
Analyzing Deep Learning Representations of Point Clouds for Real-Time In-Vehicle LiDAR Perception

Marc Uecker*

FZI Research Center for
Information Technology
uecker@fzi.de

Tobias Fleck

FZI Research Center for
Information Technology
tobias.fleck@fzi.de

Marcel Pflugfelder

Karlsruhe Institute of Technology
marcel.pflugfelder@student.kit.edu

J. Marius Zöllner

Karlsruhe Institute of Technology,
FZI Research Center for
Information Technology
zoellner@fzi.de

Abstract

LiDAR sensors are an integral part of modern autonomous vehicles as they provide an accurate, high-resolution 3D representation of the vehicle’s surroundings. However, it is computationally difficult to make use of the ever-increasing amounts of data from multiple high-resolution LiDAR sensors. As frame-rates, point cloud sizes and sensor resolutions increase, real-time processing of these point clouds must still extract semantics from this increasingly precise picture of the vehicle’s environment. One deciding factor of the run-time performance and accuracy of deep neural networks operating on these point clouds is the underlying data representation and the way it is computed. In this work, we examine the relationship between the computational representations used in neural networks and their performance characteristics. To this end, we propose a novel computational taxonomy of LiDAR point cloud representations used in modern deep neural networks for 3D point cloud processing. Using this taxonomy, we perform a structured analysis of different families of approaches. Thereby, we uncover common advantages and limitations in terms of computational efficiency, memory requirements, and representational capacity as measured by semantic segmentation performance. Finally, we provide some insights and guidance for future developments in neural point cloud processing methods.

1 Introduction

Thanks to a large amount of investment into research from many stakeholders, the field of autonomous driving is progressing rapidly [1]. One area where this is particularly evident is the field of LiDAR processing, which has recently been attracting increasing attention from the computer vision and deep learning community [2]. Meanwhile, the sensor hardware is also evolving. As major players in the industry drive demand for lower-cost, high-resolution sensors, they are becoming more affordable and increasingly widespread [3]. Following this trend, recent research vehicles and prototypes are often equipped with multiple high-resolution LiDAR sensors [4]. Modern sensors are capable of delivering millions of points per second, at frame-rates at or above 10 Hz for each sensor [5].

*Corresponding author

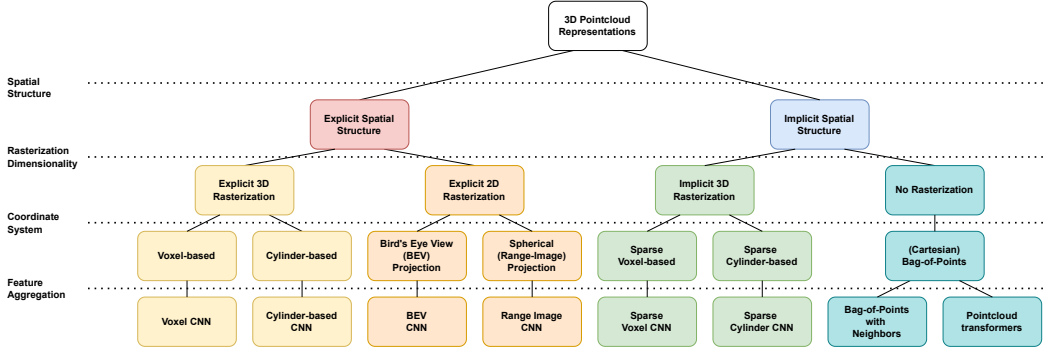


Figure 1: A categorization of common families of approaches using our proposed taxonomy for neural LiDAR pointcloud representations.

These fast and high-resolution sensors produce large amounts of data, which must be processed in real-time to be useful for the autonomous driving functions [6]. For many perception tasks which require semantic or geometric reasoning, such as object detection and semantic segmentation, only deep learning methods provide state-of-the-art processing capabilities [7–10]. However, many deep learning approaches which could be used to process LiDAR point clouds of such scale do not fulfill the real-time inference latency requirements for in-vehicle deployment [8]. We conjecture that the most important design decisions for inference run-time performance hinge on the underlying learned data representation. Multiple papers categorize approaches as either point-based, projection-based or sometimes voxel-based to simplify comparison against state-of-the-art approaches [11–13]. However, this categorization does not capture the full diversity of design decisions made in the development of new architectures. We also found no substantive, objective analysis or comparison of the impact of these design decisions on run-time performance, as each paper focuses on the approach presented. In this work, we present a taxonomy of different architecture designs, based on design decisions regarding the point cloud data representation. We categorize approaches by their choice of explicit or implicit spatial structure, by their choice of internal representation dimensionality, their choice of coordinate space, and finally by their chosen method of feature aggregation.

This taxonomy is described in detail in section 2. Using the introduced taxonomy, we analyze the impact of these design decisions on the run-time performance characteristics in section 3. Finally, based on this analysis, we also provide insights and recommendations for future work in section 4.

2 Taxonomy of neural representations for LiDAR point clouds

In this section we describe our proposed taxonomy in detail. The taxonomy is centered around the design decisions during development, which lead to a final representation of the point cloud inside a deep neural network. In figure 1 we illustrate the categorization of common pointcloud representations using our taxonomy. Notably, the categorization shown in figure 1 is not exhaustive, as there are many possible combinations of choices between the presented design decisions.

