



M Ű E G Y E T E M 1 7 8 2

Budapest University of Technology and Economics
Faculty of Electrical Engineering and Informatics
Department of Control Engineering and Information Technology

Oszkár Józsa

ANALYSIS OF 3D DYNAMIC URBAN SCENES BASED ON LIDAR POINT CLOUD SEQUENCES

SUPERVISOR:

Dr. Csaba Benedek

Budapest, 2012

Table of Contents

1	Introduction	3
2	Light Detection and Ranging	5
2.1	LiDAR technology	5
2.2	The Velodyne Sensor	7
2.2.1	Technical parameters	8
2.2.2	Scanning configurations	9
2.3	Data	10
3	Problem formation	11
4	Related work	13
5	Class segmentation	15
5.1	Method	15
5.1.1	Classes	16
5.2	Results	17
6	Point cloud registration	19
6.1	Normal Distributions Transform	19
6.2	Related work	20
6.3	Results	20
7	Framework for displaying and processing data	24

7.1	Development tools	24
7.1.1	Point Cloud Library	24
7.2	Implemented framework	25
8	Further results and possible improvements	27
8.1	Towards high level scene analysis and reconstruction	27
8.2	On the 3D reconstruction of urban scenes	28
8.2.1	Facade approximation with point clouds and triangle meshes .	29
8.2.2	Analysing moving and static objects on merged point clouds .	31
8.3	Higher level object analysis	33
8.3.1	Vegetation detection	33
8.3.2	Traffic sign detection	34
8.3.3	Higher level traffic interpretation	34
9	Conclusion	35
9.1	Output quality	35
9.2	Processing speed	35
10	Acknowledgment	37

Kivonat

Az elmúlt években a háztartási, kutatási és hadi robotok, valamint az automata járművek megjelenésével és kezdeti elterjedésével a hagyományos képi szenzorinformációk mellett egyre hangsúlyosabb szerepet kapnak a háromdimenziós érzékelők, köztük a LiDAR eszközök. Egy modern járműfedélzeti LiDAR szenzor elegendő információt szolgáltat a jármű körüli térrész értelmezésére akár egész utcákra vonatkozólag is.

Tudományos diákköri munkám célja ezen dinamikus utcai környezetek 3D rekonstrukciója és virtuális megjelenítése. Az eljárás bemenetét egy autótetőre szerelhető LiDAR készülék szolgáltatja, aminek a segítségével nagy pontosságú 3D pontfelhő szekvencia nyerhető egy akár 100 méter feletti sugarú területről 15Hz rögzítési sebességgel. Az adatok rögzítése valós, Budapesti utcai környezetben történik, míg a feldolgozás offline, CPU-n megvalósított algoritmusokkal, szem előtt tartva a későbbi GPU-ra való optimalizálást a valós idejű futás elérésének érdekében.

A cél az egyes területosztályok elkülönítése (pl.: úttest, házfal, növényzet), és meghatározott objektumok észlelése (emberek, járművek, közlekedési táblák) a pontfelhőkben egy előszegmentációs lépés során. Az osztályozáshoz különböző jellemzők együttes felhasználása szükséges, magában foglalva az objektumok alakját és színét, lehetséges kölcsönhatásaik modelljeit valamint a LiDAR adatfolyam alapján számolt különféle időbeli jellemzőket. Ezt követően az előszegmentált pontfelhők alapján lehetőség nyílik az épülethomlokzatok 3D rekonstrukciójára és az észlelt objektumok követésére, valamint az így kapott feldolgozott, értelmezett ponthalmazt későbbiekben illeszteni lehet az optikai kamerák 2D képeihez is, és ezek alapján színezett, élethű háromdimenziós helyszínmodellek származtathatók.

Az adatok rögzítése és feldolgozása egy Velodyne HDL-64E nagy teljesítményű LiDAR eszközzel történik a Magyar Tudományos Akadémia Számítástechnológiai és Automatizálási Kutatólaborjának Elosztott Események Elemzése kutatócsoportjában.

Abstract

In recent years, the spread of domestic, industrial and research robots and the robotic vehicles made 3D sensory information (such as LiDAR sensors) more and more widespread opposed to the conventional 2D image sensory information. A modern LiDAR sensor on board of a vehicle provides enough data to make us able to interpret the space surrounding the vehicle even at large scenes containing a whole street segment.

In my work for the students' scientific competition, I aim to reconstruct and display these 3D environments. The input of this procedure is a LiDAR sensor mounted on a car, which outputs a data stream covering more than 100 meters radius of space, collecting data at 15 Hz. The recording is in real environment on the streets of Budapest in real time, while the processing is offline, implemented on CPU keeping in mind the future implementation on GPUs to reach real time data processing.

The aim is to segment several region classes (such as roads, building walls, vegetation) and to identify specified objects (such as people, vehicles, traffic signs) in the point clouds through a presegmentation step. To achieve this classification, we need several features such as the color and geometrical properties of the specified objects and their possible geometrical and physical interactions. Also, we need to take into account the time domain features calculated based on the LiDAR data stream. After this presegmentation step we are able to reconstruct building facades in 3D and to track the detected objects in the 3D space. Also, in the future, this processed data set can be registered against 2D images provided by conventional cameras to reproduce realistic, colored 3D virtual spaces.

The data is provided by a Velodyne HDL-64E high performance LiDAR device, collected of Distributed Events Analysis Research Laboratory (DEVA) of the Computer and Automation Research Institute (MTA-SZTAKI).

1. Introduction

Analysis of 3D spaces comes from the need to understand the environment surrounding us and to build more and more precise virtual representations of that space. Remote sensing is a widely researched field for many decades and so is scene interpretation. But in the last decade, as the three dimensional (3D) sensors begun to spread and the commercially available computing capacity has grown big enough to be sufficient for large scale 3D data processing, new methods and applications were born. In many real world problems, 3D sensing and scene interpretation started to take the place of the conventional 2D image processing based on two dimensional imagery data.

In recent years, more and more technologies started to appear that heavily rely on these new 3D methods. Robotic cars, ships and airplanes are successfully tested and in some cases are already used on everyday bases. As Luettel et. al.[1] state in their summatory article, automation of our vehicles are the "path to the future".

But not only the automation industry depends heavily on 3D mapping and scene interpretation. There are applications of this technology in industrial manufacturing, geology and mapping, architecture, disaster and crime prevention and control, space exploration, military intelligence and virtual gaming. All these fields have their own special tasks, where 3D data can help both the professionals and the users.

