

Automatic Feature Extraction (AFE)

Innovation in the cost-effective production of numerous 3-D features

Dr. Bingcai Zhang, Engineering Fellow

BAE Systems Geospatial eXploitation Products

10920 Technology Place

San Diego, CA 92127-1874

bingcai.zhang@baesystems.com

Table of contents

Executive summary.....	3
Introduction.....	3
Principles of AFE.....	4
Using AFE	5
Minimizing editing by multiple runs with different parameters.....	7
Coastal Oregon LiDAR case study.....	10
Availability of AFE.....	28
Conclusions.....	28
References.....	28

Executive summary

Automatic feature extraction is considered to be the Holy Grail by many photogrammetrists. Automatic processing of imagery is well established for image registration and digital terrain generation, but extraction of 3-D buildings is not so successful. The radiometric properties of these features are very complex and variable. BAE Systems takes a new approach by extracting 3-D features from point clouds, either from LiDAR or generated by NGATE¹ from stereo images. BAE Systems has introduced Automatic Feature Extraction (AFE), an effective production-capable system that can extract massive amounts of 3-D features from 3-D point clouds. AFE is easy to use and can achieve success rates above 90% with dense, accurate LiDAR point clouds. AFE is designed to minimize manual editing by multiple runs with different sets of parameters. Tests indicate that AFE's high success rate and optimized workflow represent considerable progress in terms of productivity and generate cost savings sufficient to provide ample return on investment.

Introduction

Over the past few decades, attempts to extract 3-D features (buildings, houses, single trees, etc.) automatically from imagery have not generated a system robust enough for use in production. The radiometric properties or spectral characteristics of 3-D features are very complex and variable. Figure 1 indicates the difficulty for any algorithm automatically to extract six buildings from imagery alone. The six buildings have different colors and patterns. Algorithms that work well with one set of images and 3-D features may not work at all with a different set of images and 3-D features, because the radiometric properties are often sufficiently different.



Figure 1. Six different building colors and patterns. A supervised building region growing classification would need six signatures. The upper-right and lower-middle signatures cannot be used because they are not homogeneous.

LiDAR data has unique properties for automatic extraction of 3-D features. The most important and *invariant* property of a 3-D feature in LiDAR data is its three-dimensionality: in other words,

¹ NGATE is software for automatic extraction of elevation by matching of multiple overlapping images. It is available as a module for BAE Systems' commercial-off-the-shelf SOCET GXP[®] and SOCET SET[®] products.

the very availability of Z distinguishes features better than a 2-D image view. We can use this property to identify, extract and label 3-D features automatically. To identify a feature in digital images, it is crucial to use a feature property that does not change, i.e., is *invariant*. The 3-D properties of a 3-D feature are ideal. As shown in Figure 2, terrain shaded relief (TSR) highlights 3-D features in a digital surface model (DSM). In this case the DSM was photogrammetrically derived from stereo imagery by means of NGATE (Zhang and Walter, 2009). All of the 3-D features have one common property — they are above the ground. Modern stereo image matching algorithms and LiDAR provide very dense, accurate DSMs, which can then be used for automatic extraction of 3-D features (Zhang and Smith, 2010). In 2009, Dr. Bingcai Zhang began to work on AFE using this invariant 3-D property.

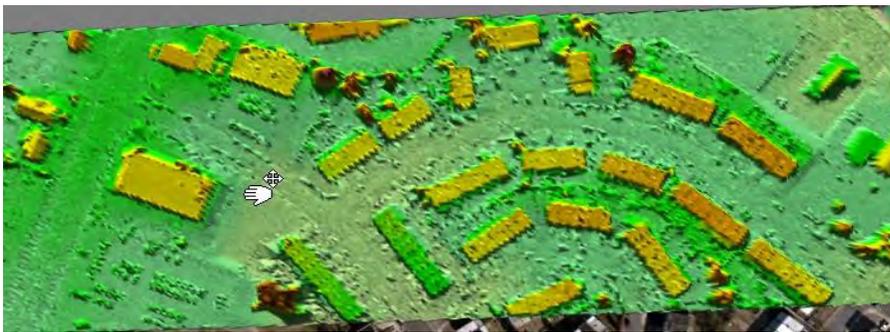


Figure 2. Terrain shaded relief of digital surface model. The existence and approximate shapes of 3-D features are immediately apparent. It is much easier to classify TSR for 3-D features than to classify digital images.

Principles of AFE

The theory behind AFE and the algorithms associated with it have been described in a series of published papers, for example Zhang and Smith (*ibid.*). AFE consists of 30 algorithms that determine the 3-D planes of a building roof and their intersections, roof polygon segments and their intersections, roof polygon dominant direction, 3-D mesh of a roof, and segmentation of each 3-D feature.

The foundation of AFE's algorithms is as follows:

- The roof of a building consists of a set of 3-D planes including vertical 3-D planes. The 3-D planes intersect and form a complete surface. The 3-D planes are arbitrary and random and there are no pre-defined roof templates. AFE uses a number of algorithms to determine the set of 3-D planes and their intersections.
- The rooftop polygon of a building consists of a set of 3-D line segments, including vertical line segments. A line segment has two 3-D endpoints. The line segments intersect and form a polygon in 3-D space. The line segments may or may not be parallel or perpendicular to each other. AFE uses a number of algorithms to determine the line segments and their intersections.

- The roof of a building has one dominant direction. This is computed based on each line segment direction. When all line segments are parallel or perpendicular to each other, the rooftop polygon has a perfect dominant direction and most likely has all 90-degree corners.
- Buildings are segmented on the basis of elevation differences between DSM and DEM, which are compared to a minimum building height parameter.

Using AFE

There are multiple ways to launch AFE. The user interface is illustrated in Figure 3. Users with LiDAR LAS files typically add multiple LAS files as input terrain. AFE will then ingest these LAS files and merge them into a single GRID DSM. Each LAS file is imported individually and all returns are written into a single TIN DSM. Users must ensure that multiple LAS files are adjacent to or overlapping each other since they will be merged into a single DSM. When merging the TIN DSMs into a GRID DSM, AFE has logic to select the optimal post spacing. Terrain Operation Merge can take a long time when the dataset is huge (more than several hundred million posts). Users must ensure that the total number of points in all the input LAS files is less than 100 million. When working with large projects with thousands of LAS files and billions of points, users must divide the project area into sub-regions such that each sub-region has less than 100 million points. When there is only a single LAS file, AFE will import it into a GRID DSM without Terrain Operation Merge.

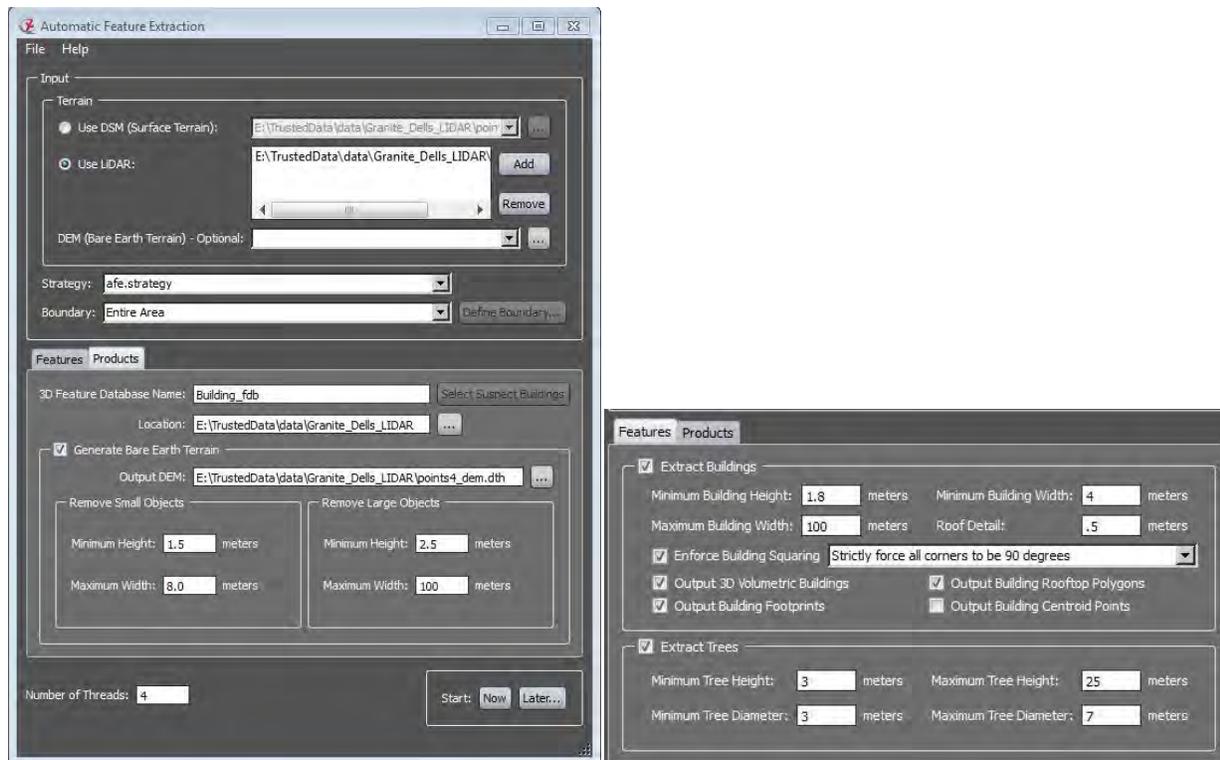


Figure 3. AFE user interface. The input terrain or 3-D point clouds can be from either LiDAR LAS files or DSMs generated by NGATE from stereo images. AFE can handle multiple LAS files. When a bare-earth terrain model (DEM) is not available, AFE automatically transforms a DSM into a DEM. There are two tabs, Features and Products. Under the Features tab, users specify building parameters as well as tree parameters, which control 3-D feature extraction. Under the Products tab, users specify parameters that control DSM to DEM transformation.

