



HAL
open science

PanoRoom: From the Sphere to the 3D Layout

Clara Fernandez-Labrador, José M Fácil, Alejandro Perez-Yus, Cédric
Demonceaux, Jose J Guerrero

► **To cite this version:**

Clara Fernandez-Labrador, José M Fácil, Alejandro Perez-Yus, Cédric Demonceaux, Jose J Guerrero.
PanoRoom: From the Sphere to the 3D Layout. ECCV 2018 Workshops, Sep 2018, Munich, Germany.
hal-02158448

HAL Id: hal-02158448

<https://hal.science/hal-02158448>

Submitted on 18 Jun 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PanoRoom: From the Sphere to the 3D Layout

Clara Fernandez-Labrador^{1,2}, José M. FÁCIL¹, Alejandro Perez-Yus¹, Cédric Démonceaux², and Jose J. Guerrero¹

¹Instituto de Investigación en Ingeniería de Aragón (I3A), Universidad de Zaragoza, Spain. {cfernandez, jmfacil, alperetz, josechu.guerrero} @unizar.es.

²Le2i VIBOT ERL-CNRS 6000, Université de Bourgogne Franche-Comté, France {cedric.demonceaux@ubfc.fr}

Abstract. We propose a novel FCN able to work with omnidirectional images that outputs accurate probability maps representing the main structure of indoor scenes, which is able to generalize on different data. Our approach handles occlusions and recovers complex shaped rooms more faithful to the actual shape of the real scenes. We outperform the state of the art not only in accuracy of the 3D models but also in speed.

1 Introduction

The problem of recovering the 3D layout of a cluttered indoor scene goes back to the early days of computer vision and still is a core research topic, as it is a key technology in several emerging application markets like augmented and virtual reality, indoor navigation, SLAM and robotics in general. While witnessing the rapid progress on layout recovery methods from perspective images with both geometry and deep learning techniques, the expansion to omnidirectional vision is yet limited. Panoramic images have broken the barriers of performance on this task. PanoContext [1] was the first to extend the frameworks designed for perspective images to panoramas. They recover both the layout, assumed as a 3D box (4 walls), and bounding boxes of the main objects inside the room. With the motivation to leave behind the assumption of simple boxes, [2] generates layout hypotheses by geometric reasoning from structural corners obtained through geometry and deep learning combination. Most recent work is LayoutNet [3], which generates 3D layout models by training a network with panoramas and vanishing lines to obtain edge and corner maps.

Here, we propose a Fully Convolutional Neural Network (FCN) that faces the 3D layout recovery problem from panoramas with the following contributions: first, we introduce a fully convolutional architecture using ResNet-50 [4], pretrained on ImageNet, as backbone. We include a single decoder that jointly predicts edge and corner maps, which requires fewer parameters and computing time than the state of the art while gaining accuracy. Notwithstanding that the great majority of the images are labeled as 4-walls rooms, our network predicts room solutions without assuming box-type rooms any more, from which we produce geometrically consistent room layouts. Our 3D models outperform existing methods on standard benchmark datasets [5,6].

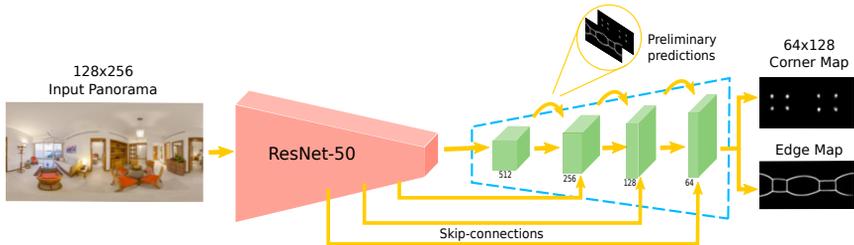


Fig. 1: PanoRoom architecture: from ResNet-50, we build a single decoder that jointly predicts edge and corner maps already refined

2 Deep learning structure perception

The proposed FCN follows the encoder-decoder structure and builds upon ResNet-50 [4]. We replace the final fully-connected layer with a decoder that jointly predicts layout edges and corners locations already refined. We illustrate the proposed architecture in Fig. 1. Our ground truth (GT) for every panorama consists of two probability maps. One represents the room edges whereas the other encodes the corner locations. Every pixel has a value of 1 if it contains structural information or 0 if it is background. Line thickening and Gaussian blur are both employed to make the model training easier to converge since the natural distribution of pixels is unbalanced, *e.g.* 95% is background. Instead, similar approaches usually need to divide the GT into different classes, [7,8,9].

