

# Inside Out Visual Place Recognition

Sarah Ibrahim  
Nanne van Noord  
Tim Alpherts  
Marcel Worring  
{s.ibrahimi,n.j.e.vannoord,t.o.l.alpherts,m.worring}@uva.nl

University of Amsterdam  
Science Park 904  
Amsterdam  
The Netherlands

## Abstract

Visual Place Recognition (VPR) is generally concerned with localizing outdoor images. However, localizing indoor scenes that contain part of an outdoor scene can be of large value for a wide range of applications. In this paper, we introduce *Inside Out Visual Place Recognition (IOVPR)*, a task aiming to localize images based on outdoor scenes visible through windows. For this task we present the new large-scale dataset *Amsterdam-XXXL*, with images taken in Amsterdam, that consists of 6.4 million panoramic street-view images and 1000 user-generated indoor queries. Additionally, we introduce a new training protocol *Inside Out Data Augmentation* to adapt Visual Place Recognition methods for localizing indoor images, demonstrating the potential of Inside Out Visual Place Recognition. We empirically show the benefits of our proposed data augmentation scheme on a smaller scale, whilst demonstrating the difficulty of this large-scale dataset for existing methods. With this new task we aim to encourage development of methods for IOVPR. The dataset and code are available for research purposes at <https://github.com/saibr/IOVPR>.

## 1 Introduction

In Visual Place Recognition the goal is to match a query image, for which the location is unknown, through Instance Search to a gallery of images with known geolocations. Application areas for this task include autonomous driving [4], SLAM [18], and matching historical architectural images to modern images [29]. Research to develop VPR models has focused on several directions ranging from recognizing popular landmarks [15, 19, 20] to arbitrary street scene localization with street-view images [4, 26, 27]. Until now place recognition methods have been exclusively developed with outdoor images as queries, yet often times indoor images provide information about the outside world through a window. Current methods fail to localize these images. We introduce the new task *Inside Out Visual Place Recognition (IOVPR)*. The goal for this task is to match the outdoor content that is visible in an indoor query image to a gallery of outdoor images, making it possible to localize images taken in an indoor setting. Figure 1(a) shows such indoor queries with their ground truth location.

IOVPR can be employed for a wide range of applications, such as automated tagging of user-images on e.g. travel platforms or social media. More specifically it is useful for forensics, where localizing crime scenes is of great importance. Photographic evidence of indoor environments is common, since these criminal activities usually take place behind

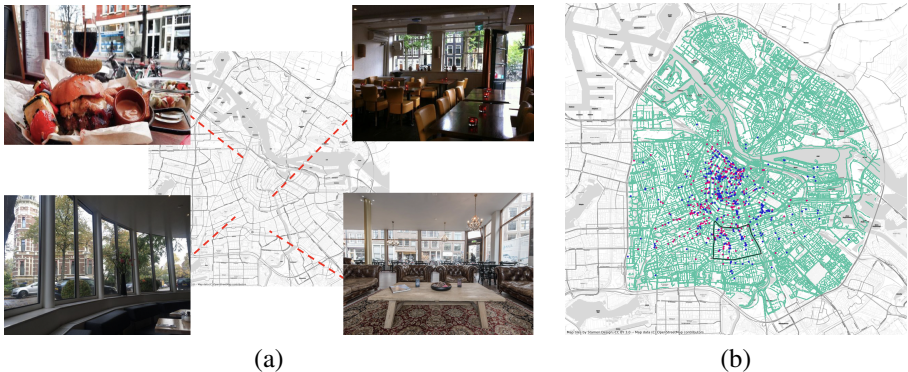


Figure 1: (a) Examples of images taken in Amsterdam that are localized in the Inside Out Visual Place Recognition task. (b) Visualization of the Amsterdam-XXXL dataset. The green points indicate the Outdoor-Ams locations. The blue and pink points are the locations of the query images from Indoor-Ams validation set and test set respectively. The black boundary indicates the neighbourhood for the Ams30k partition.

closed doors. In some cases, part of the outdoor environment is visible through a window. Organizations such as Bellingcat and Homeland Security Investigations have been involved in cases where images taken in an indoor environment as a house or car that contain a window have contributed to solving the case [6, 10]. The amount of available relevant indoor images in a case varies, but as shown in the Sunflower Case [10], one suitable image can already make a difference. For human trafficking or child abuse human lives can be saved when images depicting victims are localized [24]. Another forensic application for IOVPR is the detection of housing fraud, where house owners exceed legal rental restrictions. Since address information is only released to tenants, IOVPR can be used to localize publicly advertised images. Localizing these images takes a lot of manual effort and time that we aim to reduce by introducing IOVPR. To support the investigative efforts in this direction, which are not supported by regular Visual Place Recognition, we explicitly focus on indoor environments by releasing a new dataset *Amsterdam-XXXL*, as well as a novel training protocol *Inside Out Data Augmentation*, that makes it possible to train Instance Search models for IOVPR. We empirically show the benefits of this training protocol on a smaller scale and show the difficulty of this large-scale dataset for existing methods.

