

# DR-SPAAM: A Spatial-Attention and Auto-regressive Model for Person Detection in 2D Range Data

Dan Jia<sup>1</sup>, Alexander Hermans<sup>1</sup>, and Bastian Leibe<sup>1</sup>

**Abstract**—Detecting persons using a 2D LiDAR is a challenging task due to the low information content of 2D range data. To alleviate the problem caused by the sparsity of the LiDAR points, current state-of-the-art methods fuse multiple previous scans and perform detection using the combined scans. The downside of such a backward looking fusion is that all the scans need to be aligned explicitly, and the necessary alignment operation makes the whole pipeline more expensive – often too expensive for real-world applications. In this paper, we propose a person detection network which uses an alternative strategy to combine scans obtained at different times. Our method, Distance Robust SPatial Attention and Auto-regressive Model (DR-SPAAM), follows a forward looking paradigm. It keeps the intermediate features from the backbone network as a template and recurrently updates the template when a new scan becomes available. The updated feature template is in turn used for detecting persons currently in the scene. On the DROW dataset, our method outperforms the existing state-of-the-art, while being approximately four times faster, running at 87.2 FPS on a laptop with a dedicated GPU and at 22.6 FPS on an NVIDIA Jetson AGX embedded GPU. We release our code in PyTorch and a ROS node including pre-trained models.

## I. INTRODUCTION

Detecting persons in the surrounding environment is a key requirement for many robotic applications including search and rescue, security, and health care. Currently, this is often accomplished using multiple RGB(-D) cameras, in combination with a deep learning based object detector [1], [2], [3]. However, the limited field of view of such cameras limits the detection to a narrow frustum. Furthermore, the inaccurate depth measurements at far ranges make accurate person localization difficult in 3D space. Instead, a 2D LiDAR provides accurate range measurements with a large field of view at high acquisition rates. Thus, it is a promising sensor choice for detecting persons.

However, the limited information contained in the sparse range measurements from a 2D LiDAR is a key challenge towards reliable person detection. Recent developments have shown that it is beneficial to combine several previous scans in order to detect objects [4], [5]. In particular, Beyer *et al.* reported improved detection results by accumulating five previous scans, compared to their single scan baseline [4]. The downside, however, is the increased computational cost. Due to the ego-motion of the LiDAR, as well as the motion of objects in the scene, scans recorded at different times are not perfectly aligned, and an expensive alignment operation has to be carried out in order to fuse scans for downstream processing. In the case of [4], the alignment is done using

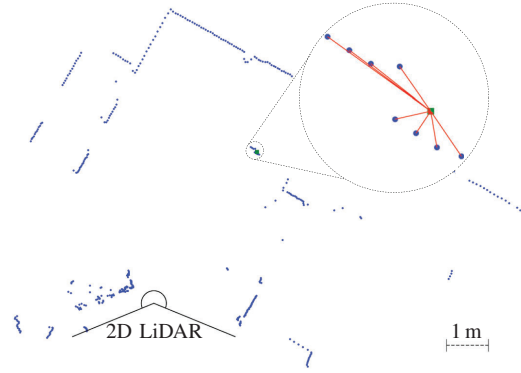


Fig. 1: The bird's-eye view of a 2D LiDAR scan (blue dots) with a person in the scene (green square). For each point DR-SPAAM outputs a classification label and an offset to the center of the person (red lines), which are then grouped into detections using a post-processing step.

the odometry information in addition to repetitive sampling on previous scans. This alignment has linear computational cost with respect to the number of scans, and using five previous scans already makes the overall detection pipeline too expensive for real-time processing on mobile platforms.

In principle, aligning and fusing previous scans follows a backward looking paradigm for aggregating temporal information. In contrast, a forward looking paradigm simply keeps a representation based on current measurements and recurrently updates this representation when a new measurement becomes available. This representation, which incorporates all previous measurements, is then used for downstream processing. An example of such a forward looking paradigm is found in the field of video object detection, where researchers have used memory modules which recurrently take input features at each frame of a video sequence and output a refined prediction at the current frame [6], [7], [8].

In this paper, we propose a person detection network which aggregates temporal information following a forward looking paradigm. Our method uses the existing architecture of the DROW detector [4], which takes as input a 2D scan and predicts for each point a classification label and an offset vector pointing towards the center of the person (Fig. 1). We augment DROW with an auto-regressive model to aggregate the intermediate features from the backbone network from each scan, allowing our method to use information from all past measurements. Instead of explicitly aligning the intermediate features from different scans, we use a spatial attention mechanism to associate features from neighboring locations based on their similarity, significantly reducing the

<sup>1</sup>All authors are with the Visual Computing Institute, RWTH Aachen University. lastname@vision.rwth-aachen.de

computational cost. We call our method Distance Robust SPatial Attention and Auto-regressive Model (DR-SPAAM). Evaluated on the original DROW dataset, DR-SPAAM outperforms the previous approaches, while running at 87.2 FPS on a laptop with a dedicated GPU, or at 22.6 FPS on an NVIDIA Jetson AGX. The high frame rate of DR-SPAAM makes it suitable for many robotic applications.

In summary, we make the following key contributions:

- We propose a spatial attention and auto-regressive model that fuses information from previous LiDAR scans without the need of explicit alignment operation.
- We propose DR-SPAAM, a fast 2D LiDAR based person detector using the spatial attention and auto-regressive model. Our proposed method outperforms the previous state-of-the-art in 2D range data based person detection both in speed and detection performance.
- We release our implementation in PyTorch, including a ROS node with pre-trained models, to facilitate easy deployment in robotic projects<sup>1</sup>.

