

Global mapping of river sediment bars

Patrice E. Carbonneau¹  | Simone Bizzi² 

¹Department of Geography, Durham University, Durham, UK

²Department of Geoscience, University of Padova, Padova, Italy

Correspondence

Patrice E. Carbonneau, Department of Geography, Durham University, Durham DH1 3LE, UK.

Email: patrice.carbonneau@durham.ac.uk

Abstract

Recently, deep learning has been increasingly applied to global mapping of land-use and land-cover classes. However, very few studies have addressed the problem of separating lakes from rivers, and to our knowledge, none have addressed the issue of mapping fluvial sediment bars. We present the first global scale inventory of fluvial gravel bars. Our workflow is based on a state-of-the-art fully convolutional neural network which is applied to Sentinel-2 imagery at a resolution of 10 m. We use Google Earth Engine to access these data for a study site that covers 89% of the Earth's surface. We count 8.9 million gravel bars with an estimated area of 41 000 km². Crucially, the workflow we present can be executed within a month of highly automated processing and thus allows for global scale, monthly, monitoring of gravel bars and associated rivers.

KEYWORDS

deep learning, Google Earth Engine, river mapping, sediment mapping, semantic classification

1 | INTRODUCTION

Our capacity to map river systems globally and our ability to reconstruct their geomorphological trajectories of the recent past (last decades/years) will determine our ability to detect freshwater biodiversity losses, associated ecosystem services availability, such as water resources for energy and food production, as well as flood mitigation strategies (Fryirs, 2017; Tickner et al., 2020). Over the last decade, emerging remote sensing technologies have boosted this capacity to observe river systems from local to global scales (Piégay et al., 2020). Global databases of river features such as river order, slope, discharge and width have been derived indirectly by integrating globally available digital elevation models (Altenau et al., 2021; Grill et al., 2019; Lehner et al., 2008), hydrological modelling, hydraulic geometry relationships and, more rarely, satellite observations (Frasson et al., 2019). These global databases have demonstrated their value allowing us, for instance, to explore the role of major dams globally affecting the natural flow of rivers. Additionally, databases of global estimates of water, sediment and nutrient river fluxes have highlighted the significant roles of human activities of altering such fluxes over the last century (Syvitski et al., 2022). However, these global databases do not yet rely on image-based information or direct observations of landforms but on topographic information used to derive fluxes and river network characteristics. Direct forms of river mapping from space at global scale have been mostly limited to

identifying water surface extent as part of comprehensive land-use classifications (e.g., Karra et al., 2021; Zanaga et al., 2022) or its changes over the last decades (e.g., Pekel et al., 2016). However, these studies are not yet able to distinguish lakes and wetlands from rivers and streams. Even more importantly, we have not yet mapped river sediment bars at global scales. This is probably due to the fact the water is easier to classify and isolate using multispectral information whereas fluvial landforms such as sediment bars have a very similar spectral signature when compared with bare soil and agriculture crops bordering river systems. Moreover, identifying water is achievable by using spectral information (Gao, 1996; Xu, 2006), whereas distinguishing rivers from lakes would require an analysis of the shape of the water surface and its interpretation. These tasks require a step forward from a simple pixel-based classification approach to more sophisticated methods based, for instance, on AI, which are able to include in their predictive power information on object shape and texture. This technology is already advanced in other disciplines of computer vision (Long et al., 2015), but it is still not fully exploited in river system mapping where such advances are needed in order to generate novel river geomorphic information capable to advance our understanding of river systems and to better support management needs. Specifically, the use of AI-based monitoring could track the extent rivers, lakes, wetlands, oxbows and so on and associated sedimentary deposits thus delivering novel information on change at global scales. Low flow channels together with sediment bars composed the active

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

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

channel and together with vegetated islands and the riparian corridor composed the river corridor (Ham & Church, 2002). Being able to map these fluvial landforms and their dynamics over time at the global scale would open to a new era in our ability to explore river forms and processes. Recently Nyberg et al. (2023) proposed a convolutional neural network (CNN) approach to identify the river channel belt, defined as the corridor of river channel migration formed during one river avulsion cycle, followed by the use of spectral indices across multiple years to inventory rivers (water inside channel belts) and lakes (water outside channel belts). They present a global map of the extent of river channel belts and distinguish between single versus multi-threaded river channels. This is indeed a river mapping contribution going beyond water surface extent mapping and which starts to provide more geomorphologically informed data. However, this approach does not explicitly map sediment bars. Failing to observe sediment bar presence and their spatial configuration over time, we neglect the possibility to generate basic information to understand river behaviours (Brierley et al., 2013).

In this letter, we present an operational workflow capable of semantic classification of streams, rivers and exposed fluvial sediment bars. For the purpose of this work, gravel bars are defined as any subaerial and unvegetated sediment deposit adjacent to a river. This workflow uses Sentinel-2 data at 10 m spatial resolution downloaded from Google Earth Engine (Gorelick et al., 2017). The core classification task is carried out with a fully convolutional network (FCN) dubbed Tiramisu (Jégou et al., 2017), a type of deep learning architecture capable of pixel-level semantic class predictions that preserve the resolution of the input image. Notably, the trained model and classification workflow separates rivers and streams from lakes with a commission error of 8%, and it separates fluvial bars from other types of bare soil with a commission error of 16%. We deploy this workflow at global scale on a study site equivalent to 89% of the Earth's land area. This allows us to present maps of global river density and, for the first time, global densities of fluvial sediment bars. The workflow can be executed in less than a month meaning that monthly monitoring of rivers and associated changes becomes possible at global scale.

