

Joint prototype and coefficient prediction for 3D instance segmentation

Remco Royen,^{1,2,✉} Leon Denis,^{1,2} and Adrian Munteanu^{1,2}

¹Department of Electronics and Informatics, Vrije Universiteit Brussel, Brussels, Belgium

²IMEC, Leuven, Belgium

✉ E-mail: remco.royen@vub.be

3D instance segmentation is crucial for applications demanding comprehensive 3D scene understanding. Here, a novel method is introduced that simultaneously learns coefficients and prototypes. Employing an overcomplete sampling strategy, the method produces an overcomplete set of instance predictions, from which the optimal ones are selected through a Non-Maximum Suppression (NMS) algorithm during inference. The obtained prototypes are visualizable and interpretable. The method demonstrates superior performance on S3DIS-blocks, consistently outperforming existing methods in mRec and mPrec. Moreover, it operates 32.9% faster than the state-of-the-art. Notably, with only 0.8% of the total inference time, the method exhibits an over 20-fold reduction in the variance of inference time compared to existing methods. These attributes render the method well-suited for practical applications requiring both rapid inference and high reliability.

Introduction: Instance segmentation is receiving considerable attention in the field of scene understanding, particularly with the rise of AI. This task distinguishes itself from semantic segmentation [1–3] by not only assigning a semantic class to each pixel but also differentiating between objects for a deeper understanding. While 2D instance segmentation has been well-explored [4–11], our focus is on 3D instance segmentation. This domain presents additional challenges due to the unstructured and unordered nature and the potential high variations in instance sizes of the point clouds.

Existing methods in this domain can be broadly classified into two categories. Proposal-based methods [12–16] involve 3D object detection followed by semantic segmentation. A notable drawback of these methods lies in the criticality of proposal accuracy, as errors in the two consecutive tasks tend to accumulate. Clustering-based methods [17–29] compute discriminative features which are employed for object clustering. Some also compute additional features such as center offsets to aid the clustering [19, 22, 26, 27]. A recent advancement, Mask3D [30], sidesteps the need for precise proposals or time-consuming clustering by refining initial low-level queries using transformers.

In this paper, we introduce a novel 3D instance segmentation method that overcomes the mentioned drawbacks. Furthermore, it omits the time-expensive query refinement steps of Mask3D by jointly learning prototypes and coefficients. By employing a linear combination of these, we generate a set of instance predictions of which the best are retained through a rapid non-maximum suppression (NMS) algorithm. A key aspect of our method is the overcompleteness of the obtained set of coefficients which are derived from a diverse set, with fixed cardinality, of sampled points in the input. This avoids the need for precise proposal prediction. Not only does our approach outperform existing methods, it also demonstrates 32.9% faster processing and exhibits a 20 times smaller variance in inference timings in comparison to existing methods, ensuring reliability and suitability for online applications. A performance and speed comparison with the state-of-the-art is visualized with a boxplot in Figure 1. In summary, our main contributions can be outlined as follows:

- A novel, end-to-end prototype-based 3D instance segmentation method is proposed. Coefficients and prototypes are learned jointly and combined to predict the instance masks.
- A thorough validation on S3DIS blocks, showing state-of-the-art performance both on 5th-fold and 6-fold cross-validation (CV) settings. Additionally, the proposed method does not only achieve the fastest inference but also the lowest variance in inference time.

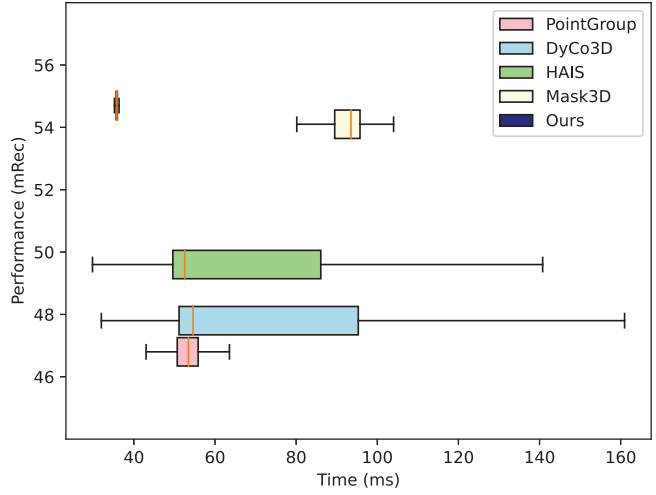


Fig. 1 Speed-performance comparison on S3DIS-blocks Area-5. The proposed method outperforms the state-of-the-art in terms of accuracy, speed and variance in inference time




