Super Sparse 3D Object Detection

Lue Fan, Yuxue Yang, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang

Abstract—As the perception range of LiDAR expands, LiDAR-based 3D object detection contributes ever-increasingly to the long-range perception in autonomous driving. Mainstream 3D object detectors often build dense feature maps, where the cost is quadratic to the perception range, making them hardly scale up to the long-range settings. To enable efficient long-range detection, we first propose a fully sparse object detector termed FSD. FSD is built upon the general sparse voxel encoder and a novel sparse instance recognition (SIR) module. SIR groups the points into instances and applies highly-efficient instance-wise feature extraction. The instance-wise grouping sidesteps the issue of the center feature missing, which hinders the design of the fully sparse architecture. To further enjoy the benefit of fully sparse characteristic, we leverage temporal information to remove data redundancy and propose a super sparse detector named FSD++. FSD++ first generates residual points, which indicate the point changes between consecutive frames. The residual points, along with a few previous foreground points, form the super sparse input data, greatly reducing data redundancy and computational overhead. We comprehensively analyze our method on the large-scale Waymo Open Dataset, and state-of-the-art performance is reported. To showcase the superiority of our method in long-range detection, we also conduct experiments on Argoverse 2 Dataset, where the perception range ($200m$) is much larger than Waymo Open Dataset ($75m$). Code is open-sourced at https://github.com/tusen-ai/SST.

Index Terms—3D object detection, LiDAR, autonomous driving, sparse, Waymo Open Dataset, instance segmentation, temporal fusion, point clustering.

✦

1 INTRODUCTION

AUTONOMOUS driving systems are eager for efficient long-range perception, especially in high-speed scenarios. Current LiDAR-based 3D object detectors usually convert sparse features into dense feature maps for further feature extraction and prediction, which we name as **dense detectors**. Dense detectors perform well on current popular benchmarks [1], [2], [3], where the perception range is relatively short (less than 75 meters). However, it is impractical to scale the dense detectors to the long-range setting (more than 200 meters, Fig. 1). In such settings, the computational and spatial complexity on dense feature maps is quadratic to the perception range. Fortunately, the sparsity of LiDAR point clouds also increases as the perception range extends (see Fig. 1), and the calculation on the unoccupied area is essentially unnecessary. Given the inherent sparsity, an essential solution for efficient long-range detection is to remove the dense feature maps and make the network architectures *fully sparse*.

However, removing the dense feature map is non-trivial since it plays an indispensable role in current designs. Commonly adopted sparse voxel encoders [5], [6], [7] only extract the features on the non-empty voxels. Without dense feature maps, the object centers are usually empty, especially for large objects. We name this issue as "**Center Feature Missing (CFM)**" (Fig. 2). CFM significantly weakens the representation power of the center voxels, even making the center feature empty in some extreme cases like super large vehicles. However, almost all popular voxel or pillar based detectors [5], [6], [8], [9], [10] adopt center-based
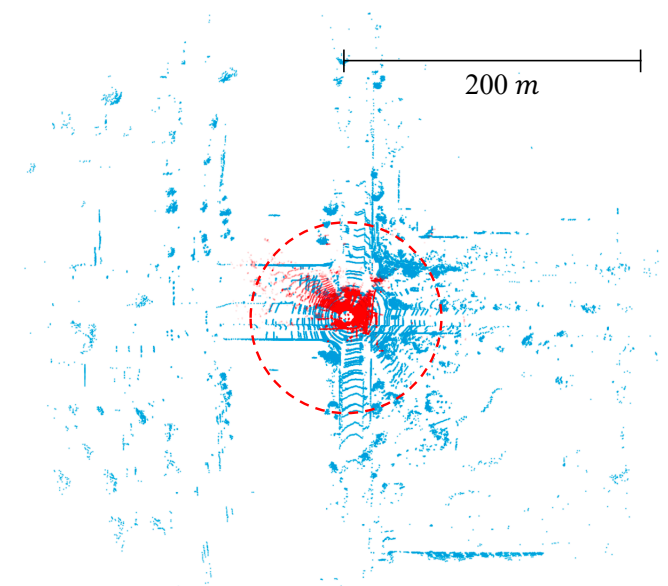

Fig. 1. Short-range point clouds (red, from KITTI [2]) v.s. long-range point clouds (blue, from Argoverse 2 [4]). The radius of the red circle is 75 meters. The sparsity quickly increases as the range extends.

assignment and rely on the center feature since it is an ideal representation of the whole object. So they have to first convert sparse voxels to dense feature maps in Bird's Eye View after the sparse voxel encoder. Then they resolve the CFM issue by applying convolutions on the dense feature maps to diffuse features to instance centers, which we name as **feature diffusion** (Fig. 2).

To properly eliminate the dense feature map, we investigate the purely point-based detectors because they are naturally fully sparse. However, two drawbacks limit the usage of point-based methods. (1) The time-consuming

- Lue Fan and Yuxue Yang and Zhaoxiang Zhang are with Center for Research on Intelligent Perception and Computing (CRIPAC), National Laboratory of Pattern Recognition (NLPR), Institute of Automation, Chinese Academy of Sciences (CASIA), Beijing 100190, China. E-mail: {fanlue2019, yangyuxue2023, zhaoxiang.zhang}@ia.ac.cn.
- Feng Wang and Naiyan Wang are with TuSimple, Beijing 100020, China. E-mail: {feng.wff, winsty}@gmail.com.
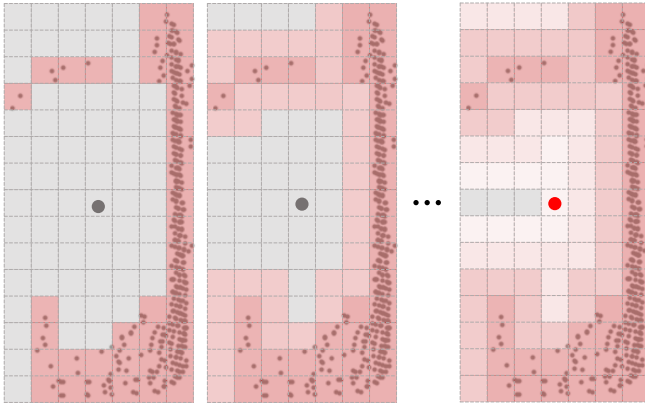
Fig. 2. Illustration of **center feature missing** and **feature diffusion** on dense feature maps from Bird's Eye View. The empty instance center (red dot) is filled by the features diffused from occupied voxels (with LiDAR points), after several convolutions.

neighborhood query [11] is the long-standing difficulty to apply it to large-scale point cloud (more than 100K points). (2) To reduce the computational overhead, point-based methods aggressively downsample the whole scene to a fixed number of points. The aggressive downsampling leads to inevitable information loss and insufficient recall of foreground objects [12], [13], especially for small ones. As a result, very few purely point-based detectors have reached state-of-the-art performance in the recent benchmarks with large-scale point clouds.

In this paper, we first propose **Fully Sparse Detector (FSD)** to sidestep the issue of center feature missing. FSD is built upon a general sparse voxel encoder [5], [6], [7] for voxel/point feature extraction. Then FSD groups the points into an instance, and further extract the instance-level feature and predict a single bounding box from the integrated instance feature, via a novel **Sparse Instance Recognition (SIR)** module. In this way, predictions are made from the whole instance feature instead of the weak or missed center feature. As a point-based module, SIR has several desired properties: (1) Unlike previous point-based modules, SIR simply treats instances as groups, and does not apply the time-consuming neighborhood query for further grouping. (2) Similar to dynamic voxelization [14], SIR leverages *dynamic broadcast/pooling* for tensor manipulation to avoid point sampling or padding. (3) Since the group in SIR covers the whole instance, it builds a sufficient receptive field regardless of the physical size of the instance.

To unleash the full potential of FSD, we further utilize temporal information and propose a **Super Sparse 3D Object Detector**, named FSD++. FSD++ is inspired by human visual behavior: human is sensitive to and focuses on dynamic parts of the physical world. In particular, FSD++ utilizes ego-motion to remove the static parts containing heavy temporal redundancy, while only retaining the informative dynamic parts. We name the detected dynamic parts as residual points since the process is similar to applying the difference between frames. In this way, we create a *super sparse* point cloud consisting of residual points and a small number of past foreground points from history predictions. FSD++ then takes the super sparse point cloud as input, achieving a very

efficient detection framework with temporal fusion. We owe the credit of the high efficiency to the synergy of the fully sparse characteristic and the super sparse input. We list our contributions as follows.

- We introduce the concept of Fully Sparse Detector (FSD), which is the essential solution for efficient long-range LiDAR detection. We further propose Sparse Instance Recognition (SIR) to sidestep the issue of Center Feature Missing (CFM) in sparse feature maps. Combining SIR with general sparse voxel encoders, we develop an efficient and effective FSD implementation.
- Based on FSD, we further present the FSD++ framework, which aggregates a super sparse point cloud from multi-frames as input, yet removing the temporal redundancy of point clouds. The proposed framework uncovers the untapped potential of sparse architecture. We hope our efforts attract the attention of the community to fully sparse architecture.
- FSD achieves state-of-the-art performance on the competitive Waymo Open Dataset. Besides, we further apply our method to the recently released Argoverse 2 dataset to demonstrate the superiority of FSD in long-range detection, where FSD is much more efficient than its dense counterparts. FSD++ achieves comparable performance with mainstream state-of-the-art multi-frame detectors with minimal additional overhead compared with single-frame input.

2 RELATED WORK

In reviewing the evolution of LiDAR-based 3D object detectors, the previous methods could be categorized into three types by their spatial sparsity: dense detectors, sparse detectors, and semi-dense detectors. Below, we provide a brief revisit of previous arts according to spatial sparsity.

2.1 Voxel-based Dense Detectors

Pioneering work 3DFCN [15] and VoxelNet [16] use dense convolution for voxel feature extraction. They bring convolutional neural networks to the field of LiDAR-based 3D object detection and achieve competitive results at the time. However, it is inefficient to apply dense convolution to 3D voxel representation. MV3D [17], PIXOR [18], and PointPillars [19] adopt 2D dense convolution in Bird's Eye View (BEV) feature maps achieving significant efficiency improvement. We refer to such detectors as **dense detectors** since they convert the sparse point cloud into dense feature maps.

2.2 Point-based Sparse Detectors

Since PointNet [20] and PointNet++ [11] shed light on the deep learning for 3D point sets, a series of point-based detectors have emerged. These purely point-based detectors are born to be fully sparse. PointRCNN [21] is the pioneering work of this line of work. 3DSSD [12] accelerates the point-based method by removing the feature propagation layer and refinement module. VoteNet [22] first makes a center voting and then generates proposals from the voted center achieving better accuracy. Albeit many methods [12], [13], [23] have tried to accelerate the point-based method, the

time-consuming point sampling and neighborhood query are still unaffordable in large-scale point clouds (more than 100k points per scene). So current benchmarks [1], [3] with large-scale point clouds are still dominated by voxel-based dense/semi-dense detectors [10], [24], [25].

2.3 Semi-dense Detectors

Different from dense detectors, **semi-dense detectors** incorporate both sparse features and dense features. SECOND [5] employs sparse convolution to extract the sparse voxel features in 3D space, which then are converted to dense feature maps in BEV to enlarge the receptive field and integrate with 2D detection head [26], [27], [28]. Based on SECOND-style semi-dense detectors, a series of work [29], [30], [31] made further improvements on the single-stage paradigm. And other methods attach a second stage for fine-grained feature extraction and proposal refinement [7], [8], [10], [32], achieving superior performance. Although semi-dense detectors become dominating in academia and industry, related research has stagnated here because the semi-dense detector cannot be trivially lifted to be fully sparse as we discussed in Sec. 1.

3 FSD: FULLY SPARSE 3D OBJECT DETECTION

3.1 Overall Architecture

Following the motivation of instances as groups, we have four steps to build the fully sparse detector (FSD): 1) We first utilize a sparse voxel encoder [5], [6], [7] to extract voxel features and casts votes for object centers(Sec. 3.2). 2) Instance Point Grouping groups foreground points into instances based on the voting results (Sec. 3.2). 3) Given the grouping results, Sparse Instance Recognition (SIR) module extracts instance/point features and generates proposals (Sec. 3.3). 4) The proposals are utilized to correct the point grouping and refine the proposals iteratively (Sec. 3.4).