In a typical case, no accurate DEM is available in the project area and AFE will generate a DEM from the DSM using the parameters specified in the Products tab. The Features and Products tabs include the Maximum Building Width and Remove Large Objects Maximum Width parameters respectively. Operators should use a value of Maximum Building Width greater than Remove Large Objects Maximum Width, for example, 120 m and 80 m respectively. When the project area is relatively flat, operators can use larger values for these two parameters. When the terrain is hilly and undulating, operators should use a smaller value for the Remove Large Objects Maximum Width.

The Maximum Tree Height in the Features tab is used to separate trees from buildings. In areas such as Oregon and Lake Tahoe where there are tall trees, users must use a large value such as 60 m. In areas such as downtown San Francisco and San Diego, users should use a smaller value such as 25 m. AFE uses the elevation variation to segment trees from buildings. The roofs of many downtown tall buildings, however, may have complex artifacts such as TV antennae, air conditioners, or even potted plants. Any point with a height above the Maximum Tree Height is

considered a building post regardless of its elevation variation.

The vast majority of buildings have 90-degree corners. Users should run AFE twice. In the first run, users should select Enforce Building Squaring – Strictly force all corners to be 90 degrees. Thereafter, operators should query all buildings with Percent_of_parallel_or_perpendicular_sides ≤ 99 and delete them. In the second run, users should not select Enforce Building Squaring and use the same feature database as in the first run, so that the second run appends the missing features into the same feature database. AFE has logic not to extract the features that are already in the feature database. This logic speeds up the second run and also avoids duplicated features.

AFE's success rate depends on the quality of the bare-earth model or DEM. In areas of hilly and undulating terrain, the automatic transformation from DSM to DEM can be challenging. In this case, operators should then use terrain editing tools to edit the DEM and ensure that it is of high quality. If stereo images are not available, operators can generate a terrain shaded relief (TSR) image from the DSM, display it as a background image and use the Dense Trees/Buildings with Polygon on Ground tool to edit the DEM. From the TSR image, it is easy to identify the bare ground and use terrain tracking to obtain the ground elevation when digitizing the polygon. In many production situations, a DEM is a valuable product by itself. Thus this DEM editing effort is needed anyway.

Minimizing editing by multiple runs with different parameters

One of the main design concepts of AFE is to increase productivity by minimizing manual editing. The accuracy and success rate of building extraction are dependent on the parameters that users specify. Users can run AFE multiple times with different parameters. In this section, we demonstrate a recommended workflow to achieve high success rate and accuracy through multiple runs of AFE with different parameters. The data (Granite Dells, AZ) is provided by National Center for Airborne Mapping, hosted at OpenTopography Facility, San Diego Supercomputer Center. The LiDAR density is 7 points per m^2 and the data is of high precision.

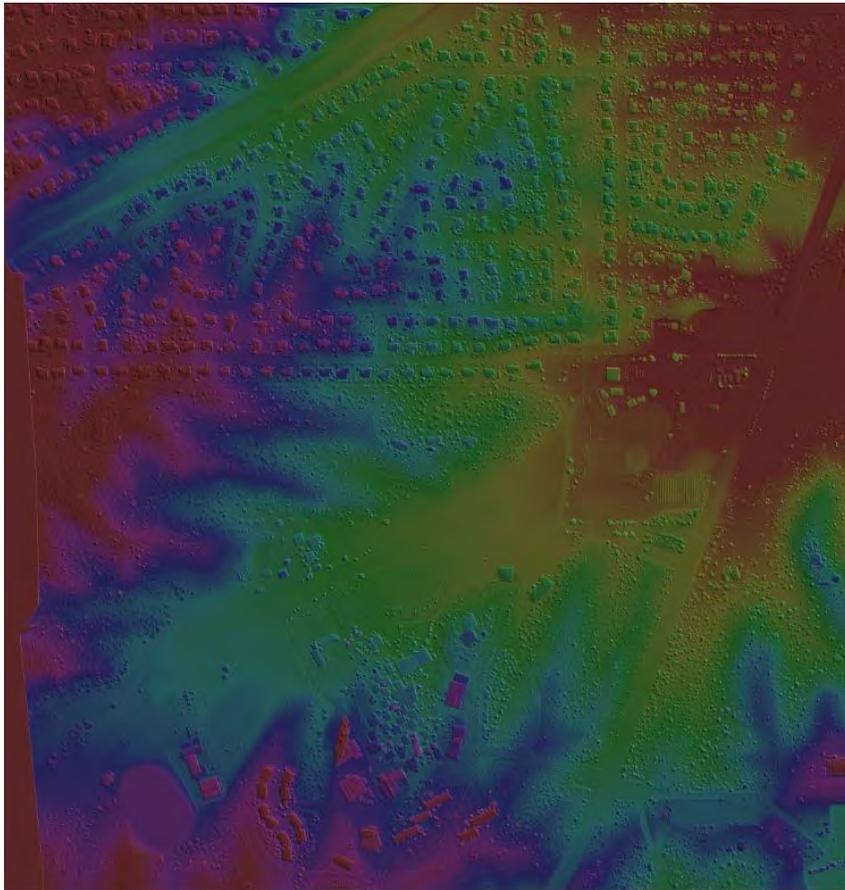


Figure 4. This terrain shaded relief image indicates clearly that most 3-D features are houses with 90-degree corners. Only a few large buildings in the lower portion of the image have non-parallel sides. We ran AFE twice. The first time, we used Enforce Building Squaring, which strictly forces all corners to be 90 degrees. The second time, we did not enforce building squaring.

We used the following set of parameters for the first run: Remove Small Objects Minimum Height 1.5 m; Remove Small Objects Maximum Width 8.0 m; Remove Large Objects Minimum Height 2.5 m; Remove Large Objects Maximum Width 100 m; Minimum Building Height 1.8 m; Minimum Building Width 4 m; Maximum Building Width 100 m; Roof Detail 0.5 m; Minimum Tree Height 3.0 m; Maximum Tree Height 25.0 m; Minimum Tree Diameter 3 m; Maximum Tree Diameter 7 m; Enforcing Building Squaring - Strictly force all corners to be 90 degrees. A total of 533 houses and buildings were extracted. Among them, 19 houses and buildings were incorrectly extracted because they do not have all right-angled corners. There were 11 buildings and houses missing due to DEM errors. The DEM was automatically transformed from LIDAR point clouds and is not error free. The success rate is $1.0 - (11+19)/533 = 94.4\%$.