Fig. 2 The proposed architecture consists of four main parts: (1) A feature extractor which retrieves per-point features. (2) The sampling of a diverse set of K points (3) A PointNet++ network that generates prototypes from the per-point features (4) In parallel, a PointConv network that computes coefficients for each sampled point. Instance predictions are obtained by linearly combining the coefficients and prototypes from which the optimal ones are selected during inference

Proposed method: The proposed method’s architecture is illustrated in Figure 2. Let $\mathbf{X} \in \mathbb{R}^{N \times I}$ be the input point cloud, where N and I represent the number of input points and channels, respectively. Per-point features $\mathbf{F} \in \mathbb{R}^{N \times F}$ for \mathbf{X} are computed, with F denoting the feature dimension. These features are utilized in two parallel branches. The first branch samples a set $\mathcal{S} = \{s_i\}_{i=1}^K$ of K points $s_i \in \mathbb{R}^3$ from the input data, and transforms the features based on the sampling into a coefficient set $\mathbf{C} \in \mathbb{R}^{K \times M}$, where M is the number of prototypes. Simultaneously, the second branch computes M prototypes, $\mathbf{P} \in \mathbb{R}^{N \times M}$, from \mathbf{F} . The linear combination of coefficients and prototypes enables the retrieval of the overcomplete set of instance masks $\mathbf{Y} \in \mathbb{R}^{K \times N}$. Mathematically, the combination of the coefficients and prototypes can be expressed as follows,

$$\mathbf{y}_k^q = \sum_{i=1}^M c_{k,i}^q \mathbf{p}_i^q, \quad (1)$$

with $\mathbf{y}_k^q = \mathbf{y}_k(\mathbf{X}_q) \in \mathbb{R}^N$ the k^{th} prediction mask in \mathbf{Y} for input sample \mathbf{X}_q ; $c_{k,i}^q = c_{k,i}(\mathbf{X}_q) \in \mathbb{R}$ the i^{th} coefficient for the k -th sampled point and $\mathbf{p}_i^q = \mathbf{p}_i(\mathbf{X}_q) \in \mathbb{R}^N$ the i^{th} prototype for input sample \mathbf{X}_q .

During inference, \mathbf{Y} is thresholded, and the K' optimal predictions are selected using a non-maximum suppression (NMS) algorithm, resulting in the final instance masks $\mathbf{Y}' \in \mathbb{R}^{K' \times N}$. Subsequently, we will delve into the details of the coefficients and prototypes prediction.

Coefficients: The coefficient computation begins with the sampling of K points from the input point cloud using Farthest Point Sampling (FPS). Unlike traditional proposal-based methods, our approach doesn’t necessitate precise proposal predictions but aims to acquire an overcomplete set of points showcasing large diversity in local neighborhoods. A PointConv network [31] calculates coefficients for each sampled point based




Fig. 3 Prototype p_i activation for different PartNet samples x_i

on the features within their local neighborhood, contributing to a diverse set of coefficients. After a linear combination with prototypes, this yields a great variety in instance masks from which the NMS algorithm selects the optimal predictions. The PointConv network uses *tanh* as its output activation function, accommodating negative coefficients and enabling the cancellation of specific regions among prototypes.

Prototypes: A PointNet++ network [32] utilizes the shared features F to compute a set of prototypes $P \in \mathbb{R}^{N \times M}$, consisting of M prototypes $p_i \in \mathbb{R}^N$. It is noteworthy that these prototypes maintain the point dimension of the input, effectively forming vectors consisting of scores which are associated to each input point. Visual representations of prototype scores on input point clouds from the PartNet-chair dataset [33] are presented in Figure 3. In these visualizations, white and black points signify high and low scores in a given prototype, respectively. Notably, each prototype, that is, row in Figure 3, emphasizes a specific region of the input point cloud, encapsulating instance information. The linear combinations of such prototypes enable the generation of instance masks.

