

Improving river network completion under absence of height samples using geometry-based induced terrain approach

Tsz-Yam Lau and W. Randolph Franklin

Autocarto 2012, 18 Sept 2012

ABSTRACT: We discuss how accuracy of broken river segment reconnections can be improved by incorporating river segment geometry into the induced terrain approach. The reconnection problem is important because in aerial photography, canopies and clouds cover parts of rivers. Yet a complete and hydrologically consistent river network is necessary for transportation, land use planning and flood plain control. The previously-presented induced terrain approach guarantees hydrologically consistent reconnections. However, it assumes rich reliable height samples (density = 10%) which may be unavailable under extensive obstacles or gentle surfaces. We propose exploiting river segment geometry in its terrain reconstruction process. Since the ends of the segments tend to be where a river is broken, we assign relatively small height values to locations radiated from those ends, so as to favor water flowing towards them in its subsequent river derivation process. This scheme recovers 40% of what we can correct with rich height samples.

KEYWORDS: River segments, River network completion, Height awareness, segment geometry, Hydrology

1. Introduction

This paper illustrates how river segment geometry can be combined with our previously-presented induced terrain approach to yield better fragmentary river network reconstruction when rich terrain height samples are not available. The river network completion problem is important since many cross-terrain transportation issues, such as ship routing, pollutant monitoring and flood plain control, rely on full network information to work on. The main source of hydrography network data, namely aerial surveys, fails to offer full data due to classification errors and view obstacles. This is not a problem if there are complementary ground surveys, or ample time is provided to complete the reconnections manually. Otherwise, say if the hydrography survey is for outer-space planets, or we are monitoring a large quantity of data for potentially changing hydrography in real time (for example, those in a tropical rain forest), automatic bridging of those broken segments is needed.

Figure 1, left, visualizes the river segment reconnection problem. The problem works on an $n \times n$ square grid of terrain \mathbf{R} . Each of these n^2 cells r_{ij} has a binary value indicating if it is on a river or not (1 for a river location, 0 otherwise). Our task is to find out those missing river locations to recover the original complete river network.

Previously, we presented the induced terrain approach, as shown in Figure 1, right (Lau and Franklin, 2011). Referencing given height samples and river network topology, this approach features higher accuracy over conventional approaches which treat the issue as a typical line-joining problem. We will detail this approach in Section 2.

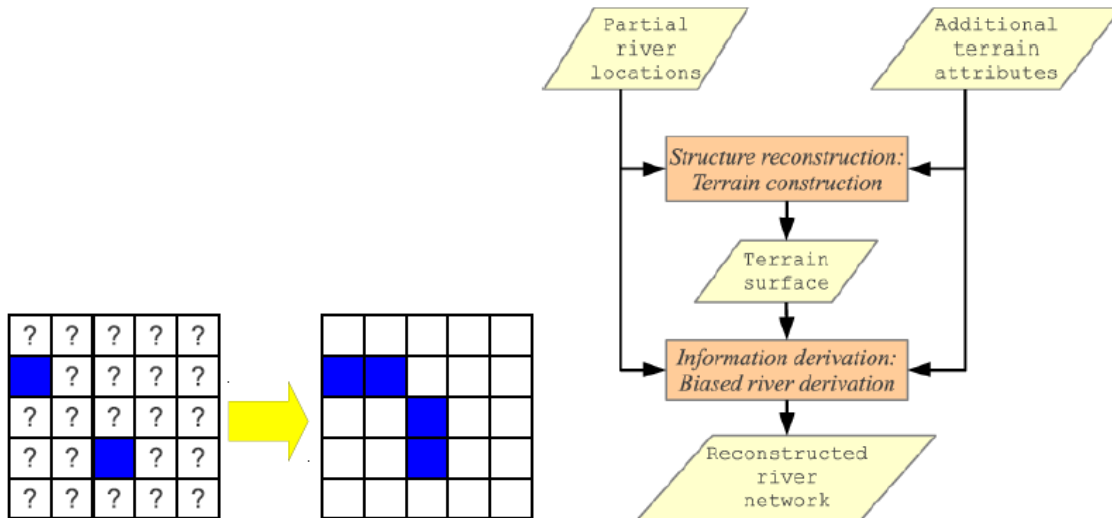


Figure 1: Left: The river segment reconnection problem. On the left is the input partial river locations (with river locations shaded). The goal is to determine if each '?' cell is a river location, so as to reconstruct the original river network shown on the right. Right: The induced terrain approach.

This paper, in contrast, works on situations in which the height samples are too sparse, or when we are working on a relatively flat surface, say in Amazonia where the slope may be just 2.7cm/km (Irion and Kalliola, 2010), in which height data are sensitive to noise. In both cases, a small noise in a few measured height values could change the topology of the reconstructed river network drastically. We purpose exploiting segment geometry, that is making use of the shape of a given river segment to improve our reconnection results. In Section 3, we will summarize a few heuristics that are often used by human when reconnecting broken segments, look into a previous attempt called axis-aligned linking which incorporates those heuristics in automated river network reconnection, and finally discuss how those heuristics could be integrated into our induced terrain approach. We will conclude our paper in Section 4.