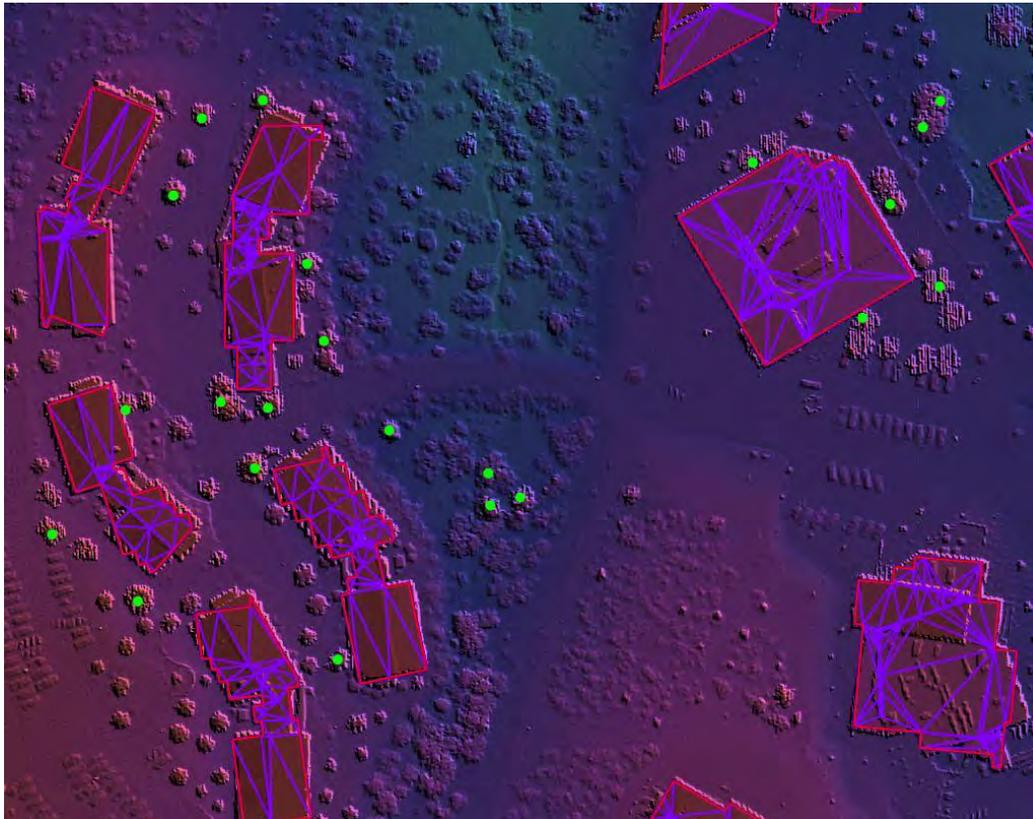


Figure 5. Buildings with non-parallel sides were incorrectly extracted because we used the Enforce Building Squaring, which strictly forces all corners to be 90 degrees parameter.

Using SOCET GXP's attribute query capability, we selected all buildings with attribute "Percent_of_parallel_or_perpendicular_sides <= 99". A total of 20 buildings and houses did not have 100% parallel or perpendicular sides. This query selected and deleted 13 of the 19 houses and buildings that were incorrectly extracted.

In the second AFE run, we did not select Enforce Building Squaring and used an edited DEM. We used the same feature database as the 3-D feature database. AFE has the logic of not extracting existing features in a database and appending only new features. Therefore, the second run was much faster and added 24 new features, 11 of them the features missing due to DEM error, and 13 the features deleted because they did not have 100% parallel sides. Figure 6 shows that some of the 13 deleted buildings were extracted correctly without the Enforce Building Squaring option.

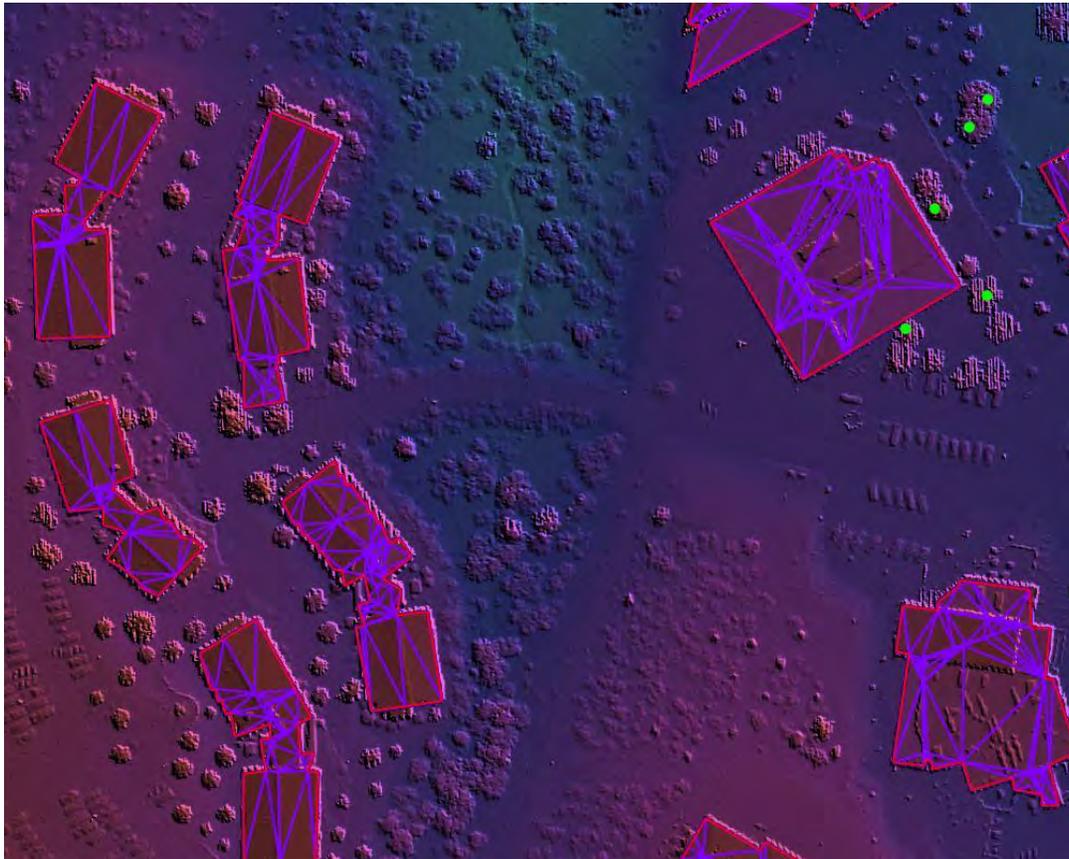


Figure 6. Buildings extracted without the Enforce Building Squaring option. The five buildings on the left side are more accurate than in Figure 5.

Using the above procedures, AFE extracted a total of 544 buildings, only six of which needed manual editing, a success rate of $1 - 6/544 = 99\%$. This is an ideal case. The terrain is relatively flat with few trees. The LiDAR point cloud is of high precision and density. Based on our tests, AFE's success rate depends on all of the above conditions. Relatively flat terrain helps the DSM to DEM transformation, which is critical for the success rate since AFE uses the differences between the DSM and DEM for the initial identification of the approximate locations and shapes of 3-D buildings and houses. High precision and high density LiDAR point clouds are two dominant factors for the success rate. AFE consists of 30 algorithms. Two groups of algorithms use the Random Sample Consensus approach to determine the rooftop 3-D planes and rooftop polygons. These algorithms require dense, accurate 3-D point clouds.

Coastal Oregon LiDAR case study

We used 10 LiDAR datasets from the Oregon coastal LiDAR project to evaluate AFE. OpenTopography's full catalog of DOGAMI Oregon LiDAR data (currently 12,678 km²) can be accessed via the Oregon Department of Geology and Mineral Industries LiDAR Program Data page on the OpenTopography Facility website (OpenTopography Facility, undated). The LiDAR

data is of high density and high precision (8 points per m², RMSE_Z < 12 cm). This data, totaling more than 100 billion LiDAR returns, covers the full extent of the Oregon coast, from the California border to the mouth of the Columbia River. The 10 datasets are for Prineville, Twin Rocks, Brookings, Coos Bay, North Bend, Redmond, Newport, Woodburn, Astoria, and Seaside as in Table 1. A total of 473 million LiDAR points was used in this case study. AFE extracted 37,409 buildings and 193,375 trees. We used the following set of parameters to run AFE just once for each dataset: Remove Small Objects Minimum Height 1.8 m; Remove Small Objects Maximum Width 10.0 m; Remove Large Objects Minimum Height 2.5 m; Remove Large Objects Maximum Width 80 m; Minimum Building Height 1.8 m; Minimum Building Width 3 m. Maximum Building Width 100 m; Roof Detail 0.5 m; Minimum Tree Height 3.0 m; Maximum Tree Height 60.0 m; Minimum Tree Diameter 3 m; Maximum Tree Diameter 7 m; Enforcing Building Squaring - Strictly force all corners to be 90 degrees.

As indicated in Table 1, AFE achieved an average success rate of 90% in a single run. As shown in Figures 7 to 23, the main cause of errors is trees overhanging buildings. In Oregon, trees are tall and can be taller than adjacent houses and buildings. We did not have success rates for Newport, Woodburn, Astoria, and Seaside owing to limited resources to measure success rate manually. Unlike NGATE, AFE does not yet have a quality statistics tool to measure success rate or accuracy.

	Points (m)	Buildings	% Success	Trees
Prineville	36	2792	90	6449
TwinRocks	32	1858	93	11852
Brookings	43	2970	89	8154
CoosBay	73	3812	87	26515
NorthBend	106	4392	91	15938
Redmond	39	1780	92	5352
Newport	37	3058		8239
Woodburn	26	1901		3407
Astoria	28	1813		4907
Seaside	53	3033		12562
Total	473	37409		193375
Average			90	

Table 1. Oregon Coastal LiDAR case study. AFE achieved an average success rate of 90% in a single AFE run based on 473 million LiDAR points covering 10 cities with a total area of 59 km².