Encoder. Most of deep-learning approaches facing layout recovery problem have made use of the VGG16 as encoder [7,9]. Instead, [8] builds their model over ResNet-101 [4] outperforming the state of the art. Here, we propose ResNet-50 [4], pre-trained on the ImageNet dataset, which leads to a faster convergence due to the general low-level features learned from ImageNet. Residual networks allow us to increase depth without increasing the number of parameters with respect to their plain counterparts. This leads, in ResNet-50, to capture a receptive field of 483×483 , enough for our input resolution 128×256 .

Decoder. Most recent work [7,3] build two output branches for multi-task learning which leads to more computation time and more parameters. We instead propose a unique branch whose output has two channels, corners and edges maps, which helps to reinforce the quality of both map types. In the decoder, we combine two different ideas. First, skip-connections [10] from the encoder to the decoder. Specifically, we concatenate ‘up-convolved’ features with their corresponding features from the contracting part. Second, we perform preliminary predictions in different resolutions which we also concatenate and feed back to the network following the spirit of [11], see Fig. 1.

Loss functions. Edge and corner maps are learned through a pixel-wise sigmoid cross-entropy loss function. Since we know a priori that these maps have an extremely unbalanced distribution of edge and corner labels, we introduce the ponder factors λ_1 and λ_0 . Where 1 and 0 are the positive and negative classes respectively and $\lambda_c = \frac{N}{N_c}$, being N the total number of pixels and N_c amount

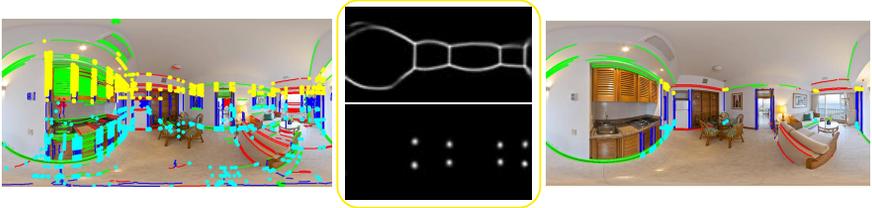


Fig. 2: **Left:** Vanishing lines and corners obtained only with geometric reasoning. **Center:** Edge and Corner Maps predicted by our FCN. **Right:** Vanishing lines and corners after combining both approaches.

of pixels of class c . The total sigmoid cross-entropy loss is the mean of the per-pixel loss: $\mathcal{L}_{M_i} = \lambda_1(y_i(-\log(S(\hat{y}_i)))) + \lambda_0((1-y_i)(-\log(1-S(\hat{y}_i))))$ for every pixel i , where y is the GT, \hat{y} is the predicted map and S is the sigmoid function. Inspired by [12], we also define a perceptual loss function that measure high-level perceptual differences between images. We make use of an auto-encoder with same structure as the main network and we train this auto-encoder to encode the GT. Apart from encouraging the output images, \hat{y} , to exactly match the target images, y , we also encourage them to have similar feature representations. Hence, having y as GT and \mathcal{I} as input image $\mathcal{L}_P(y, \mathcal{I}) = \|\hat{\phi}_j(\mathcal{I}) - \phi_j(y)\|_2^2$, where ϕ_j is the feature map on the j^{th} hidden layer for the auto-encoder and $\hat{\phi}_j$ on our network.

3 Layout recovery

While deep-learning approaches have shown tremendous success and provide a deeper understanding of the scene, their output alone is insufficient as it does not enforce geometric constraints and priors. For this reason, we take advantage of our FCN output to produce geometrically consistent room layouts by optimizing over the deep learning clues under Manhattan World assumption, whereby there exist three orthogonal main directions that define the scene. Here, we do not assume the scene as a 4 wall box any more. Lines and vanishing points (VP) extraction in perspective images has been satisfactorily faced so that many proposals working with panoramas [1,13] sample perspective subviews to make use of them, with the subsequent increase of computation time. Recently, some approaches [14,2] proposed methods to obtain lines and VP directly on panoramas improving the overall efficiency of the method. Here, we use the RANSAC-approach of [2] that has demonstrate also to be faster than other methods. Each extracted line is associated to a probability given by the sum of probabilities of the pixels it occupies in the edge map. This allows to remove those lines with 0 probability, leading to an optimal subset of accurate lines that allow our proposal to obtain final results with very few layout hypotheses generated. To make the most of our network output, we intersect the structural lines to obtain candidate corners that are scored with the corner map. See Fig. 2. The layout generation process

follows the idea of [2] but we use a different approximation to select the best hypothesis solution. Let’s define each layout hypothesis as a combination of its edges and corners retrieved, $L^h(L_E^h, L_C^h)$. We associate them with the probability sum of the pixels they occupy in the corresponding predicted edge and corner maps, we define the per-pixel probability as P_{edge} and P_{corner} respectively. In this manner, we select the layout that maximizes the match between these two sources of information, $L^{Best} = \arg \max(w_e \sum P_{edge}(L_E^h) + w_c \sum P_{corner}(L_C^h))$, where w_e and w_c are the term weights used to give them same importance.

