ASPRS Competition: Unveiling Insights with Geiger-mode Lidar Data

Dale Fried
dale.fried@3deolidar.com

Trevor Mangum
trevor.mangum@3deolidar.com

October 2, 2024

1 Introduction

3DEO is proud to announce its partnership with The American Society for Photogrammetry and Remote Sensing (ASPRS) to sponsor the 2024 ASPRS competition. Developed primarily for the defense and security communities, Geiger-mode lidar has emerged as a powerful tool in the collection of high-resolution, three-dimensional point clouds. This cutting-edge technology offers high sensitivity and high data collection rates, resulting in affordable data products with high detail and point density. To increase awareness of the full potential of Geiger-mode lidar, 3DEO and APSRS are launching a competition designed to challenge university students and young professionals to push the boundaries of innovation.

Participants in the competition are tasked with developing novel algorithms, exploitation techniques, and applications that can reveal new insights and explore the possibilities embedded within Geiger-mode lidar data. This competition not only provides students with a unique opportunity to work with state-of-the-art data but also encourages the development of creative solutions that could have significant implications in fields such as environmental monitoring, urban planning, disaster management, forestry, archaeology, and beyond.

The remainder of this document serves to briefly explain what Geiger-mode lidar is, introduce the three datasets for the competition, and suggest a few useful tools that could be used to analyze the data. Specific details about the competition can be found elsewhere.

2 What is Geiger-mode Lidar?

Lidar, which stands for Light Detection and Ranging, is a remote sensing technology that uses laser pulses to measure distances and create detailed 3D maps of the environment. It works by emitting rapid pulses of laser light and measuring the time it takes for these pulses to bounce back after hitting objects. The name "Geiger-mode" refers to the physical process by which the lidar detects the faint light pulses received by the lidar from the ground. Some of the history of Geiger-mode lidar technology is referenced in a paper[1] describing the system that collected some of the data made available in this competition. As each individual photon is received, an electrical pulse is digitized and time-stamped, analogous to the clicks of a Geiger-counter detecting radioactivity. The "clicks" are processed into point clouds using the time of flight of the laser pulse. What is different from more common "linear mode" lidar is that a point in the product point cloud is typically established using only 5 - 10 "clicks" of the GM photodetector, arising from 15 - 30 photons received by the lidar. In contrast, a typical linear mode system needs a minimum of approximately 300 received photons to reliably differentiate a real surface from the noise in the system. Because of this 10× lower light requirement, for a given average laser power and aperture GM lidar systems can collect data at high rates from higher altitudes, simplifying operations in rugged terrain and collecting wider swaths on the ground.

Geiger-mode (GM) lidar technology has been operationally proven by the US military since 2010[2], and is seeing renewed interest in the commercial geospatial world. The next-generation systems from 3DEO include the core GM technology advantages of fast measurement rates and exquisite sensitivity, and also embody significant advances over previous generations. A primary innovation is 3DEO's patent-pending agile geo-referenced scanning, which directs the full capability of the lidar into only specific areas of interest, such as a narrow winding corridor or a campus, and which enables high angular diversity to mitigate shadowing. In addition, streamlined data processing workflows enable raw data from many different scans to be aggregated so that weak signals from dark or significantly occluded surfaces can be turned into deliverable data products with high information utility. All combined, these innovations allow lidar operators to collect point clouds with data densities commonly associated with drone lidar collections, but at a scale associated with larger aircraft. This large-scale, high-density data is a critical input for applications such as wildfire prediction, forestry, urban mapping, and disaster response.

3DEO's commercial GM lidars typically use a 32x128 array of these sensitive photodetectors, all operating in parallel on every laser pulse. More advanced GM lidars utilize multiple arrays and/or larger arrays[3]enabling even higher area collection rates. 3DEO GM lidar sensors typically yield millions of interrogations per second, with a range resolution of 15 - 20 cm, and angular resolution of 76 μrad; from an altitude of 3,000 m that angular resolution corresponds to a ground spatial resolution of about 23 cm.

Essential to making these high point densities useful is collecting the point clouds from many different viewpoints to overcome the common shortcoming of shadowing and occlusions in lidar products. As mentioned above, 3DEO has developed an agile geo-referenced scanning method[4] that provides the flexibility required to achieve high angular diversity. This diversity

---

[1] Call, B. R., et. al., "Low SWaP, commercially-available Geiger-mode lidar system," in Laser Radar Technology and Applications XXVII, 12110, 58–68, SPIE (2022).