2 | METHODS

2.1 | Model objectives and semantic classes

Our aim is to produce global scale maps that can identify rivers distinctly from lakes and associated fluvial sediment deposits and thus highlight spatial and temporal patterns. We choose the following semantic classes. Class 0 is the background. The type of deep learning model we intend to use requires a background class which is in essence the non-detections. In our case, this will include vegetation, urban areas and barren areas. Class 1 is rivers. Here, we require the presence of water, and we include artificial canals. Dry ephemeral channels will not be classified as rivers. Class 2 is lakes. Mapping lakes is considered beyond the remit of this work. However, in order to teach the model to distinguish rivers from lakes, the lakes need to be in a separate semantic class. In this class, we include abandoned meanders (Oxbow lakes). One challenging scenario is where rivers emerge from lakes and reservoirs and where a boundary needed to

be set between the classes. In such cases, the user made a somewhat subjective decision based on a qualitative criteria of seeing the beginning of fluvial forms emerging from the lake. This usually corresponds to the development of curved banks and sometimes of gravel bars. Class 3 covers all forms of fluvial deposits, but we limit this to those deposits adjacent to a river. These are the four classes actually predicted by the model. We also add four so-called 'inherited' classes because they will be determined at a later stage using external data. Class 4 is the ocean. Class 5 is glaciated terrain. Class 6 is not used in the paper but is set aside for snow. Class 7 is clouds, and Class 8 is a no-data class for cases where no satellite imagery is available.

2.2 | Google Earth Engine

Our workflow is based on Sentinel-2 data, and we use Google Earth Engine as a download tool (Gorelick et al., 2017). We download Sentinel-2 Bands 8 (Near-Infrared), 4 (Red) and 3 (Green) to an expanded Google Drive cloud storage space. We keep the native spatial resolution of 10 m. In addition to the image layers, we add a cloud mask derived from the s2cloudless product (Skakun et al., 2022) as an extra band. Spatially, the download process is organised as per the Military Grid Reference System (MGRS) that further divides each UTM zone of 6° in longitude into zones with 8° of latitude. The download area is composed of the 468 MGRS zones that constitute the non-polar world (89% of the Earth's landmass, see the outline of Figure 1). Given that 1° is equivalent to 110 km at the equator but only 55 km at latitudes of ±60°, imagery and all derived products are projected to the appropriate UTM coordinate system in order to preserve the constant, latitude-invariant, 10 m spatial resolution of the original Sentinel-2 imagery. Temporally, we choose a sample duration of 1 month in order to capture changes at smaller time scale recognising that such a short temporal window is not capable of delivering global imagery which is 100% cloud free. The choice of the exact monthly period is the result of a compromise. We discounted boreal winter given that snow can cover as much as 45 000 000 km² (Déry & Brown, 2007). We also found that tropical areas of South America and Africa have persistent high levels of cloud at any time of the year. We also noted that boreal summer months with their longer days gave a better overall global coverage of low-cloud data given that most of the Earth's landmass is in the Northern Hemisphere. Consequently, we chose the month of July for the year 2021. We acknowledge that this will leave data gaps in the tropics and in the area of the Indian Monsoon. Over this period, Google Earth Engine will composite all available images with less than 20% cloud cover by using the median value of available pixels.

2.3 | External data

We also use several global databases at various stages of this work. We use HydroSHEDS vector products (Lehner et al., 2008) to define the ocean area and inform post-processing algorithms that correct errors in large lakes and large rivers. We also use the Randolph Glacier Inventory (RGI Consortium, 2017) and define glaciated terrain as a 5 km buffer extending around glaciers.




FIGURE 1 Study area. Green points show the location of the 343 labelling sites, distributed in catchments that exceed 500 000 km² (blue). The full study area is the combined blue and black outlines and represents 89% of the earth's landmass. [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

