Extraction and Modeling of Power Lines from ALS Point Clouds

Thomas Melzer, Christian Briese
Christian Doppler-Laboratory for
Spatial Data from Laser Scanning and Remote Sensing at
Institute of Photogrammetry and Remote Sensing
Vienna University of Technology
tm@ipf.tuwien.ac.at

Abstract:
Due to its ability to provide dense, fast and accurate range measurements, airborne laser scanning (ALS) is becoming increasingly popular for extensive, large area surveying tasks; however, there is still a lack of (published) methods for 3D segmentation and object modeling that exploit the full potential of this new technology. In this paper, we investigate strategies for extracting and reconstructing power lines in 3D from ALS point clouds and give preliminary results for a 1000m X 140m test scan; exact 3D-reconstructions of power lines are important for energy providers in order to assert that the physical parameters of the cable are still within a safe margin and to quickly detect anomalies or defects.

1 Introduction

Airborne laser scanning (ALS), also referred to as LIDAR, is a relatively new technology for obtaining dense (up to 10 points/m^2), accurate range measurements over large areas; for an introductory overview, see [11]. It is mainly used for the generation of detailed digital terrain models (DTMs), e.g. [7], but also for surveying and reconstruction of city areas [9, 1]. The scanner system, which is mounted on an airborne platform (plane or helicopter), actively emits laser pulses at a given rate (typically about 40 kHz), and measures the travel time between the emitted and reflected pulse, from which the distance to the illuminated spot on the ground can be computed. The laser beam is continuously moved (or deflected) across the direction of flight, resulting in a zigzag pattern of illuminated points. The polar coordinates (distance, deflection angle) thus obtained by the scanner system have to be synchronized with and transformed according to GPS/INS measurements pertaining to the position/orientation of the platform in order to obtain the final 3D cartesian coordinates of the scanned points.

From a technological point of view, commercially available ALS systems are by now relatively mature (though the technology is still rapidly evolving). However, the development of algorithms for the subsequent processing of ALS data, in particular for classification and reconstruction of complex man-made objects, is in comparison still in its early adolescence - and thus remains a very active and challenging research field - for several reasons: First, the planimetric (xy) coordinates of the points provided by an ALS scan are, by the very nature of the scanning process, *sub-randomly* distributed [1]. In particular, this prevents the direct application of image processing techniques, unless the laser data are gridded (i.e., interpolated into a regular grid) first, which, besides loss of accuracy, can lead to additional problems (e.g., 3D-information present in the original data might be "squashed" to 2.5D). Second, due the huge amount of data provided by ALS scans, specialized tools and algorithms for efficient storage management and data reduction are required. Third, the interaction of the laser beam with the environment - and hence the physical "meaning" of the backscattered pulse - is not yet sufficiently well understood, and the results of continued research in this direction are likely to have profound impact also on the later processing stages [10].

In this paper, we will investigate strategies for extracting and reconstructing power lines from 3D ALS point clouds. This task is carried out on behalf and with support of the Verbundplan Prüf- und Messtechnik (VPM), which requires this information in order to a) detect anomalies or defects in the power lines or the pylons, b) assert that line corridors have been built according to plan, and c) determine the "clearance sector" about the line corridors (i.e., the area free of vegetation and buildings). Obtaining exact 3D measurements on small diameter objects like power lines is extremely costly and time consuming using manual tachymetric techniques, and near impossible with standard photogrammetric approaches. ALS is the first technology that is able to provide dense, fast and accurate measurements on such intricate structures at reasonable cost.

In the next section, after giving a brief survey of classical 2.5D range segmentation techniques, we give a detailed discussion of our approach to 3D extraction and modeling of power lines from ALS data. Some open problems and ideas for future work are discussed in section 3.

2 Implementation

Most existing approaches for range image segmentation in computer vision are based on region-growing: 2.5D surface patches (typically planes or low-order polynomials) are fitted locally to several small seed regions, which are then grown by adding neighboring points consistent with the current surface-hypothesis (i.e., points yielding a low residual w.r.t. the estimated range value). After a region has been grown, the parameters of its corresponding surface are re-estimated. This process is iterated until a given termination condition is met. See, for example, [2] (seminal paper), [6] (overview and experimental comparison of different algorithms). Note that a segmentation thus obtained also provides a 2.5-reconstruction of the scene in terms of the instantiated primitives (i.e., surface patches). However, such a "partial reconstruction" will, in most practical applications, serve only as an intermediate representation, and additional steps - like surface grouping - are still required prior to object classification or scene interpretation.