2. Background

2.1 Induced terrain approach

Structure reconstruction. Our induced terrain approach starts with reconstructing a full height grid Z according to the given river locations and other additional terrain attributes. We emphasize the use of available height samples because the topological relationship among neighbors determines how water is routed: water flows to its lower neighbors. A hill sitting between two river segments could act as an obstacle to block water from one segment to flow to another segment.

For evenly distributed height samples, we may first reconstruct a preliminary terrain from the given partial heights with *natural neighbor interpolation* (NN) (Sibson, 1981). We then lower the heights of given river locations by a trench amount, a technique known as *stream burning* (SB) (Hutchinson, 1989), to help those river locations trap water there.

There are situations in which terrain heights may be unavailable or unreliable. Knowing the given river locations are where water is trapped, it is natural to assume they have the

lowest height in the terrain. Other locations are trying to move their water to their respective nearest river location. This is equivalent to considering each point in the river as the center of a V shape. The final terrain is the minima of all the V shapes. This leads to the baseline model, illustrated in Figure 3, which evaluates the height of a location (k, l) , $z_{k,l}$ in the terrain with the following formula.

$$z_{k,l} = \min_{(i,j):r_{i,j}=1} (z_{i,j} + \sigma\delta(p_{i,j}, p_{k,l}))$$

where $\delta(p_{i,j}, p_{k,l})$ is the distance between the two locations, probably Euclidean.

$$\delta(p_{i,j}, p_{k,l}) = \sqrt{(i - k)^2 + (j - l)^2}$$

The closer a location is from a river location, the lower its height so it is more likely to accept water and become a river location. Such a model thus favors shortest-path straight-line reconnections between break points, a major heuristic that people use to join segments manually.

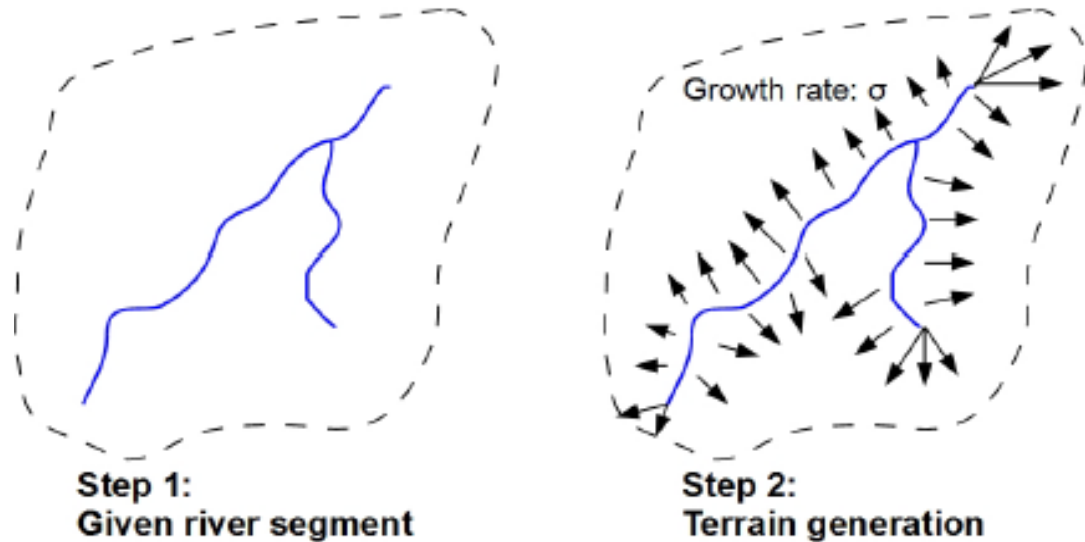


Figure 3: Baseline terrain model.

Information derivation. To reconnect the segments, we compute a river network from the above reconstructed terrain using a river derivation algorithm biased towards given river locations. Such a river derivation algorithm first computes the water flow directions of all the terrain locations based on their heights, and then finds the amount of water passing through each cell from those directions. After that, we offer each given river location an initial water amount that is exactly the critical amount (accumulation cutoff) to make sure that each location has sufficient water to be identified as river locations. Meanwhile, all other locations are allocated zero initial water. They have to receive water from the given river locations for river flow. If at the end we need to trim the river network to one-cell wide, we protect the given river locations from being removed. This ensures that the resulting river network passes through, and hence reconnects, the entire given river locations.

The use of the above river derivation algorithm helps enforce hydrological consistency in the reconnected river network. Hydrological consistency refers to the set of global constraints we expect on the complete river network. Similar to the completion of road networks, these global restrictions help eliminate invalid reconnections and thus give way to the correct ones (Steger et. al., 1997). One common constraint governs how water is routed away from sinks. A class of river derivation algorithms, such as `r.watershed` (Ehlschlaeger, 2008) available in GRASS GIS, assign flow directions such that water flows “uphill” and escapes the sink by following a least-cost path (Morris and Heerdegen, 1988}. Some others, such as *Terraflow* (Arge et. al, 2002), choose to fill all sinks by flooding so that every cell can then be assigned flow directions without routing water uphill (Jenson and Domingue, 1988). Another common constraint determines whether the network should consist of a number of tree branching structures. For example, the single-flow direction (SFD) version of `r.watershed` guarantees every location is assigned a single direction to route its water to a terrain edge or a sink in a loop-free manner. If one wishes to have a river network with loop, he may opt for the multiple-flow direction (MFD) version of `r.watershed`, which allows distribution of water from a river location to two or more neighbors and hence river branching. It is still an active topic to develop faster river derivation schemes that satisfy different sets of hydrological consistency constraints (Metz et. al., 2011; Magalhaes et. al., 2012).

