

CPO: Change Robust Panorama to Point Cloud Localization

Junho Kim¹, Hojun Jang¹, Changwoon Choi¹, and Young Min Kim^{1,2}

¹ Department of Electrical and Computer Engineering, Seoul National University

² Interdisciplinary Program in Artificial Intelligence and INMC, Seoul National University

Abstract. We present CPO, a fast and robust algorithm that localizes a 2D panorama with respect to a 3D point cloud of a scene possibly containing changes. To robustly handle scene changes, our approach deviates from conventional feature point matching, and focuses on the spatial context provided from panorama images. Specifically, we propose efficient color histogram generation and subsequent robust localization using score maps. By utilizing the unique equivariance of spherical projections, we propose very fast color histogram generation for a large number of camera poses without explicitly rendering images for all candidate poses. We accumulate the regional consistency of the panorama and point cloud as 2D/3D score maps, and use them to weigh the input color values to further increase robustness. The weighted color distribution quickly finds good initial poses and achieves stable convergence for gradient-based optimization. CPO is lightweight and achieves effective localization in all tested scenarios, showing stable performance despite scene changes, repetitive structures, or featureless regions, which are typical challenges for visual localization with perspective cameras.

Keywords: Visual Localization, Panorama, Point Cloud

1 Introduction

The location information is a crucial building block to develop applications for AR/VR, autonomous driving, and embodied agents. Visual localization is one of the cheapest methods for localization as it could operate only using camera inputs and a pre-captured 3D map. While many existing visual localization algorithms utilize perspective images [29, 31, 35], they are vulnerable to repetitive structures, lack of visual features, or scene changes. Recently, localization using panorama images [6, 7, 20, 37] has gained attention, as devices with 360° cameras are becoming more accessible. The holistic view of panorama images has the potential to compensate for few outliers in localization and thus is less susceptible to minor changes or ambiguities compared to perspective images.

Despite the potential of panorama images, it is challenging to perform localization amidst drastic scene changes while simultaneously attaining efficiency and accuracy. On the 3D map side, it is costly to collect the up-to-date 3D map

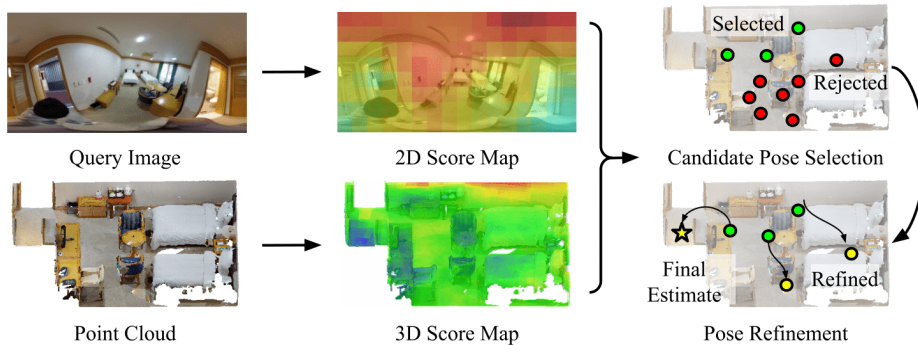


Fig. 1. Overview of our approach. CPO first creates 2D and 3D score maps that attenuate regions containing scene changes. The score maps are further used to guide candidate pose selection and pose refinement.

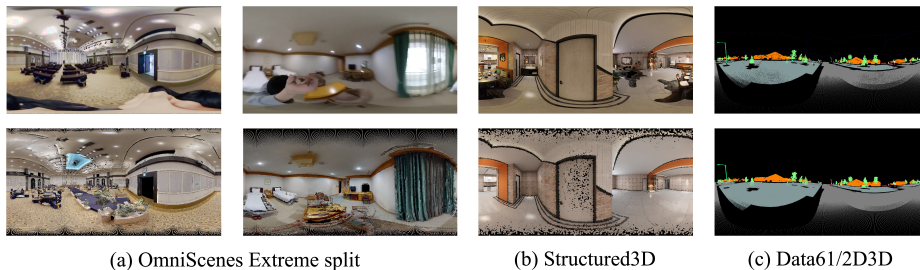


Fig. 2. Qualitative results of CPO. We show the query image (top), and the projected point cloud on the estimated pose (bottom). CPO can flexibly operate using raw color measurements or semantic labels.

that reflects the frequent changes within the scenes. On the algorithmic side, existing localization methods have bottlenecks either in computational efficiency or accuracy. While recent panorama-based localization methods [6, 7, 20, 37] perform accurate localization by leveraging the holistic context in panoramas, they are vulnerable to scene changes without dedicated treatment to account for changes. For perspective cameras, such scene changes are often handled by a two-step approach, using learning-based robust image retrieval [3, 13] followed by feature matching [30]. However, the image retrieval step involves global feature extraction which is often costly to compute and memory intensive.

We propose CPO, a fast localization algorithm that leverages the regional distributions within the panorama images for robust pose prediction under scene changes. Given a 2D panorama image as input, we find the camera pose using a 3D point cloud as the reference map. With careful investigation on the pre-collected 3D map and the holistic view of the panorama, CPO focuses on regions with consistent color distributions. CPO represents the consistency as 2D/3D score maps and quickly selects a small set of initial candidate poses from which the remaining discrepancy can be quickly and stably optimized for accurate localization as shown in Figure 1. As a result, CPO enables panorama to point

cloud localization under scene changes without the use of pose priors, unlike the previous state-of-the-art [20]. Further, the formulation of CPO is flexible and can be applied on both raw color measurements and semantic labels, which is not possible with conventional structure-based localization relying on visual features. To the best of our knowledge, we are the first to explicitly propose a method for coping with changes in panorama to point cloud localization.