Three dimensional data can be produced in several ways. It can be generated via sound propagation (sonars), radio wave propagation (radars) and light propagation (CCD devices, LiDAR devices). Light Detection and Ranging (LiDAR) devices are common tools for three dimensional mapping as they provide accurate 3D data straight out from the device (so there is no need of complex computation as via multiview cameras). These devices provide us the so called *point clouds*. A point cloud is a data set with small units of data, each representing a 3D point in the space. A point essentially has at least three information: its 3 coordinates, x , y and z (which of course can be represented in polar coordinate system or Euler angles or *latitude, longitude, altitude* in geographic data sets, etc). Additionally color, intensity information could also be provided by some devices.

These instruments can quickly produce huge amount of data which makes data processing a hard task. Some of these devices output several million data points per second. Efficient, fast methods are needed to filter the significant data out of these streams or high computing power is needed to post-process all this large amount of data. Once we are able to select the information useful to us, another interesting task is to merge this data from our whole scan sequence to even larger and more useful 3D data sets. Matching subsequent scans to each other is called point cloud registration which can provide us even more useful data sets; for example registered urban point clouds can be processed and analyzed similarly as aerial LiDAR scans but at much higher resolution since the sensor is able to scan the scene from a much smaller distance.

The analysis of 3D dynamic urban scenes can be the first step for either real time scene interpretation tasks (essential for traffic monitoring and robotic cars), or for post processing three dimensional data (building realistic virtual models of real cities).

2. Light Detection and Ranging

In this chapter, I introduce the working principles of the Light Detection and Ranging (LiDAR) technology and show examples of devices and their usage.

2.1 LiDAR technology

Light Detection and Ranging (LiDAR) technology is an optical remote sensing technology used for distance measurements. It is available in terrestrial, airborne and spaceborne devices. Wide range of instruments are available for variety of tasks - ranging from hobby devices about 100 US dollars up to military-grade high powered multi-beam ones for hundreds of thousands of US dollars.

The principals of LiDAR distance measurement is essentially the same as of a radar device but instead of radio waves, a LiDAR uses light to measure distances, thus the name *Light Detection and Ranging*. The distance of an object is calculated from the time it took the light beam to bounce and arrive back from an object. Since we know the speed of light, the total distance traveled is simply given by the time multiplied by the speed of light.

$$object\ distance = \frac{time\ of\ flight * speed\ of\ light}{2} \quad (2.1)$$

A big advantage of this technology over conventional optical (CCD) imagery is that it is not affected by lighting conditions. LiDAR is an active sensing technology, meaning that it illuminates its target so lack of external lighting (e.g.: at night) does not corrupt the measurement. Also, special types of LiDARs are able to "see" through water, thus being able to scan underwater surfaces.¹

In most cases, terrestrial LiDAR devices operate while surrounded by people. For safety reasons, these devices use eye safe (Class 1 laser) laser emitters, usually

¹So called hydrographic mapping are provided by companies such as Optech (<http://optech.ca/>) and Fugro (<http://www.fugro-uae.com/>)




Figure 2.1: The internal structure and the working principle of the Velodyne HDL-64E sensor. It has a rotating head and 64 laser beams aligned along a single line. With this line, it scans the surrounding 3D space as the head rotates around in 360° . The four shining plates on the face of the device are the laser emitters and sensors (two of each) (image source: Argocorp product description)

between 600 and 1000 nanometers of wavelength.² This light is in infrared range which makes the light beams invisible to human eye.

Different LiDARs have different operation principles. Some devices (such as the SICK LiDARs, or the industrial stationary laser scanners) have one laser beam that sweeps back and forth in one dimension, along a single line. Either the scanned object or the scanning platform has to move to for a complete 3D scan to be performed. Other types of LiDAR scanners, such as the Velodyne devices, have a revolving head which sweeps in a circle around the sensor, resulting in a full 360° scan. These sensors provide 3D images even at a stationary position, though in most cases these are used on a moving platform, also.

²Source: www.LiDAR-uk.com/how-LiDAR-works/

2.2 The Velodyne Sensor

The data processed in this thesis is provided by the *Velodyne HDL-64E* high performance LiDAR sensor. According to its website, the "HDL-64E LiDAR sensor is designed for obstacle detection and navigation of autonomous ground vehicles and marine vessels. Its durability, 360° field of view and very high data rate makes this sensor ideal for the most demanding perception applications as well as 3D mobile data collection and mapping applications."³

The typical usage of this sensor is to mount it on a moving platform, practically onto a car. The LiDAR has fast enough scanning speed to scan the whole 3D space at about every 30-60 centimeters of distance as the car travels with typical urban traffic speed. With each revolution, a full 3D scan is made, and overall the device outputs more than 1.3 million 3D points per second.

Alternatively, it can be used as a stationary sensor at outdoor or relatively large indoor scenes (e.g.: warehouses, manufacturing halls). Applications of this static configuration will not be discussed in this this thesis, instead, we will be focusing on the more challenging moving platform.

Along the 3D (x-y-z) data, the sensor also gives a fourth value for each point, called intensity. This value shows how strong the reflection was, in other words, how big fraction of the emitted light returned to the sensor. This intensity value mostly depends on the objects' surface quality: shiny, flat surfaces reflect light much better than matte, scattered surfaces such as roads or vegetation. The angle of attack can also affect this intensity value i.e. if the beam hits the surface in a high angle, the light can scatter and thus the returning intensity will be lower. Although the algorithms presented in this thesis do not exploit the intensity information, I will mention examples in Chapter 8 where it might highly support some of the proposed detection tasks.

³For full product description, visit the official product website <http://www.velodyneLiDAR.com/LiDAR/hdlproducts/hdl64e.aspx>



Figure 2.2: Velodyne HDL-64E sensor on the streets of Budapest (pictured in front of Gellért Bath)

2.2.1 Technical parameters

Specifications	
Sensor:	<ul style="list-style-type: none"> • 64 lasers/detectors • 360 degree field of view (azimuth) • 0.09 degree angular resolution (azimuth) • 26.8 degree vertical field of view (elevation) - +2° up to -24.8° down with 64 equally spaced angular subdivisions (approximately 0.4°) • <2 cm distance accuracy (one sigma) • 5-15 Hz field of view update (user selectable) • 50 meter range for pavement (~0.10 reflectivity) • 120 meter range for cars and foliage (~0.80 reflectivity) • >1.333 M points per second • Operating temperature - 10° to 50° C • Storage temperature - 10° to 80° C
Laser:	<ul style="list-style-type: none"> • Class 1 - eye safe • 4 x 16 laser block assemblies • 905 nm wavelength • 5 nanosecond pulse • Adaptive power system for minimizing saturations and blinding
Mechanical:	<ul style="list-style-type: none"> • 15V ± 1.5V @ 4 amps • <29 lbs. • 10" tall cylinder of 8" OD diameter • 300 RPM - 900 RPM spin rate (user selectable) • Environmental Protection IP67
Output:	<ul style="list-style-type: none"> • 100 MBPS UDP Ethernet packets

