

CrowdDriven: A New Challenging Dataset for Outdoor Visual Localization

Ara Jafarzadeh¹ Manuel López Antequera² Pau Gargallo² Yubin Kuang² Carl Toft¹
 Fredrik Kahl¹ Torsten Sattler³

¹Chalmers University of Technology ²Facebook ³Czech Technical University in Prague

Abstract

Visual localization is the problem of estimating the position and orientation from which a given image (or a sequence of images) is taken in a known scene. It is an important part of a wide range of computer vision and robotics applications, from self-driving cars to augmented/virtual reality systems. Visual localization techniques should work reliably and robustly under a wide range of conditions, including seasonal, weather, illumination and man-made changes. Recent benchmarking efforts model this by providing images under different conditions, and the community has made rapid progress on these datasets since their inception. However, they are limited to a few geographical regions and often recorded with a single device. We propose a new benchmark for visual localization in outdoor scenes, using crowd-sourced data to cover a wide range of geographical regions and camera devices with a focus on the failure cases of current algorithms. Experiments with state-of-the-art localization approaches show that our dataset is very challenging, with all evaluated methods failing on its hardest parts. As part of the dataset release, we provide the tooling used to generate it, enabling efficient and effective 2D correspondence annotation to obtain reference poses.

1. Introduction

Visual localization is the problem of estimating the position and orientation from which an image was taken, *i.e.*, its camera pose, with respect to the scene. Visual localization is a vital part of many computer vision and robotics applications such as self-driving cars, service robots such as gardening robots, and augmented/mixed/virtual reality.

Most visual localization methods rely heavily on local descriptors for pose estimation, and finding 2D-3D matches between images is a fundamental part of it. However, the trade-off between the local descriptors' discriminative power and their invariance limits their performance under changing conditions. On the other hand, in practice, changes in the scene are unavoidable, and there is a need for visual localization methods to be robust to them. Traditionally, ground truth poses for localization have been obtained via Structure-from-Motion (SfM) [35, 62, 66]. Yet, SfM it-

self relies on local descriptors and matching. This makes it extremely difficult to generate localization benchmarks where feature matching using traditional feature descriptors does not work, *e.g.*, in the presence of day/night and strong viewpoint changes. Benchmark datasets for localization under changing conditions such as Aachen, CMU Seasons, and Robotcar [58] rely on manual annotations to be able to provide ground truth poses. While these datasets provide interesting challenges, they have mostly been captured in controlled conditions. However, there is little control over the capture in many applications such as Autonomous Driving, collaborative AR/MR, or crowd-sourced mapping.

In this paper, we have constructed a dataset revisiting the common challenges that could be seen in different environments. We first actively mined a crowd-sourced database for image sequences where classic SfM approaches fail. To generate reliable poses, we have relied on human annotations and have created 40 sets of image sequences with diverse visual changes. The dataset and tools used for its creation are available at mapillary.com.

This paper makes the following contributions: (1) a workflow to mine and annotate challenging image sequences for benchmarking visual localization. Our approach explicitly takes pose uncertainty into account during the annotation process. (2) the CrowdDriven dataset, a geographically diverse and challenging dataset with reliable poses that covers various scenarios in illumination, weather, seasonal, and viewpoint changes. (3) experiments with state-of-the-art baselines showing CrowdDriven presents challenges that existing methods cannot handle.

2. Related Work

Visual localization methods aim to estimate the full camera pose and can be categorized based on their map representation: (1) **Image-based representations** encode each image as a feature vector [3, 17, 29, 54, 72, 73]. Using image retrieval and the known poses of the database images, the pose of a test image can either be approximated via the poses of the top-retrieved database images [61, 83] or computed precisely based on relative poses [59, 85]. (2) **Local feature-based representations** rely on classical or learned local feature descriptors and multiple view ge-

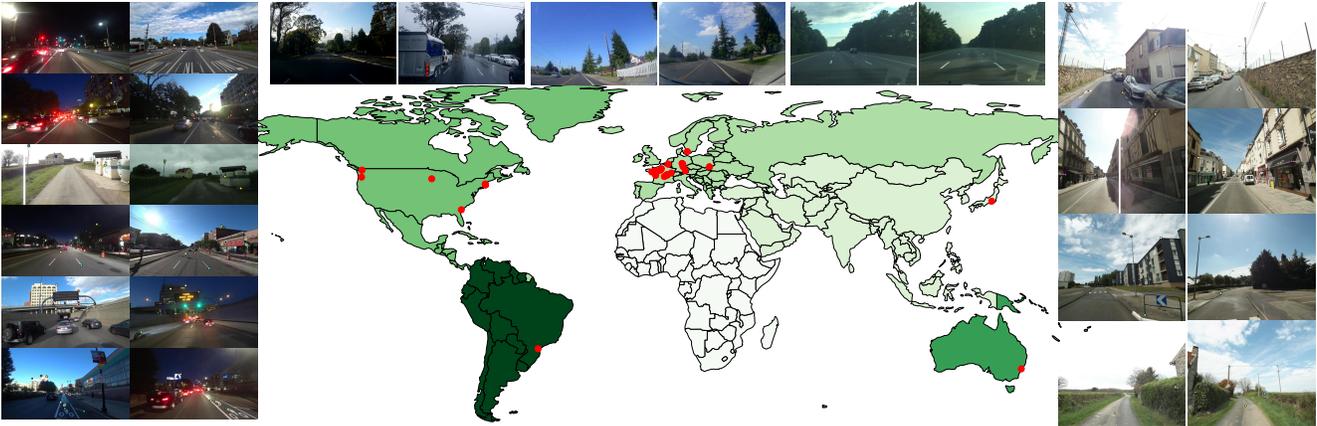


Figure 1. Easy (top), Medium (left) and Hard (right) subsets of our CrowdDriven dataset. Our dataset includes reference poses derived from hand-annotated control points, enabling the benchmarking of localization algorithms in new scenarios such as extreme viewpoint (180°) changes. CrowdDriven includes sequences from over the globe captured with a diversity of cameras.

ometry [43, 44, 55–57, 64, 68, 84]. These approaches build a 3D model of the scene and then estimate the camera pose based on 2D-3D matches between local features in a test image and the 3D model [23, 32, 41]. Such approaches, especially using learned features, constitute the current state-of-the-art in terms of visual localization in changing conditions [22, 30, 51, 52, 63] and we evaluate and analyze their performance on our benchmark. (3) **Learning-based representations** typically represent the scene using convolutional neural networks (CNNs). Pose regression techniques directly regress the camera pose for a given input image [10, 38, 39, 79]. These methods have been shown not to perform better than image retrieval methods [61]. Learning-based approaches that do not learn the full localization pipeline but only the 2D-3D matching part [6–9, 14, 15, 21, 53, 65] achieve state-of-the-art performance in small scenes [5, 80]. While most of these approaches struggle to handle conditions not seen during training, [53] generalizes well.

