

Pix2Point: Learning Outdoor 3D Using Sparse Point Clouds and Optimal Transport

R. Leroy,
P. Trouvé-Peloux, F. Champagnat
ONERA, DTIS
F-91123 Palaiseau, France
remy.leroy@onera.fr

B. Le Saux
ESA/ESRIN Φ -lab
Frascati (RM), Italy
bertrand.le.saux@esa.int

M. Carvalho
UPCITY
Montreuil, France
marcela.carvalho@upciti.com

Abstract

Good quality reconstruction and comprehension of a scene rely on 3D estimation methods. The 3D information was usually obtained from images by stereo-photogrammetry, but deep learning has recently provided us with excellent results for monocular depth estimation. Building up a sufficiently large and rich training dataset to achieve these results requires onerous processing. In this paper, we address the problem of learning outdoor 3D point cloud from monocular data using a sparse ground-truth dataset. We propose Pix2Point, a deep learning-based approach for monocular 3D point cloud prediction, able to deal with complete and challenging outdoor scenes. Our method relies on a 2D-3D hybrid neural network architecture, and a supervised end-to-end minimisation of an optimal transport divergence between point clouds. We show that, when trained on sparse point clouds, our simple promising approach achieves a better coverage of 3D outdoor scenes than efficient monocular depth methods.

1 Introduction

Geometry estimation is a prior requirement for autonomous agents to comprehend their environment, progress and interact within it. This fundamental task leads to various estimation methods for 3D reconstruction. For long, the parallax between images has been used [10, 14], either between two cameras in stereo mode or between multiple acquisitions after displacement. Recently, deep learning techniques have revolutionised 3D estimation from images, allowing even to obtain excellent results with a single view [1, 5, 8, 12, 15]. These impressive results rely upon large and highly accurate databases like KITTI [13] that involve the simultaneous collection of stereo pairs and LiDAR data, post-processing and temporal integration in order to provide accurate, reliable and dense ground-truth for learning purposes. The overall process requires large scale cooperation and is therefore lengthy. In this paper, we address the question of performing 3D estimation from monocular data with much lighter requirements in terms of training data input. Besides, if monocular point sets prediction has been

proposed in the literature, it is only for the reconstruction of a single 3D object [9, 16, 18].

In contrast, we propose here (i) one of the first approaches to reconstruct a 3D point cloud for an entire outdoor scene given only a single image using a 2D-3D hybrid neural network architecture, illustrated in Fig. 1; (ii) an end-to-end learning scheme of the hybrid model using a sparse point clouds dataset; (iii) a cost function from the optimal transport (OT) that makes possible to obtain good coverage of the scenes on the KITTI dataset [13] (iv) better performances than state-of-the-art monocular depth map prediction methods when trained with sparse data. This illustrates the advantage of the direct prediction of point clouds, in the case of a low-density ground-truth 3D dataset.

2 Related Work

The monocular 3D estimation task has been addressed in terms of depth maps prediction with the following works [2, 4, 8, 12, 15, 25, 21] but as already mentioned, they are trained with large and rich datasets. Recently, 3D estimation has also been addressed in terms of 3D point set with PSGN by Fan *et al.* [9], a method that aims to predict the envelope of an object as an unordered set of points using a single view of it and its location in the image. Mandikal and Babu [16] address the limitations of PSGN regarding the poor number of predicted points with DensePCR, a pyramidal structure allowing to enlarge the number of points. Xia *et al.* [24], also tackles the generation of a monocular point cloud for objects using prior knowledge over their shapes, making it robust to occlusions and varying poses. A generative flow-based model allowing single view object point cloud prediction has recently been proposed with C-flow [18]. It leverages a back-and-forth prediction loop from image to point cloud, then to image for consistency.

The aforementioned point cloud works only consider the reconstruction of a single 3D object model, *i.e.*, on data that are obtained through demanding procedures, either scanned objects using RGB-D sensors or laser scanners, or handcrafted models. These procedures do not apply to real-life outdoor scenes with various settings and where lies multiple objects.