Although the general approach outlined above has also been applied to the extraction of buildings from LIDAR data, (see [9]), it does not seem well suited for the extraction of geometrically more intricate structures like power lines. Since, for example, we must also be able to detect all of several different lines located above each other, the distribution of the data cannot simply be assumed (up to noise) as intrinsically 2-dimensional[1)], but the model(s) used must respect the inherently 3-dimensional nature of the data. For the same reason, *gridding* (i.e., interpolating and resampling the data into a regular grid) - which is a necessary prerequisite for the application of any image processing technique (e.g., morphological filtering) in general, and the segmentation approaches discussed above in particular - should be used with extreme caution or, better, not at all.

In our approach, after preprocessing, groups of parallel power lines (corridors) are located in the projection of the point cloud onto the x-y-plane using a 2D Hough transform (HT, see, e.g.,[4]) . A 3D fit for each power line is then computed locally within its corresponding corridor. In the rest of this section, we discuss our implementation and illustrate the major steps with results obtained from a test scan provided by VPM, which covers a 1000 m flight segment recorded with 140 m swath (scan line) width, which is shown in Fig. 1a).

2.1 Preprocessing

A DTM (digital terrain model) is computed and the terrain (i.e., non-object) points removed. For this purpose, we use the robust interpolation method for DTM-generation proposed by Kraus et al. [7]. Their approach is related to iteratively reweighted least squares M-estimation, but uses an asymmetric weight function that penalizes points above the fitted surface more than points under the surface.

In order to alleviate subsequent computations, a Karhunen-Loeve transformation is performed: the xy-coordinates of the original data are mean-normalized, and the x-axis is aligned with the axis of least inertia (first principal eigenvector); after this transformation, the x-axis will point approximately in the general flight direction.

[1)] This is, in essence, the so called *surface coherence assumption* [2] underlying all region-growing approaches based on surface fitting.

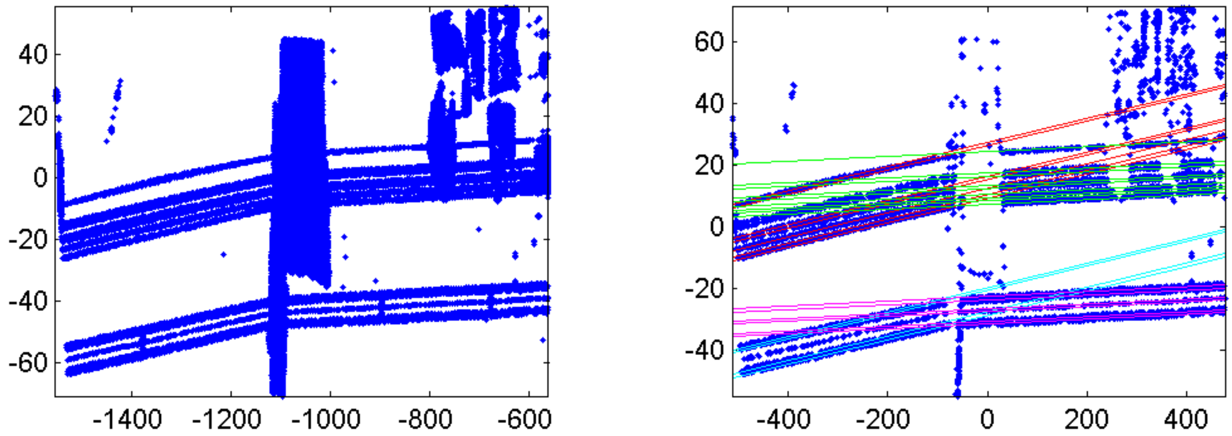

Figure 1: a) Planimetric view of test scan with terrain points already removed. The high density area in the middle part of the image contains mainly vegetation, but also the pylons. b) Data set after filtering and mean normalization, 4 clusters of parallel power lines (corridor segments) extracted by the HT can be clearly distinguished.

2.2 Culling (Filtering)

Even with the terrain points removed, the HT cannot be applied immediately, since there are still too many non-line points present.[2)] A particular problem in our case are vegetation areas with high local point density, since these tend to clutter up the accumulator. Hence, we devised a culling mechanism, which first subdivides the x-y-plane into 1 m $\times$ 1 m cells, then binarizes the grid by assigning 1 to all non-empty cells and 0 to all others, and finally removes all 1-cells having either no (singular area) or more than 3 (high density area) 8-neighbors in the grid-topology. This operation can be thought of as a kind of *erosion* that tries to preserve low-dimensional structures; it seems to work well in practice and, as can bee seen also from Fig. 1b), gets rid of (or, at least, yields a substantial thinning of) the most troublesome non-line regions with high point density.

2.3 Iterative HT

The HT is performed on the points of the original data set belonging to cells which have not been removed during the previous filtering step.

