

STROBE: Streaming Object Detection from LiDAR Packets

Davi Frossard^{1,2} Simon Suo^{1,2} Sergio Casas^{1,2} James Tu^{1,2}

Rui Hu¹ Raquel Urtasun^{1,2}

¹ Uber Advanced Technologies Group ² University of Toronto
{frossard, suo, sergio.casas, james.tu, rui.hu, urtasun}@uber.com

Abstract: Many modern robotics systems employ LiDAR as their main sensing modality due to its geometrical richness. Rolling shutter LIDARs are particularly common, in which an array of lasers scans the scene from a rotating base. Points are emitted as a stream of packets, each covering a sector of the 360° coverage. Modern perception algorithms wait for the full sweep to be built before processing the data, which introduces an additional latency. For typical 10Hz LiDARs this will be 100ms. As a consequence, by the time an output is produced, it no longer accurately reflects the state of the world. This poses a challenge, as robotics applications require minimal reaction times, such that maneuvers can be quickly planned in the event of a safety-critical situation. In this paper we propose STROBE, a novel approach that minimizes latency by ingesting LiDAR packets and emitting a stream of detections without waiting for the full sweep to be built. STROBE reuses computations from previous packets and iteratively updates a latent spatial representation of the scene, which acts as a memory, as new evidence comes in, resulting in accurate low-latency perception. We demonstrate the effectiveness of our approach on a large scale real-world dataset, showing that our approach far outperforms the state-of-the-art when latency is taken into account, and matches the performance in the traditional setting.

1 Introduction

Perceiving the world is a critical task in modern robotics applications. Self-driving vehicles must first process sensory information to perform object detection and estimate the free space before attempting to plan a safe and comfortable maneuver towards the goal. LiDAR has become the main sensing modality in most self-driving vehicles due to the geometrical richness it provides. Most prevalent LiDAR sensors operate by collecting a rotating scan of the environment, typically completing revolutions at a 10hz rate. However, as the sensor rotates, observations arrive as a stream of spatio-temporal points (x, y, z, t) grouped in fine-grained *packets*, each spanning approximately 10ms. This gives rise to a rolling shutter effect shown in [Figure 1](#), where objects in different locations are observed asynchronously.

Modern autonomous systems accumulate the LiDAR packets into a full 360° sweep before running perception. This waiting time adds significant latency to the pipeline, particularly for objects that were seen in the earlier packets in the sweep. It also introduces an erroneous assumption that all observations in the full sweep are made synchronously. In reality, when the perception model receives the input, there is already a discrepancy between the outdated observations and the true state of the world, illustrated in [Figure 1](#). Furthermore, there is a temporal discontinuity in the sweep where the earliest and the latest packets meet which creates artifacts in the point cloud.

For safety-critical applications like self-driving, even minimal delays may result in catastrophic outcomes. For example, in the presence of high-speed vehicles, building a sweep from a 10Hz LiDAR introduces a latency of 100ms, which translates to several meters of error in free space estimation. Having lower latency is crucial in safety-critical situations where the vehicle must quickly perceive and react to avoid harmful events. Therefore, it is important to process incoming sensory information with minimal latency.




Figure 1: Objects are observed at different times when building a full LiDAR sweep, indicated as the solid boxes. If a full sweep from a 10hz LiDAR is accumulated before detection, a latency of 100ms will be introduced and by the time a detection is available (Packet 10) it no longer reflects the state of the world.

Processing individual LiDAR packets can be challenging, since only a small sector of the scene is observable as illustrated in Figure 1. Objects of interest are often fragmented across different LiDAR packets, particularly when close to the sensor. Coincidentally, that is also when high accuracy and low latency are the most important as *close range* objects are typically the most critical to safety. Thus, individual packets alone may be insufficient for high quality detections, making it necessary to incorporate past observations.

