

MIT Open Access Articles

*Searching Efficient 3D Architectures
with Sparse Point-Voxel Convolution*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Tang, Haotian, Liu, Zhijian, Zhao, Shengyu, Lin, Yujun, Lin, Ji et al. 2020. "Searching Efficient 3D Architectures with Sparse Point-Voxel Convolution." Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 12373.

As Published: 10.1007/978-3-030-58604-1_41

Publisher: Springer International Publishing

Persistent URL: <https://hdl.handle.net/1721.1/143668>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Searching Efficient 3D Architectures with Sparse Point-Voxel Convolution

Haotian Tang^{1,*}, Zhijian Liu^{1,*},
Shengyu Zhao^{1,2}, Yujun Lin¹, Ji Lin¹, Hanrui Wang¹, and Song Han¹

¹ Massachusetts Institute of Technology

² IIIS, Tsinghua University

Abstract. Self-driving cars need to understand 3D scenes efficiently and accurately in order to drive safely. Given the limited hardware resources, existing 3D perception models are not able to recognize small instances (*e.g.*, pedestrians, cyclists) very well due to the low-resolution voxelization and aggressive downsampling. To this end, we propose *Sparse Point-Voxel Convolution (SPVConv)*, a lightweight 3D module that equips the vanilla Sparse Convolution with the high-resolution point-based branch. With negligible overhead, this point-based branch is able to preserve the fine details even from large outdoor scenes. To explore the spectrum of efficient 3D models, we first define a flexible architecture design space based on SPVConv, and we then present *3D Neural Architecture Search (3D-NAS)* to search the optimal network architecture over this diverse design space efficiently and effectively. Experimental results validate that the resulting SPVNAS model is fast and accurate: it outperforms the state-of-the-art MinkowskiNet by 3.3%, ranking **1st** on the competitive SemanticKITTI leaderboard upon publication. It also achieves **8×** computation reduction and **3×** measured speedup over MinkowskiNet still with higher accuracy. Finally, we transfer our method to 3D object detection, and it achieves consistent improvements over the one-stage detection baseline on KITTI.

1 Introduction

3D deep learning has received increased attention thanks to its wide applications: *e.g.*, it has been used in LiDAR perception that serves as the eyes of autonomous driving systems to understand the semantics of outdoor scenes. As the safety of the passenger is the top priority of the self-driving cars, 3D perception models are required to achieve high accuracy and low latency at the same time. However, the hardware resources on the self-driving cars are tightly constrained by the form factor (since we do not want a whole trunk of workstations) and heat dissipation. Thus, it is crucial to design efficient and effective 3D neural network models with limited computation resources (*e.g.*, memory).

Researchers have primarily exploited two 3D data representations: point cloud and rasterized voxel grids. As analyzed in Liu *et al.* [32], point-based methods [40, 43, 28] waste up to 90% of their runtime on structuring the irregular data, not on

* indicates equal contributions; order determined by a coin toss.




Fig. 1. Small instances (e.g., pedestrians and cyclists) are hard to be recognized at a low resolution (due to the coarse voxelization or the aggressive downsampling).

the actual feature extraction. On the other hand, voxel-based methods usually suffer from severe information loss: *i.e.*, the resolution of dense voxels [36, 32] is strictly constrained by the memory; the sparse voxels [14, 9] require aggressive downsampling to achieve larger receptive field, leading to low resolution at deeper layers. With low resolution (see Figure 1), multiple points or even multiple small objects might be merged into one grid and become indistinguishable. In this case, small instances (e.g., pedestrians and cyclists) are at a disadvantage compared with large objects (e.g., cars). Therefore, the effectiveness of previous 3D modules is discounted when hardware resources are limited and the resolution is low.

To tackle these problems, we propose a novel 3D module, *Sparse Point-Voxel Convolution (SPVConv)* that introduces a low-cost high-resolution point-based branch to the vanilla Sparse Convolution, which helps to capture the fine details. On top of SPVConv, we further present *3D Neural Architecture Search (3D-NAS)* to search an efficient 3D model. We incorporate fine-grained channel numbers into the search space to increase the diversity and introduce the progressive depth shrinking to accelerate the training. Experimental results validate that our model is fast and accurate: it outperforms MinkowskiNet by 3.3% in mIoU with lower latency. It also achieves $8\times$ computation reduction and $3\times$ measured speedup over MinkowskiNet, while providing higher accuracy. We further transfer our method to KITTI for 3D object detection, and it achieves consistent improvements over the previous one-stage detection baseline.

The contribution of this paper has three aspects:

1. We design a lightweight 3D module, SPVConv, that boosts the performance on small objects, which used to be challenging under limited hardware resource.
2. We introduce the first AutoML framework for 3D scene understanding, 3D-NAS, that offers the best 3D model given a specific resource constraint.
3. Our method outperforms all previous methods with a large margin and ranks 1st on the competitive SemanticKITTI leaderboard* upon publication. It can also be transferred to object detection and achieves consistent improvements.

* <http://semantic-kitti.org/tasks.html#semseg>

2 Related Work

2.1 3D Perception Models

Increased attention has been paid to 3D deep learning, which is important for LiDAR perception in autonomous driving. Early research [7, 36, 42, 64, 77] relied on the volumetric representation and vanilla 3D convolution to process the 3D data. Due to the *sparse* nature of 3D representation, the *dense* volumetric representation is inherently inefficient, and it also inevitably introduces information loss. Therefore, researchers have proposed to directly learn on the 3D point cloud representation using the symmetric function [40]. To improve the neighborhood modeling capability, researchers have defined point-based convolutions on either geometric [28, 35, 43, 52, 55, 56, 67, 69] or semantic [63] neighborhood. There are also 3D models tailored for specific tasks such as detection [38, 39, 41, 47–49, 70, 73] and instance segmentation [16, 21, 23, 71] built upon these modules.

Recently, a few researchers started to pay attention to the efficiency aspect of 3D deep learning. Riegler *et al.* [45], Wang *et al.* [60, 61] and Lei *et al.* [25] proposed to reduce the memory footprint of the volumetric representation using octrees where areas with lower density occupy fewer voxel grids. Liu *et al.* [32] analyzed the bottlenecks of point-based and voxel-based methods, and proposed Point-Voxel Convolution. Graham *et al.* [14] and Choy *et al.* [9] proposed Sparse Convolution to accelerate the volumetric convolution by keeping the activation sparse and skipping the computations in the inactive regions.

2.2 Neural Architecture Search

To alleviate the burden of manually designing neural networks [18, 46, 33, 75, 20], researchers have introduced neural architecture search (NAS) to automatically architect the neural network with high accuracy using reinforcement learning [79, 80] and evolutionary search [29]. A new wave of research started to design efficient models with neural architecture search [53, 66, 54] for mobile deployment. However, conventional frameworks require high computation cost and considerable carbon footprint [51]. To tackle these, researchers have proposed different techniques to reduce the search cost, including differentiable architecture search [30], path-level binarization [6], single-path one-shot sampling [15, 8, 4], and weight sharing [50, 4, 57]. Besides, neural architecture search has also been used in compressing and accelerating neural networks, including pruning [17, 31, 5, 27] and quantization [58, 15, 59, 62]. Most of these methods are tailored for 2D visual recognition, which has many well-defined search spaces [44]. Lately, researchers have applied neural architecture search to 3D medical image segmentation [78, 22, 72, 2, 65, 74] and 3D shape classification [34, 26]. However, they are not directly applicable to 3D scene understanding since 3D medical data are still in the similar format as 2D images (which are entirely different from 3D scenes), and 3D objects are of much smaller scales than 3D scenes (which makes them less sensitive to the resolution).