While localization approaches use metric maps, **visual place recognition** methods typically rely on topological scene representations [19, 20, 27, 28, 49, 50, 78], and thus do not directly provide a camera pose estimate. Using image retrieval and related techniques, they identify the *places* depicted in test images, where *places* are defined as sets of database images. Place recognition can be used to guide visual localization by identifying which parts of a scene are visible in a test image [37, 51, 60]. For an overview over place recognition techniques, we refer the reader to [26, 47].

Tab. 1 provides an overview over **datasets commonly used to measure localization and place recognition performance under changing conditions**. Coarse-scale location information for place recognition datasets such as Nordland [67], Pittsburgh [74], Tokyo 24/7 [72], and Mapillary Street-Level Sequences [81] can be obtained relatively easily via GPS measurements. In contrast, obtaining 6DOF poses for localization requires considerable manual effort

as the classical approaches used to automatically obtain ground truth fail under challenging conditions. RIO10 [80] focuses on changes in indoor scenes. Long-term outdoor localization datasets such as Aachen Day-Night [58, 60], CMU Seasons [4, 58], and RobotCar Seasons [48, 58] only cover a few geographical locations and using a small number of cameras. In contrast, even though our new benchmark dataset is not the largest, it has much more geographical diversity. While the other datasets were captured using only a few cameras, our image sequences are taken by a large number of different camera types and photographers.

Other localization datasets include the indoor 7 Scenes [65], 12 Scenes [76], and InLoc [70] datasets and the outdoor Dubrovnik [43], Rome [25], Vienna [37], San Francisco [16, 44], and Cambridge Landmarks [39] datasets. None of these datasets is designed to measure the impact of changing conditions on localization performance [58].

The focus of recent datasets from self-driving car companies such as Lyft [40], Waymo [2], Aptiv (nuScenes) [11], and Baidu (Apolloscape) [36] is benchmarking 2D/3D object detection based on multimodal sensor data in limited geographical areas, semantic segmentation, and depth estimation. An exception is Baidu’s Apolloscape which provides a localization benchmark track in small urban areas.

Another shortcoming of the datasets mentioned above is that the changes in viewpoints are limited. The query images are typically taken from similar vantage points as the reference images. Our dataset contains image sequences that are, for instance, taken in opposite directions with different cameras. This case commonly happens during multi-session or collaborative capture of man-made environments, where paths can be traversed in two directions.

Based on the shortcomings of current localization datasets, we generate a diverse dataset for long-term visual localization (1) containing different challenging scenarios such as day-night, seasonal, daylight illumination, strong

Dataset	Scene Type					#images		condition changes								6DOF query poses	# locations
	urban	suburban	natural	country road	indoors	reference	query	weather	seasonal	strong viewpoint	day/night	intrinsic	snow	rain	sequential		
Nordland [67]	✓	✓	✓			14k	16k		✓						✓		1
Pittsburgh [74]	✓					254k	24k								✓		1
Tokyo 24/7 [72]	✓					174k	1k		✓						✓		1
NCLT [13]	✓		✓		✓	3.8M			✓		✓				✓		1
Extended CMU Seasons [4, 58]	✓	✓	✓	✓		61k	57k		✓						✓		1
RobotCar Seasons [48, 58]	✓	✓				20k	12k	✓	✓		✓		✓		✓		1
Aachen Day-Night [58, 60]	✓					3k	922				✓			✓	✓		1
RIO10 [80]					✓	53k	200k								✓		1
7-scenes [31]					✓	26k	17k								✓		7
12-scenes [34]					✓	17k	5.8k								✓		12
Cambridge [39]	✓					8.4k	4.8k								✓		3
Dubrovnik [43]	✓					6k	0.8k								✓		1
San Francisco [16]	✓					610k	0.4k								✓		1
Rome [43]	✓					15k	1k								✓		1
Vienna [37]	✓					1k	0.2k								✓		1
InLoc [70, 82]	✓				✓	9.9k	0.3k								✓		5
CrowdDriven	✓	✓	✓	✓	✓	1.3k	1.7k	✓	✓	✓	✓	✓	✓	✓	✓	✓	26

Table 1. Comparing localization datasets: CrowdDriven is the most diverse in term of scene types and changes in viewing conditions.

viewpoint and man-made changes, (2) with reliable 6DOF camera poses based on human annotations, (3) using crowd-sourced data to imitate real driving situations, (4) with reliable camera calibration and geographical information.

3. CrowdDriven: Dataset Creation

Historically, visual localization algorithms have been evaluated on crowd-sourced datasets collected from internet photo collection websites such as Flickr, *e.g.*, the Dubrovnik [43], Rome [43], and Landmarks 1k [44] datasets. This way to acquire datasets scales very well to the use of different camera types and geographical locations. However, since the images have been taken by handheld cameras around famous landmarks, these datasets are not suitable for measuring visual localization performance in an autonomous driving scenario, which is our main focus (even though a small part of our dataset was captured by pedestrians and bicycles). In contrast, recently released datasets tailored for this task, *e.g.*, the RobotCar Seasons [48, 58] and the (extended) CMU Seasons [4, 58] datasets, contain only one or a few larger locations, use only a few cameras, and have been taken by experts. Our new dataset, CrowdDriven, is designed to cover a wide range of locations, visual conditions (*e.g.*, seasonal, day-night changes or changes in facades, presence/absence of humans (man-made changes)), and camera types by tapping into a large database of crowd-sourced images.

3.1. Data Source

To maximize the diversity and geographical coverage, we use **Mapillary**, a collaborative street-level imagery platform that hosts more than 1 billion images collected by members of their community while driving or walking on public spaces and roads. It covers most countries with hundreds of camera models in varying times of day and weather conditions. Its data is thus well-suited to evaluate problems in conditions similar to those faced in self-driving scenarios, since most images are captured with consumer-grade devices such as smartphones, action cameras, and dashcams.