Existing LiDAR object detectors generally assume access to a full 360 degree sweep, or a large subregion (e.g., front view) that spans all objects of interest. As such, these models do not explicitly reason about objects split across multiple observations. As shown in our experiments, directly adopting full-sweep models for processing individual LiDAR packets is not a good solution due to the partial observation and lack of global context. Conversely, exploiting multiple sweeps [1, 2] provides richer geometrical evidence as more LiDAR points are collected over time. However, most current solutions are computationally inefficient as each packet would be processed as many times as the duration of the history. As such, naively aggregating historical sensory information at the input level is not amenable to emitting low latency object detections from fine-grained LiDAR packets.

In this paper we propose STROBE, a novel detection model which exploits the sequential nature of LiDAR observations and efficiently reuses past computation to stream low latency object detections from LiDAR packets. Our approach voxelizes the incoming LiDAR packets into a Bird’s-Eye View (BEV) grid, and uses an efficient convolutional backbone to process only the relevant region. Furthermore, we introduce a multi-scale spatial memory that is read and updated with each LiDAR packet. This allows us to reuse past computation, and make the incremental processing of incoming LiDAR packets lightweight and efficient. Importantly, we achieve an end-to-end latency of 21 ms (from observing an actor to emitting a detection) on an NVIDIA 2080Ti: 10ms for accumulating a packet and 11ms for model inference. In contrast, even fast full sweep detectors [3] operate at an order of magnitude higher latencies: Taking 100ms to accumulate the sweep and another 28ms for model inference, for a total of 128ms.

Our second contribution is a novel large scale benchmark for evaluating streaming object detection from LiDAR packets. Unlike existing public datasets, PACKETATG4D contains LiDAR data at the packet level, along with accurate ego-pose and associated object bounding box annotations at the same temporal resolution (i.e., 100Hz). We also propose a novel metric *latency-aware mAP* to explicitly take latency into account when evaluating perception. We show that our approach far outperforms the state-of-the-art when the data buffering latency is taken into account, while still matching the performance in the conventional setting.

2 Related Work

3D object detection has made tremendous progress in recent years due to the advances of deep learning and the availability of large-scale labeled datasets. The topic of how to effectively process LiDAR data has received significant attention and many approaches have been proposed. Point clouds have been processed in perspective format using a range image [4, 5]. By converting the point cloud into an image, these approaches can leverage the vast body of knowledge on 2D object detection to build good architectures for the task. However, such methods suffer from the same challenge present in 2D detection: high variance in receptive field requirements as a function of depth.




Figure 2: Existing point cloud perception methods wait 100ms to accumulate the full sweep (**right**). STROBE (**left**) is able to process each packet and emit new detections with high accuracy and minimal latency, while leveraging global context by continuously updating a spatial memory that keeps track of previously seen packets.

To tackle these issues, some methods perform 3D detection directly on the unstructured 3D points. This is usually achieved through first extracting local signatures with a fully connected layer [6, 7, 8, 9] or by using deformable filters [10]. An alternative framework is to voxelize the points into a regularly spaced 3D grid, making reasoning on point clouds amenable to convolutional architectures. Early works [11, 12, 13, 2] leverage 3D convolutions, but they are memory intensive. Others [14, 15, 16] exploit the sparsity of point clouds to reduce redundant computation and make higher resolution processing feasible. BEV detectors [3, 17, 1] avoid heavy computation by exploiting efficient 2D convolutions over a top-down pseudo-image of the scene. Other methods have leveraged hybrid representation of points and voxels [18, 19, 20, 21, 22] to exploit the benefits of both representations.

However, the aforementioned methods assume a full sweep is available, which requires the sensor to complete a full rotation and incurs latency. Previous works have explored the problem of latency in different settings, for instance on the effect of model runtime for 2D object detection [23], or how the temporal aspect of point clouds is relevant for odometry and mapping [24, 25]. Concurrent work [26] has considered streaming object detections from a rolling shutter LiDAR. However, their model uses an LSTM to maintain the state, which does not leverage the spatial nature of the problem. Furthermore, their evaluation does not capture the impact latency has on the accuracy of state estimation.

3 Low Latency Detection on Streaming LiDAR

In this paper, we propose STROBE, a low-latency object detector that emits detections from streaming LiDAR observations. As illustrated in Figure 2, as the LiDAR sensor spins, it yields data in sector packets (each roughly spanning 36° in our 10Hz LiDAR). As opposed to previous models, which buffer this data into a full sweep before processing, our proposed method operates at the packet level. In doing so, we lower our latency by 90ms. A fundamental component to our approach is a novel spatial memory module design to reuse past computation, and make incremental processing of incoming LiDAR packets lightweight and effective.

