This CVPR 2021 workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Scaled 360 layouts: Revisiting non-central panoramas

Bruno Berenguel-Baeta, Jesus Bermudez-Cameo, Jose J. Guerrero I3A, University of Zaragoza, Spain

berenguel@unizar.es, bermudez@unizar.es, josechu.guerrero@unizar.es

Abstract

From a non-central panorama, 3D lines can be recovered by geometric reasoning. However, their sensitivity to noise and the complex geometric modeling required has led these panoramas being very little investigated. In this work we present a novel approach for 3D layout recovery of indoor environments using single non-central panoramas. We obtain the boundaries of the structural lines of the room from a non-central panorama using deep learning and exploit the properties of non-central projection systems in a new geometrical processing to recover the scaled layout. We solve the problem for Manhattan environments, handling occlusions, and also for Atlanta environments in an unified method. The experiments performed improve the state-of-the-art methods for 3D layout recovery from a single panorama. Our approach is the first work using deep learning with non-central panoramas and recovering the scale of single panorama layouts.

1. Introduction

Layout recovery and 3D understanding of indoor environments is a hot topic in computer vision research [11]. Most recent approaches for layout recovery use different neural network architectures to recover the structural elements of an indoor environment. In this context, the use of omnidirectional images has important advantages to retrieve the shape of a whole room from a single image. In order to handle the heavy distortions that introduce the omnidirectional representations, we can find different approaches in the state of the art. Dula-Net [9] extracts from the equirectangular image a perspective view of the ceiling and extrude the layout using a dual-branched architecture. Similarly but with a different network architecture, AtlantaNet [6] obtains the floor plan dividing the panorama into two perspective views of ceiling and floor separately to adjust the height of the room. On a different approach,



Figure 1: Central (up-left) and non-central panoramas (bottom-left) have similar appearance but there are subtle differences in favor of the second if we want to obtain 3D information including the scale. On the right, the scaled layout obtained from a non-central panorama in an Atlanta environment with our solution.

Corners for Layouts (CFL) [2] and HorizonNet [8] aim to extract the boundaries of the structural lines and corners of the room directly from the equirectangular panorama, obtaining an up-to-scale layout in a post-processing.

In this paper we propose a new method for layout recovery from single panoramas. We revisit the non-central panoramas [4] in order to obtain scaled 3D lines [1] and vertical planes from a single image. As the equirectangular panorama, the non-central circular panorama provides 360 information of the environment, but the image distortion in non-central panoramas includes subtle differences allowing geometric 3D reasoning [5]. This characteristic is a clear advantage with regard equirectangular panoramas since allows to recover the scale of the environment without any prior knowledge. In Fig. 1 we have an equirectangular panorama and a non-central panorama of the same scene. Even though both images look similar allowing to recover the shape of the room, only from the non-central circular panorama we can recover the scale without any assumption.

In our proposal, we adapt the neural network architecture of HorizonNet [8] to non-central circular panoramas for the extraction of structural lines of indoor environments. Besides, we propose a geometric pipeline with two new linear

Work supported by Project RTI2018-096903-B-I00 (AEI/FEDER, UE).



Figure 2: Pipeline of the proposed method. The non-central circular panorama is processed by the fine-tuned network. The network provides the pixel information of the structural lines and a per-column probability of a wall-wall intersection. Then the proposed geometric pipeline, including the new solvers, gives the final scaled layout.

solvers that jointly obtain the room height and vertical walls location for global layout extraction in Manhattan and Atlanta world assumptions. Our experiments show that our proposal improves the state of the art solutions, being ours the first that obtains the scale of the layout without using additional assumptions.

2. Proposed method

In order to recover the layout of a room from a single non-central circular panorama, we propose a new pipeline, shown in Fig. 2. We use a neural network as a line extractor from the image to reconstruct the scaled layout in a new geometrical processing that includes two new linear solvers for non-central projection systems.

2.1. Neural network as line extractor

For the first part of our pipeline, we propose a neural network as boundary extractor for structural lines. In classical approaches, lines are extracted with Hough transform and vanishing points, followed by an hypothesis generationverification algorithm [10]. This approach consumes lots of time and resources. On the other hand, neural networks have proven that can obtain patterns on images with high accuracy in a short time and, therefore, current approaches for layout recovery rely on the use of neural networks.