3.2 Instance Point Grouping

Classification and Voting   We first extract voxel features from the point cloud with a sparse voxel encoder, such as sparse attention blocks in SST [6] or sparse convolution encoder. Then we build point features by concatenating voxel features and the offsets from points to their corresponding voxel centers. These point features are passed into two heads for foreground classification and center voting. The voting is similar to VoteNet [22], where the model predicts the offsets from foreground points to corresponding object centers. L1 loss [27] and Focal Loss [33] are adopted as voting loss L_{vote} and semantic classification loss L_{sem}.

Connected Components Labeling (CCL)   To group points into instances, we regard all the predicted centers (red dots in Fig. 3) as vertices in a graph. Two vertices are connected if their distance is smaller than a certain threshold. Then a connected component in this graph can be viewed as an instance, and all points voted to this connected component share a group ID. Unlike the ball query in VoteNet, our CCL-based grouping avoids fragmented instances in most cases. Although there are many elaborately designed instance grouping methods [34], [35], [36], we opt for the simple CCL because it is adequate in our design and can be implemented by the efficient Union-Find algorithm [37] in parallel.

3.3 Sparse Instance Recognition

3.3.1 Preliminaries: Dynamic Broadcast/Pooling

Given `N` points belong to `M` groups, we define their corresponding group ID array as I in the shape of $[\text{N},]$ and their feature array as F in the shape of $[\text{N}, \text{C}]$, where `C` is the feature dimensions. $F^{(i)}$ is the feature array of points belonging to the i-th group. Dynamic pooling aggregates each $F^{(i)}$ into one group feature $\boldsymbol{g}_i$ of shape $[\text{C},]$. Thus we have $\boldsymbol{g}_i = p(F^{(i)})$, where p is a symmetrical pooling function. The dynamic pooling on all group features G of shape $[\text{M}, \text{C}]$ is formulated as $G = p(F, I)$. The dynamic broadcast can be viewed as the inverse operation to dynamic pooling, which broadcasts $\boldsymbol{g}_i$ to all the points in the i-th group. Since the broadcasting is essentially an indexing operation, we use the indexing notation $[\ ]$ to denote it as $G[I]$, which is in the shape of $[\text{N}, \text{C}]$. Dynamic broadcast/pooling is very efficient because it can be implemented with high parallelism on modern devices and well fits the sparse data with dynamic size. We provide an efficient implementation and runtime evaluation in Appendix A.

The prerequisite of dynamic broadcast/pooling is that each point *uniquely* belongs to a group, i.e. groups should not overlap with each other. Thanks to the fact that there is no overlap among instances in the real 3D world, the groups do not overlap with each other naturally.

3.3.2 Formulation of Sparse Instance Recognition

After grouping points into instances in Sec. 3.2, we can directly extract instance features via some basic point-based networks like PointNet, DGCNN, etc. There are three elements to define a basic point-based module: *group center*, *pair-wise feature* and *group feature aggregation*.

Group center   The group center is the representative point of a group. For example, in the ball query, it is the local origin of the sphere. In SIR, the group center is defined as the centroid of all voted centers in a group.

Pair-wise feature   defines the way to pair group center and neighbor points input for group-aware neighbor point feature extraction. SIR adopts two kinds of pair-wise features: 1) the relative coordinate between the group center and each point, 2) the concatenation of the group feature and each point feature. Taking feature concatenation as an example and using the notations in 3.3.1, the pair-wise feature can be denoted as $\text{CAT}(F, G[I])$, where CAT is channel concatenation.

Group feature aggregation   In a group, a pooling function is used to aggregate neighbor features. SIR applies dynamic pooling to aggregate feature array F. Following the notations in 3.3.1, we have $G = p(F, I)$, where G is the aggregated group features.

Integration   Combining the three basic elements, we could build many variants of point-based operators, such as PointNet [20], DGCNN [38], Meta-Kernel [39], etc. Fig. 4 illustrates the basic idea of how to build an instance-level point operator with dynamic broadcast/pooling. In our design, we adopt the formulation of VFE [16] as the basic structure of SIR layers, which is basically a two-layer PointNet. In the l-th layer of SIR module, given the input point-wise feature array
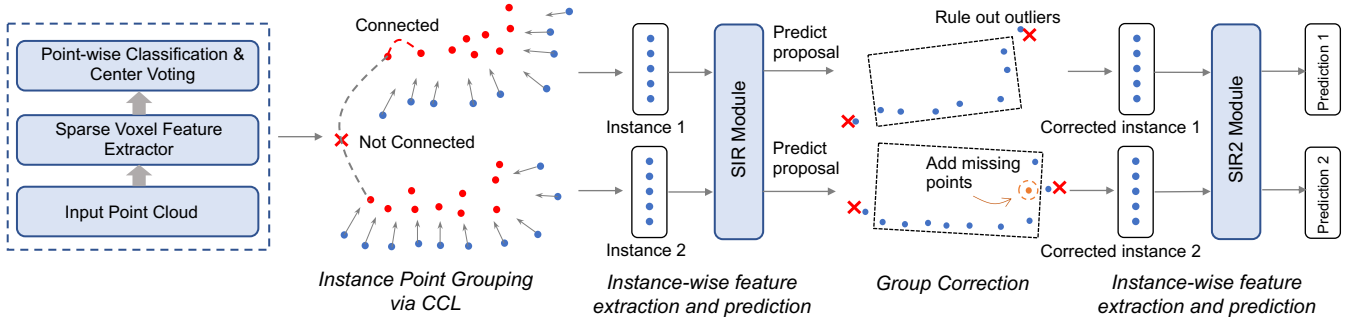
Fig. 3. Overall architecture of FSD. For simplicity, we only use two instances to illustrate the pipeline. Red dots are the voted centers from each LiDAR point (blue dots). The SIR module and the SIR2 module all contain 3 SIR layers.
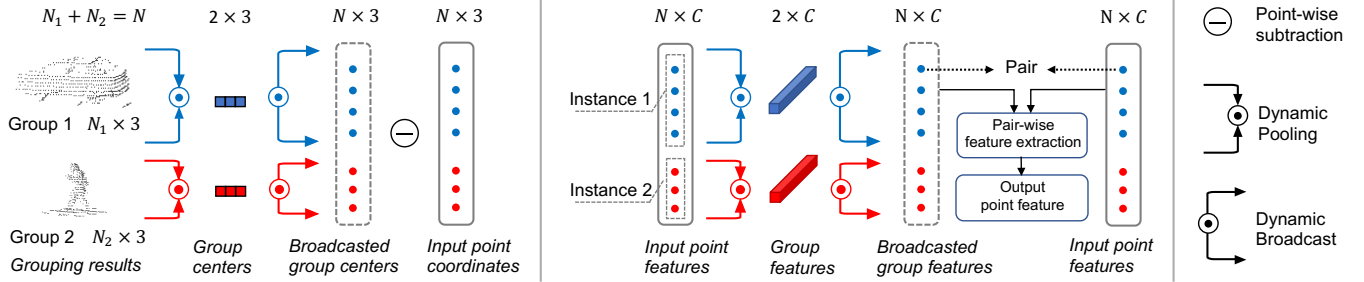

Fig. 4. Illustration of building instance-level point operators with dynamic broadcast/pooling. Best viewed in color. Left: calculating center-to-neighbor offsets given raw point clouds. Right: updating point features. Note that the operation is parallel among all instances.

F_l, point coordinates array X, the voted center X' and group ID array I, the output of l-th layer can be formulated as:

$$F_l' = \texttt{LinNormAct}\left(\texttt{CAT}\left(F_l, X - p_{\text{avg}}(X', I)[I]\right)\right), \quad (1)$$

$$F_{l+1} = \texttt{LinNormAct}\left(\texttt{CAT}\left(F_l', p_{\text{max}}(F_l', I)[I]\right)\right), \quad (2)$$

where $\texttt{LinNormAct}$ is a fully-connected layer followed by a normalization layer [40] and an activation function [41]. The p_{avg} and the p_{max} are average-pooling and max-pooling function, respectively. The output F_{l+1} can be further used as the input of the next SIR layer, so our SIR module is a stack of a couple of basic SIR layers.

3.3.3 Sparse Prediction

With the formulation in Eqn. 1 and Eqn. 2, SIR extracts features of all instances dynamically in parallel. And then SIR makes **sparse prediction** for all groups. In contrast to two-stage sparse prediction, our proposals (i.e., groups) do not overlap with each other. Unlike one-stage dense prediction, we only generate a single prediction for a group, which significantly reduces the cost of prediction head. It is noteworthy that the fully sparse architecture may face a severe imbalance problem: short-range objects contain much more points than long-range objects. Some methods [39], [42] use hand-crafted normalization factors to mitigate the imbalance. Instead, SIR avoids the imbalance because it only generates a single prediction for a group regardless of the number of points in the group. In most cases, a group corresponds to only a single ground truth box.

Specifically, for each SIR layer, there is a $G_l = p_{max}(F_l', I)$ in Eqn. 2, which can be viewed as the group features. We concatenate all G_l from each SIR layer in channel dimension and use the concatenated group features to predict bounding boxes and class labels via MLPs. All the groups whose centers fall into ground-truth boxes are positive samples. For positive samples, the regression branch predicts the offsets from group centers to ground-truth centers and object sizes and orientations. L1 loss [27] and Focal Loss [33] are adopted as regression loss L_{reg} and classification loss L_{cls}, respectively.

3.4 Group Correction

There is inevitable incorrect grouping in the Instance Point Grouping module. For example, some foreground points may be missed, or some groups may be contaminated by background clutter. So we leverage the bounding box proposals from SIR to correct the grouping. The points inside a proposal belong to a corrected group regardless of their previous group IDs. Since a few points may fall into multiple proposals, we simply make copies for these points along with their features and assign different copies to difference proposals. After correction, we apply an additional SIR to these new groups. To distinguish it from the first SIR module, we denote the additional SIR module as SIR2.

SIR2 predicts box residual from the proposal to its corresponding ground-truth box, following many two-stage detectors. To make SIR2 aware of the size and location of a proposal, we adopt the offsets from inside points to proposal boundaries as extra point features following [43]. The regression loss is denoted as $L_{res} = L1(\Delta_{res}, \widehat{\Delta_{res}})$, where Δ_{res} is the ground-truth residual and $\widehat{\Delta_{res}}$ is the predicted residual. Following previous methods [7], [8], the 3D Intersection over Union (IoU) between the proposal and ground-truth serves as the soft classification label in SIR2. Specifically, the soft label q is defined as $q = \min(1, \max(0, 2IoU - 0.5))$, where IoU is

the area of Intersection over Union (IoU) between proposals and corresponding ground-truths. Then cross entropy loss is adopted to train the classification branch, denoted as L_{iou}. Taking all the loss functions in grouping (Sec. 3.2) and sparse prediction into account, we have

$$L_{total} = L_{sem} + L_{vote} + L_{reg} + L_{cls} + L_{res} + L_{iou}, \quad (3)$$

where we omit the weight of each term for simplicity.

3.5 Discussion

The center voting in FSD is inspired by VoteNet [22], while FSD has two essential differences from VoteNet.

- After voting, VoteNet simply aggregates features around the voted centers without further feature extraction. FSD goes beyond this and builds a highly efficient SIR module taking advantage of dynamic broadcast/pooling, allowing for further instance-level feature extraction. Thus, FSD extracts more powerful instance features, which is experimentally demonstrated in Sec. 5.5.
- VoteNet is a typical point-based method. As we discussed in Sec. 1, it aggressively downsamples the whole scene to a fixed number of points for efficiency, causing inevitable information loss. Instead, the dynamic characteristic and efficiency of SIR enable fine-grained point feature extraction from any number of input points without any downsampling. In Sec. 5.5, we showcase the efficiency of our design in processing large-scale point clouds and the benefits of fine-grained point representation.