3.1 Streaming Object Detection

The overall architecture of our model is illustrated in Figure 3. The network takes as input a LiDAR packet and an HD map, which is useful as a prior on the location of actors (e.g., a vehicle is more likely to be on the road than on the sidewalk). For each packet we first voxelize the points and rasterize into a BEV pseudo-image with height as the channel dimension [3]. Following [17, 1], we also rasterize the map into a BEV pseudo-image, where each channel corresponds to a different layer of the map (e.g., crosswalks, roads, etc). We then extract features using our novel regional convolutions (Figure 3 – a, b), which only compute features in the rectangular area defined by the packet. A latent spatial representation of the scene is then maintained using a memory module (Figure 3 – c, d, e). Lastly, we channel-wise concatenate multi-scale features and regress detection parameters using our output header (Figure 3 – f).

Regional Convolution Layer: To reduce latency while leveraging the proven strength of BEV representations and 2D convolutions, we propose to process the input with a local operator, which we call regional convolution. Specifically, for an input x and coordinates x_0, x_1, y_0 and y_1 , we extract features y only on the region $\mathbf{x}[x_0 : x_1, y_0 : y_1]$, where the brackets denote indexing at the rectangle defined by the coordinate ranges. This allows us to leverage locality to minimize wasted computation.

$$y = f_{\text{region}}(\mathbf{x}[x_0 : x_1, y_0 : y_1], \mathbf{w}) \quad (1)$$

In practice, f_{region} is a sequence of 2D convolution, ReLU activation and Group Normalization [27]. Furthermore, for both the LiDAR packet and HD map, the region coordinates are defined as the minimal rectangle that fully encloses all points in the LiDAR packet. This is illustrated in Figure 3 – a, b.

Spatial Memory: While regional convolutions allow us to efficiently ingest packets, independently processing them is not sufficient for accurate perception since objects will often be fragmented across many packets. Furthermore, a single observation of an object far away will typically yield few points due to the sparsity of the sensor at range. We would thus like to leverage information from previous scans of the region. However, naively processing the history of observations every time we receive a packet results in redundant computation and slow inference. Instead, our approach iteratively builds a global spatial memory from a series of partial observations while at the same time producing new detections with each LiDAR packet, Figure 3 – c. This enables us to re-use past computation and produce low-latency and accurate detections. Importantly, the LiDAR points are registered on a consistent coordinate frame defined by a continuous ego-pose. The memory is aligned with this pose by bilinearly resampling its features to account for ego-motion with every new packet (Figure 3 – c, d). This guarantees that the LiDAR and map features are consistently aligned with the spatial memory in the same coordinate frame.

Memory Update: As each LiDAR packet arrives, the spatial memory is incrementally updated with new local features to reflect the latest state (Figure 3 – d). Each update step is done through aggregation of the current memory state \mathbf{m} and the incoming local features y . Specifically, we employ a channel reduction with learned parameters \mathbf{w} as follows

$$\mathbf{m}'[x_0 : x_1, y_0 : y_1] = f_{\text{memory}}(\mathbf{m}[x_0 : x_1, y_0 : y_1], y, \mathbf{w}). \quad (2)$$

In practice, f_{memory} channel-wise concatenates \mathbf{m} and y , resulting in a tensor with $2c$ channels, then applies two blocks of 2D convolution, ReLU activation and Group Normalization, with the second block bringing the number of channels back to c . This is illustrated as the red dotted arrows in Figure 3 – e.

Multi-Scale Backbone: In order to leverage the semantic representations of feature maps at different scales (i.e., richer geometry on higher resolutions; richer semantics on lower) we employ a multi-scale backbone for the extraction of both LiDAR and HD map features. Together with the spatial memory at each scale, the benefits of this are twofold: It allows the model to regress accurate and low latency detections from partial observations by remembering the features from immediately preceding packets. It also makes it possible for the network to persist long term features that are useful to detect objects through occlusion over multiple sweeps as well as overwrite previous features when stronger evidence is available.

