




SPVLoc: Semantic Panoramic Viewport Matching for 6D Camera Localization in Unseen Environments

Niklas Gard^{1,2}, Anna Hilsmann¹, and Peter Eisert^{1,2}

¹ Fraunhofer Heinrich Hertz Institute, HHI, Berlin, Germany
{first.last}@hhi.fraunhofer.de

² Humboldt University of Berlin, Germany

Abstract. In this paper, we present SPVLoc, a global indoor localization method that accurately determines the six-dimensional (6D) camera pose of a query image and requires minimal scene-specific prior knowledge and no scene-specific training. Our approach employs a novel matching procedure to localize the perspective camera’s viewport, given as an RGB image, within a set of panoramic semantic layout representations of the indoor environment. The panoramas are rendered from an untextured 3D reference model, which only comprises approximate structural information about room shapes, along with door and window annotations. We demonstrate that a straightforward convolutional network structure can successfully achieve image-to-panorama and ultimately image-to-model matching. Through a viewport classification score, we rank reference panoramas and select the best match for the query image. Then, a 6D relative pose is estimated between the chosen panorama and query image. Our experiments demonstrate that this approach not only efficiently bridges the domain gap but also generalizes well to previously unseen scenes that are not part of the training data. Moreover, it achieves superior localization accuracy compared to the state of the art methods and also estimates more degrees of freedom of the camera pose. Our source code is publicly available at: <https://fraunhoferhhi.github.io/spvloc>.

Keywords: Localization · 6D pose estimation · Multimodal matching

1 Introduction

Indoor camera localization is a fundamental challenge of computer vision, aiming to precisely determine a camera’s position and orientation within an indoor environment. Applications include navigation [38, 42], augmented reality (AR) [2], building management, and building digitization [26]. Current methods often depend on extensive scene-specific data, acquired through manual creation of databases with localized images, depth information, or point clouds [22, 38, 43]. However, this limits their applicability, especially in situations where exploring the building is impractical or data acquisition is costly and time-consuming.

Recent methods incorporate synthetic data during training, but these localization models require retraining or fine-tuning for new scenes using detailed BIM

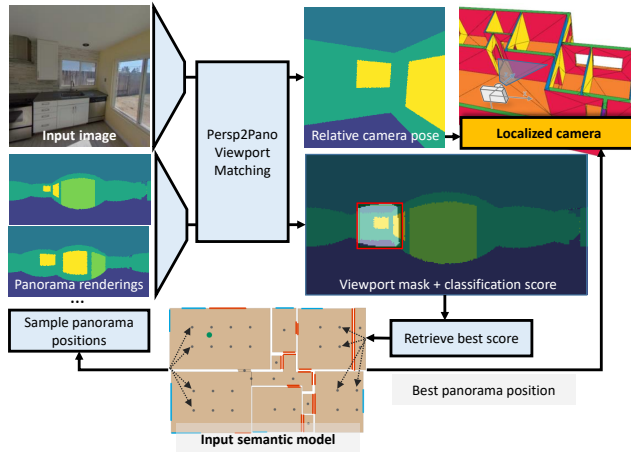


Fig. 1: SPVLoc for 6D indoor localization. Our method calculates the indoor 6D camera pose by determining the image position and orientation relative to synthetic panoramas. The best panoramic match is found through semantic viewport matching.

(Building Information Modeling) models [1, 3]. The adoption of the Structured3D format [46] streamlines and standardizes 3D structure annotation by simplifying complex models into their most fundamental geometric representations. Methods that match a query image to this simplified structural representation exhibit robustness in unexplored, dynamic environments [17, 30]. However, existing methods [9, 30] mainly focus on 2D perspective camera localization with respect to floor plans, missing two degrees of freedom for rotation and one degree of freedom for position. Also, some methods solely rely on panoramic queries [16, 17, 23]. These queries are typically captured under controlled conditions, involving a tripod-mounted camera with a precise horizontal alignment, and a fixed camera height [17]. However, in AR applications or scenarios where archived photos must be accurately registered, e.g. to annotate visible building details into a digital building model, idealized capture conditions cannot always be guaranteed. In such cases, the demand for precise 6D localization becomes imperative.

In this paper, we present SPVLoc, a novel method for indoor 6D camera localization in unseen environments. It represents a 3D scene with a set of synthetically generated omnidirectional panorama images depicting fundamental semantic labels such as door, window, wall, floor, and ceiling. The core idea involves correlating the deep embedding of a perspective RGB camera image with the embeddings of panoramas to estimate the camera viewport – the precise region visible to the camera within the panorama. Leveraging a RetinaNet-like architecture [25], the 2D-bounding box around the viewport is predicted and classified. Retrieval of the closest panorama is performed by selecting the panorama with the globally highest classification score. The exact 6D pose is then found via relative pose estimation with an MLP starting from the panorama’s position. Our method offers several advantages. It relies solely on a minimalist

3D semantic building model [46], which could even be automatically extracted from simple 2D plans [8, 28, 34]. By bridging the domain gap between semantic representation and real images, it facilitates matching in unseen environments. Our approach (Fig. 1) learns wide baseline relative pose estimation and accurately predicts poses with few reference renderings.

In our experiments, we evaluate our method using publicly available datasets, including *Structured3D* [46] and *Zillow Indoor* [10], encompassing both photorealistic synthetic data and real-world data. By sampling perspective query images from equirectangular panoramas, we efficiently augment the dataset, generating an almost infinite amount of training data. Compared to the state of the art, our method excels in estimating all degrees of freedom of the camera pose, enabling precise localization even in uncontrolled recording conditions. To summarise, we make the following contributions.

1. We introduce a model-based 6D camera pose estimation system for unseen indoor environments without the need for scene-specific training.
2. We present a novel perspective-to-panoramic image matching concept with high retrieval accuracy even under wide camera baselines.
3. Demonstrating superior localization accuracy compared to the state of the art, our method estimates more degrees of freedom concurrently.