The dataset introduced in this paper, *Amsterdam-XXXL*, consists of a spatially dense collection of 6.4 million panoramic street-view images and user photos from indoor locations, as visualized in Figure 1(b). This dataset is 80 times larger on city-level than widely used VPR datasets, to provide a full coverage of the central part of Amsterdam. A second important aspect is the domain shift between the queries and the gallery. We are the first to present a large-scale VPR dataset with such a domain shift. Additionally, we propose a novel data augmentation technique to adapt VPR models such that they are suitable for this new task. Specifically, we propose *Inside Out Data Augmentation* which we use to overlay indoor scenes onto street-view images with known geo-locations, synthetically creating ‘indoor’ query images with recognizable outdoor content. In this work we explore whether this data augmentation scheme is suitable for synthesizing training data for focusing the attention of existing VPR models on the window content. By using synthesized content it is possible to avoid the costly data collection and annotation process.

Our contributions are the following. (i) We introduce *Inside Out Visual Place Recogni-*

tion as a new task to localize indoor images through Instance Search. (ii) We release a new dataset, *Amsterdam-XXXL*, to encourage the development of new and improved techniques for IOVPR. (iii) We propose *Inside Out Data Augmentation* to train VPR models on synthetically generated indoor images and present baseline results to show the potential of our method and challenges on the large-scale aspects of this task.

## 2 Related Work

**Datasets.** Popular datasets that accelerated the development of Visual Place Recognition models for landmarks are Oxford Buildings [20], Paris [21] and Holidays [11]. Each of these datasets consists of a few thousand images and were suitable for early retrieval techniques because of their size and large number of positives for each query [22]. Deep Learning techniques required larger datasets which resulted in an extension of the Oxford and Paris datasets by the use of 1M distractor images [22]. Other landmark datasets that have been created for retrieval purposes are the Google Landmarks Dataset v1 [19] with 1M images and 15k distinct landmarks and its second version with over 5M images and over 200k distinct landmarks [6]. Another fruitful direction in VPR research has been related to the use of (panoramic) street-view images. Starting from NetVLAD[1], street-view datasets such as Pittsburgh 250k and Tokyo 24/7 use large amounts of images depicting the same places over time with the help of the Google Street View Time Machine. The Mapillary Street-level Sequences dataset is a more diverse dataset for VPR, with 1.6 million images from different cameras, captured over different times [6]. With data from different cities and continents, it aims to create models that generalize better to new environments. A recent variation of VPR is ground-to-aerial image geo-localization, where ground-view query images are compared to aerial image in a reference set [2, 9, 16]. Popular datasets for this task, e.g. CVUSA [5] and Vo and Hays’ dataset [23], use street-view images as query images.

**Methods.** For VPR, we distinguish supervised and weakly-supervised methods. Supervised methods rely on a strong ground truth label, such as the class label for a landmark, to perform the optimization. These methods typically use local-feature based descriptors [22, 25], global descriptors [8, 24], or both [3, 19, 22]. Differently, when dealing with street-view images, the geolocation is used to select nearby images, but within this subset there is no ground-truth for visual matches. Therefore it is necessary to employ weakly-supervised methods. A common aspect of weakly-supervised models for VPR is the use of a triplet-based ranking loss. As one of the first methods that used triplet-based ranking for VPR, NetVLAD [1] uses a learnable layer based on VLAD [22] to aggregate local descriptors and assign them to cluster centers. Extensions of NetVLAD are SARE [17], CRN [24], SFRS [8], and Patch-NetVLAD [7]. SARE [17] adds additional constraints on intra-place and inter-place feature embeddings, and CRN [24] uses context-aware feature reweighting to select the most relevant contexts for localization through efficient hard negative mining. SFRS [8] uses the NetVLAD backbone with self-supervised image-to-region similarities. These similarities return soft labels that further help the system find difficult positive images. This can be seen as additional, weak, training supervision which could help in refining fine-grained similarities. Patch-NetVLAD [7] is a patch-based local descriptor method that reranks the top-100 results retrieved by NetVLAD. Our work adapts existing VPR models, making them suitable for the localization of indoor images, demonstrating the potential of IOVPR.