4 Experiments

Experiments of the proposed approach are conducted within two public datasets that comprise several indoor scenes, SUN360 [5] and Stanford (2D-3D-S) [6]. We take advantage of the ~ 500 panoramas from the SUN360 dataset labeled by [1] but, since panoramas were all labeled as box-type rooms, we hand label and substitute 35 panoramas representing more faithfully the actual shapes of the rooms. We split the raw dataset consisting in 85% training and 15% test scenes. For experiments with the Stanford 2D-3D dataset, we use same testing images (area 5) and GT provided by [3].

The input to the network is a single panoramic image of resolution 128×256 , unlike [3] that uses also vanishing lines as input. The outputs, edge and corner maps, have resolution 64×128 . We apply horizontal mirroring as well as horizontal rotation from 0° to 360° of input images during training as data augmentation. We minimized the cross-entropy loss using Adam, regularized by penalizing the loss with the sum of the L2 of all weights. Initial learning rate is $2.5e^{-4}$ and exponentially decayed by a rate of 0.995 every epoch. We apply 0.3 dropout rate and $5e-6$ weight decay. Although we label some panoramas accurately to their actual shape, we still have a big unbalanced dataset. In order to overcome this problem, we choose a batch size of 16 and we force it always to include one example between those panoramas hand labeled by us (not box-type). This favors the learning of more complex rooms despite having few examples. The network is implemented in TensorFlow and trained with NVIDIA Geforce GTX 1080.

FCN Evaluation. After training our network on the SUN360 dataset, we first carried out an experiment to choose more precisely the weights leading to the best performance. To do that, we evaluate the predicted maps with their corresponding GT at different stages of the training on the same dataset. We saw that with 200 epochs (1 epoch = 25 iterations), our network reaches the best performance with a training time of ~ 1 hour and a half. Testing takes about 0.7s per image. These results are collected in Table 1, first row. Additionally, we find interesting to also evaluate results on the Stanford 2D-3D dataset to see how well our FCN is able to generalize, this results are shown in the second row. As it was expected, results are not as good when testing with a different dataset. However, we demonstrate in next experiments that it is enough to get layout reconstructions with little error and thus, our network is able to generalize.

In Fig. 3 we show 3 examples of our predicted maps with different number of

Table 1: Evaluation of our FCN trained on SUN360 dataset for 200 epochs

Dataset	Precision(%)	Recall(%)	F1 Score(%)	Accuracy(%)
SUN360	0.8338	0.8564	0.8449	0.9615
Stnfd.2D-3D	0.6309	0.5904	0.6113	0.9160

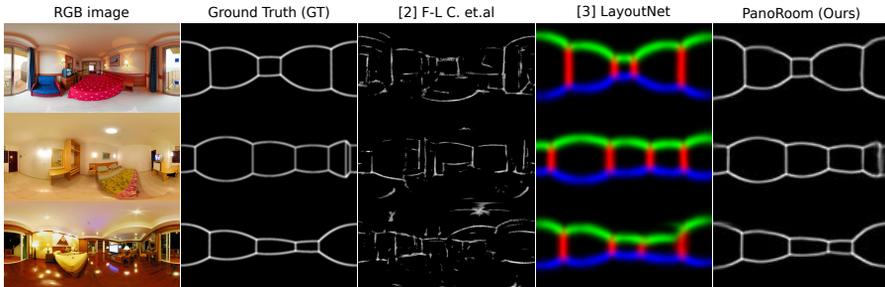


Fig. 3: **Predicted Edge Maps.** Our FCN outperforms last works from the state of the art [2,3]. We predict cleaner edges around the boundaries and recover faithful edge maps even when indoor scenes are not simple 4-walls rooms

Table 2: Performance benchmarking for SUN360 and Stanford 2D-3D datasets training on SUN360 data. *SS*: Simple Segmentation (3 categories): ceiling, floor and walls [3]. *CS*: Complete Segmentation: ceiling, floor, wall₁,..., wall_n [2].