2 Related Work

Indoor Camera Localization. Indoor camera localization is a longstanding challenge in the computer vision community and in recent years, with the advance of neural networks, several deep learning based approaches have been proposed. Typically, these methods rely on prior scene acquisition to build a detailed scene representation and estimate a relation to a query image. Some methods use a 3D model generated via Structure-from-Motion [5] from registered camera images of the indoor environment and calculate the pose from 2D-3D correspondences, which might be either sparse [27, 43] or dense [38]. Alternatively, encoding the entire scene within a neural network and performing regression to determine the absolute pose is a frequently used strategy [31].

By leveraging synthetic renderings from a detailed 3D BIM model, Acharya *et al.* eliminated the need for real-world reference data. Their approaches achieve real-time indoor localization [1] and an accuracy of approximately 1 meter [4] but are limited to the building of their origin, lacking generalizability to other structures. Liu *et al.* [26] pioneered floor plan-based image localization, predicting the facing wall within a cuboidal layout without estimating a full 3D pose. Howard *et al.* achieved 2D panorama localization in unseen scenes by aligning inferred room layouts with panoramic reference renderings [17] or encoded floor plans [16].

Several methods estimate the 2D position and 1D rotation of a perspective camera relative to a floor plan. PF-net [19] uses a recurrent neural network optimized for a differentiable particle filter, designed for sequential updating but suboptimal for single images [30]. F³Loc [9] improves sequential 2D localization

against a floorplan by integrating multiview geometry cues. LASER [30] localizes the camera with respect to a semantic floor plan, and considers both perspective and panoramic queries. The rendering code book scheme efficiently encodes floor plan information but is limited to 2D data, which constrains the network’s ability to estimate the remaining degrees of freedom.

To sum up, existing methods either depend on accurate scene specific data for a detailed 3D representation or do not estimate the 6D camera pose. In contrast, our method neither depends on real-world reference images, a detailed 3D model, or scene specific training but is still able to estimate the 6D perspective camera pose by introducing a new cross-domain panoramic-to-perspective image registration approach between RGB images and semantic panoramas.

Relative Pose Regression. Relative pose regression (RPR) aims at estimating the relative motion between a query image and reference images with known poses. These frameworks, relying on image comparisons, often enhance generalizability. Typically, they directly utilize localized examples at inference time [12].

Retrieval-based methods like NetVLAD [7] and DenseVLAD [40] start by identifying the most similar images in a database, providing a starting point for RPR. PixLoc [35] predicts multiscale deep features for direct alignment of reference and query image with a CNN and uses geometric optimization for calculating the pose. Turkoglu *et al.* [41] employ a GNN to re-localize cameras with respect to multiple frames. A graph links the query image to training counterparts, and a GNN refines representations for a consistent camera pose estimate. Our method shares similarities with [41] but goes further by addressing the domain gap between simplified synthetic and real data, and by incorporating multimodal matching between panoramic reference and perspective query images.

Perspective-to-Panorama Matching. Matching between perspective query and panoramic reference images is a crucial aspect of localization, yet it has received limited attention in the literature. Most systems avoid multimodal matching and instead either match perspective crops to perspective queries [7, 40] or panoramic to panoramic images [45].

While the past discussed geometric fundamentals for integrating catadioptric and perspective camera systems into hybrid stereo setups, including concepts like epipolar constraints [37] and calibration [14], these ideas have not yet found their way into contemporary AI-based localization systems. In the domain of CNNs and panoramic image processing, distortion-aware convolutional filters were proposed by Tateno *et al.* [39] and further refined by Fernandez *et al.* [13], particularly for indoor equirectangular images. Unlike traditional convolutions, these distortion-aware filters target the differences between perspective and panoramic images, making them particularly valuable for cross-modal matching.

Recently, perspective-to-panoramic matching advanced, due to sliding window approaches in feature space [32, 36]. It improved visual recognition compared to methods that simply sampled perspective reference images from panoramas [44]. Orhan *et al.* [33] explored leveraging semantic segmentation similarity to address appearance variations between reference and query images. Our approach

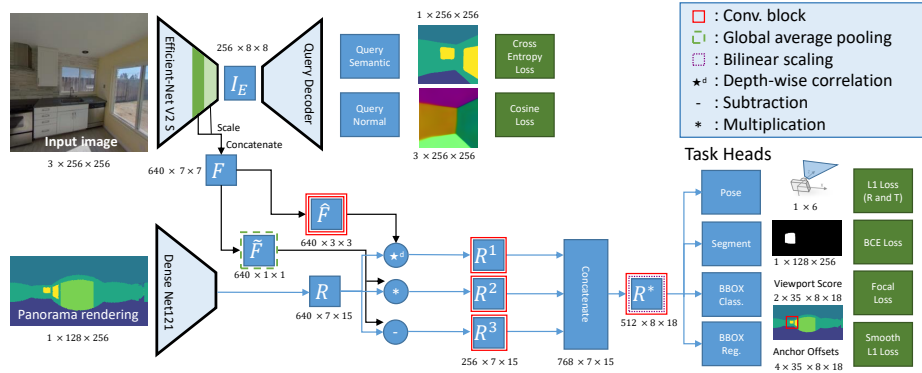


Fig. 2: Network architecture overview. The information from the query branch is fed into the panorama branch by depth-wise correlation. Three task heads predict the corresponding viewport, one task head predicts the relative pose offset.

enhances perspective-to-panorama matching using virtual scene renderings as panoramic references. Inspired by multimodal zero-shot retrieval [18] and object detection, particularly Mercier *et al.*'s work [29], it utilizes the panorama as a global context for query matching.