[2] M. J. Khan et al., "Remote Sensing Impact of Single-Photon Sensitive Airborne Lidar Systems Based on Geiger-Mode Avalanche Photodiode Arrays," IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium, Athens, Greece, 2024, pp. 2455-2459, doi: 10.1109/IGARSS53475.2024.10641168

[3] *ibid.*

[4] Brandon R. Call, Dale G. Fried, David Kelley, Christopher Reichert, "Dynamic geo-referenced scanning of aerial lidar systems," Proc. SPIE 12110, Laser Radar Technology and Applications XXVII, 121100A (3

is crucial in highly foliated scenes to "poke through" the foliage to the ground layer[5]. The combined laser and camera optical paths may be scanned anywhere in the 40 deg × 40 deg field of regard of the system. In a mapping style collection, the large area to be mapped is partitioned into multiple swaths, each with sufficient overlap. Each swath is broken up into polygons of several hundred meters size, and each polygon is scanned multiple times over the course of the flight path. In total, a given polygon might be scanned many tens of times, from many tens of different viewing geometries. By engineering the polygons appropriately, the amount of angular diversity per pass can be matched to the needs of the project, from just a few scans to many tens of scans. For foliage penetration and urban mapping, typically high angular diversity is chosen, whereas for wide area mapping of unforested scenes only a single scan of each polygon is needed. A video showing the operation of 3DEO's Agile Geo-Reference Scanning[6] explains further.

To make use of the high angular diversity from many views, 3DEO has developed a processing workflow and algorithms to coalesce all the information into one representation of the scene. The raw sensor data is extracted to create a raw 3D point cloud for each individual scan. The scans are aligned, combined, and then processed together in order to find surfaces. Some surfaces may have had lidar echos that were too weak to detect from any single scan, but which become noticeable when many scans are combined. The level of detail in the scenes of the aggregated point cloud can be compared to point densities obtained with traditional drone lidar but obtained at a much higher area collection rate. From the aggregated point clouds, standard derived products may be obtained such as bare earth digital terrain models (DTMs), digital surface models (DSMs), and point clouds annotated with height above ground (HAG) for each point.

The common terms used in the Geiger-mode world for these various stages of processing are:

- **L0** Level Zero data is the raw data written to disk during the data collection flight.
- **L1** Level One data is a 3D point cloud created from all the L0 data for a given scan, and includes 3D points for each photodiode firing event ("click").
- **L3** Level Three data is a point cloud created after combining many scans together to determine the location of real surfaces within the scene. Noise is largely eliminated. Multiple photodetection events for a common surface are typically coalesced into one or a few points which may typically be more accurately placed in 3D space due to averaging the L0 points.
- **L4** Level Four data includes raster and vector products generated from the L3 point clouds.

June 2022); `https://doi.org/10.1117/12.2619105`

[5] Brandon R. Call, Dale G. Fried, David B. Kelley, Kimberly Reichel-Vischi, Christopher Reichert, "Assessment of foliage poke-through capabilities of an airborne Geiger-mode lidar," Proc. SPIE 13049, Laser Radar Technology and Applications XXIX, 130490I (5 June 2024); `https://doi.org/10.1117/12.3025035`

[6] `https://www.youtube.com/watch?v=9saTWdozGsc`

3 The Datasets

For the competition, 3DEO is providing L3 files for three datasets of varying scene content. Each of these datasets was collected using 3DEO sensors and agile scanning capability, as described above. In addition to the mapping style collections described above common in the lidar mapping world, two of the datasets were collected in a mode where the entirety of the scanning is directed into just a single desired area of interest, which we call "Target Mode". In "Target Mode", the area of interest is typically smaller than the projection of the system field of regard onto the ground, allowing the entire target to be scanned many times over the course of a single flight line. Multiple flight lines can be flown, with the aircraft situated on various sides of the target area, in order to increase angular diversity.

The data processing used to create these data products relied upon a 3D voxelization scheme. L1 points from each scan are binned into a 3D data structure (voxel space). Voxels with significant occupation and that exhibit a peak in the local point density are written to the L3 output file, along with metadata fields that can be useful in interpreting the points. All the provided datasets were processed with voxel dimensions of 10 cm in the x and y directions, and 15 cm in the z direction.

The number of L1 points placed in each voxel is recorded in a metadata field called "Hits". A high number of hits in a voxel typically indicates a high confidence that a real surface is present, and not just noise. The "Hits" field is a floating point number, not an integer, due to a convolution step conducted after the aggregation. To view just the most significant features in the dataset and eliminate noise points, one typically filters the data on the "Hits" metadata field to include only the larger values. On the other hand, to see weaker features, such as faint wires, one typically filters on a lower value of "Hits". Of course, more volume noise will be kept in the point cloud if the "Hits" threshold is lower.

Information about which scans contributed L1 points to a given voxel is encoded in a metadata field called "Bitfield_Scans" in the LAZ files and "bitfield:Scans" in the BPF files. Each scan corresponds to one bit in the field. A value of 1 in the bit position for a given scan indicates that the scan contributed at least one L1 point to the voxel. This metadata field is experimental and may not be useful in this competition.