3 SPVConv: Designing Effective 3D Modules

We first revisit two recent 3D modules: Point-Voxel Convolution [32] and Sparse Convolution [9] and analyze their bottlenecks. We observe that both of them suffer

	Input	Voxel Size (m)	Latency (ms)	Mean IoU
PVCConv [32]	Sliding Window	0.05	35640	–
	Entire Scene	0.78	146	39.0
SPVConv (Ours)	Entire Scene	0.05	85	58.8

Table 1. Point-Voxel Convolution [32] is not suitable for large 3D scenes. If processing with sliding windows, the large latency is not affordable for real-time applications. If taking the whole scene, the resolution is too coarse to capture useful information.

from information loss (caused by coarse voxelization or aggressive downsampling) when the memory is constrained. To this end, we introduce *Sparse Point-Voxel Convolution (SPVConv)*, to effectively process the large 3D scene (as in Figure 2).

3.1 Point-Voxel Convolution: Coarse Voxelization

Liu *et al.* [32] proposed Point-Voxel Convolution that represents the 3D input data in points to reduce the memory consumption, while performing the convolutions in voxels to reduce the irregular data access and improve the locality. Specifically, its point-based branch transforms each point individually, and its voxel-based branch convolves over the voxelized input from the point-based branch.

PVCNN (which is built upon Point-Voxel Convolution) can afford at most 128^3 voxels in its voxel-based branch on a single GPU (with 12 GB of memory). Given a large outdoor scene (with size of $100\text{m} \times 100\text{m} \times 10\text{m}$), each voxel grid will correspond to a fairly large area (with size of $0.8\text{m} \times 0.8\text{m} \times 0.1\text{m}$). In this case, the small instances (*e.g.*, pedestrians) will only occupy a few voxel grids (see Figure 1). From such few points, PVCNN can hardly learn any useful information from the voxel-based branch, leading to a relatively low performance (see Table 1). Alternatively, we can process the large 3D scene piece by piece so that each sliding window is of smaller scale. In order to preserve the fine-grained information (*i.e.*, voxel size is smaller than 0.05m), we have to run PVCNN once for each of the 244 sliding windows. This takes 35 seconds to process a single scene, which is not affordable for most real-time applications (*e.g.*, autonomous driving).

3.2 Sparse Convolution: Aggressive Downsampling

Volumetric convolution has always been considered inefficient and prohibitive to be scaled up. Lately, researchers proposed Sparse Convolution [14, 9] that skips the non-activated regions to significantly reduce the memory consumption. More specifically, it first finds all active synapses (denoted as *kernel map*) between the input and output points; it then performs the convolution based on this kernel map. To keep the activation sparse, it only considers these output points that also belong to the input. We refer the readers to Choy *et al.* [9] for more details.

As such, Sparse Convolution can afford a much higher resolution than the vanilla volumetric convolution. However, the network cannot be very deep due to the limited computation resource. As a result, the network has to downsample very aggressively in order to achieve a sufficiently large receptive field, which is very lossy. For instance, the state-of-the-art MinkowskiNet [9] gradually applies




Fig. 2. Overview of Sparse Point-Voxel Convolution (SPVConv): it equips the sparse voxel-based branch with a lightweight, high-resolution point-based branch which can capture fine details in large scenes.

four downsampling layers to the input point cloud, after which, the voxel size will become $0.05 \times 2^4 = 0.8\text{m}$. Similar to Point-Voxel Convolution, this resolution is too coarse to capture the small instances (see Figure 4).

3.3 Solution: Sparse Point-Voxel Convolution

In order to overcome the limitations of both modules, we present Sparse Point-Voxel Convolution (SPVConv) in Figure 2: the point-based branch always keeps the high-resolution representation, and the sparse voxel-based branch applies Sparse Convolution to model across different receptive field size. Two branches communicate at a negligible cost through sparse voxelization and devoxelization.

Data Representation. Our Sparse Point-Voxel Convolution operates on:

- sparse voxelized tensor $\mathbf{S} = (\{\mathbf{p}_m^s, \mathbf{f}_m^s\}, v)$, where $\mathbf{p}_m^s = (x_m^s, y_m^s, z_m^s)$ is the 3D coordinate and \mathbf{f}_m^s is the feature vector of the m^{th} nonzero grid, and v is the voxel size (*i.e.*, side length) for one grid in the current layer;
- point cloud tensor $\mathbf{T} = \{\mathbf{p}_k^t, \mathbf{f}_k^t\}$, where $\mathbf{p}_k = (x_k, y_k, z_k)$ is the 3D coordinate and \mathbf{f}_k is feature vector of k^{th} point.

Sparse Voxelization. In the upper sparse voxel-based branch, we first transform the high-resolution point cloud tensor \mathbf{T} to a sparse tensor \mathbf{S} :

$$\hat{\mathbf{p}}_k^t = (\hat{x}_k^t, \hat{y}_k^t, \hat{z}_k^t) = (\text{floor}(x_k^t/v), \text{floor}(y_k^t/v), \text{floor}(z_k^t/v)), \quad (1)$$

$$\mathbf{f}_m^s = \frac{1}{N_m} \sum_{k=1}^n \mathbb{I}[\hat{x}_k^t = x_m^s, \hat{y}_k^t = y_m^s, \hat{z}_k^t = z_m^s] \cdot \mathbf{f}_k^t, \quad (2)$$

where $\mathbb{I}[\cdot]$ is the binary indicator of whether $\hat{\mathbf{p}}_k^t$ belongs to the voxel grid \mathbf{p}_m^s , and N_m is the normalization factor (*i.e.*, the number of points that fall into the m^{th} nonzero voxel grid). Such formulation, however, requires $\mathcal{O}(mn)$ complexity where $m = |\mathbf{S}|$ and $n = |\mathbf{T}|$. With typical values of m, n at the order of 10^5 , the naive implementation is impractical for real-time applications.

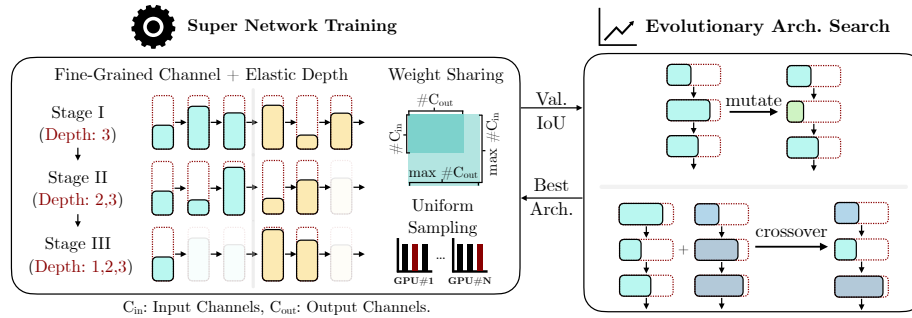


Fig. 3. Overview of 3D Neural Architecture Search (3D-NAS): we first train a super network composed of multiple SPVConv’s, supporting fine-grained channel numbers and elastic network depths. Then, we perform the evolutionary architecture search to obtain the best candidate model under a given computation constraint.

To this end, we propose to use the GPU hash table to accelerate the sparse voxelization and devoxelization. Specifically, we first construct a hash table for all activated points in the sparse voxelized tensor \mathcal{S} , which can be completed in $\mathcal{O}(n)$ time. After that, we iterate over all points in \mathcal{T} , and for each point, we use its voxelized coordinate as the key to query the corresponding index in the sparse voxelized tensor. As the lookup over the hash table requires $\mathcal{O}(1)$ time [37], this query step will in total take $\mathcal{O}(m)$ time. Therefore, the total time of coordinate indexing will be reduced from $\mathcal{O}(mn)$ to $\mathcal{O}(m + n)$.

Feature Aggregation. We then perform the neighborhood feature aggregation on the sparse voxelized tensor using a sequence of residual Sparse Convolution blocks [9]. We parallelize the kernel map operation in Sparse Convolution on GPU with the same hash table implementation as in sparse voxelization, which offers $1.3\times$ speedup over the implementation from Choy *et al.* [9]. Note that both our method and the baseline have been upgraded to this accelerated implementation.

Sparse Devoxelization. With the aggregated features (which are in the form of sparse tensors), we transform them back to the point-based representation so that the information from both branches can be fused together. Similar to Liu *et al.* [32], we choose to interpolate each point’s feature with its 8 neighbor voxel grids using trilinear interpolation instead of the naïve nearest interpolation.