4 FSD++: FSD WITH SUPER SPARSE INPUT

It is well known that aggregating multiple frames as an input benefits performance. However, naive aggregation could result in a denser point cloud, which slows down the algorithm significantly, especially in the architecture with sparse operations. This motivates us to pursue more sparse input data by removing temporal redundancy from the original point cloud stream. Thanks to the fully sparse characteristic, the fully sparse model could greatly benefit from the increase of sparsity after redundancy removal. Thus a natural question arises: *How can we remove the redundancy while retaining the informative parts in advance?* The similarities between consecutive point cloud frames offer us a potential solution to this question.

In particular, the spatial distribution of points varies continuously and smoothly in a sequence. We name the points that change between consecutive frames as *residual points*. The residual points are informative since they represent new observations in a time step. Combining the residual points and history predictions, detectors have sufficient knowledge to infer about current objects. In this paradigm, the residual points and previous foreground points together form a *super sparse point cloud*. FSD could directly take them as input for much more efficient object detection.

4.1 Residual Points Probing

LiDAR sensors capture plenty of newly observed foreground points at each time step. These points can be attributed to two main sources: (1) objects moving to new positions; (2) occluded regions becoming visible. These newly observed

points are referred to as residual points. Residual points are critical to locate moving objects and detect recently emerged objects. As we mentioned before, the residual points could be detected from the changes of point spatial distribution.

The residual point detection algorithm must fulfill two key requirements. (1) The algorithm is supposed to be robust to tiny disturbances of points, which might be caused by sensing noise or tiny ego-motion estimation error. It is unexpected that such point disturbances are detected as residual points. (2) The algorithm should be highly efficient to handle millions of points from multiple frames. In particular, each frame in WOD contains up to 200,000 points.

Several straightforward solutions meet the first requirement, i.e. ball query or voxelization into dense occupancy maps. A point can be viewed as a residual point if no previous points fall into its neighborhood defined by the ball query radii. The residual points can also be detected by the simple difference between the two dense occupancy maps. Although proper ball query radii or voxel sizes bring robustness to point disturbances, these solutions still come with either high computational complexity ($O(N^2)$) or a huge memory footprint.

To fulfill both of the demands outlined above, we resort to hashing and design an algorithm shown in Algo. 1, named *Residual Points Probing (RPP)*. RPP consists of two steps. (1) It first quantizes the point coordinates into integers. The granularity of quantization controls the robustness to point disturbances. (2) For each point, RPP verifies if it is a residual point by hash probing. Specifically, RPP first builds a hash table from previous quantized points. The key set of the hash table is denoted as $K \subset \mathbb{Z}^3$, which is the quantized integer coordinates. And value set of the hash table is denoted as $V = \{1, 0\}$, where 1 indicates the slot is occupied and 0 indicates the slot is unoccupied. RPP then uses current quantized coordinates to probe the hash table. If a current point hits an unoccupied slot, it is treated as a residual point. Here is a hidden assumption in RPP that we assume two points are the same if they occupy the same voxel after quantization. We adopt the well-known *open addressing* for probing and the *double hashing* as the hash function to reduce hash collisions.

Algorithm 1: Efficient Residual Points Probing

Input: current points P_{cur}, previous points P_{pre}, voxel size s, load factor α, hash function h

Output: Residual points of current frame ΔP_{cur}

$\widehat{P}_{cur} \leftarrow \texttt{Quantize}(P_{cur}, s)$;
$\widehat{P}_{pre} \leftarrow \texttt{Quantize}(P_{pre}, s)$;
Initialize empty hash table T of length $|\widehat{P}_{pre}|/\alpha$;
Initialize empty residual point set ΔP_{cur};
foreach $\widehat{p}_i$ *in* $\widehat{P}_{pre}$ **do**
    $slot_i \leftarrow \texttt{Probe}(T, h(\widehat{p}_i))$;
    if $slot_i$ *is not occupied* **then**
        $slot_i \leftarrow \texttt{occupied flag}$;
foreach $\widehat{p}_i$ *in* $\widehat{P}_{cur}$ **do**
    $slot_i \leftarrow \texttt{Probe}(T, h(\widehat{p}_i))$;
    if $slot_i$ *is not occupied* **then**
        Add p_i to ΔP_{cur};
return ΔP_{cur}

(a) Single-frame points

(b) Multi-frame points

(c) Residual points with change blindness

(d) Residual points w/o change blindness

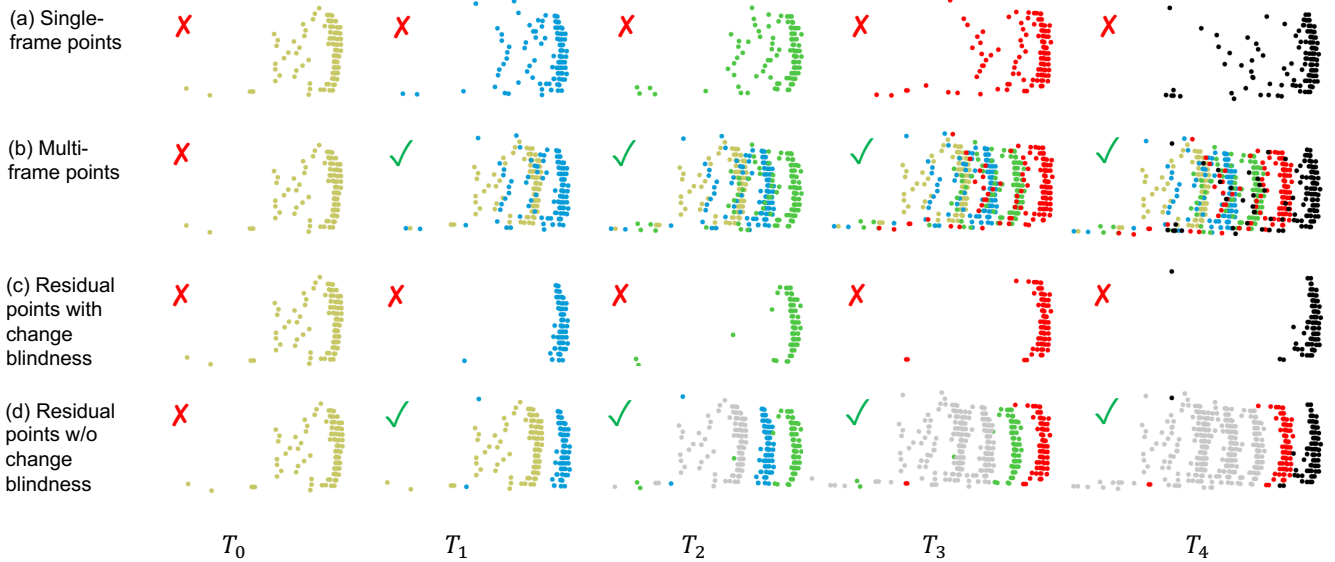T_0 T_1 T_2 T_3 T_4

Fig. 5. Illustration of temporal point manipulations. Different colors indicate points from different time steps. Gray points are the sampled skeleton points. Red cross ✗ means the observation is too weak to be recognized as foreground objects (not really, just for illustration). Green check mark ✓ means the observation is strong enough. (a) Points from a single frame are too weak. (b) After multi-frame point aggregation, the detector generates true positives from time step T_1. (c) Single-frame residual points suffer from change blindness. Since the detector cannot generate a true positive in T_0 for skeleton point sampling, we still only have the weak residual points in T_1. In this way, the detector outputs false negatives all the time. (d) Detector makes the right prediction in T_1 with residual points from two frames (max age is 2). So in the T_2, the detector could leverage the prediction in T_1 for skeleton point sampling, alleviating the change blindness.

Residual Point Probing is efficient in terms of both memory and speed. For instance, we assume there are N unique quantized coordinates in a point cloud clip. If we expect the collision rate less than α, the length of the hash table should be N/α. Empirically, N is around 500,000 in a 5-frame point cloud in WOD. A slot state can be represented as a single bit, and let α equal to 0.1. We have that the memory cost of this hash table is around 0.6MB. Moreover, the probing of each point is independent with each other, allowing for high parallelism in GPU.

Formally, we denote the RPP process as follows:

$$\Delta P_t = P_t - \bigcup_{i=1}^{B} P_{t-i}, \qquad (4)$$

where ΔP_t is the detected residual points in time step t and P_t is the raw points in time step t. The notation "$X - Y$" means removing the intersection of X and Y from X, equivalent to $X \setminus (X \cap Y)$. And the union means point cloud concatenation. B is the number of previous frames used in RPP, which we term as *base frames*.

4.2 Skeleton Point Sampling

Since residual points contain only new observations of the current time step, detectors require additional information from previous frames for sufficient input. To incorporate this historical data, we use previously predicted boxes to crop previous foreground points, while discarding others outside of the boxes. The cropped points are placed into the current frame after ego-motion compensation. However, the foreground points from multiple previous frames are still essentially redundant. Especially, the quite many points on short-range objects from multiple frames could lead to unnecessary overhead.

To reduce the redundancy from multi-frame foreground points, we further sample within these cropped points. Intuitively, we expect the sampled points contain the minimal information models need to make proper predictions. In this sense, we refer to such a minimal subset of cropped points as *skeleton points*, because they depict the basic structure or "skeleton" of objects. Specifically, we try three kinds of sampling methods: *random sampling*, *farthest points sampling*, and *voxel sampling*. All the sampling methods are applied inside the previously predicted bounding boxes. For random sampling and farthest points sampling, we adopt a predefined maximal point threshold N_T. We sample N_T points inside the bounding boxes which contain points more than N_T. For voxel sampling, we adopt dynamic voxelization [14] to voxelize points. All the points falling into a voxel are reduced to a single point by average pooling.

4.3 Treatment to Change Blindness

Theoretically, by combining the skeleton points and residual points, a model is able to make predictions in current frames. However, a phenomenon known as "change blindness" can hinder performance. Change blindness refers to the human visual system's tendency to overlook progressive small changes in a scene, even if the aggregated changes of multiple time steps are significant. A similar issue can occur in our case. Thinking of a vehicle nearly entering into the sensing range of LiDARs in time step t, only a small part of the vehicle can be observed. The detector is very likely to recognize it as background, so RPP will remove these points in time step $t + 1$ and only keep a small number of new points of the vehicle as residual points. In this way, if the vehicle appears slowly, the detector might never recognize it. Fig. 5 demonstrates the change blindness.

To remedy the change blindness, we introduce a *max age* M for residual points. In other words, the detector takes residual points from at most M steps as input. Formally, the detector takes $\bigcup_{i=0}^{M-1} \Delta P_{t-i}$ as accumulated residual points for input in time step t.

4.4 Integrated Super Sparse Input

The input point clouds consist of two parts: previous skeleton points and residual points from multiple time steps. Formally, for an N-frame FSD++ detector, we have the final input points in time step t as follows:

$$P_t^{in} = \left(\bigcup_{i=1}^{N} P_{t-i}^s \right) \cup \left(\bigcup_{i=0}^{M-1} \Delta P_{t-i} \right), \qquad (5)$$

where P_t^s is the skeleton points at time step t. P_t^{in} is much more sparse than the raw point clouds and directly sent into the FSD detector. Fig. 7 shows examples of P^s, ΔP and P^{in}.

4.5 Training and Inference Pipeline

The training and inference pipeline for FSD++ differs from the standard approach due to its use of history predictions and temporal information. Fig. 6 summarizes the overall pipeline.

4.5.1 Training

To utilize history predictions, the input point cloud stream must be arranged in the temporal order. However, the ordered input stream affects the model training due to a lack of data shuffling. Another problem is that the history predictions are not reliable in the early training stages. Considering these two issues, we use a well-trained FSD detector to generate offline predictions of the entire training set. During training, for every sample, we load it along with its previous offline predictions to sample skeleton points. Ground truth boxes seem to be an alternative to the offline predictions. However, the distribution gap between ground truth used in training and predicted boxes used in inference is considerable. Thus we adopt the offline predictions instead of ground-truth boxes.

