



RivaMap: An automated river analysis and mapping engine



Furkan Isikdogan^{a,b}, Alan Bovik^a, Paola Passalacqua^{b,*}

^a Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX, USA

^b Department of Civil, Architectural and Environmental Engineering, and Center for Research in Water Resources, The University of Texas at Austin, Austin, TX, USA

ARTICLE INFO

Article history:

Received 4 June 2016

Received in revised form 24 February 2017

Accepted 30 March 2017

Available online 13 April 2017

Keywords:

Remote sensing

Image processing

Landsat

Large-scale mapping

River delineation

Width estimation

ABSTRACT

Rivers are essential to the Earth's water cycle and deeply impact many human societies and ecosystems, yet they are currently monitored poorly at the global scale. In-situ gauging stations are distributed sparsely and heterogeneously and do not cover much of the world, whereas remotely sensed images are spatially and temporally dense and available globally. Remotely sensed multispectral images, such as the ones acquired by Landsat missions, are available to enable the analysis and surveying of rivers using suitable algorithms. However, existing algorithms are limited in ways that restrict the coverage of the produced results and that prevent the automated analysis of river networks at large scales over short periods of time. Ideally, river maps should be as "live" as possible, e.g., computed quickly and continuously as new Earth imaging data becomes available. Towards advancing progress on this problem, we describe an automated river analysis and mapping engine, RivaMap, that enables the computation of large-scale hydrography data sets from remotely sensed data in a short period of time. RivaMap facilitates water resource management by providing tools to delineate rivers and to estimate their width. As a practical application of RivaMap, we present a continental-scale centerline and width data set of North American rivers, that is automatically computed on Landsat data. We validate our mapping engine by comparing the RivaMap-generated data to a similar data set, NARWidth, and also to in-situ measurements. Our experimental results show that RivaMap is able to efficiently and accurately extract rivers from remotely sensed images at large scales. The outcomes of this research, the software, and the computed exemplary data set are publicly available.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

River networks are the veins of the Earth, delivering fresh water to ecosystems and societies. Despite the crucial importance of fresh water for life, obtaining accurate global knowledge of the spatiotemporal dynamics of surface freshwater remains challenging¹ (Alsdorf et al., 2007). In-situ networks of gauges are rare and distributed sparsely, even in well-monitored countries (Gleason & Smith, 2014; Pavelsky et al., 2014). Satellite imagery, on the other hand, provides broad coverage of the Earth at high spatial and temporal resolutions. Landsat missions, for example, have provided multispectral images of the Earth with global coverage since 1972 (Kovalsky & Roy, 2013; Loveland & Dwyer, 2012; USGS Landsat Missions, 2016). Remotely sensed images can be utilized to develop large-scale data

sets that can help enable the monitoring of hydrological change in space and time.

Algorithms that create hydrological data sets in a fully automated manner have great potential value when studying rivers and identifying the causes and consequences of global environmental change. Existing methods have demonstrated the effectiveness of using remote sensing data to create large-scale hydrological data sets, such as the HydroSHEDS (Hydrological data and maps based on Shuttle Elevation Derivatives at multiple Scales) (Lehner et al., 2008), NARWidth (North American River Width) (Allen & Pavelsky, 2015), and GWD-LR (Global Width Database for Large Rivers) (Yamazaki et al., 2014) data sets. HydroSHEDS used a steepest descent based flow direction algorithm to trace rivers on Shuttle Radar Topography Mission (SRTM) data, whereas NARWidth used a software tool, RivWidth (Pavelsky & Smith, 2008), to delineate and estimate the widths of the North American rivers on Landsat images. These remotely sensed data sets provide valuable information regarding river width and flow direction. However, the input data and aspects of the algorithms that were used to create these data sets imposed some limitations on their accuracy and usability. For example, HydroSHEDS and NARWidth contain limited or no data on certain regions, including coastal areas which are hotspots for environmental change. Difficulties in

* Corresponding author.

E-mail address: paola@austin.utexas.edu (P. Passalacqua).

¹ Very recently, Pekel et al. (2016) described an advanced tool to capture the spatiotemporal dynamics of global surface water, which appeared as the current manuscript was nearing publication.

defining the boundary between rivers and oceans led to the exclusion of some rivers in coastal regions, since they were treated as oceans. Furthermore, the creation of these data sets required laborious manual inspection and corrections. For example, an extensive manual correction protocol was required to improve the quality of HydroSHEDS (Lehner et al., 2008). NARWidth used an algorithm that required a binary water mask indicating water and land pixels (Pavelsky & Smith, 2008) as input. These land-water masks were obtained by an automatic water classification procedure (Li & Sheng, 2012) followed by visual inspection and correction. To automatically derive a global database for large rivers, GWD-LR used the SRTM Water Body Data (SWBD) and the HydroSHEDS flow direction map. As a result of using SWBD as a baseline water mask, rivers narrower than 300 m have limited coverage in the GWD-LR data set.

The amount of time and effort needed to create large-scale data sets using methods that require manual corrections greatly limits the analysis of river networks over time. A completely automatic data set creation engine, on the other hand, would significantly reduce the effort, and the temporal resolution of the obtained results would be limited only by the frequency of image acquisition on a given region. For example, the temporal resolution can be as low as 8 days for Landsat-derived data sets, since the Landsat 7 and 8 satellites complete imaging the entire Earth every 16 days with an 8-day offset from each other (USGS Landsat 8, 2016). Automatically deriving surface water data sets from Landsat images as they are acquired would help facilitate the “live” analysis of water resources over time.

The results produced by such an algorithm could be used in a wide variety of studies that help improve life on Earth. For example, an automatically created continental-scale spatiotemporal data set of hydrological measurements would allow monitoring water resources over vast spatial extents and over time. Such a data set could offer crucial information to scientists towards better understanding the implications of human-induced activities and global environmental change on ecosystems. Subsequently, the outcomes could be used to inform policy makers and governments to act more effectively.