The key to fast and stable localization is the efficient color histogram generation that scores the regional consistency of candidate poses. Specifically, we utilize color histograms generated from synthetic projections of the point cloud and make comparisons with the query image. Instead of extensively rendering a large number of synthetic images, we first cache histograms in a few selected views. Then, color histograms for various other views are efficiently approximated by re-using the histograms of the nearest neighbor from the pre-computed color distribution of overlapping views. As a result, CPO generates color histograms for millions of synthetic views within a matter of milliseconds and thus can search a wide range of candidate poses within an order-of-magnitude shorter runtime than competing methods. We compare the color histograms and construct the 2D and 3D score maps, as shown in Figure 1 (middle). The score maps impose higher scores in regions with consistent color distribution, indicating that the region did not change from the reference 3D map. The 2D and 3D score maps are crucial for change-robust localization, which is further verified with our experiments.

We test our algorithm in a wide range of scenes with various input modalities where a few exemplar results are presented in Figure 2. CPO outperforms existing approaches by a large margin despite a considerable amount of scene change or lack of visual features. Notably, CPO attains highly accurate localization, flexibly handling both RGB and semantic labels in both indoor and outdoor scenes, without altering the formulation. Since CPO does not rely on point features, our algorithm is quickly applicable in an off-the-shelf manner without any training of neural networks or collecting pose-annotated images. We expect CPO to be a lightweight solution for stable localization in various practical scenarios.

2 Related Work

In this section, we describe prior works for localization under scene changes, and further elaborate on conventional visual localization methods that employ either a single-step or two-step approach.

Localization under Scene Changes Even the state-of-the-art techniques for visual localization can fail when the visual appearance of the scene changes. This is because conventional localization approaches are often designed to find similar visual appearances from pre-collected images with ground-truth poses. Many visual localization approaches assume that the image features do not significantly change, and either train a neural network [19, 19, 22, 35] or retrieve image features [14, 18, 23, 31, 32]. Numerous datasets and approaches have been presented

in recent years to account for change-robust localization. The proposed datasets reflect day/night [25, 35] or seasonal changes [5, 25, 33] for outdoor scenes and changes in the spatial arrangement of objects [34, 36, 38] for indoor scenes. To cope with such changes, most approaches follow a structure-based paradigm, incorporating a robust image retrieval method [3, 10, 13, 17] along with a learned feature matching module [9, 29, 30, 39]. An alternative approach utilizes indoor layouts from depth images, which stay constant despite changes in object layouts [16]. We compare CPO against various change-robust localization methods, and demonstrate that CPO outperforms the baselines amidst scene changes.

Single-Step Localization Many existing methods [6, 7, 37] for panorama-based localization follow a single-step approach, where the pose is directly found with respect to the 3D map. Since panorama images capture a larger scene context, fewer ambiguities arise than perspective images, and reasonable localization is possible even without a refinement process or a pose-annotated database. Campbell *et al.* [6, 7] introduced a class of global optimization algorithms that could effectively find pose in diverse indoor and outdoor environments [4, 26]. However, these algorithms require consistent semantic segmentation labels for both the panorama and 3D point cloud, which are often hard to acquire in practice. Zhang *et al.* [37] propose a learning-based localization algorithm using panoramic views, where networks are trained using rendered views from the 3D map. We compare CPO with optimization-based algorithms [6, 7], and demonstrate that CPO outperforms these algorithms under a wide variety of practical scenarios.

Two-Step Localization Compared to single-step methods, more accurate localization is often acquired by two-step approaches that initialize poses with an effective search scheme followed by refinement. For panorama images, PICCOLO [20] follows a two-step paradigm, where promising poses are found and further refined using sampling loss values that measure the color discrepancy in 2D and 3D. While PICCOLO does not incorporate learning, it shows competitive performance in conventional panorama localization datasets [20]. Nevertheless, the initialization and refinement is unstable to scene changes as the method lacks explicit treatment of such adversaries. CPO improves upon PICCOLO by leveraging score maps in 2D that attenuate changes for effective initialization and score maps in 3D that guide sampling loss minimization for stable convergence.

For perspective images, many structure-based methods [13, 29] use a two-step approach, where candidate poses are found with image retrieval [3] or scene coordinate regression [22] and further refined with PnP-RANSAC [11] from feature matching [29, 30, 39]. While these methods can effectively localize perspective images, the initialization procedure often requires neural networks that are memory and compute intensive, trained with a dense, pose-annotated database of images. We compare CPO against prominent two-step localization methods, and demonstrate that CPO attains efficiency and accuracy with an effective formulation in the initialization and refinement.

3 Method

Given a point cloud $P = \{X, C\}$, CPO aims to find the optimal rotation $R^* \in SO(3)$ and translation $t^* \in \mathbb{R}^3$ at which the image I_Q is taken. Let $X, C \in \mathbb{R}^{N \times 3}$ denote the point cloud coordinates and color values, and $I_Q \in \mathbb{R}^{H \times W \times 3}$ the query panorama image. Figure 1 depicts the steps that CPO localizes the panorama image under scene changes. First, we extensively measure the color consistency between the panorama and point cloud in various poses. We propose fast histogram generation described in Section 3.1 for efficient comparison. The consistency values are recorded as a 2D score map $M_{2D} \in \mathbb{R}^{H \times W \times 1}$ and a 3D score map $M_{3D} \in \mathbb{R}^{N \times 1}$ which is defined in Section 3.2. We use the color histograms and score maps to select candidate poses (Section 3.3), which are further refined to deduce the final pose (Section 3.4).

3.1 Fast Histogram Generation

Instead of focusing on point features, CPO relies on the regional color distribution of images to match the global context between the 2D and 3D measurements. To cope with color distribution shifts from illumination change or camera white balance, we first preprocess the raw color measurements in 2D and 3D via color histogram matching [1, 8, 15]. Specifically, we generate a single color histogram for the query image and point cloud, and establish a matching between the two distributions via optimal transport. While more sophisticated learning-based methods [12, 24, 40] may be used to handle drastic illumination changes such as night-to-day shifts, we find that simple matching can still handle a modest range of color variations prevalent in practical settings. After preprocessing, we compare the intersections of the RGB color histograms between the *patches* from the query image I_Q and the synthetic projections of the point cloud P .