Architecture Details: We employ a BEV grid with resolution of 0.2m for each voxel. This grid then goes through 4 blocks of [2, 2, 3, 6] Regional Convolution layers with [24, 64, 128, 256] channels, followed by Max Pooling with a stride of 2. Each block has a corresponding Spatial Memory that holds the pre-pooling state of the features. In parallel, features are extracted from the HD map with a backbone that consists of a sequence of 4 blocks with [2, 2, 3, 3] Regional Convolution layers with [16, 32, 64, 128] channels. After each block, Max Pooling with a stride of 2 is employed. The feature maps from each block of both the LiDAR and HD map backbones are then bilinearly resized to a common resolution of 0.8m, channel-wise concatenated, and processed by one last block of 4 Regional Convolutions with 256 channels.

Detection Header: We perform multi-class BEV detection for vehicles, cyclists, and pedestrians via a single-stage detection header consisting of 2 convolutional layers that predict the classification




Figure 3: STROBE performs regional convolution on LiDAR packets and HD maps, using a multi-scale spatial memory for global reasoning. Δ is interpolation and \parallel channel-wise concatenation.

and regression targets for each cell in the fused feature map (hereinafter referred to as "anchors"). All objects are defined via their centroid (b_x, b_y) and confidence σ , whereas cyclists and vehicles also have length, width, and heading (b_l, b_w, b_ϕ) in BEV. For the confidence, we predict its logit $\log \frac{\sigma}{1-\sigma}$. We define the centroid of the box (b_x, b_y) as an offset $(\Delta x, \Delta y)$ from the coordinates of the center point of its anchor pixel (a_x, a_y) :

$$(b_x, b_y) = (a_x + \Delta x, a_y + \Delta y). \quad (3)$$

For the vehicle dimensions we predict $[\log l, \log w]$, which encourages the network to learn a prior on the dimension of the boxes (low variance should be expected from the dimension of vehicles). The heading b_ϕ is parameterized by the tangent value. In particular, we predict a signed ratio so that the specific quadrant can be retrieved:

$$b_\phi = \arctan \frac{\theta_1}{\theta_2}. \quad (4)$$

3.2 Learning

Following common practice in object detection [28], we employ a multi-task loss over classification and bounding box regression to optimize the model (using $\alpha = 2.0$), i.e:

$$\mathcal{L} = \mathcal{L}_{\text{reg}} + \alpha \mathcal{L}_{\text{cls}}. \quad (5)$$

Regression Loss: It is defined as the weighted sum of the smooth ℓ_1 loss [29] between the ground truth vector of detection parameters $\hat{\mathbf{y}} = [\Delta x, \Delta y, \log w, \log l, \theta_1, \theta_2]$ and predictions \mathbf{y} with $\gamma = [1, 1, 1, 1, 2, 2]$. Note that $\log w, \log l, \theta_1$ and θ_2 are not considered for pedestrians since we are only concerned with predicting their centroid.

$$\mathcal{L}_{\text{reg}}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{N} \sum_{i=0}^N \gamma \cdot \text{smooth}_{\ell_1}(\mathbf{y}_d^i - \hat{\mathbf{y}}_d^i) \quad (6)$$

Classification Loss: It is defined as the binary cross entropy between the predicted scores and the ground truth. Due to severe class imbalance between positive $\hat{\mathbf{y}}_{\text{pos}}$ and negative $\hat{\mathbf{y}}_{\text{neg}}$ anchors given that most pixels in the BEV scene do not contain an object, we employ hard negative mining:

$$\mathcal{L}_{\text{cls}}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{N} \sum_{i=0}^N \hat{\mathbf{y}}_{\text{pos}}^i \log \mathbf{y} + \frac{1}{K} \sum_{i=0}^N \mathbb{1}[i \in \mathcal{N}_K] (1 - \hat{\mathbf{y}}_{\text{neg}}^i) \log(1 - \mathbf{y}) \quad (7)$$

where \mathcal{N}_K is a set containing K hard negative anchors. This is obtained by sampling 750 anchors for vehicles, 1500 for cyclists and pedestrians, and picking the 20 with highest loss for each class.

