Bridge Detection By Road Detection

Jeff Kaufman
cbr@sccs.swarthmore.edu

1 Introduction

It is useful to be able to determine water flow patterns over terrain. The raw data for this task is usually collected via airborne laser range finding, or LIDAR. This yields a point cloud representing the uppermost surface of the terrain. This cloud is interpolated onto a grid where each grid cell represents a square portion of the earth's surface and it's value is the average hight of that portion. Water flow patterns can then be calculated by looking at the direction of greatest descent from each cell.

There is a problem, however, with local minima: cells from which there is no direction of descent. A common cause of this is bridges. The area over which a bridge passes shows up in the digital elevation as being of a height greater than the surrounding terrain, but for the purposes of water flow it is not there. A water flow simulation model will treat those grid cells as indicating a barrier, then, when there is in fact no impediment to water flow. The goal of this project is to identify bridges to aid in the accurate determination of water flow patterns

2 Past Attempts and Related Work

A common method for dealing with local minima is flooding. In this all grid cells, excluding those at the edges of the grid, where there is no lower adjacent cell are raised up to the height of their lowest neighbor. Repeating this will eventually rid the map of local minima. Identical results to this naive method can be achieved efficiently with a plane sweep algorithm using topological persistence. (Edelsbrunner et al., 2000) Unfortunately, as Soille et. al. (Soille et al., 2003) recognize, this loses information. Large flat areas, a common result of flooding, retain no information about their original low points and do not provide information actual flow patterns. One can still determine a possible flow pattern through it, but that pattern may be totally different from the true one.

Soille et. al. were working with a very low resolution (250m) grid elevation model to determine the water flow pattern for Europe and the problems they ran into are somewhat different from those encountered with higher resolution data. When they encountered a local minimum, specifically, it was usually because some small stream or channel went undetected with the coarse sampling. Their replacement of flooding, carving paths from local minima to the nearest lower area, makes sense for dealing with the missing channels but can of course get things wrong. As the local minima in higher resolution data are much more likely to be products of human terrain manipulation, generally in the creation of roads, a system that tags and removes bridges ought to come closer to true water flow paths than either flooding or carving.

For last year's senior conference, Manfredi and Pshenichkin (Manfredi and Pshenichkin, 2006) used a series of classifiers to tag bridges. They had a series of simple criteria that a window had to match to be tagged a bridge. They were able to detect many of the larger bridges but missed some smaller ones, as well as complex structures such as highway interchanges. Their system also had a large number of false positives, detecting vegetation and other small artifacts as bridges. They rightly point out, however, that there is not too much harm in removing them along with bridges as they are also not really there from the perspective of water flow.

One feature that they did not take advantage of is the tendency of bridges to be part of the road network. All of their classification work considered only the window that potentially contained the bridge. There has been some work on detection of of roads from LIDAR data (Clode et al., 2005), and while the final detected roads may not be completely accurate, for this task we don't need perfection. Instead we just need a rough idea of how likely a region is to be part of the road network, which can then be input to the bridge detection system.

3 Bridge Detection via Road Detection

For this project I locate bridges in two stages. In the first I determine an approximate map of the road network, a map that should generally be best in areas where the roads are in high relief. These areas correspond well to those where bridges are likely, so it should be a well suited map for the task. Second I identify local minima that are near the computed roads in order to tag road sections as bridges. Input consists of a digital elevation model in the form of a grid of floats indicating hight. Output consists of five similar models with floats indicating likelihood of being a bridge, with calibration required for the particular data set.

3.1 Road Detection

Roads are places in the terrain that are flat. Any flat area could be part of a road. Areas that are linearly flat, however, are much more likely to be roads. These would be areas where lines in one direction are flat while in other directions are not. Finally, roads tend not to bend sharply, so if there is a road in a direction we treat grid cells in that direction as being more likely to be a road.

This yields four indicators of bridge-likeness. All are computed on a series of cells representing a potential road. For every cell c in the grid we calculate 32 potential roads of a configurable length running through that point. We then find which of those series of cells has the lowest average change in steepness and call that the 'best road' centered on that cell. We also find the set of cells representing a line perpendicular to the best road and call that the 'perpendicular'. Each indicator acts on one of these two roads and yields a value attributed to c.