The efficient generation of color histograms is a major building block for CPO. While there could be an enormous number of poses that the synthetic projections can be generated from, we re-use the pre-computed histograms from another view to accelerate the process. Suppose we have created color histograms for patches of images taken from the original view I_o , as shown in Figure 3. Then the color histogram for the image in a new view I_n can be quickly approximated without explicitly rendering the image and counting bins of colors for pixels within the patches. Let $\mathcal{S}_o = \{S_i^o\}$ denote the image patches of I_o and $\mathcal{C}_o = \{c_i^o\}$ the 2D image coordinates of the patch centroids. \mathcal{S}_n and \mathcal{C}_n are similarly defined for the novel view I_n . For each novel view patch, we project the patch centroid using the relative transformation and obtain the color histogram of the nearest patch of the original image, as described in Figure 3. To elaborate, we first map the patch centroid location c_i^n of $S_i^n \in \mathcal{S}_n$ to the original image coordinate frame,

$$p_i = \Pi(R_{\text{rel}}\Pi^{-1}(c_i^n) + t_{\text{rel}}), \quad (1)$$

where $R_{\text{rel}}, t_{\text{rel}}$ is the relative pose and $\Pi^{-1}(\cdot) : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ is the inverse projection function that maps a 2D coordinate to its 3D world coordinate. The color

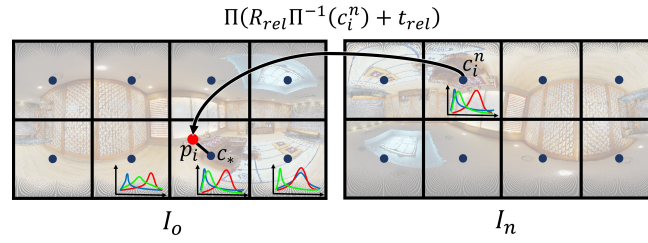


Fig. 3. Illustration of fast histogram generation. For each image patch in the novel view I_n , we first project the patch centroid c_i^n to the view of the original image I_o . The color histogram of the patch in the novel view is estimated as the histogram of image patch c_* in the original view that is closest to the transformed centroid p_i .

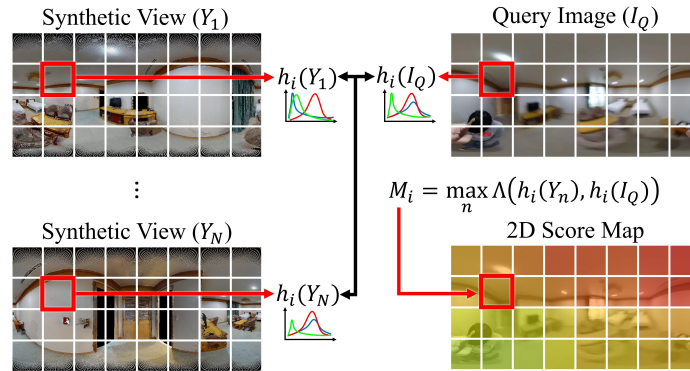


Fig. 4. Illustration of 2D score map generation. The 2D score map for the i^{th} patch M_i is the maximum histogram intersection between the i^{th} patch in query image I_Q and the synthetic views $Y_n \in \mathcal{Y}$.

histogram for S_i^n is assigned as the color histogram of the patch centroid in I_o that is closest to p_i , namely $c_* = \arg \min_{c \in C_o} \|c - p_i\|_2$.

We specifically utilize the cached histograms to generate histograms with arbitrary rotations at a fixed translation. In this case, the camera observes the same set of visible points without changes in occlusion or parallax effect due to depth. Therefore the synthetic image is rendered only once and the patch-wise histograms can be closely approximated by our fast variant with $p_i = \Pi(R_{rel}\Pi^{-1}(c_i^n))$.

3.2 Score Map Generation

Based on the color histogram of the query image and the synthetic views from the point cloud, we generate 2D and 3D score maps to account for possible changes in the measurements. Given a query image $I_Q \in \mathbb{R}^{H \times W \times 3}$, we create multiple synthetic views $Y \in \mathcal{Y}$ at various translations and rotations within the point cloud. Specifically, we project the input point cloud $P = \{X, C\}$ and assign the

measured color $Y(u, v) = C_n$ at the projected location of the corresponding 3D coordinate $(u, v) = \Pi(R_Y X_n + t_Y)$ to create the synthetic view Y .

We further compare the color distribution of the synthetic views $Y \in \mathcal{Y}$ against the input image I_Q and assign higher scores to regions with high consistency. We first divide both the query image and the synthetic views into patches and calculate the color histograms of the patches. Following the notation in Section 3.1, we can denote the patches of the query image as $\mathcal{S}_Q = \{S_i^Q\}$ and $\mathcal{S}_Y = \{S_i^Y\}$ for each synthetic view. Then the color distribution of patch i is recorded into a histogram with B bins per channel: $h_i(\cdot) : \mathbb{R}^{H \times W \times 3} \rightarrow S_i \rightarrow \mathbb{R}^{B \times 3}$. The consistency of two patches is calculated by finding the intersection between two histograms $\Lambda(\cdot, \cdot) : \mathbb{R}^{B \times 3} \times \mathbb{R}^{B \times 3} \rightarrow \mathbb{R}$. Finally, we aggregate the consistency values from multiple synthetic views into the 2D score map for the query image M_{2D} and the 3D score map for the point cloud M_{3D} . We verify the efficacy of the score maps for CPO in Section 4.3.

2D Score Map The 2D score map $M_{2D} \in \mathbb{R}^{H \times W}$ assigns higher scores to regions in the query image I_Q that are consistent with the point cloud color. As shown in Figure 4, we split M_{2D} into patches and assign a score for each patch. We define the 2D score as the maximum histogram intersection that each patch in the input query image I_Q achieves, compared against multiple synthetic views in \mathcal{Y} . Formally, denoting $\mathcal{M} = \{M_i\}$ as the scores for patches in M_{2D} , the score for the i^{th} patch is

$$M_i = \max_{Y \in \mathcal{Y}} \Lambda(h_i(Y), h_i(I_Q)). \quad (2)$$

If a patch in the query image contains scene change it will have small histogram intersections with any of the synthetic views. Note that for computing Equation 2 we use the fast histogram generation from Section 3.1 to avoid the direct rendering of Y . We utilize the 2D score map to attenuate image regions with changes during candidate pose selection in Section 3.3.