2.4 | Deep learning and image processing pipeline

Deep learning is rapidly establishing itself as the de facto method for water mapping at large scales (e.g., Isikdogan et al., 2017, 2020; Moortgat et al., 2022). We use a FCN that can deliver pixel-level semantic classifications (Long et al., 2015). Within this family of models, several architectures exist. We have chosen to use a fully convolutional DenseNet architecture proposed by Jégou et al. (2017) and dubbed 'Tiramisu'. This architecture delivers state-of-the-art results that outperform other models (e.g., SegNet, Badrinarayanan et al., 2017; DeepLab, Chen et al., 2018) against established benchmarks (Brostow et al., 2009), while having a low number of trainable parameters and corresponding training times. Crucially, fully open-source implementations of the model are readily available. We use this architecture with the common image chip size of 224 × 224 pixels because (1) this results in models that are cross-comparable with a wider range of other models (e.g., VGG16, Simonyan & Zisserman, 2015 and ResNET, He et al., 2016) and (2) the final models are more transferable and can be run on more modest GPUs with 8 GB of memory. Our final model is the result of two training phases. First, we use a large, but more error-prone, training dataset of 740 000 samples derived from an automated labelling of ESRI 2020 land-cover data (Karra et al., 2021). Second, we fine-tune the model with another training phase that uses the manually labelled data from 293 randomly selected locations of Figure 1, the remaining 50 locations are set aside for validation. For each of these sites, we extract an image of 15 × 15 km and manually label these images. We then use data augmentation in order to get 59 000 very high quality samples of 2.24 by 2.24 km. These data are used to train and validate the model (see the Supporting Information for more detail). A key step of our workflow is a set of novel image post-processing filters. Deep learning semantic classification models produce prediction probabilities for each class with each pixel being assigned to the class of highest probability. However, we find the value of the second highest probability can be combined with contextual and external information to improve our predictions. We postulate that if two adjacent pixels have a different class, there could be a classification error. An examination of the probabilities for the most-likely class (the attributed class) and the second most-likely class is then used to determine if an

error has occurred. Additionally, we use vector datasets from the HydroSHEDS database to provide further contextual corrections for large lakes in excess of 5 km² and large rivers with widths potentially in excess of 2.24 km. Full code and details on the post-processing filters can be found in Carbonneau (2023a). Figure 2 summarises our deep learning workflow.

2.5 | Validation

First, we use 50 manual validation sites set aside before the training process. The data from these tiles have 112 million background pixels, six million river pixels, nine million lake pixels and 770 000 bar pixels. These data are used to produce confusion matrices and estimate the user's and producer's accuracies. They are also used to produce validation results for some intermediate processing steps (see the Supporting Information). We also merge the river and lakes classes for both truth and model data in order to produce error estimates for a 'water body' class that are more comparable to existing work. Second, we build a validation dataset with 10 000 randomly selected river locations across the full study site. We repeat this random selection for the lakes class and bar class (30 000 locations in total). We then use Google Earth Engine to extract an image for each of these samples. We edit the pixel values at the sample location to create a blue cross composed of five 4-connected pixels, and we save the image in a folder corresponding to the class. We then examine each image in each class folder and assess if the predicted classification of the point marked by the blue cross was correct or incorrect. This procedure will estimate the user's accuracy for the river, lake and bar classes at global scales.

3 | RESULTS

3.1 | Validation

Table 1 gives the validation outcomes for the manually labelled data. Table S1 gives validation results for the intermediate steps in our processing pipeline. Accuracy metrics for the background class are in

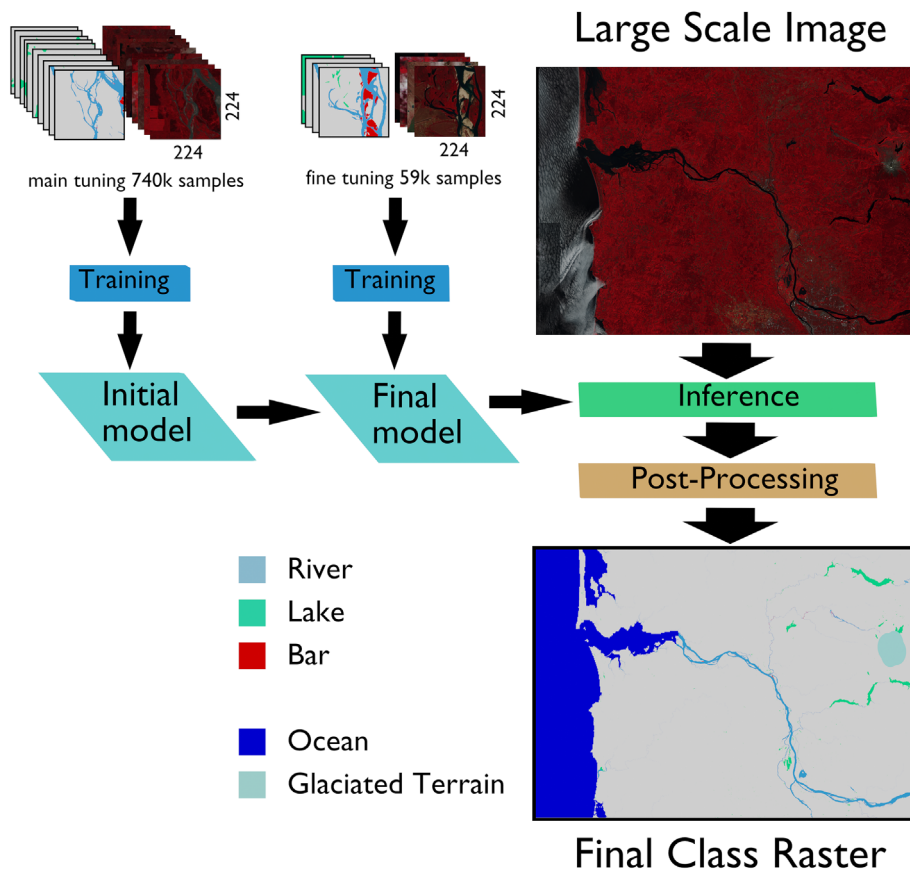


FIGURE 2 Deep learning workflow chart. We use a main training phase with 740K medium-quality, automatically generated samples followed by a fine tuning training phase with 59K high-quality samples. The final model is then used for inference (i.e., prediction) followed by our post-processing filters. The model predicts three classes: rivers, lakes and bars. External is then used to define the ocean and glaciated terrain. [Color figure can be viewed at wileyonlinelibrary.com]