3.2. Sequence Selection

Images in Mapillary are grouped into ordered sequences based on the photographer ID and the capture time. Each

image’s GPS position and compass are available, and the database can be queried to find images near a particular position. To generate a dataset that is challenging for current state-of-the-art localization algorithms, we attempt to find pairs of neighboring sequences to be localized with respect to each other. We categorize each sequence pair into a preliminary difficulty level based on the success of traditional SfM algorithms on jointly reconstructing the pair.

We begin by querying the Mapillary database to find pairs of sequences covering a wide range of geographical locations and appearances. The pairs are selected to satisfy the following criteria: 1. Sequence length between 40 and 60 images. 2. Minimum sequence density of 0.2 images/m. 3. Maximum sequence-to-sequence distance of 3m. 4. Each sequence can be reconstructed individually using SfM.

Note that we deliberately chose to focus on small scenes rather than attempting to collect data for a larger spatial area. In practice, pose priors such as GPS are used to limit the search space during localization. Rather than artificially making the problem harder by ignoring such priors, we are interested in finding realistic hard examples with small scenes. Our dataset might not be as large as previous benchmarks. However, a large fraction of test images in previous benchmarks can be localized accurately. In contrast, our dataset contains many very challenging scenarios that current state-of-the-art methods cannot handle.

After collecting the sequence pairs, we run SfM on each of them. We assign a difficulty rating to each pair depending on the result of SfM and the relative orientation of the sequences: Sequence pairs whose joint reconstruction succeeds¹ are categorized as **easy** and are shown in Fig. 1 (top). If the joint reconstruction fails, we look at the average view direction of the sequences to categorize them into **medium** (Fig. 1 - left) if they have similar orientations (less than a 45° difference), or **hard** (Fig. 1 - right) if they have different orientations (more than a 45° difference).

This preliminary categorization is based on the following observations: sequences with similar orientations are difficult to match mostly due to changes in appearance (*e.g.* illumination or seasonal change) and state-of-the-art methods

¹All images are included in the reconstruction with a sufficiently large number of well-distributed inliers

such as [51–53] are able to handle such changes quite well. In contrast, sequence pairs with different view directions are harder to match due to the addition of dramatic appearance changes of scene parts due the large viewpoint change and low visual overlap [63]. As our experiments in Sec. 5 show, this preliminary classification aligns well with how current state-of-the-art methods perform on our dataset.

Tab. 2 shows statistics over the sequence pairs in Crowd-Driven. Many capture conditions that are challenging for current localization algorithms are included in Crowd-Driven. The medium category consists of different scene types shown under conditions such as day-night changes, rain, snow and other seasonal variations that impact the scene geometry. The main focus of the hard category is the combination of large viewpoint changes with less extreme variations in illumination, seasons, and weather conditions.

On the dataset size. Similar to existing datasets [58,70,80], CrowdDriven only provides reference images taken under a single condition for each scene with test images.² As shown in Tab. 1, CrowdDriven contains fewer images than most such datasets, *e.g.*, the CMU and RobotCar Seasons datasets from [58]. However, this does not imply that our dataset is too small for learning-based methods such as camera pose [10,38,39,79] or scene coordinate [6–9,14,15,65] regression. These approaches regress the camera pose from a given image and a 3D point coordinate from an image patch, respectively. Both regression tasks are instance-level problems. As such, the number of images in local scene parts is more important for their performance in these parts than the absolute number of images in a dataset. Extending CrowdDriven in terms of the absolute number of images by adding additional scenes is easy. However, doing so is unlikely to improve the performance in existing scenes. Extending CrowdDriven by adding more images to existing scenes is hard as such images are simply not available: while crowd-sourced image capture allows us to obtain images from a diverse set of scenes from multiple continents, there is no control over how much data is acquired per scene. Still, the number of reference images per scene should be sufficient for learning-based techniques.

3.3. Reference Pose Generation

To use the sequences for benchmarking visual localization algorithms, two processing steps are required: 1. estimating intrinsic camera calibrations and reference poses for all images in a common coordinate system, such that distances can be measured in meters. 2. subdividing the datasets into reference (training) images and test images. For (2), we use a simple strategy: for each pair of sequences, the larger sequence defines the reference images, and the images from the smaller sequence are used for testing.

²The extended CMU and RobotCar Seasons provide training images taken under multiple conditions for a set of locations that does not overlap with the scene parts depicted in the test images.

Easy Datasets are those where both sequences can be reconstructed in a common coordinate frame by off-the-shelf SfM using SIFT [46] features. We follow common practice [25,39,44,58] and use the camera poses and intrinsics estimated during the reconstruction process to define our reference poses. After running SfM on the sequences, we visually inspect the 3D models. If the two sequences are *aligned*, *i.e.*, there are no duplicate 3D points, and the camera poses look visually correct, we scale the SfM poses based on available GPS data to (approximately) recover the scale of the scene. For this category, we expect state-of-the-art localization approaches to perform well given that SIFT features are sufficiently powerful to register the sequences.

We used OpenSfM [1] as our SfM pipeline. Experiments with COLMAP [62] on our data showed similar results.

Tab. 2 provides statistics over the 13 sequence pairs in the easy category. As can be seen, diversity in scenes such as weather and illumination changes is covered in different environments such as roads, suburban and urban areas. Fig. 1 (top) shows sample images from this category.

Medium and Hard Datasets are those where SfM fails to register the sequence pairs due to large changes in appearance (medium) and/or viewpoint (hard). Both medium and hard datasets are thus significantly more challenging for existing visual localization algorithms. Thus, these datasets will be most interesting to the community.

To obtain reference poses, we first reconstruct each sequence individually using SfM and recover the scale of the models using known GPS data from Mapillary. Next, we manually annotate corresponding pixel positions between images from the two sequences (Figs. 2 & 3). These annotations define manual tracks within each sequence and between sequences, which we will call control points (CPs). On average, 12 different control points have been annotated for each dataset (sequence pair), where each such point has been observed on average in 10 images.

To align the sequences, we first obtain the 3D position of the CPs in the reference frame of each reconstruction by triangulating the 2D annotations. The corresponding 3D points (as per the annotation) provide us with 3D-3D matches that we use to compute an initial registration as a similarity transformation, bringing both reconstructions into a common reference frame. Given this initial alignment, we perform bundle adjustment [75] over all camera poses and intrinsics, the 3D points triangulated from SIFT features, and the 3D points triangulated from annotations.