## II. RELATED WORKS

### A. Person Detection from 2D Range Data

Person detection from 2D range data has a long standing history in the robotics community. While early methods are mainly based on heuristics to find specific shapes in range data, the most common paradigm is to segment a scan into connected segments, compute a set of hand-crafted features for each of these, and finally classify them to create detections [9], [10], [11]. Common methods can be divided into approaches that detect, and then track, individual legs in order to detect persons [10], [11], or approaches that directly aim to include both legs in one segment [9]. Optimally, one would learn the representation of persons directly from data to avoid making such hard design choices. The DROW detector [12] was the first deep learning based walking aid detector working on 2D range data and was later extended to additionally detect persons [4]. One key aspect that improved the person detection results was the integration of temporal information. However, this significantly increases the runtime, making it infeasible for online use. We propose a new person detector that combines the existing architecture of the DROW detector with a forward looking temporal integration module, outperforming the original DROW version both in speed and in detection quality.

Others have used deep learning based methods on 2D range data. Ondruška *et al.* [13] create an occupancy grid from a stationary 2D LiDAR and predict class labels and future grid configurations based on RNNs for every grid cell. In a later version they extend this approach to work with moving LiDARs [14]. However, in both cases they do not create object detections or tracks and as such cannot be compared to a person detector.

### B. Object Detection in 3D Point Clouds

Many works have focused on detecting objects in point clouds obtained from a 3D LiDAR, since such a task plays an important role in autonomous driving applications. Point clouds, as a data representation, inherently lack the definition of structure and neighborhoods, and thus prohibit the use of popular CNN architectures. To solve this problem, earlier works have leveraged image-based object detection methods, either by projecting the point cloud onto an image plane, or by popping up 2D detections made on the RGB image with known extrinsic calibration [15], [16], [17], [18]. The runtime and accuracy of these methods, however, are bottlenecked by the employed 2D object detector. Later developments in the field, including current state-of-the-art methods, utilize the full point cloud without projection. Methods like VoxelNet [19], or SECOND [20] run (sparse) convolutions [21], [22] on structured 3D voxel grids converted from point clouds, while other methods like PointRCNN [23] directly process unstructured point clouds, using PointNet [24] inspired backbones [25], [26], [27]. Yet others combine both design schemes into two-stage detectors [28], [29], [30].

A particularly interesting work is the VoteNet proposed by Qi *et al.* [31]. Given a 3D point cloud, the first stage of the VoteNet regresses for each point an offset vector towards the object center (a vote), similar to the DROW detector, but instead of using a post-processing step, it utilizes another sub-network to group the per-point predictions into bounding box proposals. Thus, the whole network is end-to-end trainable. Similarly, 3D-MPA [32] uses voting for instance proposals which are grouped into point-level instance masks. This differentiable vote aggregation could also be an interesting approach for 2D LiDARs. However, our focus lies on the temporal integration of a sequence of scans, and these two approaches are orthogonal.

Although from a hardware point of view, a 2D LiDAR does share similar working principles with its 3D counterpart, the methods operating on 3D point clouds cannot be naively applied to range data obtained from a 2D LiDAR. To the best of our knowledge, no work exists that directly applies existing 3D methods to 2D range data, and given the sparsity of information in 2D range scans, it remains to be seen if a naive adaptation is possible at all.

### C. Video Object Detection

Video object detection is a special object detection task, where the input is a video sequence showing the same objects in multiple frames. It is thus possible to utilize information from one frame to aid detection in another. This temporal propagation of information is important, since objects can undergo large appearance changes caused by fast motion, occlusion, or change of camera angle. Earlier approaches [33], [34], [35] detect objects in each frame independently and apply sequence-level post-processing on the obtained bounding boxes. Such approaches cannot be optimized in an end-to-end fashion. Later approaches focus on directly aggregating features across frames, either explicitly

<sup>1</sup><https://github.com/VisualComputingInstitute/DR-SPAAM-Detector>

aligning features using optical flow [36], [37], [38], or using a memory network to aggregate features [6], [7], [8].

Similar to these approaches, our method uses a network to aggregate features across consecutive scans. Instead of using a designated memory module, we use an auto-regressive model, which propagates information from the previous scan to the next with an exponentially decaying weight. In video object detection, the key is to be able to aggregate long term features, since neighboring frames often have similar appearance and introduce little new information. Thus, a memory network is often used. Instead, we focus on aggregating short term features from consecutive scans, with the goal of enriching the information available for detection. Hence, an auto-regressive model is more suitable compared to a more complex memory module. Furthermore, we propose to use a similarity based spatial attention model [39], [40] to first fuse nearby spatial information before the temporal aggregation.

### III. METHOD

In this section we describe in detail our proposed method for person detection. We first describe our baseline architecture – the DROW detector [4], followed by a discussion on different paradigms of aggregating temporal information. In the end, we introduce our proposed DR-SPAAM detector, which in contrast to [4], uses a forward looking approach to aggregate information over time.

#### A. DROW Detector

The DROW detector [12], [4] was the first deep learning-based approach that detects persons from 2D range data. It consists of three stages. First, in order to normalize the appearance across different depths, the raw scans are pre-processed into small per-point windows, which are referred to as *cutouts*. These cutouts are then separately classified by a network as either object or background, and a possible object center is regressed for every cutout. Finally, all regressed object centers (referred to as *votes*) are collected and aggregated to a final set of detections.