TABLE 1 Confusion matrices, producer's accuracy and user's accuracy for the manually labelled validation data.

	Background	River	Lake	Bar	Producer's accuracy (%)
Background	110 828 808	309 448	1 563 293	53 938	98
River	206 589	5 543 802	63 103	33 273	95
Lake	403 899	114 083	8 507 836	140	94
Bar	233 256	60 752	7811	465 477	61
User's accuracy (%)	99	92	84	84	

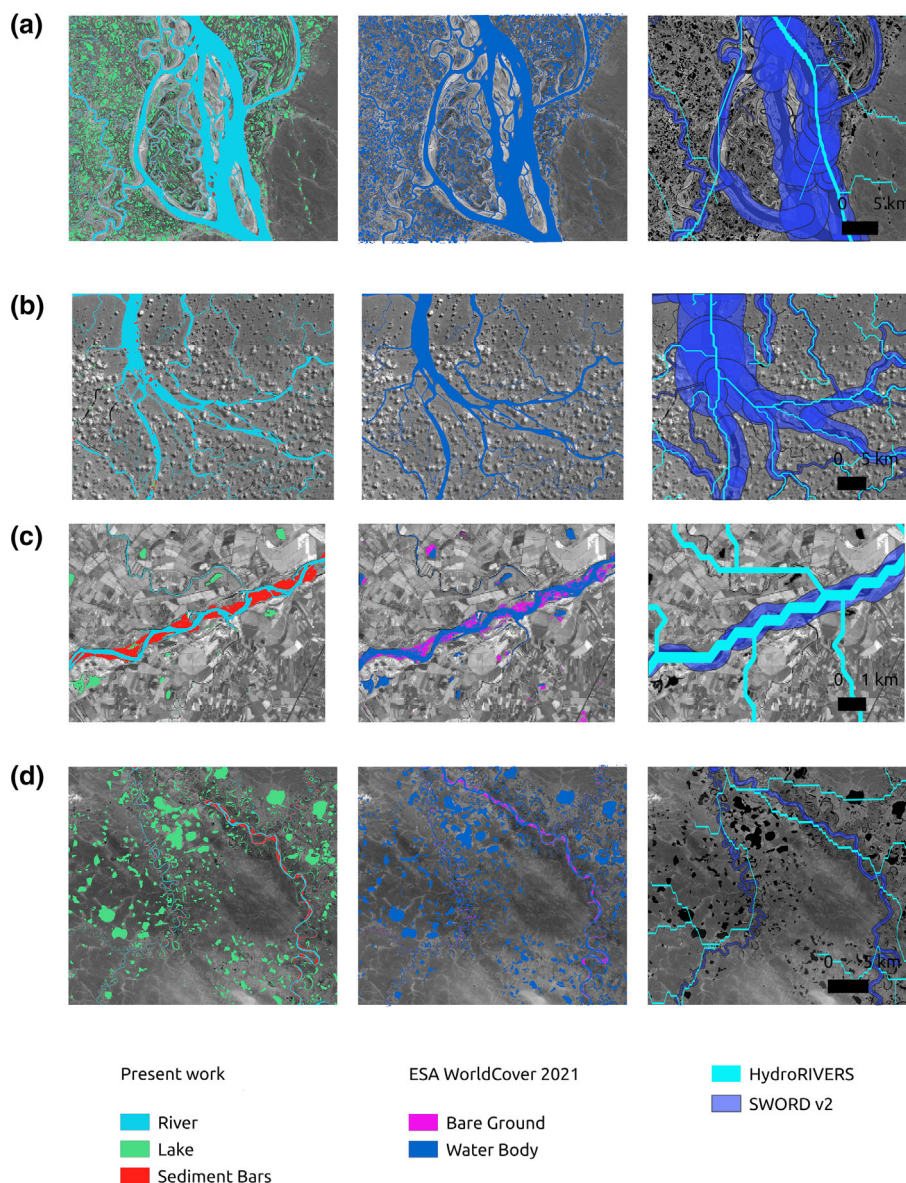
excess of 98%. Any averaging of the accuracy statistics across classes would be artificially high, especially if weighted by pixel counts, and it is therefore important to consider class accuracies individually. The confusion matrices show that the main source of error for the river class is confusion with lake pixels with producer's and user's accuracies of 95% and 92%, respectively. In the case of bars, user's accuracy is 84% with a lower producer's accuracy of 61%. Here, the main source of confusion is the background class. This means that pixels predicted as bars are reliable, but overall, we can expect a significant portion of global bar area to be absent from the survey. Our bar results should therefore be seen as a lower bound accounting for a little over half the global population of bars. Validation results for intermediate stages of our workflow are available in the [Supporting Information](#).

