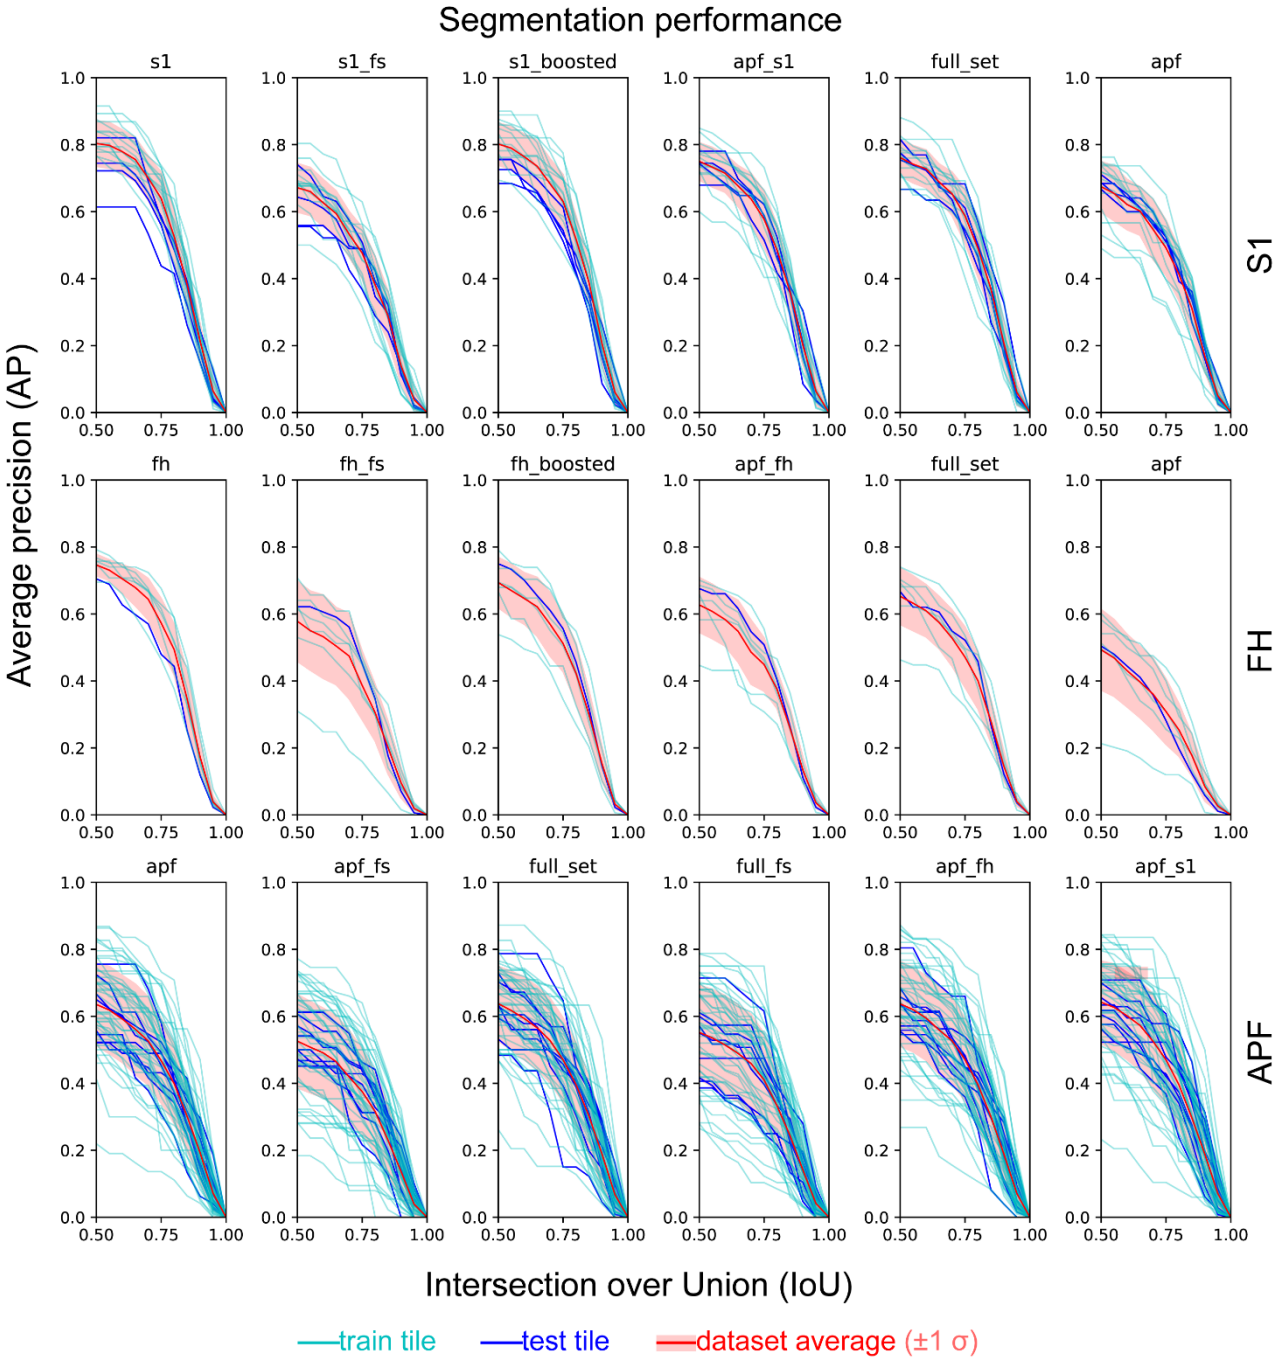

790 **Figure S10:** Detailed segmentation performance for all our models in the respective datasets. The uncertainty range of the mean (1σ = one sigma standard deviation) is calculated for the whole set, i.e., including train and test tiles. We suspect over-fitting to the trained data occurs when the segmentation performance on the test tiles is lower than the average performance on all tiles (e.g., models *fh*, *s1+*, and potentially *s1*).