Figure 2.3: Data sheet of the Velodyne HDL-64E LiDAR (image source: data sheet on the product website)

The sensor provides the data in an UDP stream through a standard Ethernet cable. The packet structure is detailed by the manual as follows:

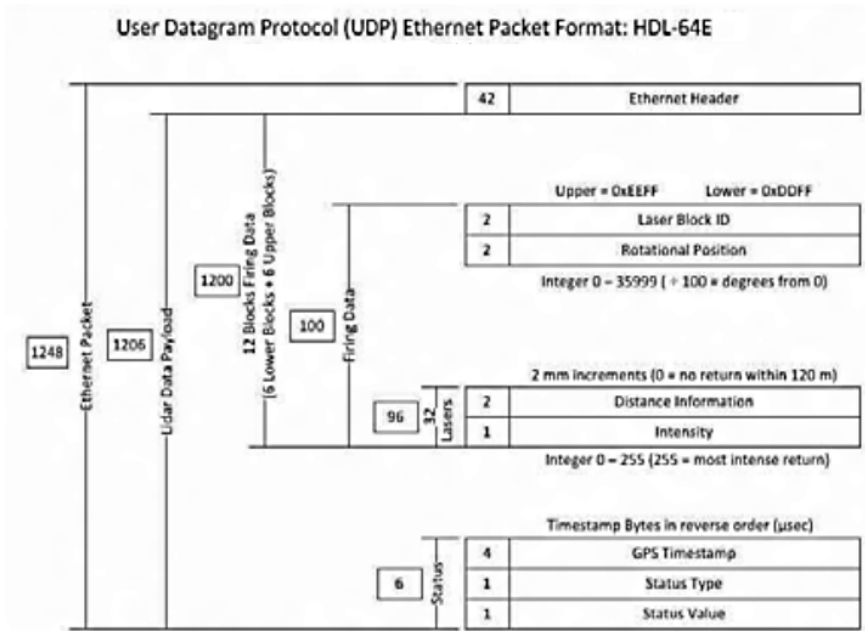


Figure 2.4: Data sheet of the Velodyne HDL-64E LiDAR (image source: product manual sheet)

2.2.2 Scanning configurations

The sensor has two main configurations when mounted atop of a vehicle. It can be either facing forward and thus, looking down in an angle of about 25 degrees or it can be tilted to face upwards. The first, forward facing configuration is more suitable for road mapping, traffic monitoring, object detection and tracking, etc. The latter configuration, because of its upward looking nature, is more suitable for scanning building facades. A typical angle in the tilted configuration is 45° which allows the sensor to see up several stories high from street level.



(a) Forward facing LiDAR

(b) Tilted LiDAR

Figure 2.5: The two scanning configurations

Though this thesis focuses on the forward facing configuration, the Velodyne

HDL-64E LiDAR can be used in both configurations and some results of the different measurements and the possible uses of them will be discussed later.

2.3 Data

All the data had been recently recorded in urban areas, on the streets of Budapest. Several types of streets have been scanned to test the algorithms as well as bridges and large squares. The types of streets range from avenues to small hillside roads.

The data recording has been done with a software tool developed earlier at the Distributed Events Analysis Laboratory (DEVA). This tool saves the data into the standard *pcap* (packet capture) format which contains the raw UDP stream with timestamp and header information along with the 3D data. Also, this program can export the data into the simple *.pcd* file format provided by the Point Cloud Library (for more detail, see Chapter 7). In this format, the recording is a file sequence with every file containing one 360° scan.

The processing framework uses the PCD sequence format. A typical sequence contains a few hundred scans and covers a whole street or a segment from an avenue. A scan recorded with 10-15 Hz of rotation has a size about 2 to 6 megabytes so a sequence can reach multiple gigabytes in size. We call the acquired sequence 4D data since each frame is a 3D point cloud and the time domain represents the fourth dimension.

3. Problem formation

Though the LiDAR sensor provides high amount of data with each revolution, often this 3D information is not sufficient to understand a whole scene which sometimes spans over 200 meters. Typical challenging problems are occlusion (and thus, large "shadows" on the occluded background surface), sparse data and noise. Though the sensor has one sigma accuracy, even a few centimeters of fluctuating noise can damage the surface reconstruction and make identification tasks challenging.

Also, the density of the provided point cloud is highly inhomogeneous: the further an object from the sensor is, the sparser its 3D scan will be. This fact hurts lots of existing processing methods developed for the analysis of homogeneous aerial LiDAR scans will not give the same results for the same object when it is close to the sensor and when it is far away.

As the sensor does not see directly down to the ground, each scan has a "hole" of about 3 meters radius in the center. Some registration algorithms consider this "edge" in the data as a significant feature. Once identifying these points as keypoints, they register two point clouds along these points which obviously results in a false registration since the "hole" in the data moves along with the car and should not be aligned in consecutive point clouds.

In addition to noise, moving points can cause hard challenges for the processing algorithms even more. In a realistic street scene, there are number of moving objects: people, vehicles, vegetation. All the points belonging to these objects change from frame to frame. A false registration can easily occur when the algorithm falsely detects moving points as significant points and tries to align consecutive point clouds along these moving points.

If we were able to register all the clouds in a sequence against each other and concatenate them into a large point cloud (within the same reference coordinate system), huge resolution could be achieved and most of the problems mentioned in this chapter could be overcome. Once we have the registered cloud, the resolution will be uniform and most of the occlusion will be gone due to the moving platform which helps the sensor to measure objects from different angles. Also, in this significantly

larger resolution, effective noise reduction can be applied and detection algorithms can perform much better.

In summary: a data preprocessing method is needed that helps the point cloud registration process. In particular, a segmentation method, which removes (and optimally classifies) all the points that harden the registration task such as the road surface, moving objects and noise. Our aim is to select a set of points that are stable, static in each frame and can be confidently used as keypoints for point cloud registration.

The key contributions of this thesis are as follows. In Chapter 5, I present my presegmentation method which aims to help resolve the basic point cloud classification problem, i.e. separating ground, wall, object etc. regions in the point cloud. In Chapter 6, I introduce a new algorithm to transform the consecutive Velodyne frames into a joint coordinate system, in such difficult scenarios, when both the LiDAR sensor is moving and in and several moving objects are also present on the street. In addition, we do not use any special sensor information (IMU) about vehicle speed and acceleration for matching the frames. In Chapter 7, I present the framework I developed for the workflow (preprocessing, registration and visualization of the data). In Chapter 9, I present my results on these registered data and some future plans for improving the current results. Finally, In Chapter 8, I will draw the conclusion and analyze the results.