We found that the standard HT did not work well, even on the filtered data. The reason is that the point density among the power lines themselves can be very inhomogeneous, and

[2)] Although the Hough transform is a robust technique, it clearly cannot give reasonable results for data sets consisting mostly of outliers. For this reason, the HT is normally applied to features indicating the presence of the primitive of interest (e.g., edge images), and not directly to the original data.

lines generating a large number of return pulses (these are, for example, lines located above or nearly above each other) tend to clutter up the accumulator in a fashion similar to high density vegetation areas. The effect is that such areas tend to produce many, only slightly different line hypotheses in their vicinity, each yielding a high number of votes.[3)] This effect can be countered, to some degree, by reducing the resolution of the accumulator, but unfortunately this also reduces the accuracy of the estimates. Another possibility would be to perform some kind of non-maxima suppression, but this has not been explored yet. Instead, we employed an iterative version of the HT, which, after each run, removes all data points supporting the hypothesis with the highest number of votes (i.e., these points are not allowed to vote again). This approach yields, in our case, better result than the standard HT, and, in addition, allows us to determine the number of lines in a semi-adaptive fashion, e.g., "stop after more than 80% of points are accounted for by line hypotheses". The results of the HT on the filtered image are shown in Fig. 1b).

2.4 Line Clustering

The lines detected by the HT are grouped into corridor segments using a *minimum linkage hierarchical clustering* approach [3], and the intersections of these segments are computed. This step reduces the amount of data to be considered during subsequent computations, gives an initial estimate of the position of the pylons and allows to analyze each corridor segment in turn. A typical corridor segment is shown in Fig. 2.

2.5 Local Analysis and Catenary Fitting

During this step, the scan points belonging to a single corridor segment (between two pylons) are segmented, i.e., assigned to individual power lines, while the line parameters are estimated. The form of a hanging cable supported at both ends acted upon by gravity is given by the so called *catenary* curve:

$$z = 0.5\ a\ (e^{x/a} + e^{-x/a}) = a \cosh(x/a), \tag{1}$$

whereby a is the "elasticity" constant of the cable. We model the power lines as 3D catenaries with 6 parameters (position, elasticity, roll and yaw angle). However, in order to compute the actual - non-linear - fit, one has to estimate $N + 6$ parameters[4)] simultaneously using an iterative optimization technique (e.g., Gauss-Newton). Similarly, computing the shortest distance between N points and a given catenary involves non-linear minimization (although,

---

[3)] At a more profound level, this problem is due to limited spatial resolution Δx of the laser scanning system: sensed lines have a non-negligible width (up to $2\Delta x$), and nearby singular lines become indistinguishable.

[4)] Since the ALS measurements have noise in all 3 coordinates, it is not sufficient to minimize only the height residuals. Instead, one has to minimize the distance of the points from the catenary, resulting, in our case, in N additional parameters for the catenary "parameter" x.

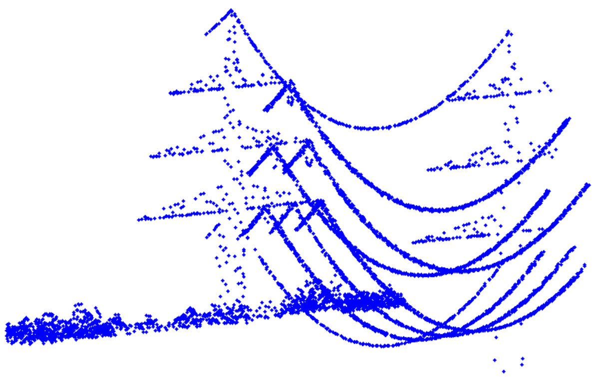

Figure 2: 3D view of the points belonging to complete corridor segment, including pylons and vegetation, but terrain points removed.

in this case, the minimization(s) can be performed sequentially) and thus is computationally rather costly. Since the catenary parameters are estimated robustly using RANSAC [5] (which requires repeated computation of the support and thus of the distances from the hypothesized catenaries), it is crucial to have good initial segmentation hypotheses, i.e., subsets of points belonging, with high probability, to the same power line. We obtain these by training a Neural Gas Network on the whole corridor segment, and then performing a *Competitive Hebbian Learning* (CHL) step [8]. Neural Gas is a vector quantization algorithm which approximates the distribution of the ALS points with a small number (200 in our case) of reference vectors. CHL connects neighboring reference vectors whose corresponding second-order Voronoi polyhedra are non-empty[5)], thus yielding a subset of the Delauny triangulation of the reference vectors. Provided the number of reference vectors is large enough, CHL will not connect reference vectors belonging to different power lines; however, a larger number of reference vectors leads to longer training times and tends to produce multiple connected regions per power line. The connected components determined by CHL now serve as pool for generating minimal subsets of size 6, which are used in a RANSAC competition in order to robustly fit the catenaries; an example of a fitted catenary is shown in Fig. 3.

---

5)The second-order Voronoi polygon V_{ij} of two units i, j is the set of all points $\mathbf{v}$ for which the corresponding reference vectors $\mathbf{w}_i, \mathbf{w}_j$ are the two nearest neighbors.