In the case of a merged water body class, we find user's accuracy of 87% and producer's accuracy reaching 95%. Specifically for the water body class, Zanaga et al. (2022) report user's accuracy of 89% and producer's accuracy of 86% for the latest version of the WorldCover product. Karra et al. (2021) report user's accuracies of 83% and producer's accuracies ranging from 86% to 94%. Isikdogan

et al. (2017) report commission errors of 8% (92% user's accuracy) and omission errors of 13% (producer's accuracy of 87%). However, we stress that none of these comparators studies present data where lakes and rivers constitute distinct semantic classes. Nyberg et al. (2023) do provide a semantic class product with distinct rivers and lakes, but they only validate the channel belt and active channel width predictions resulting from their CNN model. User's and producer's accuracies, or any other form of accuracy assessment (e.g., F1 scores and kappa scores), are not available for their global semantic class product.

At larger scales, we found that a global random check of 10 000 samples from the fine-tuned filtered predictions resulted in slightly lower user's accuracies of 85%, 78% and 77% for rivers, lakes and bars, respectively. We followed up this validation by a full visual inspection of all the data. For rivers, the main source of error remains confusion with lake pixels, most often occurring with elongated lakes or reservoirs. For lakes, the main source of error is shadows. In the case of bars, we observe two sources of error worth noting: (1) faint clouds that have not been detected by s2cloudless and occurring over rivers and (2) arid areas adjacent to

FIGURE 3 Result examples of our data from July 2021 with comparison to ESA WorldCover 2021, HydroRIVERS and SWORD v2 data. (a) River McKenzie, Canada; (b) Jacaré River (Amazon Basin), Brazil; (c) Po River, Italy; and (d) Popigay River, Russia. Greyscale background is Band 8 (NIR) of Sentinel 2. [Color figure can be viewed at wileyonlinelibrary.com]



rivers and canals with the same spectral characteristics as bars but not of fluvial origin.

Figure 3 shows four examples of local results. In these four locations, we also show the equivalent WorldCover 2021 data (Zanaga et al., 2022), the HydroRIVERS vectors (Lehner et al., 2008) and the SWORD v2 vectors (Altenau et al., 2021). In the case of HydroRIVERS, we used the Strahler order in the attributes data and the estimated widths from Downing et al. (2012) to buffer the vector lines. In the case of the SWORD v2 data, we used the channel width information present in the attributes. Figure 3 illustrates our final model's ability to effectively detect rivers and bars. However, we also note some issues. In row B, we can see that our model has falsely classified small clouds as bars. Also, some very small channels are detected as water bodies in WorldCover data but as background (undetected) in our data. In the case of the vector data, the rightmost column of Figure 3 shows reasonable correspondence and highlights the challenges of using these vector databases in combination with raster-based semantic classifications. Even if these vectors have information in their attributes that relates to local channel widths, this information was never intended to accurately map channel boundaries. It is therefore not

possible to perform a direct quality estimation of our results based on these data.

The compressed global set of classification rasters at full 10 m resolution requires 10 GB of storage (Carbonneau, 2023b). Figure 4 shows resulting global maps for the surface densities for rivers and bars. On the top, we show River and Stream Surface Area (RSSA) expressed in ppm of the surface area (10 000 ppm = 1%). The map further divides the study area in grid squares of $2^{\circ} \times 2^{\circ}$. Within each square, we used Python to extract and inventory the semantic classes. If we consider the total inventory of river area, we observe a total river area of 433 544 km². If we assume that cloudy areas in a given $2^{\circ} \times 2^{\circ}$ square have the same river density as the non-cloudy area, and if we include cloud-driven no-data values, then we can correct for clouds and we find an estimated total river area of 562 925 km². On the bottom of Figure 4, we show the global densities of bars. This reveals a high concentration of bars at high northern latitudes with additional hotspots South-West of the Amazon, in South East Africa and on either side of the Himalayas. Visual examination revealed that these observations are valid. However, it also confirmed that apparent bar density hotspots around the Nile valley and North of the Arabian




FIGURE 4 Global maps of River and Stream Surface Area (RSSA, top) and Exposed Sediment Bar Surface Area (ESBSA, bottom) for July 2021 and expressed as parts per million of surface area (104 ppm = 1%) for each $2^\circ \times 2^\circ$ area of the study area. Areas left blank have more than 20% cloud coverage. [Color figure can be viewed at wileyonlinelibrary.com]




FIGURE 5 Global scaling relationships for bar sizes. (a) Magnitude frequency plot of bar objects. (b) Scatter plot in log scale of River and Stream Surface Area versus Exposed Sediment Bar Area for all $2^\circ \times 2^\circ$ squares with non-zero bar densities.

Peninsula are in fact arid lands adjacent to canals falsely classified as bars. In total we find 8.9 million gravel bars. Total bar area is estimated as 41 165 and 33 665 km² with and without cloud correction, respectively.