Fig. 2 shows the quality of the alignment when using only GPS constraints (middle) and when using manual annotations (right). For clarity of visualization, we show the scenes as dense Multi-View Stereo point clouds, although the registration is performed using only sparse points.

We do not use dense point clouds for alignment as they cannot always be obtained solely from images, *e.g.*, due to

scene type	identifier	CP reproj. err.	pos. STD (m)	# test images	# ref. images.	ref. conditions	test conditions	considerable changes	foliage
road	Sydney	-	-	14	28	day, partly cloudy	day, rain,	illumination	
	Massachusetts1	-	-	25	49	day, partly cloudy	day, overcast,	illumination	✓
	Poing	-	-	20	56	day, clear sky	day, overcast	illumination	✓
	Washington	-	-	10	29	day, clear sky	day, cloudy	illumination	
	Melbourne	-	-	12	36	day, cloudy	day, overcast	illumination	
	Burgundy2	0.03% (0.21 px)	-	50	50	day, sunny	day, rain	illumination, rain	
	Thuringia	0.07% (0.46 px)	0.13	11	17	day, sunny	day, clear sky	illumination	✓
	Massachusetts2	0.07% (0.45 px)	0.07	24	35	day, overcast	night	illumination, day-night	
	Besançon2	0.01% (0.07 px)	0.06	50	50	day, overcast	day, cloudy	illumination, strong viewpoint	✓
	Besançon4	0.02% (0.13 px)	0.06	50	50	day, cloudy	day, overcast	strong viewpoint, illumination	✓
	Besançon3	0.04% (0.24 px)	0.04	50	50	day, overcast	day, vegetated, strong viewpoint	illumination	✓
	Brittany	0.11% (0.69 px)	0.19	31	53	day, sunny	day, partly cloudy	strong viewpoint	✓
	suburban	Portland	-	-	21	41	day, clear sky	day, overcast	illumination
Curitiba		-	-	19	20	day, cloudy	day, overcast	illumination	
Tsuru		-	-	9	26	day, cloudy	day, overcast	illumination	
Clermont-Ferrand		-	-	15	21	day, sunny	day, overcast	illumination	
Savannah		-	-	18	56	day, clear sky	day, cloudy	illumination	
Subcarpathia		-	-	17	32	day, cloudy	day, overcast	snow, seasonal	
Massachusetts3		0.03% (0.17 px)	0.04	44	56	day, clear sky	night	day-night	✓
Skåne		0.02% (0.10 px)	0.03	20	24	day, cloudy	day	small viewpoint, illumination	
Angers2		0.04% (0.26 px)	0.05	46	47	day, cloudy	day, clear sky	strong viewpoint, illumination	
Ile-de-France		0.02% (0.13 px)	0.04	50	50	day, sunny	day, cloudy	strong viewpoint, illumination	
Orleans2		0.04% (0.22 px)	0.05	31	31	day, clear sky	day, sunny	strong viewpoint, illumination	
Pays de la Loire		0.05% (0.32 px)	0.03	42	58	day, cloudy	day, overcast	strong viewpoint	✓
Bourges		0.05% (0.34 px)	0.06	22	23	day, partly cloudy	day, clear sky, illumination	strong viewpoint	✓
Nouvelle-Aquitaine2	0.05% (0.34 px)	0.05	45	46	day, sunny	day, sunny	strong viewpoint	✓	
urban	Muehlhausen	-	-	10	22	day, cloudy	day, overcast	slight illumination	
	Bayern	-	-	26	26	day, cloudy	day, overcast	illumination	
	Boston5	0.04% (0.27 px)	0.08	34	47	day, sunny	night	day-night	
	Boston1	0.04% (0.27 px)	0.05	47	53	day, sunny	night	day-night	
	Boston3	0.12% (0.78 px)	0.52	31	40	day, clear sky	day, clear sky	small viewpoint	
	Massachusetts4	0.06% (0.40 px)	0.06	39	45	day, clear sky	night	day-night	
	Boston2	0.03% (0.19 px)	0.03	49	51	day, clear sky	night	day-night	
	Boston4	0.04% (0.27 px)	0.05	34	41	day, sunny	day, sunny	day-night	
	Cambridge	0.11% (0.69 px)	0.05	33	35	night	day, sunny	day-night	
	Le-Mans	0.04% (0.28 px)	0.04	49	51	day, overcast	day, overcast	illumination, strong viewpoint	
	Nouvelle-Aquitaine1	0.06% (0.40 px)	0.05	50	50	day, overcast	day, overcast	strong viewpoint, seasonal, snow	
	Angers1	0.08% (0.51 px)	0.08	47	49	day, cloudy	day, clear sky	strong viewpoint	
	Orleans1	0.05% (0.29 px)	0.11	33	34	day, sunny	day, clear sky	strong viewpoint, illumination	
Leuven	0.08% (0.50 px)	0.04	20	21	day, cloudy	day, cloudy	strong viewpoint		

Table 2. Statistics of CrowdDriven: scene type, identifier, median reprojection error of control points (CPs), number of test and reference images, reference and test conditions and changes. Preliminary categorization: easy: light gray; medium: gray; hard: dark gray.

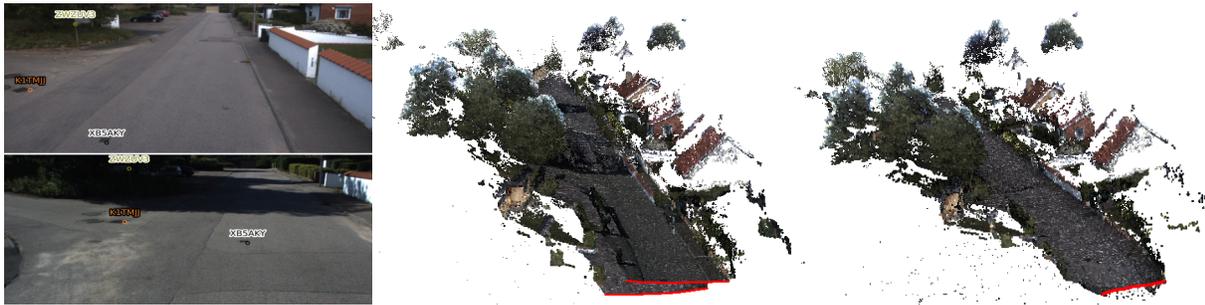


Figure 2. Left: Annotated correspondences; Middle: Initial dense reconstruction of the scene; Right: Dense reconstruction with refinement after registration. The edge of the road surface is highlighted to indicate misalignment.