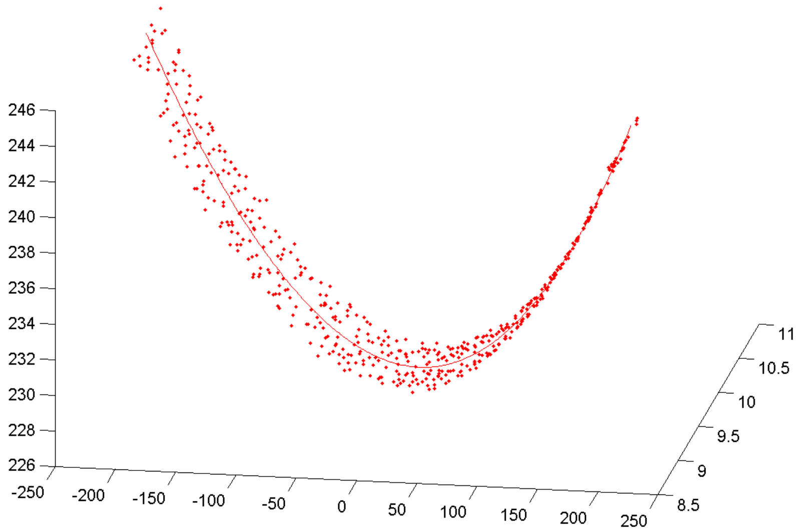

Figure 3: Fitted catenary plus consensus set (support). Note that for visualization purposes, the axes have been scaled differently (the catenary is about 500 m long and only 30 m high).

3 Conclusion and Outlook

We have presented an approach for 3D extraction and recognition of power lines from ALS data. Although the preliminary results are encouraging, there still remain several open issues. One problem that has not been tackled yet is the "concatenation" of single corridor segments (including the fitted power lines) into corridors. Also pending is the improvement of several of the already implemented parts, namely the filtering (culling) algorithm and the iterated HT, discussed in sections 2.2 and 2.3, respectively.

The filtering algorithm seems to work reasonably well for scanner systems that provide an approximately equal spacing of the points in the in-flight (x) and cross-flight (y) dimensions (e.g., TopScan); unfortunately, this is not true, for example, for the popular TopoSys scanners (which use a higher sampling frequency in the in-flight direction). A possible remedy is to compute the first derivatives within a scan line in order to identify candidate points, which are then passed to the HT. Unfortunately, the topological structure (scan line neighborhood) is often not preserved in the recorded scans (as in the case of our test scan).

As stated earlier, we had to resort to an iterative version of the HT in order to compensate for accumulator clutter. However, the iterative HT is a somewhat dangerous approach, in that the removal of the support set of one false line hypothesis will affect adversely all subsequent hypotheses . We intend to obviate the need for an iterative version of the HT by either employing a model selection approach (e.g., MDL-based) or by detecting maxima (modes) in the accumulator (using, e.g., the mean shift algorithm).

References

[1] Peter Axelsson. Processing of laser scanner data - algorithms and applications. *ISPRS International Journal of Photogrammetry and Remote Sensing*, 54:138–147, 1999.

[2] Paul J. Besl and Ramesh C. Jain. Segmentation through variable-order surface fitting. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 10(2):167–191, 1988.

[3] Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification*. John Wiley & Sons, 2nd edition, 2000.

[4] Rafael Gonzalez and Paul Wintz. *Digital Image Processing*. Addison Wesley, 2nd edition, 1987.

[5] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.

[6] Adam Hoover, Gillian Jean-Baptiste, Xiaoyi Jiang, Patrick Flynn, Horst Bunke, Dmitry Goldgof, Kevin Bowyer, David Eggert, Andrew Fitzgibbon, and Robert Fisher. An experimental comparison of range image segmentation algorithms. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 18(7), 1996.

[7] K. Kraus and N. Pfeifer. Determination of terrain models in wooded areas with airborne laser scanner data. *ISPRS International Journal of Photogrammetry and Remote Sensing*, 53:193–203, 1998.

[8] Thomas Martinetz and Klaus Schulten. Topology representing networks. *Neural Networks*, 7(3):507–522, 1994.

[9] F. Rottensteiner and C. Briese. A new method for building extraction in urban areas from high-resolution LIDAR data. In *International Archives of Photogrammetry and Remote Sensing*, volume XXXIV/3A, pages 295–301, 2002.

[10] Wolgang Wagner, Andreas Ulrich, and Christian Briese. Der Laserstrahl und seine Interaktion mit der Erdoberfläche. *VGI, Österreichische Zeitschrift für Vermessung und Geoinformation*. To appear in 2004.

[11] Aloysius Wehr and Uwe Lohr. Airborne laser scanning - an introduction and overview. *ISPRS International Journal of Photogrammetry and Remote Sensing*, 1999.