4.5.2 Inference

In the online inference phase, the input point cloud stream is naturally in the temporal order. For a point cloud sequence, the predictions of its first frame are from the well-trained FSD detector. These predictions are regarded as the "previous predictions" of the first frame, which are called the *seed predictions* of a sequence. We maintain several queues to cache some historical data that could be used more than once. For example, in a N-frame FSD++ pipeline, the raw points and skeleton points of time step t could be reused from time step $t+1$ to $t+N-1$.

5 EXPERIMENTS

5.1 Setup

5.1.1 Dataset

Waymo Open Dataset (WOD) In our experiments, we use WOD [1] as the primary dataset to evaluate the performance of our proposed method. WOD is the most trustworthy benchmark for LiDAR-based 3D object detection. With 1150 sequences and more than 200,000 frames, WOD is currently the largest dataset of its kind. Among them, 798 sequences are used for training, 202 for validation, and 150 for testing. The detection range in WOD is 75 meters (cover area of $150m \times 150m$).

Argoverse 2 (AV2) We further conduct long-range experiments on the recently released Argoverse 2 dataset [4] to demonstrate the superiority of FSD in long-range detection. AV2 has a similar scale to WOD, and it contains 1000 sequences in total, 700 for training, 150 for validation, and 150 for testing. In addition to *average precision* (AP), AV2 adopts a *composite score* as an evaluation metric, which takes both AP and localization errors into account. The perception range in AV2 is 200 meters (cover area of $400m \times 400m$), which is much larger than WOD. Such a large perception range leads to a huge memory footprint for dense detectors.

5.1.2 Model Configuration

To demonstrate the generality of SIR, we build two FSD variants. FSD_{sst} adopts the emerging single stride sparse transformer [6] as the sparse voxel feature extractor. FSD_{spconv} is built upon sparse convolution based U-Net in PartA2 [7]. In the experiments of FSD, we use FSD_{sst} in the experiments unless otherwise specified. In the experiments of FSD++, we use FSD_{spconv} as our detector since the highly optimized engineering of SpConv makes it more efficient than the SST backbone with multi-frame input.

5.1.3 Implementation Details

Our implementation is based on popular MMDetection3D(v0.15) [44]. In FSD_{sst}, we use 4 sparse regional attention blocks [6] as our voxel feature extractor. The SIR module and SIR2 module consist of 3 and 6 SIR layers, respectively. A SIR layer is defined by Eqn. 1 and Eqn. 2. Our SST-based model converges much faster than SST, so we train our models for 6 epochs for ablation study, instead of the $2\times$ schedule (24 epochs) in SST. For FSD_{spconv}, in addition to the 6-epoch schedule, we adopt a longer schedule (12 epochs) for better performance. Different from the default setting in MMDetection3D, we decrease the number of pasted instances in the CopyPaste augmentation. In FSD, some scarce classes like cyclist prone to be over-fitted with too many pasted instances. All experiments in Argoverse 2 dataset adopt a 12-epoch schedule. The models for performance analysis (Sec. 5.3 $\sim$ Sec. 5.6) are trained on 8 RTX 2080Ti GPUs with batch-size 2. And the models in Table 1 are trained on 8 RTX 3090 GPUs with batch-size 2. More details can be found in our released code.

5.2 Main Results of FSD and FSD++

We first compare FSD with state-of-the-art detectors and our baseline in Table 1 and Table 2. In the validation split, FSD/FSD++ achieves state-of-the-art average performance (L2 mAPH) in single-frame/multi-frame settings, respectively. In test split, FSD achieves the best performance on all classes among all single-frame detectors. Meanwhile, FSD++ with 7-frame input surpasses all detectors with up to 100-frame input, in terms of average metric.
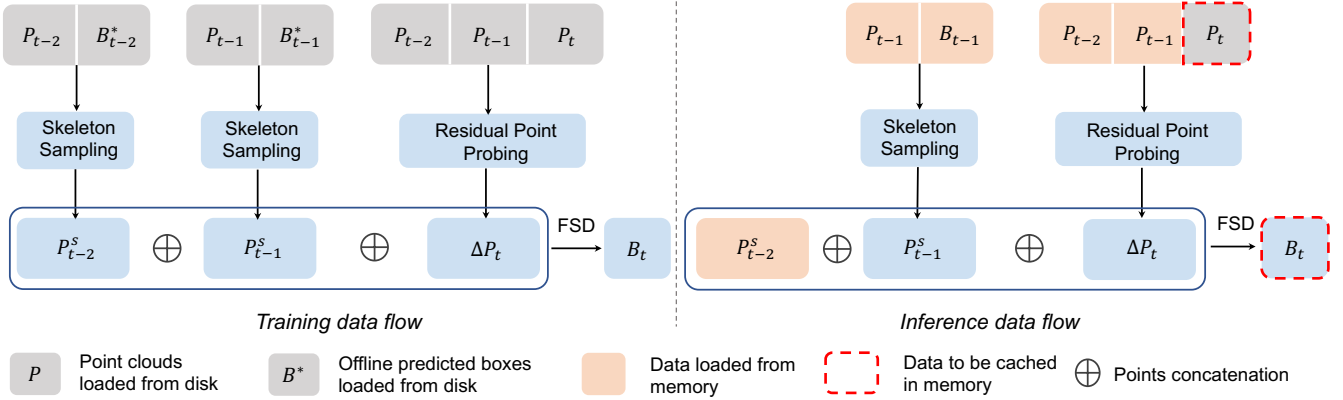
Fig. 6. The overall architecture of FSD++. In training, we adopt offline predictions to approximate history predictions. During inference, the detector uses previous online predictions for skeleton point sampling. When the max age is larger than one, there will be some other ΔP_{t-i} from time step $t-i$. For simplicity, we only present ΔP_t here.

(a) Moving cars with static ego-vehicle   (b) Moving cars with moving ego-vehicle   (c) Static cars and moving pedestrian
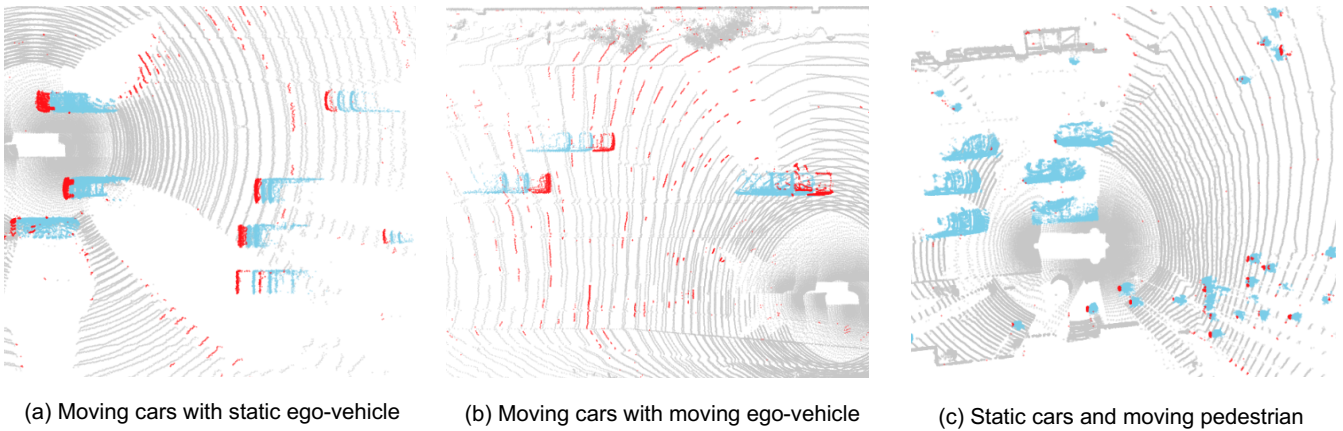
Fig. 7. Examples of super sparse input point clouds. Residual points are colored in red. Previous foreground points are in blue. Gray points will not be sent into the detector. (a) Moving cars cause apparent residual points. And some occluded points in the ground plane become visible due to car movement. (b) The ego-vehicle is moving, which causes some points in the ground plane to be detected as residual points. (c) Residual points are detected on the moving pedestrian instead of the static cars.

It is also noteworthy that FSD and FSD++ are much more efficient than most of the previous arts, especially in the multi-frame setting and long-range setting. We elaborate this in Sec. 5.4 and Sec. 5.7.

5.3 Study of Treatments to Center Feature Missing

5.3.1 Quantitative Experiments

In what follows, we conduct experiments on WOD to investigate the issue of **Center Feature Missing (CFM)**. We first develop several models with different characteristics. Note that all the following models adopt the same voxelization resolution, so they face the same degree of CFM at the beginning.

- **FSD$_{plain}$**: After the sparse voxel encoder, FSD$_{plain}$ directly predicts the box from each voxel. The voxels inside ground-truth boxes are assigned as *positive*. Although FSD$_{plain}$ uses the most straightforward solution for CFM, it suffers from the large variance of regression targets and low-quality predictions from informative voxels.
- **SST$_{center}$**: It replaces the anchor-based head in SST with CenterHead [9], [28]. Based on the sparse voxel encoder,

SST$_{center}$ converts sparse voxels into dense feature maps and applies several convolutions to diffuse features to the empty object centers as in Fig. 2. Then it makes predictions from the diffused center feature.
- **FSD$_{nogc}$**: It removes the group correction and SIR2 module in FSD.
- **CenterPoint-PP**: It does not resort to any sparse voxel encoders. Instead, it applies multiple dense convolutions soon after voxelization for feature diffusion, greatly eliminating CFM. It is also equipped with CenterHead to avoid large variance of regression targets.

There is usually a quite large unoccupied area around the centers of large vehicles. Thus the performance of large vehicles is an appropriate indicator that reveals the effect of CFM. So we build a customized evaluation tool, which breaks down the object length following the COCO evaluation [59]. Then we use it to evaluate the performance of vehicles with different lengths. Table 3 shows the results, and we list our findings as follows.

- Comparing FSD$_{plain}$ with SST$_{center}$, they share the same attention-based sparse voxel encoder. However, the trend is totally opposite w.r.t vehicle size. With feature diffusion,