Sequential Training: Due to the sequential nature of the memory, the model is trained sequentially through examples that contain 50 packets (corresponding to 0.5s). Back-propagation through time is used to compute gradients across the memory. Furthermore, the model is trained to remember by supervising it on objects with 0 points, as long as the object was seen in any of the previous packets. In practice, due to GPU memory constraints, we only compute the forward pass in the first 40 packets to warm-up the memory, then forward and backward through time in the last 10 packets.




Figure 4: In contrast to the commonly used mAP, our proposed metric takes into account the latency between observation time and detection emission time.

	Accumulation (ms)	Inference (ms)	Total Latency (ms)
PointRCNN [18]	100	390	490
PointPillars [21, 16]	100	37	137
HDNET [17]	100	28	128
Our STROBE	10	11	21

Table 1: **End-to-end Latency:** We report the end-to-end latency in ms of the models as defined by the time it takes to accumulate the data, run inference and decode detections. Accumulation considers a LiDAR operating at 10hz and inference is done with a NVIDIA 2080Ti.

4 Experimental Evaluation

We evaluate our model on a real world dataset for 3D object detection. In particular, we compute mean average precision (mAP) in the default detection setting (using full 360° sweeps) and propose a new metric that takes into account the latency incurred by different input granularities (i.e., per-packet processing vs. sweep building). Our experimental results show that our model far outperforms the baselines in the per-packet setting while remaining competitive with the state-of-the-art in the full sweep setting. Furthermore, our latency evaluation also uncovers a problem with the mAP metric in the default setting as it does not accurately measure real world performance.

Dataset: Since there is no public available dataset that provides packets, we collect a new dataset, PACKETATG4D, containing 6500 snippets with diverse conditions (e.g., geographical, lighting, road topology, vehicle types). The LiDAR rotates at a rate of 10hz and emits new packets at 100Hz – each roughly covering a 36° region – for a total of 16,250,000 packets (1,625,000 frames). Accurate ego-pose is available for each LiDAR packet via a commercial localization system. Labels provide both the spatial extents and motion of vehicles, cyclists and pedestrians, from which we can extract accurate bounding boxes at discrete observation times as well as in continuous time through the use of a precise motion model. Note that if the observation of an instance is split across packets, each packet will have an instance of the label according to the pose at the timestamp of the packet.

Baselines: We provide a wide range of baselines that exploit different representations. **HD-NET** [17] is a detection model that processes input point clouds into occupancy voxels and performs 2D convolution in BEV using the z axis voxels and HD maps as feature channels. **PointRCNN** [30] processes raw LiDAR inputs using a PointNet [6] backbone to perform foreground segmentation and generate region-of-interest (RoI) proposals. The RoI proposals are then processed by a classification and bounding box refinement network to output 3D detections. **PointPillars** [21, 16] groups input points into discrete bins in BEV and uses PointNet [6] to extract features of each bin. The BEV features are then processed with 2D convolutions to generate detection outputs. Note that the PointRCNN and PointPillar baselines do not make use of HD maps.

Metrics: We evaluate our method using mean average precision (mAP) as our metric with IOU thresholds of [0.5, 0.7] for vehicles, [0.3, 0.5] for cyclists. For pedestrians, we use the ℓ_2 distance to centroid with thresholds [0.5m, 0.3m] since we treat the detections as circles with a fixed radius, thus only the centroid is predicted. We evaluate with latency-aware labels that take into account the delay introduced by aggregating consecutive packets (*Latency mAP*). We refer the reader to Figure 4 for an illustration of this metric. We re-define the detection label for each object in the scene as its state at detection time (green box), rather than observation time (red box), which does not accurately

Model	Packet Stream						Full Sweep					
	Vehicle		Pedestrian		Cyclist		Vehicle		Pedestrian		Cyclist	
	0.5	0.7	.5m	.3m	0.3	0.5	0.5	0.7	.5m	.3m	0.3	0.5
HDNET [17]	75.6	63.6	71.0	63.9	21.3	15.3	79.6	57.8	80.2	69.8	54.6	33.8
PointPillars [21, 16]	66.8	47.7	53.4	49.2	16.8	6.1	84.2	61.1	74.4	68.9	56.1	34.9
PointRCNN [18]	70.2	63.1	49.3	47.5	28.4	25.9	72.4	57.4	54.8	52.4	31.9	26.7
Our STROBE	91.8	80.5	80.3	72.5	60.8	40.7	86.4	66.4	76.7	67.8	61.0	39.5

Table 2: **Latency mAP**: Labels are considered at detection emission times.