3D Score Map The 3D score map $M_{3D} \in \mathbb{R}^N$ measures the color consistency of each 3D point with respect to the query image. We compute the 3D score map by back-projecting the histogram intersection scores to the point cloud locations, as shown in Figure 5. Given a synthetic view $Y \in \mathcal{Y}$, let $B_Y \in \mathbb{R}^N$ denote the assignment of patch-based intersection scores between Y and I_Q into the 3D points whose locations are projected onto corresponding patches in Y . The 3D score map is the average of the back-projected scores B_Y for individual points, namely

$$M_{3D} = \frac{1}{|\mathcal{Y}|} \sum_{Y \in \mathcal{Y}} B_Y. \quad (3)$$

If a region in the point cloud contains scene changes, one can expect the majority of the back-projected scores B_Y to be small for that region, leading to smaller 3D scores. We use the 3D score map to weigh the sampling loss for pose refinement in Section 3.4. By placing smaller weights on regions that contain scene changes, the 3D score map leads to more stable convergence.

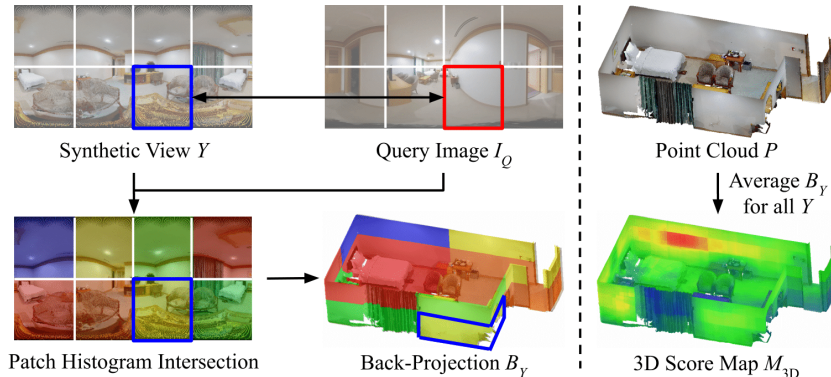


Fig. 5. Illustration of 3D score map generation. For each synthetic view $Y \in \mathcal{Y}$, the patch-wise color histogram is compared against the query image and the resulting intersection scores are back-projected onto 3D locations. The back-projected scores B_Y are averaged for all synthetic views to form the 3D score map M_{3D} .

3.3 Candidate Pose Selection

For the final step, CPO optimizes sampling loss [20] from selected initial poses, as shown in Figure 1. CPO chooses the candidate starting poses by efficiently leveraging the color distribution of the panorama and point cloud. The space of candidate starting poses is selected in two steps. First, we choose N_t 3D locations within various regions of the point cloud, and render N_t synthetic views. For datasets with large open spaces lacking much clutter, the positions are selected from uniform grid partitions. On the other hand, for cluttered indoor scenes, we propose to efficiently handle valid starting positions by building octrees to approximate the amorphous empty spaces as in Rodenberg *et al.* [28] and select centroids of the octrees for N_t starting positions.

Second, we select the final K candidate poses out of $N_t \times N_r$ poses, where N_r is the number of rotations assigned to each translation, uniformly sampled from $SO(3)$. We only render a single view for the N_t locations, and obtain patch-wise histograms for N_r rotations using the fast histogram generation from Section 3.1. We select final K poses that have the largest histogram intersections with the query panorama image. The fast generation of color histograms at synthetic views enables efficient candidate pose selection, which is quantitatively verified in Section 4.3.

Here, we compute the patch-wise histogram intersections for $N_t \times N_r$ poses where the 2D score map M_{2D} from Section 3.2 is used to place smaller weights on image patches that are likely to contain scene change. Let \mathcal{Y}_c denote the $N_t \times N_r$ synthetic views used for finding candidate poses. For a synthetic view $Y \in \mathcal{Y}_c$, the weighted histogram intersection $w(Y)$ with the query image I_Q is expressed as follows,

$$w(Y) = \sum_i M_i A(h_i(Y), h_i(I_Q)). \quad (4)$$

Conceptually, the affinity between a synthetic view Y and the query image I_Q is computed as the sum of each patch-wise intersection weighted by the corresponding patch M_i from the 2D score map M_{2D} . We can expect changed regions to be attenuated in the candidate pose selection process and therefore CPO can quickly compensate for possible scene changes.

3.4 Pose Refinement

We individually refine the selected K poses by optimizing a weighted variant of sampling loss [20], which quantifies the color differences between 2D and 3D. To elaborate, let $\Pi(\cdot)$ be the projection function that maps a point cloud to coordinates in the 2D panorama image I_Q . Further, let $\Gamma(\cdot; I_Q)$ indicate the sampling function that maps 2D coordinates to pixel values sampled from I_Q . The weighted sampling loss enforces each 3D point’s color to be similar to its 2D projection’s sampled color while placing lesser weight on points that are likely to contain change. Given the 3D score map M_{3D} , this is expressed as follows,

$$L_{\text{sampling}}(R, t) = \|M_{3D} \odot [\Gamma(\Pi(RX + t); I_Q) - C]\|_2, \quad (5)$$

where \odot is the Hadamard product and $RX + t$ is the transformed point cloud under the candidate camera pose R, t . To obtain the refined poses, we minimize the weighted sampling loss for the K candidate poses using gradient descent [21]. At termination, the refined pose with the smallest sampling loss value is chosen.