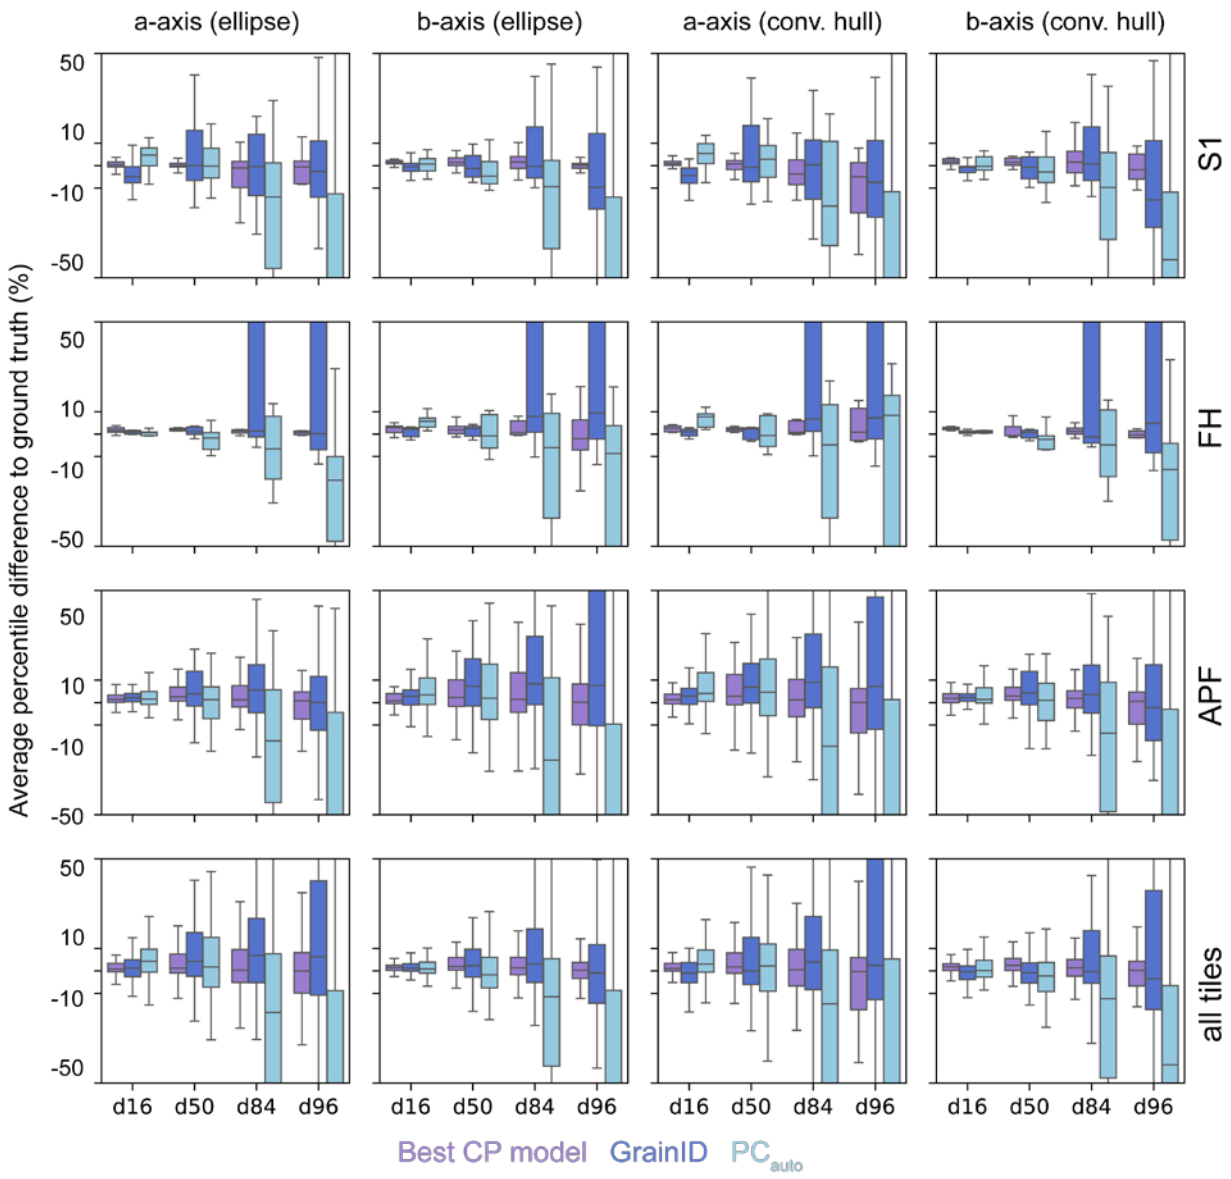

Figure S11: Difference of key grain size percentiles to grain size percentiles in the ground truth data. The average difference and its uncertainty range (1σ = one sigma standard deviation) are calculated for 81 tiles, i.e., including train and test tiles. All grain sizes are measured on filtered masks, i.e., only grains with b-axes ≥ 12px and a center-point within the central 64% of the image tile are considered. 