During preprocessing, given a scan  $S^t \in \mathbb{R}_{>0}^N$  composed of  $N$  points<sup>2</sup>  $\{s_n^t\}_{n=1}^N$  at time  $t$ ,  $N$  cutouts  $\{C_n^t\}_{n=1}^N$  are generated, each corresponding to a fixed-size window in Euclidean space around the LiDAR point  $s_n^t$ . This is done by computing an angular opening  $a_n^t$  for each point  $s_n^t$  using

$$a_n^t = 2 \cdot \arctan \frac{0.5 \cdot \bar{W}}{s_n^t}, \quad (1)$$

where  $\bar{W}$  is a hyperparameter specifying the cutout width. The points within this angular neighborhood are then re-sampled to a fixed number of  $M$  points and centered by subtracting the distance  $s_n^t$  of the central point. Background and foreground points outside a depth range of  $\pm d$  are clipped away and replaced with a constant value, based on the threshold  $d$ , and finally all values in the cutout  $C_n^t$  are normalized to  $[-1, 1]$ . The now normalized cutouts

<sup>2</sup>A 2D LiDAR scan is composed of range measurements at different angles. For simplicity, we refer to a range measurement  $s$  as a point, omitting the angular component.

$\{C_n^t\}_{n=1}^N$  are passed through a network for classification and regression. During postprocessing, votes are accumulated in a voting grid and a non-maximum suppression step is applied to obtain a set of detections, which are further refined by aggregating the class distributions of votes belonging to a detection. The results of the DROW detector show the benefit of the cutout preprocessing. In particular, it alleviates problems caused by unequal sampling densities at different distances (LiDAR points are sparser at far range), and also allows the DROW detector to work with LiDARs with different angular resolutions without requiring re-training. Furthermore, the clipping operation removes the background information, allowing the network to focus on the neighboring points that are in the same distance region. We refer the reader to [12] for more details.

#### B. Temporal Information Aggregation

Since the measurements obtained from LiDAR sensors are low in information content (especially at far range), some detectors aggregate measurements made at different times to obtain a richer representation of the space, and it has been observed that this temporal aggregation improves the performance of downstream tasks [4], [5]. Many common techniques for accumulating the temporal information follow the so-called backward looking paradigm, where measurements within the past few steps are combined together. Spatial misalignment often exists between these measurements, due to the sensor ego-motion or dynamic objects, and this misalignment has to be corrected based on odometry or point cloud registration.

Similarly, the second version of the DROW detector also accumulates temporal information by looking backward [4]. It computes cutouts on the past  $T$  scans  $\{C_n^{t-T}, \dots, C_n^t\}$  and fuses features  $\{F_n^{t-T}, \dots, F_n^t\}$ , obtained from an intermediate stage of the network, by a simple summation. The fused features are then fed into the later stage of the network for classification and regression. Due to the ego-motion of the sensor, two range measurements  $s_n^t$  and  $s_n^{t-1}$  made at the same angular index  $n$  will not correspond to a single aligned point in the world, and the DROW detector uses robot odometry to correct the misalignment before fusing the features  $F_n^t$  and  $F_n^{t-1}$ . However, odometry alone is not sufficient to compensate for the misalignment caused by dynamic objects in the scene. In the case of persons, this is especially critical, since the LiDAR ray at the same  $n^{th}$  angular index could hit the leg of a person at time  $t-1$ , while passing between the legs and hitting a distant background structure at time  $t$ , resulting in significantly different features. Thus Beyer *et al.* propose to fix the sampling location at which the cutout is centered to the location of the current point  $s_n^t$ . However, this means the cutouts of previous scans need to be recomputed at each time step. With the alignment using odometry and fixed location sampling, the DROW detector combines five scans from the past, resulting in improved detection accuracy compared to using only a single scan.

The performance gain from such a backward looking



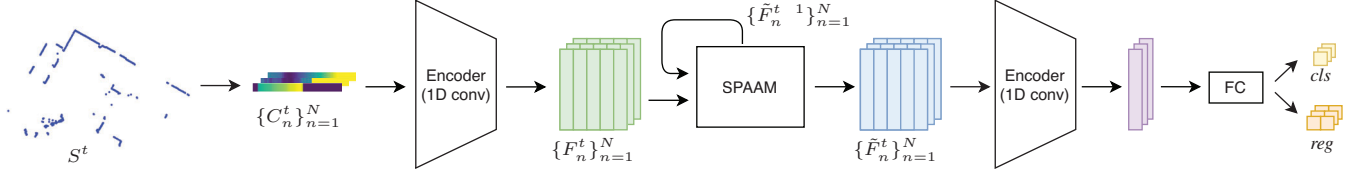


Fig. 2: Overview of the DR-SPAAM architecture. From the current LiDAR scan we create cutouts  $C_n^t$  for every point, from which the network computes intermediate representations  $F_n^t$ . Using our SPAAM module, we aggregate temporal information from previous scans (*c.f.* Fig. 3 and Sec. III-C). Based on the merged representations  $\tilde{F}_n^t$  the network outputs a classification label and predicts a relative object center for each point.

approach comes at the cost of increased computation time. The misalignment between the current measurement and each previous measurement has to be corrected, resulting in a linear increase in computation time with respect to the number of frames in the aggregation window. For the DROW detector, using only five scans already makes the overall detection pipeline too expensive for real-time applications on mobile platforms.

An alternative approach to aggregate temporal information is to follow a forward looking paradigm. Instead of explicitly aligning and combining multiple previous measurements, a forward looking approach simply keeps a representation based on the current measurements and recurrently updates the representation for each new measurement. Ideally, the update step only incurs a small computational overhead. As a result, a forward looking approach aggregates information from the past without the unfavorable runtime scaling behavior with respect to the size of the temporal window.

### C. DR-SPAAM Detector

We propose the Distance Robust Spatial-Attention and Auto-regressive Model (DR-SPAAM), which follows a forward looking paradigm to aggregate temporal information. Instead of computing spatially aligned cutouts on the past scans, we use a similarity-based spatial attention module [40], which allows the network to learn to associate misaligned features from a spatial neighborhood. Additionally, an auto-regressive model is used to update the representation, aggregating information forward through time. Our proposed detector outperforms the DROW detector, while being approximately four times faster. Diagrams of DR-SPAAM and of the proposed spatial attention and auto-regressive model are shown in Fig. 2 and 3 respectively.

