




A HIGHLY ACCURATE, DIVERSE AND DENSE ANNOTATED LASER SCANNING DATASET OF A CITY CENTRE

Downtown Dublin as a Lidar point cloud



A dense annotated ground-truth Lidar dataset of the city of Dublin is the first of its kind regarding the accuracy, density and diversity of classes. The hierarchical labels offer excellent potential for various applications in urban science.

The advance of Lidar data acquisition technologies is substantially increasing the amount of spatial data obtained and it is becoming cost-prohibitive to process it manually. Artificial intelligence (AI) has now started to offer cost-effective solutions to analyse and utilize those big datasets. AI can be employed for scene understanding, accurately detecting objects and classifying 3D assets. These are the significant fundamentals of several applications including autonomous navigation, intelligent robotics, urban planning, emergency management and even forest monitoring. The most popular types of 3D spatial information for such applications are Lidar and imagery data. This article describes the dense annotated ground-truth Lidar dataset that was generated in 2019. The labelled dataset was produced from Lidar data of Dublin that was captured alongside aerial images of Dublin in 2015. Both datasets are publicly available and their URLs are included in this article.



With two thirds of the world's population already living in urban areas, and a further increase of two billion predicted by 2050, the number of megacities (i.e. with populations of more than ten million) is predicted to increase to 41 in the next decade (UN 2014). Most of these cities had less than 2 to 3 million inhabitants in 1950, which means that their infrastructure is totally unable to support such huge growth. To plan sustainable growth of such urban areas, geometrically accurate three-dimensional (3D) models are essential for city planning. Accurate spatial modelling and interventions related to city planning are especially challenging as most parts of cities are largely undocumented, and the cost of collecting the relevant data through traditional mapping methods is typically very high. In such cases, remote-sensing technologies offer cost-effective alternatives.

Lidar laser scanning and photogrammetry are outstanding solutions amongst the current technologies. The latest Lidar scanners are able to capture around one million georeferenced points a second in the form of a point cloud. This point cloud can be acquired via three major sources: a) Terrestrial laser scanning (TLS), which is able to collect most of the vertical facade data but little of any roofs, balconies and other horizontal planes, usually from the street point of view; b) Mobile laser scanning (MLS), which is obtained by cars or other vehicles (e.g. autonomous patrolling or even boats) to survey relatively short distances up to 300m with the same characteristics of TLS; and c) [Airborne laser scanning](#) (ALS), which can contain full-waveform data with a good bird's-eye view but usually limited facade data. ALS is

generally used for obtaining data for a large area (e.g. an entire urban region).

The Dublin datasets

In 2007 and later in 2015, the Urban Modelling Group at University College Dublin captured a massive urban dataset of the Republic of Ireland's capital city, Dublin, under supervision of Prof Debra Laefer. The datasets include laser scanning point cloud (i.e. Lidar) – as shown in Figure 1 – as well as aerial imagery (i.e. vertical, oblique images and video data). As the footage in the overview of the dataset shows, the density of the dataset was hugely improved in 2015 compared to the initial capturing in 2007. The 2015 dataset is one of the densest urban aerial Lidar point clouds that has ever been collected (over 1.4 billion points) with an average point density of 250 to 348 points/m². In this project, the initial dataset consisted of 5.6km² of Dublin's city centre which was scanned with an ALS device carried out by a helicopter at an average flying altitude of 300m. The data was collected in March 2015, as there is usually minimum vegetation and hence shadows on the buildings at that time in Dublin.



Figure 1: Laser scanning point cloud of Dublin, Ireland's capital city.

While the primary output of the dataset was to generate a Lidar point cloud, the mounted cameras also captured imagery data during the flight. The imagery dataset consists of 4,471 images as georeferenced RGB images with a resolution of 9,000x6,732 pixels with a ground sampling distance of 3.4cm in TIFF format. The geographic information is given as GPS information in the EXIF metadata and the camera used for the capture was Leica RCD30. The dataset also includes 4,033 oblique JPEG images with a resolution of 7,360x4,912 that were captured by two [Nikon D800E](#) cameras. The total size of the imagery dataset is around 830GB. All Lidar and imagery data can be accessed at the [NYU data repository](#).

In addition to that dataset, in 2017 Dr Jonathan Byrne et. al. also captured aerial images of Trinity College Dublin (TCD) campus at an average altitude of around 30m by drone and generated an image-based point cloud. The dataset would be interesting for comparison of the campus from 2015 to 2017.

The reasons for annotation

The laser scanning data has no inherent classification information in the point cloud or any predefined relationships between its points. Consequently, in order to use these points in various applications, the data must be classified and the desired localized features must be extracted. Despite advances in the technology for accurately and rapidly capturing Lidar data on a large scale, automated analysis and understanding of the huge datasets obtained is still being developed.



Figure 2: A sample labelled tree that shows its trunk, branches and leaves have been precisely scanned.