Figure 1. Pix2Point: 3D point set prediction for real outdoor scenes from a single image. A first 2D CNN module encodes the RGB image for the following fully connected layer to predict a coarse 3D point cloud. Then a 3D point-wise densification module ensures to lift the number of points of this point cloud using PointNet-like MLP. All models are trained end-to-end to minimise point set distances, *e.g.*, Optimal Transport or chamfer distances.

[7] tackles the problem of monocular volumetric reconstruction with occlusion completion for indoor scenes. In addition to the input RGB image, this approach requires a corresponding normal image that is hardly obtainable for outdoor scenes. To our knowledge, we are the first to propose a deep learning method for the reconstruction of complex outdoor scenes in the form of 3D point cloud, solely conditioned by a single RGB image, and learned on sparse point clouds. Our technical contributions with respect to previous works are: a performance study of the 2D-3D hybrid neural network architecture for point cloud reconstruction of scenes “in-the-wild”, end-to-end training with OT loss optimisation on sparse point clouds up to ten thousands elements. Finally, we compare our point set prediction approach to state of the art depth map prediction approaches, when they are all trained on the same sparse dataset.

3 Method

Given a single colour image, our method, namely Pix2Point, predicts a set of 3D point coordinates with an arbitrary number of elements fixed before training. Fig. 1 shows an overview of the proposed method. Like DensePCR, our architecture consists of an encoding-decoding module to predict a first coarse point cloud that will be enriched using a densification module. Our model, unlike DensePCR, is trained to minimise an optimal transport divergence over the point coordinates in an end-to-end fashion.

Encoding The encoding block is a series of convolution, pooling and normalisation layers to extract a feature description of the full RGB input image, which is then processed by a fully connected layer to obtain a first coarse set of 3D point coordinates. We explore and compare several encoding approaches following either VGG, DenseNet and ResNet architectures in section 5. We refer to them as *backbones*.

Decoding We refer to decoding as the densification of the first coarse 3D point cloud. We duplicate every point n times, and to describe each point we concatenate:

the 3D coordinates, both global point cloud feature vector and the corresponding local feature vector, obtained using dedicated PointNet-like shared multi-layered perceptron (MLP) [19, 20], and lastly a grid alignment feature vector in order to identify every clone of the same point and to suggest geometric information between every clone. This point description is processed by another shared MLP resulting in 3D coordinates for 1 point.

Optimisation The performance of our approach mainly depends on the training loss function. Unlike the depth map prediction methods which exploit the gridded structure of the image for the evaluation of errors, our method uses distances between unordered point sets. In particular, it requires an additional, computationally expensive step to match points from the predicted and the target point clouds. We expose two usual distances for this task.

Chamfer distance The chamfer distance is the average of squared Euclidean distances to the nearest neighbour from one set to the other.

Optimal transport or OT distance To compute this distance, also known as Earth Mover’s distance, one has to find a one-to-one mapping from one set to the other that minimises the sum over each point of the squared distance between them and their corresponding image. This minimised sum is called the OT distance and it informs about the eventual discrepancy between point sets distributions.

The exact computation of an OT distance is time and memory expensive especially for several thousand elements, hence, we consider in our work an approximation of the OT distance obtained by adding a regularising term and solving the Sinkhorn-Knopp algorithm [6, 11].

4 Experimental Settings

This section presents the dataset we are considering for real scene point cloud estimation from a single image and defines several criteria for reconstruction performance evaluation.

4.1 Dataset

To assess our method, we operate on RGB image sequences of real urban scenes and corresponding LiDAR point cloud acquisitions from the KITTI depth estimation benchmark dataset [13]. Every scene point cloud is an accumulation of filtered LiDAR acquisitions over few successive time instants. We use the split defined by Eigen *et al.* [8], that is 22 600 training scenes and 697 testing scenes.

4.2 Evaluation Metrics

There is no reference value or method for the KITTI scene reconstruction task regarding 3D point clouds. Therefore, we propose to measure performances using the completeness and accuracy criteria from [22], and also propose the relative accuracy. All 3 measures are defined as follows:

Completeness is the coverage in per cent of the target point cloud by the predicted points. A target point is covered if a predicted point lies in its surrounding (*i.e.* fixed radius ball). We evaluate completeness values for radius of 50 cm, 25 cm and 10 cm.