Figure 5 shows global scaling relationships for the bars mapped in Figure 4 (bottom). In Figure 5a, we show the magnitude frequency distribution of global bar sizes in pixels. The relationship is close to a log-linear power law decay with a slope of -1.59 with a slight

upwards curvature, likely caused by errors in the range of 10 to 1000 pixels. Figure 5b shows a direct plot of RSSA and ESBSA on logarithmic axes. This shows a scattered relationship. However, we can see a central cluster of points showing rapid increase of bar densities upwards of river densities of 1000. However, we note that this graph does not include $2^\circ \times 2^\circ$ areas with no bars or no rivers due to the logarithmic axes.

4 | DISCUSSION

Using a deep learning and remote sensing workflow, we have produced the first global survey of fluvial bars. We find patterns that match large-scale areas of erosion and subsequent transport. The largest high density zone of gravel bars is found to be in Siberia and Northern Russia with sediment produced in the Siberian uplands transported north to the Laptev, East Siberian and Chukchi seas. Next, we see a high density area in Alaska where sediment is sourced in the Brooks and Alaska mountain ranges and transported to the Beaufort Sea. We also see high concentrations of sediment in Nunavut likely sourced in the Arctic cordillera mountains. The area south-west of the Andes known for high precipitation from the highest gravel densities in South America transported to the Pacific Ocean. In Africa, we note concentrations of gravels around the Licungo and the Zambezi rivers debouching into the Mozambique channel. Finally, we find high concentrations of gravels in the rivers draining the Himalayas. We estimate the total number of bars as 8.9 million, covering a total area of $\sim 41\,000\text{ km}^2$. These are the first such estimates. Considered in relation to our errors, these are likely an under-estimate. However, the slight data departure from linearity in Figure 5a suggests a large number of false positives at medium scales. At higher scales, the presence of very large bars, in excess of 10 000 pixels with a single count (10^0) suggest large patches of false positives. From the point of view of simple inventory, a false positive can compensate a false negative. Large-scale inventories can be accurate if false positives are as frequent as false negatives. Therefore, the false positives seen in our data (excess bar pixels) may compensate for the low producer's accuracy (missing bar pixels) of the bar class in Table 1. We argue that despite the relatively low producer's accuracy, our results provide an order-of-magnitude starting point for global scale gravel bar inventories. In terms of rivers, our analysis of Sentinel-2 imagery with a spatial resolution of 10 m finds, after correcting for clouds, a total river area of $562\,925\text{ km}^2$ which is broadly in agreement with other published works. Based on Landsat data with a resolution of 30 m, Allen and Pavelsky (2018) report a total RSSA of $468\,000\text{ km}^2$ through direct observation and $773\,000\text{ km}^2$ from extrapolations down to sub-pixel sizes. Downing et al. (2012) provide estimates of RSSA per river order along with an estimated width. For rivers with an averaged width above 8.3 m (fourth Strahler order), they estimate a coverage of $546\,000\text{ km}^2$. Raymond et al. (2013) propose a figure of $536\,000\text{ km}^2$ for all sizes of channels. Nyberg et al. (2023) propose $472\,000\text{ km}^2$.

Another finding of this work is the sensitivity of validation outcomes to the validation method. The practice of setting aside a portion of labelled data for the purpose of validation is common (e.g., Isikdogan et al., 2017, 2020; Nyberg et al., 2023). However, we

argue that the selection data sites for labelling leads to an implicit bias towards more favourable and useful sites for the model even when the validation sites are spatially distinct from the training sites. In contrast, when we chose locations randomly across the entire globe, our validation statistics dropped as the model was confronted with environments further removed from its training. Similarly, when validating the ESA WorldCover product, Zanaga et al. (2022) used a labour intensive validation approach with a very large number of scattered validation sites. As a result, their overall validation accuracy is cited as 74.4%. A fundamental issue is the absence of common benchmark datasets that could allow us to cross compare models and results. In computer vision, there are several high-quality benchmark datasets such as ImageNet (Deng et al., 2009) and COCO (Lin et al., 2015) that can be used to evaluate model performance. There is an urgent need to establish such benchmarks for Earth Observation.

5 | CONCLUSION

Overall, this work shows the potential of deep learning when applied to global scale studies of fluvial systems. Of particular importance is the fact that the workflow presented here required 4 weeks of computation with a very high level of automation. This means that monthly monitoring of river systems at global scale is now possible with 10 m resolution Sentinel-2 data. On this basis, future models that distinguish sedimentary features based on their formation mechanism such as point bars, alluvial fans and ephemeral channels are now within reach.

ACKNOWLEDGEMENTS

We thank two anonymous reviewers for their constructive comments. We also thank Prof. Toby Breckon for ongoing discussions on computer vision and deep learning.



CONFLICT OF INTEREST STATEMENT

The authors have no conflicts to declare.

DATA AVAILABILITY STATEMENT

The full code for the workflow is available as Carbonneau (2023a). Readers can also access the full global dataset of 468 tiles at native 10 m resolution as Carbonneau (2023b).