The traditional methods of dataset processing usually use geometric fitting (e.g. [RANSAC](#)) or Region growing. These methods are for coarse segmentation as an initial step of classification of the dataset (e.g. buildings, streets, etc.). Coarse segmentation is typically followed by explicit feature extraction based on algorithms developed to extract smaller objects (e.g. windows, doors or chimneys) for specific applications. Despite the above-mentioned available techniques, machine learning (ML) and AI appear to be more efficient and cost-effective approaches. For example, neural networks (NN) can be trained to intelligently detect and classify 3D assets in the dataset with minimum manual intervention. However, the key point for training NN models is having an accurate, diverse and well-annotated ground-truth dataset. Therefore, it is significantly important to access a full-3D, dense and non-synthetic labelled point cloud at city scale that includes a variety of urban elements (various types of roofs, building facades, windows, trees and pavements). However, the generation of such detailed labelled datasets is difficult and expensive. While there have been several attempts to generate such a labelled dataset, including by using semi-automated photogrammetric or morphological methods, a review of the commonly available datasets shows that none of them can completely satisfy all requirements.

A novel annotated massive point cloud

Iman Zolanvari et. al. provided a novel labelled dataset (Figure 3) from the above-mentioned Lidar data of Dublin at Trinity College Dublin. In this project, over 260 million laser scanning points were manually labelled into 100,000 objects within 13 classes. Those classes included a hierarchical level of detail, from coarse (i.e. buildings, vegetation and ground) to a refined level (e.g. windows, doors and trees).



Figure 3: A sample portion of the annotated dataset.

The first level produced a coarse labelling that includes four classes: Building, Ground, Vegetation and Undefined. 'Building' refers to all shapes of habitable urban structures (e.g. homes, offices, schools and libraries). 'Ground' mostly contains points that are at the ground level. The 'Vegetation' class consists of all types of separable plants. Lastly, 'Undefined' points are those of least interest to include, such as urban elements (e.g. bins, decorative sculptures, cars, benches, poles, post boxes and non-static objects). Approximately 10% of the total points were labelled as undefined and they were mostly rivers, railways and construction sites. In the second level, the first three categories from Level one were divided into a series of refined classes. Buildings are divided into roofs and facades. Vegetation is classified into separable plants (i.e. trees and bushes). Ground points are divided into street, pavement and grass. The third level includes any types of doors and windows placed in the roofs (dormers and skylights) and facades. Each class could be extracted separately or in a combination of other classes for various applications. Figure 4 shows the labelling order of the dataset.



Figure 4: Order of labels and the hierarchy.

Label generation

To generate labels, the initial Lidar dataset was divided into 13 sub-tiles of around 19 million points for annotating. The process started with importing data into the [CloudCompare](#) 2.10.1 software. Points were then coarsely manually segmented using segmentation and slicing tools into three categories (i.e. building, vegetation and ground) and labelled accordingly. Next, the process continued to the third level of the finest details (i.e. windows and doors). Hence, this pipeline produced a unique label for each point. The process took over 2,500 hours with appropriate supervision and was carefully cross-checked multiple times to minimize the degree of error.


The annotated dataset includes diverse types of historic and modern urban elements in the city centre of Dublin. Types of buildings include offices, shops, libraries and residential houses. The buildings range from detached and semi-detached to terraced houses and date from different eras (from the 17th-century Rubrics building to the 21st-century George's Quay complex). This detailed labelled dataset is the first of its kind regarding the accuracy, density and diversity of classes, particularly regarding its city-scale coverage area. The hierarchical labels offer excellent potential for various classification and semantic segmentation applications in urban science.



Figure 5: The white points show the differences between the 2017 and 2015 data of the TCD campus.

Applications and further information

The main goal of the labelled dataset is to train convolutional neural networks (CNNs) for classification of urban elements in massive point cloud data. For example, the labelled dataset can be employed to train and use [PointNet](#), PointNet++ and So-Net. These networks are able to classify urban elements for several important applications (e.g. robotic or autonomous navigation) based on semantic segmentation of the ground level which consists of streets, pavements and vegetation. These are essential elements for an autonomous navigation industry. Also, the vegetation class is highly beneficial for the detection of trees and it can be used to monitor the health of plants in an urban or even forest area. For example, a change detection technique applied to the drone and helicopter data from 2015 to 2017 clearly shows the location, size and number of removed trees. The white points (e.g. trees and the building) indicate what was removed after the 2015 scanning project (Figure 5). More information about the dataset (i.e. further description with video, link to the academic paper and a download link) can be found at bit.ly/GIM_magazine.



Further reading

- United Nations (UN) report, 2014. <http://esa.un.org/unpd/wup/highlights/WUP2014-Highlights.pdf>
- Debra F Laefer, Saleh Abuwarda, Anh-Vu Vo, Linh Truong-Hong, and Hamid Gharibi. 2015 Aerial laser and photogrammetry survey of Dublin city collection record. <https://geo.nyu.edu/catalog/nyu-2451-38684>
- NYU data repository: <https://archive.nyu.edu/handle/2451/38684>

Byrne, J., Connelly, J., Su, J., Krylov, V., Bourke, M., Moloney, D. and Dahyot, R., 2017. Trinity College Dublin drone survey dataset: <http://www.tara.tcd.ie/handle/2262/81836>

Zolanvari, S.M., Ruano, S., Rana, A., Cummins, A., da Silva, R.E., Rahbar, M. and Smolic, A., 2019. DublinCity: Annotated Lidar Point Cloud and its Applications. arXiv preprint arXiv:1909.03613.