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.

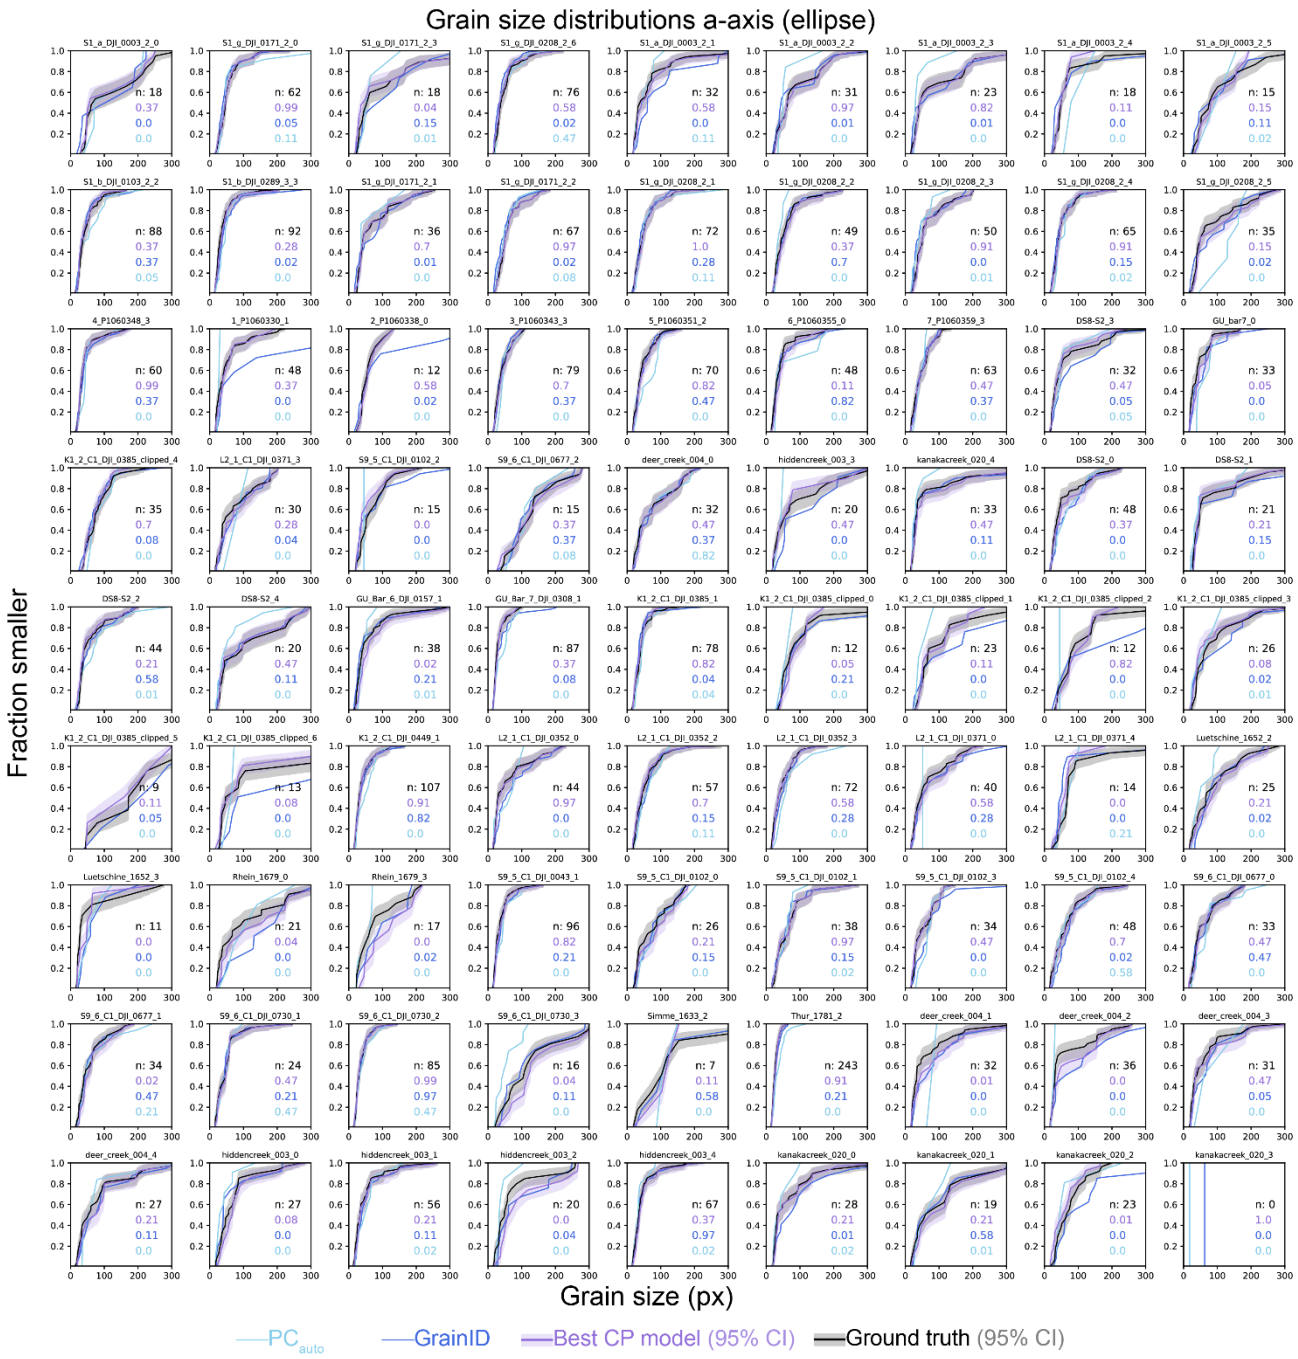

Figure S12: Grain size distributions for all tiles for the a-axis of the ellipse approximation. The statistical uncertainty is reported for 95% CI for all percentile values of our best-performing models and the ground truth masks. All grain sizes are measured on filtered masks, i.e., only grains with b-axes ≥ 12px and a center-point within the central 64% of the image tile are considered. 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. We report the number of grains in the ground truth after filtering (n). The color-coded probability of the respective results having the same distribution as the ground truth (i.e., *p* values of a two-sample Kolmogorov-Smirnov test, where $p > 0.05$ indicate no statistical difference).