3 Method

SPVLoc conducts 6D indoor localization of a 2D RGB query image within a simple semantic textureless 3D scene model. The model adheres to the Structured3D annotation format [46], containing walls, doors, windows, ceiling, and floor as planes. It can be derived from straightforward from building models [46], or from floor plans using contemporary AI tools [8, 28, 34]. SPVLoc first estimates the image's viewport within rendered panoramas from the semantic scene model using cross-domain image-to-panorama matching (Sec. 3.1). It then ranks these predictions based on classification accuracy (Sec. 3.4) and performs relative 6D pose regression with respect to the best-matching reference panorama (Sec. 3.2).

3.1 Semantic Panoramic Viewport Matching

Problem Definition. Our objective is to precisely estimate the 6D camera pose denoted as P for a given query RGB image I_q relative to a semantic 3D model S . To achieve this, we reformulate the 6D indoor localization of I_q within S as a cross-domain image-to-panorama matching problem. The model is represented by a set of N equirectangular reference layouts $L = \{L_i \mid i = 1, \dots, N\}$ consisting of semantic labels in a single channel. Each rendering comprehensively captures the entire environment around a specific vantage point. The viewport V is a distinct *cutout* section from a panorama, defining the exact region a camera positioned nearby can see. The perspective distortion of V_{P_i} contains information about

the relative pose offset P_i between L_i and I_q . Scoring and ranking viewpoints helps in choosing the optimal panorama L^* . The absolute pose P is then formed by concatenating P^* with P_i . Identifying the viewport and the corresponding pose offset poses several challenges. First, with increasing distance between the two cameras the perspective distortion between image and panorama increases. Second, the matching involves a severe domain gap as a photographic RGB image is matched against a rendered semantic panorama.

Viewport Prediction. Semantic panorama reference renderings are created utilizing a ray-tracing-based renderer [24]. The viewport mask V_P is a binary image that signifies the visibility of points within the equirectangular image from the view of the perspective camera. The calculation involves projecting all 3D points from the panorama in the perspective view. The projection is performed based on the known underlying layout during the training phase. The mask V_P is determined by identifying pixels projecting into the perspective image, taking into account the 3D camera pose, and excluding those covered by occlusion.

Additionally, the 2D bounding box v_{bb} is the minimal axis-aligned frame around V_P . This bounding box is circularly defined within the 360° panorama, allowing the overlap on the right image side to re-enter on the left side of the panorama. Each predicted bounding box is associated with a confidence score c_{bb} , which quantifies the probability of the bounding box being a valid match. The network’s objective is to predict V_P , v_{bb} and c_{bb} for each pair of input semantic panorama and query RGB image. The optimal viewpoint corresponds to the predicted bounding box with the highest overall confidence score, c_{bb} , when presented with a set of panorama images.

Network Architecture. Our matching is inspired by the object detection strategy of DTOID [29]. Figure 2 illustrates the network architecture. Initially, I_q is encoded with an *EfficientNet-S* backbone. The features obtained after the fourth and the last downsampling layer are preserved. Subsequently, both outputs are rescaled to a tensor of spatial size of 7×7 and concatenated to a tensor F .

L is encoded using a *DenseNet-121* backbone, extending up to the fifth downsampling block to extract features R . The information extracted from the input image guides the panorama branch to estimate the viewport mask and bounding box. This guidance is achieved through the application of a depth-wise correlation technique (denoted as \star_d) [6, 29]. F is processed through two consecutive convolution blocks, each incorporating convolution, batch normalization, and ELU activation, resulting in a set of 3×3 filters \hat{F} used during the depth-wise correlation $R \star_d \hat{F}$.

Furthermore, F undergoes compression via global max pooling and is both multiplied and subtracted from R , resulting in three feature tensors with injected image information. The features are individually processed with an additional convolution block and then concatenated. The concatenated features further undergo a final convolution block, leading to the correlated features R^* .

A decoder estimates the viewport of the perspective camera projected in the layout of the omnidirectional panorama. The viewport in the panorama is

estimated as a bounding box with a *BBox Class.* and *BBox Reg.* head, similar to RetinaNet [25]. The *BBox Class.* head’s output is optimized using focal loss \mathcal{L}_1^{VP} , while the *BBox Reg.* head is optimized with a *smooth l1 loss* \mathcal{L}_2^{VP} . The *BBox Class.* provides a score for a given estimation, ensuring that the bounding box with the highest probability corresponds to the best estimate. Additionally, V_P representing the exact viewport of the perspective camera is estimated with the aid of another convolutional decoder head and optimized with binary cross entropy loss \mathcal{L}_3^{VP} .

Perspective Supervision. The encoded image features I_E are connected to a convolutional decoder, outputting the underlying semantic and normal map to ensure that features closely corresponding to the image content are extracted. The semantics are learned with cross entropy loss \mathcal{L}_4^{VP} and the normals are learned with cosine loss \mathcal{L}_5^{VP} .

3.2 Feature-Correlation-based Pose Regression

The correlated features, denoted as R^* , encode information about the viewport of I_q within a L_i , but should also lead to relative pose offset P_i between them. Given that L is synthetically generated, the precise position of L_i is known. Therefore, determining the relative position of I_q in relation to L_i is equivalent to establishing its absolute position. All rendered panoramas exhibit identity orientation, due to horizontal alignment and a northward orientation. Consequently, the rotation delta directly corresponds to the absolute rotation of the camera.

An additional *Pose* head processes R^* to estimate the relative pose offset of the camera concerning a panorama rendering (Fig. 3). We first utilize three convolution blocks, whereby the first two are followed by a max-pooling layer, to compress the features. The flattened features form an 1D embedding $r^* \in \mathbb{R}^C$. Here, $C = 640$ for a panorama of dimensions 256×128 . Optionally, to handle different camera field of views during training and testing, r^* is normalized and multiplied by the horizontal field of view angle in radians. The result is fed into a four-layer MLP, producing the 3D offset to the query image location $\hat{t}_i \in \mathbb{R}^3$ and a rotation given as the logarithm of the unit quaternion $\hat{r}_i = \log(\hat{q})_{ij} \in \mathbb{R}^3$ [41].