Accuracy is the distance d , in meter, from the r -th percentile of the distances to the nearest neighbour, from the predicted point cloud to the ground-truth point cloud. It measures the greatest distance to the nearest neighbour among the predicted points closest to the ground truth. We choose $r < 90\%$ to include most of the points and discard eventual outliers.

Relative accuracy is similar to the accuracy, where every distance to their nearest neighbour is divided by the norm of the corresponding target point. It provides a higher penalty to short-range predictions.

4.3 Implementation Details

Experiments were conducted using the *Pytorch* Framework [17]. We kept the original image resolution and cropped every picture to 1224×370 pixel definition. Due to heavy computational cost for loss back-propagation, parameters were updated after every sample forward, making batch normalisation ineffective. Instance normalisation was applied instead [23]. The number of predicted points was determined to fully load the 8GB GPU during training. Therefore, 2500 elements point clouds are first predicted by the fully connected encoding module, that are then up-scaled by a DensePCR-like module making a point cloud with 10k elements.

Training vs. Testing point clouds: Using an OT loss enables, in principle, the comparison of any ground-truth point cloud to the predicted one, however, in practice, computation and optimisation of an OT loss are computationally much more efficient with point sets of equal cardinality. Therefore, ground-truth point cloud databases are randomly subsampled to 10k

points, that is as many points as Pix2Point predicts. When testing we measure performances to the whole ground-truth point cloud.

5 Experimental Results

In this section we first present the performances of Pix2Point using various encoding backbones and losses, then we compare our method to depth map prediction approaches through evaluation metrics defined in 4.2.

5.1 Network parameter study

We trained several models with varying encoding backbones and loss functions. We considered the following configurations: Pix2Point architecture with VGG backbone and training on the minimisation end-to-end of either the chamfer or OT distance, and Pix2Point with ResNet backbone and minimising the OT distance. The performances of these models are given in Table 1 respectively as P2P-VGG-C, P2P-VGG-OT and P2P-ResNet-OT. From these figures, we can notice that the minimisation of chamfer distance thrives toward predictions with low local error, and minimising the OT distance grants predictions with higher completeness, hence, better coverage of the scenes. In order to find if these distances could help each other, combinations of both distances have been tested. However, they lead to convergence issues during training and overall worst performances due to opposite objectives of the distances. Changing the backbone from VGG to ResNet has also a slight impact on the completeness and accuracy. The small gain in relative accuracy indicates that far predictions are more accurate. We also provide a comparison to a similar image-to-point-cloud approach, DensePCR [16], initially proposed for 3D graphics models.

5.2 Comparison to depth prediction approaches

The current dominant approach to 3D estimation from a single image consists of predicting corresponding depth maps by leveraging the power of image-to-image translation networks. On KITTI, these methods are trained on pseudo-dense depth maps built by accumulating several consecutive LiDAR acquisitions. For comparison in a similar setting, we train two state-of-the-art models for monocular depth estimation from the KITTI challenge¹, namely BTS [15] and AdaBins [2], on the same 10k-point-cloud as Pix2Point. At inference time, dense depth maps are projected back to 3D using known camera parameters. Performance comparison with various flavours of Pix2Point is reported in Table 1.

These results show that Pix2Point, with only 15M parameters, trained with Chamfer distance performs