Recently, we developed a method (Isikdogan et al., 2015) of automatically extracting rivers at multiple scales directly from satellite imagery, without requiring a binary water mask as input. This river extraction algorithm runs in a fully automatic manner and can quickly and accurately extract complex network structures, such as braided rivers and deltaic patterns. Adapting this algorithm can eliminate or alleviate the aforementioned limitations, while providing an automatically updating monitoring and analysis tool that can produce automatic updates whenever required.

We have significantly improved this earlier, preliminary development work on the extraction of river networks, by creating a more sophisticated model that produces better estimates of river centerlines and widths. Using this improved river network extraction algorithm, we have developed a fully automated river analysis and mapping engine, called RivaMap, that is able to create large scale maps of rivers (Section 2). We demonstrate a large-scale application of RivaMap (Section 3) by deriving a data set of North American rivers from Landsat data. We systematically compare this exemplary data set that was automatically generated by RivaMap to existing data sets to evaluate its performance.

The major contribution of this paper is a fully automatic river mapping software package that is capable of extracting rivers from Landsat imagery without any user intervention, that maps coastal systems and braided rivers, that can process Landsat images at large scales, and that can be used to monitor water resources over large spatiotemporal extents. A secondary contribution of this paper is an exemplary automatically generated width data set of North American rivers, which demonstrates the capabilities of the proposed method at continental-scale. By applying our RivaMap tool, a user may generate their own similar dataset using the most recently

available Landsat imagery whenever they require it. The software tools and the generated data set are freely available at <http://live.ece.utexas.edu/research/rivamap/>.

2. Methods

This section presents the methods employed in our river analysis engine, RivaMap. The processes composing the engine include the acquisition of the images, creating cloud-free composites, enhancing water features, extracting and delineating rivers, and estimating river width (Fig. 1).

2.1. Acquisition of Landsat images

There are several ways to obtain Landsat data. Landsat images can be accessed free of charge through USGS websites and third party environmental data catalogs: LandsatLook Viewer, USGS GloVis: The Global Visualization Viewer, USGS Earth Explorer (Landsat Science: Where to Get Data, 2016), and Google Earth Engine (2016). Google Earth Engine is a cloud computing platform that has a catalog of public data sets and a library of functions for performing analyses on the data. The platform is optimized for parallel processing to reduce computing time (Thau, 2015). The data catalog provides fast and easy access to images from Earth-observing satellites, including Landsat archives. An API allows users to access a library of functions and run scripts on the Earth Engine. We use the Earth Engine Python API to select, preprocess, and download images in batch.




Fig. 1. Flowchart of the automated river analysis and mapping engine, RivaMap.

2.2. Creating water-enhanced cloud-free composites

The quality of the automatically generated maps depends heavily on the quality of the input images. Cloud coverage, in particular, has a prominent impact on the quality of the information derived from remotely sensed imagery. For example, it might not be possible to derive useful information from images that have large cloud coverage. To reduce the impact of clouds, cloud-free scenes can be selected automatically, using the estimated percentage of cloud coverage that is provided for each scene in the Landsat archives. However, in many cases where cloud-free scenes might not be available, further processing of images is necessary to remove clouds.

We employ a multitemporal complementation based method to create cloud-free Landsat composites using the Google Earth Engine. The method creates composites that minimize cloud coverage given a set of scenes by selecting the least cloudy images for a given location and time interval and by computing the temporal sample median of those pixels that are least likely to be cloudy. To determine the cloud-likelihood of a pixel, the Earth Engine uses a simple cloud scoring algorithm that makes use of brightness, temperature, and a normalized difference snow index (Lobell et al., 2015).

To enhance water features and suppress built-up land, vegetation, and soil noise in the cloud-free composites, we employ the Modified Normalized Difference Water Index (MNDWI) (Xu, 2006), which is defined as follows:

$$MNDWI = \frac{Green - MIR}{Green + MIR} \quad (1)$$

where *MIR* is a middle infrared band (e.g., band 5 in Landsat 7 and band 6 in Landsat 8) and *Green* is a green band (e.g., band 2 in Landsat 7 and band 3 in Landsat 8). After computing the water index, we download the water-enhanced cloud-free composite images from the Earth Engine and further process the images locally to extract river centerline and width.

2.3. Curvilinear structure extraction

In our recent work (Isikdogan et al., 2015), we have demonstrated the effectiveness of a modified and elaborated version of the multiscale singularity index (Muralidhar et al., 2013a,b) for achieving fully automated extraction of river networks from remotely sensed images. The multiscale singularity index is a dimensionless ratio index that is defined as:

$$(\psi f)(x, y, \sigma) = \frac{|f_{0,\theta,\sigma}(x, y) f_{2,\theta,\sigma}(x, y)|}{1 + |f_{1,\theta,\sigma}(x, y)|^2} \quad (2)$$

where, $f_{0,\theta,\sigma}(x, y)$, $f_{1,\theta,\sigma}(x, y)$, and $f_{2,\theta,\sigma}(x, y)$ are the responses to the zero, first, and second order derivatives of Gaussians, and σ and $\theta(x, y)$ specify the scale of the Gaussians and direction of the derivatives, respectively. The directions of the derivatives are orthogonal to the curvilinear singularities in the input and are estimated by computing the direction along which the second order derivative attains a local extremum (Freeman & Adelson, 1991; Muralidhar et al., 2012). The constant a in the denominator of Eq. (2) has a fixed value of 1.7754, which achieves maximum attenuation of the side lobe response of the index as shown by Muralidhar et al. (2013b). This attenuation is important for minimizing responses to step-like image edges.