The MLP is evaluated using two loss functions. \mathcal{L}_1^P assesses the precision of the relative translation, and \mathcal{L}_2^P evaluates the accuracy of the final rotation estimation. Both losses are computed using the l_1 distance.

3.3 Optimization

The training loss is the sum of the two losses of the MLP and the five losses of the viewport estimation. With seven losses, our multi-task learning has to be balanced by using a different weight for each loss. As in [20, 21], we use an exponential mapping and weight multiple loss functions by considering the homoscedastic uncertainty of each task and learned weighting factors. The final loss is given as:

$$\mathcal{L} = \sum_{i=1}^2 \mathcal{L}_i^P e^{-\beta_i} + \beta_i + \sum_{i=1}^5 \mathcal{L}_i^{VP} e^{-\gamma_i} + \gamma_i \quad (1)$$

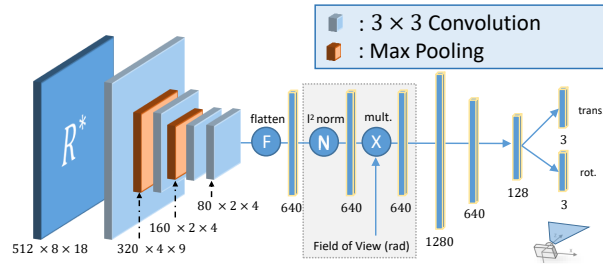


Fig. 3: The *Pose* head utilizes convolutions, an MLP, and a FoV-based side input.

3.4 Inference

During inference, panorama positions are determined from a fixed 2D grid superimposed on the floor plan, with virtual cameras set at a constant height denoted as h_{pano} , and the grid’s dimensions specified as xy_{pano} . Virtual cameras are exclusively sampled above floors to ensure valid reference images inside rooms. Setting xy_{pano} as large as possible is advantageous in order to minimize the number of references, given that feature correlation and bounding box regression are executed for every panorama.

For pose estimation, we select reference positions with the top- n highest bounding box classification scores. The absolute pose is determined by the result of the *Pose* head, where the selected reference position is added to the translation. To eliminate the necessity for rendering during inference, we propose pre-generating reference panoramas for a floor plan. Additionally, the encoding of these panoramas can be precomputed. This ensures that only the encoded features R need to be stored, significantly streamlining the inference process.

Refinement. Without altering the network architecture or training process, SPVLoc can enhance an estimated pose by rendering a new reference panorama at the estimated position and re-executing the relative pose regression. The smaller step to the actual target position in this second iteration simplifies the problem, allowing for increased precision in the refined pose estimation.

4 Experiments

To validate our method we conduct several experiments and make use of two publically available and frequently used datasets.

Structured3D (S3D) [46] consists of 3500 near-photorealistic professionally modeled indoor environments, each annotated with ground truth 3D structure information. It encompasses 21835 panoramic images, each provided in both furnished and unfurnished configurations. The authors define a unified representation that describes scenes using geometric primitives such as planes, lines, junction points, and semantic annotations including floor, ceiling, wall, door, and window classes.




Fig. 4: ZinD data preparation. Annotations generate 3D reference models (left), while resampled panoramas create perspective train and test images (right).

Zillow Indoor (ZInD) [10] includes 67,448 panorama images from 1,575 unfurnished residential homes, globally aligned and registered to floor plans. The 3D annotations of panoramas within a scene are adequate for transforming them into the S3D representation format. Before training, we transform all data to that format. Figure 4 shows an example reference model generated from ZInD annotations. Unlike S3D, ZInD includes room openings such as corridors connected to rooms without doors, enhancing structural variation in semantic renderings.

For both datasets, we employ the official train-test scene split. To enhance control over the field of view and angle variation, we sample perspective test images from the RGB panoramas. They are registered in the 3D model, so the semantics are known for each perspective image. We compile test-sets suitable for 6D pose estimation and structured evaluation. The test sets consist of square perspective images with a viewing angle of 60, 90, and 120°, respectively, along with variations in roll and pitch, and random yaw angles for ZinD and S3D. To ensure meaningful views, we apply two criteria on the sampled yaw. First, the images contain a minimum of three semantic classes, like wall-floor-ceiling or wall-window-floor. Second, the camera faces inside the room rather than towards the next wall. The criteria are applied to the 60° images (4420 for ZinD and 1133 for S3D). Subsequent variations in opening angles and roll-pitch configurations are generated from this starting point.

4.1 Details

For evaluation, we report the fraction of frames localized correctly within a threshold τ . We also list the median rotation and translation error, as well as the fraction of frames localized within τ and an additional angle threshold ν . In the retrieval process, the top-3 best matches are selected, and we report recall metrics for both top-1 and the best among the top-3.

Training. For the state of the art comparison, we train a model with a variable field of view, while other experiments use a model trained with a fixed opening angle of 90°. The model with variable viewing angles randomly samples views between 45 and 135° during training. We consistently render s randomly positioned panoramas around the query within a radius of $\pm r_1$ in xy direction and

Table 1: Comparison with state-of-the-art baseline on ZInD and S3D: Median translation (terr_{xy}) in cm, rotation error (rerr_{yaw}) in deg. for instances localized under 1m. Reporting recall at various translation accuracy levels, recall for inliers ($<1\text{m}$ and $<30^\circ$), and top-3 recall at 1m. *Ours*[†] refers to our model tested with a slightly different testset, incorporating pitch and roll rotation variation ($\pm 10^\circ$).

