Geometric Overpass Extraction from Vector Road Data and DSMs

Joshua Schpok
Google Inc.
schpok@google.com

ABSTRACT

We present a method to extract elevated road structures, typically overpassing other roads, transit lines, and watercourses. The technique uses a *digital surface model* (DSM) and roughly aligned vector road data and outputs geometry approximating the shape and elevation of the elevated road deck. Our method is robust against noise in DSM elevations and can recover elevated roads partially obscured by trees and other overpasses. We demonstrate our method parallelized over city-wide DSMs, and formulate a confidence metric ranking the fidelity of the reconstruction.

Categories and Subject Descriptors

I.4.8 [**Image Processing and Computer Vision**]: Scene Analysis—*Object Recognition*; I.4.6 [**Image Processing and Computer Vision**]: Segmentation—*Edge and feature detection*; I.4.7 [**Image Processing and Computer Vision**]: Feature Measurement—*Feature representation*

1. INTRODUCTION

Overpass structures are an important part of realistically and accurately representing roadways. Without overpass structures in a terrain model, roadways in bare *digital terrain models* (DTMs) such as in Figure 1 appear unrealistic and do not clearly convey the absence of an intersection between the crossroads. Though these artifacts may have some artistic merit [10], they can be visually distracting, difficult to parse visually, and break an immersive experience.

The shapes of elevated road structures are also useful in the generation of DTMs, since they are often misclassified [8]. These structures serve not only as a mask for non-terrain elements, but may also assume roadways *not* classified as elevated are likely on the ground, which can serve as the local terrain elevation [1].

With the relative infrequency of overpass structures across terrain, we wish to derive results from existing data instead of collecting new overpass-specific data. Our method reconstructs overpass structures from 2-meter resolution DSMs and aligned vector road paths.

Figure 1: Realism is lost without overpass structures, and it is less clear if these roads intersect.

2. RELATED WORK

Several methods have been published to extract overpasses from geospatial data. Using LIDAR instruments mounted to terrestrial vehicles [7], overpass structures can be detected. This differs from our approach of using a DSM, which has wider and more thorough coverage than just those overpasses over drivable, accessible roadways.

Work extracting elevated roadways solely from panchromatic texture [5][4][9] successfully recovers structure over heterogeneous regions, though it fails to identify if two roads of similar constitution cross.

Topographic vector data can be given elevations by fitting them to a LIDAR-derived segmentation [6]. Our datasets lack a complete object segmentation. Though this would greatly assist in measuring roads, they are limited in availability over the breadth of terrain we wish to process. Instead, we measure the road shape as a part of extraction. While this means road geometry may be less accurate, it does permit extraction in regions where only road contours exist.

Among other properties, the geometry of elevated road structures can be reconstructed solely from LIDAR and multispectral imagery [3]. The results are tuned for urban environments, where stretches of continuous flat terrain are

ACM SIGSPATIAL GIS '11, November 1–4, 2011, Chicago, IL, USA.

extracted as roads. It is unclear how the algorithm performs in windy, hilly roads, and results demonstrate some tracking limitations without vector road data.

3. METHOD

Our method can be divided into four sequential steps, outlined in Figures 2 and 5. First, vector road data is used to measure *road spans* by finding drop-offs perpendicular to the road direction. Next, road spans are clustered by a distance metric defined in Section 3.2. These clusters are merged to form continuous overpass structures in Section 3.3. Finally, the road spans are rebuilt into a geometric model in Section 3.4.

3.1 Span Measurement

For all road segments, we sample a 1-dimensional cross section perpendicular to the road direction to find a drop-off and determine if this segment could be part of an overpass. We express our surface model as $DSM : \mathbb{R}^2 \mapsto \mathbb{R}$, mapping a two-dimensional surface coordinate to its surface elevation. In our implementation, this function bilinearly interpolates a 2-meter grid. We sample a new cross section maximally at every DSM grid cell along a road segment. This sampling rate can be reduced for efficiency at the cost of modeling robustness.