2.2 Test datasets

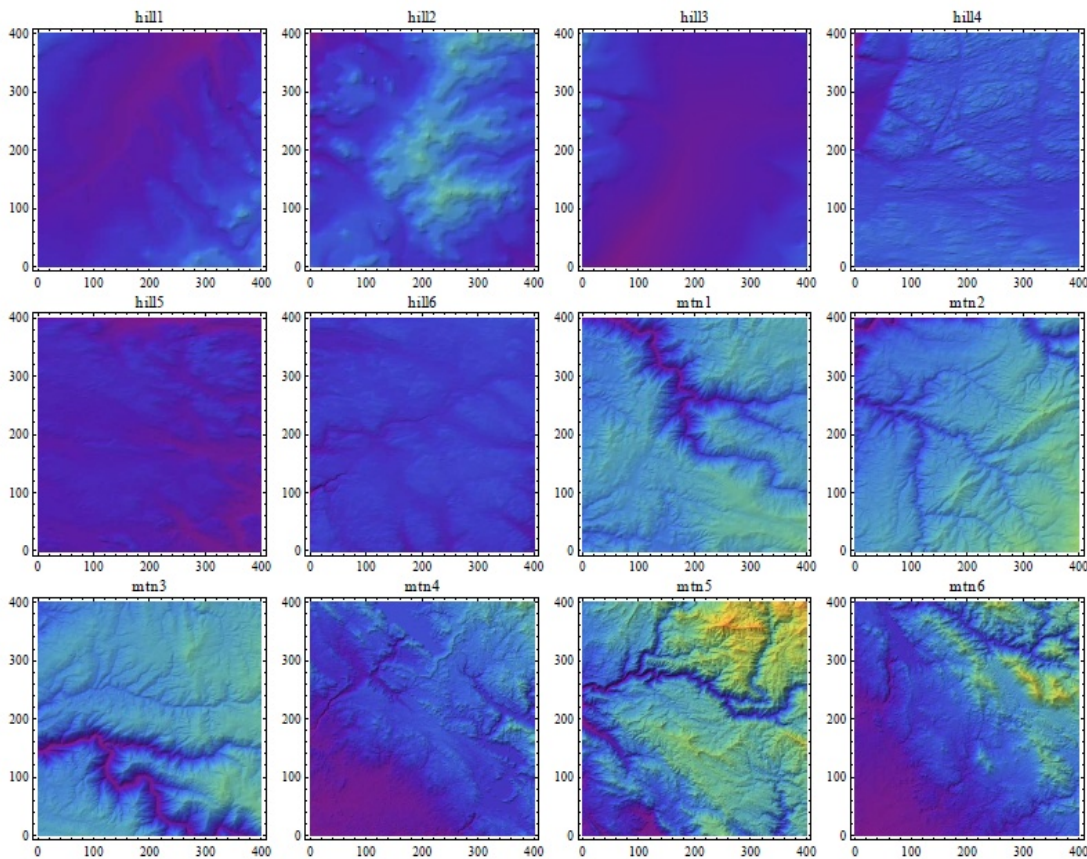


Figure 2: Test DEMs.

Table 1: Test DEM information.

DEM	Type	Name	Range	Elevation in meters		
				Mean	Std. Dev.	Range
<i>hill1</i>	SRTM1	W111N31	401:800,1:400	1251	79	1105:1610
<i>hill2</i>	SRTM1	W111N31	401:800,401:800	1548	134	1198:1943
<i>hill3</i>	SRTM1	W111N31	401:800,801:1200	1309	59	1199:1699
<i>hill4</i>	SRTM3	W60N52	1:400,1:400	441	62	140:583
<i>hill5</i>	SRTM3	W60N52	401:800,401:800	486	38	368:598
<i>hill6</i>	SRTM3	W60N52	801:1200,801:1200	447	32	229:536
<i>mtn1</i>	SRTM1	W121N38	1201:1600,1201:1600	712	146	219:1040
<i>mtn2</i>	SRTM1	W121N38	2801:3200,801:1200	847	152	330:1283
<i>mtn3</i>	SRTM1	W121N38	3201:3600,401:800	723	161	233:1021
<i>mtn4</i>	SRTM3	W121N37	1:400,401:800	294	143	45:883
<i>mtn5</i>	SRTM3	W121N37	1:400,801:1200	857	281	239:1767
<i>mtn6</i>	SRTM3	W121N37	401:800,801:1200	415	234	72:1297