Even though many state-of-the-art networks can handle omnidirectional panoramas, there are not any that have considered non-central systems. We propose to adapt the existing network architecture of HorizonNet [8] to handle noncentral panoramas. The main advantage of this architecture is that handles the information of the panorama column by column. This is particularly interesting for the non-central circular panorama since it is locally a central projection system in each column.

In order to adapt the network to the distortions of the non-central panoramas, we have fine-tuned it. However, since non-central projection systems are little used in the research community, there is no data-set available. To overcome this difficulty, we have generated a data-set of noncentral circular panoramas with ground truth information for layout recovery in synthetic indoor environments. This data-set is composed of around 650 different layouts from 6 to 10 walls and more than 2500 images. The data-set will be available under request.

2.2. Geometric solvers

In the second part of our pipeline, we take the pixel information provided by the network and reconstruct the scaled layout of the room. For that purpose, we derive two new geometric solvers to jointly obtain the room height and vertical walls location for Manhattan or Atlanta world assumptions.

We define a wall as a set of two parallel lines contained in a vertical plane (see Fig. 3). Let $\mathbf{L} = (\mathbf{l}^T, \mathbf{\bar{l}}^T)^T$ and $\mathbf{M} = (\mathbf{m}^T, \mathbf{\bar{m}}^T)^T$ be the ceiling and floor lines defined in Plücker coordinates [7] and $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ an orthonormal basis attached to the vertical wall. We define the closest points of the lines to the acquisition system \mathbf{x}_L and \mathbf{x}_M with h_c and h_f , distance from the acquisition system to the ceiling and floor planes, and d, distance to the wall plane, such that $\mathbf{x}_L = d\mathbf{e}_2 + h_c \mathbf{e}_3$ and $\mathbf{x}_M = d\mathbf{e}_2 + h_f \mathbf{e}_3$. Notice that with this description we can parameterize the Plücker coordinates of the lines as $\mathbf{l} = \mathbf{m} = \mathbf{e}_1, \mathbf{\bar{l}} = \mathbf{x}_L \times \mathbf{l} = h_c \mathbf{e}_2 - d\mathbf{e}_3$, $\mathbf{\bar{m}} = \mathbf{x}_M \times \mathbf{m} = h_f \mathbf{e}_2 - d\mathbf{e}_3$. We also define the projecting rays that intersect the ceiling and floor lines as $\mathbf{\Xi} = (\boldsymbol{\xi}^T, \mathbf{\bar{\xi}}^T)^T$ and $\mathbf{X} = (\mathbf{\chi}^T, \mathbf{\bar{\chi}}^T)^T$ respectively.

$$side(\mathbf{\Xi}, \mathbf{L}) = \boldsymbol{\xi}^T (h_c \mathbf{e}_2 - d\mathbf{e}_3) + \bar{\boldsymbol{\xi}}^T \mathbf{e}_1 = 0$$
 (1)

$$side(\mathbf{X}, \mathbf{M}) = \boldsymbol{\chi}^T (h_f \mathbf{e}_2 - d\mathbf{e}_3) + \bar{\boldsymbol{\chi}}^T \mathbf{e}_1 = 0$$
 (2)

Known the projecting rays, given by the output of the neural network, we aim to obtain the 3D lines that define each wall in the environment. The relation among the projection rays and the lines is given by their intersection, defined in equations (1) and (2). This is, in general, a non-linear problem which is difficult to tackle directly. However, we propose two new DLT-like approaches that allows to compute the solution for the layout as a linear problem.

$$\bar{\xi}_1 u_x + \bar{\xi}_2 u_y - \xi_1 v_y - \xi_2 v_x - d\xi_3 = 0 \tag{3}$$

$$\bar{\chi}_1 u_x - \bar{\chi}_2 u_y - \chi_1 w_y - \chi_2 w_x - d\chi_3 = 0 \qquad (4)$$



Figure 3: Rays and wall parameter definition. The parameters are: wall reference system $\{e_1, e_2, e_3\}$; Ξ and X define the projecting rays; (l, \overline{l}) and (m, \overline{m}) are the ceiling and floor lines that define the wall; $\mathbf{x_L}, \mathbf{x_M}$ define the closest points of the lines to the origin; h_c , h_f and d are the ceiling and floor height and distance to the wall respectively.