Point Transformation and Feature Fusion. In the lower point-based branch, we directly apply an MLP on each point to extract individual point features. We then fuse the outputs of two branches with an addition to combine the complementary information provided. Compared with the vanilla Sparse Convolution, MLP layers only cost little computation overhead (4% in terms of #MACs) but introduce important fine details into the information flow (see Figure 5).

4 3D-NAS: Searching Efficient 3D Architectures

Even with our module, designing an efficient neural network is still challenging. We need to carefully adjust the network architecture (*e.g.*, channel numbers and kernel sizes of all layers) to meet the constraints for real-world applications (*e.g.*, latency, energy, and accuracy). To this end, we introduce *3D Neural Architecture Search (3D-NAS)*, to automatically design efficient 3D models (as in Figure 3).

4.1 Design Space

The performance of neural architecture search is greatly impacted by the design space quality. In our search space, we incorporate fine-grained channel numbers and elastic network depths; however, we do not support different kernel sizes.

Fine-grained Channel Numbers. The computation cost increases quadratically with the number of channels; therefore, the channel number selection has a large influence on the network efficiency. Most existing neural architecture frameworks [6] only support the coarse-grained channel number selection: *e.g.*, searching the expansion ratio of the ResNet/MobileNet blocks over a few (2-3) choices. In this case, only intermediate channel numbers of the blocks can be changed; while the input and output channel numbers will still remain the same. Empirically, we observe that this limits the variety of the search space. To this end, we enlarge the search space by allowing all channel numbers to be selected from a large collection of choices (with size of $O(n)$). This fine-grained channel number selection largely increase the number of candidates for each block: *e.g.*, from constant (2-3) to $\mathcal{O}(n^2)$ for a block with two consecutive convolutions.

Elastic Network Depths. We support different network depth in our design space. For 3D CNNs, reducing the channel numbers alone cannot achieve significant measured speedup, which is drastically different from 2D CNNs. For example, by shrinking all channel numbers in MinkowskiNet [9] by $4\times$ and $8\times$, the number of MACs will be reduced to 7.5 G and 1.9 G, respectively. However, although the number of MACs is drastically reduced, their measured latency on the GPU is very similar: 105 ms and 96 ms (measured on a single GTX1080Ti GPU). This suggests that scaling down the number of channels alone cannot offer us with very efficient models, even though the number of MACs is very small. This might be because 3D modules are usually more memory-bounded than 2D modules; the number of MACs decreases quadratically with channel number, while memory decreases linearly. Motivated by this, we choose to incorporate the elastic network depth into our design space so that these layers with very small computation (and large memory cost) can be removed and merged into their neighboring layers.

Small Kernel Matters. Kernel sizes are usually included into the search space of 2D CNNs. This is because a single convolution with larger kernel size can be more efficient than multiple convolutions with smaller kernel sizes on GPUs. However, it is not the case for the 3D CNNs. From the computation perspective, a single 2D convolution with kernel size of 5 requires only $1.4\times$ more MACs than two 2D convolutions with kernel sizes of 3; while a single 3D convolution with

kernel size of 5 requires $2.3\times$ more MACs than two 3D convolutions with kernel sizes of 3 (if applied to dense voxel grids). This larger computation cost makes it less suitable to use large kernel sizes in 3D CNNs. Furthermore, the computation overhead of 3D modules is also related to the kernel sizes. For example, Sparse Convolution [14, 9] requires $\mathcal{O}(k^3n)$ time to build the kernel map, where k is the kernel size and n is the number of points, which indicates that its cost grows cubically with respect to the kernel size. Based on these reasons, we decide to keep the kernel size of all convolutions to be 3 and do not allow the kernel size to change in our search space. Even with the small kernel size, we can still achieve a large receptive field by changing the network depth, which can achieve the same effect as changing the kernel size.

4.2 Training Paradigm

Searching over a fine-grained design space is very challenging as it is impossible to train every sampled candidate network from scratch [53]. Motivated by Guo *et al.* [15], we incorporate all candidate networks into a single super network, and after training this super network once, each candidate network can be directly extracted with inherited weights. As such, the total training cost can be reduced from $\mathcal{O}(n)$ to $\mathcal{O}(1)$, where n is the number of candidate networks.

Uniform Sampling. At each training iteration, we randomly sample a candidate network from the super network: randomly select the channel number for each layer, and then randomly select the network depth (*i.e.* the number of blocks to be used) for each stage. The total number of candidate networks to be sampled during training is very limited; therefore, we choose to sample different candidate networks on different GPUs and average their gradients at each step so that more candidate networks can be sampled. For 3D, this is more critical because the 3D datasets usually contain fewer training samples than the 2D datasets: *e.g.* 20K on SemanticKITTI [3] *vs.* 1M on ImageNet [11].

Weight Sharing. As the number of candidate networks is enormous, every candidate network will only be optimized for a small fraction of the total schedule. Therefore, uniform sampling alone is not enough to train all candidate networks sufficiently (*i.e.*, achieving the same level of performance as being trained from scratch). To tackle this, we adopt the weight sharing technique so that every candidate network can be optimized at each iteration even if it is not sampled. Specifically, given the input channel number C_{in} and output channel number C_{out} of each convolution layer, we simply index the first C_{in} and C_{out} channels from the weight tensor accordingly to perform the convolution [15]. For each batch normalization layer, we similarly crop the first c channels from the weight tensor based on the sampled channel number c . Finally, with the sampled depth d for each stage, we choose to keep the first d layers, instead of randomly sampling d of them. This ensures that each layer will always correspond to the same depth index within the stage.

Progressive Depth Shrinking. Suppose we have n stages, each of which has m different depth choices from 1 to m . If we sample the depth d_k for each stage k randomly, the expected total depth of the network will be $\mathbb{E}[d] = \sum_{k=1}^n \mathbb{E}[d_k] =$

$n \times (m + 1)/2$, which is much smaller than the maximum depth nm . Furthermore, the probability of the largest candidate network (with the maximum depth) being sampled is extremely small: m^{-n} . Therefore, the largest candidate networks are poorly trained due to the small possibility of being sampled. To this end, we introduce progressively depth shrinking to alleviate this issue. We divide the training epochs into m segments for m different depth choices. During the k^{th} training segment, we only allow the depth of each stage to be selected from $m - k + 1$ to m . This is essentially designed to enlarge the search space gradually so that these large candidate networks can be sampled more frequently.

4.3 Search Algorithm

After the super network is fully trained, we use evolutionary architecture search to find the best architectures under a certain resource constraint.

Resource Constraints. We use the number of MACs as the resource constraint. For 3D CNNs, the number of MACs cannot be simply determined by the input size and the network architecture: *e.g.*, Sparse Convolution only performs the computation over the active synapses; therefore, its computation is also determined by the input sparsity pattern. To address this, we first estimate the average kernel map size over the entire dataset for each convolution layer, and we can then measure the number of MACs based on these statistics.

Evolutionary Search. We automate the architecture search with the evolutionary algorithm [15]. We initialize the starting population with n randomly sampled candidate networks. At every iteration, we evaluate all candidate networks in the population and select the k models with the highest mIoUs (*i.e.*, the fittest individuals). The population for the next iteration is then generated with $(n/2)$ mutations and $(n/2)$ crossovers. For each mutation, we randomly select one among the top- k candidates and alter each of its architectural parameters (*e.g.*, channel numbers, network depths) with a pre-defined probability; for each crossover, we select two from the top- k candidates and generate a new model by fusing them together randomly. Finally, the best model is obtained from the population of the last iteration. During the evolutionary search, we ensure that all the candidate networks in the population always meet the given resource constraint (we will resample another candidate network until the resource constraint is satisfied).