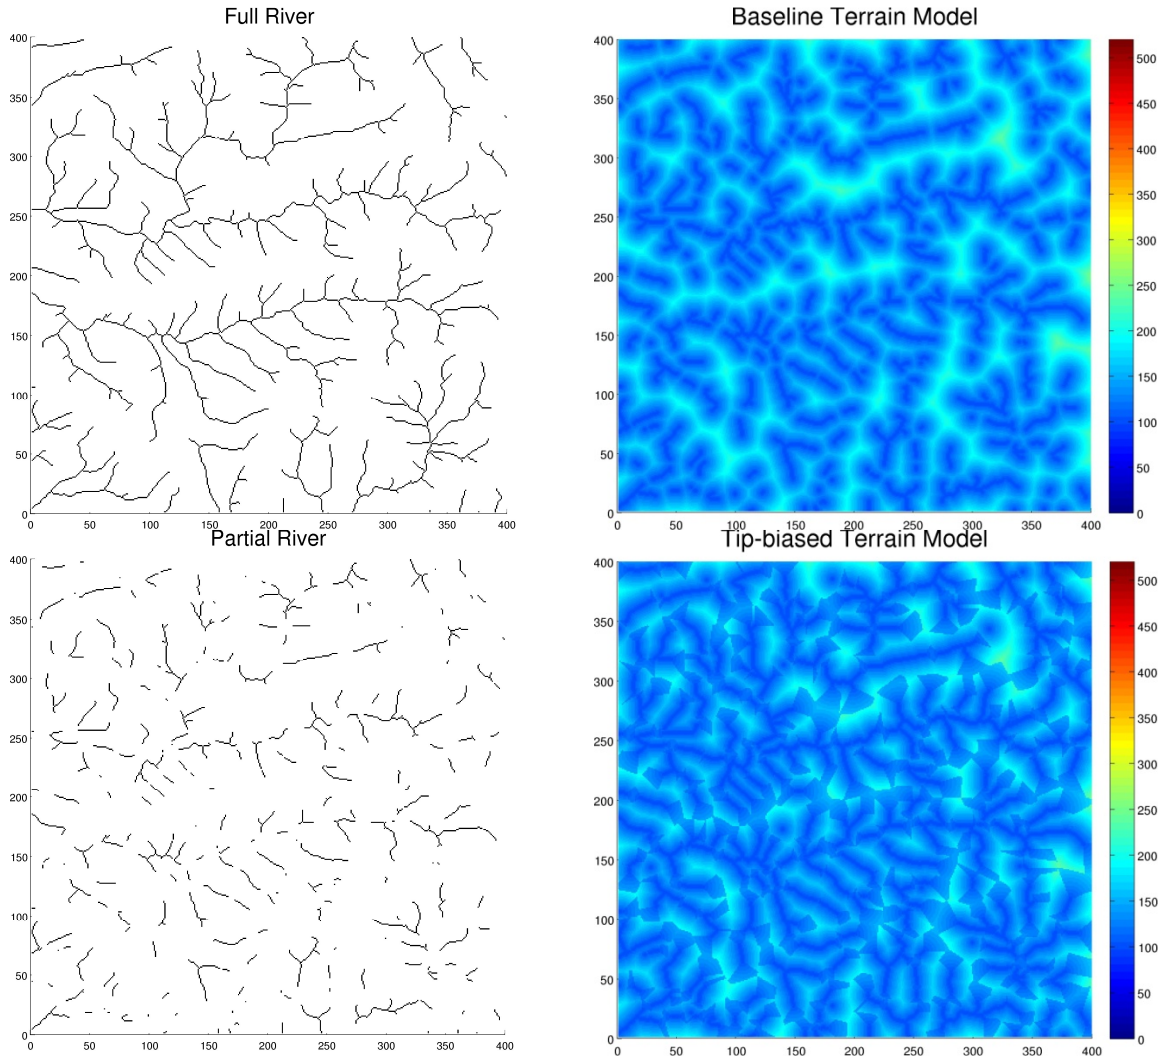


Figure 4: *mtn3*. Full river network (top left). Partial river (bottom left). Terrain reconstruction based on given river segments. Baseline model (top right). Tip-biased model (bottom right).

Our test datasets consist of the twelve 400×400 DEMs shown in Figure 2. Those DEMs are extracted from two SRTM1 cells and two SRTM3 cells, as described in Table 1. They cover a variety of mean heights, standard deviations and elevation ranges. We derive the respective full river networks with complete elevation data since we need the accurate ground-truth river networks for comparison with the reconnection results. Real partial observations rely on humans to complete the missing parts, which may mean errors.

To obtain the ground-truth drainage and river network, we run `r.watershed` with accumulation cutoff threshold = 200 and initial water amount at each location = 1 over these twelve DEMs. We pick `r.watershed` due to its readiness to accept different initial water amounts to the cells, which is necessary for the biased river derivation mentioned above. We obtain the eight-connected river networks. We sample for observed river locations as follows: first we divide the whole grid into 20×20 subgrids. In each subgrid, we randomly pick a point and mask an area of 12×12 around it. This practice mimics the occlusions in real aerial photos. Figure 4, bottom left, shows the resulting partial river locations of the *mtn3* dataset, while top right shows the terrain generated with the baseline terrain model.