Road span quantities are illustrated in Figure 3. We define the elevation cross section originating from a road point $\mathbf{x}$ as:

$$f_+(t) = DSM(\mathbf{x} + \mathbf{n}t), \quad 0 < t < b_{max} \tag{1}$$

$$f_-(t) = DSM(\mathbf{x} - \mathbf{n}t), \quad 0 < t < b_{max} \tag{2}$$

where $\mathbf{n}$ is a unit direction perpendicular to the road, and b_{max} is a practical maximum overpass breadth. We observe height discontinuities from the moving average with the function:

$$g(t) = f(t) - \frac{1}{t}\int_0^t f(s)\,\mathrm{d}s. \tag{3}$$

With j as a user-tunable drop-off threshold, we attempt to minimize t for Equations 1 and 2 such that $-j > g(t)$, where t is bounded $0 < t < b_{max}$, and there exists no $u < t$ such that $j < g(u)$, representing a height discontinuity from the rolling average upwards. In addition, this measurement may be rejected early by testing if $\mathbf{x}$ lies within a cell classified as vegetation[1]. This avoid misinterpreting the canopies of tree-lined streets as the deck of an elevated road.

Finding a minimal t for f_+ and f_- gives us values t_+ and t_-, from which we compute approximate span breadth in grid space:

$$b = t_+ + t_- \tag{4}$$

a revised road span midpoint:

$$\mathbf{m} = \mathbf{x} + \mathbf{n}\,(t_+ - t_-) \tag{5}$$

and elevation:

$$h = \frac{1}{b}\int_{-t_-}^{t_+} DSM(\mathbf{x} + \mathbf{n}s)\,\mathrm{d}s. \tag{6}$$

[1]We classify vegetative cells as having a high *normalized difference vegetation index* (NDVI), computed from existing aerial imagery.

(a) Sample points in A along road vectors.

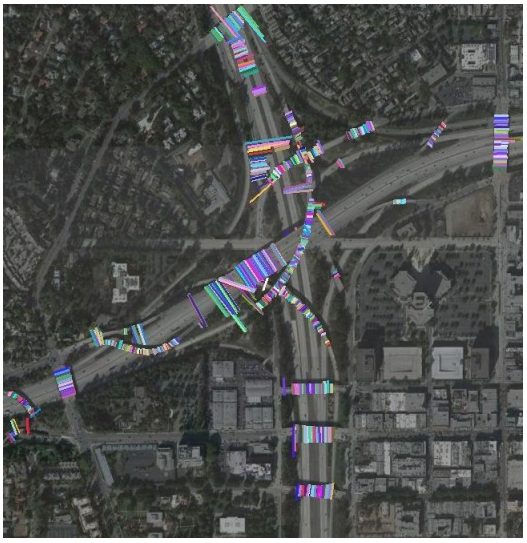

(b) Road spans with measured drop-offs M.

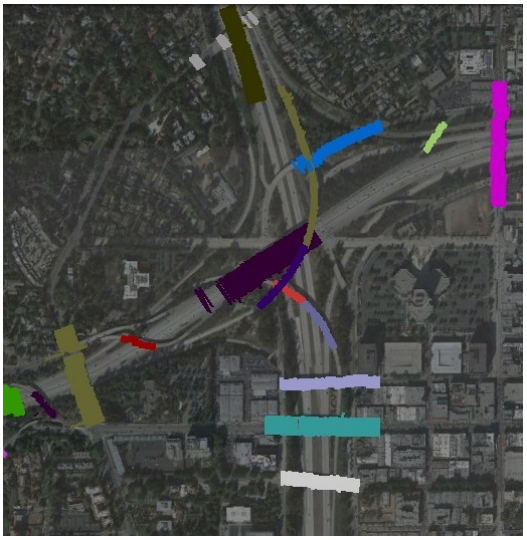

(c) Clustered and extended road spans C.

Figure 2: Stages of road span processing.