Due to the misalignment, features  $F_n^t$  and  $F_n^{t-1}$  computed at two time steps cannot be naively combined. Instead of explicitly modeling the alignment as in [4], we propose to let the network learn to associate features using a similarity-based attention mechanism. For a point  $s_n^t$ , we look at its spatial neighbors  $\{s_{n-w}^{t-1}, \dots, s_{n+w}^{t-1}\}$  at previous time  $t-1$  and compute a pairwise similarity

$$\Omega_{nj} = \psi(F_j^{t-1})^T \cdot \psi(F_n^t) \quad (2)$$

between the features extracted at each previous point  $s_j^{t-1}$  and those from the current point  $s_n^t$ . Here,  $w$  is a parameter defining the size of the neighborhood,  $F_n^t$  is the intermediate feature of  $s_n^t$ , and  $\psi$  is a generic mapping function, realized by a neural network, that maps the feature to an embedding

space. We then use a softmax function to convert the similarity into weighting factors and produce fused features  $\tilde{F}_n^{t-1}$  from the previous frame:

$$\tilde{F}_n^{t-1} = \sum_{j=n-w}^{n+w} \text{softmax}(\Omega_n)_j F_j^{t-1}. \quad (3)$$

This model gives more weight to the points with a higher similarity score, which are more likely to contain information from regions near  $s_n^t$ , while suppressing features from other irrelevant points. The fused feature  $\tilde{F}_n^{t-1}$  from the previous frame can then be combined with the current feature  $F_n^t$  and be used for further processing.

This model, however, only combines information from two consecutive scans. In order to aggregate information from previous scans further back in the past, we propose to combine Eqn. 3 with an auto-regressive approach. We treat the fused features  $\tilde{F}_n^{t-1}$  from time  $t-1$  as a template, and when the new features  $F_n^t$  at time  $t$  become available, we compute an updated template:

$$\tilde{F}_n^t = \alpha F_n^t + (1 - \alpha) \sum_{j=n-w}^{n+w} \text{softmax}(\tilde{\Omega}_n)_j \tilde{F}_j^{t-1}, \quad (4)$$

where  $\alpha \in [0, 1)$  is a parameter that controls the update rate. Here the first term is our update to the stored template, and the second term summarizes the information from the past. Notice that unlike in Eqn. 2, the term  $\tilde{\Omega}_{nj}$  here denotes the similarity between the current feature  $F_n^t$  and the neighboring features  $\{\tilde{F}_{n-w}^{t-1}, \dots, \tilde{F}_{n+w}^{t-1}\}$  from the previous templates, rather than from the previous scan, *i.e.*

$$\tilde{\Omega}_{nj} = \psi(\tilde{F}_j^{t-1})^T \cdot \psi(F_n^t). \quad (5)$$

The updated template  $\tilde{F}_n^t$  is then passed to the later stage of the network for the final classification and offset regression.

Compared to the original DROW detector, our DR-SPAAM detector has a significantly lower computational complexity, requiring neither robot odometry nor the cutout recomputation on previous scans. The auto-regressive model also allows our detector to keep only a single template per angular index, without having to store multiple past scans, while being able to accumulate information within a larger temporal window.

## IV. EVALUATION

We evaluate our methods using the DROW dataset [12], [4], which is recorded in an indoor rehabilitation facility using a SICK S300 scanner. The dataset contains 24,012

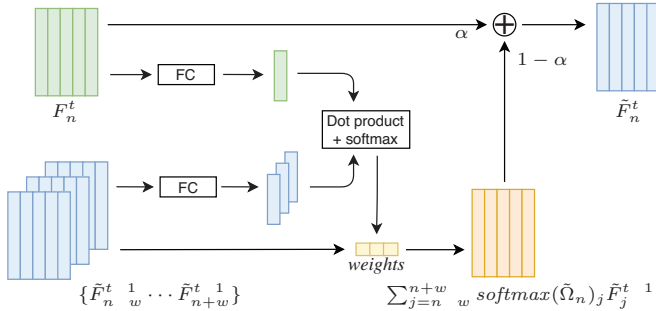


Fig. 3: The SPAAM module associates features from previous scans using similarity-based spatial attention. The similarity-weighted features are then combined with the current features in an auto-regressive fashion.

annotated scans, split into *train* (17,665), *validation* (3,919), and *test* (2,428) sets. The annotation includes the locations of three classes of objects: *wheelchair*, *walker*, and *person*. In this work we specifically focus on detecting persons and ignore the annotations of the other two categories, albeit our method should be general enough to handle other classes with adjusted hyper-parameters.

Following the standard in the object detection community, we use average precision (AP) at different association distances as our main evaluation metric.  $AP_d$  means that a detection is considered as positive if there exists an unmatched ground truth that is within  $d$  m radius of the predicted location. Notice that [12], [4] reported the area under the precision-recall curve (AUC), which is equivalent to average precision by definition. Additionally, we report the peak-F1 score (using 0.5 m association distance), which is the maximum harmonic mean of the different precision and recall values, as well as the equal error rate (EER), the value at which precision and recall are equal.