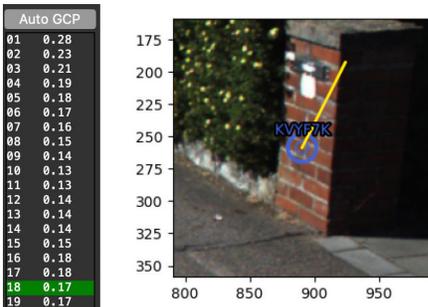


Figure 3. Our annotation UI integrates annotation metrics to enable efficient annotation and QA: (1) Reprojection error of annotated CPs (yellow line) (2) Positional std. deviation derived from bundle adjustment, highlighted in green (17 cm for this example).

over- (bright sun) or underexposure (night) or lack of texture. Moreover, they are expensive to compute and introduce some drawbacks, *e.g.*, registration failures due to missing or changed geometry (foliage, snow, *etc.*).

3.4. Quality Control

To verify the reference poses, we use the following criteria and workflow. If SfM succeeds and there are no notice-

able artefacts and errors, then it typically produces highly accurate estimates for the camera poses and the 3D structure of the scene. Thus, we directly trust the reference poses we obtained for the easy datasets.

For the medium and hard datasets, we rely on manually-annotated matches, which are susceptible to human error and might be distributed in a sub-optimal way in the scene, under-constraining the alignment. Previous work on building benchmark datasets that also uses manual annotations [58, 70] often relied on visual inspection and did not rigorously measure the uncertainty of their generated poses. In the case of the Aachen Day-Night dataset [58, 60], [86] recently showed that inaccurate poses passed visual inspection. We thus made a conscious effort to quantitatively measure the accuracy of our poses.

We use two metrics to verify that the annotations are correct and sufficient. After a few control points have been annotated, an alignment is generated by estimating a similarity transform followed by bundle adjustment (as described above). To detect potential issues, we begin by estimating the **reprojection error** of every annotation. This metric

identifies wrongly annotated points such as wrong correspondences between different objects with similar appearance or simply misclicks during annotation. All annotations with an error of more than one pixel (at VGA resolution) are flagged as incorrect and must be refined.

After all the reprojection errors are sufficiently small, we compute the **camera position covariance** for each image using bundle adjustment [24]. To fix the gauge ambiguity, we fix the poses of one of the sequences, run the bundle adjustment problem, and compute the covariance of the poses of the other sequence. We repeat this while fixing the poses of the other sequence. From the covariances, we compute the standard deviation of the camera positions³ in meters and seek for small values when annotating. This metric is used to understand if the calculated position of an image is under-constrained: all annotations can be correct with very small reprojection errors, but the poses might still be under-constrained, *e.g.*, if the annotations were only performed on far-away points, resulting in large standard deviations.

Annotation Tools. These two metrics are integrated with the annotation UI that we have developed and will release as part of this work (*c.f.* Fig. 3). Annotators can run bundle adjustment and get feedback directly in the UI, visualize wrongly annotated points and focus their efforts on annotating those frames with higher positional uncertainty, simplifying and accelerating the annotation process. Using our tool, the manual annotation (and QA) of a sequence pair reconstruction takes, on average, only 30 minutes.

After all positional std. deviations are confirmed to be under 30 cm, we perform a final check by visually inspecting the combined point cloud for the aligned models (Fig. 2). Tab. 2 shows some statistics about the annotation process, including the number of annotated points, the median reprojection error and median positional std. deviation for each dataset. For the whole set of Medium and Hard datasets, the median annotation reprojection error is 0.28 px and the median positional std. deviation is 5.4 cm.

4. Baselines

In order to show that our benchmark introduces new challenges for localization algorithms, we use a set of state-of-the-art localization methods as baselines. We focus on methods that have been shown to work well under changing conditions (based on their performance on the **benchmark** from [58]) and provide source code and trained models:

HLoc [51,52] uses learned SuperPoint [18] features and SuperGlue [52] to establish 2D-3D matches with a SfM model, which are then used for camera pose estimation.

D2-Net [22] uses a single CNN for feature detection and description. Pose estimation is implemented in

³We focus on the positional covariance and not the full pose covariance as (1) it is easier to understand and (2) the positions are less certain than the orientation estimates (which can be constrained by points at infinity) [24].

COLMAP [62]. In contrast to HLoc, which runs (close to) real-time, D2-Net requires multiple seconds per test image.

Rectified SIFT [71] uses a depth estimation network [45] to detect planar regions in the image. Warping them to remove perspective foreshortening leads to features that can be matched under strong viewpoint changes.

S2DHM [30] uses an asymmetric matching approach: reference images are represented by sparse features corresponding to the 3D points in the SfM model, while the test image is represented by densely extracted descriptors.

The above-mentioned baselines are based on matching local features to establish 2D-3D matches between the test image and an SfM model of the scene. In contrast, **PixLoc** [53] does not rely on feature matching but refines an initial pose estimate by minimizing a feature-metric cost function.

We do not evaluate camera pose and scene coordinate regression methods. Pose regressors have been shown to be significantly less accurate than other localization approaches [61], even on scenes without changing conditions. We thus see no reason why they should perform well on our more challenging benchmark. Current scene coordinate regressors seem to struggle with strong condition changes between the training and test sets: on the Aachen Day-Night dataset [58,60], where all training images are taken at day, ESAC [8] performs significantly worse for nighttime test images compared to daytime queries and is less accurate than HLoc by a wide margin. CrowdDriven only provides reference images taken under a single condition for each scene. Given the strong condition changes between training and test images, current scene coordinate regressors are unlikely to perform well on our dataset.

In our experiments, for a given test image, we consider only reference images from the same scene, *e.g.*, Boston1, and not from different scenes. Since our scenes are rather small (*c.f.* Tab. 2), using place recognition / image retrieval is not necessary. Instead, we exhaustively match the test image against all reference images in the scene.

The supp. material shows results for **sequence-based localization** methods that simultaneously estimate the poses of all images in a sequence [12,41,42,69,77]. While this improves performance, it is still not sufficient to obtain reasonable performance on our more challenging scenes.

5. Experimental Evaluation

This section shows how our dataset introduces new challenging scenarios that most baselines cannot cope with. After introducing our evaluation measures, we analyze the impact of different types of changes on the performance.