4. Related work

Hereby I present a short overview of the state of the art results on this topic. Most researches are focusing on automation, robotic vehicles only fewer researches are made on the topic of virtual city reconstruction and building facade reconstruction.

In most cases, for point cloud registration in robotic environments the ICP [2] algorithm is used for matching which performs well in homogeneous 3D scans but often fails in noisy, fast-changing real world environments or has to be modified to be scene specific. In robotic applications, the main task is to detect and track moving objects[3, 4], while our current task is to understand a scene with both moving and static objects and build realistic virtual representation of that scene.

Douillard et. al. have developed a method to extract "segments" (interest regions) from the point cloud [5] [6]. This approach is similar to ours: with our preprocessing step, we also try to find sets of points that will later help the registration step. Makadia et. al. constructed an automatic registration method using large histograms calculated from the surface normal fields[7] but is computationally expensive and works well on homogeneous scans, not on LiDAR data.

Others have taken a different approach Behley et. al. [8] constructed a multiscale 3D histogram descriptor that can be efficiently used for point matching. Quadros et. al. presented a feature called *the line image* to support point matching that outperforms the widely used NARF descriptor but requires a computationally expensive principal component analysis (PCA) calculation[9].

After the registration is done, next step is virtual reconstruction. Since the scans are noisy by nature and possibly have holes in them (e.g.: occlusion caused by trees and parking cars standing in front of buildings, etc.), reconstruction is not trivial. Recently, Shen et. al. showed an efficient method for facade reconstruction which uses an iterative adaptive partitioning and is robust against the aforementioned issues [10]. Also, Wang et. al. of the NAVTEQ Corporation showed a method to upsample sparse LiDAR data and to clean the points corresponding to building interiors (since the laser beam travels through windows also, the LiDAR also scans the interiors of buildings)[11].

Also, for already registered, large data sets, there are effective detection, segmentation and recognition algorithms available. Velizhev et. al. developed a highly effective algorithm to interpret large outdoor urban scenes[12]. Their method is proposed for mobile robot mapping; though, once we register LiDAR data with the method presented in this thesis, our data sets will be similar to theirs.

Though in our method, running time is not a bottleneck, in future development parallelisation might fasten up the algorithms greatly. Hu et. al. numerically proved that point cloud filtering and segmentation can be greatly faster with a graphics processing unit (GPU) implementation[13].

5. Class segmentation

To achieve the aforementioned goals, I developed a novel method to preprocess and process 3D LiDAR data. The key idea is to insert a simple and fast preprocessing step that will help the registration step, hence scene reconstruction and scene interpretation also.

This classification is based on local point properties. Using some statistical descriptors, I segment the data into semantic classes which later can be used together or separately for various tasks.

I have defined four classes:

- Road surface (including sidewalks)
- Short street objects (which are moving objects in lots of cases, such as cars, people)
- Walls and tall static objects (lamps posts, traffic lights, etc.)
- Sparse data / noise

5.1 Method

As mentioned before, point cloud segmentation is done by a grid based approach. First, I fit a regular 2D grid S with W_S rectangle width onto the $P_{z=0}$ plane, where s denotes a single cell. I used a W_S value between 50cm and 80cm. Smaller grid size is not viable due to the resolution; smaller cells would not have enough points in them to calculate good enough statistical information. On the other hand, larger cell size can result in larger number of falsely classified points, since within a large cell, multiple objects can occur. Near-the-center grid cells within the above specified range have hundreds of points in them (grid cells further away obviously have less and less points).

Then, we assign to each $p \in \mathcal{P}$ point of the point cloud to the corresponding cell s_p , which contains the projection of p to $P_{z=0}$. Let us denote by $\mathcal{P}_s = \{p \in \mathcal{P} : s = s_p\}$ the point set projected to cell s . $y_{\max}(s)$, $y_{\min}(s)$ and $\hat{y}(s)$ are the maximum, minimum and average of the elevation values within \mathcal{P}_s .

The first step of the process is the segmentation of the cell-map, i.e. we assign to each cell (and to each point in that cell) $s \in S$ an ω_s class label from the finite label set: $\mathcal{L} = \{l_{\text{to}}, l_{\text{so}}, l_{\text{gr}}, l_{\text{sp}}\}$, corresponding to the classes (i) *tall structure object*, (ii) *short street object*, (iii) *ground* and (iv) *sparse region*.

5.1.1 Classes

At first, sparse regions (l_{sp}) are detected. These encapsulate only a few or not any points: we can obtain only very limited, and probably misleading information from these cells regarding the scene structure and objects, therefore we will neglect these regions in the later model phases. The threshold of a cell s being considered as a sparse cell is typically 4-8 points - any cell containing less points than this threshold is considered sparse (the exact value of this threshold also depends on the sensors revolving speed, slower speeds allow to have larger threshold). Sparse data is classified first, so in later processing we will know not to check these grid cells, thus save some processing time.

A cell should belong to the l_{to} class if it contains a tall object of the street structure, such as building walls, lamp post or tree trunk. These object can be considered static in the scene, since these 3D points obviously belong to objects that are not moving and are large enough to be present in several consecutive scans. These points will be used later for point cloud registration of the consecutive time frames. The two criteria for a grid cell to belong in this class are:

$$y_{\max}(s) < 140 \ || \ y_{\max}(s) - y_{\min}(s) > 310 \tag{5.1}$$

So either the maximal height is larger than 140 centimeters, either the height difference is larger than 310 centimeters. The two criteria are similar. The first one says that if there are points in the grid cell which are higher than the sensor height plus 120 centimeters, than the cell should be considered as a tall object (the sensor has a height about 2 meters). The second criteria says that if there is a height difference of more than 310 centimeters within the cell, it should be considered as a wall. The second criteria is needed for dealing objects standing on a lower point

of the ground compared where the car travels (e.g.: elevation differences, walls seen from an overpass road, etc).

Next, the ground cells (l_{gr}) are identified. These cells contain only a surface which is not covered by any objects, but contains a considerable number of ground points. The criteria for the grid cell to belong into this class is:


$$y_{\max}(s) - y_{\min}(s) < 25 \ \&\& \ y_{\max}(s) < -50 \quad (5.2)$$

In words: if the points within a cell are "flat" enough, and the maximal height is below -100 centimeters, the grid cell is considered as road surface. The first criteria ensures the flatness or homogeneity of the points. Given a cell with 60 centimeters if width, this allows 22.6° of elevation within a cell; higher elevations are highly unrealistic in an urban scene. The second criteria ensures that this patch of flat surface is under the car; again, the sensor is at about 200 centimeters of height so the road surface directly beneath the car has a height of about -200 centimeters. Less than -50 criteria is used to deal with elevation differences that can occur within the field of view of the sensor.