### 3 Inside Out Visual Place Recognition

While Visual Place Recognition aims to recognize the location of a given outdoor query image by finding the same location in a large gallery of outdoor images, we propose to modify this task to enable the localization of indoor images. Specifically, for Inside Out Visual Place Recognition the query set consists of indoor images and the gallery of outdoor images. The challenge thus becomes to identify the portions of the indoor image that contain information that can be used to match it to an outdoor image, and to perform matching based on partial and occluded visual information.

#### 3.1 Problem Formulation

To enable image matching, Visual Place Recognition models are trained by feeding triplets consisting of a query image  $q$ , the easiest positive gallery image  $p^*$  and a set of the most difficult negative gallery images  $\{n_j | j \in \{1, \dots, N\}\}$  [0, 5, 10]. The easiest positive image  $p^*$  is obtained by ranking all gallery images within a distance of 10 meters from query  $q$  and taking the one with the highest visual similarity. The set of negatives is constructed by selecting the  $N$  hardest negatives from a pool of 1000 randomly sampled negatives, where  $N$  is usually set to 10. While all negative images are at least 25 meters away from the query, the hardest negatives are selected based on their visual similarity to the query.

For each training triplet  $\{q, p^*, \{n_j\}\}$ , the goal is to learn a mapping function  $d_\theta$ , such that the distance  $d_\theta(q, p^*)$  between the training query  $q$  and the easiest positive  $p^*$  will be smaller than the distance  $d_\theta(q, n_j)$  between the query  $q$  and all negative images  $n_j$ :

$$d_\theta(q, p^*) < d_\theta(q, n_j), \forall j. \quad (1)$$

Subsequently, the triplets of queries, positives, and negatives are used in a weakly-supervised ranking loss  $L_\theta$ :

$$L_\theta = \sum_j l(d_\theta^2(q, p^*) + m - d_\theta^2(q, n_j)) \quad (2)$$

where  $l$  is the hinge loss  $l(x) = \max(x, 0)$  and  $m$  is the margin.

Considering the Inside Out Visual Place Recognition task, most query images are indoor images where the window only covers a small part of the image. This means that the majority of the visual content of an indoor image cannot be used to identify its location, as it does not match any of the outdoor gallery images. Because of this, the largest part of the image can be considered noise, especially for a model that has never learned to distinguish information about the indoor scene from the outdoor scene this can lead to a deterioration in performance. To measure the potential of existing models for Inside Out Visual Place Recognition, we introduce Inside Out Data Augmentation  $M$  as a training protocol to enforce that models learn to differentiate between indoor and outdoor environments.  $M$  will be used to augment the query only, with the help of a layout of an indoor space  $c$  and a binary mask  $b$ . A detailed explanation of  $M$ ,  $b$ , and  $c$  will be given in Section 3.2.

The goal for Inside Out Visual Place Recognition is to learn a representation in which the augmented query is closer to the easiest positive than it is to all negative images:

$$d_\theta(M(q, b, c), p^*) < d_\theta(M(q, b, c), n_j), \forall j. \quad (3)$$



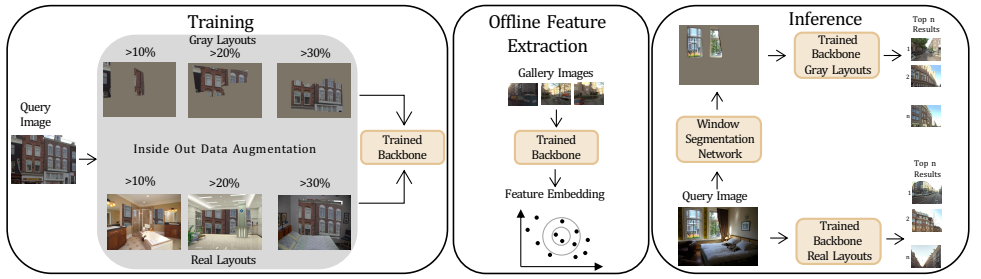


Figure 2: Our approach has three phases: training, offline feature extraction, and inference. For training, we use Inside Out Data Augmentation to create synthetic training data by overlaying queries with gray and real layouts with window proportions of  $>10\%$ ,  $>20\%$ , and  $>30\%$ . For each combination of layout type and proportion, we train three backbone models: NetVLAD, SARE, and SFRS. The training phase is followed by an offline feature extraction phase where the embeddings for the gallery images are calculated. At inference time, we distinguish between two scenarios: 1) the query is sent directly to the trained backbone for real layouts or 2) the query is first processed by a window segmentation network to predict its gray layout after which it is sent to the trained backbone for gray layout. In both scenarios, we obtain the top  $n$  results of most similar gallery images to the queries.