Once the voxel data structure is populated, peak finding algorithms are used to identify likely surfaces by identifying peaks in the local L1 data density. Information about the direction along which a peak was found is encoded in the metadata field called "Bitfield_PeakAxis" in the LAZ files and "bitfield:PeakAxis" in the BPF files. Peaks are sought in each of three orthogonal directions. A 1 in bit position 0 (value of $2^0 = 1$) indicates a peak was found when searching for peaks along the x-axis; bit 1 (value of 2^1) indicates the y-axis, and bit 3 (value of $2^2 = 4$) indicates the z-axis. A strong reflection from a single isolated point will likely register as a peak in all three directions, so the value of "Bitfield_PeakAxis" will be $2^0 + 2^1 + 2^2 = 7$. A horizontal surface will likely register as a peak in the z-direction. Cleaner imagery in a point cloud viewer often results when one filters the data to show only points that had a peak in the vertical direction, meaning "Bitfield_PeakAxis" values of 4 - 7. Vertical surfaces such as tree stems will likely have a peak in a horizontal axis and z-axis, corresponding to values of 1 - 3 and 5 - 7, respectively.

Each of the datasets are provided in both LAZ version 1.4 format and BPF version 3 format. The Binary Point File (BPF) is a standard maintained by the National Geospatial-

Intelligence Agency (NGA)[7]. 3DEO's current processing workflow performs all of the data processing in BPF format and converts the L3 BPF product into a LAS/LAZ data product at the conclusion of processing. Many of the standard channels in the LAS format are not directly applicable to Geiger-mode lidar. As such, when the BPF is converted to the LAS format, many of the standard LAS channels remain unused. The unused fields in the LAS conversion get populated with zeros. For these datasets, if a standard LAS channel is not explicitly referenced in this section, assume the channel is unused. The LAS files also contain additional custom channels created during 3DEO's standard processing, as described above. The L3 files contain data channels for the x, y, and z coordinates. These channels correspond to the (x, y, z) coordinate for each point in the UTM coordinate system. The horizontal datum for each dataset is a chosen UTM zone. The vertical datum is WGS84.

The provided datasets populate the LAZ intensity channel with values that correspond to the "Hits" meta channel. Note that the number of L1 points ("Hits") is not a measure of intensity since it has not been normalized by the amount of laser energy incident on the surface. The intensity field in the LAS file is derived from the "Hits" value and is intended solely to provide a more aesthetic presentation. There is no additional information in the intensity field beyond the information in the "Hits" field.

Point density A common figure of merit for lidar data products is the point density. Typical "linear mode" lidars scan an area at a given average number of laser pulses per square meter (PPSM). Each laser pulse results in zero, one, or more measurements of the distance from the lidar to surfaces in the scene. Very dark surfaces might result in no measurement for that pulse. Complex tree canopy, through which the laser beam might hit multiple surfaces, can result in multiple measurements. Geiger-mode lidar systems operate differently, and so the data density metrics are somewhat different. Geiger-mode systems detect and time-stamp individual photons. In data processing all of the information about the photons detected in a given voxel is combined, using observations from typically hundreds of laser pulses and multiple scans. The user of the data product is typically interested in the density of meaningful, unique information about the surfaces in the scene. To be "meaningful" the data point provided to the user must have a high confidence of representing a real surface, or put another way, it must have a very low likelihood of being noise. The presence of several detected photons ("Hits") in the same voxel gives confidence that a real surface is present. To express "unique information" in each point in the output data product the processing averages together all the raw points in a voxel and reports just a single point. Several research groups have proposed and implemented various approaches to this problem of finding and expressing meaningful and unique information in point clouds derived from Geiger-mode lidar data, and it is still an area of active research. For the data sets provided for this ASPRS competition, instead of stating a point density we instead provide curves that address the density of meaningful, unique information. The curves are the cumulative distribution function of the "Hits" field, as shown in figures 1, 4, and 7. The plots show the number of L3 points in the point cloud with "Hits" equal to or greater than the value in the abscissa. For example, in the Arboretum data set there are 3×10^7 points with at least 8 "Hits". Averaged over the entire area of the point cloud of 130,000 m^2, those points

---

[7] `https://nsgreg.nga.mil/doc/view?i=4220&month=8&day=30&year=2016`

correspond to an average point density of 230 points per square meter. The datasets are provided with many points that have lower “Hits”, causing them to initially look noisy before filtering the data on the “Hits” channel. The lower confidence points (lower “Hits” value) are included in the provided data in order to preserve weaker signals from, for example, partially occluded tree trunks. The data can be filtered on the “Hits” meta channel if a cleaner looking point cloud that does not contain as much signal is desired.

Dataset	**Sensor**	**Collection Altitude (ft AGL)**	**Angular Resolution (μrad)**	**Area Collected (km^2)**	**Number of Full Scans**	**Collection Mode**	**Total Recording Time (s)**
Arboretum	Zion-A	4,000	130	0.13	24	Target	89
Lawrence	Zion-B	8,500	76	1.0	13	Map	120
Leominster	Zion-B	6,500	76	0.25	44	Target	145

Table 1: Summary of each of the datasets with sensor information, collection altitude, angular resolution, area collected, number of scans, collection mode employed, and total recording time.