Spatial structure The first design decision we observe (fig. 1, first layer) is the choice between an explicit or implicit multi-dimensional spatial arrangement of data in memory. An explicit spatial structure directly encodes positional information in the memory layout of the represented data. Typically, a rasterized representation of a point cloud can be indexed by a point’s coordinate to receive its feature vector. In comparison, an implicit spatial structure stores the points’ feature vectors in a sparse representation. In this case, the points’ coordinates and/or a separate indexing data structure are often stored to encode positional information and accessed to extract neighborhood relations [14–16]. Figure 2 visualizes the difference for an exemplary one-dimensional point cloud.

Rasterization dimensionality The second design decision we observe (fig. 1, second layer) is the dimensionality of the internal mathematical representation of the point cloud. The main varieties we observe are three-dimensional voxel representations, two-dimensional projections of 3D space, and one-dimensional unsorted set- or list-based representations [17, 13, 18]. We refer to the one-dimensional representations as "Bag-of-Points", as their order is typically irrelevant for the operations

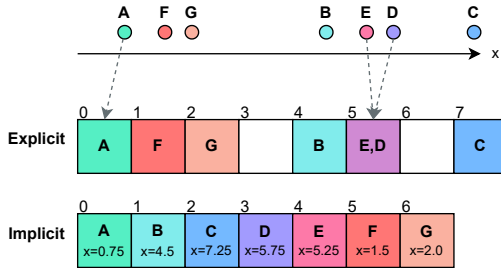


Figure 2: A comparison of explicit and implicit data representations. The example shows a 1-dimensional pointcloud on the top, and an explicit and implicit memory layout at the bottom. An example of aliasing can be observed, as points E and D collide into the same memory cell. Wasted memory space can be seen as empty memory cells.

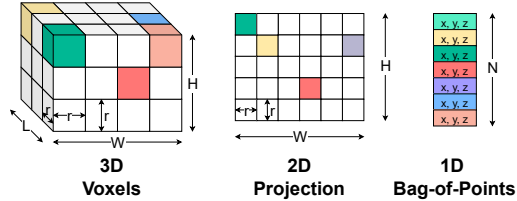


Figure 3: A comparison of 3D voxel representations, 2D Image representations and unordered 1-dimensional Bag-of-Points representations. As the resolution r becomes more fine-grained, the sparsity of rasterized representations increases. However, a coarse resolution may cause multiple points to collide within a single representation cell (as seen in the top right of the 2D projection).

performed on them. The multi-dimensional representations perform a rasterization of the space into a finite number of grid cells, aligning each memory cell to a section of 3D or 2D space. The decision of representation dimensionality is orthogonal to the memory layout, as multi-dimensional rasterizations can also be sparsely stored in one-dimensional data structures [15]. Figure 3 illustrates different rasterization dimensionalities.

Coordinate system The third design decision (fig. 1, third layer) we observe concerns the choice of coordinate system, which is used for rasterizations of multi-dimensional spaces. Rasterization divides 2D or 3D space into chunks of finite size. This partition is typically performed along regular intervals across coordinate axes. Therefore, the coordinate axes chosen for this division also impacts how the resulting representation partitions 3D space. Here, we mainly differentiate between Cartesian coordinate systems, which refer to absolute positions in 3D Euclidean space [12], and polar coordinate systems which refer to locations by combinations of angles and distance measurements [19, 20]. Figure 4 (left) illustrates how different coordinate systems can lead to different rasterizations of two-dimensional space.

Spherical coordinates are an extension of polar coordinates which uses two angles and one distance measurement to index 3D space. Projecting spherical coordinates along the radial axis results in a range image 2D representation. There are also various coordinate systems which combine polar and Cartesian geometry for different axes. Figure 4 shows an example of this: Cylinder coordinates use a polar coordinate system for the X-Y plane and a Cartesian axis for the Z-direction. Similarly, some approaches use a polar coordinate system in a 2D Bird’s eye view (BEV) projection, which projects points along a Cartesian z-axis [19]. For polar coordinate systems, the coordinate origin is typically chosen as the center of the LiDAR sensor in order to minimize aliasing [19, 20].

Feature aggregation As a final design decision (fig. 1, fourth layer), we differentiate approaches by their choice of mathematical operation to be applied to compute the resulting point cloud representation. To compute a feature representation for a point in a 3D point cloud, almost all deep learning approaches aggregate information about its local or global neighborhood. This aggregation typically requires finding other points within the neighborhood of the point whose features are to be computed. Next, these features are aggregated using parametric or non-parametric mathematical operations. The decision for a rasterized representation often directly leads to the use of convolutions [13, 11], although other operations are certainly possible. For Bag-of-Points representations, the choice of the feature aggregation method becomes the main differentiating factor. Some approaches use convolutions in non-rasterized spaces [21, 16, 22], others perform feature aggregation through different variants of weighted or non-weighted pooling of local neighbors [23, 14, 18]. For brevity, we group these into neighbor-based approaches in figure 1 (bottom right). A final group of approaches which exist for smaller point clouds, but have not yet found application in the large-scale point clouds used in autonomous driving are point cloud transformers. These approaches use local or global (point cloud-wide) attention mechanisms to exchange features of points inside a point cloud [24].