5 Experiments

Based on our lightweight 3D module, we first manually construct our backbone network (denoted as SPVCNN). Then, we leverage our neural architecture search framework to explore the best 3D model (denoted as SPVNAS). We provide more implementation details in the appendix. Evaluated on 3D semantic segmentation and 3D object detection, our proposed method consistently outperforms previous state-of-the-art models with lower computation cost and measured latency (on a single GTX1080Ti GPU).

	#Params (M)	#MACs (G)	Latency (ms)	mIoU
PointNet [40]	3.0*	–	500*	14.6
SPGraph [24]	0.3*	–	5200*	17.4
PointNet++ [43]	6.0*	–	5900*	20.1
PVCNN [32]	2.5	42.4	146	39.0
TangentConv [55]	0.4*	–	3000*	40.9
RandLA-Net [19]	1.2	66.5	880 (256+624)	53.9
KPConv [56]	18.3	207.3	–	58.8
MinkowskiNet [9]	21.7	114.0	294	63.1
SPVNAS (Ours)	2.6	15.0	110	63.7
	12.5	73.8	259	66.4

Table 2. Results of outdoor scene segmentation on SemanticKITTI. SPVNAS outperforms MinkowskiNet with $2.7\times$ measured speedup. Here, **red numbers** correspond to the computation time, and **blue numbers** correspond to the post-processing time. *: results directly taken from Behley *et al.* [3].

	#Params (M)	#MACs (G)	Latency (ms)	mIoU
DarkNet21Seg [3]	24.7	212.6	73 (49+24)	47.4
DarkNet53Seg [3]	50.4	376.3	102 (78+24)	49.9
SqueezeSegV3-21 [68]	9.4	187.5	97 (73+24)	51.6
SqueezeSegV3-53 [68]	26.2	515.2	238 (214+24)	55.9
3D-MiniNet [1]	4.0	–	–	55.8
PolarNet [76]	13.6	135.0	62	57.2
SalsaNext [10]	6.7	62.8	71 (47+24)	59.5
SPVNAS (Ours)	1.1	8.9	89	60.3

Table 3. Results of outdoor scene segmentation on SemanticKITTI. SPVNAS outperforms the 2D projection-based methods with at least $3.6\times$ model size reduction and $7.1\times$ computation reduction. Here, **red numbers** correspond to the computation time, and **blue numbers** correspond to the projection time.

5.1 3D Scene Segmentation

We first evaluate our method on 3D semantic segmentation and conduct experiments on the large-scale outdoor scene dataset, SemanticKITTI [3]. This dataset contains 23,201 LiDAR point clouds for training and 20,351 for testing, and it is annotated from all 22 sequences in the KITTI [12] Odometry benchmark. We train all models on the entire training set and report the mean intersection-over-union (mIoU) on the official test set under the single scan setting. We provide additional experimental results (on both validation and test set) in the appendix.

Results. As shown in Table 2, SPVNAS outperforms the previous state-of-the-art MinkowskiNet [9] by 3.3% in mIoU with $1.7\times$ model size reduction, $1.5\times$ computation reduction and $1.1\times$ measured speedup. Further, we downscale our SPVNAS by setting the resource constraint to 15G MACs. This offers us with a

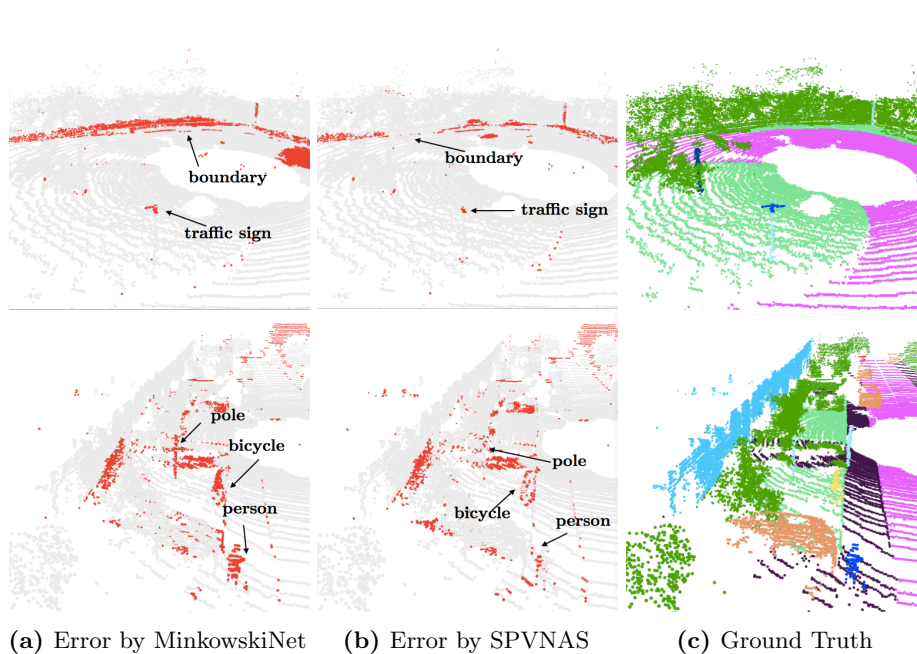


Fig. 4. MinkowskiNet has a higher error recognizing small objects and region boundaries, while SPVNAS recognizes small objects better thanks to the high-resolution point-based branch.

much smaller model that outperforms MinkowskiNet by 0.6% in mIoU with $8.3\times$ model size reduction, $7.6\times$ computation reduction, and $2.7\times$ measured speedup. In Figure 4, we also provide some qualitative comparisons between SPVNAS and MinkowskiNet: our SPVNAS has lower errors especially for small instances.

We further compare our SPVNAS with 2D projection-based models in Table 3. With the smaller backbone (by removing the decoder layers), SPVNAS outperforms DarkNets [3] by more than 10% in mIoU with $1.2\times$ measured speedup even though 2D convolutions are much better optimized by modern deep learning libraries. Compared with other 2D methods, SPVNAS achieves at least $8.5\times$ model size reduction and $15.2\times$ computation reduction while being much more accurate. Furthermore, our SPVNAS achieves higher mIoU than KPConv [56], which is the previous state-of-the-art point-based model, with $17\times$ model size reduction and $23\times$ computation reduction.

5.2 3D Object Detection

We also evaluate our method on 3D object detection and conduct experiments on the outdoor scene dataset, KITTI [12]. We follow the generally adopted training-validation split, where 3,712 samples are used for training and 3,769 samples are left for validation. We report the mean average precision (mAP) on the test set

	Car			Cyclist			Pedestrian		
	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard
SECOND [70]	84.7	76.0	68.7	75.8	60.8	53.7	45.3	35.5	33.1
SECOND (Repro.)	87.5	77.9	74.4	76.0	59.7	52.9	49.1	41.7	39.1
SPVCNN (Ours)	87.8	78.4	74.8	80.1	63.7	56.2	49.2	41.4	38.4

Table 4. Results of outdoor object detection on KITTI. SPVCNN outperforms SECOND in most categories especially for the cyclist.

	Person	Bicycle	Bicyclist	Motorcycle	Motorcyclist
MinkowskiNet [9]	60.9	40.4	61.9	47.4	18.7
SPVNAS (Ours)	65.7 (+4.8)	51.6 (+11.2)	65.2 (+3.3)	50.8 (+3.4)	43.7 (+25.0)

Table 5. Results of per-class performance on SemanticKITTI. SPVNAS has a large advantage on small objects, such as bicyclist and motorcyclist.

with 3D IoU thresholds of 0.7 for car, 0.5 for cyclist and pedestrian. We refer the readers to the appendix for more experimental results on the validation set.

Results. We compare our method against SECOND [70], the state-of-the-art single-stage model for 3D object detection. SECOND consists of a sparse encoder using 3D Sparse Convolutions and a region proposal network that performs 2D convolutions after projecting the encoded features to the bird’s-eye view (BEV). We reimplement and retrain SECOND: our implementation already outperforms the results in the original paper [70]. As for our model, we only replace these 3D Sparse Convolutions in SECOND with our SPVConv while keeping all the other settings the same for fair comparison. As summarized in Table 4, our SPVCNN achieves significant improvement in cyclist detection, for which we argue that the high-resolution point-based branch carries more information for small instances.