3.1 University of Idaho Arboretum

The Arboretum dataset was collected over the Charles Houston Shattuck Arboretum on the University of Idaho’s campus. The Arboretum dataset was collected at an altitude of 4,000 feet AGL, which corresponds to a ground spatial resolution of 17 cm with 3DEO’s first generation sensor, Zion-A. This dataset is derived from 24 scans of the entire collection area and an additional 15 partial scans of the area. All of these complete and partial scans were aggregated into the L3 product. The distribution of L1 “Hits” within the L3 dataset is shown in figure 1.

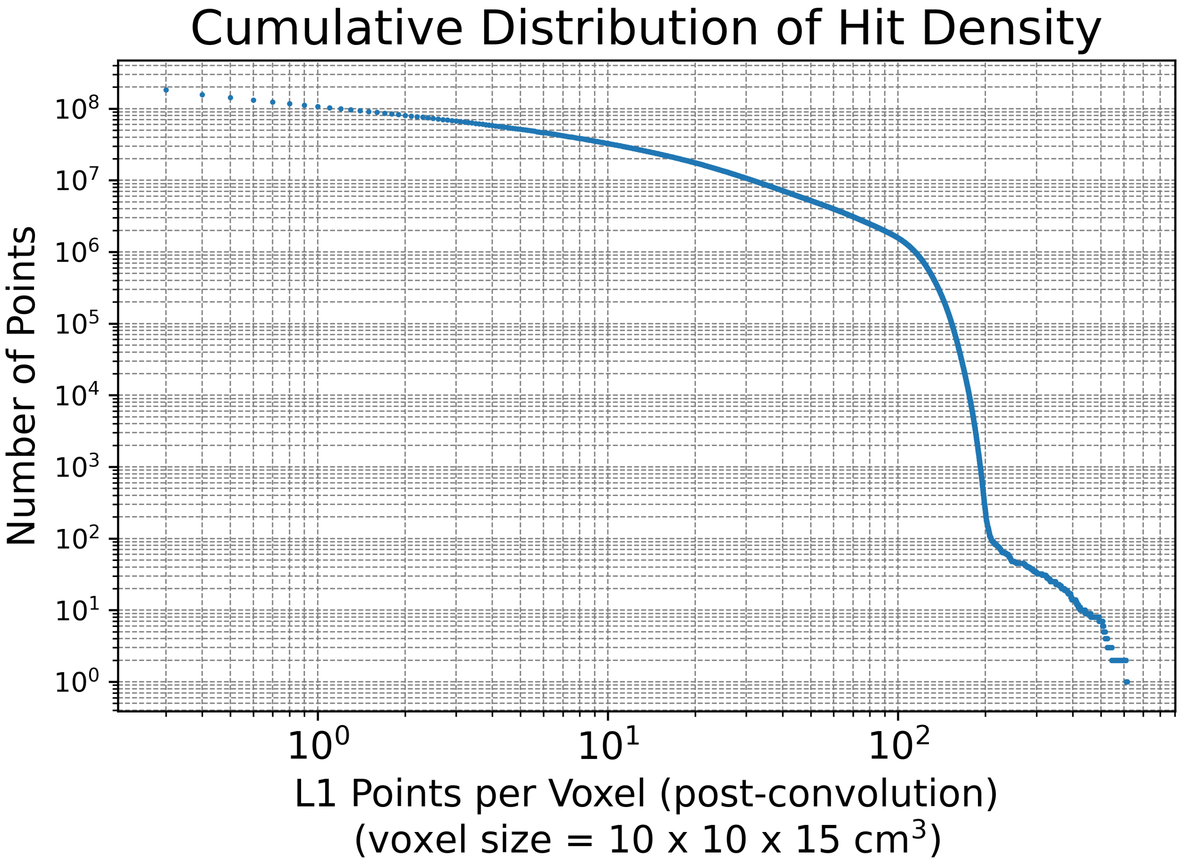

Figure 1: Cumulative distribution of L1 density in the Arboretum dataset. The abscissa is the "Hits" metadata field: the number of L1 points per voxel smoothed by a convolution function. "Hits" is indicative of the relative confidence of real surface detection. The voxel size is $10 \times 10 \times 15$ cm^3. The ordinate is the number of L3 points in the point cloud meeting or exceeding the "Hits" threshold. A point may be considered to have high confidence if the "Hits" count exceeds approximately 10. There are approximately 3×10^7 such points in this point cloud.

The Arboretum dataset is provided with many points that have lower "Hits" counts, causing it to initially look noisy before filtering the data on the "Hits" channel. The lower confidence points are included in the provided data in order to preserve weaker signals from, for example, partially occluded tree trunks. The data can be filtered on the "Hits" meta channel if a cleaner looking point cloud is desired, showing only points with higher confidence. The area encompassed by this dataset is around 130,000 m^2. The LAZ version of the Arboretum dataset is 1.31 Gigabytes and the BPF version is 4.06 Gigabytes. A variety of scene content can be seen in the Arboretum dataset including roads, buildings, light poles, a water tower, and a highly documented forest containing a wide variety of trees.

Figure 2: The bounding polygon of the target area for the Arboretum dataset.