The singularity index responds strongly to curvilinear structures, where the second derivative in the numerator is large, and weakly to edges where the first derivative in the denominator is large. Since rivers are curvilinear structures, the singularity index can be used to extract them from remotely sensed images. The input of the index can be any image that has a contrast between water and land. Such

input images can be obtained from a single band, such as a middle infrared band, or from multiple bands using a water index. The input images do not require thresholding since the algorithm does not require a binary water mask as input. We use the MNDWI to obtain the input images, since it enhances the contrast between land and water. Using MNDWI reduces the chances of detecting other curvilinear structures, such as roads, as rivers.

The singularity index responds to curvilinear structures in the input regardless of their polarity, i.e. whether the structure or its neighboring area has larger intensity values. As a result, islands in an input image could be mistaken for rivers, especially in braided rivers. The island response can be discriminated from the river response by retaining the sign of the second derivative in the singularity index (2) as follows:

$$(\psi f)(x, y, \sigma) = \frac{|f_{0,\theta,\sigma}(x, y)| f_{2,\theta,\sigma}(x, y)}{1 + |f_{1,\theta,\sigma}(x, y)|^2} \quad (3)$$

Rivers of different size can be captured by computing the singularity index over multiple scales. Rivers that are wider than σ_{min} and narrower than σ_N can be extracted by computing $(\psi f)(x, y, \sigma)$ for $\sigma_n = \sigma_{min} \sqrt{2^{(n-1)}}$, $n = 1, 2, \dots, N$. However, the singularity index response is expected to decrease with increasing scale, since the amplitudes of spatial derivatives generally decrease as the scale increases (Lindeberg, 1998). Motivated by Lindeberg's γ -normalized derivatives (Lindeberg, 1998), the drop in the response with increasing scale can be compensated by scale-normalizing the singularity index as $(\psi_{norm} f)(x, y, \sigma) = \sigma^2 (\psi f)(x, y, \sigma)$ (Muralidhar et al., 2013b). We compute the scale-normalized singularity index over the scales $\sigma_n = \sigma_{min} \sqrt{2^{(n-1)}}$, $n = 1, 2, \dots, N$, where the maximum scale σ_N is upper bounded by the image size. The number of scales N can be set to its upper bound, which is determined automatically for each input image by finding the largest scale such that the filter dimension is smaller than both image dimensions (Isikdogan et al., 2015).

The singularity index is sensitive to the intensity of the input signal. Therefore, the presence of intensity variations in the input image could lead to variations in the output. In remotely sensed images, any non-uniformity of atmospheric conditions could cause structurally irrelevant intensity variations. To provide invariance to local intensity variations, an input image should be locally debiased. This can be achieved by subtracting a Gaussian filtered version of the input image from the input. As explained in Muralidhar et al. (2013b), the standard deviation of the Gaussian filter σ_g can be chosen as $\sigma_g \geq \sigma_n \sqrt{\frac{1-\epsilon^2}{\epsilon^2}}$, where $\epsilon = 0.2$ and σ_n is the scale of the input signal. We debias the input image by using a Gaussian with $\sigma_g = 5\sigma_n > \sigma_n \sqrt{\frac{1-\epsilon^2}{\epsilon^2}}$ before computing the singularity index at each scale σ_n , $n = 1, 2, \dots, N$.

We combine the multiscale singularity index response into a single-band overall response image to highlight channelized areas in a given MNDWI image. Our initial approach to the overall response computation (Isikdogan et al., 2015) involved pooling the maximum singularity index response across scales. However, this process could create artificial ripple-like discontinuities near the channel banks, requiring post processing of the response with an adaptive smoothing algorithm. In this work, we use the Euclidean norm of the multiscale singularity index response across scales to compute the overall magnitude of the response as:

$$\|(\psi f)(x, y)\| = \sqrt{\sum_{n=1}^N (\psi f)(x, y, \sigma_n)^2} \quad (4)$$


Using the Euclidean norm instead of maximum-pooling produces a seamless response, therefore removing the need for an adaptive

smoothing algorithm. The overall magnitude of the multiscale singularity index highlights channels, helps isolate rivers from other high-MNDWI land features, and defines a soft boundary between rivers and oceans in coastal regions (Fig. 2).


2.4. River centerline delineation

River centerlines can be delineated by suppressing the non-maxima response on the overall magnitude of the multiscale singularity index response along the dominant orientation. Similar to the non-maxima suppression procedure in the Canny edge detector (Canny, 1986), the overall singularity index response can be non-maxima suppressed to obtain the river centerline. Our automated river analysis engine, RivaMap, executes a process of non-maxima suppression at each pixel along the dominant orientation on the Euclidean norm of the multiscale singularity index response. The dominant orientation θ is orthogonal to the channel direction, computed as explained in Freeman & Adelson (1991) and Muralidhar et al. (2012). The channel direction θ_{\perp} at a given spatial location denotes the elongation axis of the channel along which the width is computed (Fig. 3).

To produce a binary map of river centerlines, we apply a hysteresis threshold to the non-maxima suppressed response. Hysteresis thresholding uses two threshold levels: τ_l (low) and τ_h (high). Values below τ_l are discarded, values above τ_h are classified as centerlines,



(a)



(b)

Fig. 2. Highlighting channelized areas on MNDWI response. (a) MNDWI response, (b) overall magnitude of the multiscale singularity index response. Wax Lake and Atchafalaya Deltas, LA.



Fig. 3. A quiver plot of channel direction θ_{\perp} and magnitude $\|(\psi f)(x, y)\|$ of the overall singularity index response in the Wax Lake Delta, LA. The length and orientation of the lines show the magnitude and angle of the singularity index response at a given spatial location, respectively.

and values between τ_l and τ_h are classified as centerlines only if they are adjacent to other centerline pixels. This approach helps preserve the connectivity of rivers while eliminating spurious responses.