2.3 Evaluation criterion

We adopt an evaluation criterion called *correct adjacent segment reconnection*. It calculates the proportion of river segments which are connected back to their respective adjacent downstream segment.

As an illustration, we compare the reconnection results with and without partial height data. To emulate additional partial height input in this case, we hide about 90% of the height values in the ground truth elevation grid. The trench amount used in the NN-SB river segment reconnection scheme is arbitrarily set to 30 because the precise value above some floor is unimportant (Callow et. al., 2007). For the subsequent biased river derivation scheme, we use the same set of parameters as when we derive the theoretical river networks. According to Table 4, the mean correct adjacent segment reconnection with our dataset reaches 85%. In contrast, the results are poorer when such partial height data are not available, just 73%. There is a 14% margin from NN-SB and the difference is statistically significant.

3. Segment geometry

By exploiting segment geometry in the river segment reconnection process, we are making use of the shape of the given river segments to improve our guess on the extensions that reconnect the segment with the others.

3.1 Related human heuristics

We are aware of the following three strategies when humans join line segments:

- We tend to extend the segments from the tips. Indeed, a tip suggests that a river segment may have been blocked from view starting at there.

- We tend to join two segments which are facing each other. Humans consider that mutually facing segments are likely to be connected behind the view obstacle.
- We tend to replicate the straightness behavior observed in the given segment to its extension. In particular, if we see a given segment to be straight, our extrapolation of the segment tends to be straight as well. Otherwise, we will allow more curvature in its extension. In fact, the shape of a river tends to be locally similar due to similarity in their water speed, soil erodibility, vegetation, etc. This is just another application of the first law in geography (Tobler, 1970).

A few morphological methods do pay attention to those observations. A notable example is axis-oriented linking as shown in Figure 5 (Zhang, 2000).

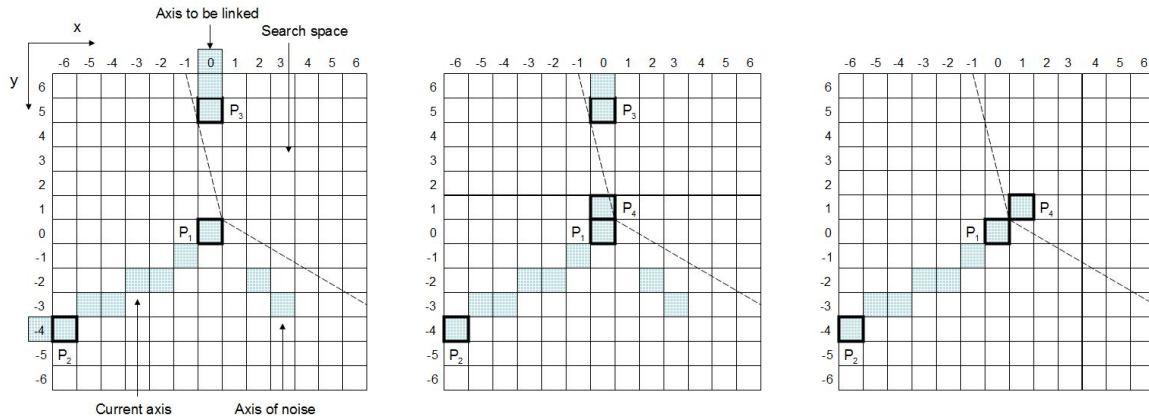


Figure 5: Axis-oriented linking. Operation window and the search space (left). A starting pixel of another segment is found in the search space (middle). No starting pixel of another segment is found in the search space (right).

Strategy 1 is realized in how the algorithm determines from where it grows a given segment: Its moves a variable operation window across all pixels of \mathbf{R} repeatedly, and extend a segment (axis) n pixels ($n \geq 1$) at a time (step length) when the algorithm sweeps its tip.

Strategy 2 and Strategy 3 are used to decide the extension direction: When the center of the operation window falls on a starting pixel of a segment P_1 , the algorithm first finds the nearest starting pixel of another segment P_3 from within a search window of the segment's forward direction. If such a P_3 is found, the algorithm extends the current segment with P_4 such that P_4 lies next to P_1 while it is closest to P_3 . Otherwise, the pixel P_2 which lies on the same segment as P_1 and is most distant from P_1 is searched for in the window. The slope between P_1 and P_2 is computed and the segment is extended according to the slope. As a result, the segment tends to grow towards another nearby segment if there is any. Otherwise, it grows according to its straightness behavior.