To optimize for this new goal we modify the loss function in equation 2 and we formulate our new loss function as follows:

$$L_{\theta} = \sum_j l \left( d_{\theta}^2(M(q, b, c), p^*) + m - d_{\theta}^2(M(q, b, c), n_j) \right). \quad (4)$$

Ideally we would have a large set of annotated indoor images to measure the potential of existing models for Inside Out Visual Place Recognition. However, such a dataset is currently not available and collecting it would be a costly process, since it should have a large variety of outdoor scenes that are visible through the window and this can only be manually determined. Our training protocol is introduced to overcome this barrier.

### 3.2 Inside Out Data Augmentation

The training protocol we propose, Inside Out Data Augmentation, consists of applying an augmentation function  $M$  to each query image  $q$  to create a synthetic indoor environment with an outdoor view. As input the augmentation function takes a query street-view image  $q$ , a layout  $c$  of an indoor space with a window, and a binary mask  $b$ . The layout and binary mask are obtained from a reference dataset for semantic segmentation with pixel level annotations for windows. The layout represents the indoor scene itself and the binary mask is retrieved from applying a binary encoder on the layout that segments the layout into *window* and *non-window* categories. To perform the augmentation, we do a pixelwise multiplication between the street-view query image and the binary mask and a pixelwise multiplication between the layout and the inverse mapping of the binary mask. By adding these two components, we obtain a new query image with an indoor layout on top of it. The augmentation function is formalized as:

$$M(q, b, c) = q \odot b + c \odot b^{-1}. \quad (5)$$

For the layout, we consider two options: a real layout and a gray layout. The real layout refers to the original indoor scene of the reference image. The gray layout uses the same reference image, but every non-window pixel is matched to the mean color of the ImageNet training set to fully discard any irrelevant visual content. Examples of both layout types can be found in the Supplementary Material A2. During inference the masks to perform gray layout data augmentation are not available, so instead we use a segmentation network to detect the window portions of the query image. An overview of our proposed method is given in Figure 2, highlighting three stages: training, offline feature extraction and inference.

## 4 The Amsterdam-XXXL Dataset

To evaluate the performance of models on IOVPR we present the Amsterdam-XXXL dataset. This dataset consists of three partitions: 1) Outdoor-Ams, 2) Indoor-Ams, and 3) Ams30k. The datapoints are presented in Figure 1(b).<sup>1</sup> Outdoor-Ams and Ams30k consist of panoramic street-view images that are selected from a set that has been made publicly available by the Municipality of Amsterdam.<sup>2</sup> Indoor-Ams consists of two sets of 500 user-generated photos, one consisting of images with creative commons licenses and the other provided by TripAdvisor.<sup>3</sup> Examples of these images for each data source are presented in the Supplementary Material A1. All partitions will be published and made available for research purposes.

**Outdoor-Ams.** This partition consists of 6,467,112 GPS-annotated street-view images taken at 269,463 locations in Amsterdam, where each panoramic image is cut into 24 partially overlapping images of  $480 \times 640$ . With 6.4M images, this is much larger on city-level than widely used VPR datasets and results in a more challenging dataset. Details regarding the coverage of the area and the cutting process are described in the Supplementary Material A3.

**Indoor-Ams.** This partition consists of two query sets for evaluation. To test the performance of our approach properly, at least a few hundred images which show an outdoor scene through a window, but which are taken indoors, are needed. We construct the Indoor-Ams set for this purpose, which consists of a validation and test set of 500 indoor images each with a view of an outdoor scene. These splits are created based on the image license, ensuring consistent licensing within each split. The validation set consists of 500 images in total from platforms with images with Creative Commons licenses (335), namely Flickr, Unsplash, and Traffickcam and self-made pictures taken from public and university buildings(165). The test set consists of 500 images that are selected from an image collection of user images provided by TripAdvisor, consisting of hotels, restaurants, and other visiting areas in Amsterdam.

Indoor-Ams consists of images taken in the time period from 2004 until 2021. We manually verified that all query images (even ones from before the panorama images were taken) can be visually matched to the appearance of the location in the panoramic images. About 5% of these images are night images, where we made sure that the outdoor area is still visible. The window proportions vary from 10-80%. Due to most neighborhoods being mainly residential, we decided not to take pictures there to respect the resident’s privacy. Therefore our dataset only contains public images from residential places with a Creative Commons license. As a consequence, the majority of the images are collected from and taken at public, commercial, and university spaces and released with a license that can be used for research

<sup>1</sup>Map tiles by Stamen Design, under CC BY 3.0. Data by OpenStreetMap.