All our models are trained on the *train* set with a batch size of 8 scans for 40 epochs. For our DR-SPAAM detector, we load 10 frames back into the past during training, since scans that are further back in time are less relevant due to the exponential decay. We use an Adam optimizer, with initial learning rate  $10^{-3}$  with an exponential decay (after each iteration) to  $10^{-6}$  during the complete training. For the classification we use the binary cross entropy loss and for the regression we use the  $L1$ -norm of the regression error. To convert the network output into detections we use the same postprocessing scheme introduced in [4]. We use *Hyperopt* [41] to optimize the hyper-parameters of the voting step by maximizing  $AP_{0.5}$  on the validation set for each model individually. During evaluation, similar as during training, we provide a temporal context of 10 past frames for each test scan. However, our approach readily generalizes to run on complete sequences.

#### A. Quantitative Results

We evaluate our proposed method using the test set and report the average precision, peak-F1 score, and equal error rate of our method at an association threshold of 0.5 m in Table I. As an additional baseline, we also report the

TABLE I: Detection accuracy on the test set with 0.5m association threshold. Note that DR-SPAAM and our re-trained baseline DROW do not use odometry information.

Method	$AP_{0.5}$	peak-F1	EER
ROS leg detector [10]	23.2	41.7	41.0
Arras (re-trained) [9]	47.6	50.3	50.1
Leigh (re-trained) [11]	57.2	64.3	62.3
DROW ( $T = 1$ ) in [4]	59.4	61.5	61.4
DROW ( $T = 5$ ) in [4]	67.0	65.9	64.9
DROW ( $T = 5$ , + odom.) in [4]	68.1	68.1	<b>67.2</b>
DROW ( $T = 1$ ) baseline	66.6	66.1	65.2
DROW ( $T = 5$ ) baseline	67.9	65.1	63.8
DR-AM (w/o spatial attention)	66.3	65.2	64.0
DR-SPA (w/o auto-regression)	68.0	67.0	66.1
DR-SPAAM	<b>70.3</b>	<b>68.5</b>	<b>67.2</b>

performance of two re-trained DROW models, using a single scan and five scans, respectively. Compared to the original implementation, our re-trained models use a smaller cutout window and more sampling points within each cutout, which we selected based on a better performance on the validation set (*c.f.* Sec. IV-C). In order to keep the comparison meaningful, the same cutout parameters are used for both the re-trained DROW baseline and DR-SPAAM, and no odometry information is used. The original DROW score reported by Beyer *et al.* on the person class in [4], as well as the score of the Leigh [11] and Arras detectors [9], re-trained on the DROW dataset, are also included in Table I.

As the results show, DR-SPAAM achieves the highest  $AP_{0.5}$  of 70.3%, which is 2.4% above the baseline model and 2.2% above the original DROW in [4]. Thus, DR-SPAAM establishes a new state-of-the-art, even though it does not use odometry information. By comparing our re-trained DROW models with the original DROW, we can also observe the effectiveness of our proposed adjustment, especially in the single scan case. Fig. 4 shows the precision-recall curves of all different models. Here we can see that DR-SPAAM outperforms all other setups, except for a small region in the high precision regime, where the DROW ( $T=5$ ) baseline scores higher.

Table I also shows the results of an ablation study that highlights the contribution of the different components in our temporal aggregation module. DR-AM corresponds to a network where the auto-regressive model is updated with the new features directly, without using the weighted sum from the spatial attention mechanism. It has a slightly worse performance compared to a single-scan DROW baseline, showing that it is not beneficial to naively combine misaligned features. DR-SPA, on the other hand, only combines the features from the current and the previous scan using spatial attention, without using the accumulated feature template from the auto-regressive model. This two-scan approach already outperforms the five-scan DROW baseline, showing the benefit of using a learning-based approach to incorporate features from the previous measurements. The full model, DR-SPAAM, outperforms the two-scan DR-SPA, showing the benefit of aggregating information over a larger temporal window.

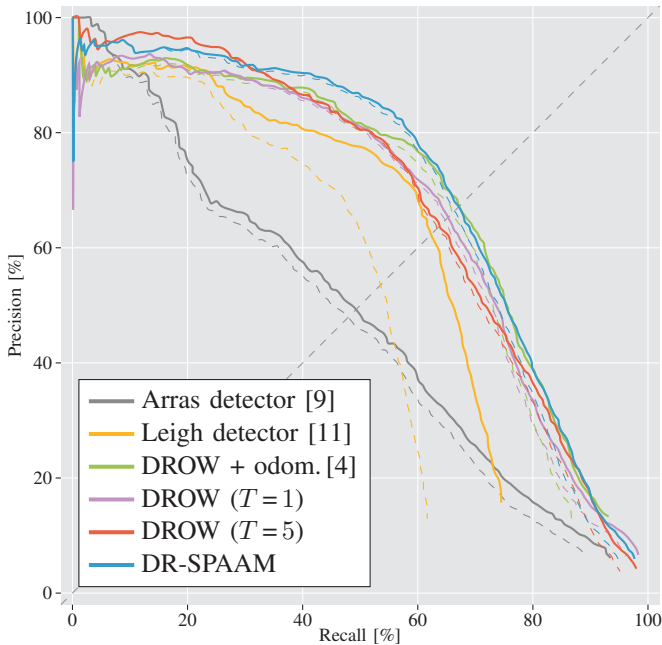


Fig. 4: Precision-recall curves for several baselines and our DR-SPAAM detector, evaluated with association distance 0.5 m (solid) and 0.3 m (dashed).