3.2 Improving baseline terrain modeling

To utilize the above three human heuristics mentioned in our induced terrain approach, we modify the baseline terrain model to the one shown in Figure 6: First we estimate the forward direction of each segment tip. Then we compute a binary matrix \mathbf{I} which

indicates the *privileged areas* which are radiated from the forward direction of the segment tip by $\pm\theta$. We apply a slope σ' which is smaller than σ in those regions.

$$z_{k,l} = \begin{cases} \min_{(i,j):r_{i,j}=1} (z_{i,j} + \sigma\delta(p_{i,j}, p_{k,l})) & \text{if } I_{k,l} = 1 \\ \min_{(i,j):r_{i,j}=1} (z_{i,j} + \sigma'\delta(p_{i,j}, p_{k,l})) & \text{otherwise} \end{cases}$$

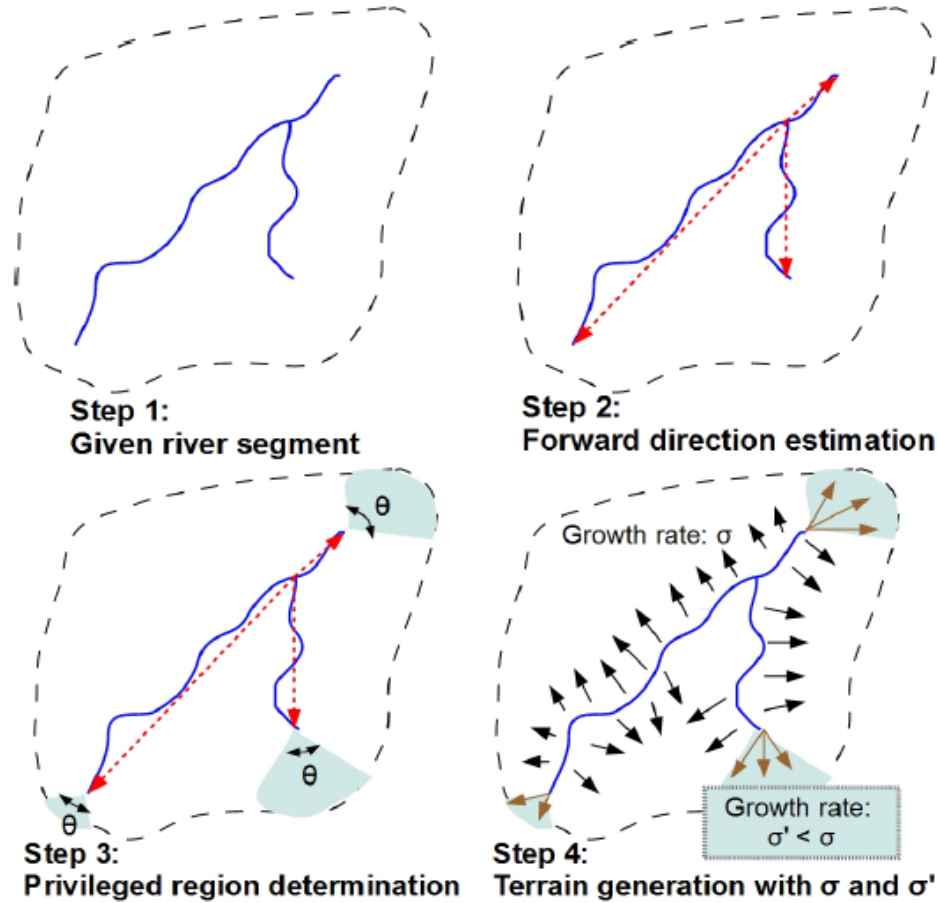


Figure 6: Tip-biased terrain model.

θ tells how much bending we are going to accept for privileged connections of mutually facing segments. A smaller θ tends to make the process selective, favoring connections of mutually facing segments. This is similar to axis-aligned linking in which the current tip is asked to look for another segment to connect in a forward search space only. An associated implicit parameter is the distance along the given segment that we use to estimate the tip forward direction. In axis-aligned linking, the forward direction depends on how far the current segment extends from the tip. If the segment does not go beyond the search window, the slope between the two ends becomes the forward direction. Otherwise, we use the tip point and the further point within the search window for slope computation. In other words, the search window size w plays a key role in determining the distance. Our scheme uses a slightly different approach to pick the distance: Starting from the tip point, we traverse along the segment until we reach the end of the segment (at which no more traversal is possible), or when we encounter a branching point (in

which case the river flow of subsequent pixels may be affected by another stream). The vector that goes from the final traverse position to the tip location then forms the forward direction.

σ' determines to what extent we favor height growing according to given segment's straightness over proximity to river locations. A small σ' relative to σ tends to favor extensions according to given segment's straightness and hence connect facing segments even if they are far apart. This is similar to the step length parameter n in axis-aligned linking. The larger is the value of n , the longer a segment is grown in each step.

The *mtn3* terrain generated with such a tip-biased model is shown in Figure 4, bottom right. As you can see, the heights of those privileged areas are smaller than those of the same distance from the nearest river location outside those regions.

3.3 Parameters setting

Table 2: Correct adjacent segment reconnection rates with different θ and σ' values.