Model	Packet Stream						Full Sweep					
	Vehicle		Pedestrian		Cyclist		Vehicle		Pedestrian		Cyclist	
	0.5	0.7	.5m	.3m	0.3	0.5	0.5	0.7	.5m	.3m	0.3	0.5
HDNET [17]	75.7	63.7	71.1	64.1	21.3	15.3	89.5	77.2	84.3	74.7	68.3	45.5
PointPillars [21, 16]	66.9	48.0	53.5	49.3	16.9	6.2	84.8	70.6	74.2	69.2	56.1	36.3
PointRCNN [18]	70.3	63.3	49.3	47.5	28.4	25.8	73.1	66.9	54.6	52.7	31.4	26.9
Our STROBE	91.8	80.5	80.3	72.5	60.8	40.7	87.4	76.1	76.9	69.0	61.3	41.4

Table 3: **Common mAP**: Labels are considered at their observation times.

reflect the current state of the world. The benefits of this metric are twofold: (1) It evaluates how well the detector models the state of the real world and the quality of the information that would be used by downstream motion planning, and (2) it allows for a direct comparison with standard detection metrics, thus making apparent the effects of latency.

End-to-end Latency: Since implementations might differ, we did not consider model inference times in the latency aware detection metric. However, it is an important factor in the end-to-end latency for safety since it indicates the minimal amount of time the system would require to be able to recognize an actor, i.e., the time taken for sensor data acquisition, model inference, and emission of a corresponding detection for the actor to downstream systems. We report end-to-end latency timings in Table 1; our approach leads to a much faster (on average 6x!) detection emission time.

Latency-aware Detection: Table 2 shows our results for PACKETATG4D. In the leftmost setting – *Packet Stream* – all models are first trained on detection using LiDAR packets (as opposed to full sweeps) and evaluated using the state of the labels at the time of detection (i.e., green box in Figure 4). Our model far outperforms the baselines, which do not have memory and struggle with partial observations (i.e., a single packet as opposed to the full sweep). In the right portion of the table – *Full Sweep* – the models are trained using a traditional full sweep setting and evaluation is done using the label states at the end of the sweep (therefore in the worst case an object could move for 100ms before evaluation).

Latency-unaware Detection: We additionally evaluate in the standard object detection setting, not taking into account the sweep building latency and using the labels for each object in the scene at the time of observation (i.e., when the LiDAR hit the object). The leftmost columns of Table 3 show the results of the models trained in a packet setting. A key takeaway from these results stems from comparing the numbers in the "Packet Stream" setting between Tables 2 and 3, which shows that the 10ms latency introduced by accumulating a single packet is negligible in the mAP settings we evaluate, since performance remains the same. Conversely, comparing the "Full Sweep" setting in Tables 2 and 3 shows considerable degradation in metrics. This indicates that the performance of full sweep models in the real world would be considerably lower.

Ablation Studies: We first ablate the memory component of the model. In particular, we evaluate two alternative approaches: (1) *No Memory*: A memoryless instantiation of our model; (2) *Attention*: A memory module that uses linear attention to update the spatial memory (see supplementary for more details). As shown in Table 4, memory is a fundamental component for effective perception from partial observations. Furthermore, the attention based memory updates are outperformed by our approach which learns the aggregation function through convolutions. We also evaluate our model without the HD map component to evaluate its importance. The results in Table 4 (*No Map* row)

	Vehicle		Pedestrian		Cyclist	
	0.5	0.7	.5m	.3m	0.3	0.5
No Memory	75.6	63.6	71.0	63.9	21.3	15.3
Attention	89.3	78.2	75.9	67.9	53.5	35.3
No Map	90.6	79.9	79.3	71.8	59.6	40.5
Our STROBE	91.8	80.5	80.3	72.5	60.8	40.7

Table 4: **Ablation studies:** Our multi-scale spatial memory is a critical component in our model. Using maps is beneficial but not critical. Labels are at detection emission time.