4 Experiments

In this section, we analyze the performance of CPO in various localization scenarios. CPO is mainly implemented using PyTorch [27], and is accelerated with a single RTX 2080 GPU. We report the full hyperparameter setup for running CPO and further qualitative results for each tested scenario in the supplementary material. All translation and rotation errors are reported using median values, and for evaluating accuracy a prediction is considered correct if the translation error is below 0.05m and the rotation error is below 5°.

Baselines We select five baselines for comparison: PICCOLO [20], GOSMA [7], GOPAC [6], structure-based approach, and depth-based approach. PICCOLO, GOSMA, and GOPAC are optimization-based approaches that find pose by minimizing a designated objective function. Structure-based approach [29,31] is one of the most performant methods for localization using perspective images. This baseline first finds promising candidate poses via image retrieval using global features [13] and further refines pose via learned feature matching [30]. To adapt structure-based method to our problem setup using panorama images, we construct a database of pose-annotated synthetic views rendered from the point cloud and use it for retrieval. Depth-based approach first performs learning-based monocular depth estimation on the query panorama image [2], and finds the pose that best aligns the estimated depth to the point cloud. The approach is similar

Table 1. Quantitative results on all splits containing changes in OmniScenes [20].

Method	t -error (m)			R -error ($^{\circ}$)			Accuracy		
	Robot	Hand	Extreme	Robot	Hand	Extreme	Robot	Hand	Extreme
PICCOLO	3.78	4.04	3.99	104.23	121.67	122.30	0.06	0.01	0.01
PICCOLO w/ prior	1.07	0.53	1.24	21.03	7.54	23.71	0.39	0.45	0.38
Structure-Based	0.04	0.05	0.06	0.77	0.86	0.99	0.56	0.51	0.46
Depth-Based	0.46	0.09	0.48	1.35	1.24	2.37	0.38	0.39	0.30
CPO	0.02	0.02	0.03	1.46	0.37	0.37	0.58	0.58	0.57

Table 2. Quantitative results on all splits containing changes in Structured3D [38].

Method	t -error (m)	R -error ($^{\circ}$)	Acc. (0.05m, 5°)	Acc. (0.02m, 2°)	Acc. (0.01m, 1°)
PICCOLO	0.19	4.20	0.47	0.45	0.43
Structure-Based	0.02	0.64	0.59	0.47	0.29
Depth-Based	0.18	1.98	0.45	0.33	0.19
CPO	0.01	0.29	0.56	0.54	0.51

to the layout-matching baseline from Jenkins *et al.* [16], where it demonstrated effective localization under scene change. Additional details about implementing the baselines are deferred to the supplementary material.

4.1 Localization Performance on Scenes with Changes

We assess the robustness of CPO using the OmniScenes [20] and Structured3D [38] dataset, which allows performance evaluation for the localization of panorama images against point clouds in changed scenes.

OmniScenes The OmniScenes dataset consists of seven 3D scans and 4121 2D panorama images, where the panorama images are captured with cameras either handheld or robot mounted. Further, the panorama images are obtained at different times of day and include changes in scene configuration and lighting. OmniScenes contains three splits (Robot, Handheld, Extreme) that are recorded in scenes with changes, where the Extreme split contains panorama images captured with extreme camera motion.

We compare CPO against PICCOLO [20], structure-based approach, and depth-based approach. The evaluation results for all three splits in OmniScenes are shown in Table 1. In all splits, CPO outperforms the baselines without the help of prior information or training neural networks. While PICCOLO [20] performs competitively with gravity direction prior, the performance largely degrades without such information. Further, outliers triggered from scene changes and motion blur make accurate localization difficult using structure-based or depth-based methods. CPO is immune to such adversaries as it explicitly models scene changes and regional inconsistencies with 2D, 3D score maps.

The score maps of CPO effectively attenuate scene changes, providing useful evidence for robust localization. Figure 6 visualizes the exemplar 2D and 3D score maps generated in the wedding hall scene from OmniScenes. The scene contains drastic changes in object layout, where the carpets are removed and

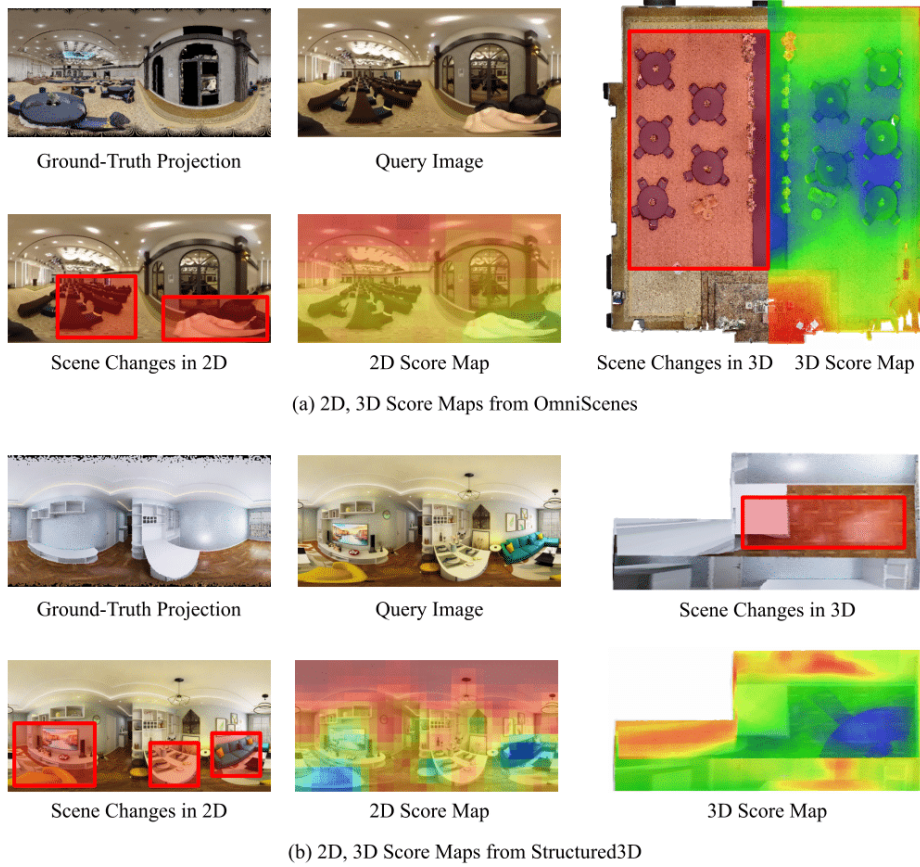


Fig. 6. Visualization of 2D, 3D score maps in OmniScenes [20] and Structured3D [38]. The 2D score map assigns lower scores to the capturer’s hand and objects not present in 3D. Similarly, the 3D score map assigns lower scores to regions not present in 2D.