The rest of the point cloud is assigned to class l_{so} . These points belong to short objects like vehicles, pedestrians, short road signs, line posts etc. These entities can be either dynamic or static, which attribute can only be determined later, after further, more complex investigation of the point cloud sequence.

5.2 Results


Here I present some images that show the results of the segmentation. For the conclusions, see Chapter 8 and for detailed description about the processing framework, see Chapter 7.



(a) Bartok Béla street

(b) Villányi street


Figure 5.1: Streets "cleaned" from noise and objects



(a) Parking cars on the side of Villányi street

(b) Traffic at red light, some parking cars and a person passing by on a small street

Figure 5.2: Streets with only the object clouds displayed



(a) Intersecting one-way streets

(b) Intersecting one-way streets

Figure 5.3: Street segments with sparse and noise data included (blue points)

Notice, that on Fig5.3, there are seemingly large amount of noise. Actually, since the majority of those points are in the sparse region, in most of the cases the amount of the noise/sparse data is below 10%.

6. Point cloud registration

A single scan has many points but since it covers a large area, the resolution is sufficiently good only within few meters of distance. Though the device has a sensing distance of more than 100 meters, the measurements at more than 15 meters of distance are too sparse for many detection or reconstruction algorithms. But since the sensor is on a moving platform, it is a good idea to register, or in other words, concatenate the consecutive scans in the sequence. Registration will help us to fill in the holes due to occlusion and we can achieve very high point density after registering several point clouds.

Although various established techniques do exist for point cloud registration, such as Iterative Closest Point (ICP) [2] and Normal Distribution Transform (NDT) [14], these methods fail, if we try to apply them for the raw Velodyne LIDAR point clouds for two reasons: 1) All moving points appear as outliers for the matching process, and since in a crowded street scene we expect several moving objects, many frames are erroneously aligned. 2) Due to the strongly inhomogeneous density of the LIDAR clouds, even the static ground points mislead the registration process. The above algorithms often match the concentric circles of the ground (See Fig. 5.1), which yields that the best match erroneously corresponds to a near zero displacement between two consecutive frames. However, I have also observed that the point density is quite uniform in local wall regions which are perpendicular to the ground.

My key idea is to utilize the point classification result from the previous chapter to support the registration process. I only use as input of the registration algorithm the points segmented as high objects, since I expect that in majority, these points correspond to stationary objects (such as buildings), thus they provide stable features for registration. Details and test results are in the following sections.

6.1 Normal Distributions Transform

This method was introduced in Martin Magnusson's doctoral thesis, *The Three-Dimensional Normal-Distributions Transform — an Efficient Representation for*

Registration, Surface Analysis, and Loop Detection [14] which is a 3D implementation of the 2D Normal Distributions Transform developed by Biber and Straßer.[15]

The Normal Distributions Transform itself uses a grid based approach also. First, it divides the space into cubes. For each cube, calculates its so-called local probability density function (PDF) to describe that cube: "Each PDF can be seen as an approximation of the local surface, describing the position of the surface as well as its orientation and smoothness." [14]

For the registration step, it uses Newton's optimization method to find the rotation and translation between the two point clouds, searching for the best match between the PDFs of the two scans. This method is robust to outliers.

6.2 Related work

In many cases, this registration is done by the help of external sensors such as Inertial Measurement Units (IMU), Global Positioning System (GPS) or other sensors on the car (e.g. speedometers, wheel rotation sensors, radars).

Han et. al. presented a method to detect road boundaries and road segments based on LiDAR data and vehicle movement information [16]. Robotic cars such as Stanley and Junior built for the DARPA Grand Challenges[17] [18], the AnnieWAY self driving car [19] and Carnegie Mellon's Car, Boss[20] for example. use a huge amount of sensors to map their environment and interpret the scene surrounding them.

All of the results above incorporate additional sensors. The novelty of my method that it does not need any supplementary sensory information, also does not need any learning algorithms or huge sample database. Once the first correspondence is found, the calculated transformation matrix is used as a prediction for the next movement step, the NDT algorithm only has to "refine" this prediction according to changes in the movement of the car.

6.3 Results

In this section, I present some results of the point cloud registration. Fig 6.1 shows an upper view of a large registered point cloud of a bridge. Fig 6.2 presents

the process of the point cloud registration on a street segment, showing how radical improvement on the point cloud is possible. Lastly, Fig 6.3 and Fig 6.4 present two registration result, both recorded on the streets of Budapest.






Figure 6.1: Liberty Bridge (Szabadság híd) with boats in the background




(d) The resulting registered scan (made out of 20 scans, not only the 3 above).

Figure 6.2: An image sequence showing the registration process.




(a) Upper view of the street. Notice, how the interiors of the building gets visible as the scanner passes the door.




(b) Zoomed in picture.

Figure 6.3: A large registration of Kende street showing only the building class. Scan sequence starts at the middle of the street and runs to its end. The point cloud is made out of 230 point clouds, the full cloud contains more than 20 million points.



(a) The whole registered street segment



(b) Part of the street, zoomed in. Note the level of detail, even objects in the store window are visible.

Figure 6.4: Registered street segment. For registration, 30 point clouds were used, final cloud contains almost 10 million points.

7. Framework for displaying and processing data

7.1 Development tools

The framework is written in the C++ programming language on the Windows platform but using only multiplatform technologies.

7.1.1 Point Cloud Library

The Point Cloud Library (PCL)¹ is a free, open source C++ library for point cloud processing. It is available for all major platforms and the source code is licensed with the BSD license. It is developed by Willow Garage and many universities around the Globe and financially supported by companies such as Google, nVidia, Trimble, Velodyne. The library provides both high and low level 3D point cloud processing algorithms. For details on this software library, please refer to [21]

The library also defines a widely used file format for storing point clouds. It is a simple format containing a header and the 3D information, functions to read and write this format are available in the library.² I store the Velodyne LiDAR data in this file format. It is flexible enough to contain supplementary information about each point such as labels for the classes, normal and curvature data, etc. The data is stored in binary format to speed up disk usage time. Reading and saving LiDAR data is done with the help of the PCL's input-output library. For the 3D transformation data, I constructed a simple file format which contains the calculated transformation matrices between the consecutive point clouds.