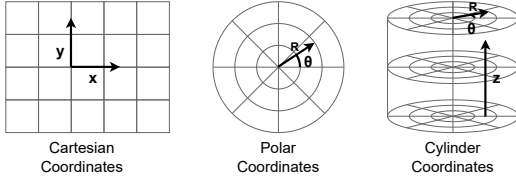


Figure 4: Cartesian and polar coordinates shown in 2-Dimensional space. Cylinder coordinates are one way to combine both approaches in 3-dimensional space.

Table 1: Invariance properties of different representation families when combined with translation-invariant feature aggregation. "Object rotation" refers to the rotation of an object around its center, whereas "Ego rotation" refers to a rotation of the sensor itself.

Representation Family	Translation	Ego rotation	Object rotation
Range Image	✗	yaw, pitch	✗
Bird's Eye View (BEV)	✓	✗	✗
Polar BEV	z-axis	yaw	✗
3D Voxel	✓	✗	✗
3D Cylinder	z-axis	yaw	✗
(Cartesian) Bag-of-Points	✓	✗	✗

We note that there are likely additional possible options for each of the design decisions presented here, and encourage future work to experiment with possible alternatives. We provide this taxonomy as a framework to analyze families of approaches used in practice. In the following, we analyze each design decision from the taxonomy on its own, before continuing to a more high-level discussion of relevant design decisions for memory, run-time and representational capacity.

3 Analysis

In an autonomous vehicle, there are two major resource constraints which affect the perception system, namely memory and computation time. The deployed hardware only provides a finite amount of memory. Therefore, the amount of memory consumed by the point cloud representation should remain within reasonable scales. More critically, computation time is severely limited in a moving vehicle, as any incurred latency makes a perception system's output outdated. There is a certain limit up to which an outdated observation remains useful [6]. Therefore, the latency should be as low as possible. However, a more practical hard constraint is given by the rate that data is produced by the vehicle's sensors. If a perception component's latency is less than the frame rate of the sensor whose data it processes (typically about 10 Hz for current LiDAR sensors [3]), the approach can be considered real-time capable [25].

These resource constraints are contrasted by the representational capacity of a perception model, i.e. its ability to correctly infer information about the environment from the sensor data. In this work, we use semantic segmentation performance as a proxy for representational capacity, based on the widely used SemanticKITTI [8] benchmark. However, the same design decisions and implications presented here can likely be generalized to 3D object detection. As a starting point for a similar analysis, multiple surveys provide an overview of recent 3D object detection methods [26–28]. Typically, larger models have higher representational capacity, but also higher resource utilization [18, 29, 11]. Therefore, the difficulty lies in creating a perception model with high representational capacity, which can still run real-time inference on in-vehicle hardware.

In the following sections, we examine each design decision from the taxonomy presented in section 2 for its impact on memory consumption, run-time latency, and representational capacity.

3.1 Choice of spatial structure (explicit vs implicit)

In this section, we analyze the impact of the choice between an explicit or implicit spatial structure in the memory alignment of the point cloud representation. The choice of spatial structure does not have a direct impact on representational capacity, as the same mathematical operations can be performed on both memory layouts. Therefore, from a mathematical point of view, this choice is an implementation detail. However, in practice, the memory layout can have a large impact on the run-time and memory consumption of a deep learning model.

One main advantage of an explicit spatial structure is that it directly encodes spatial information and therefore spatial locality through memory layout. As a result, finding the feature vector for a particular location in the local neighborhood of a point can run in $O(1)$ time complexity regardless of point cloud size. This is illustrated in figure 2, where the neighborhood of point F in cell 1 can be accessed directly by accessing cells 0 and 2 in the explicit representation (center row). In comparison, to find

Table 2: A comparison of different approaches for 3D semantic segmentation, categorized by their used point cloud representation. mIoU scores and Latency measurements are reported on the SemanticKITTI dataset. The rightmost column provides the source for timing measurements. An asterisk (*) denotes that the pre-processing time for projection (approx. 24 ms following [11]) was added to the number reported in the paper, as this is often omitted from latency comparisons for projection-based approaches [15]. A dagger (†) is used to mark approaches which have not published code, as they are not easily reproducible or usable for downstream tasks.