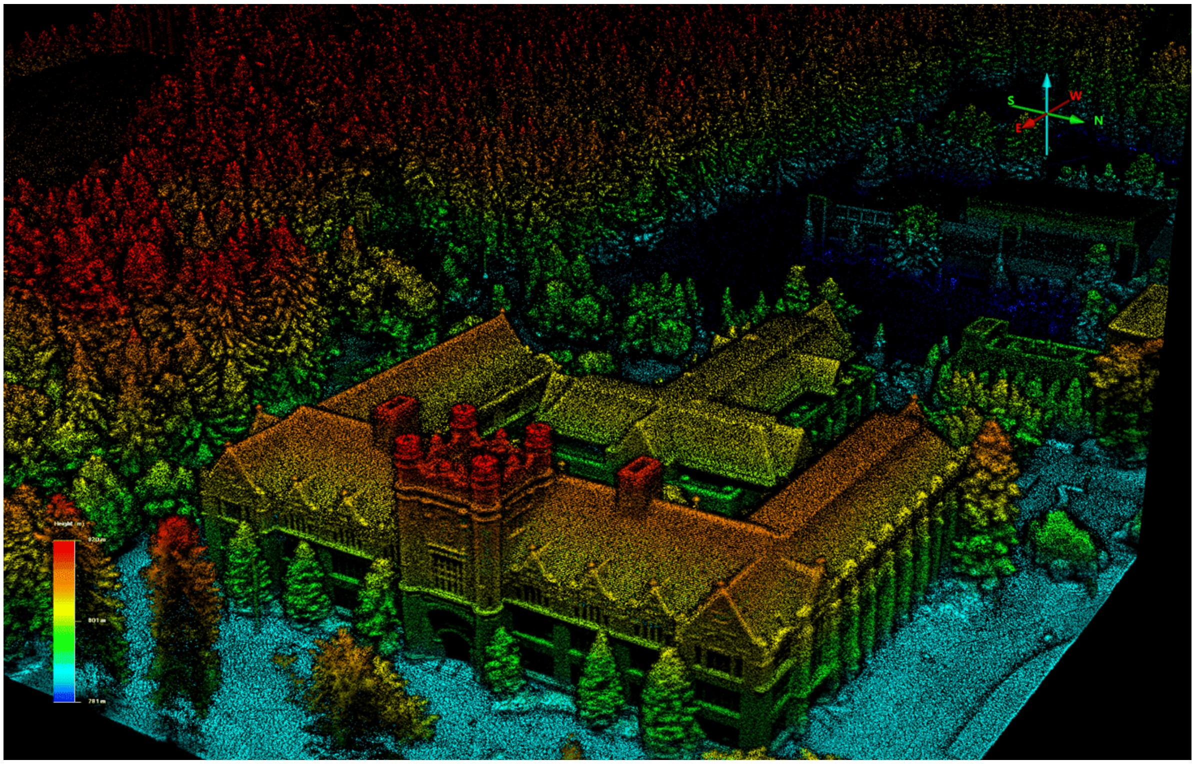

Figure 3: A screenshot of the Arboretum dataset point cloud, colored by z from 781 m to 820 m. Highlighted is the UI Administration building.

3.2 Lawrence, Massachusetts

The Lawrence dataset captures the heart of an old New England mill town along the Merrimack river, including canals and old mill buildings. The data was collected at an altitude of 8,500 feet AGL, which corresponds to a ground spatial resolution of 21 cm with 3DEO's second generation sensor, Zion-B. The city was collected in map mode. This dataset is the aggregation of 13 scans of the entire collection area and an additional 34 partial scans of the area. The distribution of L1 "Hits" within the L3 product is shown in figure 4.