In our earlier work (Isikdogan et al., 2015), we chose τ_h using Otsu's threshold selection method (Otsu, 1975) and fixed τ_l to a fraction of τ_h for each input image. This strategy works if the channelized areas in the image are sufficiently large such that the distribution of ψ is bimodal. In large-scale processing, however, the distribution of ψ is usually unimodal (e.g., Fig. 4), since the rivers constitute a small portion of the input images. Therefore, using Otsu's method to determine the thresholds for each input tile can amplify noise in images where there is little or no river presence. We address this problem by determining the thresholds based on the overall distribution of ψ at continental-scale, which was estimated using the Gaussian kernel density estimation function in R (R Documentation, 2016). To determine the thresholds, we use a simple geometrical




Fig. 4. Determining the thresholds: lower threshold τ_l at the peak of the distribution (0.012), middle threshold τ_m at the corner of the distribution (0.061), and upper threshold τ_h at $\tau_m + (\tau_m - \tau_l)$.

method that is based on Rosin's unimodal thresholding algorithm (Rosin, 2001). Rosin's algorithm consists of drawing a straight line from the peak to the end of the distribution and selecting a threshold at the point where the perpendicular distance between the line and the distribution function is maximum (Fig. 4). Using this method, we first find a middle threshold τ_m at the corner of the distribution. Then, we choose the lower and upper thresholds as $\tau_m \pm \epsilon_T$, where ϵ_T is a tolerance variable that is empirically set to the distance between the peak of the distribution and τ_m . More specifically, we choose the lower threshold τ_l at the peak of the distribution and the upper threshold τ_h at $\tau_m + (\tau_m - \tau_l)$, bisecting the middle threshold between τ_l and τ_h . In our experiments, we estimated the overall distribution of ψ by computing the density of 10^8 centerline points (i.e., $\psi > 0$) obtained by computing ψ over randomly sampled cloud-free composite image patches from North America. Our analysis yielded the lower and upper thresholds $\tau_l = 0.012$ and $\tau_h = 0.11$, respectively. By following this procedure, RivaMap is able to reliably detect rivers in a variety of environments, including coastal areas (Fig. 5).

2.5. Width estimation

Locating the dominant scale at which the singularity index response is the largest is useful for estimating the river width at a given spatial coordinate. A simple approach to estimate the width would be determining the dominant scale where the singularity index response is maximum, $\sigma_d = \operatorname{argmax}_{\sigma} (\psi f)(x, y, \sigma)$, then correlating σ_d with the width. However, such an approach can only provide a broad description of the width while failing to capture sub-scale width differences between rivers. The maximum singularity index response in scale can instead be estimated by interpolation (Fig. 6). Estimating the accurate scale that maximizes the singularity index response is particularly important at coarser scales, since the step size between scales grows exponentially.

To estimate the dominant scale at sub-scale accuracy, we first locate σ_d , the scale that has the highest singularity index response. Then, we fit a quadratic curve to the singularity index response at σ_d and its neighboring scales. Finally, we locate the dominant scale as the vertex of the quadratic function. Formally, the dominant scale $\hat{\sigma}$ is computed as:

$$\hat{\sigma} = \frac{1}{2} \frac{\sigma_{d+1}^2 \Delta\psi_1 + \sigma_d^2 \Delta\psi_2 + \sigma_{d-1}^2 \Delta\psi_3}{\sigma_{d+1} \Delta\psi_1 + \sigma_d \Delta\psi_2 + \sigma_{d-1} \Delta\psi_3} \quad (5)$$




Fig. 6. An example of estimation of scale for a given river centerline point. The vertical axis shows the magnitude of the singularity index at a given scale σ . The dashed line indicates the estimated location of the scale that maximizes the singularity index response. The horizontal axis σ is plotted on a logarithmic scale.

where $\sigma_{d\pm 1}$ are the neighbors of σ_d in scale and

$$\begin{aligned} \Delta\psi_1 &= (\psi f)(x, y, \sigma_{d-1}) - (\psi f)(x, y, \sigma_d) \\ \Delta\psi_2 &= (\psi f)(x, y, \sigma_{d+1}) - (\psi f)(x, y, \sigma_{d-1}) \\ \Delta\psi_3 &= (\psi f)(x, y, \sigma_d) - (\psi f)(x, y, \sigma_{d+1}) \end{aligned} \quad (6)$$

at a given spatial coordinate (x, y) .

To compute an accurate estimate of the river width, the algorithm needs the dominant scale to be larger than or equal to the second finest scale. As a result, widths of narrow rivers get overestimated if the dominant scale is smaller than the second finest scale. To estimate the widths of narrow rivers, we make use of the magnitude of the overall singularity index response $\|(\psi f)\|$, which correlates well with river width for small values of $\|(\psi f)\|$ ($\lesssim 0.25$) (Section 3, Fig. 10). We pass $\|(\psi f)\| / 0.25$ through a hyperbolic tangent function and use the result as a scale factor. The hyperbolic tangent function is used as a blending function for the scale factor, limiting the range of input values that correlate with the width to $[0, 1]$. Finally, the river




Fig. 5. Automatically delineated rivers in the Wax Lake and Atchafalaya Deltas, LA.

width w is estimated as $\hat{w} = k \times \tanh(\|\psi f\|/0.25) \times \hat{\sigma}$, where k is a scalar variable that depends on the resolution of the input image, allowing conversion from scale in pixels to width in meters.

2.6. Direct use of intermediate outputs

The final products of RivaMap consist of river centerline and width measurements on a given input image. In addition to these final products, the responses produced by intermediate processes, such as the overall magnitude of the multiscale singularity index response, can be used to analyze and quantify the dynamics of rivers. These intermediate outputs can be used to directly develop surface water and inundation metrics to assist in the analysis of patterns in river networks, particularly in coastal systems. For example, the overall magnitude of the positive and negative polarity singularity index (Eq. (3)) response can be used to quantify the islandness and channelness of a river delta (e.g. Fig. 2b highlights the channelized areas). Subsequently, channel and island gain/loss in response to inundation can be measured as the difference of channelness and islandness over time. Furthermore, channel and island instability in a delta can be formulated as the variance of channelness/islandness over time, showing the parts of the delta that changed the most over a given period of time (Fig. 7). Other metrics, such as river width distribution, sinuosity, and symmetry (Liang et al., 2016), can be derived from the river width, centerline, and orientation information.