Figure 7 shows the rooftop polygons of 3-D buildings extracted by AFE. The LiDAR dataset was downloaded from OpenTopography Facility, San Diego Supercomputer Center. There is extensive LiDAR coverage of Prineville, Oregon. We selected a small area with many buildings and houses to evaluate AFE. The LiDAR dataset is available to the public.

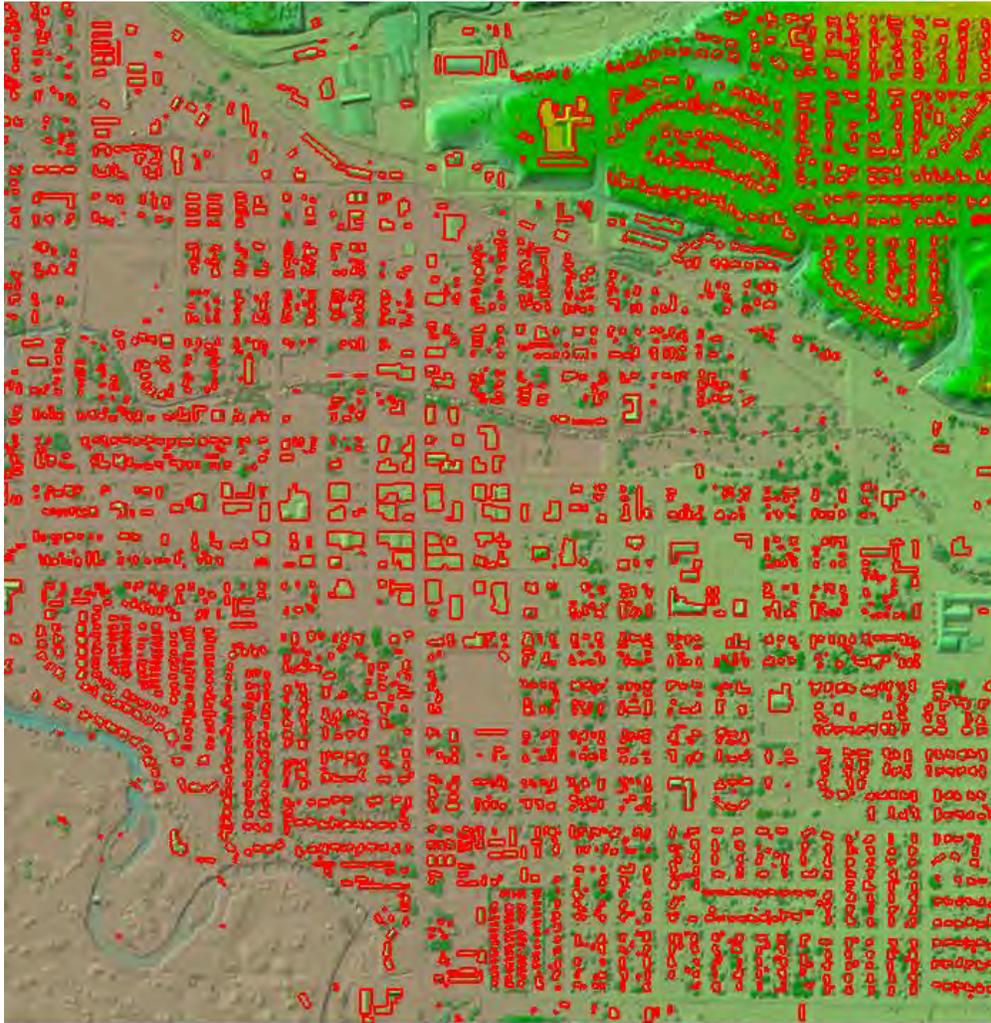


Figure 7. Prineville, Oregon. This LiDAR dataset has 36 million points covering an area of 4.5 km². The background image is terrain shaded relief generated from a 0.59 foot gridded DSM. 2792 houses and buildings were extracted by AFE with a success rate of 90%. Most of the incorrectly extracted houses and missing houses are due to overhanging trees.

Figures 8 and 9 are two subsets from Figure 7 for measuring success rate since the entire area is too large. The selection of these two subsets is random to ensure that the success rate is representative of the entire area. The red lines are rooftop polygons in Figures 7, 8, and 9. In Figure 8, the success rate is 92%; in Figure 9, 88%. The main cause of error is overhanging trees.



Figure 8. There are a total of 52 buildings. AFE extracted 49 buildings and houses correctly. Two houses were extracted incorrectly and one house is missing. AFE achieved a success rate of 92%. The two incorrect houses are due to overhanging trees. The missing house is due to the size of the house being too small. These three errors are all in the upper-left corner where there are tall trees.

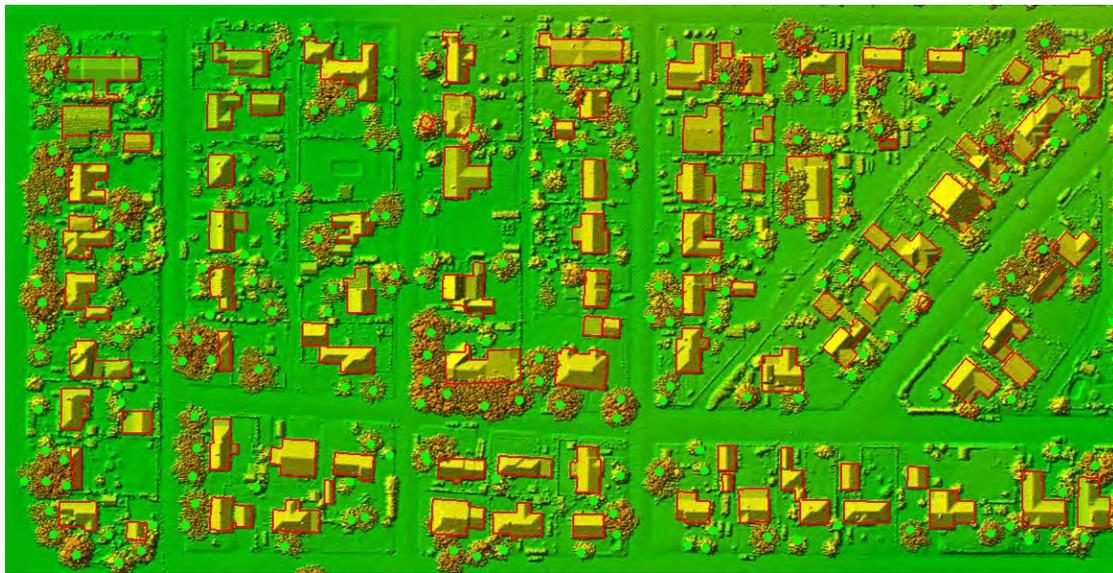


Figure 9. There are a total of 67 buildings. AFE extracted 59 correctly. Eight houses were extracted incorrectly due to overhanging trees. AFE achieved a success rate of 88%.

Figure 10 shows a complex building roof and its 3-D meshes. The number of 3-D meshes is controlled by the Roof Detail parameters and the complexity of the roof itself. AFE uses as few triangles as possible to model the roof to the specified accuracy (Roof Detail). AFE may not, however, select the same optimal set of triangles to model a roof top as a human operator. This is an area that needs more research and enhancement. In general, if all 3-D planes of a roof have sufficient points on them, AFE can accurately determine all the 3-D planes and their intersections, and select an optimal set of triangles to model the roof as shown in Figure 10. If there are not enough points on each 3-D plane, or the points are noisy, AFE may not be able to select the optimal set of triangles.



Figure 10. AFE uses a 3-D mesh to model a complex roof. On the right is an image captured from Google Earth at 44° 18' 08.64" N, 120° 50' 25.53" W. Purple lines are 3-D meshes. Red lines are the rooftop polygon. Yellow dots are trees. Image © 2012 DigitalGlobe and 2012 Google.

Figure 11 shows part of Twin Rock, Oregon. AFE extracted over 11,000 trees, which are not shown in the figure because they would cloud it.