Method	Category in taxonomy	mIoU [%]	Latency [ms]	# params	Memory Complexity	Source
PointNet [23]	Bag of Points (BoP)	14.6	500	3M	$O(N \cdot d)$	[8]
PointNet++ [18]	BoP + Neighbors	20.1	5900	6M		[8]
TangentConv [21]	BoP + Neighbors	40.9	3000	0.4M	$O(N \cdot d)$	[8]
RandLA-Net [14]	BoP + Neighbors	53.9	880	1.2M		[11]
KPConv [16]	BoP + Neighbors	58.8	279	18.3M		[11]
PVCNN [17]	Voxel CNN	39.0	146	2.5M	$O(\frac{L \cdot W \cdot H}{r^3})$	[11]
MinkowskiNet [30]	Sparse Voxel CNN	63.1	294	21.7M		[11]
SPVCNN [15]	Sparse Voxel CNN	58.8	85	21.8M		[11]
SPVNAS [15]	Sparse Voxel CNN	60.3	89	1.1M	$O(N \cdot d)$	[11]
SPVNAS-A [11]	Sparse Voxel CNN	63.7	110	2.6M		[11]
SPVNAS-B [11]	Sparse Voxel CNN	66.4	259	12.5M		[11]
Cylinder3D [31]	Sparse Cylinder-based	67.8	170		$O(N \cdot d)$	[32]
SqueezeSegV2 [33]	Range Image CNN	39.7	*44	0.9M		[13]
SqueezeSegV3-21 [34]	Range Image CNN	51.6	97	9.4M		[15]
SqueezeSegV3-53 [34]	Range Image CNN	55.9	238	26.2M		[15]
RangeNet53 [13]	Range Image CNN	49.9	102	50.4M	$O(\frac{W \cdot H}{r^2} \cdot d)$	[11]
RangeNet53++ [13]	Range Image CNN	52.2	*126			[13]
SalsaNext [35]	Range Image CNN	59.5	71	6.7M		[11]
Lite-HDSeg† [36]	Range Image CNN	63.8	*74			[36]
PolarNet [19]	Polar BEV CNN	57.2	62	13.6M	$O(\frac{W \cdot H}{r^2} \cdot d)$	[11]
RPVNet† [12]	BoP + Voxel + Range Image CNN	68.3	111	24.8M	$O((N + \frac{W \cdot H}{r^2}) \cdot d)$	[12]
CPGNet† [37]	BoP + BEV + Range Image CNN	68.3	*91			[37]

the neighborhood of point F in the implicit representation (bottom row), one must either traverse the entire list of N points (an $O(N)$ operation), or construct and navigate some external indexing structure, such as a KD-tree [16] or hashmap [15]. We list this neighbor lookup time as t_{lookup} in table 3. As an additional benefit, spatial locality in the memory layout often leads to better cache locality, as neighboring data is often already cached when its neighborhood has been accessed [17]. On GPU accelerators, which are typically used for these models, accessing neighboring memory cells in parallel can also be much faster than random accesses due to architectural details [38, 39].

The main disadvantage of an explicit spatial structure is given through the spatial sparsity of LiDAR data itself. As the LiDAR points are not evenly distributed in 3D space, an explicit representation often incurs a significant fraction of empty memory cells, which represent sections of space with no LiDAR points. Depending on the choice of data representation, a large percentage of memory can be occupied by empty spatial cells, where no LiDAR points reside [17, 13]. By default, implementations of operations such as convolutions do not treat these empty cells any differently than populated cells. Therefore, not only memory, but also computation time is wasted on empty space. This problem of wasted memory and computation cycles grows very quickly as the resolution r of grid cells becomes finer, as can be inferred from figure 3. For 3D voxel-based methods, the number of spatial locations (i.e. voxels) M in the representation, which dictates memory and run-time requirements, grows in the order of $O(\frac{L \cdot W \cdot H}{r^3})$. Here, r is the side-length of voxels, and L , W , and H are the side-lengths of the total volume to be encoded. For 2D pixel-based methods, this scaling is similarly $O(\frac{W \cdot H}{r^2})$, where r is the pixel size, and W and H are again the side-lengths of the 2-dimensional spatial region to be represented. As the resolution r tends to zero, the memory cost and run-time of both approaches grow rapidly. Particularly for explicit 3D voxel approaches, the memory requirements tend to explode very quickly, limiting the practically usable resolution significantly [17, 11].

In summary, as deep neural networks perform aggregation of local neighborhoods of points, the choice of explicit or implicit representation is currently a trade-off between increasing either the number of locations M whose local context is aggregated, or the cost of neighborhood queries t_{lookup} , and therefore neighborhood aggregation for an individual location. We believe future work could de-correlate the memory and time costs of an explicit representation, by operating sparsely on

an explicit memory layout. This would bring both the advantage of spatial locality as well as fast neighborhood queries, at the cost of large memory overhead.

3.2 Choice of rasterization dimensionality (3D voxels vs 2D image vs 1D Bag-of-Points)

The choice of dimensionality for the internal – mathematical – point cloud representation is very important for both run-time and representational capacity. If an explicit representation is chosen, the rasterization dimensionality of the internal representation also affects the scaling of memory costs, as explained in section 3.1.

Constructing and accessing both rasterized and non-rasterized representations takes a non-negligible amount of time, especially in a low-latency use case. The construction and access times of representations depend significantly on the rasterization dimensionality, which we compare qualitatively in table 3.

As described in section 2, deep neural networks on point clouds typically aggregate neighborhood information to infer the features of a point. Depending on whether the representation is rasterized (2D or 3D) or not (Bag of Points), finding neighbors can either be a lookup in a data structure, or a nearest-neighbor search, which significantly affects the neighbor lookup latency t_{lookup} . To reduce neighborhood lookup times, Bag-of-Points representations are typically combined with search tree index structures such as KD-trees [16]. However, even in these tree structures, neighborhood search remains expensive, and their construction takes significant time as well. As a result, non-rasterized representations (BoP + Neighbors in tab. 2) typically take much longer to compute and process their internal representation than rasterized methods (range image or sparse voxel CNN in tab. 2) for similar numbers of parameters and mIoU scores [12]. While accessing rasterized representations is much faster, their construction takes a non-negligible amount of time. For 2D projection-based approaches, Liu et al. [11] report approximately 24 milliseconds of latency for the projection computation alone². For sparse 3D rasterizations, this equates to a hashmap insertion for each point [15].