3. Experiments

To evaluate the effectiveness of the RivaMap engine for large-scale data set creation, we first automatically generated a map of the North American rivers and small lakes using RivaMap. Then, we compared the resultant data against similar data sets. In our experiments, we empirically set the minimum scale $\sigma_{min} = 1.2$ pixels, as the corresponding output of the finest scale singularity index response captured marginally more detail than its recommended value $\sigma_{min} =$

1.5 (Isikdogan et al., 2015; Muralidhar et al., 2013b). We also set the width scalar k to $2L_m$, where the resolution of the input Landsat images is $L_m = 30$ m/pixel.

3.1. Developing a width data set of North American rivers

To create a data set that spans the entire North America, we first created a grid that covers North America with cells of 1-degree height and width. To prevent boundary artifacts, the cells were buffered 50% to overlap. For each cell in the grid, we generated cloud-free composite images using Landsat-8 data drawn from the last 3 years. We created a map of river centerline and width for each cell using RivaMap. Processing of each cell ($\sim 6000 \times 6000$ pixel image) took ~ 2 min on average on a single core of a desktop CPU. We processed all the cells in parallel on 8 cores of an Intel(R) Core(TM) i7 CPU, which took about 1 day to complete. Finally, we aggregated the results we obtained from individual cells to create a continental-scale data set.

The data set consists of 4.8×10^8 data points, including measurements of rivers (Fig. 8). Each data point in the data set contains latitude, longitude, and the corresponding width and channel orientation of a centerline point at any given spatial location. Also included in the data set are reservoirs and small and mid-size lakes. Rivers in the data set can be differentiated from other types of surface water by their length to width ratio. The data set excludes bodies of surface water larger than the tile size. The tile size could be increased to include larger water bodies, such as the great lakes. However, larger tiles would require more memory.

3.2. Comparison to other data sets

We compared our RivaMap-generated data set to the NARWidth data set and a set of in-situ measurements (Allen & Pavelsky, 2015). Rivers wider than 2000 m were excluded in NARWidth as they made up a small portion of all measurements but skewed the results significantly. Consequently, we compared only rivers narrower than




Fig. 7. Channel (blue) and island (red) instability in Wax Lake and Atchafalaya Deltas between 1985 and 2015 (light: stable, dark: unstable).




Fig. 8. A simple visualization of the RivaMap-generated data set showing the overall coverage of the dataset. Channels that have a strong overall singularity index response ($\psi > \tau_h$) and that are longer than 10 km among the overlapping points from the NARWidth and RivaMap results are drawn thicker. The data points in the figure are heavily subsampled for visualization purposes only.

2000 m, although wider rivers are included in our data set. We used the same set of in-situ river width records that were used in the validation of the NARWidth data set. The set included 1049 measurements obtained from the U.S. Geological Survey (USGS) and the Water Survey of Canada (WSC). In our comparisons, we included 902 of the 1049 in-situ measurements, which were detected automatically by our algorithm. The remaining 147 measurements were not detected by RivaMap as the overall magnitude of the singularity index response fell below the thresholds for those measurements. The in-situ gauging stations were located near the river banks, whereas RivaMap produces width estimates for every centerline sample point. To estimate the readings of width at a given in-situ gauging station, we computed the average of the width estimates for the centerline points that were within a given distance from the station, where the distance is determined by the width of the river measured at the gauging station.

We calculated Spearman's rank-order correlation coefficients (Spearman, 1904) to compare the RivaMap river width estimates to the NARWidth estimates and the in-situ measurements. We found the RivaMap width estimates to be similar to the NARWidth estimates at the in-situ measurement locations (Spearman correlation of 0.85, 0.83, and 0.81 between RivaMap and in-situ, NARWidth and in-situ, and RivaMap and NARWidth measurements, respectively, Fig. 9). Incorporating the magnitude of the overall singularity index response (Fig. 10) as a scaling factor into the width estimation formula in RivaMap improved the correlation value from 0.74 to 0.85 (Fig. 9a and b). Without the scaling factor, the algorithm could not estimate the sub-scale width at the finest scale, resulting in a threshold (Fig. 9b). Although the scaling factor helped improve the accuracy of the estimates of the widths of narrow rivers, RivaMap overestimated them. This overestimation is due to the capability of our algorithm to capture rivers of width less than the lower bound required for width estimation, such as sub-pixel wide rivers, when there is a visible contrast between the river and its surrounding area in the input image. Similarly, NARWidth also overestimated the width of narrow rivers having widths close to the resolution of the Landsat imagery (Allen & Pavelsky, 2015). RivaMap estimates showed little mean bias (-1.94 m) and root mean square error

(RMSE) (65.1 m) with respect to the in-situ measurements, larger than but comparable to the NARWidth measurements (mean bias: -0.35 m, RMSE: 38.0 m).

NARWidth reports the widths of the rivers that were likely at mean discharge, as it was derived from Landsat scenes sampled at mean discharge. The RivaMap-generated data set, on the other hand, does not guarantee that the extracted widths correspond to mean discharge, since the data set was derived from imagery selected based on minimal cloud coverage rather than on consistent hydrological conditions. Therefore, the correlations with the in-situ data and the NARWidth data set might be negatively impacted by the time difference between the compared data samples. Of course, RivaMap could be easily applied to hydrologically consistent data.