### B. Inference Time

DR-SPAAM not only achieves better detection, it also has a significantly lower computational complexity, since it eliminates the need to perform an expensive recomputation of past cutouts. We implement all networks in Python and PyTorch without using any inference time acceleration framework (*e.g.*, TensorRT), and profile the run time of different components of the complete pipeline. The timing results for two mobile platforms are reported in Table II. Here we use a laptop equipped with a mobile NVIDIA GeForce RTX 2080 Max-Q GPU and an Intel-i7 9750H CPU, as well as a Jetson AGX Xavier. The table shows that DR-SPAAM has a computation time similar to that of a single-scan DROW method. This is expected, since at each time step only the current scan needs to be processed and the more expensive fusion of DR-SPAAM adds a small overhead. However, even though the methods are similar in speed, the single-scan DROW method has a by far lower detection accuracy. DROW with temporal integration, on the other hand, has a runtime that scales linearly with the number of scans due to the required cutout recomputation on all past scans, and already becomes significantly slower when using five scans. Considering the Jetson AGX platform, a frame rate of 2.6 FPS is too slow for a real-time application, whereas our 9.7 FPS is still well within a usable range. On the laptop, we can in fact run all models faster than real-time, given that the DROW dataset was recorded at a frame rate of roughly 13 FPS. However, DR-SPAAM can run at significantly higher frame rates if needed, providing detections with a lower latency and consuming less power, which is a relevant aspect for mobile platforms.

The cutout generation and the voting are the expensive

TABLE II: Computation time (in milliseconds) and frame rate of different setups on two different mobile platforms.

Method	Laptop (RTX 2080)				Jetson AGX			
	AP <sub>0.5</sub>	cutout	net	vote FPS	cutout	net	vote FPS	
DROW ( $T=1$ )	66.6	7.0	1.4	6.1 68.6	63.3	4.8	29.3 10.4	
DROW ( $T=5$ )	67.9	34.3	1.5	19.2 18.2	306.3	5.1	78.1 2.6	
DR-SPAAM	<b>70.3</b>	7.0	2.0	7.7 59.8	62.0	6.9	33.6 9.7	
DR-SPAAM*	<b>71.8</b>	1.1	1.9	8.5 87.2	4.2	7.7	32.4 22.6	

steps in the whole pipeline. To further increase the frame rate, we resort to a faster implementation of the cutout generation and the resulting model, DR-SPAAM\*, runs at 87.2 FPS on a laptop with a dedicated GPU, or at 22.6 FPS on a Jetson AGX, well-beyond the requirement for many real-time applications. The new implementation increases the network performance thanks to the improved numerical precision.

### C. Hyperparameter Selection

**The cutout** operation is parameterized by the width ( $\bar{W}$ ) and depth ( $D$ ) of the window, as well as the number of points ( $N$ ) used for re-sampling. In [4] a larger ( $1.66\text{ m} \times 2.0\text{ m}$ ) sized window with 48 points was used in order to cope with the bigger walking-aid classes. Since we are only concerned with detecting persons, we propose to use a smaller sized window that tightly fits the footprint of a person, thus reducing distracting information from the surroundings. We conduct an experiment on the size of the cutout window on the validation set using a DROW network with a single scan. The results are shown in Table III. Based on these results, we set our cutout window to ( $1.0\text{ m} \times 1.0\text{ m}$ ) with 56 points and use these cutout parameters for all models we train.

One can observe that the scores on the validation set are significantly lower than those on the test set. The original DROW dataset is created for detecting three classes: *wheelchair*, *walker*, and *person*. In this work, we have only kept the annotations for the *person* class, and we observed that the validation set happens to contain more person annotations at farther distances. Even though we use a distance robust preprocessing, at farther distances information becomes so sparse that the detection will always become less reliable, thus rendering the validation set more challenging.

**The Spatial-Attention and Auto-Regressive Model** are parameterized by the update rate  $\alpha$  and the size of the search window  $W$ . Although the meaning of these two parameters is intuitively clear, it is not a trivial task to select the proper combination. We train multiple DR-SPAAM networks with different combinations of  $\alpha$  and  $W$ . Based on the validation set results (Table IV), we choose to use  $W=11$  and  $\alpha=0.5$  for our final model. No clear pattern can be seen in these results, neither larger search windows nor specific update rates consistently work better. A more thorough search through the parameter space could potentially result in models that perform even better.

TABLE III: Validation set scores of DROW ( $T=1$ ) detectors with different cutout parameters.

$\overline{W}$	$D$	$N$	AP <sub>0.3</sub>	AP <sub>0.5</sub>	peak-F1	EER
1.66	2.0	48	41.9	43.0	48.1	47.6
1.66	1.0	48	42.6	43.4	49.2	48.6
1.0	2.0	48	43.6	44.8	50.7	50.4
1.0	1.0	48	44.0	45.0	50.3	50.2
<hr/>						
1.0	1.0	32	42.0	43.0	49.1	48.8
1.0	1.0	40	43.1	44.1	50.0	49.6
1.0	1.0	48	44.0	45.0	50.3	50.2
1.0	1.0	56	<b>45.1</b>	<b>46.3</b>	<b>50.9</b>	<b>50.8</b>
1.0	1.0	64	43.8	45.1	50.7	50.4

TABLE IV: Validation set scores of DR-SPAAM with different window sizes and update rates.

$W$	$\alpha$	AP <sub>0.3</sub>	AP <sub>0.5</sub>	peak-F1	EER
7	0.3	45.0	46.2	52.5	52.5
7	0.5	49.5	50.9	54.6	53.6
7	0.8	46.8	48.3	54.1	54.0
<hr/>					
11	0.3	51.5	53.0	56.8	56.4
11	0.5	<b>52.7</b>	<b>53.9</b>	<b>57.3</b>	<b>57.3</b>
11	0.8	47.4	48.7	53.6	53.2
<hr/>					
15	0.3	51.5	52.8	56.1	55.3
15	0.5	50.7	52.1	55.0	54.7
15	0.8	47.0	48.2	53.1	53.0

#### D. Sampling Rate

Different LiDARs often have different sampling rates. Since a detection network is likely to be deployed on different LiDAR sensors, its robustness against varying sensor specifications should be examined. We take two networks, DROW ( $T=5$ ) and DR-SPAAM, and evaluate them on the test set using temporally sub-sampled sequences, simulating different sampling rates. Table V reports the detection accuracies at different temporal strides (a stride of  $n$  means keeping only every  $n^{\text{th}}$  scan).