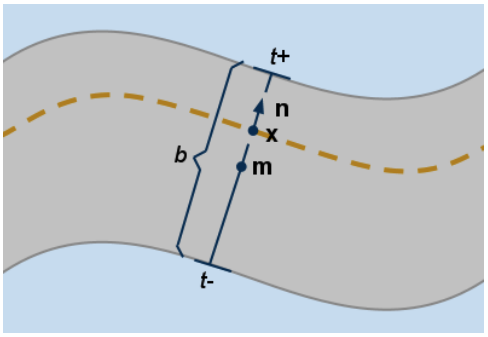

Figure 3: Road span quantities. A point x along a road vector (dashed, yellow) has a perpendicular vector n. If drop-offs (solid, gray) are found along $\mathbf{x} + \mathbf{n}t$ for t_+ and t_-, breadth b and midpoint m can be computed.

These measurements are attempted for all road spans A, constructing a subset of spans M where t_+ and t_- are successfully found.

3.2 Span Clustering

The next step is to divide the set of measured road spans M into subsets C where the spans form a continuous length. We achieve this with single-linkage clustering: Beginning with every road span in its own set, sets C_ϕ and C_θ are merged if the minimum distance between any two elements of either set is less than 1:

$$\min\{d(c_\phi, c_\theta) : c_\phi \in C_\phi, c_\theta \in C_\theta\} < 1 \tag{7}$$

Our distance metric measures difference in position, direction, and breadth from the terms of Section 3.1:

$$d_{cluster}(\phi, \theta) = \max\{d_m, d_n, d_b\} \tag{8}$$
$$d_m = q_m \left\|\mathbf{m}_\phi - \mathbf{m}_\theta\right\| \tag{9}$$
$$d_n = q_n \left|1 - \mathbf{n}_\phi \cdot \mathbf{n}_\theta\right| \tag{10}$$
$$d_b = q_b \left(b_\phi - b_\theta\right) \tag{11}$$

The q scalars adjust the sensitivity of each of these components. Ideally q_n is small enough to avoid clustering with intersecting roads heading other directions, yet large enough to follow the road as it curves. Similarly, q_b is small enough to avoid merging ramps with wide freeway decks, yet large enough to accommodate some minor width fluctuation.

We formulate the distance Equation 8 without road network connectivity to combine disjoint roadways sharing a common overpass structure. This is the case in divided bridges, as illustrated in Figure 4.

Noise, misclassified vegetation, and other elevation errors can produce erroneous road spans. Our distance metric prevents outlying spans from polluting our cluster trend, and will usually be left in sets with low cardinality. After clustering, we cull sets $|C_\xi| < \epsilon$, erasing its spans' measures and removing them from M, but still remaining as members of A. This permits these spans to still contribute to the path of the overpass in Section 3.3 without affecting its breadth, offset, and elevation measures.

Figure 4: Vector roadways are reconstructed from DSM measurements instead of the original vector data because multiple vector thoroughfares may be sharing the same structure. Here, vehicular (yellow) and pedestrian (red) traffic disjointly share the Colorado Boulevard Bridge in Pasadena, California.

3.3 Overpass Growing

After clustering measured road spans, we grow each subset C to include certain road spans from outgoing edges. From an outgoing edge $\{c_\gamma, a\}$ $c_\gamma \in C_\gamma, a \in A \setminus C_\gamma$, a road span a is added to C_γ if its direction is similar by Equation 10, and the path along the vertices C_γ to the nearest member of M is within a growth threshold, q_g. This criterion can be expressed as single-linkage clustering using a distance function of:

$$d_{grow}(c_\gamma, a) = \max\{d_m, d_g\} \tag{12}$$
$$d_g = \min\{w(C_\gamma, a, c_m) : c_m \in M \cap C_\gamma\}/q_g \tag{13}$$

where w is graph length using the Euclidean distance of source points $\mathbf{x}$ between directly-connected vertices. We consider road span a a continuation of C_γ and add it to that set when $d_{grow} < 1$.

As per single-linkage clustering, if a candidate span a belongs to another cluster C_β, sets C_γ and C_β are merged. The scalar q_g adjusts how far overpass segments will grow, not only to connect broken regions of a common overpass structure, but to extend these structures into berms. This is a useful result to connect overpasses into terrain where a clear drop-off will not be observed.