Let the main direction of a wall be horizontal and described by the vector $\mathbf{u} = (u_x, u_y)^T$ such that $\mathbf{l} = \mathbf{m} = (u_x, u_y, 0)^T$. We can define vectors $\mathbf{v} = h_c \mathbf{u}$ and $\mathbf{w} = h_f \mathbf{u}$ such that expressions (1) and (2) become linear obtaining a set of expressions ((3) and (4)) depending on the unknown wall homogeneous vector $\mathbf{W} = (\mathbf{u}^T, \mathbf{v}^T, \mathbf{w}^T, d)^T$.

$$\lambda(\mathbf{v_1} - h_c \mathbf{u_1}) = h_c \mathbf{u_0} - \mathbf{v_0}$$
(5)

$$\lambda(\mathbf{w_1} - h_f \mathbf{u_1}) = h_f \mathbf{u_0} - \mathbf{w_0} \tag{6}$$

However in this linear system $A\mathbf{W} = 0$, \mathbf{u} , \mathbf{v} and \mathbf{w} are independent variables which are non-parallel. In order to impose the parallelism of these vectors we compute the null space of the system with a Singular Value Decomposition (SVD) obtaining a parametric solution which is the linear combination of the singular vectors with a set of parameters λ_i . Two horizontal lines contained in a vertical plane have 4 degrees of freedom. A minimal solution would need 2 rays for each line of the wall, describing the null-space with three singular vectors and two parameters λ_1 and λ_2 . By solving a system of two quadratic equations for λ_1 and λ_2 (with resultants, action matrices or as a polynomial eigenvalue problem [3]) we obtain a set of 4 different solutions which should be discriminated. Since the network provides enough robust information, instead of the minimal solution, we propose to solve the over-determined case (with a minimum of 3 rays lying to each line) with a linear combination involving two singular vectors and single parameter λ (such that $\mathbf{W} = \mathbf{W}_0 + \lambda \mathbf{W}_1$) obtaining two uncoupled quadratic equations (5) and (6) respectively. These equations provide two solutions, where only one of them sets the ceiling line above the floor line.

Notice that with Manhattan world assumption there is a set of walls sharing the wall direction $\mathbf{u} = (u_x, u_y)^T$

Table 1: Comparison of different methods of 3D layout recovery.

	Manhattan World assumption			
	3D IoU	3D IoU	CEN	CE
	(u2s)	50 100	CLIV	CL
CFL [2]	78.87	-	0.75	-
HorizonNet [8]	82.66	-	0.69	-
AtlantaNet [6]	83.94	-	0.71	-
Ours	93.88	86.18	0.787	0.223
	Atlanta World assumption			
HorizonNet [8]	73.53	-	-	-
AtlantaNet [6]	80.01	-	-	-
Ours	91.67	76.17	1.335	0.513
	higher is better		smaller is better	

and the complementary set of walls share the orthogonal direction vector $\mathbf{u}_{\perp} = (-u_y, u_x)^T$. Since all the walls share the ceiling height h_c and the floor h_f , we extend the DLT-like fitting to the whole set of walls, by computing the null-space of $A\mathfrak{L}_M = 0$ where \mathfrak{L}_M is the layout vector $\mathfrak{L}_M = (\mathbf{u}^T, \mathbf{v}^T, \mathbf{w}^T, d_1, \cdots, d_N)^T$ where N is the number of walls and the matrix A is full-filed with relations (3) and (4). The same reasoning as in the case of a single wall can be used to enforce parallelism among $\mathbf{u}, \mathbf{v}, \mathbf{w}$.

$$\bar{\xi_1}' + h_c \xi_2' - d\xi_3' = 0 \tag{7}$$

$$\bar{\chi_1}' + h_f \chi_2' - d\chi_3' = 0 \tag{8}$$

For Atlanta world assumption each wall can have a different horizontal direction, therefore we have to extract each wall independently. Notice that in this case we are not imposing that h_c and h_f are common to all the walls. However, if the direction of each wall is known (for example extracting each wall independently) we can derive a new solution for the whole layout. Assuming that wall directions are known, we can express the projecting rays of each wall in its own local reference system and then equations (1) and (2) become (7) and (8) respectively, where Ξ' and X'are the projecting rays in each wall reference system. Then, we can solve the null-space of a system of linear equations $A\mathfrak{L}_A = \mathbf{0}$ with $\mathfrak{L}_A = (1, h_c, h_f, d_1, \cdots, d_N)$ where A is composed from equations (7) and (8). The main advantage of this second approach is that can be used for Manhattan as well as Atlanta world environments whenever the layout has only one ceiling and floor heights.