<sup>2</sup><https://api.data.amsterdam.nl/>

<sup>3</sup>A dataset provided by © 2020 TripAdvisor LLC was used for analysis and visualisation.

purposes. Therefore Indoor-Ams has a spatial bias to the center of Amsterdam. However, this spatial bias should not affect the methods that are trained on Amsterdam-XXXL, since the query sets only serve as a validation and test set.

The annotation process has been done by three annotators who live in Amsterdam. They manually inspected each image and its tag. Details about the selection and annotation process are described in the Supplementary Material A4.

**Ams30k.** This is created for training purposes and is for this reason modelled after Pitts30k [10], where 30k represents the number of gallery images. The total number of panoramas in Ams30k is 2443 which results in almost 59k street-view images. Details about its splits can be found in the Table 1 in the Supplementary Material A5. To guarantee variation in viewpoint, light, and occlusion, the gallery images are from 2018 and the query images from 2019. We select the geographical region 'the Pijp' in Amsterdam as a region for this dataset with good coverage, since it has a large variety in architectural styles, a good balance of houses, restaurants and stores, and green areas (i.e. parks). We define the train, validation, and test set through geographical boundaries, enforcing that the three sets have a similar surface area.

## 5 Experimental Setup

**Inside Out Data Augmentation.** To perform Inside Out Data Augmentation as explained in Section 3.2, we take a subset of ADE20K. This dataset has images that cover a range of scenes and objects and are pixel annotated with 150 different classes. Our subset consists of indoor images with a window, more specifically 4687 images in the train set and 461 images in its validation set. After reshaping these images to  $480 \times 640$ , we apply our layout function from Equation 5 and create two sets by using two layouts, the original layout and a gray layout. Examples of the output are shown in Supplementary Material A2. To measure the effect of the percentage of available window pixels for training and inference, we create three different sets of images per data augmentation type depending on the window proportion of the layouts, namely  $>10\%$ ,  $>20\%$ , and  $>30\%$ . In this setup,  $> k\%$  stands for  $k\%-100\%$ .

**Backbones.** We evaluate four model architectures as possible backbones: NetVLAD [11], SARE [12], SFRS [5], and Patch-NetVLAD which were briefly discussed in Section 2. NetVLAD is a trainable VLAD layer, inspired by the Vector of Locally Aggregated Descriptors [13], a popular descriptor pooling method for instance retrieval before the Deep Learning era. It generates global features of images by aggregating feature maps by summing residuals between features and cluster centers. SARE [12] which stands for Stochastic Attraction-Repulsion Embedding aims to capture the intra-place attraction and inter-place repulsion in the embedding space. This is achieved by minimizing the KL divergence between the learned and the actual probability distributions. SFRS [5] refers to Self-supervising Fine-grained Region Similarities. It splits image feature maps in 4 half and 4 quarter regions and computes similarity scores between these regions and the query during training. These similarities are used to supervise the network through a soft cross-entropy loss on top of the triplet loss, to enhance the learning of local features. Patch-NetVLAD [14] combines global with local descriptor methods. It uses NetVLAD descriptors to retrieve the top- $k$ , with default  $k = 100$ , most likely matches and performs local matching of patch-level descriptors at multiple scales to reorder this top- $k$ . Then a RANSAC-based scoring method is used for higher retrieval performance. We test a new model architecture by combining the patch-based local feature method of Patch-NetVLAD with the SFRS backbone and show initial

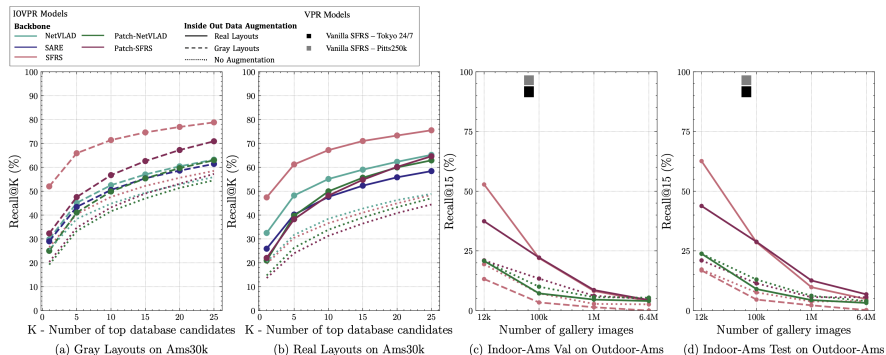


Figure 3: Results of Inside Out Visual Place Recognition on the Amsterdam-XXXL dataset. Figure (a) and (b) present the results for gray and real layouts on Ams30k, where the dotted pink line indicates the results of not using Inside Out Data Augmentation (i.e., no layouts). (c) and (d) show the results on Indoor-Ams for the val and test set respectively, evaluated on subsets of Outdoor-Ams. This is best viewed in color.