3.4 Geometric Reconstruction

Each set C now contain all the road spans we consider composing that overpass. We choose a subset of spans to construct simple geometry approximating the shape of the overpass as seen in Figure 5. Since the spans are more dense than necessary, we choose a target density, τ spans per meter. Our process is simply to traverse the road spans from one end to the other, choosing a span every $1/\tau$ units. Unfortunately, since spans can originate from a collection of disjoint roads, we lack an obvious traversal order.

To construct a traversal order, we begin a sequence with the furthest span from a randomly selected span in the set. We

Figure 5: Overpass geometry is approximated with a rectangular prism. Orientation, breadth and elevation are derived from road spans, while the road depth is defined by the user.

then eliminate all spans within $1/\tau$ units, and enqueue the next closest. We continue the process until all the spans have been eliminated or enqueued.

Recall that not all spans belong to set M, which has valid breadth and elevation measurements, and a corrected midpoint. Because constant width is an assumption of our model, the road geometry's breadth is simply the average of the breadths measured in $C \cap M$. All of the geometry's elevations are spline-fit to suppress any noise and interpolate missing elevation values in $C \setminus M$. To approximate the midpoints for spans in $C \setminus M$, we interpolate the value of $\mathbf{m} - \mathbf{x}$ from the neighboring spans $C \cap M$.

3.5 Confidence Metric

Though it is sufficient for a majority of our input, our overpass model is unrealistically limited for a variety of overpass structures. We attempt to characterize the overpasses that fit poorly in our model so that we might consider the better fits first in a quantity of output. We linearly combine several measurements to formulate an objective function:

- Standard deviation in b: An overpass with a range of breadths will not be realistically approximated as constant by our model.
- Range of orientations in $\mathbf{n}$: Our span ordering algorithm is not robust against structures that turn tightly such that endpoints are not the furthest-distanced from any random span.
- Length of overpass relative to breadth: We limit unusually long or short overpasses, weighted by breadth. Longer elevated roads will likely be wider (such as large bridges), while shorter structures will likely be narrower (such as pedestrian overpasses).

4. RESULTS

We scale our algorithm by dividing the work over many machines, each processing a small tile of a large DSM. To avoid the complexity of sharing road networks and stitching geometry between tiles, we impose a maximum overpass length which we use as an overlapping margin between DSM tiles. This guarantees overpass structures will be completely contained within a tile.

Spatially sorting the road network to overlapping DSM tiles adapts well to a MapReduce framework [2]. A collection of map workers traverse the entire road network emitting key/value pairs, where the key is a DSM tile index, and the value is a road segment. Note a road segment may be emitted multiple times with different indices where tiles overlap.

A reduce worker iterates over all the values of a single key. This means a reduce worker only needs to load a single DSM tile for its road segments, populating its local set A. Once all the road segments for a key are consumed, the reducer runs our algorithm. The results of overpass extraction over several metropolitan areas using 2000 machines are summarized in Table 1. We threshold against the metric described in Section 3.5 to filter reconstructions of sufficient quality for 2D segmentation and 3D models. These thresholds results in only a few false positives, but at the expense of as much as 50% false negatives.

The following figures demonstrate our algorithm, untextured for illustrative purposes. Figure 6 demonstrates faithful reconstruction of multiple overlapping overpasses. Though portions of overpasses could not be measured, our algorithm reconstructs the completed structures.

Figure 7 illustrates realistic extraction with vector data that does not precisely conform to the overpass structure. Here, Interstate 5's overpass over Artesia Boulevard was widened, but vector data was previously aligned to the older structure. Our clustering algorithm reconstructs the road deck by aligning structures measured directly from the DSM.

Our algorithm fails to produce representative geometry where elevated roads vary in width, or vector lines fall outside the road in the DSM. Figure 8 shows such a failure where where multiple overpass structures connect. Because the road decks vary in width where lanes bifurcate and merge, our road growth algorithm constructs separate sets at these junctions. Though our confidence metric assists in identifying these poor fits, we aspire to capture these structures in future work.