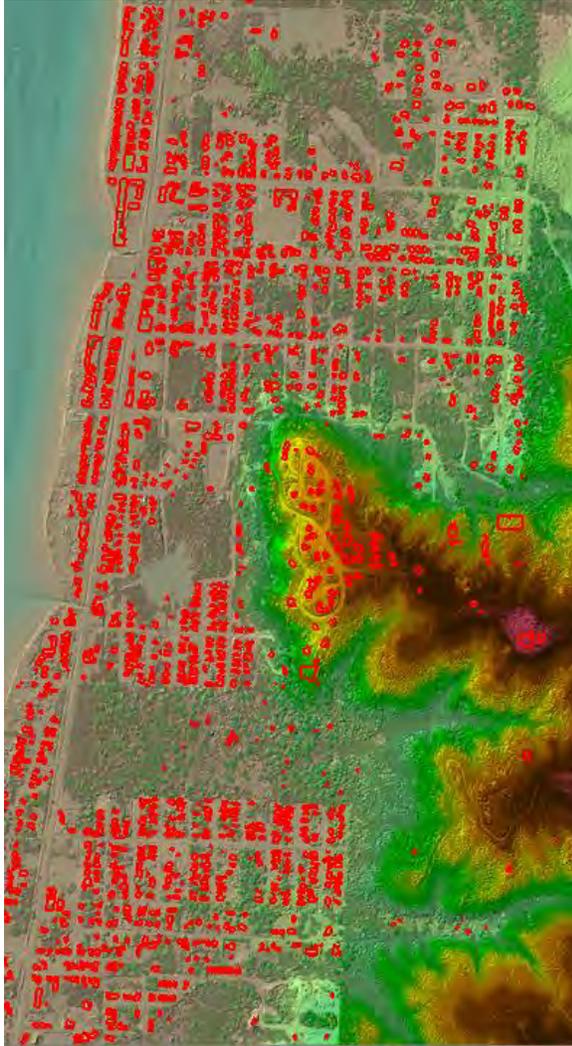


Figure 11. Twin Rocks, Oregon. This LIDAR dataset has 32 million points covering an area of 4.0 km². The background image is terrain shaded relief generated from a 0.59 feet gridded DSM. 1858 houses and buildings were extracted by AFE with a success rate of 93%. Most of the incorrectly extracted houses and missing houses are due to overhanging trees.

Figure 12 is a subset from Figure 11 for measuring AFE success rate since the entire area is too large. The selection is random. In Figure 12, the success rate is 93%. The main cause of error is overhanging trees.



Figure 12. Twin Rocks, Oregon. There are a total of 186 houses. AFE extracted 174 buildings and houses correctly. Several trees were incorrectly extracted as houses. Several houses have incorrect boundaries due to overhanging trees. AFE achieved a success rate of 93%.

Figure 13 shows Brookings, Oregon. Zwagg Island on the lower-left corner caused some difficulty for the DSM to DEM transformation. As a result, there are a few incorrectly extracted irregular polygons. AFE uses the differences between DSM and DEM for initially locating the approximate locations and shapes of 3-D features. When the DEM has errors, the extracted buildings will have concomitant errors. Users can generate DEMs without first extracting any 3-D features, then use SOCET GXP's terrain editing tools to fix DEM errors before running AFE.

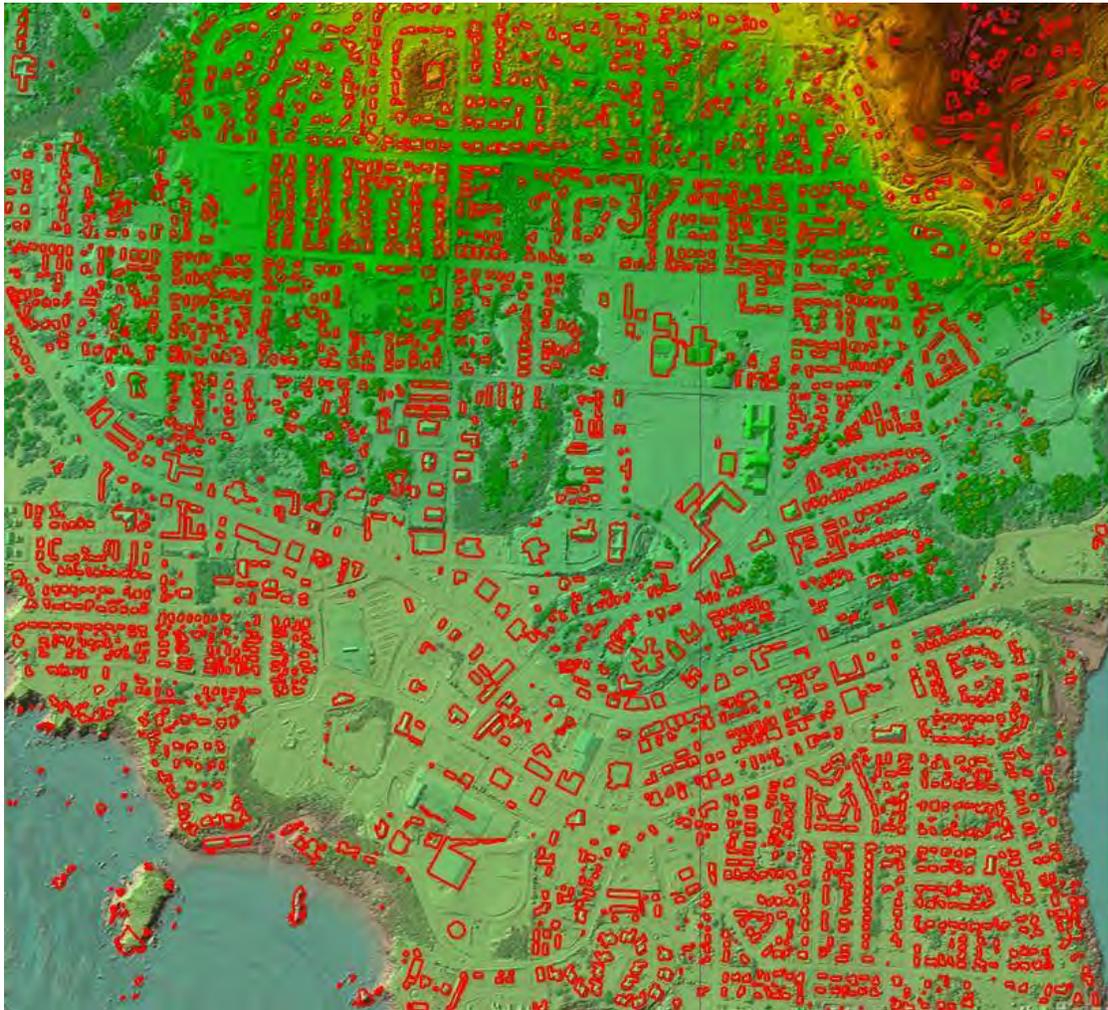


Figure 13. Brookings, Oregon. This LiDAR dataset has 42 million points covering an area of 5.4 km². The background image is terrain shaded relief generated from a 0.59 feet gridded DSM. 2970 houses and buildings were extracted by AFE with a success rate of 89%.

Figure 14 is a sub-region from Figure 13 for accuracy analysis. A few trees were incorrectly extracted as houses. Expert users may wish to modify the parameters in the `afe.strategy` file such that these trees would not be extracted as houses.

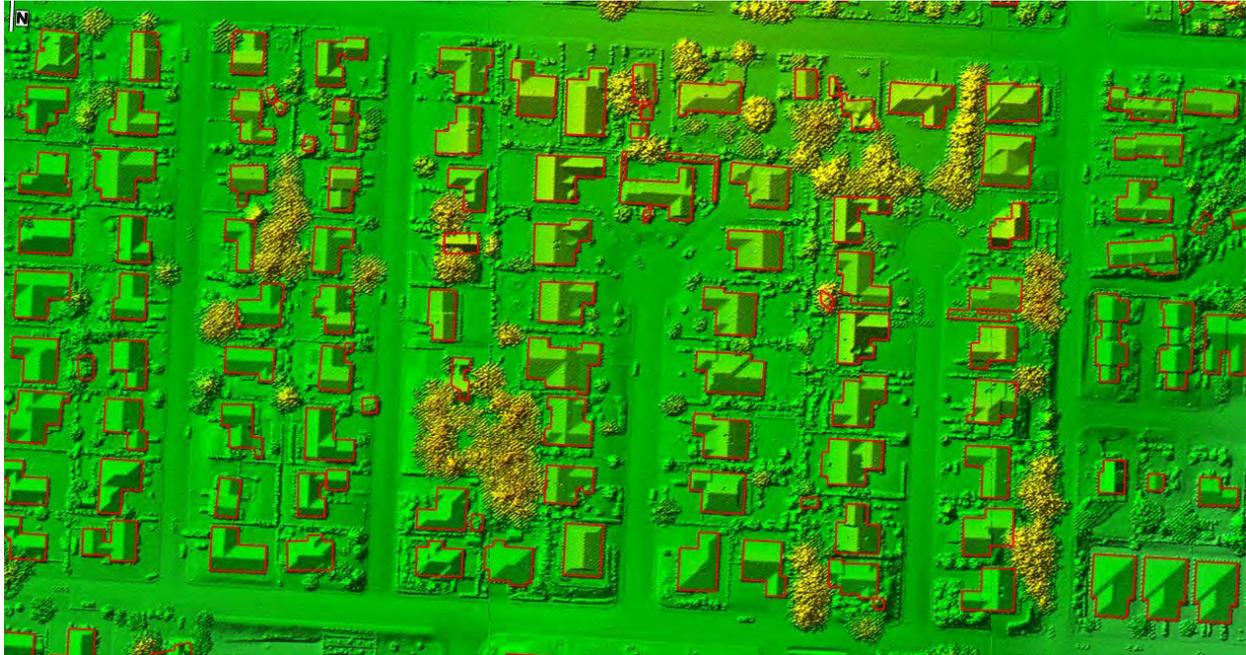


Figure 14. Brookings, Oregon. There are a total of 83 houses. AFE extracted 74 buildings and houses correctly. A few trees were incorrectly extracted as houses. A few houses have incorrect boundaries due to overhanging trees. AFE achieved a success rate of 89%.