Although the RivaMap-generated data set was not based on the same Landsat scenes as NARWidth, it produced similar results. The RivaMap-generated data set overlapped with over 5.7×10^6 out of 7.7×10^6 NARWidth measurements of width < 2000 m. On these overlapping measurements, Spearman's correlation between the RivaMap estimates and NARWidth was 0.80. To compare the distributions of width in the overlap, we computed histograms of overlapping RivaMap and NARWidth river width estimates. Comparison of the histograms showed that the distributions of the two data sets were similar (Fig. 11).

3.3. Visual evaluation of results

In addition to this systematic evaluation, we evaluated the results visually over regions selected to have different characteristics. The selected regions included braided rivers and deltas, which are usually excluded or not represented accurately by existing methods. To visualize the automatically extracted rivers, we drew a line of length \hat{w} orthogonal to the river orientation at a given centerline point for every data point in the selected regions (Fig. 12). Our visual evaluation showed that RivaMap can properly extract rivers at different scales and represent braided patterns (e.g. Fig. 12a), oxbow lakes (e.g. Fig. 12a and c), and delta networks (e.g. Fig. 12b and d) correctly.

We observed some false negative and false positive centerlines, especially at high latitudes (Fig. 13). False positive centerlines were

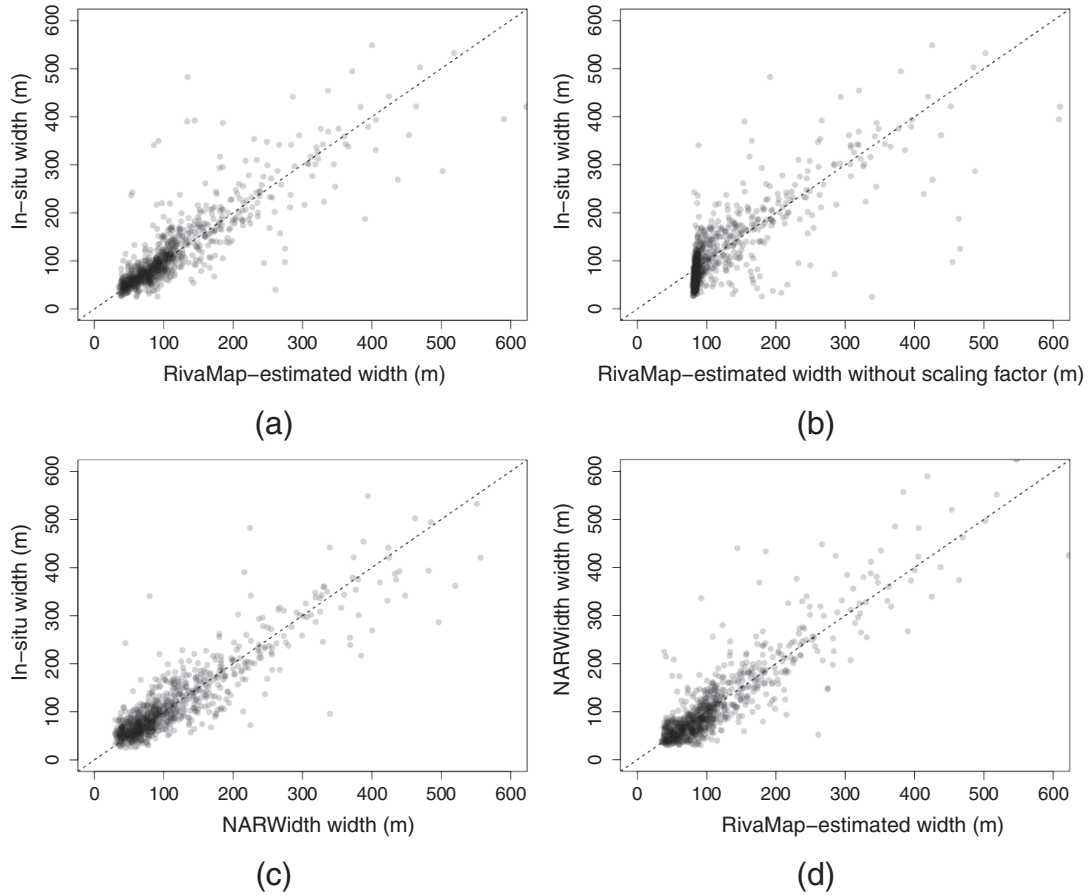


Fig. 9. Comparison of: (a) RivaMap estimates of river width against in-situ measurements (Spearman correlation: 0.85); (b) RivaMap estimates without scaling factor against in-situ measurements (Spearman correlation: 0.74); (c) NARWidth river width against in-situ measurements (Spearman correlation: 0.83); and (d) RivaMap estimates of river width against NARWidth estimates (Spearman correlation: 0.81).

mostly caused by false positive responses in the MNDWI images. MNDWI poorly differentiates between water, snow, and terrain shadows (Feng et al., 2015). As a result, false positives arose in the presence of snow and strong curvilinear shadows, where the singularity index response was strong enough to clear the thresholds. Similarly, false negatives occurred where the singularity index response fell below the thresholds. A river may produce a weak singularity index response due to several factors, including a low contrast between land and water in the input image, a width that

exceeds the input image dimensions, large bifurcation points, and obstacles breaking the continuity of the river. Most of these could be ameliorated by improving the input data, or by preprocessing it.

3.4. Limitations

The coverage and resolution of the RivaMap-produced data sets are limited by the input and the methods used. Rivers with a width narrower than or close to the resolution of the input images are




Fig. 10. Correlation between the magnitude of the overall singularity index response $\|(\psi f)\|$ and the river width measurements is observed where $\|(\psi f)\| \lesssim 0.25$.