5. CONCLUSIONS AND FUTURE WORK

We present a scalable algorithm to extract and geometrically approximate elevated road structures using a basic vector road network and 2-meter DSM. Although our road model is limited, it captures a majority of elevated road structures.

A confidence score indicates when our model or supporting data is insufficient to accurately capture an overpass. This helps us adjust how conservatively we accept models depending on our application.

6. ACKNOWLEDGMENTS

The author wishes to thank the Claus Brenner, Spudde Childs, Aleksey Golovinskiy, Mike Lin, Brett Miller, and Tilman Reinhardt for technical guidance and suggestions, and Google for supporting this research.

Metro Area	Coverage	Wall Time	# Extracted	Segmentation	3D Modeling
Chicago	4,231 sqkm	4.46 min.	1,054	689 (65%)	476 (45%)
Denver	4,350 sqkm	4.26 min.	1,233	814 (66%)	505 (41%)
Houston	7,183 sqkm	6.64 min.	2,218	1,564 (70%)	1,226 (55%)
Las Vegas	1,206 sqkm	3.45 min.	496	317 (64%)	218 (44%)
Los Angeles	4,087 sqkm	4.47 min.	3,882	2,270 (58%)	1,398 (36%)
Salt Lake City	3,297 sqkm	4.74 min.	617	417 (69%)	261 (43%)
San Diego	2,810 sqkm	4.55 min.	1,319	901 (68%)	564 (43%)
Seattle	1,792 sqkm	3.52 min.	1,096	627 (57%)	410 (37%)

Table 1: Overpass extraction statistics using Mapreduce with 2,000 machines. We filter the structures sufficient for 2D segmentation and 3D modeling by a confidence metric in Section 3.5.

Figure 6: Successful reconstruction of multiple overlapping overpasses of the I-105/I-710 Interchange in Compton, California.

Figure 7: Geometric reconstruction of expanded I-5/Artesia Boulevard overpass structure from older vector road data.

Figure 8: Overpasses that abruptly change width at the I-5/California State Route 91 interchange are poorly captured by our method.

7. REFERENCES

[1] N. Akel, K. Kremeike, S. Filin, M. Sester, and Y. Doytsher. Dense DTM generalization aided by roads extracted from LiDAR data. pages 54–59, 2005.

[2] J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51:107–113, January 2008.

[3] W. A. Harvey and J. McKeown, David M. Automatic compilation of 3D road features using LIDAR and multi-spectral source data. In *ASPRS 2008 Annual Conference*, Portland, Oregon, April 28 - May 2 2008.

[4] M. T. Jong-Hyeok Jeong. Extraction of bridge positons from IKONOS images for accuracy control of bridge database.

[5] N. Loménie, J. Barbeau, and R. Trias-Sanz. Integrating textural and geometric information for an automatic bridge detection system. In *Proc. of the 2003 International Geosciences And Remote Sensing Symposium (IGARSS 2003)*, Toulouse, France, July 2003.

[6] S. J. Oude Elberink. *Acquisition of 3D topography : automated 3D road and building reconstruction using airborne laser scanner data and topographic maps.* PhD thesis, University of Twente, Enschede, April 2010.

[7] C. Qian, B. Gale, and J. Bach. Earth documentation: Overpass detection using mobile LiDAR. In *Image Processing (ICIP), 2010 17th IEEE International Conference on*, pages 3901–3904, September 2010.

[8] G. Sithole and G. Vosselman. Experimental comparison of filter algorithms for bare-earth extraction from airborne laser scanning point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing*, 59(1-2):85–101, 2004.

[9] R. Trias-Sanz and N. Loménie. Automatic bridge detection in high-resolution satellite images. In *Proceedings of the 3rd international conference on Computer vision systems*, ICVS'03, pages 172–181, Berlin, Heidelberg, 2003. Springer-Verlag.

[10] C. Valla. Postcards from Google Earth, bridges, 2010. http://clementvalla.com/index.php?/work/bridges/.