Loss: The proposed method is trained end-to-end. As loss-function, we employ a binary cross entropy loss between the instance masks Y_i and the ground-truth instance masks GT_i . To avoid instability and accelerate training, we introduce a spatial constraint where the predicted instance mask associated to the sampled point s_i , $Y(s_i)$, should match the ground-truth instance that contains the sampled point, $GT(s_i)$. Mathematically, this is expressed as follows:

$$J = \sum_{i=1}^{I_{PR}} L_{BCE}(Y(s_i), GT(s_i)), \quad (2)$$

with I_{PR} the number of predicted instances.

Experiments: To assess the efficacy of our proposed method, we conduct qualitative and quantitative experiments on the S3DIS dataset [34], encompassing 6 large-scale indoor areas with a total of 271 rooms. Each point in the dataset is annotated with one of 13 semantic categories. Our evaluation strategy and metrics align with established practices in the literature [25].

In the literature, two primary approaches for handling large-scale and dense datasets such as S3DIS are recognized. One, utilized by [17, 24, 25, 35, 36], firstly segments scenes into overlapping blocks. These are processed separately to instance masks, and merged using the BlockMerge algorithm [17, 36]. The other, found in [22, 26, 27, 37], processes the entire scene at once without merging. While the latter offers a higher segmentation performance as it avoids partial objects and provides a global overview, it incurs significant memory requirements, as depicted in Figure 4. PointGroup, favoring whole-scene processing, experiences escalating memory demands with an increasing number of points, exceeding 20 GB for the largest scenes. Notably, even on our powerful RTX 3090 GPU, four scenes were unable to be processed as a whole, necessitating downsampling to obtain results. In contrast, block-based methods, like ours, maintain a consistent GPU memory requirement for processing point clouds of arbitrary sizes. As example, our method requires a constant 2.4 GB peak GPU memory. Block-based (BB) methods




Fig. 4 Peak memory usage of the proposed method (block-based) and PointGroup [22] (full-scene) for different number of input points and number of instances (intensity of marker). The experiment is conducted on the scenes of S3DIS Area-5 [34].

are also suitable for online applications where the scene has to be processed during capturing. While not reaching the same performance as full-scene methods, block-based approaches remain relevant for practical applications, particularly on embedded devices or in online scenarios. Due to computational constraints, ProtoSeg couldn't be applied to the entire scene, and we present our results solely using a block-based approach (S3DIS-blocks).

We follow the experimental settings of [25, 38] and divide each room in $1\text{m} \times 1\text{m}$ blocks with stride 0.5 such that an overlap is created. In each block 4096 points are randomly sampled during training. Each input point is represented by a 9-dimensional vector: XYZ, RGB and the normalized location of the point in the whole room. We have used a PointNet [38] feature extractor, a PointConv without densitynet [31] to compute the coefficients and PointNet++ [32] for the prototypes. Semantic segmentation is obtained, in a block-based manner, by a PointTransformer [39] network. An Adam optimizer [40] is used with learning rate 0.001 and the network is trained with batch size 16 for 65 epochs. We have empirically selected the number of features, prototypes, sampled points and the sampling method as 64, 128, 64 and 'FPS on xyz', respectively, as it led to the best results. During inference, we employ a threshold of 0.3 and use a regular NMS to retrieve the instance mask predictions for each block. The same BlockMerge algorithm as proposed in [17] is employed to ensure a fair comparison between block-based methods.

Experimental results on S3DIS-blocks: Quantitative results on S3DIS-blocks, along with comparisons with existing methods, are presented in Table 1. Note that for a fair comparison, we retrained methods that do not operate on blocks using their official, publicly available code.

On the Area-5 validation, our method surpasses existing works in three out of the four metrics and secures the second-best result in the remaining one. More specifically, we achieve a 0.6% higher mean recall (mRec) and a notable 2.5% increase in mean precision (mPrec) over the closest competitor. Similar trends are observed in the 6-fold cross-validation, where our method slightly outperforms Mask3D by 0.3% in mRec and significantly surpasses the second-best by 2.1% in mPrec. Regarding timing, the proposed method stands out with a total inference time of 35.7 ms, making it 32.9% faster than the current fastest method. Additionally, our overcomplete approach significantly reduces the variance in inference time to just 0.8% of our total inference time, or 0.3 ms. This is over 20 times smaller than the smallest variance observed in existing methods.