Due to these run-time constraints, sparse voxel-based approaches tend to have less than ten million parameters (15-20 billion multiply-accumulate (MAC) operations as reported in [15]) before exceeding 100ms of latency. In comparison, 2D CNNs with an explicit representation have a much higher limit of tens of millions of parameters (approximately 200 billion MACs according to [15]). Therefore, sparse 3D convolutional neural networks appear to have much lower parameter budgets than 2D CNNs for similar latency constraints. However, the IoU scores for these 3D models are typically higher than 2D CNNs with comparable numbers of parameters (compare range image vs sparse voxel CNNs in tab. 2). This indicates that the 3D representation enables a higher representational capacity for similar parameter budgets. This is supported by an analysis performed by Triess et al. [20], who find that 2D Range image based CNNs quickly reach diminishing returns with regards to scaling parameter counts. We conjecture that this may be due to inherent invariances present in varying representations, and discuss this point further in section 3.3.

As the dimensionality of a rasterized representation is changed, this also affects both the size and structure of a point’s neighborhood. In a 3D voxel representation, each voxel has up to six direct neighbors, and a neighborhood of 27 voxels is considered for a 3x3x3 convolution. In 2D, each pixel has four direct neighbors, and a 3x3 convolution considers a neighborhood of 9 pixels. The time required for a typical neighborhood aggregation operation scales with the number of accessed neighbors ($n_{\text{neighbors}}$ in tabs. 3 and 4), which is significantly affected by the rasterization dimensionality [15]. This effect on neighborhood size is an important consideration for inference time.

Another important consideration for the choice of rasterization dimensionality is the resulting sparsity of the representation. A typical three-dimensional voxel space will be considerably more sparsely populated with points than a 2D range image projection [17, 20]. This sparsity also affects the choice of explicit or implicit representation, as discussed in section 3.1.

²We suspect this projection time was measured on the CPU and GPU implementations may be faster. However, it is the only source we found for this projection latency.

Table 3: Effects of representation dimensionality and memory layout on run-time determining factors. Actual timings depend on implementation and hardware details.

Representation	M	construction time	t_{lookup}	$n_{\text{neighbors}}$
Explicit 3D	very high	high	low	cubic
Explicit 2D	medium-high	medium	low	quadratic
Implicit 3D	low-medium	low-medium	medium	cubic
Bag-of-Points	low-medium	medium-high	high	configurable

Table 4: Run-time Latency and Parameter scaling of different feature aggregation operations.

Aggregation	Latency	Parameters
Convolution	$O(M \cdot n_{\text{neighbors}} \cdot (t_{\text{lookup}} + d^2))$	$n_{\text{neighbors}} \cdot d^2$
Depth-wise Convolution	$O(M \cdot n_{\text{neighbors}} \cdot (t_{\text{lookup}} + d))$	$n_{\text{neighbors}} \cdot d$
Pooling	$O(M \cdot n_{\text{neighbors}} \cdot (t_{\text{lookup}} + d))$	0
Attention	$O(M \cdot n_{\text{neighbors}} \cdot (t_{\text{lookup}} + d))$	0
Fully-connected	$O(M \cdot (t_{\text{lookup}} + d^2))$	d^2

3.3 Choice of coordinate system (Cartesian vs polar)

The choice of mathematical coordinate system is very important for representational capacity. As the coordinate system imparts a notion of distance, this choice affects what is regarded as the local neighborhood of a point. Since deep learning methods typically aggregate features over this local neighborhood, this also affects representation capacity and, to a lesser extent, run-time latency.

Given that almost all rasterization-based approaches use convolutions to aggregate features, they inherit translation invariance in their internal coordinate system. Depending on this coordinate system, this may lead to different invariances in 3D space, which again affects representational capacity. Table 1 compares the invariances of various rasterized coordinate systems when utilized with a translation-invariant feature aggregation method. There currently appears to exist a mutual exclusivity for translation invariance and invariance against rotations around the sensor. Additionally, no listed rasterized representation provides rotation invariance around arbitrary axes. Based on representational capacity, we find no clearly preferred coordinate system. However, recent results indicate that utilizing multiple representations including coordinate systems with both rotational and translational invariances may provide noticeable improvements in semantic segmentation performance [12, 37].

Since LiDAR sensors scan the environment using radial beams, the density of the resulting point clouds decreases with increasing distance [5]. In Cartesian coordinate systems, this leads to large differences in the local density of points, which may impede training stability. Polar coordinate systems do not suffer from this imbalance, at the cost of utilizing non-uniform spatial sections, which may again negatively affect representational capacity.

Angle-based representations such as range image projections are also sensitive to the choice of coordinate origin. This sensitivity makes their use in multi-sensor applications more challenging, as the points are no longer evenly distributed around a radial center.