		ZInD								Structured3D (Furnishing-Level : Full)							
Query	Method	<1m terr (cm)	med <1m rerr (deg)	10cm (%)	50cm (%)	1m (%)	1m & 30° (%)	top-3 1m (%)	<1m terr (cm)	med <1m rerr (deg)	10cm (%)	50cm (%)	1m (%)	1m & 30° (%)	top-3 1m (%)		
Perspective 60° FoV	LASER [30]	33.53	1.03	4.05	47.24	64.52	63.53	78.64	17.60	0.84	16.46	63.22	68.39	67.34	86.08		
	Ours	17.16	1.50	21.04	81.33	86.76	86.09	94.12	12.25	1.21	34.94	87.13	89.49	88.79	96.15		
	Ours [†]	17.77	1.63	19.59	81.13	87.10	86.43	94.86	13.31	1.24	30.82	84.68	88.18	87.48	95.80		
Perspective 90° FoV	LASER [30]	26.28	0.77	8.69	67.01	80.90	80.25	89.73	14.60	0.67	23.47	72.77	76.80	76.18	91.94		
	Ours	13.63	1.34	30.57	85.59	90.84	90.34	96.36	9.17	0.83	50.09	89.67	92.12	91.86	97.37		
	Ours [†]	14.25	1.37	29.59	85.90	91.33	90.95	96.97	10.09	0.91	45.62	89.32	92.12	91.94	96.85		
Perspective 120° FoV	LASER [30]	23.10	0.78	12.69	75.09	85.63	85.20	92.38	12.90	0.76	29.07	77.76	82.14	81.35	94.13		
	Ours	14.78	1.45	29.05	85.14	91.72	91.27	96.92	11.69	0.87	39.49	88.00	92.12	91.94	97.37		
	Ours [†]	15.13	1.53	28.12	86.13	92.71	92.35	97.24	11.77	0.88	38.18	89.67	93.08	92.91	97.55		

$\pm r_2$ in z (upward) direction. For a single query, the viewport and relative pose offset is therefore calculated against multiple references. In our experimental setup, we set s to 4, r_1 to 1.4 meters, and r_2 to 30 centimeters. Additionally, one random *negative* example is generated outside r_1 , likely in a different room, with the target to enhance the network’s ability to distinguish subtle room differences. Images are sampled with random yaw and $\pm 10^\circ$ random pitch and roll angle. The batch size is set to 40, comprising 40 query images and 200 panoramas, and the training is performed on a single NVIDIA A100 GPU. Query images with fewer than three semantic classes are disregarded during the loss calculation process. Training runs for approximately 42,000 steps with an initial learning rate of 2.5×10^{-4} , halved two times during training.

Testing. During testing, panoramas are sampled on a 1.2×1.2 m grid, requiring significantly fewer samples than other rendering-based approaches such as LaLaLoc [17], which uses a 0.5×0.5 m grid. Additionally, for a scene, all panoramic views can be encoded initially, and during inference, only the query image needs to be encoded. If not otherwise noted the studies are conducted on ZInD with fixed 90 degree FoV, reporting the 3D rotation and translation error.

4.2 Comparison with the State of the Art

To our knowledge, SPVLoc is the first method to perform 6D localization in unseen scenes with respect to semantic 3D models. The state of the-art method for scene-independent global localization, LASER [30], takes perspective views as input and matches them against a semantic floor plan. Hence, it estimates only two position and one rotation degree of freedom. Nonetheless, it serves as a strong baseline to assess the localization accuracy of our method.

Table 1 shows the results on ZInD and S3D. Given that we compare against a 2D method, our evaluation also includes recall based on 2D distance and angular error derived from the yaw angle. The yaw angle is extracted from the estimated rotation matrix. Notably, SPVLoc without refinement outperforms LASER in

Table 2: Ablation study. SPVLoc variations compared on ZInD with 90° FoV. In contrast to Tab. 1, terr denotes 3D position errors, and rerr denotes 3D rotation errors.

Abl. Type	Method	<1m med terr (cm)	<1m med rerr (deg)	10cm (%)	50cm (%)	1m (%)	1m & 30° (%)	top-3 1m (%)
(Ours)	base	14.31	2.05	26.88	86.92	91.90	91.61	96.56
	+ refine	12.32	1.96	33.69	90.84	92.76	91.74	-
Loss	- query dec.	15.70	2.23	21.97	85.50	90.79	90.43	96.45
	- view dec.	15.19	2.21	23.60	85.36	90.45	89.82	96.49
	- view/query dec.	16.22	2.19	21.56	84.50	89.59	89.12	96.36
Training	- neg. samples	15.06	2.28	22.60	76.65	82.19	81.74	92.99
Arch.	Δ Resnet18 (query)	15.56	2.56	20.88	86.56	90.72	90.32	96.56
	Δ Equiconv (pano)	14.98	2.13	25.07	87.44	91.70	91.31	96.97
Input	+ normal	14.68	2.28	25.09	85.93	90.84	90.23	96.47
	+ depth	14.06	2.10	27.60	88.05	92.15	91.79	96.74

terms of localization accuracy, exhibiting higher recall across all categories, an increased inlier rate, and significantly reduced median translation error. In terms of 2D rotation, our method is slightly less precise. This is because LASER estimates yaw directly while we estimate a full rotation matrix. Estimating camera height and roll and pitch angle is an advantage of our method. In each last row (*Ours^t*), we evaluate our method on an alternate test set with roll and pitch variations ($\pm 10^\circ$). The results suggest that these variations do not substantially impact the accuracy of our method for 2D localization.

Regarding further methods, LASER has showcased a significant performance improvement over PF-net [19] and non-AI-based Monte-Carlo Localization (MCL) [11]. PF-net, originally designed for sequential time updating, does not perform well in single-query localization scenarios. MCL neglects semantic annotations, uses LIDAR instead of images, which also has been shown to drastically reduce performance [30]. Our substantial improvement over LASER therefore also illustrates a performance gain over aforementioned methods. Another experiment using the pre-rendered perspective images of S3D, including a comparison to F³Loc [9], is part of the supplementary material.