Qualitative results are showcased in Figure 5. Overall, one can observe highly accurate segmentation, with nearly all instances detected in the corridor. Similarly, for the office, the segmentation results are notably accurate, encompassing the detection of chairs, walls, and apparel.

Conclusion: In conclusion, this paper presents a novel 3D instance segmentation method designed for accurate, fast and reliable scene

Table 1. Instance segmentation result on S3DIS-blocks for both Area-5 and 6-fold cross validation and timing. The methods denoted with * were retrained on blocks. Instance segmentation inference time and standard deviation are reported for S3DIS-blocks Area-5. For fair comparison, all runtimes were measured on the same CPU and NVIDIA GeForce RTX 3090. The '-' symbol is employed for non-public codebases, impeding the measurement of inference time. The '/' symbol indicates that a method does not incorporate that specific step.

Method	5th-fold				6-fold CV				Inference time (ms)			
	mCov	mWCov	mRec	mPrec	mCov	mWCov	mRec	mPrec	Network	NMS	Grouping	Total
JSNet [25]	48.7	51.5	46.9	62.1	54.1	58.0	53.9	66.9	34.0 ± 11.3	/	96.2 ± 70.4	133.6 ± 71.0
IAM [24]	49.9	53.2	48.5	61.3	54.5	58.0	51.8	67.2	—	—	—	—
MPNet [35]	50.1	53.2	49.0	62.5	55.8	59.7	53.7	68.4	—	—	—	—
PointGroup* [22]	43.1	47.8	46.8	55.1	48.4	54.4	55.3	59.4	47.1 ± 6.6	0.4 ± 0	4.8 ± 1.5	53.2 ± 6.3
DyCo3D* [26]	44.4	50.0	47.8	59.3	47.9	54.5	52.4	58.9	63.3 ± 28.7	0.5 ± 0.7	6.5 ± 7.3	72.7 ± 35.4
HAIS* [27]	45.4	50.6	49.6	58.4	48.2	55.4	55.0	60.7	57.4 ± 23.3	/	9.1 ± 10.0	69.3 ± 30.5
SoftGroup* [37]	40.6	45.6	43.9	42.6	45.0	51.2	48.4	52.3	43.8 ± 3.5	/	5.6 ± 2.7	55.4 ± 6.0
JSPNet [29]	50.7	53.5	48.0	59.6	54.9	58.8	55.0	66.5	—	—	—	—
Mask3D* [30]	51.4	56.7	54.1	59.9	52.3	61.0	56.2	67.1	73.8 ± 4.7	/	/	92.5 ± 8.3
Ours	50.8	57.0	54.7	65.0	52.5	59.9	56.5	70.5	35.0 ± 0.3	0.4 ± 0	/	35.7 ± 0.3




Fig. 5 Visualization of S3DIS results. Note that different colors represent different instances and that the same instance may have a different color in the ground-truth and prediction

understanding. The proposed approach leverages an overcomplete sampling strategy and simultaneously learns coefficients and prototypes. Through experimentation on the S3DIS-blocks dataset, our method demonstrates superior performance, consistently outperforming existing methods in key metrics such as mean recall and mean precision. Notably, with a total inference time of 35.7 ms, our method is 32.9% faster than the current fastest method. An additional strength of our method lies in its ultra-low variance in inference time, ensuring exceptional consistency and reliability.

In comparison with existing methodologies, our approach stands out as a promising solution for real-world applications demanding both accuracy and computational efficiency. We believe our contribution paves the way for further advancements in 3D instance segmentation, particularly in practical scenarios such as embedded devices or online applications.

Author contributions: **Remco Royen:** Conceptualization; data curation; formal analysis; funding acquisition; investigation; methodology; project administration; resources; software; validation; visualization; writing—original draft; writing—review and editing. **Leon Denis:** Conceptualization; formal analysis; investigation; methodology; supervision; visualization; writing—original draft; writing—review and editing. **Adrian Munteanu:** Conceptualization; formal analysis; funding acquisition; investigation; methodology; resources; supervision; writing—review and editing.

Acknowledgements: This work is funded by Fonds Wetenschappelijk Onderzoek (FWO) - 1S89420N and Innoviris within the research project SPECTRE.