3.4 Choice of feature aggregation method

As explained in section 2, deep learning methods often compute the feature representation of a 3D location by aggregating features from its local neighborhood. For Bag-of-Points representations, the choice of feature aggregation method often becomes a distinguishing factor, whereas rasterized (2D or 3D) representations are typically used in conjunction with convolutions. The choice of aggregation method significantly affects the representational capacity and run-time latency.

We observe that each feature aggregation method brings its own trade-off between latency and representational capacity. The representational capacity of a model is typically correlated with its number of parameters [29]. This is also reflected in table 2, as a higher number of parameters is typically correlated with higher performance in semantic segmentation. Therefore, the ratio between the number of parameters and its run-time latency may be a useful indication of the suitability of an aggregation method for the low-latency use case.

In table 4, we list the theoretical run-time and parameter scaling of various common feature aggregation methods used in deep neural networks. Here, M is the number of locations for which features are aggregated, d is the dimensionality of the feature vector for each location, and $n_{\text{neighbors}}$ denotes the number of locations in the neighborhood whose features are aggregated. t_{lookup} represents the time required to find a neighboring location in the representation or an external indexing structure, and load its feature vector from memory. From this table we can observe that convolutions and fully-connected layers have the highest ratio of parameter count to induced latency, especially when using representations with high memory lookup times (t_{lookup}). Fully-connected layers do not aggregate features across neighbors. As such, they need to be combined with a separate neighborhood

aggregation method, which affects the run-time cost of the combination. Another important consideration for selecting a feature aggregation method is the resulting invariance properties of the representation. For instance, convolutions impart translation invariance into a representation, which might positively affect representational capacity [40]. We encourage future work to experiment with aggregation methods which might provide invariances against point cloud scale, density, or rotation around arbitrary axes.

4 Discussion

In section 3, we examined each design decision in the proposed taxonomy in isolation for their effects on memory consumption, latency, and representational capacity. In this section, we take an opposite view point and examine for each outcome – i.e memory consumption, run-time latency, and representational capacity – by which design decisions they are most affected.

Memory The memory consumption of point cloud representations are most prominently affected by the choice of explicit or implicit memory layout. As listed in table 2, the memory consumption characteristics of implicit representations are approximately identical, scaling linearly with the number of points. Therefore, by choosing an implicit memory layout, any point cloud representation can be reduced to $O(N \cdot d)$ memory complexity, excluding any indexing structures. Here, N is the number of points in the point cloud, and d is the dimensionality of the feature vector for each point. As discussed in section 3.1, representations with an explicit spatial structure are often less memory-efficient due to empty spatial sections being represented in memory. The memory consumption of rasterized representations is typically dominated by the most fine-grained grid resolution used. Therefore, it may be useful for certain applications to mix sparse implicit memory layouts for fine-grained rasterizations and use explicit low-resolution feature maps for coarse grid representations, which may be faster to operate on while incurring little sparsity. As mentioned in section 3.2 and section 3.3, some explicit representations incur much less sparsity than others, and therefore can represent the entire point cloud more precisely without much additional memory cost. Namely, range image based 2D rasterizations can perform their projection in such a way that their memory representation incurs little to no sparsity [20]. In contrast, Cartesian representations typically incur the highest sparsity, as spatial cells in regions far away from the sensor – i.e. with low point density – are still represented by fine-grained grid cells. Therefore, the choice of coordinate system should also be considered for explicit 2D or 3D rasterizations.

Latency The run-time latency incurred during the computation of point cloud representations is affected by several design decisions. As features are aggregated element-wise for each spatial location in the point cloud representation, the number of spatial locations is an important factor. For implicit spatial representations, this is typically the number of points in the point cloud, whereas for explicit representations it depends on the rasterization resolution. Given a sufficiently fine spatial resolution, an implicit representation will typically incur far less spatial locations, and thus be much faster to compute.

The next factor to consider is the initial latency overhead for the construction of the representation. This overhead includes the cost of projection, rasterization, but also the cost of constructing additional indexing structures such as KD-trees or hashmaps for implicit representations [16, 15]. As the overhead is typically unaffected by the scale of the following model, it is particularly important to consider for small models, where it may cause a significant fraction of total incurred latency [11]. Another important consideration for latency is the duration of feature aggregation, as discussed in section 3.4. This time is affected by the number of neighbors to aggregate, i.e. representation dimensionality and receptive field for aggregation. However, it is also significantly affected by neighborhood lookup times. This can be seen particularly for the un-structured Bag-of-Points representations in table 2. Therefore, the choice of explicit vs implicit representation again significantly affects run-time latency. Finally, the run-time latency also scales with the size of the model used. As models utilize larger feature dimensionalities, more parameters and more computation, the run-time scales accordingly. As seen in table 2, the scaling of run-time against the number of parameters is highly affected by the choice of specific representation. As we observe from this comparison, rasterized 2D CNNs typically have some initial overhead, but their latency scales least steeply when increasing their parameter count. If slightly higher latency is acceptable, sparse 3D CNNs also provide a highly performant alternative [20, 11].