6 Analysis

Our SPVNAS achieves higher accuracy and better efficiency than the previous state-of-the-art MinkowskiNet. In this section, we provide more detailed analysis to better understand the contributions of SPVConv and 3D-NAS.

6.1 Sparse Point-Voxel Convolution (SPVConv)

From Table 5, our SPVNAS has a very large advantage (up to 25%) on relatively small objects such as pedestrians and cyclists. To explain this, we train SPVCNN on SemanticKITTI with the sequence 08 left out for visualization. In Figure 5, we highlight the points with top 5% feature norm within the point-based branch (in the final SPVConv). Clearly, the point-based branch learns to attend to small instances such as pedestrians, cyclists, trunks and traffic signs, which echos with our superior performance on these classes.




Fig. 5. The point-based branch learns to put its attention on small instances (*i.e.*, pedestrians, cyclists, traffic signs). Here, the points in red are the ones with the top 5% largest feature norm in the point-based branch.




Fig. 6. Average percentage of activated points in point-based and sparse voxel-based branches for all 19 classes in SemanticKITTI [3]: the point-based branch attends to smaller objects as the red bars are much higher.

Further, we quantitatively analyze the feature norms from both point-based and sparse voxel-based branches. Specifically, we first rank the points from both branches separately based on their feature norms, and then, we mark these points with top 10% feature norm in each branch as activated. From Figure 6, there are significantly more points in the point-based branch being activated for small instances: *e.g.*, more than 80% for bicyclist. This suggests that our observation in Figure 5 generally holds.

6.2 3D Neural Architecture Search (3D-NAS)

In Figure 7, we present both mIoU *vs.* #MACs and mIoU *vs.* latency trade-offs, where we uniformly scale the channel numbers in MinkowskiNet and SPVCNN down as our baselines. It can be observed that a better 3D module (SPVConv) and a well-designed network architecture (3D-NAS) are equally important to the final performance boost. Remarkably, SPVNAS outperforms MinkowskiNet by more than 6% in mIoU at 110 ms latency. Such a large improvement comes from the non-uniform channel scaling and elastic network depth. In these manually-

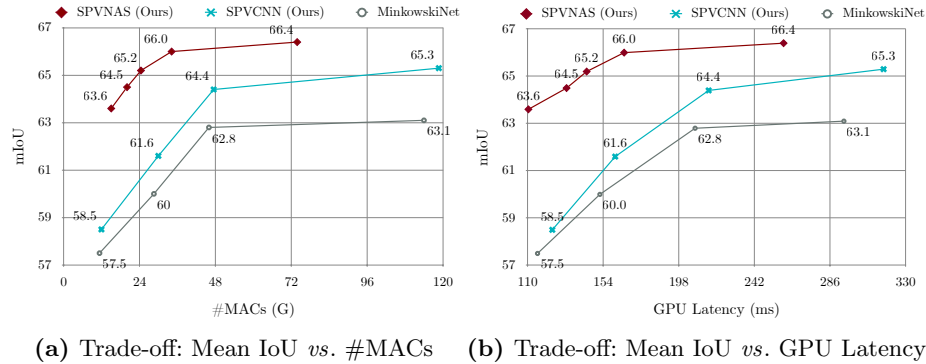


Fig. 7. An efficient 3D module (SPVConv) and a well-designed network architecture (3D-NAS) are equally important to the final performance of SPVNAS: **7.6** \times computation reduction and **2.7** \times measured speedup over MinkowskiNet.

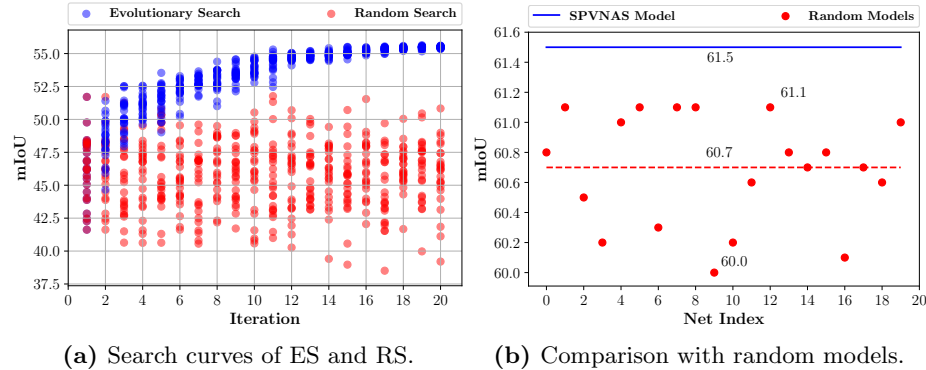


Fig. 8. Evolutionary Search (ES) is more sample-efficient than Random Search (RS).

designed models (MinkowskiNet and SPVCNN), 77% of the total computation is distributed to the upsampling stage. With 3D-NAS, this ratio is reduced to 47-63%, making the computation more balanced and the downsampling stage (*i.e.*, feature extraction) more emphasized.

We also compare our evolutionary search with the random architecture search to show that the success of 3D-NAS does not entirely come from the search space. As in Figure 8a, random architecture search has poor sample efficiency: the best model at the 20th generation performs even worse than the best model in the first generation. In contrast, our evolutionary search is capable of progressively finding better architecture, and the final best architecture performs around 3% better than the one in the first generation. We also retrain 20 random models sampled from our search space and compare them with SPVNAS in Figure 8b. As a result, our SPVNAS performs 0.8% better compared with the average performance of these random models.

7 Conclusion

We present Sparse Point-Voxel Convolution (SPVConv), a lightweight 3D module specialized for small object recognition. With SPVCNN built upon SPVConv, we solve the problem that Sparse Convolution cannot always keep high-resolution representation and Point-Voxel Convolution does not scale up to large 3D scenes. Furthermore, we introduce 3D-NAS, the first architecture search framework for 3D scene understanding that greatly improves the efficiency and performance of SPVCNN. Extensive experiments on outdoor 3D scene benchmarks demonstrate that the resulting SPVNAS model is lightweight, fast and powerful. We hope that this work will inspire future research on efficient 3D deep learning.

Acknowledgements. We thank Nick Stathas and Yue Dong for their feedback on the draft. This work is supported by MIT Quest for Intelligence, MIT-IBM Watson AI Lab, Xilinx and Samsung. We thank AWS Machine Learning Research Awards for providing the computational resource.

References

1. Alonso, I., Riazuelo, L., Montesano, L., Murillo, A.C.: 3D-MiniNet: Learning a 2D Representation from Point Clouds for Fast and Efficient 3D LIDAR Semantic Segmentation. arXiv (2020)
2. Bae, W., Lee, S., Lee, Y., Park, B., Chung, M., Jung, K.H.: Resource Optimized Neural Architecture Search for 3D Medical Image Segmentation. In: MICCAI (2019)
3. Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In: ICCV (2019)
4. Cai, H., Gan, C., Wang, T., Zhang, Z., Han, S.: Once for All: Train One Network and Specialize it for Efficient Deployment. In: ICLR (2020)
5. Cai, H., Lin, J., Lin, Y., Liu, Z., Wang, K., Wang, T., Zhu, L., Han, S.: AutoML for Architecting Efficient and Specialized Neural Networks. IEEE Micro (2019)
6. Cai, H., Zhu, L., Han, S.: ProxylessNAS: Direct Neural Architecture Search on Target Task and Hardware. In: ICLR (2019)
7. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: ShapeNet: An Information-Rich 3D Model Repository. arXiv (2015)
8. Chen, Y., Yang, T., Zhang, X., Meng, G., Xiao, X., Sun, J.: DetNAS: Backbone Search for Object Detection. In: NeurIPS (2019)
9. Choy, C., Gwak, J., Savarese, S.: 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. In: CVPR (2019)
10. Cortinhal, T., Tzelepis, G., Aksoy, E.E.: SalsaNext: Fast, Uncertainty-aware Semantic Segmentation of LiDAR Point Clouds for Autonomous Driving. arXiv (2020)
11. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR (2009)
12. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets Robotics: The KITTI Dataset. IJRR (2013)
13. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In: CVPR (2012)
14. Graham, B., Engelcke, M., van der Maaten, L.: 3D Semantic Segmentation With Submanifold Sparse Convolutional Networks. In: CVPR (2018)