Conflict of interest statement: The authors declare no conflicts of interest.

Data availability statement: The data that support the findings of this study are available from the corresponding author upon reasonable request.

© 2024 The Authors. *Electronics Letters* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology.

This is an open access article under the terms of the Creative Commons Attribution-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited and no modifications or adaptations are made.

Received: 30 November 2023 Accepted: 15 February 2024
doi: 10.1049/ell2.13137

References

- Lu, J., Cheng, H., Luo, M., Liu, T., Zhang, K.: Puconv: Upsampling convolutional network for point cloud semantic segmentation. *Electron. Lett.* **56**(9), 435–438 (2020)
- Joukovsky, B., Hu, P., Munteanu, A.: Multi-modal deep network for rgb-d segmentation of clothes. *Electron. Lett.* **56**(9), 432–435 (2020)
- Royen, R., Munteanu, A.: Rascal3d: Resolution scalable 3d semantic segmentation of point clouds. In: 2023 IEEE International Conference on Image Processing (ICIP), pp. 2775–2779. IEEE, Piscataway (2023)
- He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2961–2969. IEEE, Piscataway (2017)
- Bolya, D., Zhou, C., Xiao, F., Lee, Y.J.: Yolact: Real-time instance segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 9157–9166. IEEE, Piscataway (2019)
- Wang, X., Kong, T., Shen, C., Jiang, Y., Li, L.: Solo: Segmenting objects by locations. In: European Conference on Computer Vision, pp. 649–665. Springer, Berlin (2020)
- Tian, Z., Shen, C., Chen, H.: Conditional convolutions for instance segmentation. In: European Conference on Computer Vision, pp. 282–298. Springer, Berlin (2020)
- Wang, X., Zhang, R., Kong, T., Li, L., Shen, C.: Solov2: Dynamic and fast instance segmentation. *Adv. Neural Inf. Process. Syst.* **33**, 17721–17732 (2020)
- Fang, Y., Yang, S., Wang, X., Li, Y., Fang, C., Shan, Y., et al.: Instances as queries. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 6910–6919. IEEE, Piscataway (2021)
- Ke, L., Danelljan, M., Li, X., Tai, Y.W., Tang, C.K., Yu, F.: Mask transfiner for high-quality instance segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4412–4421. IEEE, Piscataway (2022)
- Zhu, C., Zhang, X., Li, Y., Qiu, L., Han, K., Han, X.: Sharpcontour: A contour-based boundary refinement approach for efficient and accurate instance segmentation. In: Proceedings of the IEEE/CVF Conference