Methods	#. frames	mAP/mAPH L2	*Vehicle* 3D AP/APH L1	L2	*Pedestrian* 3D AP/APH L1	L2	*Cyclist* 3D AP/APH L1	L2
SECOND [5]	1	61.0/57.2	72.3/71.7	63.9/63.3	68.7/58.2	60.7/51.3	60.6/59.3	58.3/57.0
MVF [14]	1	-/-	62.9/-	-/-	65.3/-	-/-	-/-	-/-
AFDet [45]	1	-/-	63.7/-	-/-	-/-	-/-	-/-	-/-
Pillar-OD [46]	1	-/-	69.8/-	-/-	72.5/-	-/-	-/-	-/-
RangeDet [39]	1	65.0/63.2	72.9/72.3	64.0/63.6	75.9/71.9	67.6/63.9	65.7/64.4	63.3/62.1
PointPillars [19]	1	62.8/57.8	72.1/71.5	63.6/63.1	70.6/56.7	62.8/50.3	64.4/62.3	61.9/59.9
Voxel RCNN [32]	1	-/-	75.6/-	66.6/-	-/-	-/-	-/-	-/-
RCD [42]	1	-/-	69.0/68.5	-/-	-/-	-/-	-/-	-/-
VoTr-TSD [47]	1	-/-	74.9/74.3	65.9/65.3	-/-	-/-	-/-	-/-
LiDAR-RCNN [43]	1	65.8/61.3	76.0/75.5	68.3/67.9	71.2/58.7	63.1/51.7	68.6/66.9	66.1/64.4
Pyramid RCNN [48]	1	-/-	76.3/75.7	67.2/66.7	-/-	-/-	-/-	-/-
Voxel-to-Point [49]	1	-/-	77.2/-	69.8/-	-/-	-/-	-/-	-/-
3D-MAN [50]	16	-/-	74.5/74.0	67.6/67.1	71.7/67.7	62.6/59.0	-/-	-/-
M3DETR [51]	1	61.8/58.7	75.7/75.1	66.6/66.0	65.0/56.4	56.0/48.4	65.4/64.2	62.7/61.5
Part-A2-Net [7]	1	66.9/63.8	77.1/76.5	68.5/68.0	75.2/66.9	66.2/58.6	68.6/67.4	66.1/64.9
CenterPoint-Pillar [9]	1	-/-	76.1/75.5	68.0/67.5	76.1/65.1	68.1/57.9	-/-	-/-
CenterPoint-Voxel [9]	1	69.8/67.6	76.6/76.0	68.9/68.4	79.0/73.4	71.0/65.8	72.1/71.0	69.5/68.5
IA-SSD [13]	1	62.3/58.1	70.5/69.7	61.6/61.0	69.4/58.5	60.3/50.7	67.7/65.3	65.0/62.7
PV-RCNN [8]	1	66.8/63.3	77.5/76.9	69.0/68.4	75.0/65.6	66.0/57.6	67.8/66.4	65.4/64.0
RSN [52]	1	-/-	75.1/74.6	66.0/65.5	77.8/72.7	68.3/63.7	-/-	-/-
SST_TS [6]	1	-/-	76.2/75.8	68.0/67.6	81.4/74.0	72.8/65.9	-/-	-/-
SST [6]	1	67.8/64.6	74.2/73.8	65.5/65.1	78.7/69.6	70.0/61.7	70.7/69.6	68.0/66.9
AFDetV2 [24]	1	71.0/68.8	77.6/77.1	69.7/69.2	80.2/74.6	72.2/67.0	73.7/72.7	71.0/70.1
PillarNet-34 [53]	1	71.0/68.5	79.1/78.6	70.9 / 70.5	80.6/74.0	72.3/66.2	72.3/71.2	69.7/68.7
PV-RCNN++ [10]	1	68.4/64.9	78.8/78.2	70.3/69.7	76.7/67.2	68.5/59.7	69.0/67.6	66.5/65.2
PV-RCNN++(center) [10]	1	71.7/69.5	79.3 / 78.8	70.6/70.2	81.3/76.3	73.2/68.0	73.7/72.7	71.2/70.2
CenterFormer [54]	8	75.1/73.7	78.8/78.3	74.3/73.8	82.1/79.3	77.8/75.0	75.2/74.4	73.2/72.3
INT [55]	10	-/73.6	-/-	-/73.3	-/-	-/71.9	-/-	-/75.6
MPPNet [56]	16	75.6/74.9	82.7 / 82.3	75.4 / 75.0	84.7/82.3	77.4/75.1	77.3/76.7	75.1/74.5
FSD$_{spconv}$ (ours)	1	71.9/69.7	77.8/77.3	68.9/68.5	81.9/76.4	73.2/68.0	76.5/75.2	73.8/72.5
FSD$_{sst}$ (ours)	1	71.5/69.2	76.8/76.3	67.9/67.5	81.3/75.3	72.5/67.0	77.2/76.0	74.4/73.2
FSD$_{spconv}$ (ours)†	1	72.9 / 70.8	79.2/78.8	70.5/70.1	82.6 / 77.3	73.9 / 69.1	77.1 / 76.0	74.4 / 73.3
FSD++ (ours)†	7	76.8 / 75.5	81.4/80.9	73.3/72.9	85.1 / 82.2	78.2 / 75.4	81.2 / 80.3	78.9 / 78.1

TABLE 1
Performances on the Waymo Open Dataset validation split. All reported results are from single model without any test-time augmentations. †: Longer schedule (12 epochs). We mark the best single-frame results and multi-frame results with gray boxes and cyan boxes, respectively.

SST$_{center}$ attains much worse performance than FSD$_{plain}$ on large vehicles. It suggests feature diffusion is a sub-optimal solution for CFM in the case of large objects. For those large objects, the features may not be diffused to the centers or the diffused features are too weak to make accurate predictions.

- However, FSD$_{plain}$ obtains the worst performance among all detectors on vehicles with normal sizes. Note that the CFM issue is minor for the normal-size vehicles. So, in this case, the center-based assignment in SST$_{center}$ shows its superiority to the assignment in FSD$_{plain}$. It suggests the solution for CFM in FSD$_{plain}$ is also sub-optimal, even if it achieves better performance in large objects.

- Comparing FSD$_{nogc}$ with SST$_{center}$, they share the same sparse voxel encoder while FSD$_{nogc}$ replaces the dense part in SST$_{center}$ with SIR. The huge improvements of FSD$_{nogc}$ on large vehicles fairly reveal that SIR effectively resolves CFM and is better than feature diffusion.

- CenterPoint-PP suffers much less from CFM because it leverages dense feature maps from the very beginning of the network. It is also equipped with advanced center-based assignment. Even so, FSD$_{nogc}$ and FSD still outperform CenterPoint-PP, especially on large vehicles.

5.3.2 Qualitative Analysis

In addition to the quantitative experiments, we demonstrate the qualitative effect of CFM and our treatment, shown in Fig. 8. Fig. 8 showcases the voted centers of FSD and the
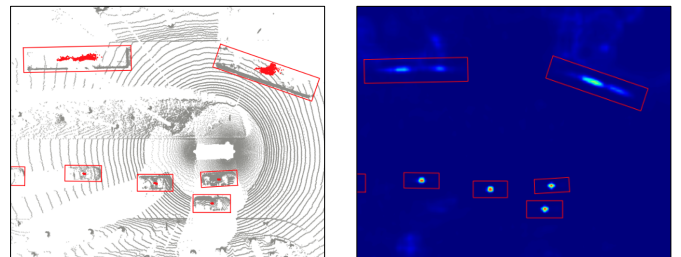
Fig. 8. An intuitive illustration of the center feature missing. **Left**: Voted centers of FSD. **Right:** Predicted heatmap of SST$_{center}$.

predicted heatmap of center-based SST. Both of them yield high-quality predictions for vehicles of normal size, but their predictions (votes) are usually ambiguous for large vehicles. Center-based dense detectors make predictions from such ambiguous heatmaps, so they are prone to make flawed final predictions. Although the center voting of FSD on large vehicles is also mediocre, FSD only uses the votes to obtain point groups (i.e., instance segmentation), which does not

Methods	#. frames	mAP/mAPH L2	Vehicle 3D AP/APH L1	L2	Pedestrian 3D AP/APH L1	L2	Cyclist 3D AP/APH L1	L2
CenterPoint [9]	1	-/69.0	-/-	-/71.9	-/-	-/67.0	-/-	-/68.2
AFDetV2-lite [24]	1	72.2/70.0	80.5/80.0	73.0/72.6	79.8/74.3	73.7/68.6	72.4/71.2	69.8/69.7
PV-RCNN [8]	1	71.3/68.8	80.6/80.1	72.8/72.4	78.2/72.0	71.8/66.0	71.8/70.4	69.1/67.8
PV-RCNN++ [10]	1	72.4/70.2	81.6/81.2	73.9/73.5	80.4/75.0	74.1/69.0	71.9/70.8	69.3/68.2
Graph R-CNN [57]	1	73.8/71.6	83.6 / 83.1	76.0 / 75.6	81.9/76.5	75.6/70.5	72.5/71.3	69.8/68.7
AFDetV2 [24]	2	74.6/73.1	81.7/81.2	74.3/73.9	81.3/78.0	75.5/72.4	76.4/75.4	74.1/73.0
CenterPoint++ [9]	3	74.2/72.8	82.8/82.3	75.5/75.1	81.0/78.2	75.1/72.4	74.4/73.3	72.0/71.0
BEVFusion* [58]	3	77.7/76.3	85.0 / 84.6	77.9/77.5	84.7 / 82.0	79.1/76.4	78.5/77.5	76.0/75.1
DeepFusion* [25]	5	76.9/75.5	83.3/82.8	76.1/75.6	84.6/81.8	79.2/76.4	77.8/76.8	75.5/74.5
CenterFormer [54]	16	76.9/75.6	84.7/84.4	78.1 / 77.7	84.6/81.8	79.4 / 76.6	75.5/74.5	73.3/72.4
INT [55]	100	76.6/75.2	84.7/84.3	78.0/77.6	82.4/79.7	76.6/74.0	77.4/76.3	75.2/74.1
MPPNet [56]	16	76.9/75.7	84.3/83.9	77.3/76.9	84.1/81.5	78.4/75.9	77.1/76.4	74.9/74.2
FSD$_{spconv}$ (ours)†	1	74.4 / 72.4	82.7/82.3	74.4/74.1	82.9 / 77.9	75.9 / 71.3	75.6 / 74.4	72.9 / 71.8
FSD++ (ours)†	7	78.4 / 77.1	84.5/84.1	77.1/76.7	84.5/81.7	79.0/76.2	81.4 / 80.5	79.2 / 78.3

TABLE 2
Performances on the Waymo Open Dataset test split. All results are in single-model setting without ensemble or test-time augmentations. *: Multi-modal methods with camera information. We mark the best single-frame results and multi-frame results with gray boxes and cyan boxes, respectively. †: 12-epoch schedule.

Methods	[0, 4)	[4, 8)	[8, 12)	[12, +∞)	Official*
		Vehicle length (m)			
CenterPoint-PP†	34.3	69.3	42.0	43.6	66.2
FSD$_{plain}$	32.2	64.6	41.3	42.2	62.3
SST$_{center}$ [6]	36.0	69.4	33.7	30.5	66.3
FSD$_{nogc}$	33.5 ↓2.5	68.2 ↓1.2	47.7 ↑14.0	47.9 ↑17.4	65.2 ↓1.1
FSD	36.7 ↑0.7	71.0 ↑1.6	51.3 ↑17.6	53.7 ↑23.2	69.3 ↑3.0

TABLE 3
Vehicle detection with vehicle length breakdown. †: re-implemented ourselves. *: official Waymo L2 overall metric. Arrows indicate the performance changes from SST$_{center}$.

necessitate perfect center voting. The final predictions of FSD are derived from the complete point groups rather than the weak center features, thereby sidestepping the issue of CFM.
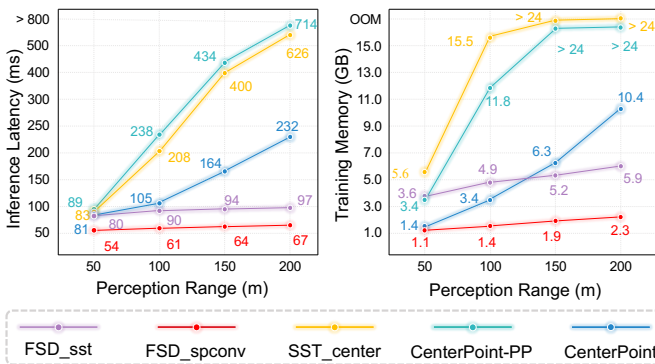

Fig. 9. Memory footprints and inference latency in different perception ranges. Statistics are obtained on a single 3090 GPU with batch size 1. Inference latency is evaluated by the standard benchmark script in MMDetection3D without any test-time optimization. CenterPoint-PP and SST$_{center}$ are defined in Sec. 5.3. Best viewed in color.

5.4 Long-range Detection

Several widely adopted 3D detection benchmarks [1], [2], [3] have relatively short perception range. To unleash the potential of FSD, we conduct long-range detection experiments

on the recently released Argoverse 2 dataset (AV2). AV2 has a perception range up to 200 meters, making it an ideal testbed for our method. In addition, AV2 contains objects in 30 classes, exhibiting the long-tail distribution, which is also another challenge for FSD.

5.4.1 Main results

We first list the main results of FSD on AV2 in Table 4. The authors of AV2 provide a baseline CenterPoint model, but the results are mediocre. To make a fair comparison, we re-implement a stronger CenterPoint model on the AV2 dataset. The re-implemented CenterPoint adopts the same training scheme with FSD, including ground-truth sampling to alleviate the long-tail issue. FSD outperforms CenterPoint in the average metric. It is noteworthy that FSD significantly outperforms CenterPoint in some tiny objects (e.g., Pedestrian, Construction Cone) as well as some objects with extremely large sizes (e.g., Articulated Bus, School Bus). We owe this to the virtue of instance-level fine-grained feature extraction in SIR.