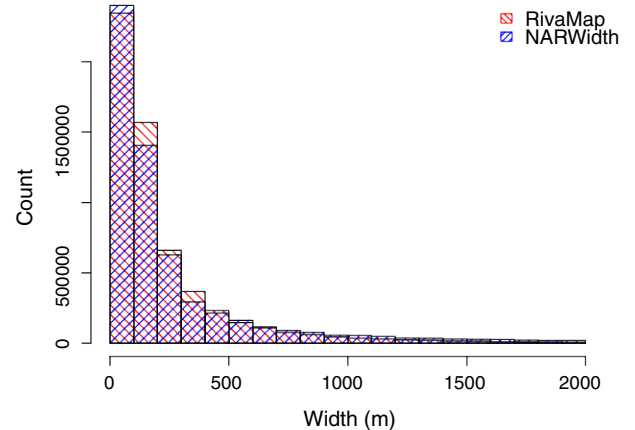


Fig. 11. Comparison of river width distributions.

either not detected or have their width overestimated. Additionally, creating completely cloud-free composites might not be possible for some regions, especially if the region had constantly cloudy weather conditions for a given time interval. Since the cloud-free composites are created by sampling the least cloudy pixels, the measurements are not guaranteed to capture rivers at mean discharge for a given




Fig. 12. Automatically extracted rivers overlaid on Google Earth imagery. The extracted rivers are visualized by drawing lines of length \hat{w} orthogonal to the channel orientation for every data point in the selected regions. (a) Mississippi and Iowa rivers near Keithsburg, IL, (b) Wax Lake Delta, LA, (c) Greater Austin Area, TX, (d) Mississippi River Delta, LA.

time period. Rivers could be extracted for a specific time by filtering the set of input images by date. However, a narrow range of dates may lead to a larger cloud coverage in the resultant composites.

We used 3-year Landsat composites as input to generate an exemplary data set having minimal cloud coverage. The time interval between the composites can be narrowed to produce results that are as “live” as possible. Technically, RivaMap does not require inputs to be composite images. Even a single Landsat scene can be used to produce river maps. From a computational perspective, RivaMap can process Landsat images as soon as they become available. However, creating live maps in certain regions, or during certain seasons, might not be possible, since clouds may eliminate the required image.

RivaMap relies on MNDWI to enhance water features. Therefore, snow and mountain shadows sometimes can be mistaken for water bodies and detected as rivers. Although the curvilinear structure extraction process helps isolate rivers using shape information, non-river features can still be mistaken for rivers if they have a high MNDWI response and curvilinear shape (Fig 13). Supplementary data such as digital elevation models and water body data sets (Feng et al., 2015; Mueller et al., 2016; Verpoorter et al., 2014; Yamazaki et al., 2015) can be used to eliminate these false positives. This could be done by replacing MNDWI with a more sophisticated algorithm that uses multiple data sources to separate snow, ice, cloud, and shadow from water. Very recently, Pekel et al. (2016) developed an effective rule-based algorithm to detect surface water. Their algorithm separates non-valid observations (snow, ice, cloud or sensor-related issues) from water using ancillary data. This new algorithm could easily be used to supply improved input data to RivaMap as an alternative to MNDWI, to produce cleaner results.

Currently, RivaMap produces a set of points rather than a fully-connected network of rivers. Therefore, there is room for improvement in the representation of river networks and further work is needed to improve river connectivity. It is trivial to connect neighbor points, where the distance between the points is equal or nearly equal to the resolution of the input data. However, several factors, such as bridges, dams, and imaging artifacts in the input images, may cause gaps between points along a river. Furthermore, RivaMap extracts the centerlines of rivers and smaller tributaries, but does




Fig. 13. An example of false positive and false negatives. Extracted rivers are overlaid on inverted MNDWI image. Mountains near Kluane Lake, Yukon Territory.

not connect the centerlines of tributaries to the centerlines of the channels they flow into. As a consequence, gaps occur at the points where channels having different widths meet. Channel orientation and width information could be used to connect broken links. For example, pairs of unconnected centerline points having similar orientation can be connected if they are within the river width. The tributaries could be connected to the main channels within the width of the larger river and within a certain range of angles between the vector of the tributary and the one of the main channels. In addition, RivaMap estimates the channel orientation, along which a channel is elongated at a given spatial location (Fig. 3), but it does not provide any information about the flow direction. Complementary data might be necessary to define upstream-downstream relationships. We plan to explore ideas, including the aforementioned ones, to improve connectivity. Having the connectivity information would help analyze global and regional statistics descriptive of river length, bifurcations, junctions, and outlets.

4. Conclusions and future work

We have proposed an automated river analysis and mapping engine, RivaMap, which is capable of creating large-scale data sets of river widths and centerlines. We derived a continental-scale river width data set from Landsat data by running the engine over the entire North America, demonstrating the potential of the engine. We validated our methodology by comparing our automatically created data set to in-situ measurements and a similar data set, NARWidth, which was also created using Landsat images. Although we used only Landsat images in this work, the engine we proposed can be extended to process other types of remotely sensed images, including digital elevation models, synthetic aperture radar images, and aerial photographs.

In the future, we plan to use the RivaMap engine and Landsat archives for developing a global-scale spatiotemporal river data set. The data set would allow monitoring of hydrological changes over space and time, providing both current and retrospective information regarding river networks. Such a data set would offer crucial information to scientists towards better understanding the implications of human-induced activities and global environmental change on ecosystems.

Acknowledgments

The authors would like to acknowledge support from the National Science Foundation, grant numbers CAREER/EAR-1350336, FESD/EAR-1135427, and OCE-1600222 awarded to P.P. The authors would also like to thank the anonymous reviewers for their constructive comments on an earlier version of this manuscript.