It is important to note that latency is also dependent on the hardware configuration deployed in the vehicle. This includes both the amount of parallel processing available in GPU-based, FPGA or ASIC hardware for processing, but also includes any ASICs for pre-processing on the sensor's chip itself. As some sensors provide different "raw" representations in their data output, transformations to particular input representations for a deep learning model may be more or less costly.

Capacity The representational capacity of point cloud representations is significantly affected by the number of parameters in the model, which can be scaled using the dimensionality d of feature vectors in the representation. Another important consideration however is how directly the feature aggregation methods can infer the underlying 3D structure from the representation. As seen in table 2, some sparse 3D voxel-based approaches require less parameters to achieve the same segmentation performance as projection-based approaches. To resolve this issue, some 2D projection-based approaches have added three-dimensional post-processing steps, which can provide improvements to segmentation performance, but also incur run-time latency [13]. Therefore, the choice of representation dimensionality and coordinate system is important for representational capacity. As discussed in section 3.3, different coordinate systems also provide different invariances when combined with translation-invariant operations. Whether polar or Cartesian coordinates are best for representational capacity is still not fully clear. However, recent results suggest that there may be added benefit from combining polar and Cartesian representations to benefit from invariances to both rotation and translation (see bottom rows in tab. 2) [12, 37].

When using a vehicle setup with multiple LiDAR sensors, one should also carefully consider the coordinate origin and rasterization resolution of the chosen coordinate system. This is especially the case for spherical range image projections and other polar coordinate systems.

Orthogonal factors such as loss functions and data augmentation have also been shown to bring measurable benefits to representational capacity for little inference time cost [41, 20, 35, 32]. Future work is therefore still required in order to disentangle the representational capacity of point cloud representations from their training methodology.

5 Conclusion

In this work, we presented a taxonomy covering the most significant design decisions regarding the LiDAR point cloud representation used in deep neural networks. By categorizing recent state-of-the-art methods for 3D semantic segmentation, we uncover common trends which can be traced back to common design decisions regarding the representational capacity. We dive deeply into the design decisions made during the development of point cloud representations, and analyze the impact of each design decision on memory consumption, run-time latency and representational capacity. Finally, we discuss for each of the limitations present in the autonomous vehicle use-case, which design decisions can be altered to lighten the load of latency- or memory-constraints, or increase representational capacity. The presented taxonomy can be used as a structural basis for future research in the domain of deep learning on point clouds. It can further be used to classify different model designs and to reason about their memory consumption, inference time and representational capacity during the design phase of the model.

Acknowledgments and Disclosure of Funding

We would like to thank Daniel Bogdoll and Sven Ochs for for their helpful feedback and discussions during the development of this work.

The research leading to these results was conducted within the project KISME (Artificial Intelligence for selective near-real-time recordings of scenario and maneuver data in testing highly automated vehicles) and was funded by the German federal ministry for economic affairs and energy. Responsibility for the information and views set out in this publication lies entirely with the authors.

References

- [1] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A Survey of Autonomous Driving: Common Practices and Emerging Technologies. *IEEE access : practical innovations, open solutions*, 8, 2020.
- [2] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep Learning for 3D Point Clouds: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43, 2021.
- [3] Ricardo Roriz, Jorge Cabral, and Tiago Gomes. Automotive LiDAR Technology: A Survey. *IEEE Transactions on Intelligent Transportation Systems*, 23, 2022.
- [4] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next Generation Datasets for Self-Driving Perception and Forecasting. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS Datasets and Benchmarks 2021)*, 2021.
- [5] Alexander Carballo, Jacob Lambert, Abraham Monrroy, David Wong, Patiphon Narksri, Yuki Kitsukawa, Eijiro Takeuchi, Shinpei Kato, and Kazuya Takeda. LIBRE: The Multiple 3d Lidar Dataset. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, 2020.
- [6] Kieran Strobel, Sibozhu, Raphael Chang, and Skanda Koppula. Accurate, Low-Latency Visual Perception for Autonomous Racing: Challenges, Mechanisms, and Practical Solutions. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [7] Anh Nguyen and Bac Le. 3D Point Cloud Segmentation: A Survey. In *2013 6th IEEE Conference on Robotics, Automation and Mechatronics (RAM)*, 2013.
- [8] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A Dataset for Semantic Scene Understanding of Lidar Sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [9] Yuxing Xie, Jiaojiao Tian, and Xiao Xiang Zhu. Linking Points With Labels in 3D: A Review of Point Cloud Semantic Segmentation. *IEEE Geoscience and Remote Sensing Magazine*, 8, 2020.
- [10] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. Nuscenes: A Multimodal Dataset for Autonomous Driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [11] Zhijian Liu, Haotian Tang, Shengyu Zhao, Kevin Shao, and Song Han. PVNAS: 3D Neural Architecture Search with Point-Voxel Convolution. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 2021.
- [12] Jianyun Xu, Ruixiang Zhang, Jian Dou, Yushi Zhu, Jie Sun, and Shiliang Pu. RPNNet: A Deep and Efficient Range-Point-Voxel Fusion Network for LiDAR Point Cloud Segmentation. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [13] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. RangeNet ++: Fast and Accurate LiDAR Semantic Segmentation. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [14] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [15] Haotian Tang, Zhijian Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching Efficient 3D Architectures with Sparse Point-Voxel Convolution. In *European Conference on Computer Vision*, 2020.
- [16] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, Francois Goulette, and Leonidas J. Guibas. KPConv: Flexible and Deformable Convolution for Point Clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [17] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-Voxel CNN for Efficient 3D Deep Learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