5.4.2 Range Scaling

To demonstrate the efficiency of FSD in long-range detection, we depict the trend of training memory and inference latency of three detectors when the perception range increases in Fig. 9. Fig. 9 shows that dense detectors experience a dramatic increase in latency and memory footprint as the perception range grows. Designed to be fully sparse, the resource needed for FSD is roughly linear to the number of input points, so its memory and latency only slightly increase as the perception range extends.

5.5 Performance Inspection of FSD

5.5.1 Effectiveness of Components

In addition to FSD$_{plain}$ and FSD$_{nogc}$ (Sec. 5.3), we also degrade FSD to FSD$_{agg}$ to gain insights into its mechanism. In FSD$_{agg}$, we aggregate grouped point features by dynamic pooling after Instance Point Grouping, and then directly

Methods	Average	Vehicle	Bus	Pedestrian	Stop Sign	Box Truck	Bollard	C-Barrel	Motorcyclist	MPC-Sign	Motorcycle	Bicycle	A-Bus	School Bus	Truck Cab	C-Cone	V-Trailer	Sign	Large Vehicle	Stroller	Bicyclist
Precision																					
CenterPoint† [9]	13.5	61.0	36.0	33.0	28.0	26.0	25.0	22.5	16.0	16.0	12.5	9.5	8.5	7.5	8.0	8.0	7.0	6.5	3.0	2.0	14
CenterPoint*	22.0	67.6	38.9	46.5	16.9	37.4	40.1	32.2	28.6	27.4	33.4	24.5	8.7	25.8	22.6	29.5	22.4	6.3	3.9	0.5	20.1
FSD	24.0	67.1	39.8	57.4	21.3	38.3	38.3	38.1	30.0	23.6	38.1	25.5	15.6	30.0	20.1	38.9	23.9	7.9	5.1	5.7	27.0
FSD‡	28.2	68.1	40.9	59.0	29.0	38.5	41.8	42.6	39.7	26.2	49.0	38.6	20.4	30.5	14.8	41.2	26.9	11.9	5.9	13.8	33.4
Composite Score																					
CenterPoint*	17.6	57.2	32.0	35.7	13.2	31.0	28.9	25.6	22.2	19.1	28.2	19.6	6.8	22.5	17.4	22.4	17.2	4.8	3.0	0.4	16.7
FSD	19.1	56.0	33.0	45.7	16.7	31.6	27.7	30.4	23.8	16.4	31.9	20.5	12.0	25.6	15.9	29.2	18.1	6.4	3.8	4.5	22.1
FSD‡	22.7	57.7	34.2	47.5	23.4	31.7	30.9	34.4	32.3	18.0	41.4	32.0	15.9	26.1	11.0	30.7	20.5	9.5	4.4	11.5	28.0

TABLE 4

Performance in Argoverse 2 validation split. †: provided by authors of AV2 dataset. ‡: Weak CopyPaste augmentation for preventing overfitting (one instance per class). ∗: re-implemented by ourselves. C-Barrel: construction barrel. MPC-Sign: mobile pedestrian crossing sign. A-Bus: articulated bus. C-Cone: construction cone. V-Trailer: vehicular trailer. We omit the results of dog, wheelchair and message board trailer because these categories contain very few instances. The average results take all categories into account, including the omitted categories. We mark the categories attaining notable improvements in **bold**.

make predictions from the pooled features. FSD_{agg} is similar to the way in VoteNet [22] as we discussed in Sec. 3.5. Thus, FSD_{agg} can explicitly leverage instance-level features other than the point-level features in FSD_{plain}, mitigating the issue of CFM. However, FSD_{agg} cannot take advantage of further point feature extraction in SIR. As can be seen in Table 5, the improvement is limited if we only apply grouping without SIR. The combination of grouping and SIR, on the other hand, yields significant improvements.

	Grouping	SIR	Group Correction	Vehicle	L2 3D APH Pedestrian	Cyclist
FSD_{plain}				62.29	64.31	64.49
FSD_{agg}	✓			63.13	65.13	64.52
FSD_{nogc}	✓	✓		65.20	67.39	67.78
FSD	✓	✓	✓	69.30	69.30	69.60

TABLE 5

Ablation of design factors in SIR. Performances are evaluated on Waymo validation split.

5.5.2 Downsampling in SIR

The efficiency of SIR makes it feasible to extract fine-grained point features without any point downsampling. This is another notable difference between FSD and VoteNet. To demonstrate the superiority, we apply voxelization on the raw points before the SIR module and treat the centroids of voxels as downsampled points. We conduct experiments on the AV2 dataset because it contains a couple of categories in a tiny size, which may be sensitive to downsampling. As expected, small objects have notable performance loss when adopting downsampling, and we list some of them in Table 6. We also evaluate the inference latency of the SIR module on a 3090 GPU, which is highly efficient.

5.5.3 HD Map-assisted Detection

Argoverse 2 dataset provides a highly reliable HD map, which could be utilized as a prior to remove uninterested regions making the scene more sparse. Thus we proceed with experiments removing some uninterested regions to show the advantages of FSD in more sparse scenarios. The results are summarized in Table 7. FSD has a significantly lower memory footprint and latency with an acceptable precision

Voxel size	CC	Bollard	AP Bicyclist	Stop Sign	Latency (ms)†
30cm	35.4	36.5	24.6	18.3	3.5
20cm	37.3	37.3	26.4	20.0	4.1
10cm	38.9	38.3	27.0	21.3	4.5
Point	39.3	38.6	27.1	21.5	6.3

TABLE 6

Performances with different representation granularity. †: Latency of SIR module.

	FSD Mem.	FSD Latency(ms)	mAP	CenterPoint Mem.	CenterPoint Latency(ms)	mAP
all	5.9	97	24.0	10.4	232	22.0
only RoI†	3.2 ↓45.8%	81↓16.5%	23.2	9.9↓4.8%	227↓2.2%	21.5
w/o ground	2.3 ↓61.0%	74↓25.8%	21.0	9.7↓6.7%	217↓6.4%	19.8

TABLE 7

Performance with different detection areas. †: Region of Interest is defined by the HD map in AV2 dataset.

loss after removing the uninterested regions. On the contrary, the efficiency improvement of CenterPoint is minor. It reveals that FSD benefits more from the increase of data sparsity, which is another advantage of the fully sparse architecture.

5.6 Comprehensive Analysis of FSD++

5.6.1 Preliminary Settings for the Analysis of FSD++

In this section, we conduct extensive experiments to reveal the inner workings of FSD++. Here we first present the setting of our *baseline* model in this section, which is slightly different from the best FSD++ model in Table. 1 and Table. 2. Unless otherwise specified, the default hyper-parameters of all the FSD++ models in Sec. 5.6 are listed in the first column of Table 8.

The model latency reported in this section is measured on a single RTX 3090 GPU with a mini-batch size of 1 in float32 precision. To ensure accuracy, we only consider the latency of the model in all evaluations, excluding the latency of IO, which is potentially unstable in the multi-frame setting.

It is worth noting that we have observed run-to-run variation in the performance of cyclist class, likely due to its low number in the dataset. As a result, we mark the

	Baseline FSD++	Best FSD++	Multi-frame FSD
Schedule	6 epochs	12 epochs	6 epochs
SPS†	Random	Random	-
#. frames	6	7	6
Max age	2	2	-
Backbone∗	SpUNet-base	SpUNet-large	SpUNet-base
#. layers in SIR2	3	3	3
RPP size	(0.25, 0.25, 0.4)	(0.25, 0.25, 0.4)	-

TABLE 8
Basic hyper-parameter choice of models adopted in this section
(Sec. 5.6) †: SPS stands for skeleton point sampling. ∗: SpUNet-large
has one more stage than SpUNet-base [7] and the number of channels
of its first is doubled.

performance of this class in gray in some experiments to indicate that it may not be reliable.

5.6.2 Skeleton Point Sampling

Table 9 shows the performance with different skeleton point sampling strategies. We find that there are no significant differences between the three strategies considered. However, using random skeleton sampling considerably reduces the latency of FSD++ without sacrificing performance. And skeleton sampling consistently boosts the performance of the cyclist class, which suggests that appropriate sampling alleviates the overfitting for rare classes. The performance of pedestrian is better without sampling, which reveals that more points might be helpful for small objects. In practice, although different sampling strategies could be adopted for different classes, we use the random sampling for all classes for simplicity and generality.

		L2 3D AP/APH		Latency
	Mean	Vehicle	Pedestrian	Cyclist	(ms)
Random	76.10/74.73	72.20/71.74	76.93/74.11	79.20/78.33	68.7
Object FPS	75.56/74.21	72.06/71.59	76.94/74.14	77.68/76.88	72.3
Voxel Sampling	75.76/74.40	72.10/71.66	76.80/74.05	78.37/77.49	71.4
None	75.23/73.91	71.93/71.50	77.54/74.72	76.33/75.51	73.9

TABLE 9
Effectiveness of different skeleton point sampling.

5.6.3 Different Number of Frames

FSD++ samples skeleton points from multiple previous frames. Table 10 showcases how the number of used frames affects its performance. There are two interesting findings.

- Performance becomes better as the number of frames grows. In the meantime, the latency does not significantly increase. We owe the credit to residual point probing, which removes most of the background. It could offer even more clean residual points if more base frames (Eqn. 4) are used.
- FSD++ outperforms FSD with the same number of frames in vehicle and cyclist class. We also intuitively owe it to RPP since it removes most background clutter and eases the burden of the segmentation. The slightly lower performance of pedestrian suggests it might be better to retain all points for pedestrian. However, the performance loss is acceptable since FSD++ achieves better average performance and much lower latency.

#. frames		L2 3D AP/APH			Latency
	Mean	Vehicle	Pedestrian	Cyclist	(ms)
2	73.39/71.83	69.54/69.12	74.68/71.35	75.94/75.02	66.1
3	75.20/73.74	70.95/70.52	76.13/73.09	78.51/77.62	67.0
4	75.44/74.03	71.67/71.21	76.48/73.50	78.18/77.37	67.3
5	75.13/73.72	71.50/71.04	76.71/73.83	77.17/76.29	68.7
6	76.10/74.73	72.19/71.74	76.92/74.11	79.20/78.33	68.7
6 (FSD)†	75.65/74.28	71.54/71.07	78.04/75.22	77.37/76.54	116.2

TABLE 10
Performance of FSD++ with the different number of frames. Since
FSD++ uses previous foreground points, it needs at least two frames. †:
multi-frame FSD model with simple point concatenation. Performance is
unstable in the scarce cyclist class, so we mark the numbers in gray.

5.6.4 Drifting Analysis

It would be a major concern if FSD++ suffers from the drifting error given its reliance on history predictions. In particular, if the detector makes inaccurate predictions at time step t, it is likely that the detector becomes worse at time step $t + 1$ since the predictions in $t + 1$ rely on the predictions from t (skeleton point sampling).

To prevent potential drifting, we insert some keyframes at regular intervals during the inference of a sequence. At keyframes, we use the predictions from standard FSD for skeleton point sampling, which could be viewed as a rectification of the potential drifting. Table. 11 shows that FSD++ achieves competitive results without any keyframes. And the minor gap between the first row and the last row confirms that the drifting of FSD++ is negligible.

Gap between key frames		L2 3D AP/APH		
	Mean	Vehicle	Pedestrian	Cyclist
5	76.05/74.67	72.23/71.77	76.83/74.02	79.08/78.21
10	76.03/74.66	72.20/71.75	76.88/74.07	79.00/78.15
20	76.07/74.68	72.20/71.74	76.90/74.08	79.11/78.23
50	76.10/74.72	72.20/71.74	76.91/74.10	79.20/78.32
None	76.10/74.73	72.20/71.74	76.93/74.11	79.20/78.33

TABLE 11
The performance of different keyframe gaps. "None" means using only
initial predictions.