The evaluation results show that DR-SPAAM is very robust against changing sampling rate. Even at a five times lower scanning frequency (roughly 2 Hz), the AP<sub>0.3</sub> only reduces by 2.1%. This result shows the benefit of a learned spatial attention module, which combines information based on appearance similarity without relying on a fixed temporal context window. Hence, DR-SPAAM can be deployed on LiDARs with a wide range of sampling rates, or operate with a reduced sampling rate if the computation capacity is limited. On the other hand, the DROW detector performance degrades rapidly with increased temporal stride. Larger time differences between consecutive scans lead to greater motion induced misalignments, which in turn hurt the accuracy of the DROW detector.

#### E. Temporal Association

During the temporal aggregation step, DR-SPAAM computes the similarity between the aggregated template and the latest scan features (Eqn. 5). This similarity can be further exploited for associating points across different scans. Here we provide a preliminary example. We take 200 consecutive frames (roughly spanning 15 s) from a sequence and detect

TABLE V: Test set results with different temporal strides.

Stride	DROW ( $T=5$ )				DR-SPAAM			
	AP <sub>0.3</sub>	AP <sub>0.5</sub>	p-F1	EER	AP <sub>0.3</sub>	AP <sub>0.5</sub>	p-F1	EER
1	<b>66.6</b>	<b>67.9</b>	<b>65.1</b>	<b>63.8</b>	68.5	70.3	68.5	67.2
2	59.3	60.5	60.1	59.3	69.3	70.8	<b>68.8</b>	<b>67.6</b>
3	54.3	55.8	56.8	56.7	<b>69.4</b>	<b>70.9</b>	68.1	66.5
4	53.6	55.1	56.0	55.7	67.7	69.1	66.4	64.9
5	51.5	53.4	54.6	54.3	66.4	67.7	65.5	64.5

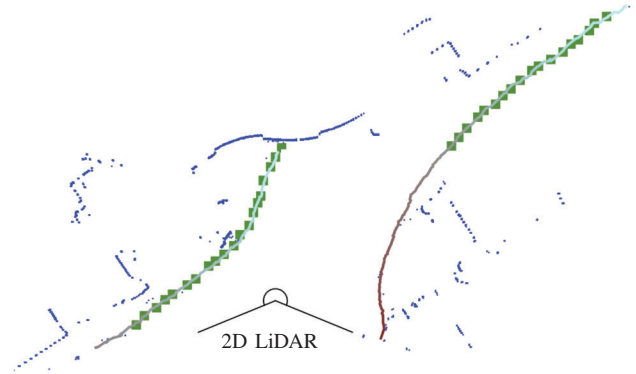


Fig. 5: Tracklets generated by DR-SPAAM using 200 consecutive scans. The blue points are the overlaid scans, and the green squares are the ground truth annotations. Notice that for clarity, we omit the points that have been classified as persons from plotting. The colored lines are the tracklets, and the coloring encodes the time of detections. The sequence is taken from the training set, since the validation set does not have any sequence recorded using a stationary LiDAR, which is needed for plotting overlaid scans. Nevertheless, most of the scans have not been annotated (shown as the missing annotations along the tracklets) and have not been exposed to the network during training.

persons using DR-SPAAM for each individual frame. For each detection  $D_i^t$ , we find its corresponding points in the previous scan, simply by selecting the ones with highest similarity score. If these corresponding points were grouped into a detection  $D_j^{t-1}$ , and if the distance between the two detections are smaller than a threshold (0.5 m), we group both detections into a tracklet. Otherwise a new tracklet is started using  $D_i^t$ . After 200 frames, we compute the confidence of each tracklet as the mean of all its detections. In Fig. 5 we plot all tracklets that have a confidence greater than 0.35 and that are composed of at least five detections. The trajectories of persons in the scene are clearly visible. These associations can provide extra information for tracking algorithms. Their full potential is yet to be explored in future research. The velocity and movement direction of the persons can also be derived from the associated detection pairs, and this information can be helpful for motion planning.

## V. CONCLUSION

We propose the DR-SPAAM person detector that combines the distance robust detection scheme of the DROW detector with a powerful spatial attention and auto-regressive



temporal integration model. The spatial attention is able to associate misaligned features from different frames using their appearance similarity, while the auto-regressive model aggregates temporal information forward through time. Compared to the previous state-of-the-art approaches, DR-SPAAM achieves higher detection accuracy, while being significantly faster and able to run in real-time even on low-powered mobile platforms. Experiments show that DR-SPAAM generalizes well to LiDARs with different temporal sampling rates, and with our provided code and ROS node, we expect that our model will be useful for many robotic applications.

**Acknowledgements:** We thank Francis Engelmann for his valuable feedback. This project was funded, in parts, by the EU H2020 project "CROWDBOT" (779942) and the BMBF project "FRAME" (16SV7830). Most experiments were performed on the RWTH Aachen University CLAIX 2018 GPU Cluster (rwth0485).