- on Computer Vision and Pattern Recognition, pp. 4392–4401. IEEE, Piscataway (2022)
- 12 Yang, B., Wang, J., Clark, R., Hu, Q., Wang, S., Markham, A., et al.: Learning object bounding boxes for 3d instance segmentation on point clouds. In: *Advances in Neural Information Processing Systems*, vol. 32. MIT Press, Cambridge (2019)
 - 13 Yi, L., Zhao, W., Wang, H., Sung, M., Guibas, L.J.: Gspn: Generative shape proposal network for 3d instance segmentation in point cloud. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3947–3956. IEEE, Piscataway (2019)
 - 14 Liu, S.H., Yu, S.Y., Wu, S.C., Chen, H.T., Liu, T.L.: Learning gaussian instance segmentation in point clouds. *arXiv preprint, arXiv:200709860* (2020)
 - 15 Engelmann, F., Bokeloh, M., Fathi, A., Leibe, B., Nießner, M.: 3d-mpa: Multi-proposal aggregation for 3d semantic instance segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9031–9040. IEEE, Piscataway (2020)
 - 16 Sun, W., Rebain, D., Liao, R., Tankovich, V., Yazdani, S., Yi, K.M., et al.: Neuralbf: Neural bilateral filtering for top-down instance segmentation on point clouds. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 551–560. IEEE, Piscataway (2023)
 - 17 Wang, W., Yu, R., Huang, Q., Neumann, U.: Sgpn: Similarity group proposal network for 3d point cloud instance segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2569–2578. IEEE, Piscataway (2018)
 - 18 Wang, X., Liu, S., Shen, X., Shen, C., Jia, J.: Associatively segmenting instances and semantics in point clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4096–4105. IEEE, Piscataway (2019)
 - 19 Lahoud, J., Ghanem, B., Pollefeys, M., Oswald, M.R.: 3d instance segmentation via multi-task metric learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9256–9266. IEEE, Piscataway (2019)
 - 20 Pham, Q.H., Nguyen, T., Hua, B.S., Roig, G., Yeung, S.K.: Jsis3d: Joint semantic-instance segmentation of 3d point clouds with multi-task pointwise networks and multi-value conditional random fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8827–8836. IEEE, Piscataway (2019)
 - 21 Elich, C., Engelmann, F., Kontogianni, T., Leibe, B.: 3d bird's-eye-view instance segmentation. In: *German Conference on Pattern Recognition*, pp. 48–61. Springer, Berlin (2019)
 - 22 Jiang, L., Zhao, H., Shi, S., Liu, S., Fu, C.W., Jia, J.: Pointgroup: Dual-set point grouping for 3d instance segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4867–4876. IEEE, Piscataway (2020)
 - 23 Han, L., Zheng, T., Xu, L., Fang, L.: Occuseg: Occupancy-aware 3d instance segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2940–2949. IEEE, Piscataway (2020)
 - 24 He, T., Liu, Y., Shen, C., Wang, X., Sun, C.: Instance-aware embedding for point cloud instance segmentation. In: *European Conference on Computer Vision*, pp. 255–270. Springer, Berlin (2020)
 - 25 Zhao, L., Tao, W.: Jsnet: Joint instance and semantic segmentation of 3d point clouds. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 12951–12958. AAAI Press, Menlo Park, CA (2020)
 - 26 He, T., Shen, C., van den Hengel, A.: Dyco3d: Robust instance segmentation of 3d point clouds through dynamic convolution. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 354–363. IEEE, Piscataway (2021)
 - 27 Chen, S., Fang, J., Zhang, Q., Liu, W., Wang, X.: Hierarchical aggregation for 3d instance segmentation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15467–15476. IEEE, Piscataway (2021)
 - 28 Zhang, B., Wonka, P.: Point cloud instance segmentation using probabilistic embeddings. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8883–8892. IEEE, Piscataway (2021)
 - 29 Chen, F., Wu, F., Gao, G., Ji, Y., Xu, J., Jiang, G.P., et al.: Jspnet: Learning joint semantic & instance segmentation of point clouds via feature self-similarity and cross-task probability. *Pattern Recogn.* **122**, 108250 (2022)
 - 30 Schult, J., Engelmann, F., Hermans, A., Litany, O., Tang, S., Leibe, B.: Mask3d: Mask transformer for 3d semantic instance segmentation. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8216–8223. IEEE, Piscataway (2023)
 - 31 Wu, W., Qi, Z., Fuxin, L.: Pointconv: Deep convolutional networks on 3d point clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9621–9630. IEEE, Piscataway (2019)
 - 32 Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: *Advances in Neural Information Processing Systems*, vol. 30. MIT Press, Cambridge, MA (2017)
 - 33 Mo, K., Zhu, S., Chang, A.X., Yi, L., Tripathi, S., Guibas, L.J., et al.: Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 909–918. IEEE, Piscataway (2019)
 - 34 Armeni, I., Sener, O., Zamir, A.R., Jiang, H., Brilakis, I., Fischer, M., et al.: 3d semantic parsing of large-scale indoor spaces. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1534–1543. IEEE, Piscataway (2016)
 - 35 He, T., Gong, D., Tian, Z., Shen, C.: Learning and memorizing representative prototypes for 3d point cloud semantic and instance segmentation. In: *European Conference on Computer Vision*, pp. 564–580. Springer, Berlin (2020)
 - 36 Denis, L., Royen, R., Munteanu, A.: Improved block merging for 3d point cloud instance segmentation. In: *2023 24th International Conference on Digital Signal Processing (DSP)*, pp. 1–5. IEEE, Piscataway (2023)
 - 37 Vu, T., Kim, K., Luu, T.M., Nguyen, T., Yoo, C.D.: Softgroup for 3d instance segmentation on point clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2708–2717. IEEE, Piscataway (2022)
 - 38 Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 652–660. IEEE, Piscataway (2017)
 - 39 Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16259–16268. IEEE, Piscataway (2021)
 - 40 Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)