results on Outdoor-Ams. We name this method *Patch-SFRS*.

**Window Segmentation Network.** To obtain a gray layout at inference time for the gray layout model, we finetune the network UPerNet50 [43] for window segmentation. For finetuning we use the same subset of the ADE20K dataset as for Inside Out Data Augmentation [44]. After finetuning UPerNet50 on two classes (i.e., window and non-window) for three epochs we obtain an accuracy score of 95% on the validation set of ADE20K.

**Evaluation Criteria.** As is typical for VPR we report our results using the Recall@K metric [43]. More specifically, the query image is seen as correctly localized if at least one of the top K retrieved gallery images is within a radius of  $n$  meters from the ground truth position of the query. The radius for VPR is usually set to 25 meters. Since the street-view images can be taken further away from the building, we increase the distance to 50 meters for our experiments on Indoor-Ams. Regarding the recall, we set our K values for our small sets to 1, 5, 10, 15, 20, and 25. For our large set of 6.4M images, we add 50, 75, and 100.

## 6 Results

**Inside Out Data Augmentation.** In our experiments, we analyze the effect of Inside Out Data Augmentation on our queries during training. In particular we study the difference between using gray and real layouts. In addition, we take models trained on Ams30k without layouts, to measure the contribution of Inside Out Data Augmentation. In Figure 3, the results are presented on Ams30k for gray (a) and real layouts (b) with a window proportion of  $>20\%$ . The results for  $>10\%$  and  $>30\%$  and the corresponding tables can be found in Supplementary Material B Tables 2-3. The best performing model in both setups is SFRS with a window proportion of  $>20\%$ , where the model for gray layouts (66% Recall@5) performs slightly better than for real layouts (61% Recall@5). Comparing this to SFRS that has not been trained with Inside Out Data Augmentation, we notice a difference of 25% with the best model for gray layouts and 31% with real layouts. We conclude that the use of Inside Out Data Augmentation significantly helps in localizing indoor images from Ams30k and that the benefits of the gray and real layouts are similar.

**Window Proportion.** We study the effect of the proportion of available window pixels in images by setting a lower bound on the amount of visible window pixels. For both the gray and the real layouts of Ams30k we compare three lower bounds, namely  $>10\%$ ,  $>20\%$ , and  $>30\%$  of the full image. We study how well these different training setups generalize by training on these three different sets and testing on a test set with a window proportion of at least  $>10\%$  of the full image. Considering the different window proportions, Figure B.1 in the Supplementary Material indicates that the  $>10\%$  proportion performs worst out of the three proportions. We expect that for this proportion there is insufficient information for models to learn about outdoor environments. We also see that for gray and real layouts, the models trained on a window proportion of  $>20\%$  and  $>30\%$  perform well on the test set, despite the test set containing images with a window proportion of  $>10\%$ . Overall the highest performance is obtained by models trained on images with a  $>20\%$  proportion. For this reason we recommend training on images with at least a decent window proportion of 20% to make sure that the network has enough opportunities to optimize for this task.

**Backbone Comparison.** For each of the chosen backbones NetVLAD, SARE, and SFRS we measure the performance on different layouts and window proportions by employing three different backbones, two types of layouts, and three window proportions. Additionally, we evaluate Patch-NetVLAD and Patch-SFRS on Ams30k queries with gray layouts and real layouts. With respect to the choice of backbone, we see that a strong model backbone contributes more to the performance than the window proportion that is used during training, except for when the model has not seen any layouts during training. In general we see that the SFRS model outperforms the other backbones in every setup. An unexpected result is that Patch-SFRS performs worse than SFRS on Ams30k for both real and gray layouts. When not using any layouts during training, Patch-SFRS does also not improve the performance. From Figure 3(d) we see that Patch-SFRS performs worse than SFRS on a small scale subset. We expect that the higher overall performance of SFRS might be due to its fine-grained approach, which utilizes region-similarities, allowing SFRS to focus on the information rich window parts of the images. Hence we recommend future approaches to focus on fine-grained techniques.