Functions I used from the PCL library:

- Input - output module

¹Webpage: <http://pointclouds.org/>

²The detailed description can be found on the webpage

- Kd-trees for search within clouds
- Statistical outlier removal
- Visualization module
- Random Sample Consensus (RANSAC) model for comparison for road detection
- Normal Distributions Transform to register point clouds
- Euclidean Cluster Extraction for clustering the objects in a point cloud

7.2 Implemented framework

The core module of the framework is the cell-based simple segmenter class. It is responsible to build the grid and to calculate the statistical information about the point cloud and each cell. Also, it is responsible to segment the points (cells) into classes and color them accordingly for visualization.

The framework has a simple display module for visualizing data and the results of the algorithms. It is built atop of PCL's visualization library (which uses the Visualization Toolkit library's core functions). The display module also uses the Boost library's parallelization library; I implemented visualization and keyboard handling to run in separate threads.

The display module can be turned off via starting the application with the *-novis* option. This way, the application will behave as a command line processing tool (additional command line options are: *-help* for displaying usage information *-verbose* for displaying statistical information about the processed data and *-o* for turning on a basic statistical outlier removal function on the point cloud).

Display modes:

- **Stream mode** to display data sequence without processing or coloring. This mode is useful for playback of 4D data
- **Simple Segment** class coloring. In this mode, the grid based segmentation will be performed and the point cloud will be colored according the pictures shown before.

- Linear coloring based on **average height** in a cell. In this mode, the processing function calculates the average height within a grid cell and colors the point cloud accordingly (similarly to a heatmap)
- Linear coloring based on the **point density** in a cell. In this mode one can easily check visually the point density distribution in a point cloud.
- Linear coloring based on **minimum-maximum height difference** in a cell. Similarly to the previous two, the processing module calculates a statistical information about each cell and colors the cloud accordingly. Useful for visually evaluating the height differences.

The framework has some additional functionalities for evaluation and helping development. If the user clicks on a point, it displays information about the point and the grid cell the point belongs to. Also, the threshold parameters can be changed during 4D playback for faster evaluation. If the file save flag is set to true, the framework outputs the data into several files after processing the input:

- A colored point cloud. The colors correspond to the four classes.
- A point cloud containing only the walls and high objects.
- A point cloud containing only the street objects.

And finally, the registration workflow is as follows: first, the point clouds containing only the static wall points are loaded. Then going through this list of point clouds each pair is registered each against with the Normal Distribution Transform. Once this is done, we know all the transformation matrices that represent a single transformation between the coordinate systems of consecutive scans. Using these matrices, we can transform any point cloud (the full cloud, cloud with only the walls, etc) into one mutual coordinate system. After this, we only have to concatenate the point clouds and save the result onto the hard drive.

8. Further results and possible improvements

In this chapter I present various algorithms and test results of my thesis work and some future plans, which point towards complex scene analysis and reconstruction. First, I introduce a point filtering algorithm, which is able to refine the grid based coarse point cloud classification step (Sec. 8.1). Then, I present my results on surface reconstruction and object detection in street scenes. Finally, models on higher level object and scene analysis are discussed.

8.1 Towards high level scene analysis and reconstruction

As mentioned before, the grid is defined as a cuboid with a base of 50x50 to 80x80 centimeters. A classification is done for the whole cell, which means that if a wall is (correctly) classified as wall, though some road points are (incorrectly) also classified as wall at the foot of the wall.




Figure 8.1: Cars with a layer of "carpet" under them. Blue color means it is classified as sparse/noise data

As the picture shows, currently this small patch of misclassified road is simply removed from the data by cutting down a few centimeters of each cell that has object class. It should be further investigated and reclassified to road class if possible.




Figure 8.2: Cars without the detected road patches. It can be seen that this simple removal method is not sufficient to remove all the misclassified points

8.2 On the 3D reconstruction of urban scenes

As mentioned before, the long-term aim is to build realistic, virtual 3D models of the streets of Budapest. For this task, the existing triangulation methods should

be investigated and tested, as well as data smoothing methods.

According to plans, later an RGB camera will be available - either a 360° field-of-view panoramic camera or a 360° looking multicamera system such as the PointGrey Ladybug camera¹. Once such a camera is available, the point cloud can be colored and that additional color information could be used as a base of the 3D surface texturing methods.

8.2.1 Facade approximation with point clouds and triangle meshes

After registering several point clouds against each other, the resolution can be high enough and uniform enough to create realistic building facade reconstruction. As the car passes by a building, it collects data from several point of view so most of the wholes on the walls due to occlusion can be filled in. Also, after concatenating a few dozen scans, the resolution of the data will be significantly higher which results in a precise 3D reconstruction of wall surfaces and more efficient noise reduction also.

The current workflow I developed is as follows. After the transformation matrices are calculated via the NDT algorithm, they are applied only to the point clouds that contain the static tall points. After transforming these clouds into a mutual reference frame, we will get a large point cloud, which only contains walls, lamp posts, trees and similar tall objects. On this larger, registered pointcloud a clusterization is done which separates the aforementioned objects into separate point clouds.

¹Product description at <http://www.ptgrey.com/products/spherical.asp>




Figure 8.3: A clusterization result. Each color represents a separate cluster (colors do not have any particular meaning, they were assigned randomly).

Once we have the separate point clouds, each representing an object, we need to "clean" it and to smooth it. Cleaning is necessary because most likely there are outliers around the object and smoothing is required because of the noise. Since the resolution is high (millions of data points along a single facade), the noise can be eliminated and the reconstructed surface will be smooth enough.




Figure 8.4: A triangulation result. It depicts a building on Moricz square.

For this task, a tilted LiDAR configuration is even more suitable than the forward facing configuration. In this case, the data is much more unfitting for traffic monitoring but more suitable for scanning buildings. As the sensor looks up it can scan the buildings up to their roofs. When registering this data, we can achieve incredibly high-detailed building facades.



Figure 8.5: An example of the tilted configuration. It depicts the Klotild Palace, near the Elizabeth bridge in Budapest

The results shown above are proofs of concept, the framework will be extended to automatically process the aforementioned steps.

8.2.2 Analysing moving and static objects on merged point clouds

General 3D object reconstruction is a much harder task than reconstructing walls or road surface. While these have a well defined, and mostly flat surface, a general street object can have virtually any shape and due to small resolution and the fact that the LiDAR only scans one side of the object, the reconstruction is sometimes impossible.




Figure 8.6: An example of how moving persons look like in the registered point cloud (the two blue boxes) versus a person who was standing still during the whole scanning process (red box)




Figure 8.7: Another example of moving objects appearing as a smudge in the registered cloud. Here, a car turned onto the main street after stopping at the red light

Moving objects are even harder to reconstruct based solely on LiDAR data. As these objects (typically vehicles, people) are moving through the scene, they appear at different spots in each scan which make them appear like a long-drawn shadow in the registered point cloud.