Evaluation Measures. We follow common evaluation protocols [25,39,58,70] and report median position (in meters) and orientation errors (in degrees), as well as the percentage of test images with poses that differ within certain error bounds from their reference poses. The position error is

name	D2-Net			S2DHM			HLoc			Rectified SIFT			PixLoc			Changes
	pos. err.	rot. err.	% of localized 0.5/0.5/0/100 (m) 25/10/20 (°)	pos. err.	rot. err.	% of localized 0.5/0.5/0/100 (m) 25/10/20 (°)	pos. err.	rot. err.	% of localized 0.5/0.5/0/100 (m) 25/10/20 (°)	pos. err.	rot. err.	% of localized 0.5/0.5/0/100 (m) 25/10/20 (°)	pos. err.	rot. err.	% of localized 0.5/0.5/0/100 (m) 25/10/20 (°)	
Angers1	28.02	177.43	F	97.81	171.26	F	46.39	161.65	F	191.62	148.27	F	21.03	175.11	F	🌧️
Angers2	35.51	165.56	F	174.61	153.78	F	68.34	122.82	0/0/0/6.52	437.63	132.38	F	45.26	173.18	F	🌧️
Bayern	0.03	0.06	96.15/ 96.15/ 96.15/ 96.15	0.09	0.30	80.77/ 80.77/ 80.77/ 80.77	0.02	0.07	80.77/ 80.77/ 80.77/ 80.77	0.04	0.11	80.77/ 84.62/ 84.62/ 84.62	0.09	0.12	61.54/ 61.54/ 65.38/ 73.08	💡 🌧️
Besaçon2	81.45	160.22	F	-	-	-	59.02	152.43	F	128.70	121.79	F	34.30	169.03	F	🌧️
Besaçon3	48.16	162.30	F	258.71	148.52	F	71.61	162.18	F	107.51	148.63	F	36.86	168.47	F	🌧️
Besaçon4	117.25	151.57	F	-	-	-	108.59	141.24	F	287.57	134.02	F	69.80	172.84	F	🌧️
Boston1	28.78	4.99	0/0/4.26/ 23.40	239.57	140.05	F	31.75	8.47	0/0/0/ 10.64	125.13	129.79	F	27.09	15.55	0/0/0/ 2.13	🌧️
Boston2	6.46	0.96	0/0/24.49/ 97.96	496.16	86.71	0/0/8.16/ 16.33	4.68	0.82	0/0/63.27/ 95.92	87.21	150.42	F	13.26	7.22	0/0/6.12/ 38.78	🌧️
Boston3	6.66	4.20	0/0/29.03/ 51.61	196.98	114.24	F	27.57	32.83	0/0/12.90/ 19.35	111.70	155.53	F	20.72	16.12	0/0/0/ 19.35	🌧️
Boston4	12.94	2.51	0/0/20.83/ 41.67	-	-	-	15.57	6.19	0/0/26.47/ 38.24	-	-	-	18.90	5.81	0/0/11.76/ 26.47	🌧️
Boston5	18.08	2.26	0/0/11.76/ 17.65	97.70	74.40	0/0/0/ 5.88	16.59	4.35	0/0/26.47/ 26.47	91.74	157.72	F	13.36	12.36	0/0/0/ 26.47	🌧️
Brittany	14.74	147.24	F	177.38	143.03	F	36.44	137.46	0/3.23/ 3.23/ 3.23	305.68	121.02	F	14.96	162.08	F	🌧️
Brouages	31.84	153.74	F	-	-	-	22.24	153.54	F	57.10	97.23	F	14.64	177.38	F	🌧️
Burgundy2	4.41	3.72	0/4.00/ 60/ 76.00	57.78	34.59	F	7.32	5.36	0/4.00/ 40/ 58.00	554.19	166.83	F	20.97	14.82	0/0/0/ 26.00	🌧️
Cambridge	0.50	0.87	51.52/ 90.91/ 93.94/ 96.97	94.58	82.57	9.09/ 12.12/ 12.12/ 12.12	0.37	0.43	69.70/ 100/ 100/ 100	58.58	135.56	3.03/ 6.06/ 18.18/ 21.21	30.06	15.24	0/0/3.03/ 6.06	💡 🌧️
Clermont-Ferrand	0.22	0.24	100/ 100/ 100/ 100	0.25	0.47	100/ 100/ 100/ 100	0.15	0.27	100/ 100/ 100/ 100	0.21	0.33	80/ 93.33/ 93.33/ 93.33	0.19	0.20	93.33/ 93.33/ 93.33/ 93.33	💡 🌧️
Curitiba	0.03	0.07	100/ 100/ 100/ 100	0.21	0.36	84.21/ 89.47/ 100/ 100	0.04	0.06	100/ 100/ 100/ 100	0.06	0.08	89.47/ 89.47/ 89.47/ 89.47	0.06	0.09	84.21/ 84.21/ 84.21/ 84.21	💡 🌧️
Ile-de-France	58.17	159.89	F	-	-	-	121.89	120.76	F	325.41	159.15	F	24.99	175.11	F	🌧️
Le-Mans	53.53	160.64	F	63.90	165.43	F	52.82	163.53	F	249.61	138.46	F	40.68	176.24	F	🌧️
Leuven	10.85	164.76	F	48.83	154.03	F	12.54	173.48	F	69.29	132.33	F	8.42	140.80	F	🌧️
Massachusetts1	0.08	0.06	100/ 100/ 100/ 100	0.59	0.36	40/ 72.00/ 96.00/ 96.00	0.12	0.07	100/ 100/ 100/ 100	0.16	0.06	96.00/ 96.00/ 100/ 100	0.20	0.10	76.00/ 84.00/ 84.00/ 84.00	💡 🌧️
Massachusetts2	5.80	0.24	0/0/0/ 100	-	-	-	18.34	9.47	0/0/8.33	111.63	168.17	F	6.70	10.76	0/0/0/ 66.67	🌧️
Massachusetts3	23.95	25.32	0/10.53/ 36.84/ 42.11	498.52	123.61	F	3.91	5.85	0/25.00/ 52.77/ 59.09	680.80	104.46	F	18.02	10.24	0/4.55/ 22.73	🌧️
Massachusetts4	1.57	1.01	0/0/97.44/ 100	69.55	24.08	0/8.11/ 32.43/ 32.43	2.23	1.14	0/0/100/ 100	40.25	115.17	0/0/0/ 4.00	10.38	5.70	0/0/30.77/ 46.15	🌧️
Melbourne	0.07	0.07	100/ 100/ 100/ 100	0.21	0.22	83.33/ 100/ 100/ 100	0.09	0.16	100/ 100/ 100/ 100	0.16	0.16	100/ 100/ 100/ 100	14.07	0.93	16.67/ 16.67/ 16.67/ 41.67	💡 🌧️
Muehlhausen	0.03	0.07	100/ 100/ 100/ 100	0.32	0.53	80/ 100/ 100/ 100	0.04	0.07	100/ 100/ 100/ 100	0.04	0.11	100/ 100/ 100/ 100	0.04	0.10	100/ 100/ 100/ 100	💡 🌧️
Nouvelle-Aquitaine1	40.23	172.90	F	328.30	150.36	F	34.02	160.82	F	418.74	137.36	F	44.17	170.42	F	🌧️
Nouvelle-Aquitaine2	81.48	126.82	F	35.89	161.41	F	67.65	147.96	F	90.10	150.32	F	28.51	169.94	F	🌧️
Orleans1	17.42	178.86	F	226.48	149.29	F	33.42	175.46	F	87.40	157.44	0/0/3.03/ 3.03	25.87	178.74	F	🌧️
Orleans2	175.99	127.58	F	-	-	-	32.95	165.35	F	359.37	154.26	F	15.30	177.55	F	🌧️
Pays de la Loire	25.32	159.23	F	52.74	156.08	F	34.11	166.72	F	131.31	135.25	0/0/0/ 4.76	17.56	175.95	F	🌧️
Poing	0.05	0.07	100/ 100/ 100/ 100	0.45	0.61	60/ 85.00/ 100/ 100	0.06	0.07	100/ 100/ 100/ 100	183.46	85.13	20/ 20/ 20/ 20.00	0.08	0.04	85.00/ 85.00/ 85.00/ 85.00	💡 🌧️
Portland	0.13	0.16	100/ 100/ 100/ 100	0.40	0.46	66.67/ 95.24/ 100/ 100	0.11	0.14	100/ 100/ 100/ 100	0.16	0.12	95.24/ 100/ 100/ 100	0.13	0.16	85.71/ 85.71/ 85.71/ 85.71	💡 🌧️
Savannah	0.08	0.05	100/ 100/ 100/ 100	0.25	0.28	83.33/ 94.44/ 100/ 100	0.08	0.05	100/ 100/ 100/ 100	0.07	0.05	100/ 100/ 100/ 100	0.09	0.08	94.44/ 94.44/ 94.44/ 94.44	💡 🌧️
Skåne	5.14	2.97	0/0/50/ 85.00	392.36	120.41	0/0/0/ 5.00	4.26	3.70	0/5.00/ 55.00/ 90.00	609.69	154.13	F	35.93	25.73	F	🌧️
Subcarpathia	0.54	0.46	47.06/ 70.59/ 88.24/ 94.12	6.39	4.09	0/0/41.18/ 58.82	0.35	0.26	70.59/ 76.47/ 100/ 100	116.17	80.70	11.76/ 11.76/ 29.41/ 35.29	66.44	9.84	5.88/ 5.88/ 11.76/ 11.76	🌧️
Sydney	0.18	0.11	100/ 100/ 100/ 100	1.39	0.82	35.71/ 35.71/ 85.71/ 92.86	0.18	0.15	85.71/ 100/ 100/ 100	2.81	1.86	0/25.00/ 75.00/ 75.00	0.18	0.20	64.29/ 71.43/ 71.43/ 71.43	🌧️
Thuringia	0.57	0.26	45.45/ 90.91/ 100/ 100	0.95	0.60	18.18/ 54.55/ 100/ 100	0.37	0.25	81.82/ 100/ 100/ 100	431.41	121.14	F	7.03	6.51	0/0/27.27/ 63.64	🌧️
Tsuru	0.01	0.04	100/ 100/ 100/ 100	0.06	0.30	100/ 100/ 100/ 100	0.03	0.03	100/ 100/ 100/ 100	0.02	0.04	100/ 100/ 100/ 100	0.03	0.04	100/ 100/ 100/ 100	💡 🌧️
Washington	1.00	0.42	0/50/ 100/ 100	3.96	0.66	0/0/ 100/ 100	1.07	1.47	10/ 30/ 100/ 100	0.90	0.68	40/ 50/ 90/ 100	2.46	1.20	0/0/ 70.00/ 70.00	💡 🌧️