the arrangement of chairs has largely changed since the 3D scan. As shown in Figure 6, the 2D score map assigns smaller scores to new objects and the capturer’s hand, which are not present in the 3D scan. Further, the 3D score map shown in Figure 6 assigns smaller scores to chairs and blue carpets, which are present in the 3D scan but are largely modified in the panorama image.

Structured3D We further compare CPO against PICCOLO in Structured3D, which is a large-scale dataset containing synthetic 3D models with changes in object layout and illumination, as shown in Figure 2. Due to the large size of the dataset (21845 indoor rooms), 672 rooms are selected for evaluation. For each room, the dataset contains three object configurations (empty, simple, full) along with three lighting configurations (raw, cold, warm), leading to nine configurations in total. We consider the object layout change from empty to full, where illumination change is randomly selected for each room. We provide further de-

Table 3. Quantitative results on Stanford 2D-3D-S [4], compared against PICCOLO (PC), structure-based approach (SB), and depth-based approach (DB).

Area	t -error (m)				R -error ($^{\circ}$)				Accuracy			
	PC	SB	DB	CPO	PC	SB	DB	CPO	PC	SB	DB	CPO
Area 1	0.02	0.05	1.39	0.01	0.46	0.81	89.48	0.25	0.66	0.51	0.28	0.89
Area 2	0.76	0.18	3.00	0.01	2.25	2.08	89.76	0.27	0.42	0.41	0.14	0.81
Area 3	0.02	0.05	1.39	0.01	0.49	1.01	88.94	0.24	0.53	0.50	0.24	0.76
Area 4	0.18	0.05	1.30	0.01	4.17	1.07	89.12	0.28	0.48	0.50	0.28	0.83
Area 5	0.50	0.10	2.37	0.01	14.64	1.31	89.88	0.27	0.44	0.47	0.18	0.73
Area 6	0.01	0.04	1.54	0.01	0.31	0.74	89.39	0.18	0.68	0.55	0.29	0.90
Total	0.03	0.06	1.72	0.01	0.63	1.04	89.51	0.24	0.53	0.49	0.23	0.83

tails about the evaluation in the supplementary material. The median errors and localization accuracy at various thresholds is reported in Table 2. CPO outperforms the baselines in most metrics, due to the change compensation of 2D/3D score maps as shown in Figure 6.

4.2 Localization Performance on Scenes without Changes

We further demonstrate the wide applicability of CPO by comparing CPO with existing approaches in various scene types and input modalities (raw color / semantic labels). The evaluation is performed in one indoor dataset (Stanford 2D-3D-S [4]), and one outdoor dataset (Data61/2D3D [26]). Unlike OmniScenes and Structured3D, most of these datasets lack scene change. Although CPO mainly targets scenes with changes, it shows state-of-the-art results in these datasets. This is due to the fast histogram generation that allows for effective search from the large pool of candidate poses, which is an essential component of panorama to point cloud localization given the highly non-convex nature of the objective function presented in Section 3.

Localization with Raw Color We first make comparisons with PICCOLO [20], structure-based approach, and depth-based approach in the Stanford 2D-3D-S dataset. In Table 3, we report the localization accuracy and median error, where CPO outperforms other baselines by a large margin. Note that PICCOLO is the current state-of-the-art algorithm for the Stanford 2D-3D-S dataset. The median translation and rotation error of PICCOLO [20] deviates largely in areas 2, 4, and 5, which contain a large number of scenes such as hallways that exhibit repetitive structure. On the other hand, the error metrics and accuracy of CPO are much more consistent in all areas.

Localization with Semantic Labels We evaluate the performance of CPO against algorithms that use semantic labels as input, namely GOSMA [7] and GOPAC [6]. We additionally report results from PICCOLO [20], as it could also function with semantic labels. To accommodate for the different input modality, CPO and PICCOLO use color-coded semantic labels as input, as shown in Figure 2(c).

Table 4. Localization performance using semantic labels on a subset of Area 3 from Stanford 2D-3D-S [4]. Q_1 , Q_2 , Q_3 are quartile values of each metric.

	t -error (m)			R -error ($^\circ$)			Runtime (s)		
	Q_1	Q_2	Q_3	Q_1	Q_2	Q_3	Q_1	Q_2	Q_3
PICCOLO	0.00	0.01	0.07	0.11	0.21	0.56	14.0	14.3	16.1
GOSMA	0.05	0.08	0.15	0.91	1.13	2.18	1.4	1.8	4.4
CPO	0.01	0.01	0.02	0.20	0.32	0.51	1.5	1.6	1.6

Table 5. Localization performance on all areas of the Data61/2D3D dataset [26].

Method	t -error (m)			R -error ($^\circ$)		
	GOPAC	PICCOLO	CPO	GOPAC	PICCOLO	CPO
Error	1.1	4.9	0.1	1.4	28.8	0.3

We first compare CPO with PICCOLO and GOSMA on 33 images in Area 3 of the Stanford 2D-3D-S dataset following the evaluation procedure of Campbell *et al.* [7]. As shown in Table 4, CPO outperforms GOSMA [7] by a large margin, with the 3rd quartile values of the errors being smaller than the 1st quartile values of GOSMA [7]. Further, while the performance gap with PICCOLO [20] is smaller than GOSMA, CPO consistently exhibits a much smaller runtime.