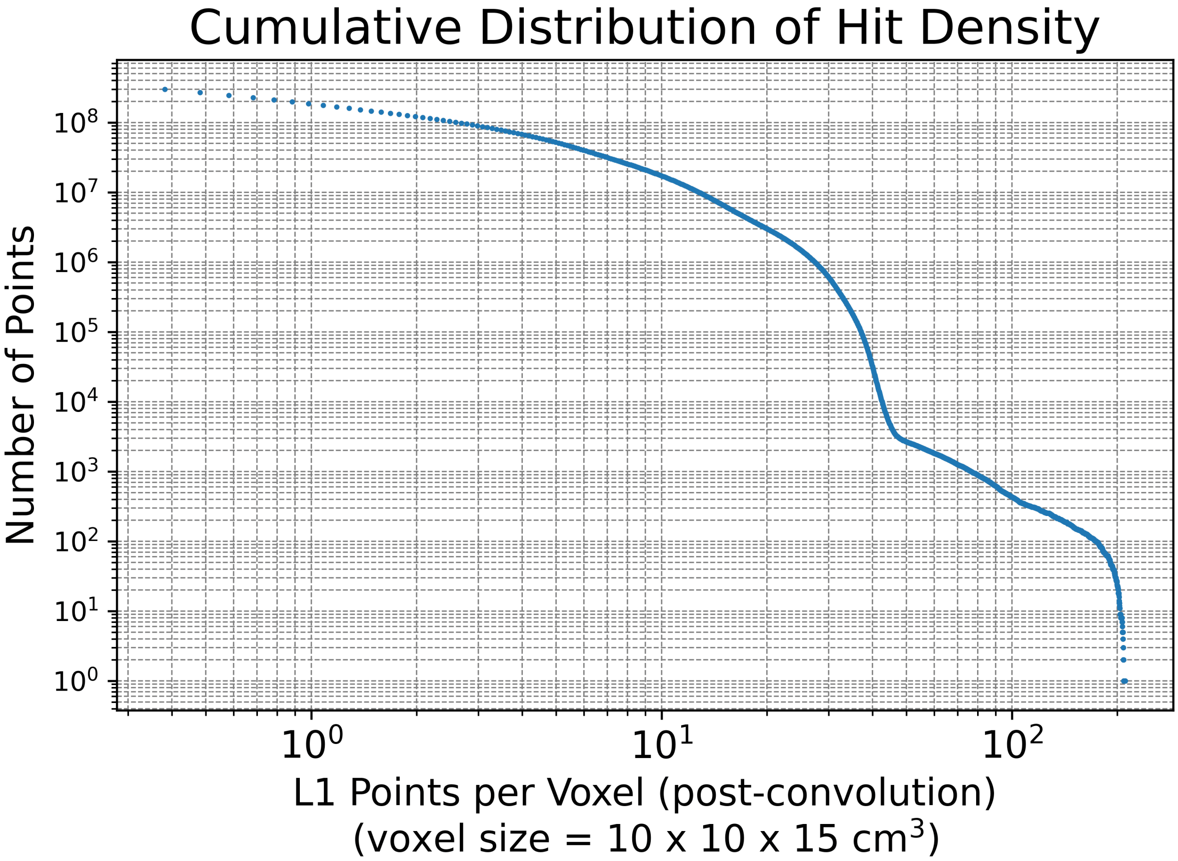

Figure 4: Cumulative distribution of L1 density in the Lawrence dataset. The caption of Figure 1 highlights additional information for interpreting this figure. There are approximately 2×10^7 points in this point cloud above the "Hits" threshold of 10.

The area included in this dataset is around 1 km^2. The LAZ version of the Lawrence dataset is 2.09 Gigabytes and the BPF version is 6.65 Gigabytes. A variety of scene content can be found in the Lawrence data including roads, bridges, powerlines, cars, and typical urban scenery. It should also be noted that a flock of birds (likely crows) appears to have been present during several scans in the center-east section of the dataset. Large flock of crows are known to roost in the area in the winter[8].

---

[8] https://www.wgbh.org/news/local/2024-03-27/thousands-of-crows-in-lawrence-are-a-marvel-of-nature-and-test-case-for-new-imaging-technology

Figure 5: The bounding polygon of the target area for the Lawrence dataset.