Dataset	Method	3DIoU(%)	CE(%)	PE ^{SS} (%)	PE ^{CS} (%)
SUN360	PanoContext [1]	67.22	1.60	4.55	10.34
	F-L C. <i>et al.</i> [2]	-	-	-	7.26
	LayoutNet [3]	74.48	1.06	3.34	-
	Ours	76.82	0.79	2.59	3.13
Stnfd.2D-3D	F-L C. <i>et al.</i> [2]	-	-	-	12.1
	Ours	70.64	1.15	3.95	4.98

walls compared to the GT and to other two approaches [2,3]. Our proposal is able to directly handle network outputs not limited to 4-wall rooms. Here we demonstrate that is possible to train strategically in a way that the network takes full advantage of the few different data that we have at our disposal.

3D Layout Evaluation. We evaluate our layout recovering approach on three standard metrics, 3D intersection over union (3DIoU), corner error (CE) and pixel error (PE), and compare ourselves with three works from the state of the art [1,3,2]. Results are averaged across all images. For all experiments, only SUN360 dataset is used for training. Table 2 shows the performance of our proposal testing on both datasets. We show that results on SUN360 dataset demonstrate better performance as the FCN has been trained on the same dataset, however results on Stanford 2D-3D dataset are also very competitive.

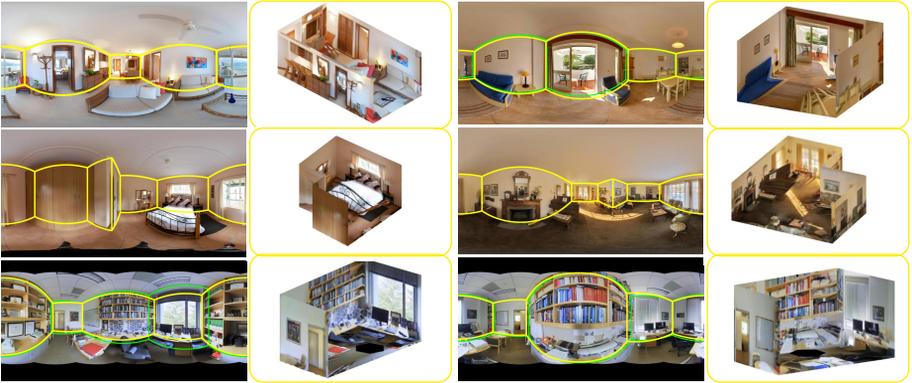


Fig. 4: Qualitative results on both datasets. Yellow: PanoRoom, green: GT.

References

1. Zhang, Y., Song, S., Tan, P., Xiao, J.: PanoContext: A whole-room 3D context model for panoramic scene understanding. In: IEEE ECCV. (2014) 668–686
2. Fernandez-Labrador, C., Perez-Yus, A., Lopez-Nicolas, G., Guerrero, J.J.: Layouts from panoramic images with geometry and deep learning. arXiv:1806.08294 (2018)
3. Zou, C., Colburn, A., Shan, Q., Hoiem, D.: Layoutnet: Reconstructing the 3d room layout from a single rgb image. arXiv:1803.08999 (2018)
4. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE CVPR. (2016) 770–778
5. Xiao, J., Ehinger, K., Oliva, A., Torralba, A.: Recognizing scene viewpoint using panoramic place representation. In: IEEE CVPR. (2012) 2695–2702
6. Armeni, I., Sax, A., Zamir, A.R., Savarese, S.: Joint 2D-3D-Semantic Data for Indoor Scene Understanding. arXiv:1702.01105 (2017)
7. Mallya, A., Lazebnik, S.: Learning informative edge maps for indoor scene layout prediction. In: IEEE ICCV. (2015) 936–944
8. Zhao, H., Lu, M., Yao, A., Guo, Y., Chen, Y., Zhang, L.: Physics inspired optimization on semantic transfer features: An alternative method for room layout estimation. arXiv:1707.00383 (2017)
9. Lee, C., Badrinarayanan, V., Malisiewicz, T., Rabinovich, A.: RoomNet: End-to-end room layout estimation. In: IEEE ICCV. (2017)
10. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI. (2015) 234–241
11. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., van der Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. In: IEEE ICCV. (2015) 2758–2766
12. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV. (2016) 694–711
13. Yang, H., Zhang, H.: Efficient 3d room shape recovery from a single panorama. In: IEEE CVPR. (2016) 5422–5430
14. Bazin, J.C., Seo, Y., Demonceaux, C., Vasseur, P., Ikeuchi, K., Kweon, I., Pollefeys, M.: Globally optimal line clustering and vanishing point estimation in manhattan world. In: IEE CVPR. (2012) 638–645