4.3 Performance Studies

Ablation Study. We commence our experimentation with the full network (**base**) configuration. Subsequently, we modify the network by removing specific components (Tab. 2). Initially, we test the influence of perspective supervision (**-query dec.**) and the view segment task head (**-view dec.**) by excluding them during training. Their exclusion reduces network performance, likely because the network’s ability to understand the image content and establish correspondences is diminished. Removing both (**-view/query dec.**) further reduces accuracy.

We then exclude the negative samples from different rooms during training (**-neg. samples**). This significantly reduces the amount of correctly localized images. It supports the hypothesis that teaching the network to recognize spaces where there is no match helps to improve the accuracy of predictions.

To modify the architecture, the image encoder *EfficientNet-S* is replaced with a smaller *ResNet-18* [15] (**Δ Resnet18**), resulting in a performance drop. This




Fig. 5: Qualitative localization results. Top to bottom - query, rendering with top-1 estimated pose, panorama with estimated viewport, map. Green box: success for top-1 match. Yellow box: success for top-2 match. Red box: failure case.

suggests an increased demand for network capacity to analyze the images. In the panorama encoder, replacing all convolutional layers with *Equiconv* [13] (Δ **Equiconv**), a specialized convolution for panorama images that remaps query locations from Euclidean to spherical space, does not result in a performance gain. This suggests that the capacity of the small *DenseNet-121* is sufficient for analyzing the low-resolution semantic inputs.


Lastly, we incorporate additional input modalities for the panoramas into the panorama encoder. While the addition of world coordinate normals (+**normals**) did not improve the result, including an extra depth map slightly enhanced the outcome (+**depth**). This suggests that there are cases where additional information is encoded in the depth that cannot be inferred from semantics alone.

Qualitative Results. Figure 5 shows selected SPVLoc outcomes. The top row displays camera images, and the second row shows renderings from the semantic 3D model with estimated top-1 poses. All green boxed images are correctly localized for the top-1 match. In the yellow box, the top-1 match fails, but the top-2 match is accurate. The red box example cannot be localized. Across all examples, a pose with high visual similarity is found. Ambiguous examples arise from similar layouts occurring in multiple rooms. Failure cases are attributed to either images lacking sufficient structural information for matching or outliers in the data, such as those taken on balconies or in dark cellars. The third row displays the chosen reference panorama, the selected viewport bounding box corresponding to the camera image, and a semi-transparent overlay representing the estimated viewport mask. Additional detailed visualizations are part of the supplementary material.

Influence of Sampling Grid Sizes and Refinement. Reference positions result from overlaying a grid onto the floorplan, and the grid density significantly

Table 3: Grid size vs. accuracy: Deviation of 10cm and 1m recall from (boxed) baseline.

Sample dist.	Samples per m^2	Adv.	10cm (%)	10cm + refine	1m (%)	1m + refine
0.7	2.07		+1.00	+7.49	+1.86	+2.38
0.9	1.25		+0.95	+6.99	+1.79	+2.08
1.2	0.70		26.88	+6.81	91.90	+0.86
1.2	0.48	✓	+0.57	+6.95	-0.63	+0.61
1.5	0.45		-1.00	+6.63	-2.47	-1.02
1.5	0.31	✓	-1.04	+6.47	-2.67	-0.38
2.0	0.26		-4.19	+3.60	-9.05	-5.79
2.0	0.20	✓	-1.76	+5.81	-3.91	-1.18

**Fig. 6:** Reference positions for global grid of size 0.7m, 1.5m and local grid of size 1.5m (from left to right).

affects SPVLoc’s efficiency and inference time. Minimizing references reduces rendering effort and requires the model to estimate larger offsets. E.g. to increase the grid from $0.7m$ to $2m$ reduces the number of samples per m^2 by a factor of 10. We vary the grid size to analyze its impact on performance. We also experiment with a local grid per room substituting the global grid (**Adv.**) reducing the risk of missing rooms entirely. The local grid is centered on the room outline and additionally places one camera in the middle of each opening. Figure 6 shows different grids for an example plan. Table 3 illustrates the relationships among grid size, global and local sampling, and refinement. Refinement notably improves 10cm recall across configurations. For the largest local grid with refinement, the 1m recall only decreases by 1.18% compared to the baseline configuration. In summary, a trade-off between refinement and dense/sparse scene sampling exists, yet our method consistently delivers strong results with very sparse sampling, making it suitable for large scenes.

Influence of Focal Length. Table 4 depicts the performance difference between a variant trained with a known camera focal length and a focal length-independent network trained with varying focal lengths (ranging from 45 to 135 degrees). The fixed focal length network performs slightly better with matching images but loses precision with test images of different focal lengths. In contrast, the focal length-independent network demonstrates the capability to handle such differences. We argue that the significantly enhanced flexibility of the latter justifies the minor performance degradation in practical, real-world scenarios.

Influence of Camera Rotation Angles. The range of pitch and roll angles the network is able to handle is determined by the variations of angles seen during training. We train the network for ZInD using a angle variation of $\pm 22^\circ$ and test with images where the range of variation increases from 0° to 20° in steps of $\pm 5^\circ$. Table 5 shows only a slight increase in median errors for rotation and translation as the rotation angles in test images increase. The chosen ranges represent a significant coverage and offer substantial improvements over existing methods. In indoor scenes, larger camera tilt often results in images dominated by the ceiling or floor, rendering them unsuitable for semantic matching.

Table 4: The fixed FoV network excels at a specific angle but loses precision with changing FoV. The proposed network remains stable.