ORCID

Patrice E. Carbonneau  <https://orcid.org/0000-0001-8246-9491>
Simone Bizzi  <https://orcid.org/0000-0002-0588-826X>

REFERENCES

- Allen, G.H. & Pavelsky, T.M. (2018) Global extent of rivers and streams. *Science*, 361, 585–588. Available from: <https://doi.org/10.1126/science.aat0636>
- Altenau, E.H., Pavelsky, T.M., Durand, M.T., Yang, X., Frasson, R.P.d.M. & Bendezu, L. (2021) SWOT River Database (SWORD) (Version v0) [Data set]. Zenodo. Available from: <https://doi.org/10.5281/zenodo.3898570>
- Badrinarayanan, V., Kendall, A. & Cipolla, R. (2017) SegNet: a deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12), 2481–2495. Available from: <https://doi.org/10.1109/TPAMI.2016.2644615>

- Brierley, G., Fryirs, K., Cullum, C., Tadaki, M., Huang, H.Q. & Blue, B. (2013) Reading the landscape: integrating the theory and practice of geomorphology to develop place-based understandings of river systems. *Progress in Physical Geography*, 37(5), 601–621. Available from: <https://doi.org/10.1177/0309133313490007>
- Brostow, G.J., Fauqueur, J. & Cipolla, R. (2009) Semantic object classes in video: a high-definition ground truth database. *Pattern Recognition Letters*, 30(2), 88–97. Available from: <https://doi.org/10.1016/j.patrec.2008.04.005>
- Carbonneau, P.E. (2023a) *Semantic Classification of Rivers and Bars* (v1.3). Zenodo. Available from: <https://doi.org/10.5281/zenodo.10016467>
- Carbonneau, P.E. (2023b) *Global Mapping of Rivers and Gravel Bars [Data set]*. Durham, UK: Durham University. Available from: <https://doi.org/10.15128/r17w62f824x>
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K. & Yuille, A.L. (2018) DeepLab: semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4), 834–848. Available from: <https://doi.org/10.1109/TPAMI.2017.2699184>
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K. & Fei-Fei, L. (2009) ImageNet: a large-scale hierarchical image database. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. Miami, FL, USA: IEEE. Available from: <https://doi.org/10.1109/CVPRW.2009.5206848>
- Déry, S.J. & Brown, R.D. (2007) Recent Northern Hemisphere snow cover extent trends and implications for the snow-albedo feedback. *Geophysical Research Letters*, 34(22), L22504. Available from: <https://doi.org/10.1029/2007GL031474>
- Downing, J.A., Cole, J.J., Duarte, C.M., Middelburg, J.J., Melack, J.M., Prairie, Y.T. et al. (2012) Global abundance and size distribution of streams and rivers. *Inland Waters*, 2(4), 229–236. Available from: <https://doi.org/10.5268/IW-2.4.502>
- Frasson, R.P.d.M., Pavelsky, T.M., Fonstad, M.A., Durand, M.T., Allen, G.H., Schumann, G. et al. (2019) Global relationships between river width, slope, catchment area, meander wavelength, sinuosity, and discharge. *Geophysical Research Letters*, 46(6), 3252–3262. Available from: <https://doi.org/10.1029/2019GL082027>
- Fryirs, K.A. (2017) River sensitivity: a lost foundation concept in fluvial geomorphology. *Earth Surface Processes and Landforms*, 42(1), 55–70. Available from: <https://doi.org/10.1002/esp.3940>
- Gao, B. (1996) NDWI—a normalized difference water index for remote sensing of vegetation liquid water from space. *Remote Sensing of Environment*, 58(3), 257–266. Available from: [https://doi.org/10.1016/S0034-4257\(96\)00067-3](https://doi.org/10.1016/S0034-4257(96)00067-3)
- Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D. & Moore, R. (2017) Google Earth Engine: planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202, 18–27. Available from: <https://doi.org/10.1016/j.rse.2017.06.031>
- Grill, G., Lehner, B., Thieme, M., Geenen, B., Tickner, D., Antonelli, F. et al. (2019) Mapping the world's free-flowing rivers. *Nature*, 569(7755), 215–221. Available from: <https://doi.org/10.1038/s41586-019-1111-9>
- Ham, D. & Church, M. (2002) *Channel island and active channel stability in the lower Fraser River gravel reach*. Fraser River Research Group Report. Vancouver, Canada: Department of Geography, University of British Columbia.
- He, K., Zhang, X., Ren, S. & Sun, J. (2016) Deep residual learning for image recognition. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, NV, USA: IEEE, pp. 770–778. Available from: <https://doi.org/10.1109/CVPR.2016.90>
- Isikdogan, F., Bovik, A.C. & Passalacqua, P. (2017) Surface water mapping by deep learning. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 10(11), 4909–4918. Available from: <https://doi.org/10.1109/JSTARS.2017.2735443>