We further compare CPO with PICCOLO and GOPAC [6] in the Data61/2D 3D dataset [26], which is an outdoor dataset that contains semantic labels for both 2D and 3D. The dataset is mainly recorded in the rural regions of Australia, where large portions of the scene are highly repetitive and lack features as shown in Figure 2(c). Nevertheless, CPO exceeds GOPAC [6] in localization accuracy, as shown in Table 5. Note that CPO only uses a single GPU for acceleration whereas GOPAC employs a quad-GPU configuration for effective performance [6]. Due to the fast histogram generation from Section 3.1, CPO can efficiently localize using a smaller number of computational resources.

4.3 Ablation Study

In this section, we ablate key components of CPO, namely histogram-based candidate pose selection and 2D, 3D score maps. The ablation study for other constituents of CPO is provided in the supplementary material.

Histogram-Based Candidate Pose Selection We verify the effect of using color histograms for candidate pose selection on the Extreme split from the OmniScenes dataset [20]. CPO is compared with a variant that performs candidate pose selection using sampling loss values as in PICCOLO [20], where all other conditions remain the same. As shown in Table 6, a drastic performance gap is present. CPO uses patch-based color histograms for pose selection and thus considers larger spatial context compared to pixel-wise sampling loss. This allows for CPO to effectively overcome ambiguities that arise from repetitive scene structures and scene changes that are present in the Extreme split.

We further validate the efficiency of histogram-based initialization against various initialization methods used in the baselines. In Table 7, we report the

Table 6. Ablation of various components of CPO in OmniScenes [20] Extreme split.

Method	t -error (m)	R -error ($^{\circ}$)	Acc.
w/o Histogram Initialization	3.29	75.60	0.20
w/o 2D Score Map	0.10	1.19	0.48
w/o 3D Score Map	0.03	1.56	0.55
Ours	0.03	0.37	0.57

Table 7. Average runtime for a single synthetic view in Room 3 from OmniScenes [20].

Method	PICCOLO	Structure-Based	Depth-Based	CPO
Runtime (ms)	2.135	38.70	2.745	0.188

average runtime for processing a single synthetic view in milliseconds. The histogram based initialization used in CPO exhibits an order-of-magnitude shorter runtime than other competing methods. The effective utilization of spherical equivariance in fast histogram generation allows for efficient search within a wide range of poses and quickly generate 2D/3D score maps.

Score Maps We validate the effectiveness of the score maps for robust localization under scene changes on the Extreme split from the OmniScenes dataset [20]. Recall that we use the 2D score map for guiding candidate pose selection and the 3D score map for guiding pose refinement. We report evaluation results for variants of CPO that do not use either the 2D or 3D score map. As shown in Table 6, optimal performance is obtained by using both score maps. The score maps effectively attenuate scene changes, leading to stable pose estimation.

5 Conclusion

In this paper, we present CPO, a fast and robust algorithm for 2D panorama to 3D point cloud localization. To fully leverage the potential of panoramic images for localization, we account for possible scene changes by saving the color distribution consistency in 2D, 3D score maps. The score maps effectively attenuate regions that contain changes and thus lead to more stable camera pose estimation. With the proposed fast histogram generation, the score maps are efficiently constructed and CPO can subsequently select promising initial poses for stable optimization. By effectively utilizing the holistic context in 2D and 3D, CPO achieves stable localization results across various datasets including scenes with changes. We expect CPO to be widely applied in practical localization scenarios where scene change is inevitable.

Acknowledgements This work was partly supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government(MSIT) (No. 2020R1C1C1008195), Creative-Pioneering Researchers Program through Seoul National University, and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2021-0-02068, Artificial Intelligence Innovation Hub).

References

1. Affi, M., Barron, J.T., LeGendre, C., Tsai, Y.T., Bleibel, F.: Cross-camera convolutional color constancy. In: The IEEE International Conference on Computer Vision (ICCV) (2021)
2. Albanis, G., Zioulis, N., Drakoulis, P., Gkitsas, V., Sterzentsenko, V., Alvarez, F., Zarpalas, D., Daras, P.: Pano3d: A holistic benchmark and a solid baseline for 360° depth estimation. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 3722–3732 (2021). <https://doi.org/10.1109/CVPRW53098.2021.00413>
3. Arandjelović, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: NetVLAD: CNN architecture for weakly supervised place recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (2016)
4. Armeni, I., Sax, S., Zamir, A.R., Savarese, S.: Joint 2d-3d-semantic data for indoor scene understanding. arXiv preprint arXiv:1702.01105 (2017)
5. Badino, H., Huber, D., Kanade, T.: The CMU Visual Localization Data Set. <http://3dvis.ri.cmu.edu/data-sets/localization> (2011)
6. Campbell, D., Petersson, L., Kneip, L., Li, H.: Globally-optimal inlier set maximisation for camera pose and correspondence estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence p. preprint (June 2018). <https://doi.org/10.1109/TPAMI.2018.2848650>
7. Campbell, D., Petersson, L., Kneip, L., Li, H., Gould, S.: The alignment of the spheres: Globally-optimal spherical mixture alignment for camera pose estimation. In: Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). p. to appear. IEEE, Long Beach, USA (June 2019)
8. Coltuc, D., Bolon, P., Chassery, J.M.: Exact histogram specification. IEEE transactions on image processing : a publication of the IEEE Signal Processing Society **15**, 1143–52 (06 2006). <https://doi.org/10.1109/TIP.2005.864170>
9. Dong, S., Fan, Q., Wang, H., Shi, J., Yi, L., Funkhouser, T., Chen, B., Guibas, L.J.: Robust neural routing through space partitions for camera relocalization in dynamic indoor environments. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8544–8554 (June 2021)
10. Dusmanu, M., Rocco, I., Pajdla, T., Pollefeys, M., Sivic, J., Torii, A., Sattler, T.: D2-Net: A Trainable CNN for Joint Detection and Description of Local Features. In: Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (2019)
11. Fischler, M.A., Bolles, R.C.: Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM **24**(6), 381–395 (1981), <http://dblp.uni-trier.de/db/journals/cacm/cacm24.htmlFischlerB81>
12. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
13. Ge, Y., Wang, H., Zhu, F., Zhao, R., Li, H.: Self-supervising fine-grained region similarities for large-scale image localization. In: European Conference on Computer Vision (2020)
14. Gee, A.P., Mayol-Cuevas, W.W.: 6d relocalisation for RGBD cameras using synthetic view regression. In: Bowden, R., Collomosse, J.P., Mikolajczyk, K. (eds.) British Machine Vision Conference, BMVC 2012, Surrey, UK, September 3-7, 2012. pp. 1–11. BMVA Press (2012). <https://doi.org/10.5244/C.26.113>, <https://doi.org/10.5244/C.26.113>