Figure 15 shows nine triangles used to model a building. AFE can model roofs accurately only when all 3-D planes of a roof have enough LiDAR points on them and the LiDAR points are accurate enough, for example $RMSE_z < 12$ cm.

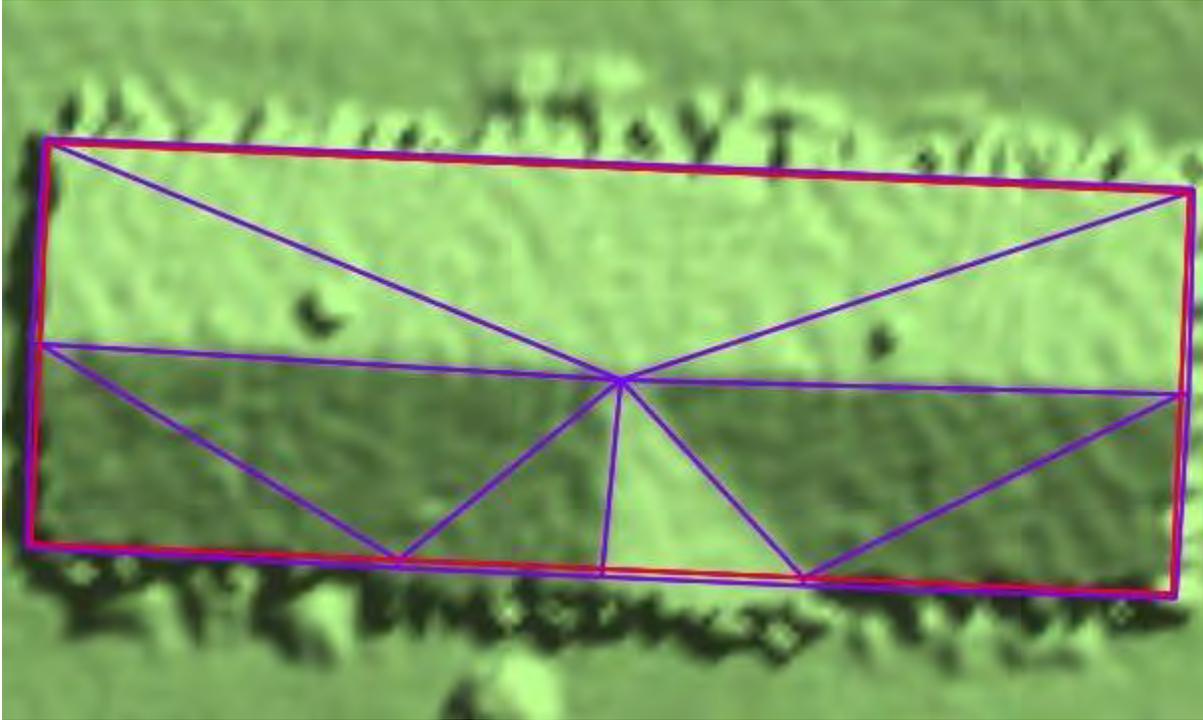


Figure 15. Brookings, Oregon. Nine triangles or 3-D meshes are used to model a house precisely.

Figure 16 shows North Bend, Oregon. On the upper-left corner is the Southwest Oregon Regional Airport. It is interesting to notice that three bridges are incorrectly extracted as buildings. In the SOCET GXP v4.0 AFE release, there is no logic to handle bridges.

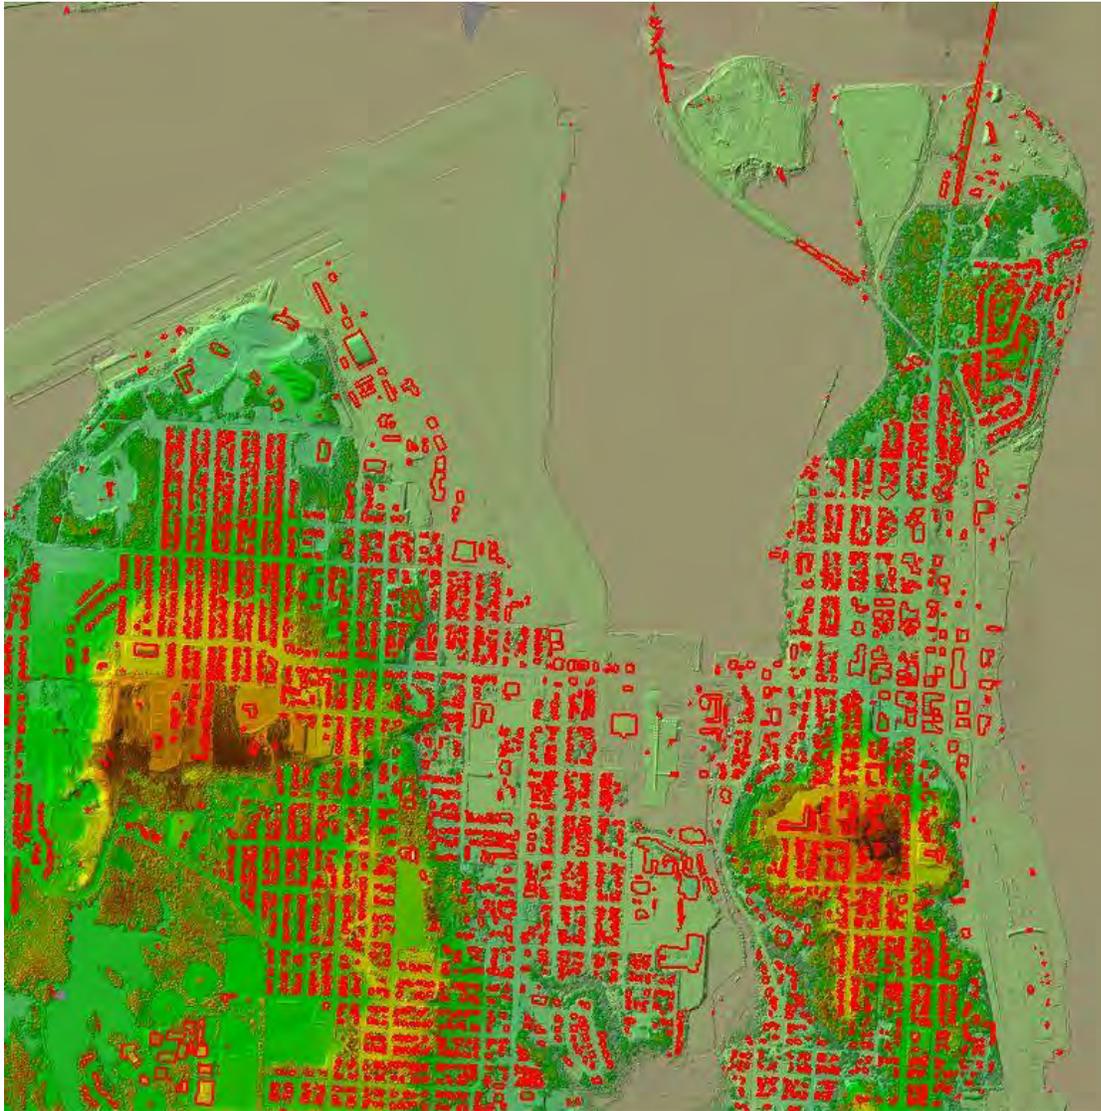


Figure 16. North Bend, Oregon. This LiDAR dataset has 73 million points covering an area of 9.1 km². The background image is terrain shaded relief generated from a 0.59 feet gridded DSM. 4392 houses and buildings were extracted by AFE with a success rate of 91%. Most of the incorrectly extracted and missing houses are due to overhanging trees.

Figure 17 is a random subset of Figure 16 for accuracy analysis. A few trees were incorrectly extracted as buildings. AFE uses elevation variation and a number of LiDAR returns to classify trees as opposed to buildings. There are a few parameters in the afe.strategy file that users can modify such that these trees would be correctly extracted. The afe.strategy is an ASCII file with sufficient explanations for expert users.