**Large Scale Visual Place Recognition.** To get an understanding of how well our proposed approach scales to a large gallery, we evaluate the performance of using indoor images captured by users. Some success and failure cases from our best performing model on the test set of Indoor-Ams are presented in Figure 4. The difficulty of VPR is directly related to the number of gallery images, due to the potential for false positives (i.e. distractor images being ranked high) increasing with gallery size. With 6.4M images Amsterdam-XXXL is 80 times larger than widely used VPR datasets on city-level (e.g., Tokyo 24/7, Pitts250k, and city-splits of Mapillary Street-level Sequences Dataset) and thus includes many more distractors. To show the effect of distractors we zoom in on the reported results by creating three subsets of Outdoor-Ams of 12k, 100k, and 1M images. In Figure 3(c) we present the Recall@15 results of SFRS, Patch-NetVLAD, and Patch-SFRS trained with and without our data augmentation methods of 20% for the validation set of Indoor-Ams, Figure 3(d) shows the same results for the test-set. As demonstrated in this figure, when evaluated with fewer distractors (as is typical for VPR) SFRS with real layouts shows a large improvement over the SOTA of up to 50% and 30% for the 12k and 100k subsets. However, if the number of distractors is increased to well beyond what is common for standard VPR, the performance for both methods collapses. For a high number of distractors we see that reranking the top-100 results improves the performance and therefore Patch-SFRS performs best.



Figure 4: Qualitative results of Inside Out Visual Place Recognition success and failure cases. These results demonstrate the potential of our approach in returning the correct location (sometimes even multiple times), but also how challenging it is to find the correct match when very little is visible of the outdoor scene.

To highlight the differences between IOVPR and VPR, we add a black and gray box, for vanilla SFRS models trained on Tokyo 24/7 and Pitts250k respectively as reported in [5], to Figure 3. On VPR this model almost achieves 100% correctness. The same model has a score of nearly 8% on a similarly sized subset of Outdoor-Ams. Although our proposed modifications boost this to almost 30%, existing VPR methods are clearly not suited for IOVPR. Given these results we conclude that the scale of Outdoor-Ams remains an interesting challenge for Visual Place Recognition.

## 7 Conclusion

In this work, we have introduced the novel task Inside Out Visual Place Recognition and a large-scale dataset Amsterdam-XXXL to train and evaluate this task on. Additionally, we proposed Inside Out Data Augmentation as a new training protocol to extend Visual Place Recognition for localizing indoor images. The results on Ams30k and subsets of Outdoor-Ams show the potential of Inside Out Data Augmentation, on both synthetic and real indoor images. However, this task remains a challenge for existing Visual-only Place Recognition methods on a large-scale dataset with millions of images. This opens up a range of possible future directions, including multimodal approaches.

We are aware of the fact that techniques for Inside Out Visual Place Recognition have the potential to cause harm related to the privacy of individuals when misused. While we condemn any non-ethical applications, we hope to inspire the research community to work on this task for good purposes and to help advance the state-of-the-art for Visual Place Recognition on indoor environments.



## Acknowledgements

We thank Jos de Winde and TripAdvisor for their help in creating the dataset. This research was supported by the Nationale Politie. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

## References

- [1] R. Arandjelović, P. Gronat, A. Torii, T. Pajdla, and J. Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *CVPR*, 2016.
- [2] Sudong Cai, Yulan Guo, Salman Khan, Jiwei Hu, and Gongjian Wen. Ground-to-aerial image geo-localization with a hard exemplar reweighting triplet loss. In *ICCV*, 2019.
- [3] Bingyi Cao, Andre Araujo, and Jack Sim. Unifying deep local and global features for efficient image search. *CoRR*, abs/2001.05027, 2020. URL <https://arxiv.org/abs/2001.05027>.
- [4] Dzung Doan, Yasir Latif, Tat-Jun Chin, Yu Liu, Thanh-Toan Do, and Ian D. Reid. Scalable place recognition under appearance change for autonomous driving. In *ICCV*, 2019.
- [5] Yixiao Ge, Haibo Wang, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-supervising fine-grained region similarities for large-scale image localization. In *ECCV*, 2020.
- [6] Carlos Gonzales. Creating Impact: A Year On Stop Child Abuse — Trace An Object. <https://www.bellingcat.com/news/uk-and-europe/2020/04/22/creating-impact-a-year-on-stop-child-abuse-trace-an-object/>, 2020. [Accessed: 2010-01-25].
- [7] Stephen Hausler, Sourav Garg, Ming Xu, Michael Milford, and Tobias Fischer. Patchnetvlad: Multi-scale fusion of locally-global descriptors for place recognition. In *CVPR*, 2021.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [9] Sixing Hu, Mengdan Feng, Rang M. H. Nguyen, and Gim Hee Lee. Cvm-net: Cross-view matching network for image-based ground-to-aerial geo-localization. In *CVPR*, 2018.
- [10] U.S. Immigration and Customs Enforcement. 123 sexually exploited children identified by HSI during 'Operation Sunflower'. <https://www.ice.gov/news/releases/123-sexually-exploited-children-identified-hsi-during-operation-sunflower>, 2015. [Accessed: 2021-01-25].
- [11] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. Hamming embedding and weak geometric consistency for large scale image search. In David A. Forsyth, Philip H. S. Torr, and Andrew Zisserman, editors, *ECCV*, 2008.