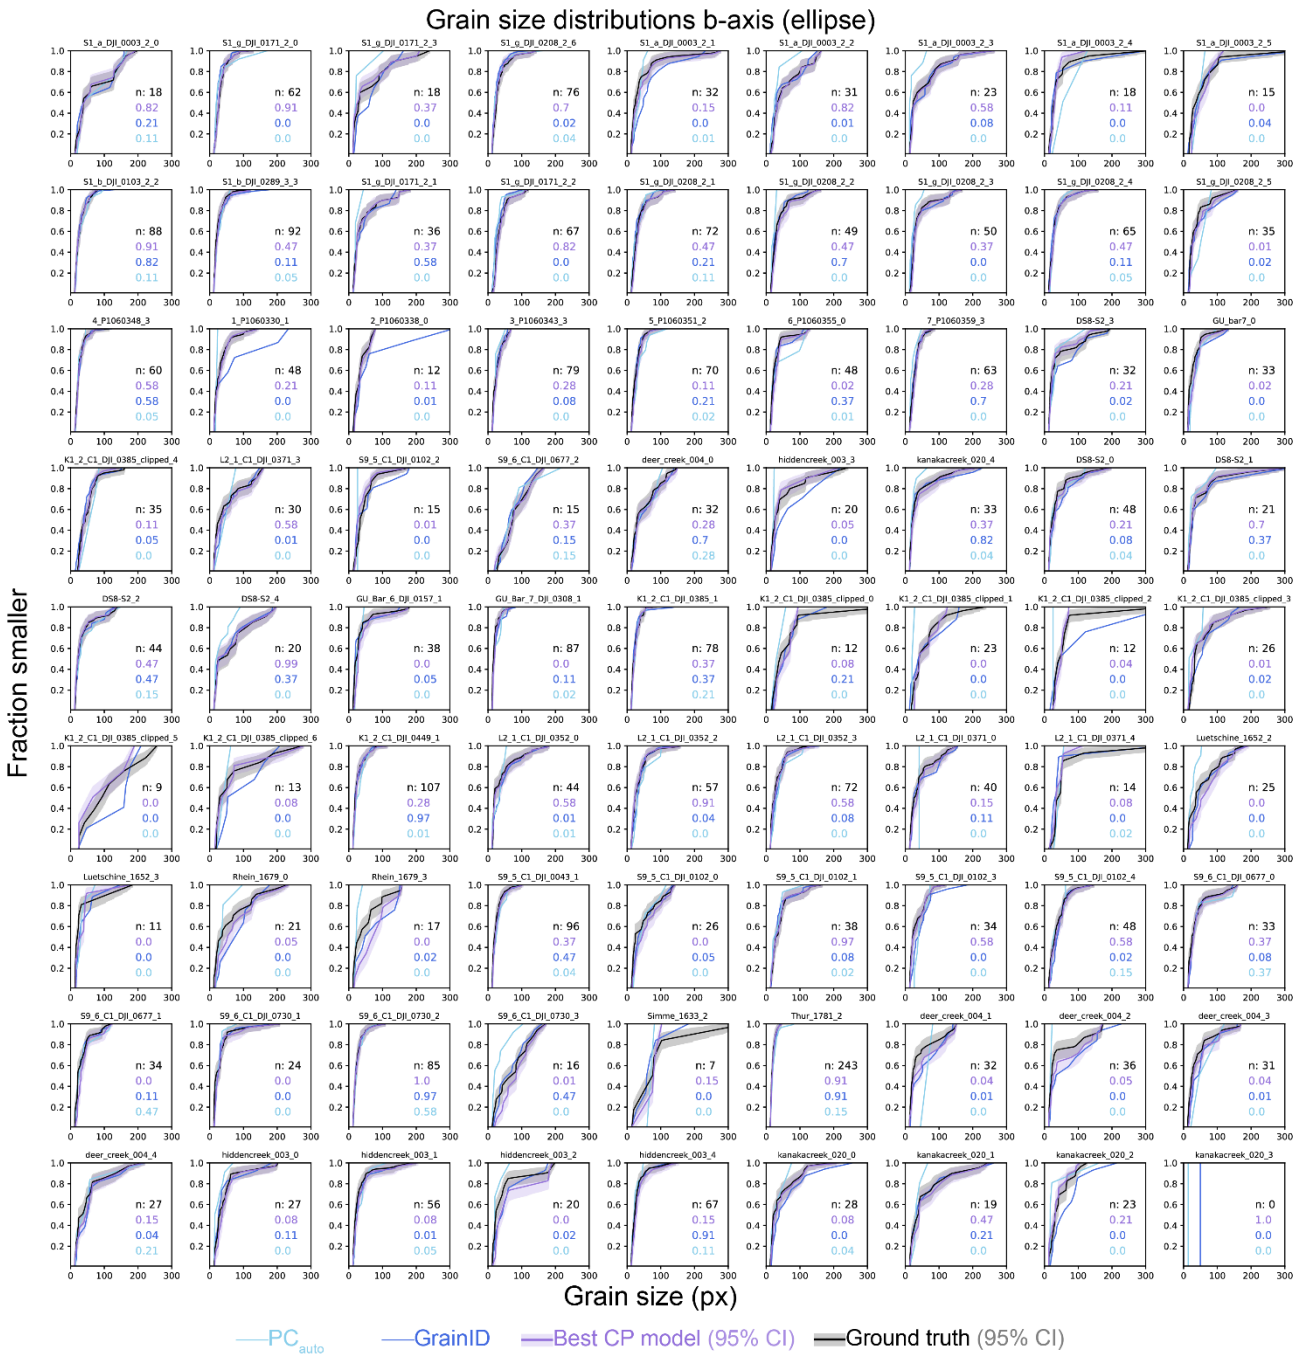

Figure S13: Grain size distributions for all tiles for the b-axis of the ellipse approximation. The statistical uncertainty is reported for 95% CI for all percentile values of our best-performing models and the ground truth masks. All grain sizes are measured on filtered masks, i.e., only grains with b-axes ≥ 12px and a center-point within the central 64% of the image tile are considered. 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. We report the number of grains in the ground truth after filtering (n). The color-coded indicates the probability of the respective results having the same distribution as the ground truth (i.e., *p* values of a two-sample Kolmogorov-Smirnov test, where $p > 0.05$ indicate no statistical difference).