Table 3. Localization performance on our CrowdDriven benchmark. We report the median position (in meters) and orientation (in degrees) errors, and the percentage of test images localized within certain error bounds on the position and orientation errors. Easy, medium, and hard datasets are color-coded in light, standard, and dark gray, respectively. We also provide information about the type of change between the training and test sequences: illumination: 💡, overcast: ☁, foliage: 🌳, snow: ❄, seasonal: 🍂, day-night: 🌑, small viewpoint: 🏠, rain: 🌧️, strong viewpoint: 🏠, man-made changes: 🏠. 'F' stands for failure to localize any image within the coarsest precision regime.



Figure 4. Inlier plots for selected scenes. None of the evaluated methods is able to robustly localize the *medium* and *hard* datasets. This is made apparent by the fraction of localized images under the *medium-precision* threshold for the **best** method, shown in below each column.

measured as the Euclidean distance between the reference and estimated position. To measure the orientation error, we compute the minimum rotation angle α that aligns the estimated rotation matrix R_{est} with the reference rotation R_{ref} as $2\cos(|\alpha|) = \text{trace}(R_{\text{ref}}^T R_{\text{est}}) - 1$ [33]. Inspired by [58],

we use three error bounds: high precision (images localized within 0.5m and 2° of their reference poses), medium precision (1m, 5°), and coarse precision (5m, 10°). In addition, we introduce a very coarse precision regime (10m, 20°).

We also report the changes in conditions between the test

and training images: (slight) illumination changes (il.), *e.g.*, between a cloudy and a sunny day, foliage / no foliage on the vegetation (fo.), snow / no snow on the ground (sn.), other seasonal changes (se.) (summer/autumn), day-night changes (ng.), rain / no rain (rn.), small (sm. v.) and strong viewpoint (st. v.), and man-made (*e.g.*, appearance / disappearance of cars) changes. Tab. 3 summarizes the results obtained by evaluating the baselines. In the following, we focus our analysis on three types of changes: **slight illumination**, **day-night**, and **strong viewpoint** changes, which are the dominant types of changes.