DEM	$\sigma' = 0.5\sigma$				$\theta = \frac{\pi}{4}$		
	$\theta = \frac{\pi}{8}$	$\theta = \frac{\pi}{4}$	$\theta = \frac{3\pi}{4}$	$\theta = \frac{\pi}{2}$	$\sigma' = 0.25\sigma$	$\sigma' = 0.50\sigma$	$\sigma' = 0.75\sigma$
<i>hill1</i>	70.81%	70.33%	64.11%	64.11%	72.72%	70.33%	67.46%
<i>hill2</i>	79.65%	81.39%	77.92%	76.19%	82.68%	81.39%	79.22%
<i>hill3</i>	62.87%	63.24%	56.62%	56.62%	68.01%	63.24%	58.09%
<i>hill4</i>	75.91%	73.36%	70.07%	67.15%	70.80%	73.36%	70.44%
<i>hill5</i>	81.38%	78.54%	77.73%	74.49%	79.76%	78.54%	77.73%
<i>hill6</i>	79.30%	81.25%	77.34%	73.83%	79.68%	81.25%	81.64%
<i>mtn1</i>	81.25%	77.08%	77.92%	76.67%	74.58%	77.08%	80.00%
<i>mtn2</i>	75.66%	82.30%	79.65%	76.55%	75.66%	82.30%	79.20%
<i>mtn3</i>	77.68%	78.57%	77.68%	72.77%	76.34%	78.57%	77.23%
<i>mtn4</i>	72.79%	76.19%	77.49%	72.11%	74.49%	76.19%	72.45%
<i>mtn5</i>	76.83%	79.15%	75.68%	69.11%	75.29%	79.15%	78.38%
<i>mtn6</i>	79.12%	79.92%	78.31%	77.51%	79.12%	79.92%	78.31%

Paired diff. (θ)	$\{\frac{\pi}{8}\}-\{\frac{\pi}{4}\}$	$\{\frac{\pi}{8}\}-\{\frac{\pi}{2}\}$	$\{\frac{\pi}{4}\}-\{\frac{\pi}{2}\}$	$\{\frac{3\pi}{4}\}-\{\frac{\pi}{2}\}$	$\{\frac{3\pi}{4}\}-\{\frac{\pi}{4}\}$	$\{\frac{3\pi}{4}\}-\{\frac{\pi}{2}\}$
Average	-0.67%	4.68%	5.35%	2.54%	-2.82%	-2.14%
95% C.I. min	-1.41%	3.94%	4.74%	2.05%	-3.36%	-2.95%
95% C.I. max	0.07%	5.42%	5.96%	3.02%	-2.27%	-1.34%

Paired diff. (σ')	$\{0.25\sigma\}-\{0.50\sigma\}$	$\{0.25\sigma\}-\{0.75\sigma\}$	$\{0.50\sigma\}-\{0.75\sigma\}$
Average	-1.02%	0.75%	1.76%
95% C.I. min	-1.78%	-0.31%	1.24%
95% C.I. max	-0.25%	1.80%	2.29%

In an attempt to find a suitable θ for the average case, we first fix σ' to 0.50σ and set θ at four different levels: $\pi/8$, $\pi/4$, $3\pi/4$ and $\pi/2$ for all the segment tips in all the terrains. Table 2, left half, displays the respective segment reconnection results. Results with $3\pi/4$ and $\pi/2$ are statistically significantly less accurate than those with $\pi/4$ and $\pi/8$.

Then we fix θ to $\pi/4$ and have σ' set at three different levels: 0.25 , 0.50 and 0.75 of σ for all the segment tips in all the terrains. Setting σ' to 0.50σ gives statistically significantly better results than the other settings, as shown in Table 2, right half.

3.4 Evaluation

We compare the tip-biased model results using the best parameter setting ($\sigma' = 0.50\sigma$, $\theta = \pi/4$) with those by the baseline terrain model and NN-SB. As shown in Table 3, NN-SB offers the highest correct reconstruction rate in almost all the DEMs under test. This is due to the use of actual partial height data. When height data are not available, our tip-biased terrain models perform 5.20% better on average than the unbiased counterparts, which is 40% of the improvement achieved with partial height data. This supports the use of our tip-biased model over the baseline model. The few reconnections that are made wrong due to the tip-biased model are generally outnumbered by the corrected counterparts.

Table 4: Correct adjacent segment reconnection rates with baseline, tip-biased and NN-SB schemes.

DEM	Baseline	Tip-biased	Diff. from Baseline	NN-SB	Diff. from Baseline
<i>hill1</i>	65.07%	70.33%	5.26%	83.98%	18.91%
<i>hill2</i>	76.62%	81.39%	4.77%	83.98%	7.36%
<i>hill3</i>	56.25%	63.24%	6.99%	58.45%	2.20%
<i>hill4</i>	67.52%	73.36%	5.84%	78.10%	10.58%
<i>hill5</i>	72.87%	78.54%	5.67%	87.04%	14.17%
<i>hill6</i>	75.00%	81.25%	6.25%	85.16%	10.16%
<i>mtn1</i>	77.92%	77.08%	-0.84%	91.25%	13.33%
<i>mtn2</i>	75.22%	82.30%	7.08%	95.13%	19.91%
<i>mtn3</i>	72.32%	78.57%	6.25%	91.96%	19.64%
<i>mtn4</i>	70.41%	76.19%	5.78%	79.59%	9.18%
<i>mtn5</i>	72.97%	79.15%	6.18%	93.44%	20.47%
<i>mtn6</i>	76.71%	79.92%	3.21%	92.77%	16.06%
Average			5.20%		13.50%
95% C.I. min			3.98%		12.05%
95% C.I. max			6.42%		14.94%

4. Conclusion

We have discussed how segment geometry may be used instead of the unreliable or unavailable partial height data in the induced terrain approach of completing fragmentary river networks. Instead of growing height uniformly to all directions of the river segments, we reduce the growth rate of location radiated from the segment tips. Our prototype using the same set of parameters for all segments improves reconnection results by around 5 percentage points on average, while with partial height we achieve around 14 percentage points.