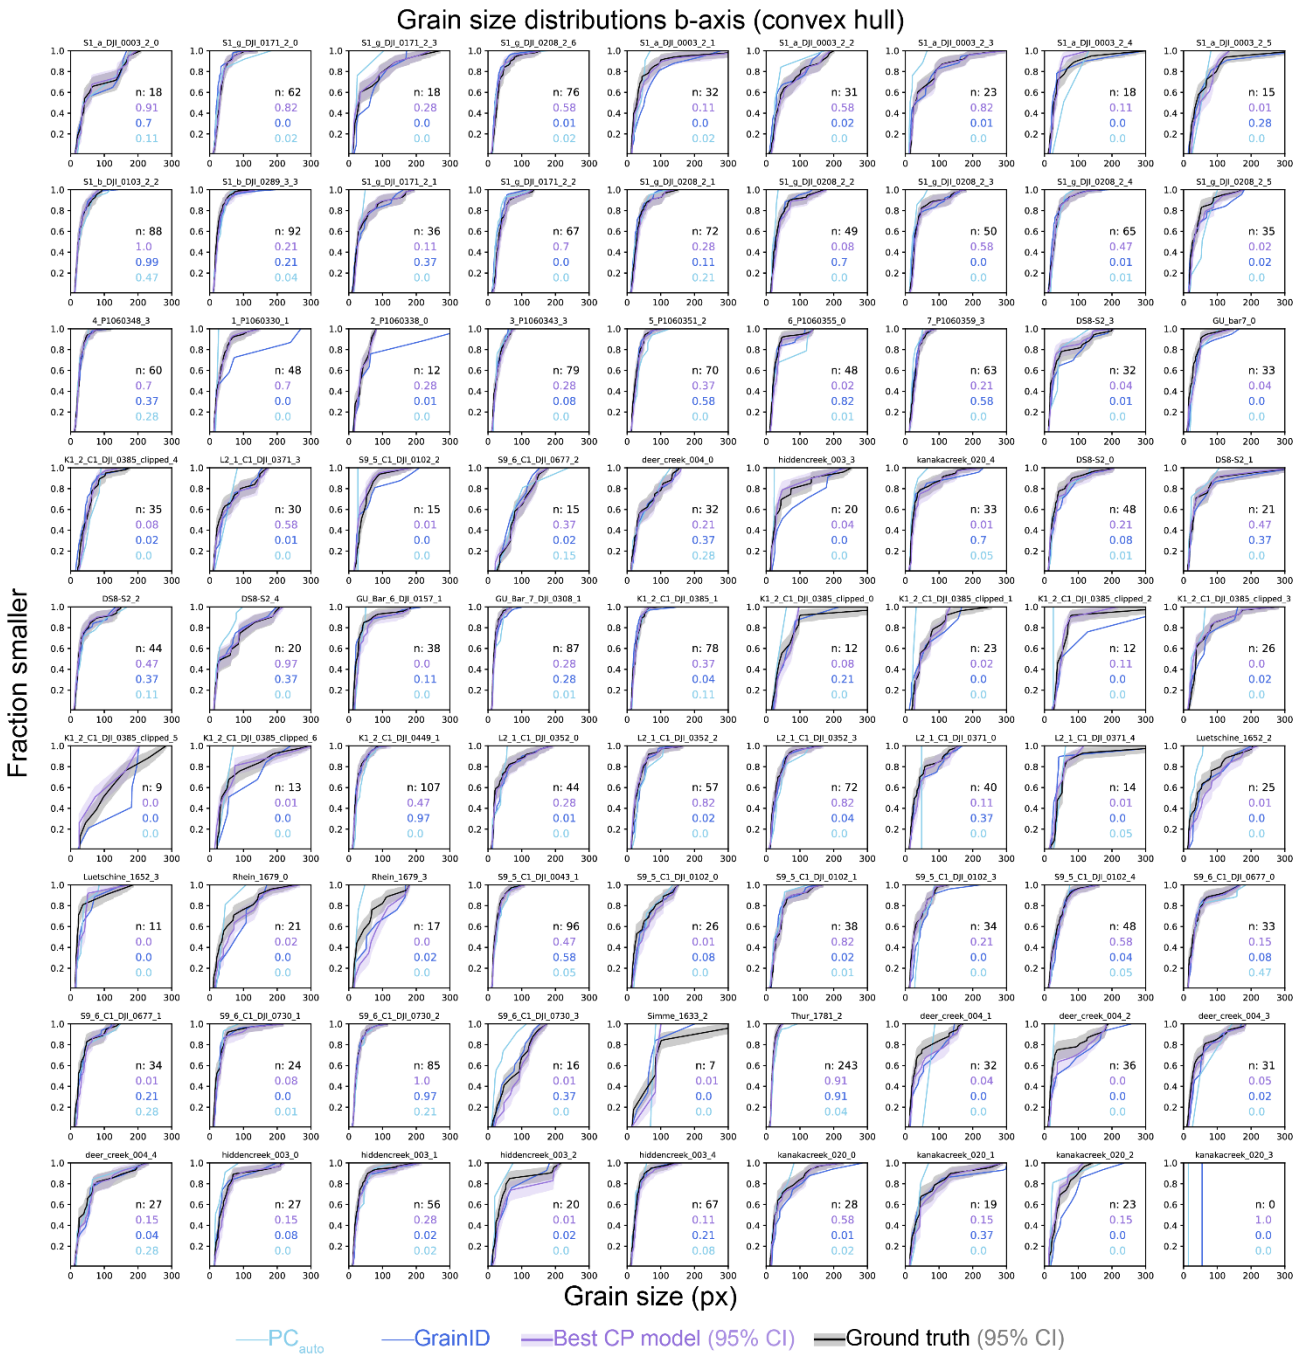

Figure S14: Grain size distributions for all tiles for the a-axis of the convex hull approximation. The statistical uncertainty is reported for 95% CI for all percentile values of our best-performing models and the ground truth masks. All grain sizes are measured on filtered masks, i.e., only grains with b-axes ≥ 12px and a center-point within the central 64% of the image tile are considered. 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. We report the number of grains in the ground truth after filtering (n). The