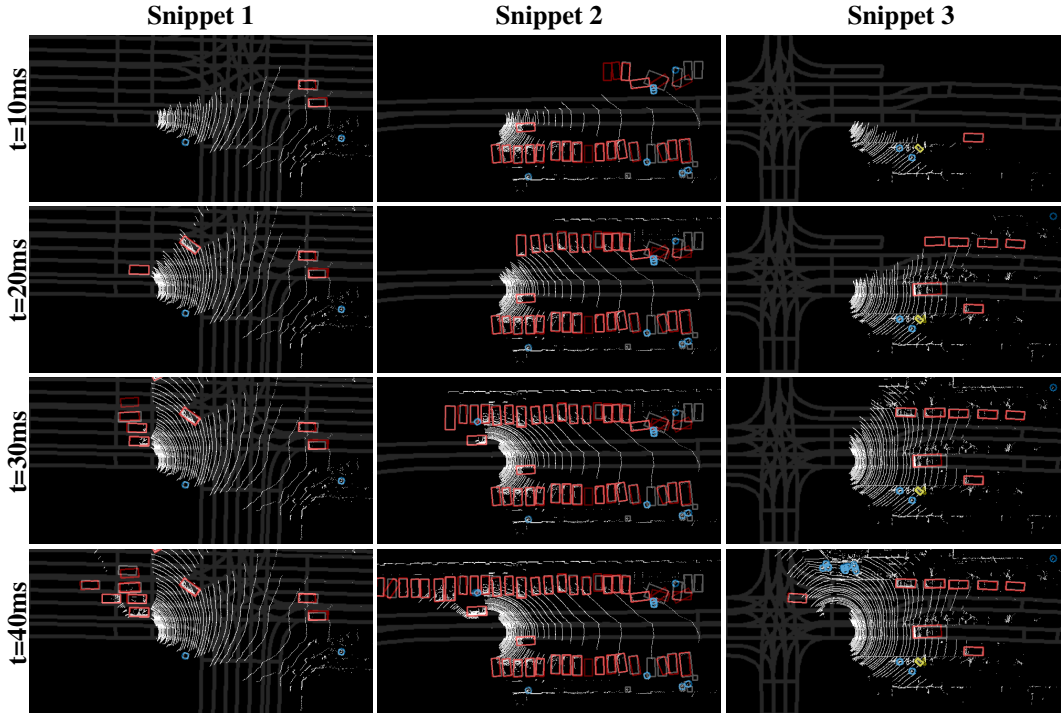


Figure 5: Qualitative results of STROBE. Each column is a sequence of packets from the same snippet. Detected vehicles are shown in red, cyclists in yellow and pedestrians in blue.

show that while the map backbone proved to be overall beneficial to the model, is not a fundamental component as its removal does not lead to major degradations in metrics.

Qualitative Results: The qualitative results in Figure 5 show the predictions of the model over 4 consecutive packets in 3 snippets. The network is able to predict boxes even before points are visible due to the memory module. It can also update the positions of detections as new points arrive to best exploit the evidence.

5 Conclusion

We have proposed a novel method for perception of point cloud streaming data. Our approach produces highly accurate object detections at very low latency by using regional convolutions over individual LiDAR packets alongside a spatial memory that keeps track of previous observations. We also introduced a new latency-aware metric that quantifies the cost of data buffering, and how that affects the quality of the models in the real world, which are inevitably affected by latency. Results on the large-scale PACKETATG4D show that our approach far outperforms the state-of-the-art in the packet setting that takes into account latency, while being competitive in the commonly adopted full sweep setting. For future work, we intend to expand the use of the memory module for long term tracking through occlusion and motion forecasting.