Note that for the segment geometry extension, first we set the same θ and the σ' for all tips in all terrains in an attempt to find a universal setting. However, we realize that different local regions achieve largest improvement at different settings. It may be interesting to investigate how these parameters can be adjusted for individual tips. It is also beneficial to test with many other datasets, especially the real ones. Second, we are not using any height information in the process. It will be useful to find out a threshold partial height data density or terrain steepness when this is applicable. It will also be exciting to develop a scheme that combines segment geometry and partial height data, because we are potentially complementing one source of information with the other. For both extensions, although the modified terrains achieve better river segment reconnection results, they are not in good natural shape. Work may be done to improve this situation.

Ultimately, we would like to port the same solution framework to complete some other 2D networks like road networks, and extend it to solve 3D network completion problems like fragmentary dendrite networks. Preliminary work will include studying how similar the river reconnection issue is with those problems, and figuring out the appropriate clues in those new domains.

References

- Arge L., Chase J.S., Halpin P., Toma L., Vitter J. S., Urban D., and Wickremesinghe R. (2003) Efficient flow computation on massive grid terrain datasets. *Geoinformatica*, 7, 4, pp.283-313.
- Arge L., Mitasova H. and Toma L. (2002) r.terraflow.
http://www.cs.duke.edu/geo*/terraflow/r.terraflow.html Last Visit 2/18/2011.
- Asante K. and Maidment D. (1999) Creating a river network from the arcs in the digital chart of the world.
<http://www.ce.utexas.edu/prof/maidment/grad/asante/dcw/rivernet.htm> Last Visit 6/23/2010.
- Callow J. N., Niel K. P. V., and Boggs G.S. (2007) How does modifying a DEM to reflect known hydrology affect subsequent terrain analysis? *Journal of Hydrology*, 332, pp.30-39.
- Ehlschlaeger C. (2008) GRASS GIS: r.watershed.
http://grass.itc.it/gdp/html_grass63/r.watershed.html Last Visit 6/23/2010.
- Hough P. V. C. (1962) Method and Means for Recognizing Complex Patterns. US Patent 3069654.
- Hutchinson M. F. (1989) A new procedure for gridding elevation and stream line data with automatic removal of spurious pits. *Journal of Hydrology*, 106, pp.211-232.
- Irion G. and Kalliola R. (2010). Long-term landscape development processes in amazonia. In Hoorn C. and Wesselingh F., editors, *Amazonia, Landscape and Species Evolution: A Look into the Past*. 1ed. Wiley-Blackwell, pp.185-200.
- Jenson S. K. and Domingue J. O. (1988). Extracting topographic structure from digital elevation data for geographic information system analysis. *Photogrammetric Engineering and Remote Sensing*, 54, 11.
- Lau T.-Y. and Franklin W. R. (2011). Completing fragmentary river networks via induced terrain. *Cartography and Geographic Information Science*, 38, 2, pp.162-174.
- Magalhaes S. V. G., Andrade M. V. A., Franklin W. R., and Pena G. C. (2012) A new method for computing the drainage network based on raising the level of an ocean

- surrounding the terrain. In *AGILE'2012 15th AGILE international conference on geographic information science*, Avignon, 24-27 April 2012.
- Metz M., Mitasova H., and Harmon R. S. (2011) Efficient extraction of drainage networks from massive, radar-based elevation models with least cost path search. *Hydrology and Earth System Sciences*, 15, 2, pp.667-678.
- Morris D. and Heerdegen R.(1988). Automatically derived catchment boundary and channel networks and their hydrological applications. *Geomorphology*, 1, pp.131-141.
- Noble J. A (1996) The effect of morphological filters on texture boundary localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18, pp.554-561.
- Sibson R. (1981) A brief description of natural neighbor interpolation. In Barnett V., editor, *Interpreting Multivariate Data*. John Wiley & Sons, New York, chapter 2, pages 21-36.
- Steger C., Mayer H., and Radig B. (1997). The role of grouping for road extraction. In *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, *Birkh auser Verlag Basel*, pp.245-256. Verlag.
- Tobler W. (1970) A computer movie simulating urban growth in the Detroit region. *Economic Geography*, 46, 2, pp.234-240.
- Zhang Y. (2000) A Method for Continuous Extraction of Multispectrally Classified Urban Rivers. *Photogrammetric Engineering and Remote Sensing*, 66, 8, pp.991-1000.

Tsz-Yam Lau, PhD candidate, Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180. Email <laut@cs.rpi.edu>

W. Randolph Franklin, Professor, Department of Electrical, Computer and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180. Email mail@wrfranklin.org

This material is based upon work supported by the National Science Foundation under Grants No. 0835762 and 1117277.