1. Maximum gradient. For the best road, the likelihood of it being actually a road is inversely proportional to the greatest difference between adjacent cells in the road.

2. Average gradient. Like the previous, except the average absolute difference is calculated instead of the maximum one.

3. Maximum gradient of perpendicular. The likelihood of the best road being a road instead of just a cornfield is indicated by the unroadlikeness of the perpendicular. This is calculated as for the maximum gradient.

4. Standard deviation of gradient. Even when not level, roads tend to be flat. That is, while they might sometimes have high gradients the (root mean square) standard deviation should be low.

3.2 Local Minima

A maximally simple algorithm for determination of local minima turns out to be quite effective as the data is not very noisy. For every grid cell, if no neighbor is smaller, then that cell is a local minimum. With worse data we might have a large number of these places and a small amount of flooding might be worth while. After flooding an amount small enough not to overflow bridges we have not lost much flow information and now should generally have local minima just in places where there are bridges.

4 Results

These four indicators were tested on two different examples of roads. Figures 1 and 2 show the LIDAR-derived input grids. Figures 3 and 4 show the maximum gradient indicator. Figures 5 and 6 show the average gradient indicator. Figures 7 and 8 show the maximum gradient of perpendicular indicator. Figures 9 and 10 show the standard deviation of gradient indicator.

All four indicators appear to capture an element of bridge detection. One important aspect of these indicators is that they most strongly label cells as indicating bridges when those cells are in places where they would be incorrectly