¹http://www.cvlibs.net/datasets/kitti/eval_depth.php?benchmark=depth_prediction

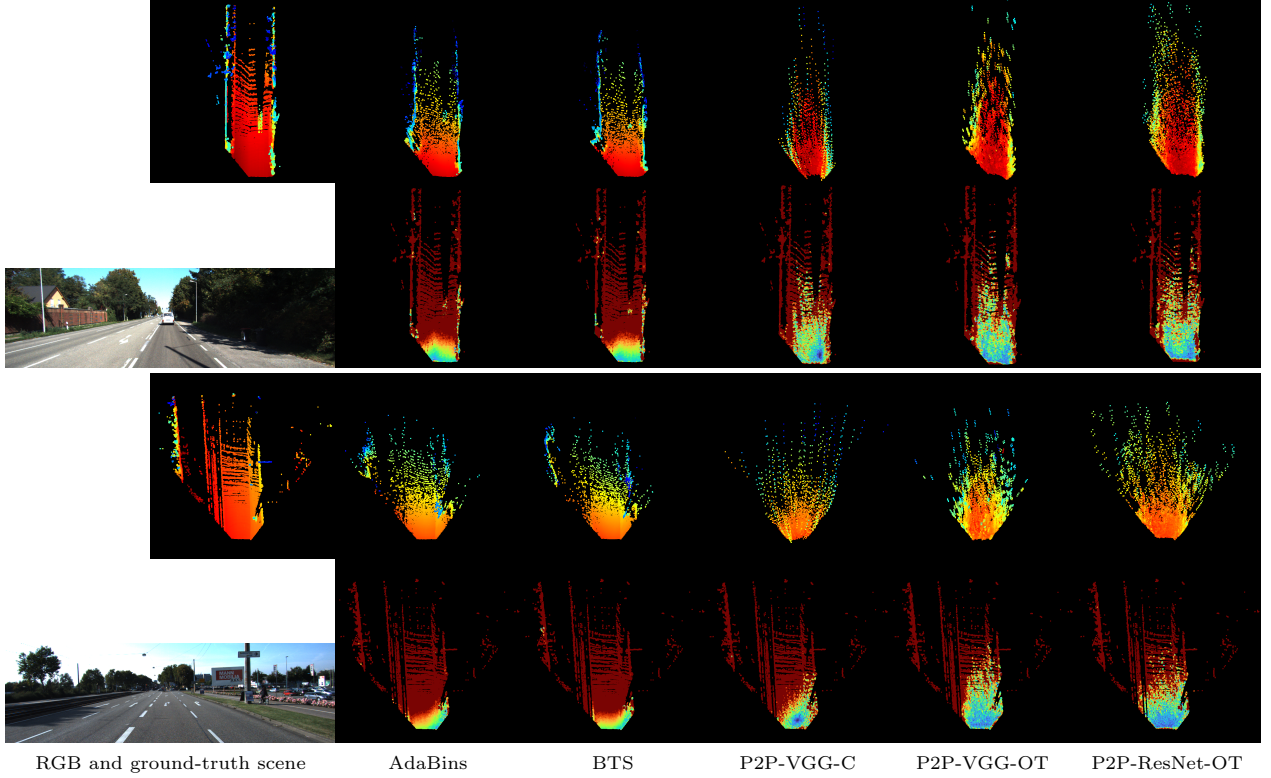


Figure 2. For each scene, first row: 3D ground-truth and predictions for the RGB image according to AdaBins [2], BTS [15], our Pix2Point VGG-chamfer, VGG-OT and ResNet-OT, all trained on 10k points. We follow the 3D representation of [3]. Bird’s eye view where the colour encodes the altitude. Second row: the input RGB image and the ground-truth-to-prediction error map for each method. Errors are from 0 (blue) to 50cm (red).

better than BTS and Adabins, respectively 45M and 78M parameters. Moreover, when trained with the OT distance, Pix2Point accuracy decreases but it covers three times more points than depth map approaches for the closest neighbourhoods. These observations can be made through Figure 2 where we show point cloud predictions and the coverage error map for each method (for comparison all predictions are visualized with 10k points). This error map displays for each ground-truth point the distance to its nearest predicted point. We display the error from 0 to 50cm using the jet colormap. While AdaBins and BTS preserve fine features, all Pix2Point variants achieve better coverage of the scene and a lower error, especially for far-away elements and small size features, that are not retrieved by AdaBins and BTS (see for instance the right part of the bottom scene).

6 Conclusion

We proposed in this work an innovative approach to tackle point clouds reconstruction for complex outdoor scenes from a unique RGB image, using a light 2D-3D hybrid neural network. It recovers properly distributed

Table 1. Comparison of 3D scene reconstructions on KITTI. We report completeness and accuracy. All methods are trained with 10k point clouds.