15. Guo, Z., Zhang, X., Mu, H., Heng, W., Liu, Z., Wei, Y., Sun, J.: Single Path One-Shot Neural Architecture Search with Uniform Sampling. In: ECCV (2020)
16. Han, L., Zheng, T., Xu, L., Fang, L.: OccuSeg: Occupancy-aware 3D Instance Segmentation. In: CVPR (2020)
17. He, Y., Lin, J., Liu, Z., Wang, H., Li, L.J., Han, S.: AMC: AutoML for Model Compression and Acceleration on Mobile Devices. In: ECCV (2018)
18. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. arXiv (2017)
19. Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A.: RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds. In: CVPR (2020)
20. Iandola, F.N., Han, S., Moskewicz, M.W., Ashraf, K., Dally, W.J., Keutzer, K.: SqueezeNet: AlexNet-Level Accuracy with 50x Fewer Parameters and < 0.5MB Model Size. arXiv (2016)
21. Jiang, L., Zhao, H., Shi, S., Liu, S., Fu, C.W., Jia, J.: PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation. In: CVPR (2020)
22. Kim, S., Kim, I., Lim, S., Baek, W., Kim, C., Cho, H., Yoon, B., Kim, T.: Scalable Neural Architecture Search for 3D Medical Image Segmentation. In: MICCAI (2019)
23. Lahoud, J., Ghanem, B., Pollefeys, M., Oswald, M.R.: 3D Instance Segmentation via Multi-Task Metric Learning. In: ICCV (2019)
24. Landrieu, L., Simonovsky, M.: Large-Scale Point Cloud Semantic Segmentation With Superpoint Graphs. In: CVPR (2018)
25. Lei, H., Akhtar, N., Mian, A.: Octree Guided CNN With Spherical Kernels for 3D Point Clouds. In: CVPR (2019)
26. Li, G., Qian, G., Delgadillo, I.C., Muller, M., Thabet, A., Ghanem, B.: SGAS: Sequential Greedy Architecture Search. In: CVPR (2020)
27. Li, M., Lin, J., Ding, Y., Liu, Z., Zhu, J.Y., Han, S.: GAN Compression: Efficient Architectures for Interactive Conditional GANs. In: CVPR (2020)
28. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: PointCNN: Convolution on \mathcal{X} -Transformed Points. In: NeurIPS (2018)
29. Liu, C., Zoph, B., Neumann, M., Shlens, J., Hua, W., Li, L.J., Fei-Fei, L., Yuille, A., Huang, J., Murphy, K.: Progressive Neural Architecture Search. In: ECCV (2018)
30. Liu, H., Simonyan, K., Yang, Y.: DARTS: Differentiable Architecture Search. In: ICLR (2019)
31. Liu, Z., Mu, H., Zhang, X., Guo, Z., Yang, X., Cheng, K.T., Sun, J.: MetaPruning: Meta Learning for Automatic Neural Network Channel Pruning. In: ICCV (2019)
32. Liu, Z., Tang, H., Lin, Y., Han, S.: Point-Voxel CNN for Efficient 3D Deep Learning. In: NeurIPS (2019)
33. Ma, N., Zhang, X., Zheng, H.T., Sun, J.: ShuffleNet V2: Practical Guidelines for Efficient CNN Architecture Design. In: ECCV (2018)
34. Ma, Z., Zhou, Z., Liu, Y., Lei, Y., Yan, H.: Auto-ORVNet: Orientation-Boosted Volumetric Neural Architecture Search for 3D Shape Classification. IEEE Access (2020)
35. Mao, J., Wang, X., Li, H.: Interpolated Convolutional Networks for 3D Point Cloud Understanding. In: ICCV (2019)
36. Maturana, D., Scherer, S.: VoxNet: A 3D Convolutional Neural Network for Real-Time Object Recognition. In: IROS (2015)
37. Pagh, R., Rodler, F.F.: Cuckoo Hashing. Journal of Algorithms (2001)
38. Qi, C.R., Chen, X., Litany, O., Guibas, L.J.: ImVoteNet: Boosting 3D Object Detection in Point Clouds with Image Votes. In: CVPR (2020)

39. Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep Hough Voting for 3D Object Detection in Point Clouds. In: ICCV (2019)
40. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In: CVPR (2017)
41. Qi, C.R., Liu, W., Wu, C., Su, H., Guibas, L.J.: Frustum PointNets for 3D Object Detection from RGB-D Data. In: CVPR (2018)
42. Qi, C.R., Su, H., Niessner, M., Dai, A., Yan, M., Guibas, L.J.: Volumetric and Multi-View CNNs for Object Classification on 3D Data. In: CVPR (2016)
43. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In: NeurIPS (2017)
44. Radosavovic, I., Johnson, J., Xie, S., Lo, W.Y., Dollár, P.: On Network Design Spaces for Visual Recognition. In: ICCV (2019)
45. Riegler, G., Ulusoy, A.O., Geiger, A.: OctNet: Learning Deep 3D Representations at High Resolutions. In: CVPR (2017)
46. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: MobileNetV2: Inverted Residuals and Linear Bottlenecks. In: CVPR (2018)
47. Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H.: PV-RCNN: Point-Voxel Feature Set Abstraction for 3D Object Detection. In: CVPR (2020)
48. Shi, S., Wang, X., Li, H.: PointRCNN: 3D Object Proposal Generation and Detection from Point Cloud. In: CVPR (2019)
49. Shi, S., Wang, Z., Shi, J., Wang, X., Li, H.: PV-RCNN: Point-Voxel Feature Set Abstraction for 3D Object Detection. TPAMI (2020)
50. Stamoulis, D., Ding, R., Wang, D., Lymberopoulos, D., Priyantha, B., Liu, J., Marculescu, D.: Single-Path NAS: Designing Hardware-Efficient ConvNets in less than 4 Hours. arXiv (2019)
51. Strubell, E., Ganesh, A., McCallum, A.: Energy and Policy Considerations for Deep Learning in NLP. In: ACL (2019)
52. Su, H., Jampani, V., Sun, D., Maji, S., Kalogerakis, E., Yang, M.H., Kautz, J.: SPLATNet: Sparse Lattice Networks for Point Cloud Processing. In: CVPR (2018)
53. Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: MnasNet: Platform-Aware Neural Architecture Search for Mobile. In: CVPR (2019)
54. Tan, M., Le, Q.V.: EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In: ICML (2019)
55. Tatarchenko, M., Park, J., Koltun, V., Zhou, Q.Y.: Tangent Convolutions for Dense Prediction in 3D. In: CVPR (2018)
56. Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: KPConv: Flexible and Deformable Convolution for Point Clouds. In: ICCV (2019)
57. Wang, H., Wu, Z., Liu, Z., Cai, H., Zhu, L., Gan, C., Han, S.: HAT: Hardware-Aware Transformers for Efficient Natural Language Processing. In: ACL (2020)
58. Wang, K., Liu, Z., Lin, Y., Lin, J., Han, S.: HAQ: Hardware-Aware Automated Quantization with Mixed Precision. In: CVPR (2019)
59. Wang, K., Liu, Z., Lin, Y., Lin, J., Han, S.: Hardware-Centric AutoML for Mixed-Precision Quantization. IJCV (2020)
60. Wang, P.S., Liu, Y., Guo, Y.X., Sun, C.Y., Tong, X.: O-CNN: Octree-based Convolutional Neural Networks for 3D Shape Analysis. In: SIGGRAPH (2017)
61. Wang, P.S., Liu, Y., Guo, Y.X., Sun, C.Y., Tong, X.: Adaptive O-CNN: A Patch-based Deep Representation of 3D Shapes. In: SIGGRAPH Asia (2018)
62. Wang, T., Wang, K., Cai, H., Lin, J., Liu, Z., Wang, H., Lin, Y., Han, S.: APQ: Joint Search for Network Architecture, Pruning and Quantization Policy. In: CVPR (2020)

63. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic Graph CNN for Learning on Point Clouds. In: SIGGRAPH (2019)
64. Wang, Z., Lu, F.: VoxSegNet: Volumetric CNNs for Semantic Part Segmentation of 3D Shapes. TVCG (2019)
65. Wong, K.C., Moradi, M.: SegNAS3D: Network Architecture Search with Derivative-Free Global Optimization for 3D Image Segmentation. In: MICCAI (2019)
66. Wu, B., Dai, X., Zhang, P., Wang, Y., Sun, F., Wu, Y., Tian, Y., Vajda, P., Jia, Y., Keutzer, K.: FBNet: Hardware-aware Efficient Convnet Design via Differentiable Neural Architecture Search. In: CVPR (2019)
67. Wu, W., Qi, Z., Fuxin, L.: PointConv: Deep Convolutional Networks on 3D Point Clouds. In: CVPR (2019)
68. Xu, C., Wu, B., Wang, Z., Zhan, W., Vajda, P., Keutzer, K., Tomizuka, M.: SqueezeSegV3: Spatially-Adaptive Convolution for Efficient Point-Cloud Segmentation. In: ECCV (2020)
69. Xu, Y., Fan, T., Xu, M., Zeng, L., Qiao, Y.: SpiderCNN: Deep Learning on Point Sets with Parameterized Convolutional Filters. In: ECCV (2018)
70. Yan, Y., Mao, Y., Li, B.: SECOND: Sparsely Embedded Convolutional Detection. Sensors (2018)
71. Yang, B., Wang, J., Clark, R., Hu, Q., Wang, S., Markham, A., Trigoni, N.: Learning Object Bounding Boxes for 3D Instance Segmentation on Point Clouds. In: NeurIPS (2019)
72. Yang, D., Roth, H., Xu, Z., Milletari, F., Zhang, L., Xu, D.: Searching Learning Strategy with Reinforcement Learning for 3D Medical Image Segmentation. In: MICCAI (2019)
73. Yang, Z., Sun, Y., Liu, S., Shen, X., Jia, J.: STD: Sparse-to-Dense 3D Object Detector for Point Cloud. In: ICCV (2019)
74. Yu, Q., Yang, D., Roth, H., Bai, Y., Zhang, Y., Yuille, A., Xu, D.: C2FNAS: Coarse-to-Fine Neural Architecture Search for 3D Medical Image Segmentation. In: CVPR (2020)
75. Zhang, X., Zhou, X., Lin, M., Sun, J.: ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices. In: CVPR (2018)
76. Zhang, Y., Zhou, Z., David, P., Yue, X., Xi, Z., Gong, B., Foroosh, H.: PolarNet: An Improved Grid Representation for Online LiDAR Point Clouds Semantic Segmentation. In: CVPR (2020)
77. Zhou, Y., Tuzel, O.: VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. In: CVPR (2018)
78. Zhu, Z., Liu, C., Yang, D., Yuille, A., Xu, D.: V-NAS: Neural Architecture Search for Volumetric Medical Image Segmentation. In: 3DV (2019)
79. Zoph, B., Le, Q.V.: Neural Architecture Search with Reinforcement Learning. In: ICLR (2017)
80. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning Transferable Architectures for Scalable Image Recognition. In: CVPR (2018)

A.1 Implementation Details

We provide more implementation details on how to build our backbone network (SPVCNN), and train the super network and search for the best model (3D-NAS).

A.1.1 SPVCNN: Backbone Network

Based on MinkowskiNet [9], we build our backbone network by wrapping residual Sparse Convolution blocks with the high-resolution point-based branch. Specifically, the first SPVConv voxelizes before the first layer and devoxelizes after the stemming stage (*i.e.*, before the first downsampling). The second SPVConv voxelizes right after the first SPVConv and devoxelizes after all four downsampling stages. The final two SPVConv’s each wraps around two upsampling stages.

Also, we design a smaller backbone based on PVCNN [32] by directly replacing each volumetric convolution with one convolution layer (followed by normalization and activation layers) and two residual Sparse Convolution blocks.

A.1.2 3D-NAS: Architecture Search

We train the super network for 15 epochs that supports the fine-grained channel setting with a starting learning rate 0.24 and cosine learning rate decay. Then, we train for another 15 epochs to incorporate elastic network depth with a starting learning rate 0.096 and cosine learning rate decay. After that, we perform evolutionary architecture search with a population of 50 candidates for 20 generations on the official validation set (sequence 08). Best architecture is directly extracted from the super network and submitted to the test server after finetuning for 10 epochs with a starting learning rate of 0.032 and cosine learning rate decay.

A.2 More Results

A.2.1 3D Scene Segmentation

We present more detailed results of MinkowskiNet [9], SPVCNN and SPVNAS on both the official test set and validation set (sequence 08) of SemanticKITTI [3] in Table 6 and Table 7. For the results on the validation set in Table 7, we run the whole architecture search pipeline again on sequences 00-07 and 09, leaving sequence 10 out as the mini-validation set, and report the results on sequence 08. We observe similar trends on both test and validation results: both a better 3D module (SPVConv) and our 3D Neural Architecture Search (3D-NAS) pipeline improve the performance of MinkowskiNet [9].

A.2.2 3D Object Detection

We further provide the results of SPVCNN on the validation set of KITTI [13] in Table 8. We train both SECOND [70] and our SPVCNN on the training set for three times, and we report the average results to reduce the variance. Similar to the results on the test set, SPVCNN also has consistent improvement in almost all classes on the validation set.

	#Params (M)	#MAdds (G)	Latency (ms)	Mean IoU
MinkowskiNet [9]	2.2	11.3	115.7	57.5
SPVCNN (Ours)	2.2	11.9	124.3	58.5
SPVNAS (Ours)	2.6	15.0	110.4	63.6
MinkowskiNet [9]	5.5	28.5	152.0	60.0
SPVCNN (Ours)	5.5	30.0	160.9	61.6
SPVNAS (Ours)	4.2	20.0	132.6	64.5
MinkowskiNet [9]	8.8	45.9	207.4	62.8
SPVCNN (Ours)	8.8	47.4	214.3	64.4
SPVNAS (Ours)	5.1	24.4	144.3	65.2
MinkowskiNet [9]	21.7	113.9	294.0	63.1
SPVCNN (Ours)	21.8	118.6	317.1	63.8
SPVNAS (Ours)	7.5	34.1	166.1	66.0
SPVNAS (Ours)	12.5	73.8	259.9	66.4

Table 6. Results of 3D scene segmentation on the test set of SemanticKITTI [3].

	#Params (M)	#MAdds (G)	Latency (ms)	Mean IoU
MinkowskiNet [9]	5.5	28.5	152.0	58.9
SPVCNN (Ours)	5.5	30.0	160.9	60.7
SPVNAS (Ours)	4.5	24.6	158.1	62.9
MinkowskiNet [9]	8.8	45.9	207.4	60.3
SPVCNN (Ours)	8.8	47.4	214.3	61.4
SPVNAS (Ours)	7.0	34.7	175.8	63.5
MinkowskiNet [9]	21.7	113.9	294.0	61.1
SPVCNN (Ours)	21.8	118.6	317.1	63.8
SPVNAS (Ours)	10.8	64.5	248.7	64.7

Table 7. Results of 3D scene segmentation on the validation set of SemanticKITTI [3].

A.3 More Visualizations

A.3.1 3D Scene Segmentation

We provide more visualizations for MinkowskiNet [9] and SPVNAS in Figure A1 to demonstrate that the improvements brought by SPVConv on small objects and region boundaries are general. For instance, in the first row, MinkowskiNet segments the entire traffic sign and bicycle instances incorrectly; however, our SPVNAS is capable of making almost no mistakes on these very small objects. Also, we observe in next two rows that MinkowskiNet does not perform well on the sidewalk-building or sidewalk-vegetation boundaries where our SPVNAS has a clear advantage.