15. Gonzalez, R.C., Woods, R.E.: Digital image processing. Prentice Hall, Upper Saddle River, N.J. (2008), <http://www.amazon.com/Digital-Image-Processing-3rd-Edition/dp/013168728X>
16. Howard-Jenkins, H., Ruiz-Sarmiento, J.R., Prisacariu, V.A.: Lalaloc: Latent layout localisation in dynamic, unvisited environments. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10107–10116 (October 2021)
17. Humenberger, M., Cabon, Y., Guerin, N., Morat, J., Revaud, J., Rerole, P., Pion, N., de Souza, C., Leroy, V., Csurka, G.: Robust image retrieval-based visual localization using kapture (2020)
18. Irschara, A., Zach, C., Frahm, J., Bischof, H.: From structure-from-motion point clouds to fast location recognition. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 2599–2606 (2009). <https://doi.org/10.1109/CVPR.2009.5206587>
19. Kendall, A., Grimes, M., Cipolla, R.: Posenet: A convolutional network for real-time 6-dof camera relocalization (2015)
20. Kim, J., Choi, C., Jang, H., Kim, Y.M.: Piccolo: Point cloud-centric omnidirectional localization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3313–3323 (October 2021)
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings (2015), <http://arxiv.org/abs/1412.6980>
22. Li, X., Wang, S., Zhao, Y., Verbeek, J., Kannala, J.: Hierarchical scene coordinate classification and regression for visual localization. In: CVPR (2020)
23. Li, Y., Snavely, N., Huttenlocher, D.P.: Location recognition using prioritized feature matching. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) Computer Vision – ECCV 2010. pp. 791–804. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
24. Luan, F., Paris, S., Shechtman, E., Bala, K.: Deep photo style transfer. arXiv preprint arXiv:1703.07511 (2017)
25. Maddern, W., Pascoe, G., Gadd, M., Barnes, D., Yeomans, B., Newman, P.: Real-time kinematic ground truth for the oxford robotcar dataset. arXiv preprint arXiv:2002.10152 (2020), <https://arxiv.org/pdf/2002.10152>
26. Namin, S., Najafi, M., Salzmann, M., Petersson, L.: A multi-modal graphical model for scene analysis. In: 2015 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1006–1013. IEEE Computer Society, Los Alamitos, CA, USA (jan 2015). <https://doi.org/10.1109/WACV.2015.139>, <https://doi.ieeecomputersociety.org/10.1109/WACV.2015.139>
27. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems 32, pp. 8024–8035. Curran Associates, Inc. (2019), <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
28. Rodenberg, O.B.P.M., Verbree, E., Zlatanova, S.: Indoor A* Pathfinding Through an Octree Representation of a Point Cloud. ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences **IV21**, 249–255 (Oct 2016). <https://doi.org/10.5194/isprs-annals-IV-2-W1-249-2016>

29. Sarlin, P.E., Cadena, C., Siegwart, R., Dymczyk, M.: From coarse to fine: Robust hierarchical localization at large scale. In: CVPR (2019)
30. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: SuperGlue: Learning feature matching with graph neural networks. In: CVPR (2020)
31. Sattler, T., Leibe, B., Kobbelt, L.: Improving image-based localization by active correspondence search. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) Computer Vision – ECCV 2012. pp. 752–765. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
32. Sattler, T., Leibe, B., Kobbelt, L.: Efficient & effective prioritized matching for large-scale image-based localization. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(9), 1744–1756 (2017)
33. Sattler, T., Maddern, W., Toft, C., Torii, A., Hammarstrand, L., Stenborg, E., Safari, D., Okutomi, M., Pollefeys, M., Sivic, J., Kahl, F., Pajdla, T.: Benchmarking 6DOF Outdoor Visual Localization in Changing Conditions. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
34. Taira, H., Okutomi, M., Sattler, T., Cimpoi, M., Pollefeys, M., Sivic, J., Pajdla, T., Torii, A.: InLoc: Indoor Visual Localization with Dense Matching and View Synthesis. In: CVPR 2018 - IEEE Conference on Computer Vision and Pattern Recognition. Salt Lake City, United States (Jun 2018), <https://hal.archives-ouvertes.fr/hal-01859637>
35. Walch, F., Hazirbas, C., Leal-Taixe, L., Sattler, T., Hilsenbeck, S., Cremers, D.: Image-based localization using lstms for structured feature correlation. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (Oct 2017)
36. Wald, J., Sattler, T., Golodetz, S., Cavallari, T., Tombari, F.: Beyond controlled environments: 3d camera re-localization in changing indoor scenes. In: ECCV (2020)
37. Zhang, C., Budvytis, I., Liwicki, S., Cipolla, R.: Rotation equivariant orientation estimation for omnidirectional localization. In: ACCV (2020)
38. Zheng, J., Zhang, J., Li, J., Tang, R., Gao, S., Zhou, Z.: Structured3d: A large photo-realistic dataset for structured 3d modeling. In: Proceedings of The European Conference on Computer Vision (ECCV) (2020)
39. Zhou, Q., Sattler, T., Leal-Taixe, L.: Patch2pix: Epipolar-guided pixel-level correspondences. In: CVPR (2021)
40. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Computer Vision (ICCV), 2017 IEEE International Conference on (2017)