References

- [1] S. Casas, W. Luo, and R. Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *Conference on Robot Learning (CoRL)*, pages 947–956, 2018.
- [2] W. Luo, B. Yang, and R. Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3569–3577, 2018.
- [3] B. Yang, W. Luo, and R. Urtasun. Pixor: Real-time 3d object detection from point clouds. In *conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7652–7660, 2018.
- [4] B. Li, T. Zhang, and T. Xia. Vehicle detection from 3d lidar using fully convolutional network. In *Robotics: Science and Systems (RSS)*, 2016.
- [5] G. P. Meyer, A. Laddha, E. Kee, C. Vallespi-Gonzalez, and C. K. Wellington. Lasernet: An efficient probabilistic 3d object detector for autonomous driving. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12677–12686. IEEE, 2019.
- [6] C. R. Qi, H. Su, K. Mo, and L. 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*, pages 652–660, 2017.
- [7] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen. Pointcnn: Convolution on x-transformed points. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 820–830, 2018.
- [8] B.-S. Hua, M.-K. Tran, and S.-K. Yeung. Pointwise convolutional neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 984–993. IEEE, 2018.
- [9] S. Wang, S. Suo, W.-C. Ma, A. Pokrovsky, and R. Urtasun. Deep parametric continuous convolutional neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2589–2597. IEEE, 2018.
- [10] Y. Xiong, M. Ren, R. Liao, K. Wong, and R. Urtasun. Deformable filter convolution for point cloud reasoning. In *Advances in Neural Information Processing Systems (NeurIPS) – Workshop on Sets & Partitions*, 2019.
- [11] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1912–1920, 2015.
- [12] D. Maturana and S. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 922–928, 2015.
- [13] B. Li. 3d fully convolutional network for vehicle detection in point cloud. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1513–1518. IEEE, 2017.
- [14] G. Riegler, A. Osman Ulusoy, and A. Geiger. Octnet: Learning deep 3d representations at high resolutions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3577–3586, 2017.
- [15] M. Ren, A. Pokrovsky, B. Yang, and R. Urtasun. Sbnet: Sparse blocks network for fast inference. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8711–8720, 2018.
- [16] Y. Yan, Y. Mao, and B. Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018.
- [17] B. Yang, M. Liang, and R. Urtasun. Hdnet: Exploiting hd maps for 3d object detection. In *Conference on Robot Learning (CoRL)*, pages 146–155, 2018.
- [18] Y. Chen, S. Liu, X. Shen, and J. Jia. Fast point r-cnn. In *International Conference on Computer Vision (ICCV)*, pages 9775–9784. IEEE, 2019.

- [19] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [20] Z. Yang, Y. Sun, S. Liu, X. Shen, and J. Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *International Conference on Computer Vision (ICCV)*, pages 1951–1960. IEEE, 2019.
- [21] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019.
- [22] 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. *arXiv preprint arXiv:1910.06528*, 2019.
- [23] M. Li, Y.-X. Wang, and D. Ramanan. Towards streaming image understanding. *European Conference on Computer Vision (ECCV)*, 2020.
- [24] H. Alismail, L. D. Baker, and B. Browning. Continuous trajectory estimation for 3d slam from actuated lidar. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6096–6101. IEEE, 2014.
- [25] J. Zhang and S. Singh. Low-drift and real-time lidar odometry and mapping. *Autonomous Robots*, 41(2):401–416, 2017.
- [26] W. Han, Z. Zhang, B. Caine, B. Yang, C. Sprunk, O. Alsharif, J. Ngiam, V. Vasudevan, J. Shlens, and Z. Chen. Streaming object detection for 3-d point clouds. *arXiv preprint arXiv:2005.01864*, 2020.
- [27] Y. Wu and K. He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.
- [28] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems (NeurIPS)*, pages 91–99, 2015.
- [29] R. B. Girshick. Fast R-CNN. *CoRR*, abs/1504.08083, 2015. URL <http://arxiv.org/abs/1504.08083>.
- [30] S. Shi, X. Wang, and H. Li. Pointtrcnn: 3d object proposal generation and detection from point cloud. *CoRR*, abs/1812.04244, 2018. URL <http://arxiv.org/abs/1812.04244>.