Instead of reconstructing these objects from the LiDAR cloud, our plan is only to detect them in the cloud. Once they are detected, they can be replaced with realistic, full-scale 3D models in virtual space. Free 3D models of vehicles, vegetation, trash

cans, etc are available to use for this purpose. Also, the Geometric Modeling and Computer Vision Laboratory of the Computer and Automation Research Institute has the capabilities to scan dynamic models, such as walking people. These dynamic models can be inserted into our virtual space to achieve a realistic 3D world.²

8.3 Higher level object analysis

8.3.1 Vegetation detection

Another well defined task is vegetation (typically: trees and bushes) detection. The reason for this is two-fold: on one hand, a popular task is "vegetation mapping" along with calculating the green area in a city. On the other hand, the removal of the detected vegetation data from the point cloud can help detection algorithms, for example in the case of trees hanging over parking cars.

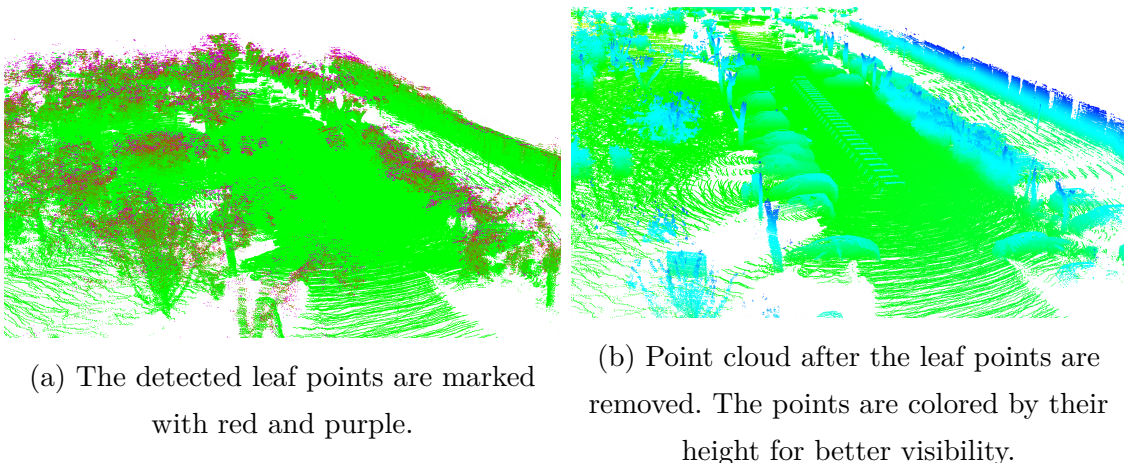


Figure 8.8: An example of vegetation removal

I have developed a vegetation removal algorithm which calculates a statistical feature for each point based on the distance and irregularity of its neighbors. As mentioned before, the intensity channel also can serve as a strong indicator of vegetation. Once the intensity and the statistical information will be used as a multifeature descriptor, it will be a strong feature for vegetation detection.

²For initial results on this topic, see the demo video at the website <http://web.eee.sztaki.hu/home4/node/14>

8.3.2 Traffic sign detection

A popular task on urban scenes is traffic sign detection. It is a crucial task in intelligent vehicles but also can be used by route authorities to inspect whether the necessary traffic signs are present where they should be.

The possibilities of traffic sign *detection* should be investigated in two aspects:

- **Intensity channel** - traffic signs have highly reflective paintings on them for better visibility. This can be a strong indicator of a traffic sign.
- **Physical constraints** - as the walls or the road surface, a traffic sign also has a well definable physical constraint. A model should be developed to represent these constraints. Also it should be investigated whether this model should be used in the existing grid-based processing or on full clouds without grid representation.

Even further development is traffic sign *identification*. For this task, a simple 3D data is not enough, color information is also needed to identify similarly shaped but different traffic signs. Once RGB referenced 3D data and the traffic sign detection method is available, the possibilities of pattern matching for identification should be investigated.

8.3.3 Higher level traffic interpretation

For higher level traffic monitoring, a better road interpretation is necessary. Detailed information of the road surface is required for this task; such as dividing the road class into smaller semantic units e.g.: sidewalks, lanes, parking lots, road intersections.

Lanes (and parking spots) can be identified by the lane markings. Lane markings are made out of bright painting which makes them fairly easily detectable. Also, later RGB information can be used to support this method.

Also, for more detailed traffic interpretation, vehicle detection is necessary. It is a highly researched area, and there are existing effective methods for aerial LiDAR data. After registering a whole street sequence, the data can be used similarly as of an aerial LiDAR data, with much higher resolution.

An overview and results on this topic on aerial LiDAR data can be found in [22]

9. Conclusion

The results above show the success of the simple, yet useful presegmentation step that has a great positive effect on the point cloud registration. Dividing the initial point cloud into multiple semantic classes and using only the relevant points results in a faster, more stable registration.

Also, a software framework has been successfully implemented atop of the PCL library for handling, processing and visualizing the Velodyne LiDAR data.

9.1 Output quality

As it can be seen in the pictures above, successful point cloud registration can be done on the segmented, preprocessed data. The registration is accurate and can be easily performed on the static points.

The big advantage of the road surface detection method over a RANSAC-based plane matching that it can deal with curvatures, elevation changes in the scene. RANSAC can find the best fitting plane for the road surface but in real world, the road surfaces are not planars, especially in cities built on or next to a hill. Using the proposed method, such difficulties can be overcome and once the surface grid cells are detected, even an elevation map can be built.

One other benefit of the grid based segmentation that is also removes noise. Grids containing too few points are not useful, they make computation harder. The segmentation filters these points out also.

9.2 Processing speed

For the segmentation step, the main aim was to implement a fast algorithm that performs a presegmentation step. This aim has been met; once the data is in the memory (reading from the disk is the slowest step), the algorithm runs in real time

on a modern CPU.¹

Besides the segmentation being fast, it fastens up the registration step also. Once all the noise and interfering moving points are removed, the registration is easier, thus faster. As the registration algorithm tries to match up points, the overall error will be smaller thanks to the "clean" point cloud and the algorithm will converge faster towards the termination criteria. Registering point clouds containing only stable, stationary points takes only 1-4 seconds (this value depends on the scanning rotation speed and on the amount of points left in the cloud after segmentation). Registering a full cloud takes 5-10 second - if it converges, since in a lot of cases the registration fails.

¹The test environment has an Intel Core i7 processor and 8 gigabytes of DDR3 RAM

10. Acknowledgment

I would like to thank my supervisor, Dr. Csaba Benedek for his support, his advises and his lecturing work. Also, I would like to thank Dömötör Molnár the software tools for capturing and converting Velodyne data. And finally, I would like to thank the Distributed Events Analysis Research Laboratory (DEVA) of the Computer and Automation Research Institute (MTA-SZTAKI) for providing the devices and place of work and all the colleagues who have inspired or supported me in any way.