References

- Allen, G.H., Pavelsky, T.M., 2015. Patterns of river width and surface area revealed by the satellite-derived North American River Width data set. *Geophys. Res. Lett.* 42, 395–402.
- Alsford, D.E., Rodríguez, E., Lettenmaier, D.P., 2007. Measuring surface water from space. *Rev. Geophys.* 45.
- Canny, J., 1986. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 679–698.
- Feng, M., Sexton, J.O., Channan, S., Townshend, J.R., 2015. A global, high-resolution (30-m) inland water body dataset for 2000: first results of a topographic-spectral classification algorithm. *Int. J. Digital Earth* 1–21.
- Freeman, W.T., Adelson, E.H., 1991. The design and use of steerable filters. *IEEE Trans. Pattern Anal. Mach. Intell.* 891–906.
- Gleason, C.J., Smith, L.C., 2014. Toward global mapping of river discharge using satellite images and at-many-stations hydraulic geometry. *Proc. Natl. Acad. Sci.* 111, 4788–4791.
- Engine, Google Earth, 2016. Google earth engine: a planetary-scale geo-spatial analysis platform. <https://earthengine.google.com>.
- Isikdogan, F., Bovik, A., Passalacqua, P., 2015. Automatic channel network extraction from remotely sensed images by singularity analysis. *IEEE Geosci. Remote Sens. Lett.* 12, 2218–2221.
- Kovalskyy, V., Roy, D., 2013. The global availability of landsat 5 TM and landsat 7 ETM+ land surface observations and implications for global 30 m landsat data product generation. *Remote Sens. Environ.* 130, 280–293.
- Landsat Science: Where to Get Data, 2016. Landsat science: where to get data. http://landsat.gsfc.nasa.gov/?page_id=2367.
- Lehner, B., Verdin, K., Jarvis, A., 2008. New global hydrography derived from space-borne elevation data. *Eos. Trans. AGU* 89, 93–94.
- Li, J., Sheng, Y., 2012. An automated scheme for glacial lake dynamics mapping using landsat imagery and digital elevation models: a case study in the Himalayas. *Int. J. Remote Sens.* 33, 5194–5213.
- Liang, M., Van Dyk, C., Passalacqua, P., 2016. Quantifying the patterns and dynamics of river deltas under conditions of steady forcing and relative sea level rise. *J. Geophys. Res. Earth Surf.*
- Lindeberg, T., 1998. Feature detection with automatic scale selection. *Int. J. Comput. Vis.* 30, 79–116.
- Lobell, D.B., Thau, D., Seifert, C., Engle, E., Little, B., 2015. A scalable satellite-based crop yield mapper. *Remote Sens. Environ.* 164, 324–333.
- Loveland, T.R., Dwyer, J.L., 2012. Landsat: building a strong future. *Remote Sens. Environ.* 122, 22–29.
- Mueller, N., Lewis, A., Roberts, D., Ring, S., Melrose, R., Sixsmith, J., Lymburner, L., McIntyre, A., Tan, P., Curnow, S., 2016. Water observations from space: mapping surface water from 25 years of landsat imagery across Australia. *Remote Sens. Environ.* 174, 341–352.
- Muralidhar, G.S., Bovik, A.C., Markey, M.K., 2012. A new singularity index. *Image Processing (ICIP), 2012 19th IEEE International Conference on*. pp. 1873–1876. IEEE.
- Muralidhar, G.S., Bovik, A.C., Markey, M.K., 2013. Noise analysis of a new singularity index. *Signal Processing, IEEE Transactions on* 61, 6150–6163.
- Muralidhar, G.S., Bovik, A.C., Markey, M.K., 2013. A steerable, multiscale singularity index. *IEEE Signal Process. Lett.* 20, 7–10.
- Otsu, N., 1975. A threshold selection method from gray-level histograms. *Automatica* 11, 23–27.
- Pavelsky, T.M., Durand, M.T., Andreadis, K.M., Beighley, R.E., Paiva, R.C., Allen, G.H., Miller, Z.F., 2014. Assessing the potential global extent of swot river discharge observations. *J. Hydrol.* 519, 1516–1525.
- Pavelsky, T.M., Smith, L.C., 2008. RivWidth: a software tool for the calculation of river widths from remotely sensed imagery. *IEEE Geosci. Remote Sens. Lett.* 5, 70–73.
- Pekel, J.-F., Cottam, A., Gorelick, N., Belward, A.S., 2016. High-resolution mapping of global surface water and its long-term changes. *Nature*.
- Documentation, R., 2016. R kernel density estimation. <https://stat.ethz.ch/R-manual/R-devel/library/stats/html/density.html>.
- Rosin, P.L., 2001. Unimodal thresholding. *Pattern Recogn.* 34, 2083–2096.
- Spearman, C., 1904. The proof and measurement of association between two things. *Am. J. Psychol.* 15, 72–101.
- Thau, D., 2015. Google earth engine: democratizing global geospatial analysis with cloud computing. 2015 AAAS Annual Meeting (12–16 February 2015).
- Landsat 8, U.S.G.S., 2016. Landsat 8. <http://landsat.usgs.gov/landsat8.php>.
- Landsat Missions, U.S.G.S., 2016. Landsat missions: imaging the Earth since 1972. http://landsat.usgs.gov/about_mission_history.php.
- Verpoorter, C., Kutser, T., Seekell, D.A., Tranvik, L.J., 2014. A global inventory of lakes based on high-resolution satellite imagery. *Geophys. Res. Lett.* 41, 6396–6402.
- Xu, H., 2006. Modification of normalised difference water index (NDWI) to enhance open water features in remotely sensed imagery. *Int. J. Remote Sens.* 27, 3025–3033.
- Yamazaki, D., O'Loughlin, F., Trigg, M.A., Miller, Z.F., Pavelsky, T.M., Bates, P.D., 2014. Development of the global width database for large rivers. *Water Resour. Res.* 50, 3467–3480.
- Yamazaki, D., Trigg, M.A., Ikeshima, D., 2015. Development of a global 90 m water body map using multi-temporal landsat images. *Remote Sens. Environ.* 171, 337–351.