color-coded indicates the probability of the respective results having the same distribution as the ground truth (i.e., *p* values of a two-sample Kolmogorov-Smirnov test, where $p > 0.05$ indicate no statistical difference).

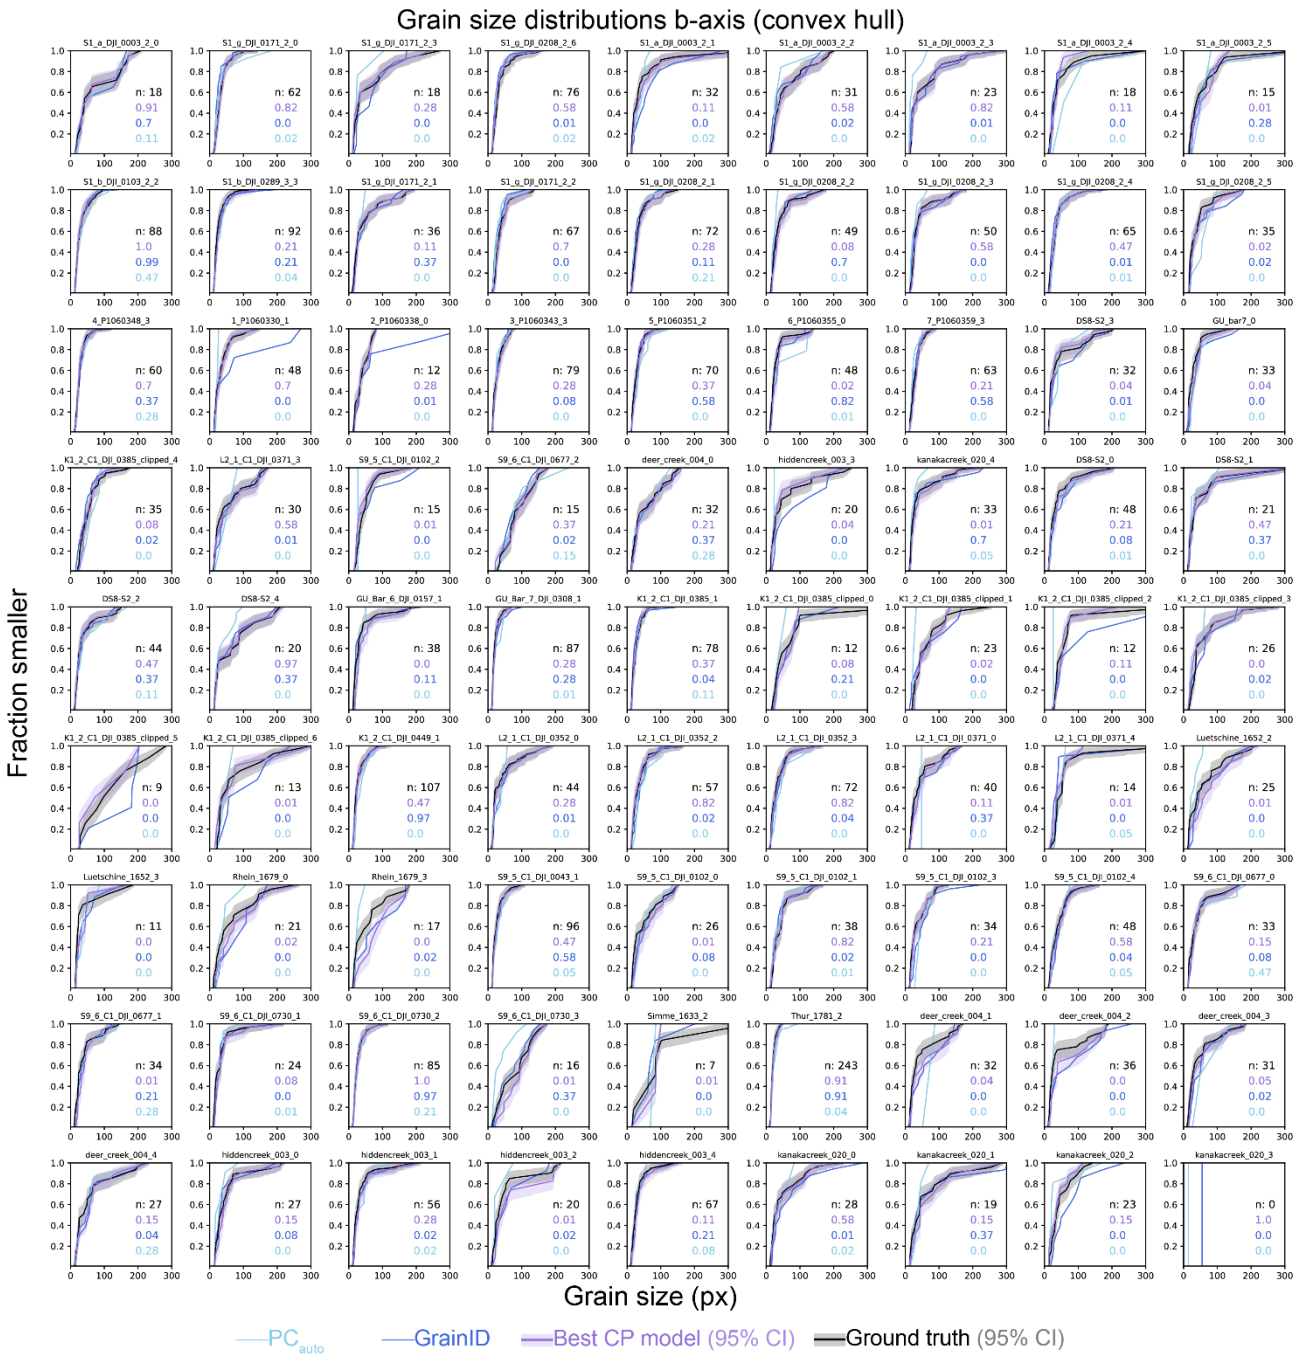

Figure S15: Grain size distributions for all tiles for the b-axis of the convex hull approximation. The statistical uncertainty is reported for 95% CI for all percentile values of our best-performing models and the ground truth masks. All grain sizes are measured on filtered masks, i.e., only grains with b-axes ≥ 12px and a center-point within the central 64% of the image tile are considered. The best Cellpose models (CP) refer to the models with the highest average AP score (0.5 IoU; see son 3.1.1) and are *full_set* for S1 and APF and *fh+* for FH, respectively. We report the number of grains in the ground truth after filtering (n). The color-coded indicates the probability of the respective results having the same distribution as the ground truth (i.e., *p* values of a two-sample Kolmogorov-Smirnov test, where $p > 0.05$ indicate no statistical difference).