5.6.5 Change Blindness Ablation

Due to the *change blindness* we discussed in Sec. 4.3, newly emerged objects might be ignored by the detector. *Max age* is proposed to mitigate the issue of change blindness, and Table 12 shows its effect. We find keep residual points for two time steps (max age 2) is enough.

Max age		L2 3D AP/APH			Latency
	Mean	Vehicle	Pedestrian	Cyclist	(ms)
1	75.17/73.82	71.38/70.94	76.74/73.97	77.40/76.56	65.3
2	76.10/74.73	72.20/71.74	76.93/74.11	79.20/78.33	68.7
3	**76.14/74.74**	**72.22/71.75**	**77.06/74.20**	79.13/78.27	70.2

TABLE 12
Different max ages of residual points. Performance is unstable in the
scarce cyclist class, so we mark the numbers in gray.

For a closer look at the issue of change blindness, we split the objects in the original WOD validation set to *emerging objects* and *existing objects* for further ablations. Emerging objects mean those objects do not appear in the first frame of a sequence, while emerging later. The results are shown in

Table 13. The emerging objects recall of FSD++(1)[1] is inferior to FSD_6f with the same number of frames. This suggests that change blindness is indeed an issue for FSD++. However, prolonging the max age makes FSD++ outperform FSD_6f in all classes, which demonstrates the proposed max age effectively mitigates change blindness.

Method	Recall of emerging objects			
	Mean	Vehicle	Pedestrian	Cyclist
FSD	73.05	66.01	72.62	80.53
FSD_6f*	77.54	69.74	78.34	84.53
FSD++(1)	75.58	69.18	77.48	83.09
FSD++(2)	77.82	70.56	78.10	84.79
FSD++(3)	78.14	70.60	78.30	85.52

TABLE 13
Performance for emerging objects. ∗: FSD with 6-frame concatenated input. In the WOD validation split, the number of emerging objects count for around 42.4%/37.1%/52.6% in all objects for vehicle / pedestrian / cyclist, respectively.

5.6.6 Robustness to Seed Quality

During inference of every point cloud sequence, FSD++ needs the predictions in the initial frame as a seed to start, as we discussed in Sec. 4.5.2. Here we figure out how the quality of seed predictions affects the performance. Concretely, we add two typical kinds of random noise to seed predictions, including random box drop and random box insertion. They are designed to simulate false negatives and false positives. All the experiments share a trained FSD++ detector. The modifications above are applied during inference and are not adopted for training-time augmentation.

Table 14 shows FSD++ is robust to both two types of noises. Particularly, for box drop, there is only marginal performance degradation even after dropping all the initial seed boxes. We explain this surprising phenomenon in two aspects: (1) FSD++ is born to be robust to the dropping of moving objects. This is because moving objects create a considerable amount of residual points and FSD++ is capable of making predictions from these residual points. (2) There is still a small number of residual points in the static objects due to the change of viewpoint. Moreover, the mechanism of max age also helps accumulate residual points on static objects.

In the case of box insertion, FSD++ is almost unaffected because they can be easily identified as background in the segmentation stage of FSD.

Noise type	L2 3D AP/APH			
	Mean	Vehicle	Pedestrian	Cyclist
None	76.10/74.73	72.20/71.74	76.93/74.11	79.20/78.33
Drop (10%)†	75.95/74.58	72.11/71.66	76.81/73.99	78.92/78.08
Drop (50%)	75.76/74.39	71.86/71.41	76.63/73.82	78.78/77.95
Drop (100%)	74.69/73.35	70.47/70.02	75.41/72.69	78.20/77.33
Insertion (10%)	76.00/74.62	72.16/71.70	76.82/73.99	79.01/78.16
Insertion (50%)	76.02/74.64	72.14/71.69	76.86/74.04	79.05/78.19
Insertion (100%)	75.98/74.61	72.16/71.71	76.84/74.02	78.95/78.10

TABLE 14
Robustness to the noisy seed predictions. †: the percentage in parentheses denotes the ratio of dropped/inserted instances.

1. The numbers in the parenthesis denote the max ages.

5.6.7 Analysis of Residual Point Probing

Quantization size and the number of base frames are two important hyper-parameters in RPP. Here we show how they affect the output residual points and performance.

Quantization size makes RPP robust to small point disturbance. We list the results of different quantization sizes in Table 15. It could be seen from the "residual point ratio" that RPP with larger quantization sizes leads to less residual points making the detector more efficient but leading to slightly lower performance.

Quantization size	L2 3D AP/APH			Residual point ratio†	Latency (ms)
	Vehicle	Pedestrian	Cyclist		
(0.15, 0.15, 0.4)	72.06/71.60	77.19/74.43	77.53/76.67	17.4%	72.3
(0.25, 0.25, 0.4)	72.20/71.74	76.93/74.11	79.20/78.33	9.6%	68.7
(0.35, 0.35, 0.4)	72.04/71.58	76.88/74.03	78.54/77.68	6.0%	65.2

TABLE 15
The effectiveness of quantization size in Residual Point Probing. †: residual point ratio means the average ratio of the residual points to the total points in a single frame.

Base frame (in Eqn. 4) also has a considerable effect on RPP. The more base frames are incorporated, the less residual points could be obtained, leading to higher efficiency. Moreover, the performance is hardly affected by the increase of base frames.

#. RPP base frames	L2 3D AP/APH			Residual point ratio†	Latency (ms)
	Vehicle	Pedestrian	Cyclist		
3	72.48/72.03	77.16/74.35	77.60/76.78	14.4%	70.2
4	72.24/71.79	77.04/74.20	78.44/77.60	10.8%	69.0
5	72.20/71.74	76.93/74.11	79.20/78.33	9.6%	68.7
6	72.22/71.77	77.41/74.59	78.60/77.74	9.4%	68.5

TABLE 16
The effectiveness of the number of base frames in Residual Point Probing. †: residual point ratio means the average ratio of the residual points to the total points in a single frame.

5.7 Detailed Runtime Evaluation

Here we elaborate on the efficiency of each component of FSD and FSD++. All evaluated models use SpUNet-large as the backbone. Evaluations are conducted on a single RTX 3090 in FP32 precision without any test-time optimizations. We only record the *single-sample* forward latency of the detector implemented with MMDetection3Dv0.15, ignoring the IO of point clouds which is unstable in the multi-frame setting. Fig. 10 shows the detailed results, which are average numbers evaluated on the first ten sequences of validation split.

As can be seen from the figure, the latency of the segmentor is greatly reduced, which consists of the sparse voxel encode (i.e., backbone) and segmentation head. As a results, FSD++ is as fast as the single-frame FSD, yet achieves better performance than FSD_6f (Table 10). It is worth emphasizing that the "others" part of latency is usually brought by some serialized operations, such as class-wise detection heads and class-wise NMS. This part of latency could be greatly reduced in deployment.

6 CONCLUSION

This paper first proposes a LiDAR-based fully sparse 3D object detection framework, namely FSD. FSD forsakes the
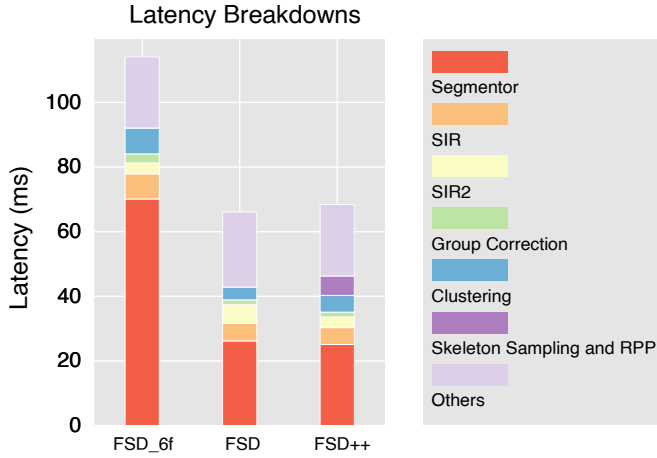
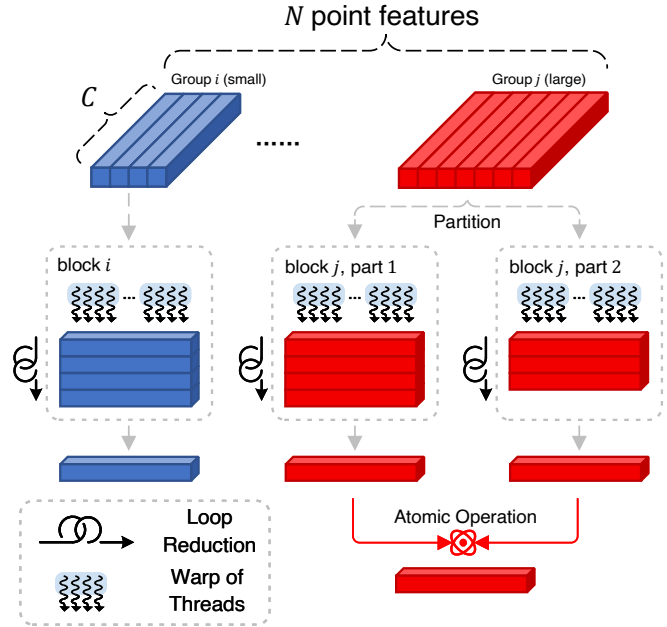Fig. 10. Latency breakdowns of multi-frame FSD, single-frame FSD and FSD++.

Fig. 11. Illustration of dynamic pooling implementation in CUDA. Best viewed in color. We take two groups with different sizes as examples. Dynamic pooling reduces each group to a single feature.

widely adopted dense BEV feature map in previous arts, which is the hindrance to making detectors fully sparse. Instead, FSD consists of a general sparse voxel encoder and a highly-efficient sparse instance recognition (SIR) module. SIR remedies the issue of the center feature missing, which is the essential difficulty of fully sparse architecture. FSD not only actualizes efficient long-range (up to 200 meters) detection on Argoverse 2 dataset, but also achieves state-of-the-art performance on the competitive Waymo Open Dataset.

To unleash the potential of FSD, we propose leveraging temporal information to remove data redundancy. The proposed Skeleton Point Sampling and Residual Point Probing offer FSD a super sparse input point cloud, which constitutes the FSD++ framework. FSD++ achieves state-of-the-art single-model performance on both validation and test split Waymo Open Dataset, and maintains high efficiency. We hope our work lightens future research direction for LiDAR-based point cloud recognition.

APPENDIX A
EFFICIENT DYNAMIC POOLING

The proposed SIR consists of three basic operations: MLP, dynamic broadcasting, and dynamic pooling. MLP is highly optimized in mainstream deep learning frameworks. Dynamic broadcasting is essentially an indexing operation, which is highly parallel in modern GPUs. The efficiency bottleneck of SIR lies in dynamic pooling. Thus we provide an efficient implementation of dynamic pooling in this section.

A.1　Implementation

The dynamic pooling implemented by PyTorch is known as scatter operation. In this implementation, each thread manages one feature and performs simple atomic operations for feature reduction. Intensive atomic operations in large groups are detrimental to parallelism.

We optimize the dynamic pooling operator in three aspects. (1) Partitioning large groups into small sub-groups with fixed sizes to balance the workload. (2) Sorting the

features with the same group ID in adjacent positions to enhance the memory locality. Note that the group IDs of each element remain unchanged throughout the SIR module, so the sort is only applied once. (3) Threads in a warp are assigned to adjacent channels for coalesced memory access without warp divergence.

In particular, each thread corresponds to a feature dimension and reduces features in a partitioned sub-group through loops. Besides, the calls to atomic operators are reduced from once per thread to once per block, which contributes to high parallelism. Fig. 11 illustrates our efficient implementation.

A.2　Runtime Evaluation

We take dynamic max-pooling as an example and evaluate the latency of our implementation and *torch_scatter* in different cases, including multiple feature dimensions, multiple group sizes, and whether the group size is balanced. The total number of groups in the evaluation is 100. The results are shown in Table 17 and Table 18. With different data sizes, our implementation achieves 2.48× to 39.58× speedup. And we have a more significant speedup with imbalanced group sizes.