Input FoV	Train FoV	<1m med terr (cm)	<1m med rerr (deg)	10cm (%)	1m (%)
90°	90°	14.31	2.05	26.88	91.90
	45 - 135°	16.11	2.20	21.81	91.45
80°	90°	37.51	2.37	1.67	90.14
	45 - 135°	16.75	2.19	19.30	90.50

Table 5: Effect of increasing angle variation in test data.

Metric	Roll/Pitch Variation				
	±0°	±5°	±10°	±15°	±20°
<1m med. terr	14.46	14.88	14.88	15.21	15.91
<1m med. rerr	2.35	2.36	2.38	2.51	2.60
10cm (%)	24.64	23.94	24.46	23.69	22.74
1m (%)	91.45	91.47	91.99	92.38	92.60

Table 6: Timing. Performance on ZInD using a NVIDIA Tesla V100, following the configurations outlined in Sec. 4.3, and compared against results from [30].

Method	Sampling Time (s)	Query Fps	Query Module Time (ms)
PF-net [19]	48.95 ± 38.95	5.06 ± 1.77	Eff.Net V2 S 4.4
LASER [30]	0.97 ± 1.09	8.31 ± 0.64	Correlation 35.7
Ours	1.66 ± 1.06	28.63 ± 11.54	Pose 1.1

Limitations. In large spaces with repeating room layouts, the approach’s effectiveness may be constrained by the limited level of detail in the semantic reference model. However, enhancing the S3D format by adding more semantic classes and structural information, such as staircases or sanitary installations, could reduce ambiguities.

Timing. We split the measured time into scene sampling/preparation time, including position sampling, panorama rendering and encoding, and inference/query time, including query image encoding, feature correlation with bounding box classification, and the *Pose* task head, but no refinement. Table 6 shows that our method exhibits a considerably lower query time than LASER [30] and PF-net [19]. This difference places our approach close to real-time for moderately sized scenes. Currently, scene sampling time is slower than [30], mainly due to rendering on the CPU without parallelization [24]. Parallelizing this process using GPU or CPU could significantly accelerate scene preparation.

5 Conclusion

This paper introduces a method for scene-independent model-based 6D localization for indoor scenes, involving a novel method for multimodal image matching (panorama to perspective, RGB to semantic). The matching and retrieval is efficient and scalable under sparse reference sampling. The localization accuracy and inference speed outperform state of the art methods and the inclusion of the 3D-model, reduces ambiguities to precisely estimate 6D poses. Future work involves combining localization and image analysis to enhance digital building models or exploring applications in augmented reality scenarios.

Acknowledgements. This work is partially funded by the German Federal Ministry of Economic Affairs and Climate Action (BIMKIT, 01MK21001H) and the German Federal Ministry for Digital and Transport (EConoM, 19OI22009C).

References

1. Acharya, D., Khoshelham, K., Winter, S.: BIM-PoseNet: Indoor camera localisation using a 3D indoor model and deep learning from synthetic images. *ISPRS journal of photogrammetry and remote sensing* **150**, 245–258 (2019)
2. Acharya, D., Ramezani, M., Khoshelham, K., Winter, S.: Bim-tracker: A model-based visual tracking approach for indoor localisation using a 3d building model. *ISPRS Journal of Photogrammetry and Remote Sensing* **150**, 157–171 (2019)
3. Acharya, D., Tatli, C.J., Khoshelham, K.: Synthetic-real image domain adaptation for indoor camera pose regression using a 3d model. *ISPRS Journal of Photogrammetry and Remote Sensing* **202**, 405–421 (2023)
4. Acharya, D., Tennakoon, R., Muthu, S., Khoshelham, K., Hoseinnezhad, R., Bab-Hadiashar, A.: Single-image localisation using 3D models: Combining hierarchical edge maps and semantic segmentation for domain adaptation. *Automation in Construction* **136**, 104–152 (2022)
5. Agarwal, S., Furukawa, Y., Snavely, N., Simon, I., Curless, B., Seitz, S.M., Szeliski, R.: Building rome in a day. *Communications of the ACM* **54**(10), 105–112 (2011)
6. Ammirato, P., Fu, C.Y., Shvets, M., Kosecka, J., Berg, A.C.: Target driven instance detection. arXiv preprint arXiv:1803.04610 (2018)
7. Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: Netvlad: Cnn architecture for weakly supervised place recognition. In: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
8. Cambeiro, A.B., Trzeciakiewicz, M., Hilsmann, A., Eisert, P.: Automatic reconstruction of semantic 3d models from 2d floor plans. In: *Proc. Int. Conf. on Machine Vision Applications (MVA)*. Hamamatsu, Japan (July 2023)
9. Chen, C., Wang, R., Vogel, C., Pollefeys, M.: F3Loc: Fusion and Filtering for Floorplan Localization. In: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2024)
10. Cruz, S., Hutchcroft, W., Li, Y., Khosravan, N., Boyadzhiev, I., Kang, S.B.: Zillow indoor dataset: Annotated floor plans with 360deg panoramas and 3d room layouts. In: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2133–2143 (2021)
11. Dellaert, F., Fox, D., Burgard, W., Thrun, S.: Monte carlo localization for mobile robots. In: *Proceedings IEEE International Conference on Robotics and Automation*. vol. 2, pp. 1322–1328 (1999)
12. Ding, M., Wang, Z., Sun, J., Shi, J., Luo, P.: CamNet: Coarse-to-fine retrieval for camera re-localization. In: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2871–2880 (2019)
13. Fernandez-Labrador, C., Facil, J.M., Perez-Yus, A., Demonceaux, C., Civera, J., Guerrero, J.J.: Corners for layout: End-to-end layout recovery from 360 images. *IEEE Robotics and Automation Letters* **5**(2), 1255–1262 (2020)
14. He, B., Chen, Z., Li, Y.: Calibration method for a central catadioptric-perspective camera system. *JOSA A* **29**(11), 2514–2524 (2012)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016)