- Isikdogan, L.F., Bovik, A. & Passalacqua, P. (2020) Seeing through the clouds with DeepWaterMap. *IEEE Geoscience and Remote Sensing Letters*, 17(10), 1662–1666. Available from: <https://doi.org/10.1109/LGRS.2019.2953261>
- Jégou, S., Drozdal, M., Vazquez, D., Romero, A. & Bengio, Y. (2017) The one hundred layers tiramisu: fully convolutional DenseNets for semantic segmentation. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. Honolulu, HI, USA: IEEE, pp. 1175–1183. Available from: <https://doi.org/10.1109/CVPRW.2017.156>
- Karra, K., Kontgis, C., Statman-Weil, Z., Mazzariello, J.C., Mathis, M. & Brumby, S.P. (2021) Global land use/land cover with Sentinel 2 and deep learning. In: *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*. Brussels, Belgium: IEEE, pp. 4704–4707. Available from: <https://doi.org/10.1109/IGARSS47720.2021.9553499>
- Lehner, B., Verdin, K. & Jarvis, A. (2008) New global hydrography derived from spaceborne elevation data. *EOS. Transactions of the American Geophysical Union*, 89(10), 93–94. Available from: <https://doi.org/10.1029/2008EO100001>
- Lin, T.-Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J. et al. (2015) Microsoft COCO: Common Objects in Context. Available from: <https://doi.org/10.48550/arXiv.1405.0312>
- Long, J., Shelhamer, E. & Darrell, T. (2015) Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Boston, MA, USA: IEEE, pp. 3431–3440.
- Moortgat, J., Li, Z., Durand, M., Howat, I., Yadav, B. & Dai, C. (2022) Deep learning models for river classification at sub-meter resolutions from multispectral and panchromatic commercial satellite imagery. *Remote Sensing of Environment*, 282, 113279. Available from: <https://doi.org/10.1016/j.rse.2022.113279>
- Nyberg, B., Henstra, G., Gawthorpe, R.L., Ravnås, R. & Ahokas, J. (2023) Global scale analysis on the extent of river channel belts. *Nature Communications*, 14(1), 2163. Available from: <https://doi.org/10.1038/s41467-023-37852-8>
- Pekel, J.-F., Cottam, A., Gorelick, N. & Belward, A.S. (2016) High-resolution mapping of global surface water and its long-term changes. *Nature*, 540, 418–422. Available from: <https://doi.org/10.1038/nature20584>
- Piégay, H., Arnaud, F., Belletti, B., Bertrand, M., Bizzi, S., Carbonneau, P. et al. (2020) Remotely sensed rivers in the Anthropocene: state of the art and prospects. *Earth Surface Processes and Landforms*, 45(1), 157–188. Available from: <https://doi.org/10.1002/esp.4787>
- Raymond, P.A., Hartmann, J., Lauerwald, R., Sobek, S., McDonald, C., Hoover, M. et al. (2013) Global carbon dioxide emissions from inland waters. *Nature*, 503(7476), 355–359. Available from: <https://doi.org/10.1038/nature12760>
- RGI Consortium. (2017) *Randolph Glacier Inventory—A Dataset of Global Glacier Outlines, Version 6 [Data set]*. Boulder, Colorado USA: National Snow and Ice Data Center.
- Simonyan, K. & Zisserman, A. (2015) Very deep convolutional networks for large-scale image recognition. Available from: <https://doi.org/10.48550/arXiv.1409.1556>
- Skakun, S., Wevers, J., Brockmann, C., Doxani, G., Aleksandrov, M., Batič, M. et al. (2022) Cloud Mask Intercomparison eXercise (CMIX): an evaluation of cloud masking algorithms for Landsat 8 and Sentinel-2. *Remote Sensing of Environment*, 274, 112990. Available from: <https://doi.org/10.1016/j.rse.2022.112990>
- Syvitski, J., Ángel, J.R., Saito, Y., Overeem, I., Vörösmarty, C.J., Wang, H. et al. (2022) Earth's sediment cycle during the Anthropocene. *Nature Reviews Earth & Environment*, 3(3), 179–196. Available from: <https://doi.org/10.1038/s43017-021-00253-w>
- Tickner, D., Opperman, J.J., Abell, R., Acreman, M., Arthington, A.H., Bunn, S.E. et al. (2020) Bending the curve of global freshwater biodiversity loss: an Emergency Recovery Plan. *Bioscience*, 70(4), 330–342. Available from: <https://doi.org/10.1093/biosci/biaa002>
- Xu, H. (2006) Modification of normalised difference water index (NDWI) to enhance open water features in remotely sensed imagery. *International Journal of Remote Sensing*, 27(14), 3025–3033. Available from: <https://doi.org/10.1080/01431160600589179>

Zanaga, D., Van De Kerchove, R., Daems, D., De Keersmaecker, W., Brockmann, C., Kirches, G. et al. (2022) *ESA WorldCover 10 m 2021 v200 (Version v200) [Data set]*. Zenodo. Available from: <https://doi.org/10.5281/zenodo.7254221>

SUPPORTING INFORMATION

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

How to cite this article: Carbonneau, P.E. & Bizzi, S. (2024) Global mapping of river sediment bars. *Earth Surface Processes and Landforms*, 49(1), 15–23. Available from: <https://doi.org/10.1002/esp.5739>