- [12] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez. Aggregating local descriptors into a compact image representation. In *CVPR*, 2010.
- [13] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE TPAMI*, 2011.
- [14] Hyo Jin Kim, Enrique Dunn, and Jan-Michael Frahm. Learned contextual feature reweighting for image geo-localization. In *CVPR*, 2017.
- [15] Jan Knopp, Josef Sivic, and Tomáš Pajdla. Avoiding confusing features in place recognition. In Kostas Daniilidis, Petros Maragos, and Nikos Paragios, editors, *ECCV*, 2010.
- [16] Tsung-Yi Lin, Yin Cui, Serge J. Belongie, and James Hays. Learning deep representations for ground-to-aerial geolocalization. In *CVPR*, 2015.
- [17] Liu Liu, Hongdong Li, and Yuchao Dai. Stochastic attraction-repulsion embedding for large scale image localization. In *ICCV*, 2019.
- [18] Raul Mur-Artal, J. M. M. Montiel, and Juan D. Tardós. ORB-SLAM: A versatile and accurate monocular SLAM system. *IEEE Trans. Robotics*, 31(5), 2015.
- [19] Hyeonwoo Noh, Andre Araujo, Jack Sim, Tobias Weyand, and Bohyung Han. Large-scale image retrieval with attentive deep local features. In *ICCV*, 2017.
- [20] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *CVPR*, 2007.
- [21] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Lost in quantization: Improving particular object retrieval in large scale image databases. In *CVPR*, 2008.
- [22] F. Radenović, A. Iscen, G. Toliás, Y. Avrithis, and O. Chum. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *CVPR*, 2018.
- [23] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [24] Abby Stylianou, Hong Xuan, Maya Shende, Jonathan Brandt, Richard Souvenir, and Robert Pless. Hotels-50k: A global hotel recognition dataset. In *AAAI*, 2019.
- [25] Giorgos Toliás and Hervé Jégou. Visual query expansion with or without geometry: Refining local descriptors by feature aggregation. *Pattern Recognit.*, 47(10):3466–3476, 2014. doi: 10.1016/j.patcog.2014.04.007. URL <https://doi.org/10.1016/j.patcog.2014.04.007>.
- [26] Akihiko Torii, Josef Sivic, Tomáš Pajdla, and Masatoshi Okutomi. Visual place recognition with repetitive structures. In *CVPR*, 2013.
- [27] Akihiko Torii, Relja Arandjelovic, Josef Sivic, Masatoshi Okutomi, and Tomáš Pajdla. 24/7 place recognition by view synthesis. In *CVPR*, 2015.
- [28] Nam N. Vo and James Hays. Localizing and orienting street views using overhead imagery. In *ECCV*, 2016.

- [29] Ziqi Wang, Jiahui Li, Seyran Khademi, and Jan van Gemert. Attention-aware age-agnostic visual place recognition. In *ICCV Workshops*, 2019.
- [30] Frederik Warburg, Soren Hauberg, Manuel Lopez-Antequera, Pau Gargallo, Yubin Kuang, and Javier Civera. Mapillary street-level sequences: A dataset for lifelong place recognition. In *CVPR*, 2020.
- [31] T. Weyand, A. Araujo, B. Cao, and J. Sim. Google Landmarks Dataset v2 - A Large-Scale Benchmark for Instance-Level Recognition and Retrieval. In *CVPR*, 2020.
- [32] Scott Workman, Richard Souvenir, and Nathan Jacobs. Wide-area image geolocalization with aerial reference imagery. In *ICCV*, 2015.
- [33] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, 2018.
- [34] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017.