16. Howard-Jenkins, H., Prisacariu, V.A.: LaLaLoc++: Global Floor Plan Comprehension for Layout Localisation in Unvisited Environments. In: Proc. European Conference on Computer Vision (ECCV). pp. 693–709. Springer (2022)
17. Howard-Jenkins, H., Ruiz-Sarmiento, J.R., Prisacariu, V.A.: LaLaLoc: Latent Layout Localisation in Dynamic, Unvisited Environments. In: Proc. IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
18. Janik, M., Gard, N., Hilsmann, A., Eisert, P.: Zero in on Shape: A Generic 2D-3D Instance Similarity Metric Learned from Synthetic Data. In: Proc. IEEE International Conference on Image Processing (ICIP). pp. 2638–2642. IEEE (2021)
19. Karkus, P., Hsu, D., Lee, W.S.: Particle filter networks with application to visual localization. In: Conference on robot learning. pp. 169–178. PMLR (2018)
20. Kendall, A., Cipolla, R.: Geometric loss functions for camera pose regression with deep learning. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5974–5983 (2017)
21. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7482–7491 (2018)
22. Kendall, A., Grimes, M., Cipolla, R.: Posenet: A convolutional network for real-time 6-dof camera relocalization. In: Proc. IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2938–2946 (2015)
23. Kim, J., Jeong, J., Kim, Y.M.: Fully geometric panoramic localization. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20827–20837 (2024)
24. Li, T.M., Aittala, M., Durand, F., Lehtinen, J.: Differentiable Monte Carlo Ray Tracing through Edge Sampling. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* **37**(6), 222:1–222:11 (2018)
25. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proc. IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2980–2988 (2017)
26. Liu, C., Schwing, A., Kundu, K., Urtasun, R., Fidler, S.: Rent3d: Floor-plan priors for monocular layout estimation. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
27. Liu, L., Li, H., Dai, Y.: Efficient global 2d-3d matching for camera localization in a large-scale 3d map. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2372–2381 (2017)
28. Lv, X., Zhao, S., Yu, X., Zhao, B.: Residential floor plan recognition and reconstruction. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16717–16726 (June 2021)
29. Mercier, J.P., Garon, M., Giguere, P., Lalonde, J.F.: Deep Template-Based Object Instance Detection. In: Proc. WACV. pp. 1507–1516 (January 2021)
30. Min, Z., Khosravan, N., Bessinger, Z., Narayana, M., Kang, S.B., Dunn, E., Boyadzhiev, I.: Laser: Latent space rendering for 2d visual localization. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11122–11131 (2022)
31. Naseer, T., Burgard, W.: Deep regression for monocular camera-based 6-dof global localization in outdoor environments. In: Proc. International Conference on Intelligent Robot Systems (IROS). pp. 1525–1530. IEEE (2017)
32. Orhan, S., Baştanlar, Y.: Efficient search in a panoramic image database for long-term visual localization. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1727–1734 (2021)

33. Orhan, S., Guerrero, J.J., Baştanlar, Y.: Semantic Pose Verification for Outdoor Visual Localization With Self-Supervised Contrastive Learning. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 3989–3998 (June 2022)
34. Park, S., Kim, H.: 3dplannet: generating 3d models from 2d floor plan images using ensemble methods. *Electronics* **10**(22), 2729 (2021)
35. Sarlin, P.E., Unagar, A., Larsson, M., Germain, H., Toft, C., Larsson, V., Pollefeys, M., Lepetit, V., Hammarstrand, L., Kahl, F., et al.: Back to the feature: Learning robust camera localization from pixels to pose. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3247–3257 (2021)
36. Shi, Z., Shi, H., Yang, K., Yin, Z., Lin, Y., Wang, K.: PanoVPR: Towards Unified Perspective-to-Equirectangular Visual Place Recognition via Sliding Windows across the Panoramic View. In: 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC). pp. 1333–1340 (2023)
37. Sturm, P.: Mixing catadioptric and perspective cameras. In: Proceedings of the IEEE Workshop on Omnidirectional Vision 2002. Held in conjunction with ECCV’02. pp. 37–44. IEEE (2002)
38. 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: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7199–7209 (2018)
39. Tateno, K., Navab, N., Tombari, F.: Distortion-aware convolutional filters for dense prediction in panoramic images. In: Proc. European Conference on Computer Vision (ECCV). pp. 707–722 (2018)
40. Torii, A., Arandjelovic, R., Sivic, J., Okutomi, M., Pajdla, T.: 24/7 place recognition by view synthesis. In: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1808–1817 (2015)
41. Türkoğlu, M.Ö., Brachmann, E., Schindler, K., Brostow, G., Monszpart, A.: Visual Camera Re-Localization Using Graph Neural Networks and Relative Pose Supervision. In: Proc. International Conference on 3D Vision (3DV) (2021)
42. Wang, S., Fidler, S., Urtasun, R.: Lost shopping! monocular localization in large indoor spaces. In: Proc. IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2695–2703 (2015)
43. Xia, J., Gong, J.: Precise indoor localization with 3d facility scan data. *Computer-Aided Civil and Infrastructure Engineering* **37**(10), 1243–1259 (2022)
44. Zamir, A.R., Shah, M.: Accurate image localization based on google maps street view. In: Proc. European Conference on Computer Vision (ECCV). pp. 255–268. Springer, Heraklion, Greece (Sep 2010)
45. Zhang, C., Budvytis, I., Liwicki, S., Cipolla, R.: Lifted semantic graph embedding for omnidirectional place recognition. In: Proc. International Conference on 3D Vision (3DV). pp. 1401–1410. IEEE (2021)
46. Zheng, J., Zhang, J., Li, J., Tang, R., Gao, S., Zhou, Z.: Structured3D: A large photo-realistic dataset for structured 3d modeling. In: Proc. European Conference on Computer Vision (ECCV). pp. 519–535. Springer, Glasgow, UK (Aug 2020)