In Figure A2, we compare our smaller SPVNAS with DarkNet53 [3]. DarkNet53 uses spherical projections to project 3D point clouds to a 2D plane such

	Car			Cyclist			Pedestrian		
	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard
SECOND [70]	89.8	80.9	78.4	82.5	62.8	58.9	68.3	60.8	55.3
SPVCNN (Ours)	90.9	81.8	79.2	85.1	63.8	60.1	68.2	61.6	55.9

Table 8. Results of 3D object detection on the validation set of KITTI [12].

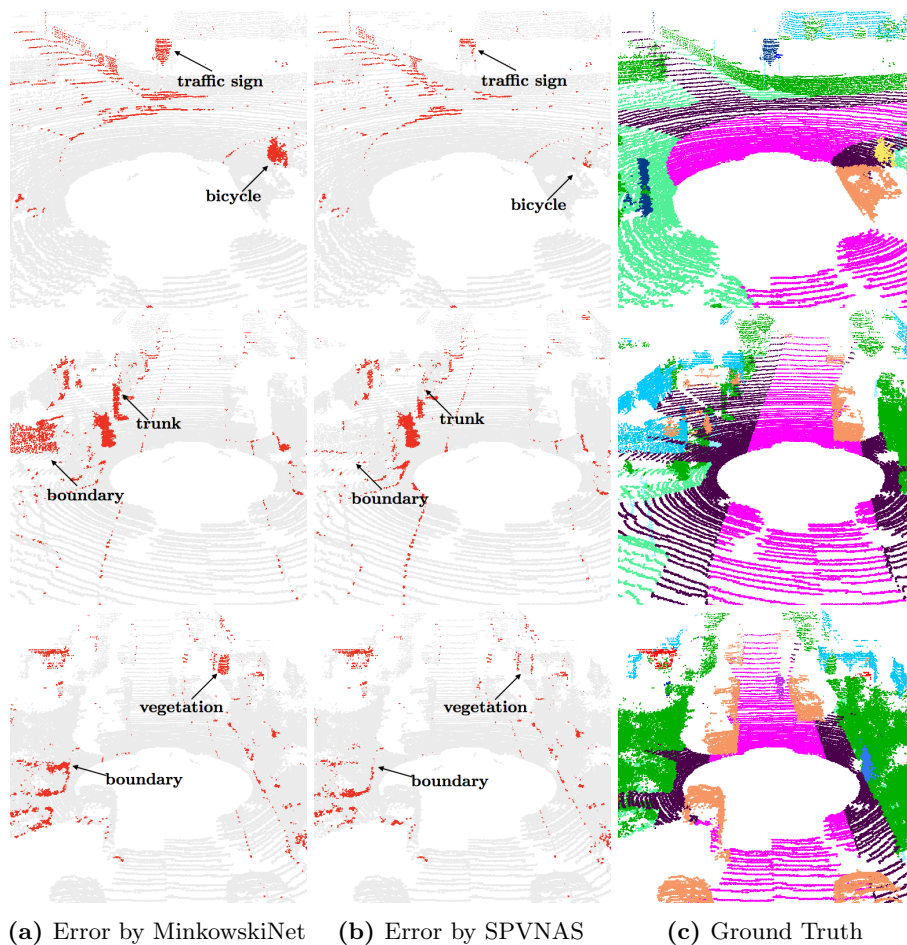


Fig. A1. Qualitative comparisons between MinkowskiNet [9] and SPVNAS.

that part of the geometry information is lost. Our SPVNAS, in contrast, directly learns on 3D data and is more powerful in the geometric modeling. We observe that SPVNAS has significant advantages in both large regions (last three rows) and smaller instances (first row).

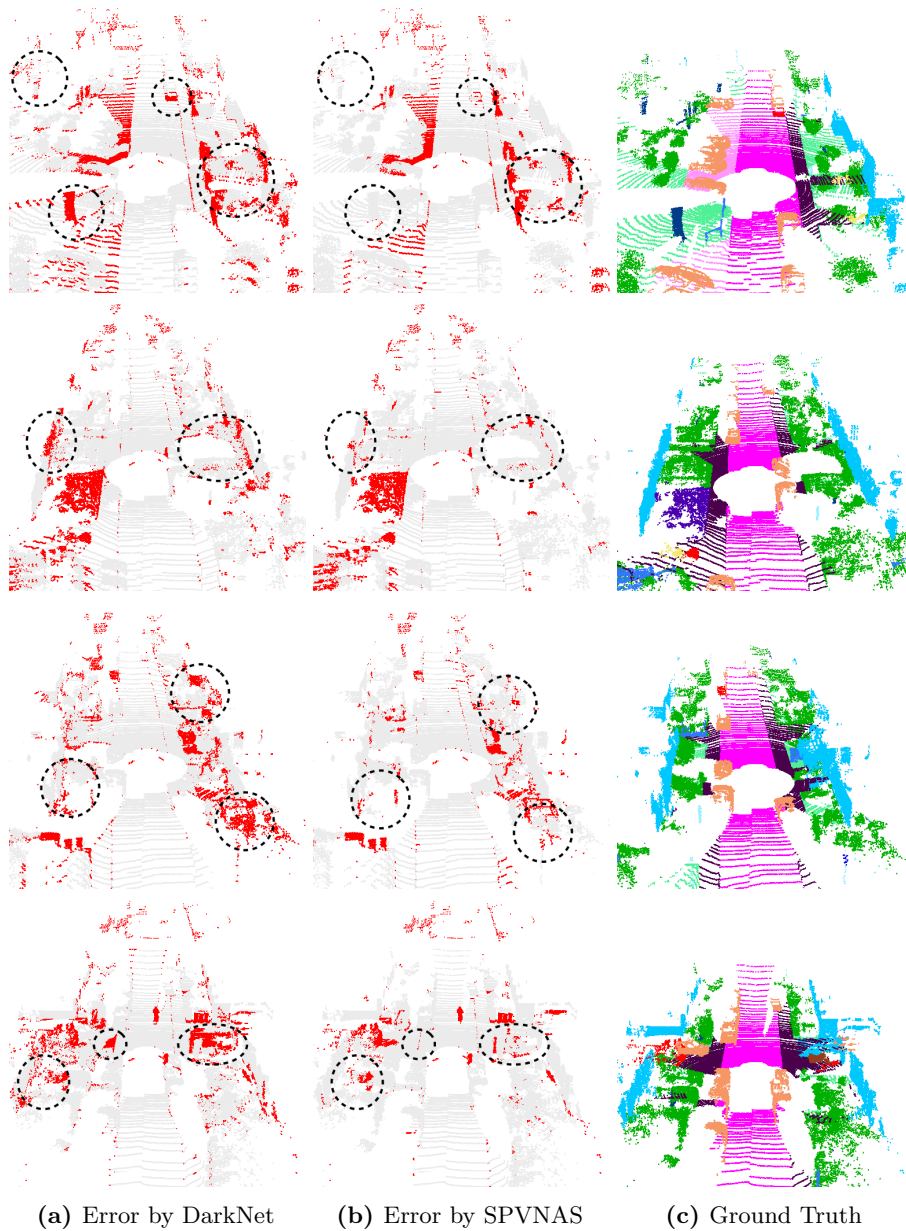


Fig. A2. Qualitative comparisons between DarkNet53 [3] and SPVNAS.

A.3.2 3D Object Detection

In Figure A3, we demonstrate the superior performance of SPVCNN over SECOND [70], which is a state-of-the-art single-stage 3D detector. Our SPVCNN




Fig. A3. Qualitative comparisons between SECOND [70] and SPVCNN.

has a large advantage in scenes with crowded small objects. In the first row of Figure A3, our SPVCNN is capable of detecting a challenging small pedestrian instance missed by SECOND, and it also avoids duplicate predictions made by SECOND in the next three rows.