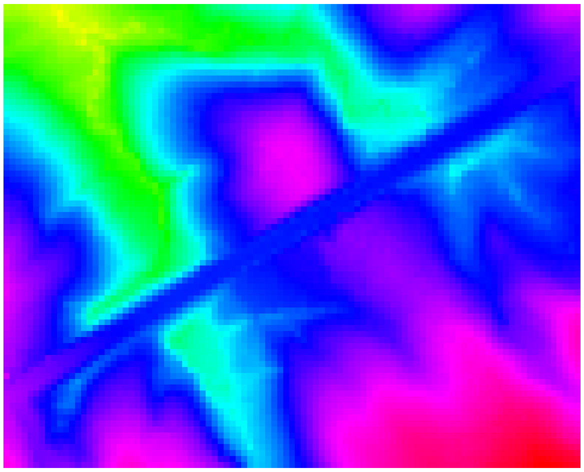

Figure 1: The input digital elevation model for the Bridge test

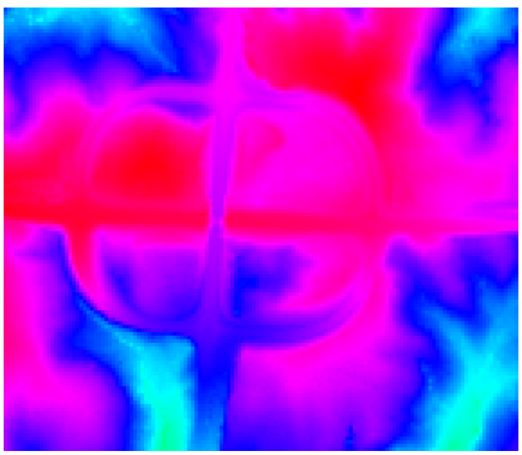

Figure 2: The input digital elevation model for the Interchange test

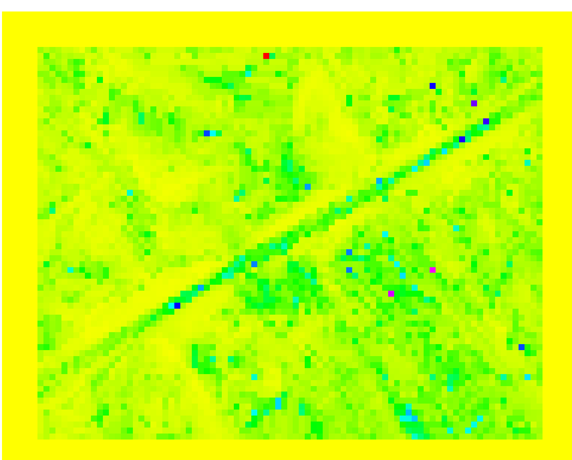

Figure 3: The maximum gradient indicator on the Bridge

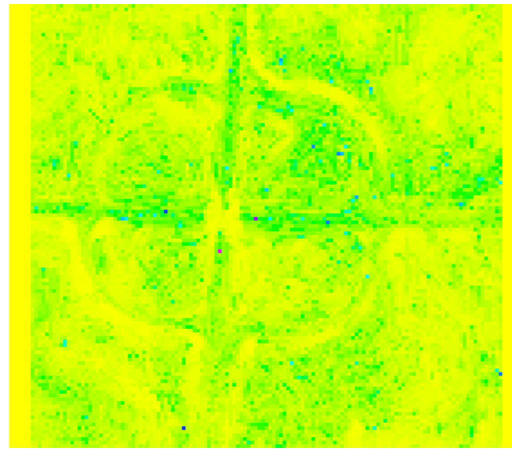

Figure 4: The maximum gradient indicator on the Interchange

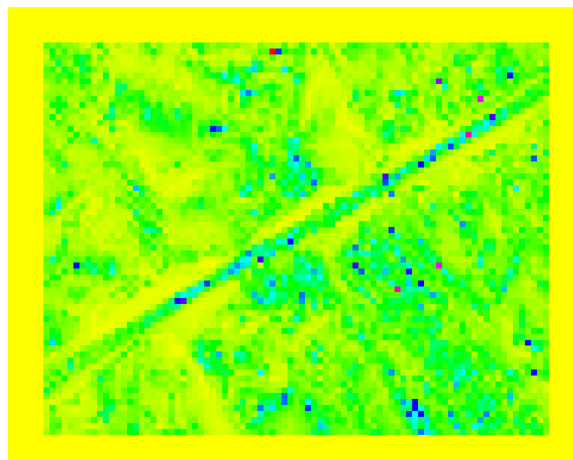

Figure 5: The average gradient indicator on the Bridge

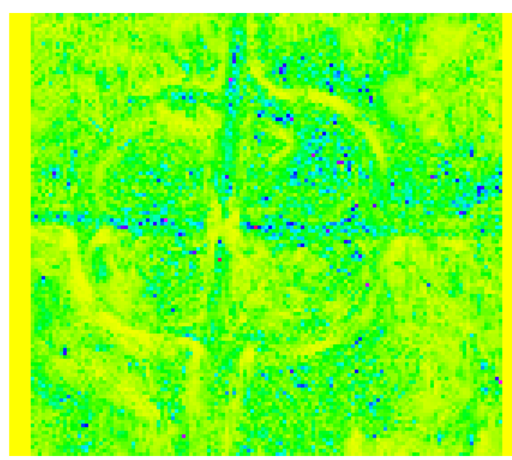

Figure 6: The average gradient indicator on the Interchange

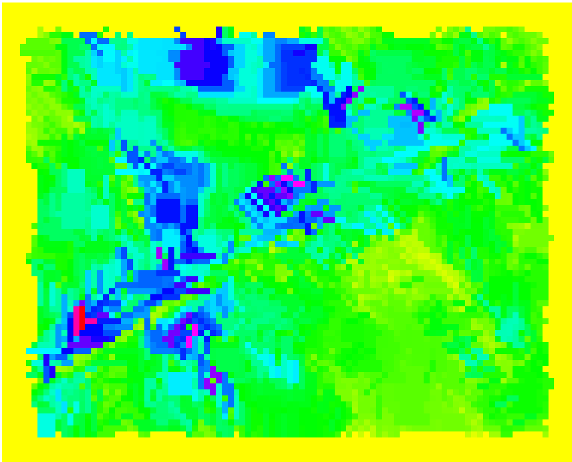

Figure 7: The maximum perpendicular of gradient indicator on the Bridge

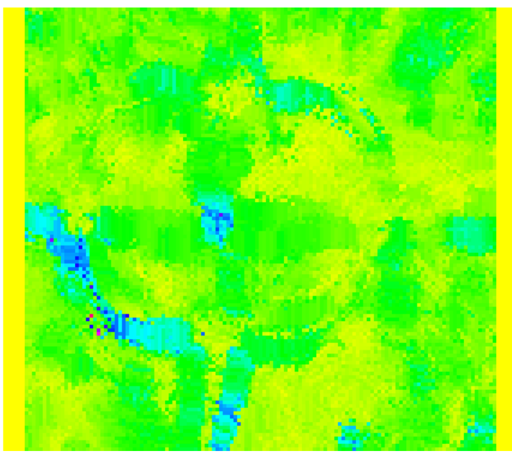

Figure 8: The maximum perpendicular of gradient indicator on the Interchange

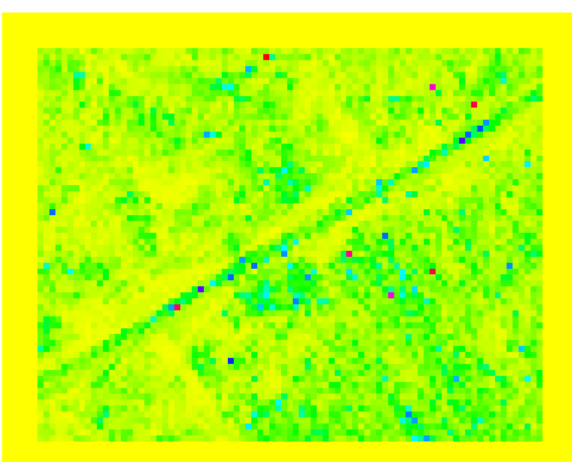

Figure 9: The standard deviation of gradient on the Bridge

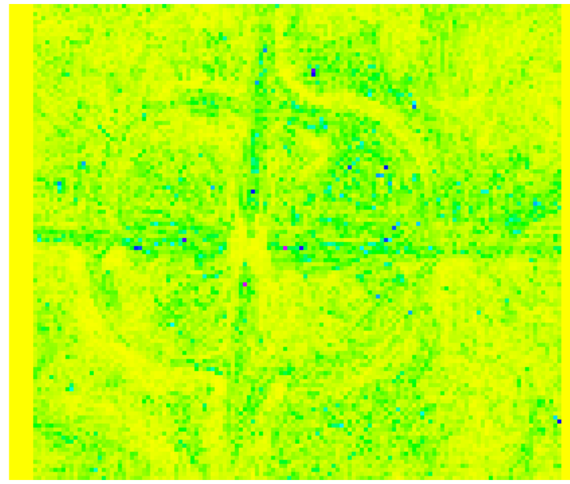

Figure 10: The standard deviation of gradient on the Interchange

impeding water flow. Actual use of these indicators on real data would require calibrating them. This would require hand tagging a small number of bridge examples and computationally determining the combination of these indicators that best fits that data.

5 Conclusions

In this paper we have shown that several relatively simple functions analyzing a digital elevation model can produce good indicators for classification of cells as to their probability of being a road. Further work would include a large scale test with calibration on a large digital elevation model. Implementation of a second pass that combined these local bridge likelihood estimates into a road network could also improve accuracy.

References

S. Clode, F. Rottensteiner, and P. Kootsookos. 2005. Improving city model determination by using road detection from lidar data. *IAPRSSIS*, XXXVI(3/W24):159–164.

H. Edelsbrunner, D. Letscher, and A. Zomorodian. 2000. Topological persistence and simplification. In *41st IEEE Symposium on Foundations Computer Science*, pages 454–463.

A. Manfredi and A. Pshenichkin. 2006. Bridge detection from elevation data using a classifier cascade. http://web.cs.swarthmore.edu/~adanner/cs97/f06/papers/bridge.pdf.

P. Soille, J. Vogt, and R. Colombo. 2003. Carving and adaptive drainage enforcement of grid digital elevation models. *Water Resources Research*, 39(12).