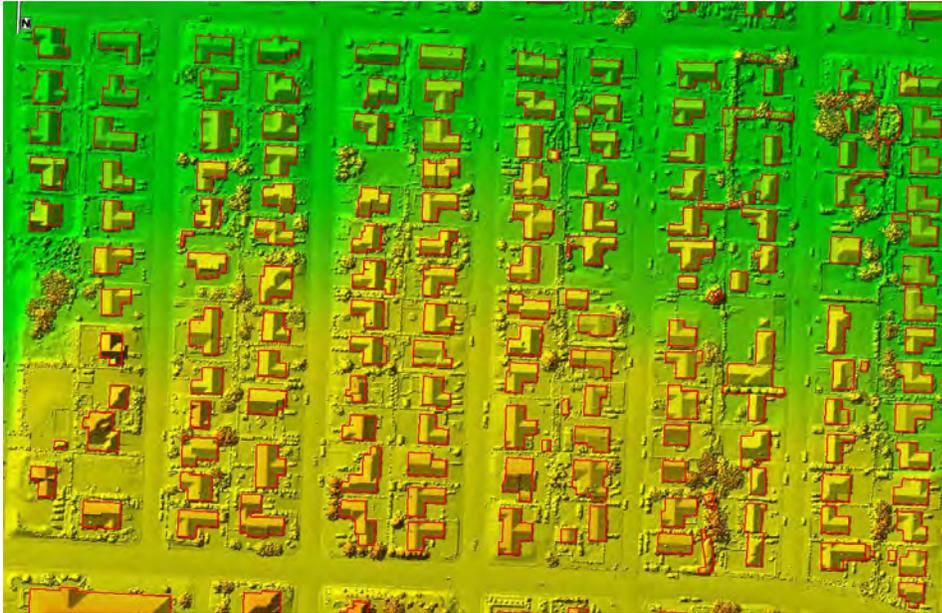


Figure 17. North Bend, Oregon. There are a total of 178 houses. AFE extracted 163 houses correctly. Several trees were incorrectly extracted as houses. Several houses have incorrect boundaries due to overhanging trees. AFE achieved a success rate of 91%.

Figure 18 shows that many 3-D meshes are used to model a complex roof. It is a very challenging issue when there are air conditioners, chimneys, and TV antennae on the rooftop.



Figure 18. North Bend, Oregon. AFE uses a 3-D mesh to model a complex roof. On the right is an image captured from Google Earth at 43° 24' 24.16" N, 124° 15' 15.67" W. Purple lines are 3-D meshes. Red lines are the rooftop polygon.

Figure 19 shows Coos Bay, Oregon. There are a few large polygons in the forested area on the left side. These large polygons are incorrectly extracted as buildings, whereas they are actually very tall trees. In the AFE GUI, the Maximum Tree Height parameter is set too low for this area. As a result, some of these really tall trees are incorrectly extracted as buildings. With the database query tool in SOCET GXP v4.0, users can easily select and delete these large polygons and run AFE a second time with a larger Maximum Tree Height parameter.



Figure 19. Coos Bay, Oregon. This LiDAR dataset has 73 million points covering an area of 9.1 km². The background image is terrain shaded relief generated from a 0.59 feet gridded DSM. 3812 buildings were extracted by AFE with a success rate of 87%. Most of the incorrectly extracted and missing buildings are due to overhanging trees.

Figure 20 is a random sub-region selection from Figure 19 for accuracy analysis. The main cause of errors is overhanging trees. Only the rooftop polygons are drawn with red lines. The 3-D meshes for some of the buildings have complex graphics such as shown in Figure 21, which could cloud the figure. Therefore, these are not drawn in Figure 20.

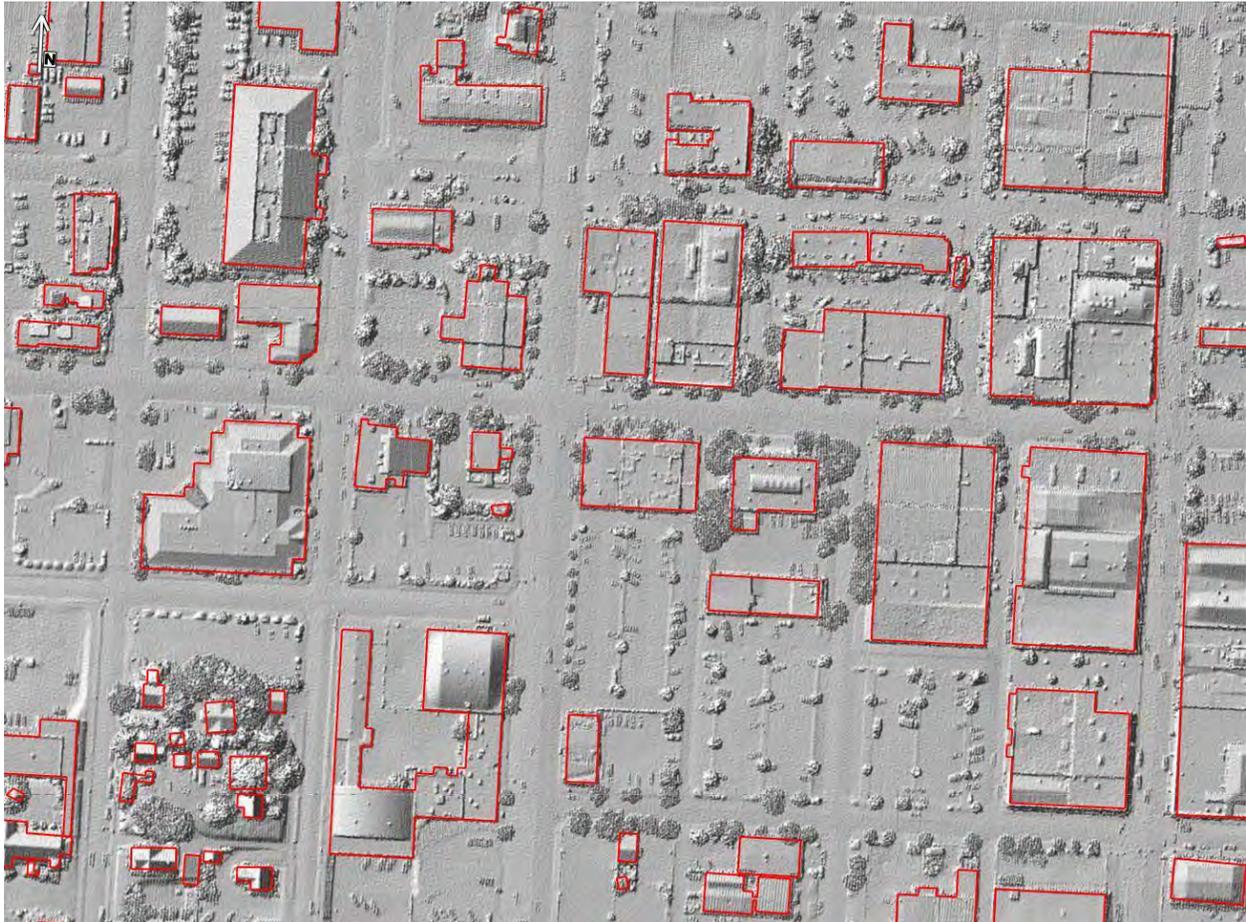


Figure 20. Coos Bay, Oregon. There are a total of 66 buildings in the above figure. AFE extracted 58 buildings correctly. A few trees were incorrectly extracted as houses. A few houses have incorrect boundaries due to overhanging trees. AFE achieved a success rate of 87%.

Figure 21 shows the Coos Bay City Hall. The left image shows the 3-D meshes in purple and rooftop polygon in red extracted by AFE. The right image is from Google Earth. Two small trees and plants on the middle, right side of the building cause ambiguities. Without clearer, higher resolution images, it is even difficult for a human operator to tell exactly where the rooftop boundary is on the middle, right side.



Figure 21. Coos Bay City Hall. On the right is an image captured from Google Earth at 43° 22' 04.52" N, 124° 13' 02.81" W. Purple lines are 3-D meshes. Red lines are the rooftop polygon.

Figure 22 shows the rooftop polygons extracted by AFE in Redmond, Oregon. We downloaded a small area where there are lots of buildings and houses from OpenTopography Facility, San Diego Supercomputer Center. The background image is terrain shaded relief generated from the same LiDAR point clouds.