List of Figures

2.1	The internal structure and the working principle of the Velodyne HDL-64E sensor. It has a rotating head and 64 laser beams aligned along a single line. With this line, it scans the surrounding 3D space as the head rotates around in 360° . The four shining plates on the face of the device are the laser emitters and sensors (two of each) (image source: Argocorp product description)	6
2.2	Velodyne HDL-64E sensor on the streets of Budapest (pictured in front of Gellért Bath)	8
2.3	Data sheet of the Velodyne HDL-64E LiDAR (image source: data sheet on the product website)	8
2.4	Data sheet of the Velodyne HDL-64E LiDAR (image source: product manual sheet)	9
2.5	The two scanning configurations	9
5.1	Streets "cleaned" from noise and objects	18
5.2	Streets with only the object clouds displayed	18
5.3	Street segments with sparse and noise data included (blue points)	18
6.1	Liberty Bridge (Szabadság híd) with boats in the background	21
6.2	An image sequence showing the registration process.	21
6.3	A large registration of Kende street showing only the building class. Scan sequence starts at the middle of the street and runs to its end. The point cloud is made out of 230 point clouds, the full cloud contains more than 20 million points.	22
6.4	Registered street segment. For registration, 30 point clouds were used, final cloud contains almost 10 million points.	23

8.1	Cars with a layer of "carpet" under them. Blue color means it is classified as sparse/noise data	28
8.2	Cars without the detected road patches. It can be seen that this simple removal method is not sufficient to remove all the misclassified points	28
8.3	A clusterization result. Each color represents a separate cluster (colors do not have any particular meaning, they were assigned randomly). .	30
8.4	A triangulation result. It depicts a building on Moricz square.	30
8.5	An example of the tilted configuration. It depicts the Klotild Palace, near the Elizabeth bridge in Budapest	31
8.6	An example of how moving persons look like in the registered point cloud (the two blue boxes) versus a person who was standing still during the whole scanning process (red box)	32
8.7	Another example of moving objects appearing as a smudge in the registered cloud. Here, a car turned onto the main street after stopping at the red light	32
8.8	An example of vegetation removal	33

Bibliography

- [1] T. Luettel, M. Himmelsbach, and H.-J. Wuensche. Autonomous ground vehicles: Concepts and a path to the future. *Proceedings of the IEEE*, 100(Special Centennial Issue):1831–1839, 13 2012.
- [2] Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *International journal of computer vision*, 13(2):119–152, 1994.
- [3] N. Wojke and M. Haselich. Moving vehicle detection and tracking in unstructured environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3082–3087, may 2012.
- [4] A. Azim and O. Aycard. Detection, classification and tracking of moving objects in a 3d environment. In *Intelligent Vehicles Symposium (IV), 2012 IEEE*, pages 802–807, june 2012.
- [5] B. Douillard, J.P. Underwood, N. Kuntz, V. Vlaskine, A. Quadros, P. Morton, and A. Frenkel. On the segmentation of 3D LIDAR point clouds. In *International Conference on Robotics and Automation (ICRA)*, Shanghai, China, 2011.
- [6] B. Douillard, A. Quadros, P. Morton, J.P. Underwood, M. De Deuge, S. Hugosson, M. Hallstrom, and T. Bailey. Scan segments matching for pairwise 3d alignment. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3033–3040, may 2012.
- [7] A. Makadia, A.I. Patterson, and K. Daniilidis. Fully automatic registration of 3d point clouds. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 1, pages 1297–1304, june 2006.
- [8] J. Behley, V. Steinhage, and A.B. Cremers. Performance of histogram descriptors for the classification of 3d laser range data in urban environments. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 4391–4398, may 2012.

- [9] A. J. Quadros, J. P. Underwood, and B. Douillard. An occlusion-aware feature for range images. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4428–4435, St. Paul, MN, USA, 2012.
- [10] C. H. Shen, S. S. Huang, H. Fu, and S. M. Hu. Adaptive partitioning of urban facades. In *SIGGRAPH Asia Conference*, pages 184:1–184:10, New York, NY, USA, 2011. ACM.
- [11] R. Wang, J. Bach, J. Macfarlane, and F.P. Ferrie. A new upsampling method for mobile lidar data. In *Applications of Computer Vision (WACV), 2012 IEEE Workshop on*, pages 17–24. IEEE, 2012.
- [12] A. Velizhev, R. Shapovalov, and K. Schindler. An implicit shape model for object detection in 3D point clouds. In *ISPRS Congress*, Melbourne, Australia, 2012.
- [13] X. Hu, X. Li, and Y. Zhang. Fast filtering of lidar point cloud in urban areas based on scan line segmentation and gpu acceleration. *IEEE Geoscience and Remote Sensing Letters*, 10(2):308 –312, March 2013.
- [14] M. Magnusson. *The Three-Dimensional Normal-Distributions Transform — an Efficient Representation for Registration, Surface Analysis, and Loop Detection*. PhD thesis, Örebro University, December 2009. Örebro Studies in Technology.
- [15] P. Biber and W. Strasser. The normal distributions transform: A new approach to laser scan matching. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 2743–2748, Las Vegas, USA, October 2003.
- [16] J. Han, D. Kim, M. Lee, and M. Sunwoo. Enhanced road boundary and obstacle detection using a downward-looking LIDAR sensor. *IEEE Transactions on Vehicular Technology*, 61(3):971 –985, march 2012.
- [17] S. Thrun, M. Montemerlo, and A. Aron. Probabilistic terrain analysis for high-speed desert driving. *Proc. Robotics Science and Systems, Philadelphia, PA, USA*, 2006.
- [18] A. Petrovskaya and S. Thrun. Model based vehicle tracking for autonomous driving in urban environments. *Proceedings of Robotics: Science and Systems IV, Zurich, Switzerland*, 34, 2008.
- [19] C. Stiller and J. Ziegler. 3d perception and planning for self-driving and cooperative automobiles. In *Systems, Signals and Devices (SSD), 2012 9th International Multi-Conference on*, pages 1–7. IEEE, 2012.

- [20] W. Zhang. LIDAR-based road and road-edge detection. In *Intelligent Vehicles Symposium (IV), 2010 IEEE*, pages 845 –848, June 2010.
- [21] R. B. Rusu and S. Cousins. 3D is here: Point Cloud Library (PCL). In *IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9-13 2011.
- [22] A. Börcs and Cs. Horváth. Városi környezet automatikus analízise és rekonstrukciója légi lidar mérések alapján, 2011. TDK Dolgozat.