3. Experiments

We have performed a set of experiments in order to evaluate our proposal and make a comparison with the state-ofthe-art methods. The metrics used for the comparison are: 3D IoU, which refer to the 3D intersection over union of the



Figure 4: Examples of 3D reconstruction from the proposed method. We shown the non-central panorama and the 3D reconstruction. The green wire frame is the real 3D layout of the room.

predicted layout and the ground truth; 3D IoU(u2s), which refer to the up-to-scale intersection over union of the layout; CEN, which refer to the Corner Error Normalized computed as the L2 distance of the corners divided by the diagonal of the layout; CE, which refers to the Corner Error computed as the L2 distance of the corners in meters.

The comparison with state-of-the-art methods is not completely fair. The datasets used for the different methods are different, so the results can depend on the dataset used and not only on the method. Besides, our proposal only uses the image information in order to recover the 3D layout with the scale while the rest of the methods in the state of the art provide up-to-scale measures, relying on some measure in the environment for the 3D reconstruction, e.g. the camera height. Nevertheless, a summary of this comparison is shown in Table 1.

These results show that our proposal outperforms the state of the art methods for Manhattan as well as Atlanta world assumptions. Besides, our method also recovers the scale of the layout without any prior assumption. Some examples of our results are shown in Fig. 4, where different layouts are tested. We demonstrate that our method can handle quite challenging layouts, in different illumination conditions and world assumptions.

4. Conclusions

In this paper we have proposed two new solvers for indoor layout recovery from a single non-central circular panorama. We use a neural network for extracting the edges of structural lines from a non-central projection system and geometrically process the output in order to recover the 3D information of the layout. Our experiments show that our approach with non-central circular panoramas has better performance than state-of-the-art methods for Manhattan and Atlanta environments. In addition, our method can extract the scale of the room without any prior knowledge.

References

- J. Bermudez-Cameo, O. Saurer, G. Lopez-Nicolas, JJ. Guerrero, and M. Pollefeys. Exploiting line metric reconstruction from non-central circular panoramas. *Pattern Recognition Letters*, 2017. 1
- [2] C. Fernandez-Labrador, JM. Facil, A. Perez-Yus, C. Demonceaux, J. Civera, and JJ. Guerrero. Corners for layout: Endto-end layout recovery from 360 images. *IEEE Robotics and Automation Letters*, 2020. 1, 3
- [3] Z. Kukelova, M. Bujnak, and T. Pajdla. Polynomial eigenvalue solutions to minimal problems in computer vision. *PAMI*, 2011. 3
- [4] Y. Li, HY. Shum, CK. Tang, and R. Szeliski. Stereo reconstruction from multiperspective panoramas. *PAMI*, 2004. 1
- [5] L. Perdigoto and H. Araujo. Reconstruction of 3d lines from a single axial catadioptric image using cross-ratio. In *ICPR*. IEEE, 2012.
- [6] G. Pintore, M. Agus, and E. Gobbetti. Atlantanet: Inferring the 3d indoor layout from a single 360 image beyond the manhattan world assumption. In *ECCV*. Springer, 2020. 1, 3
- [7] H. Pottmann and J. Wallner. *Computational line geometry*. Springer Science & Business Media, 2009. 2
- [8] C. Sun, CW. Hsiao, M. Sun, and HT. Chen. Horizonnet: Learning room layout with 1d representation and pano stretch data augmentation. In *CVPR*. IEEE, 2019. 1, 2, 3
- [9] ST. Yang, FE. Wang, CH. Peng, P. Wonka, M. Sun, and HK. Chu. Dula-net: A dual-projection network for estimating room layouts from a single rgb panorama. In *CVPR*. IEEE, 2019. 1
- [10] Y. Zhang, S. Song, P. Tan, and J. Xiao. Panocontext: A whole-room 3d context model for panoramic scene understanding. In *ECCV*. Springer, 2014. 2
- [11] C. Zou, JW. Su, CH. Peng, A. Colburn, Q. Shan, P. Wonka, HK. Chu, and D. Hoiem. Manhattan room layout reconstruction from a single 360 image: A comparative study of stateof-the-art methods. *IJCV*, 2021. 1