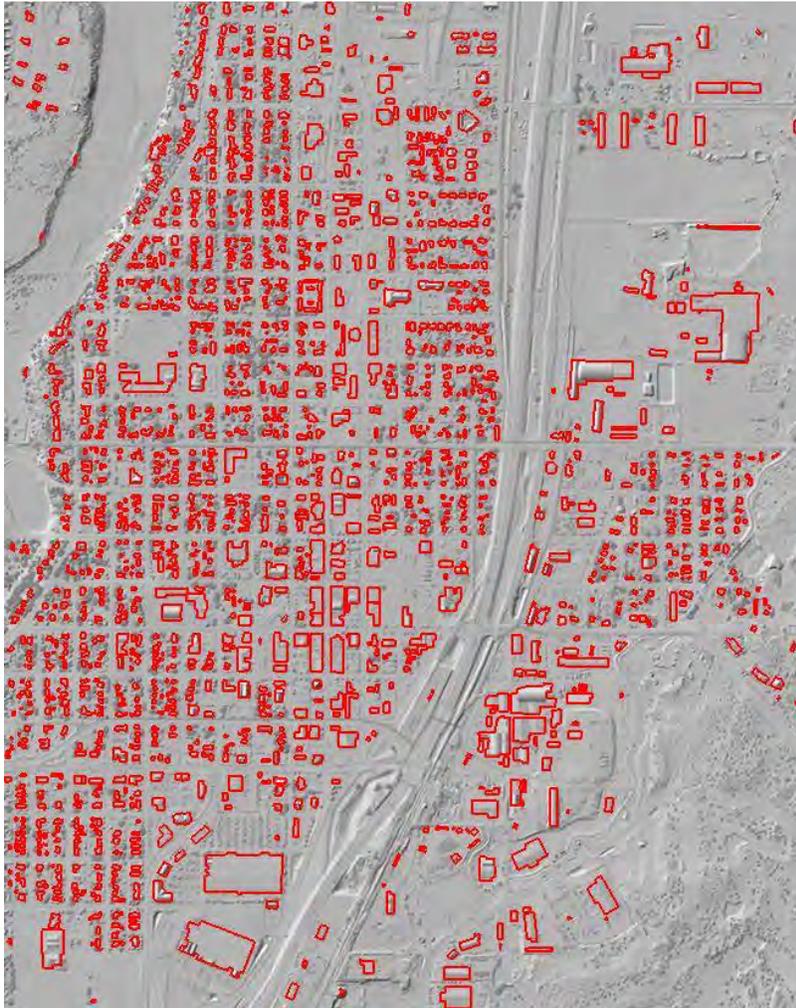


Figure 22. Redmond, Oregon. This LiDAR dataset has 39 million points covering an area of 4.9 km². The background image is terrain shaded relief generated from a 0.59 foot gridded DSM. 1780 houses and buildings were extracted by AFE with a success rate of 92%. Most of the incorrectly extracted houses and missing houses are due to overhanging trees.

Figure 23 shows a sub-region of Figure 22, randomly selected for accuracy analysis. The red polygons are rooftop polygons for houses and buildings. Since we do not have an automatic tool to evaluate AFE success rate, all the success rate measurements are done manually, which is a time consuming process, so we select a small enough region to count and analyze AFE accuracy.



Figure 23. Redmond, Oregon. There are a total of 122 buildings. AFE extracted 104 buildings correctly. A few trees were incorrectly extracted as houses. A few houses have incorrect boundaries due to overhanging trees. AFE achieved a success rate of 92%.

Figure 24 shows a 3-D view of the wireframes of buildings in Seaside, Oregon. Since no images are available, the 3-D view is not textured. The background image is terrain shaded relief generated from the same LiDAR point clouds. It is interesting to notice that a bridge is lowered on to the water in the TSR image. This is because a DEM is used to generate the 3-D view.

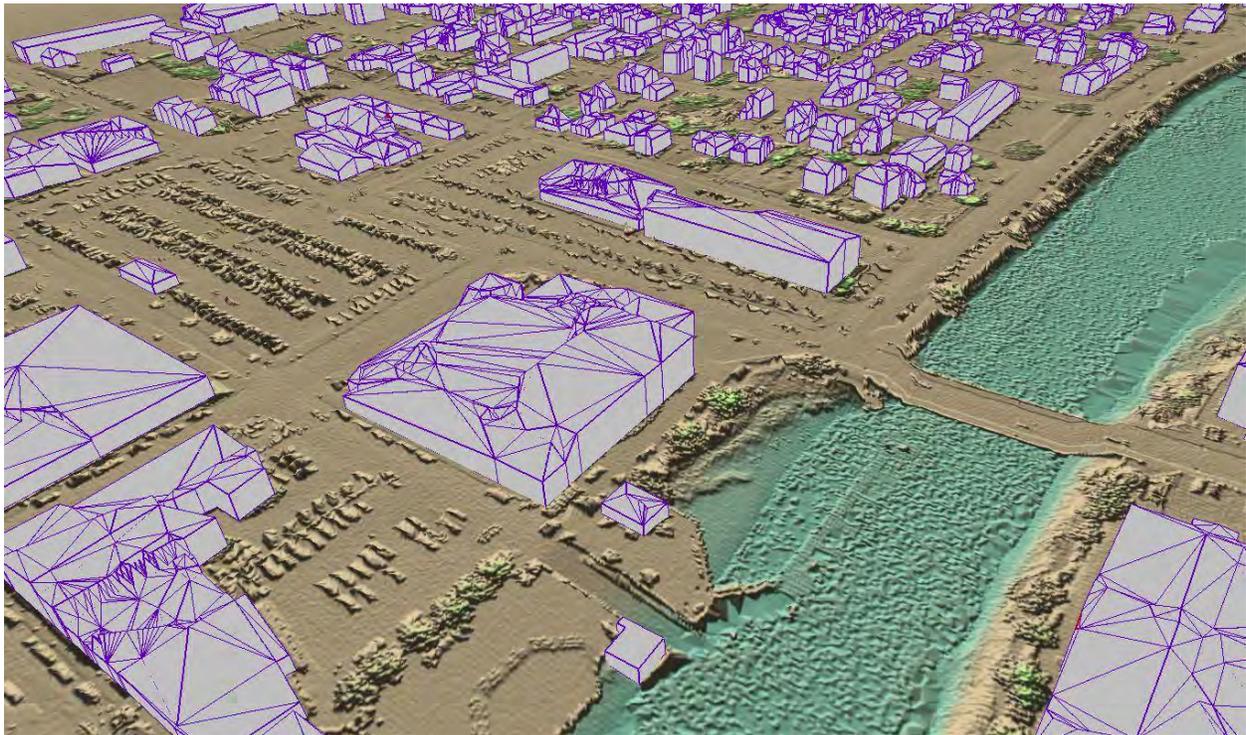


Figure 24. 3-D view of Seaside, Oregon. The building in the middle is the Seaside Civic and Convention Center. The building on the right is the Seaside Downtown Development Association. The building on the lower-left corner hosts Pizza Harbor and Beach Books. The background image is terrain shaded relief generated from LiDAR point clouds. Since images are not available, only the wireframes of the buildings are shown.

Figure 25 shows a fully textured 3-D view of a small portion of Bristol, U.K. The buildings are extracted by AFE with very dense and accurate LiDAR data, provided by Ordnance Survey. Stereo images covering the same area are accurately registered with the LiDAR point clouds. The whole process is fully automated in SOCET GXP v4.0 without any manual editing. When LiDAR point density reaches 25 or more points per m² with a precision of RMSE_z < 10 cm, it is likely that AFE can generate 3-D features accurate enough for most GIS applications.



Figure 25. Fully textured 3-D view of Bristol, U.K. The buildings were extracted by AFE using dense and accurate LiDAR point clouds. Stereo images and LiDAR point clouds were accurately registered. SOCET GXP Auto Texture was used to apply textures to the buildings. The whole process is automated without any manual editing.

Availability of AFE

AFE is included with SOCET GXP v4.0, BAE Systems' flagship image processing product that will include full photogrammetric functionality, scheduled for release in spring 2012. AFE has been thoroughly tested to process LiDAR point clouds in its initial release. In a future release, AFE will process DSMs generated by NGATE more effectively. Planned enhancements include the extraction of power lines, better tree extraction, enhancements to complex roof modeling, and donut types of buildings.

Conclusions

AFE is an innovative SOCET GXP module that can extract 3-D features automatically from 3-D point clouds. With multiple runs, AFE can achieve very high success rates, minimize manual editing and increase 3-D feature productivity. With accurate and dense LiDAR point clouds, AFE offers an automatic tool to produce vast amounts of 3-D features for 3-D GIS, 3-D visualization and robotic applications.

References

- OpenTopography Facility, undated. Oregon Department of Geology and Mineral Industries
LiDAR Program Data,
http://opentopo.sdsc.edu/gridsphere/gridsphere?gs_action=lidarDataset&cid=geonlidarframeportlet&opentopoID=OTLAS.022011.2994.1, last viewed 6 March 2012
- Zhang, B, and W. Smith, 2010. From where to what – image understanding through 3-D geometric shapes, ASPRS 2010 Annual Conference, San Diego, CA, unpaginated CD.
- Zhang, B. and M. Walter, 2009. Stereo image matching – breakthroughs in automated change detection, *Earth Imaging Journal*, 6(3), pp. 16-21.