Approaches	Completeness \uparrow (in %)			Accuracy \downarrow	
	50cm	25cm	10cm	in m	rel.
P2P-ResNet-OT	71.35	48.82	15.12	1.92	0.18
P2P-VGG-OT	67.4	47.7	14.7	1.79	0.19
P2P-VGG-C	64.4	36.0	8.0	0.85	0.05
DensePCR	59.9	23.5	3.5	1.77	0.18
BTS	67.59	31.29	6.28	1.23	0.06
AdaBins	65.86	27.52	5.71	1.25	0.06

point clouds by taking advantage of an optimal transport loss. We also provided the first benchmark for this novel task on the KITTI dataset and introduced performance metrics to assess the quality of point cloud reconstruction. We show that our method outperforms state of the art depth map prediction methods, when trained with sparse data, illustrating the interest of direct image to 3D point-cloud translation.

References

- [1] Ali Jahani Amiri, Shing Yan Loo, and Hong Zhang. Semi-supervised monocular depth estimation with left-right consistency using deep neural network. In *2019 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2019.
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [3] Lucas Caccia, Herke van Hoof, Aaron Courville, and Joelle Pineau. Deep generative modeling of LiDAR data. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [4] Marcela Carvalho, Bertrand Le Saux, Pauline Trouvé-Peloux, Andrés Almansa, and Frédéric Champagnat. Deep Depth from Defocus: How can defocus blur improve 3D estimation using dense neural networks? In *European Conf. on Computer Vision (ECCV)*, 2018.
- [5] Marcela Carvalho, Bertrand Le Saux, Pauline Trouvé-Peloux, Frédéric Champagnat, and Andrés Almansa. On regression losses for deep depth estimation. In *IEEE Int. Conf. on Image Processing (ICIP)*, 2018.
- [6] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NIPS*, 2013.
- [7] Maximilian Denninger and Rudolph Triebel. 3d scene reconstruction from a single viewport. In *European Conference on Computer Vision*, pages 51–67. Springer, 2020.
- [8] David Eigen, Christian Puhersch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in Neural Information Processing Systems 27*. 2014.
- [9] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2017.
- [10] Olivier Faugeras. *Three-Dimensional Computer Vision: A Geometric Viewpoint*. 1993.
- [11] Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trouvé, and Gabriel Peyré. Interpolating between optimal transport and MMD using sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 2019.
- [12] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *International Journal of Robotics Research (IJRR)*, 2013.
- [14] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. 2004.
- [15] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From Big to Small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019.
- [16] Priyanka Mandikal and Venkatesh Babu Radhakrishnan. Dense 3D point cloud reconstruction using a deep pyramid network. In *2019 Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2019.
- [17] Adam Paszke, Sam Gross, Francisco Massa, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*. 2019.
- [18] Albert Pumarola, Stefan Popov, Francesc Moreno-Noguer, and Vittorio Ferrari. C-flow: Conditional generative flow models for images and 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [19] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2017.
- [20] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems 30*, 2017.
- [21] Ashutosh Saxena, Sung H. Chung, and Andrew Y. Ng. Learning depth from single monocular images. In *Adv. in Neural Information Processing Systems 18*. 2006.
- [22] Radim Tylecek, Torsten Sattler, Hoang-An Le, Thomas Brox, Marc Pollefeys, Robert B. Fisher, and Theo Gevers. The second workshop on 3D Reconstruction Meets Semantics: Challenge results discussion. In *ECCV 2018 Workshops*, 2019.
- [23] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- [24] Yan Xia, Yang Zhang, Dingfu Zhou, Xinyu Huang, Cheng Wang, and Ruigang Yang. Realpoint3d: Point cloud generation from a single image with complex background. *arXiv preprint arXiv:1809.02743*, 2018.
- [25] Minghan Zhu, Maani Ghaffari, Yuanxin Zhong, Pingping Lu, Zhong Cao, Ryan M Eustice, and Hui Peng. Monocular depth prediction through continuous 3D loss. *arXiv preprint arXiv:2003.09763*, 2020.