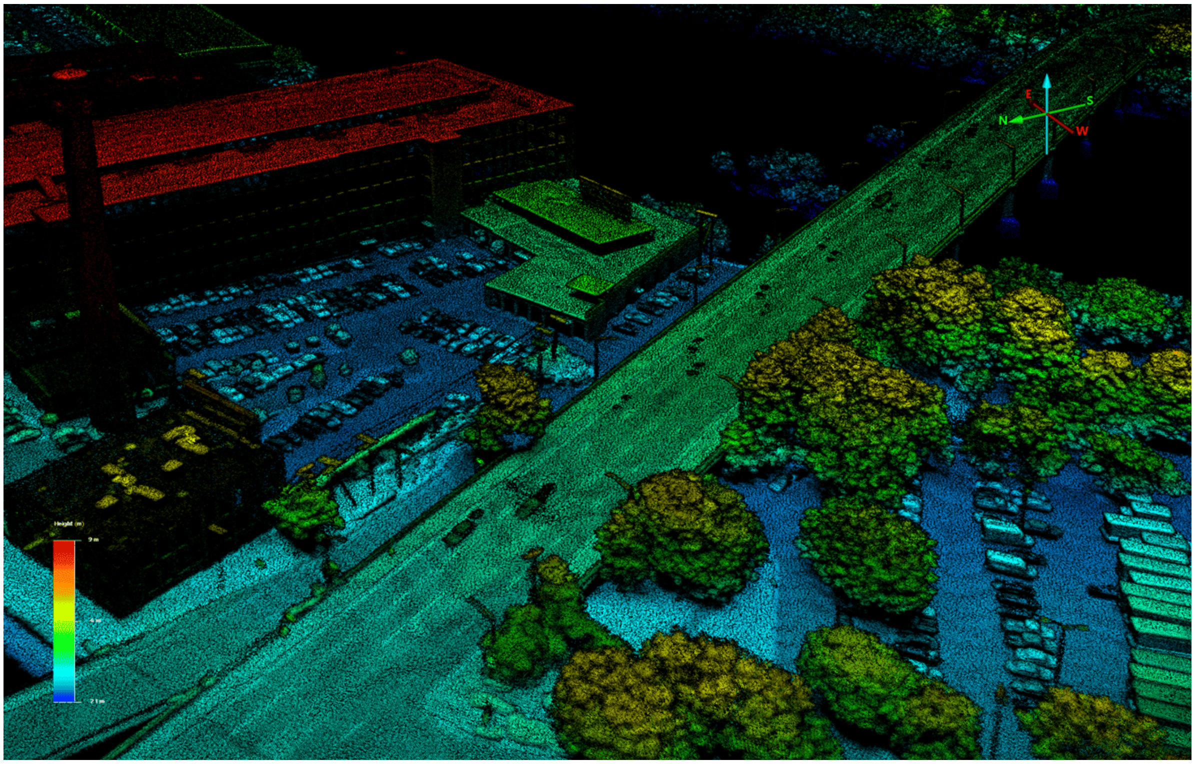

Figure 6: A screenshot of the Lawrence dataset point cloud, colored by z from -21 m to 9 m. Highlighted is the bridge on the western side of the dataset.

3.3 Leominster, Massachusetts

The Leominster dataset was collected shortly after an intense rainstorm and subsequent over-topping event at Barrett Park in Leominster, MA. Significant damage was sustained by the dam and downstream structures. The data was collected from an altitude of 6,500 feet AGL, which corresponds to a ground spatial resolution of 16 cm with 3DEO's second generation sensor Zion-B. This dataset includes 44 scans of the entire collection area and an additional 55 partial scans of the area. The distribution of L1 "Hits" in the L3 product is shownn in figure 7.

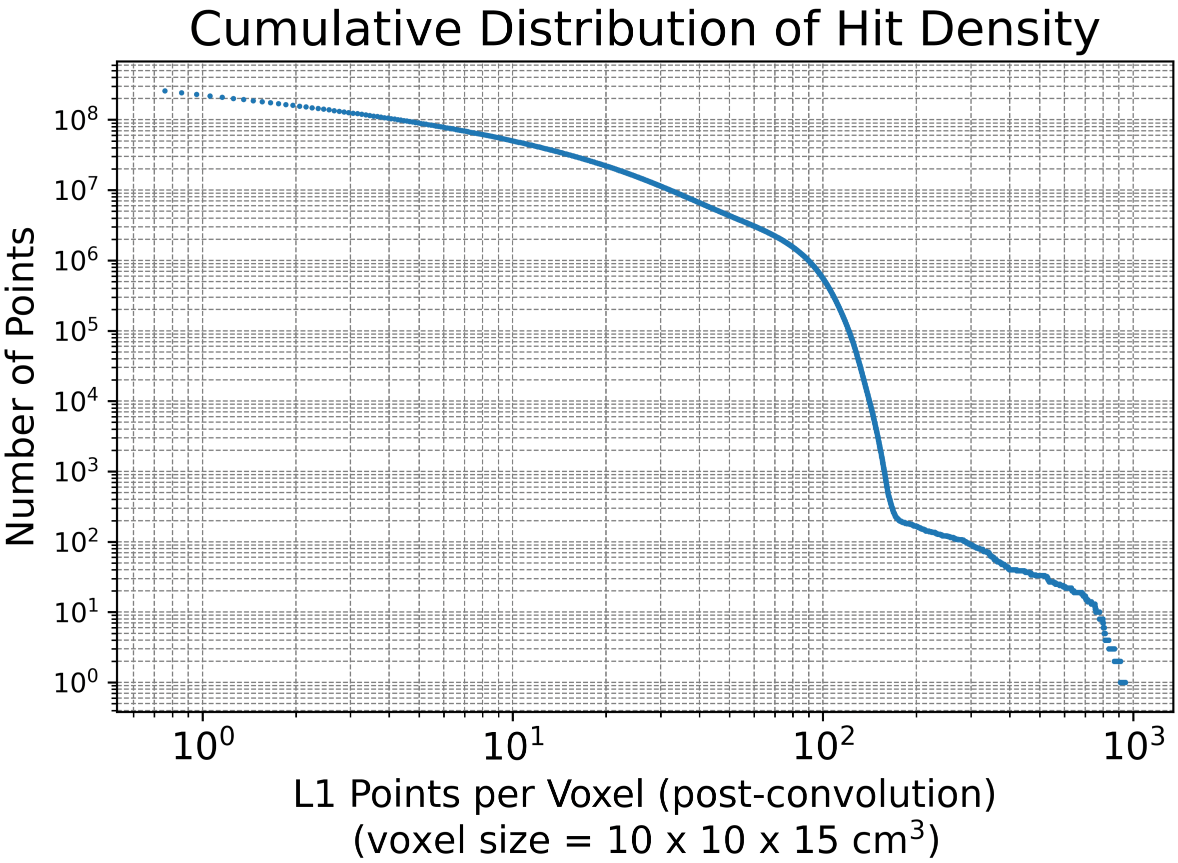

Figure 7: Cumulative distribution of L1 density in the Leominster dataset. The caption of Figure 1 highlights additional information for interpreting this figure. There are approximately 5×10^7 points in this point cloud above the "Hits" threshold of 10.