Feature dimension	Latency (ms) with Different Group Sizes			
	$[10^0, 10^1)$	$[10^1, 10^2)$	$[10^2, 10^3)$	$[10^3, 10^4)$
64	0.06/0.16	0.06/0.16	0.08/1.49	0.24/14.05
256	0.06/0.18	0.06/0.33	0.14/3.52	0.72/20.53
1024	0.09/0.16	0.10/0.85	0.37/6.53	4.82/65.28
Speedup	2.48×	5.56×	20.47×	30.53×

TABLE 17
Latency of dynamic max pooling on data with **balanced** group sizes. Each item denotes the latency of **ours/torch_scatter** in milliseconds.

Feature dimension	Latency (ms) with Different Group Sizes			
	$[10^0, 10^1]^*$	$[10^1, 10^2]^*$	$[10^2, 10^3]^*$	$[10^3, 10^4]^*$
64	0.06/0.16	0.06/0.32	0.09/2.60	0.36/20.87
256	0.06/0.15	0.08/0.78	0.20/5.44	1.19/32.75
1024	0.06/0.24	0.12/1.36	0.57/9.24	5.91/196.4
Speedup	$3.05\times$	$8.81\times$	$24.01\times$	$39.58\times$

TABLE 18
Latency of dynamic max pooling on data with **imbalanced** group sizes. Each item denotes the latency of **ours/torch_scatter** in milliseconds. $*$: The sizes of one-tenth groups are enlarged by $10\times$ to create imbalanced data.

REFERENCES

[1] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine *et al.*, "Scalability in Perception for Autonomous Driving: Waymo Open Dataset," in *CVPR*, 2020.

[2] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 2012.

[3] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *CVPR*, 2020.

[4] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, D. Ramanan, P. Carr, and J. Hays, "Argoverse 2: Next Generation Datasets for Self-Driving Perception and Forecasting," in *NeurIPS Datasets and Benchmarks 2021*, 2021.

[5] Y. Yan, Y. Mao, and B. Li, "SECOND: Sparsely Embedded Convolutional Detection," *Sensors*, vol. 18, no. 10, 2018.

[6] L. Fan, Z. Pang, T. Zhang, Y.-X. Wang, H. Zhao, F. Wang, N. Wang, and Z. Zhang, "Embracing Single Stride 3D Object Detector with Sparse Transformer," in *CVPR*, 2022.

[7] S. Shi, Z. Wang, J. Shi, X. Wang, and H. Li, "From Points to Parts: 3D Object Detection from Point Cloud with Part-aware and Part-aggregation Network," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[8] S. Shi, C. Guo, L. Jiang, Z. Wang, J. Shi, X. Wang, and H. Li, "PV-RCNN: Point-Voxel Feature Set Abstraction for 3D Object Detection," in *CVPR*, 2020.

[9] T. Yin, X. Zhou, and P. Krähenbühl, "Center-based 3D Object Detection and Tracking," *arXiv preprint arXiv:2006.11275*, 2020.

[10] S. Shi, L. Jiang, J. Deng, Z. Wang, C. Guo, J. Shi, X. Wang, and H. Li, "PV-RCNN++: Point-Voxel Feature Set Abstraction With Local Vector Representation for 3D Object Detection," *arXiv preprint arXiv:2102.00463*, 2021.

[11] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space," in *NeurIPS*, 2017.

[12] Z. Yang, Y. Sun, S. Liu, and J. Jia, "3DSSD: Point-based 3D Single Stage Object Detector," in *CVPR*, 2020.

[13] Y. Zhang, Q. Hu, G. Xu, Y. Ma, J. Wan, and Y. Guo, "Not All Points Are Equal: Learning Highly Efficient Point-based Detectors for 3D LiDAR Point Clouds," in *CVPR*, 2022.

[14] Y. Zhou, P. Sun, Y. Zhang, D. Anguelov, J. Gao, T. Ouyang, J. Guo, J. Ngiam, and V. Vasudevan, "End-to-End Multi-View Fusion for 3D Object Detection in LiDAR Point Clouds," in *CoRL*, 2020.

[15] B. Li, "3D Fully Convolutional Network for Vehicle Detection in Point Cloud," in *IROS*, 2017.

[16] Y. Zhou and O. Tuzel, "VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection," in *CVPR*, 2018.

[17] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, "Multi-View 3D Object Detection Network for Autonomous Driving," in *CVPR*, 2017.

[18] B. Yang, W. Luo, and R. Urtasun, "PIXOR: Real-time 3D Object Detection from Point Clouds," in *CVPR*, 2018.

[19] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "PointPillars: Fast Encoders for Object Detection from Point Clouds," in *CVPR*, 2019.

[20] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation," in *CVPR*, 2017.

[21] S. Shi, X. Wang, and H. Li, "PointRCNN: 3D Object Proposal Generation and Detection from Point Cloud," in *CVPR*, 2019.

[22] C. R. Qi, O. Litany, K. He, and L. J. Guibas, "Deep Hough Voting for 3D Object Detection in Point Clouds," in *ICCV*, 2019.

[23] J. Yang, L. Song, S. Liu, Z. Li, X. Li, H. Sun, J. Sun, and N. Zheng, "DBQ-SSD: Dynamic ball query for efficient 3d object detection," *arXiv preprint arXiv:2207.10909*, 2022.

[24] Y. Hu, Z. Ding, R. Ge, W. Shao, L. Huang, K. Li, and Q. Liu, "AFDetV2: Rethinking the Necessity of the Second Stage for Object Detection from Point Clouds," *arXiv preprint arXiv:2112.09205*, 2021.

[25] Y. Li, A. W. Yu, T. Meng, B. Caine, J. Ngiam, D. Peng, J. Shen, B. Wu, Y. Lu, D. Zhou *et al.*, "DeepFusion: Lidar-Camera Deep Fusion for Multi-Modal 3D Object Detection," in *CVPR*, 2022.

[26] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single Shot Multibox Detector," in *ECCV*, 2016.

[27] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-time Object Detection with Region Proposal Networks," *NeurIPS*, vol. 28, 2015.

[28] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as Points," *arXiv preprint arXiv:1904.07850*, 2019.

[29] W. Zheng, W. Tang, L. Jiang, and C.-W. Fu, "SE-SSD: Self-ensembling Single-stage Object Detector from Point Cloud," in *CVPR*, 2021.

[30] W. Zheng, W. Tang, S. Chen, L. Jiang, and C.-W. Fu, "Cia-ssd: Confident Iou-aware Single-stage Object Detector from Point Cloud," in *AAAI*, 2021.

[31] C. He, H. Zeng, J. Huang, X.-S. Hua, and L. Zhang, "Structure Aware Single-stage 3D Object Detection from Point Cloud," in *CVPR*, 2020.

[32] J. Deng, S. Shi, P. Li, W. Zhou, Y. Zhang, and H. Li, "Voxel R-CNN: Towards High Performance Voxel-based 3D Object Detection," in *AAAI*, 2021.

[33] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal Loss for Dense Object Detection," in *ICCV*, 2017.

[34] L. Jiang, H. Zhao, S. Shi, S. Liu, C.-W. Fu, and J. Jia, "Pointgroup: Dual-set point grouping for 3d instance segmentation," in *CVPR*, 2020.

[35] W. Wang, R. Yu, Q. Huang, and U. Neumann, "SGPN: Similarity Group Proposal Network for 3d Point Cloud Instance Segmentation," in *CVPR*, 2018.

[36] F. Hong, H. Zhou, X. Zhu, H. Li, and Z. Liu, "Lidar-based panoptic segmentation via dynamic shifting network," in *CVPR*, 2021.

[37] R. E. Tarjan, "A class of algorithms which require nonlinear time to maintain disjoint sets," *Journal of computer and system sciences*, 1979.

[38] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *Acm Transactions On Graphics (tog)*, vol. 38, no. 5, pp. 1–12, 2019.

[39] L. Fan, X. Xiong, F. Wang, N. Wang, and Z. Zhang, "RangeDet: In Defense of Range View for LiDAR-Based 3D Object Detection," in *ICCV*, 2021.

[40] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention Is All You Need," in *NeurIPS*, 2017.

[41] D. Hendrycks and K. Gimpel, "Gaussian Error Linear Units (GELUs)," *arXiv preprint arXiv:1606.08415*, 2016.

[42] A. Bewley, P. Sun, T. Mensink, D. Anguelov, and C. Sminchisescu, "Range Conditioned Dilated Convolutions for Scale Invariant 3D Object Detection," *arXiv preprint arXiv:2005.09927*, 2020.

[43] Z. Li, F. Wang, and N. Wang, "LiDAR R-CNN: An Efficient and Universal 3D Object Detector," in *CVPR*, 2021.

[44] M. Contributors, "MMDetection3D: OpenMMLab Next-generation Platform for General 3D Object Detection," https://github.com/open-mmlab/mmdetection3d, 2020.

[45] R. Ge, Z. Ding, Y. Hu, Y. Wang, S. Chen, L. Huang, and Y. Li, "AFDet: Anchor Free One Stage 3D Object Detection," *arXiv preprint arXiv:2006.12671*, 2020.

[46] Y. Wang, A. Fathi, A. Kundu, D. Ross, C. Pantofaru, T. Funkhouser, and J. Solomon, "Pillar-based Object Detection for Autonomous Driving," in *ECCV*, 2020.

[47] J. Mao, Y. Xue, M. Niu, H. Bai, J. Feng, X. Liang, H. Xu, and C. Xu, "Voxel Transformer for 3D Object Detection," in *ICCV*, 2021.

[48] J. Mao, M. Niu, H. Bai, X. Liang, H. Xu, and C. Xu, "Pyramid R-CNN: Towards Better Performance and Adaptability for 3D Object Detection," in *ICCV*, 2021.

[49] J. Li, H. Dai, L. Shao, and Y. Ding, "From Voxel to Point: IoU-guided 3D Object Detection for Point Cloud with Voxel-to-Point Decoder," in *ACM-MM*, 2021.

[50] Z. Yang, Y. Zhou, Z. Chen, and J. Ngiam, "3D-MAN: 3D Multi-Frame Attention Network for Object Detection," in *CVPR*, 2021.

[51] T. Guan, J. Wang, S. Lan, R. Chandra, Z. Wu, L. Davis, and D. Manocha, "M3DETR: Multi-Representation, Multi-Scale, Mutual-Relation 3D Object Detection With Transformers," in *WACV*, 2022.

[52] P. Sun, W. Wang, Y. Chai, G. Elsayed, A. Bewley, X. Zhang, C. Sminchisescu, and D. Anguelov, "RSN: Range Sparse Net for Efficient, Accurate LiDAR 3D Object Detection," in *CVPR*, 2021.

[53] G. Shi, R. Li, and C. Ma, "PillarNet: High-Performance Pillar-based 3D Object Detection," *arXiv preprint arXiv:2205.07403*, 2022.

[54] Z. Zhou, X. Zhao, Y. Wang, P. Wang, and H. Foroosh, "Center-Former: Center-based Transformer for 3D Object Detection," in *ECCV*. Springer, 2022.

[55] J. Xu, Z. Miao, D. Zhang, H. Pan, K. Liu, P. Hao, J. Zhu, Z. Sun, H. Li, and X. Zhan, "INT: Towards Infinite-frames 3D Detection with An Efficient Framework," in *ECCV*. Springer, 2022.

[56] X. Chen, S. Shi, B. Zhu, K. C. Cheung, H. Xu, and H. Li, "MPP-Net: Multi-Frame Feature Intertwining with Proxy Points for 3D Temporal Object Detection," in *ECCV*. Springer, 2022.

[57] H. Yang, Z. Liu, X. Wu, W. Wang, W. Qian, X. He, and D. Cai, "Graph r-cnn: Towards accurate 3d object detection with semantic-decorated local graph," in *European Conference on Computer Vision*. Springer, 2022, pp. 662–679.

[58] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. Rus, and S. Han, "BEVFusion: Multi-Task Multi-Sensor Fusion with Unified Bird's-Eye View Representation," *arXiv preprint arXiv:2205.13542*, 2022.

[59] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common Objects in Context," in *European conference on computer vision*. Springer, 2014.