- [18] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [19] Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, and Hassan Foroosh. Polarnet: An Improved Grid Representation for Online Lidar Point Clouds Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [20] Larissa T. Triess, David Peter, Christoph B. Rist, and J. Marius Zöllner. Scan-Based Semantic Segmentation of LiDAR Point Clouds: An Experimental Study. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, 2020.
- [21] Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. Tangent Convolutions for Dense Prediction in 3D. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [22] Wei Song, Zhen Liu, Ying Guo, Su Sun, Guidong Zu, and Maozhen Li. DGPolarNet: Dynamic Graph Convolution Network for LiDAR Point Cloud Semantic Segmentation on Polar BEV. *Remote Sensing*, 14, 2022.
- [23] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [24] Dening Lu, Qian Xie, Mingqiang Wei, Linlin Xu, and Jonathan Li. Transformers in 3D Point Clouds: A Survey. *arXiv preprint arXiv:2205.07417*, 2022.
- [25] Ilias Papadeas, Lazaros Tsochatzidis, Angelos Amanatiadis, and Ioannis Pratikakis. Real-Time Semantic Image Segmentation with Deep Learning for Autonomous Driving: A Survey. *Applied Sciences*, 2021.
- [26] Eduardo Arnold, Omar Y. Al-Jarrah, Mehrdad Dianati, Saber Fallah, David Oxtoby, and Alex Mouzakitis. A Survey on 3D Object Detection Methods for Autonomous Driving Applications. *IEEE Transactions on Intelligent Transportation Systems*, 2019.
- [27] Rui Qian, Xin Lai, and Xirong Li. 3D Object Detection for Autonomous Driving: A Survey. *Pattern Recognition*, 2022.
- [28] Georgios Zamanakos, Lazaros Tsochatzidis, Angelos Amanatiadis, and Ioannis Pratikakis. A Comprehensive Survey of LIDAR-based 3D Object Detection Methods with Deep Learning for Autonomous Driving. *Computers & Graphics*, 2021.
- [29] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A Convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [30] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4D Spatio-temporal Convnets: Minkowski Convolutional Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [31] Hui Zhou, Xinge Zhu, Xiao Song, Yuexin Ma, Zhe Wang, Hongsheng Li, and Dahua Lin. Cylinder3d: An Effective 3d Framework for Driving-Scene Lidar Semantic Segmentation. *arXiv preprint arXiv:2008.01550*, 2020.
- [32] Yuenan Hou, Xinge Zhu, Yuexin Ma, Chen Change Loy, and Yikang Li. Point-to-Voxel Knowledge Distillation for LiDAR Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [33] Bichen Wu, Xuanyu Zhou, Sicheng Zhao, Xiangyu Yue, and Kurt Keutzer. Squeezesegv2: Improved Model Structure and Unsupervised Domain Adaptation for Road-Object Segmentation from a Lidar Point Cloud. In *2019 International Conference on Robotics and Automation (ICRA)*, 2019.
- [34] Chenfeng Xu, Bichen Wu, Zining Wang, Wei Zhan, Peter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. Squeezesegv3: Spatially-adaptive Convolution for Efficient Point-Cloud Segmentation. In *European Conference on Computer Vision*, 2020.
- [35] Tiago Cortinhal, George Tzelepis, and Eren Erdal Aksoy. SalsaNext: Fast, Uncertainty-Aware Semantic Segmentation of LiDAR Point Clouds. In *International Symposium on Visual Computing*, 2020.
- [36] Ryan Razani, Ran Cheng, Ehsan Taghavi, and Liu Bingbing. Lite-HDseg: LiDAR Semantic Segmentation Using Lite Harmonic Dense Convolutions. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

- [37] Xiaoyan Li, Gang Zhang, Hongyu Pan, and Zhenhua Wang. CPGNet: Cascade Point-Grid Fusion Network for Real-Time LiDAR Semantic Segmentation. *arXiv preprint arXiv:2204.09914*, 2022.
- [38] Eladio Gutierrez, Sergio Romero, Maria A. Trenas, and Emilio L. Zapata. Memory Locality Exploitation Strategies for FFT on the CUDA Architecture. In *High Performance Computing for Computational Science - VECPAR 2008*, volume 5336. 2008.
- [39] Jun Liu, Wei Ding, Ohyoung Jang, and Mahmut Kandemir. Data Layout Optimization for GPGPU Architectures. *ACM SIGPLAN Notices*, 48, 2013.
- [40] Larissa Triess, David Peter, and Johann Marius Zöllner. Semi-Local Convolutions for LiDAR Scan Processing. In *I (Still) Can't Believe It's Not Better! NeurIPS 2021 Workshop*, 2021.
- [41] Zhi Zhang, Tong He, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of Freebies for Training Object Detection Neural Networks. *arXiv preprint arXiv:1902.04103*, 2019.