#### REFERENCES

- [1] R. B. Girshick, "Fast R-CNN," in *International Conference on Computer Vision*, 2015.
- [2] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, "Mask R-CNN," in *International Conference on Computer Vision*, 2017.
- [3] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," *arXiv:1804.02767*, 2018.
- [4] L. Beyer, A. Hermans, T. Linder, K. O. Arras, and B. Leibe, "Deep Person Detection in 2D Range Data," *IEEE Robotics and Automation Letters (RA-L)*, vol. 3, no. 3, pp. 2726–2733, 2018.
- [5] B. Zhu, Z. Jiang, X. Zhou, Z. Li, and G. Yu, "Class-balanced Grouping and Sampling for Point Cloud 3D Object Detection," *arXiv:1908.09492*, 2019.
- [6] S. Tripathi, Z. C. Lipton, S. J. Belongie, and T. Q. Nguyen, "Context matters: Refining object detection in video with recurrent neural networks," *arXiv:1607.04648*, 2016.
- [7] Y. Lu, C. Lu, and C.-K. Tang, "Online Video Object Detection Using Association LSTM," in *International Conference on Computer Vision*, 2017.
- [8] F. Xiao and Y. J. Lee, "Video Object Detection with an Aligned Spatial-Temporal Memory," in *European Conference on Computer Vision*, 2017.
- [9] K. O. Arras, Ó. M. Mozos, and W. Burgard, "Using Boosted Features for the detection of People in 2D Range Data," in *IEEE International Conference on Robotics and Automation*, 2007.
- [10] C. Pantofaru, "ROS leg\_detector package." [https://wiki.ros.org/leg\\_detector](https://wiki.ros.org/leg_detector), 2010. Accessed 2018-02-22.
- [11] A. Leigh, J. Pineau, N. Olmedo, and H. Zhang, "Person tracking and Following with 2D Laser Scanners," in *IEEE International Conference on Robotics and Automation*, 2015.
- [12] L. Beyer\*, A. Hermans\*, and B. Leibe, "DROW: Real-Time Deep Learning based Wheelchair Detection in 2D Range Data," *IEEE Robotics and Automation Letters (RA-L)*, vol. 2, no. 2, pp. 585–592, 2016.
- [13] P. Ondruska, J. Dequaire, D. Zeng Wang, and I. Posner, "End-to-End Tracking and Semantic Segmentation Using Recurrent Neural Networks," in *Robotics: Science and Systems Conference, Workshop on Limits and Potentials of Deep Learning in Robotics*, 2016.
- [14] J. Dequaire, P. Ondruška, D. Rao, D. Wang, and I. Posner, "Deep tracking in the wild: End-to-end tracking using recurrent neural networks," *International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 492–512, 2018.
- [15] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, "Multi-view 3D Object Detection Network for Autonomous Driving," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [16] C. R. Qi, W. Liu, C. Wu, H. Su, and L. J. Guibas, "Frustum PointNets for 3D Object Detection from RGB-D Data," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [17] J. Ku, M. Mozifian, J. Lee, A. Harakeh, and S. L. Waslander, "Joint 3D Proposal Generation and Object Detection from View Aggregation," in *International Conference on Intelligent Robots and Systems*, 2017.
- [18] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "PointPillars: Fast Encoders for Object Detection From Point Clouds," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [19] Y. Zhou and O. Tuzel, "VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [20] B. L. Yan Yan, Yuxing Mao, "SECOND: Sparsely Embedded Convolutional Detection," *Sensors*, 2018.
- [21] B. Graham, M. Engelcke, and L. van der Maaten, "3D Semantic Segmentation with Submanifold Sparse Convolutional Networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [22] C. Choy, J. Gwak, and S. Savarese, "4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [23] S. Shi, X. Wang, and H. Li, "PointRCNN: 3D Object Proposal Generation and Detection From Point Cloud," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [24] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [25] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space," in *Neural Information Processing Systems*, 2017.
- [26] F. Engelmann, T. Kontogianni, and B. Leibe, "Dilated Point Convolutions: On the Receptive Field Size of Point Convolutions on 3D Point Clouds," in *IEEE International Conference on Robotics and Automation*, 2020.
- [27] 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*, vol. 38, no. 5, 2019.
- [28] Y. Chen, S. Liu, X. Shen, and J. Jia, "Fast Point R-CNN," *arXiv:1908.02990*, 2019.
- [29] 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*, 2019.
- [30] 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," *arXiv:1908.02990*, 2019.
- [31] C. R. Qi, O. Litany, K. He, and L. J. Guibas, "Deep Hough Voting for 3D Object Detection in Point Clouds," in *International Conference on Computer Vision*, 2019.
- [32] F. Engelmann, M. Bokeloh, A. Fathi, B. Leibe, and M. Nießner, "3D-MPA: Multi Proposal Aggregation for 3D Semantic Instance Segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [33] K. Kang, W. Ouyang, H. Li, and X. Wang, "Object Detection from Video Tubelets with Convolutional Neural Networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [34] W. Han, P. Khorrami, T. L. Paine, P. Ramachandran, M. Babaeizadeh, H. Shi, J. Li, S. Yan, and T. S. Huang, "Seq-NMS for Video Object Detection," *arXiv:1602.08465*, 2016.
- [35] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Detect to Track and Track to Detect," in *International Conference on Computer Vision*, 2017.
- [36] X. Zhu, J. Dai, L. Yuan, and Y. Wei, "Towards High Performance Video Object Detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [37] X. Zhu, Y. Wang, J. Dai, L. Yuan, and Y. Wei, "Flow-Guided Feature Aggregation for Video Object Detection," in *International Conference on Computer Vision*, 2017.
- [38] S. Wang, Y. Zhou, J. Yan, and Z. Deng, "Fully Motion-Aware Network for Video Object Detection," in *European Conference on Computer Vision*, 2018.
- [39] H. Wu, Y. Chen, N. Wang, and Z. Zhang, "Sequence Level Semantics Aggregation for Video Object Detection," in *International Conference on Computer Vision*, 2019.
- [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 *Neural Information Processing Systems*, 2017.
- [41] J. Bergstra, D. Yamins, and D. D. Cox, "Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures," in *International Conference on Machine Learning*, 2013.