The area encompassed in this dataset is around 250,000 m^2. The LAZ version of the Leominster dataset is 1.57 Gigabytes and the BPF version is 5.74 Gigabytes. A variety of scene content can be seen in the Leominster dataset including foliage, powerlines, many buildings, and a lake.

Figure 8: The bounding polygon of the target area for the Leominster dataset.

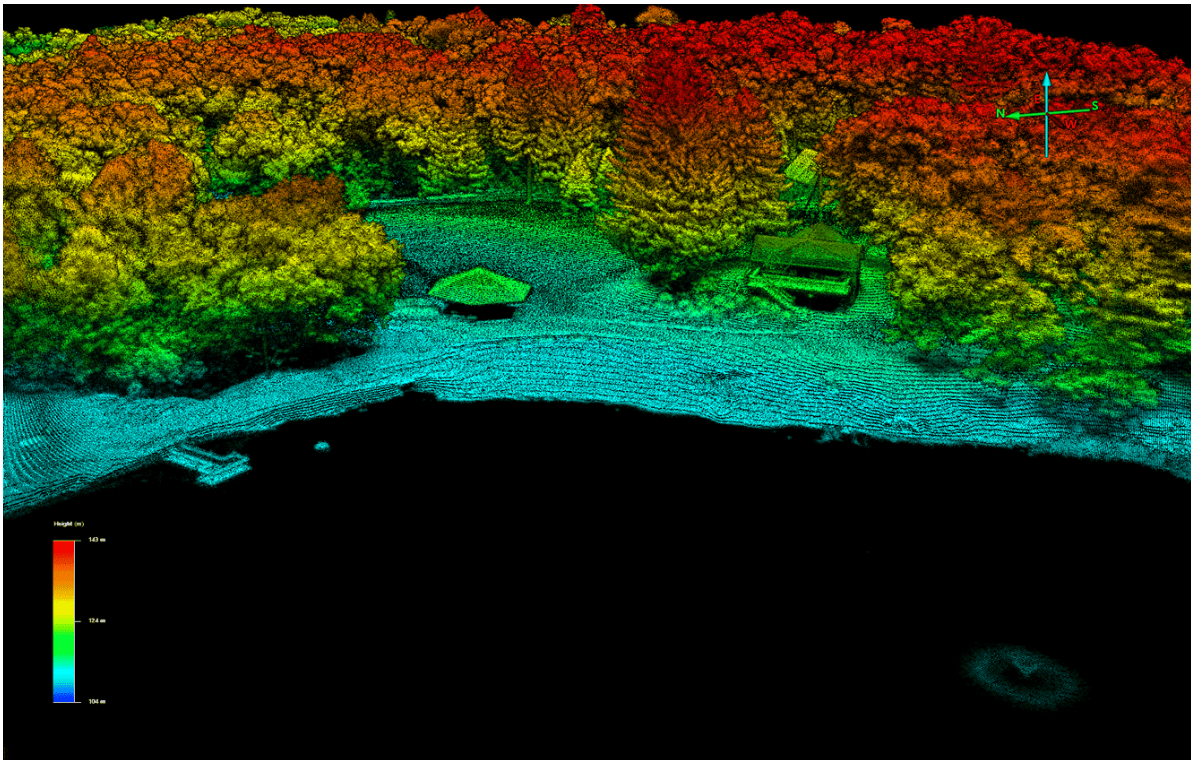

Figure 9: A screenshot of the Leominster dataset point cloud, colored by z from 104 m to 143 m. Highlighted is the gazebo near the southeast corner of the pond.

4 Tools for Analyzing the Datasets

4.1 Software for Reading Dataset

PDAL (Point Data Abstraction Library) is a powerful tool for reading and processing point clouds. PDAL provides native support for both LAZ and BPF, enabling users to easily load and manipulate point cloud data through customizable pipelines. For example, a simple pipeline can read a BPF file and extract or filter points based on spatial or attribute constraints.

Beyond reading, PDAL offers advanced functionality such as point classification, clustering, noise removal, geometric transformations, and coordinate reprojection. These features make PDAL an essential tool for efficiently managing, analyzing, and processing point cloud data. PDAL offers a command line interface as well as Python bindings. A variety of PDAL tutorials can be found online[9].

The laspy[10] Python library is also a useful tool that can be used to natively read LAZ files. Although BPF and LAZ are not natively supported, the Point Cloud Library (PCL)[11] offers a variety of useful tools for point cloud manipulation and analysis. A variety of Python or C++ readers could be used in conjunction with the PCL for possible analysis of the provided datasets.

4.2 Compatible Data Viewers

QT Modeler, QT Reader, and Global Mapper Pro offer native support of the BPF and LAZ datasets. CloudCompare, ArcGis, and a variety of other viewers can be used to view and analyze the LAZ format of the datasets.

---

[9] `https://pdal.io/en/2.7.2/tutorial/index.html`

[10] `https://laspy.readthedocs.io/en/latest/`

[11] `https://pointclouds.org/documentation/`