Slight illumination changes. As shown in Tab. 3, datasets such as Muehlhausen, Tsuru, Poing, Bayern, Savannah, Curitiba, Melbourne, Sydney, and Clermont-Ferrand are categorized as being affected solely by illumination changes seen during the day. Almost all the test images are localized with high precision by most of the methods. In this scenario, SIFT is already robust enough to slight changes in illumination. However, rectified SIFT struggles to deal with the heavy vegetation present in Poing, Sydney, and Subcarpathia. We observe that for some datasets, PixLoc performs significantly worse than the other baselines. We attribute this to different viewing condition between the training set of PixLoc and our dataset, *e.g.*, PixLoc was not trained on images showing as much snow as the Subcarpathia dataset, and the challenges of identifying the best pose from all available pose estimates (*c.f.* supp. mat.) and of pose refinement in complex scenes such as Melbourne and Sydney. Still, we conclude that slight illumination changes are not challenging for current algorithms.

Day-night changes. The significant appearance changes that occur between day and night have considerable effects on the performance of most of the methods, leading to a noticeable increase in the errors. D2-Net’s robustness to illumination changes makes the model perform well in some of the datasets such as Massachusetts4 and Boston4, allowing localization with medium precision. On the other hand, the effect of artificial lights during the night could cover the previously-seen features or considerably change them, making one part of the image sparser, and the pose estimation less accurate. In Massachusetts2, Massachusetts3, Boston2, and Boston5, the errors are larger than those of consumer-grade GPS. HLoc and S2DHM, as seen in Tab. 3, produce large errors in all of the night cases except for Cambridge, with the median errors of these methods starting at more than 30 meters. We can conclude that for night-time test images, due to the considerable changes in feature appearances, approaches such as S2DHM do not seem to improve the results as long as the extracted Hypercolumn descriptors themselves are not robust enough to the day-night changes (*c.f.* Fig. 4). Rectified SIFT inherits the limitations of SIFT features and also struggles significantly on these scenes.

Strong viewpoint changes. By far the most challenging

condition present in our dataset is strong viewpoint change, as seen in Tab. 3. The rotation errors obtained for all the methods are over 160° , indicating complete failures. Strong viewpoint changes significantly alter the appearance of objects, resulting in substantially different feature descriptors depending on the viewing angle. In theory, removing perspective distortion via rectification should help to better handle larger viewpoint changes [71]. Yet, the rectified SIFT baseline still mostly fails under strong viewpoint changes. This can be attributed to the limited visual overlap as well as inaccurate depth predictions on the distant overlapping image regions (*c.f.* Fig. 4). Unsurprisingly, PixLoc always fails under strong viewpoint changes: PixLoc refines an initial pose estimate, which is obtained from the poses of the database images. In the case of opposite viewpoints, these initial estimates are simply too different from the test poses for PixLoc to converge to a reasonable pose estimate.

Difficulty levels revisited. Sec. 3 provided a preliminary classification based on our understanding of what is challenging for current localization algorithms. Based on the measured performance of our baselines, we revisit this classification: a dataset should be considered as easy if (nearly) all baselines perform well, as medium if some baselines work well (at least for the coarsest threshold), and as hard if all methods fail. We color-coded the datasets in Tab. 3 accordingly. Comparing the color-codings in Tabs. 2 and 3 shows that the classification based on performance is consistent with our preliminary categorization for most scenes.

6. Conclusion

We have introduced the CrowdDriven dataset for visual localization, with a focus on exposing the failure cases of state-of-the-art pipelines. The dataset covers three difficulty tiers: (1) slight illumination changes, (2) severe illumination changes, (3) severe viewpoint changes.

We have evaluated the performance of the best visual localization methods available on our dataset and analyzed the failure modes for the different methods. We conclude that there is a large performance gap to be covered by future work and that can be evaluated with our dataset. Finally, we release the annotation tool we developed to create this dataset, enabling easy extension to more scenarios.

As shown concurrently to our work [5], using a feature-based approach (SfM) to generate reference poses can create a bias towards feature-based methods. We consider the problem of eliminating such a bias from the evaluation as an important direction for future work.

Acknowledgements. This work has been funded by the Chalmers AI Research Centre (VisLocLearn), the Swedish Foundation for Strategic Research (Semantic Mapping and Visual Navigation for Smart Robots), the EU Horizon 2020 project RICAIP (grant agreement No 857306), and the European Regional Development Fund under IMPACT No. CZ.02.1.01/0.0/0.0/15.003/0000468.

References

- [1] OpenSfM. <https://github.com/mapillary/OpenSfM>. 4
- [2] Waymo open dataset: An autonomous driving dataset. <https://waymo.com/open/>, 2019. 2
- [3] Relja Arandjelović, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *CVPR*, 2016. 1
- [4] Hernan Badino, Daniel Huber, and Takeo Kanade. Visual topometric localization. In *Intelligent Vehicles Symposium (IV)*, 2011. 2, 3
- [5] Eric Brachmann, Martin Humenberger, Carsten Rother, and Torsten Sattler. On the Limits of Pseudo Ground Truth in Visual Camera Re-Localization. In *International Conference on Computer Vision (ICCV)*, 2021. 2, 8
- [6] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. DSAC - Differentiable RANSAC for Camera Localization. In *CVPR*, 2017. 2, 4
- [7] Eric Brachmann, Frank Michel, Alexander Krull, Michael Ying Yang, Stefan Gumhold, and Carsten Rother. Uncertainty-driven 6d pose estimation of objects and scenes from a single rgb image. In *CVPR*, 2016. 2, 4
- [8] Eric Brachmann and Carsten Rother. Expert Sample Consensus Applied to Camera Re-Localization. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 4, 6
- [9] Eric Brachmann and Carsten Rother. Visual Camera Re-Localization from RGB and RGB-D Images Using DSAC. *arXiv preprint arXiv:2002.12324*, 2020. 2, 4
- [10] Samarth Brahmabhatt, Jinwei Gu, Kihwan Kim, James Hays, and Jan Kautz. Geometry-aware learning of maps for camera localization. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 4
- [11] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. *CoRR*, abs/1903.11027, 2019. 2
- [12] Federico Camposeco, Torsten Sattler, and Marc Pollefeys. Minimal Solvers for Generalized Pose and Scale Estimation from Two Rays and One Point. In *ECCV*, 2016. 6
- [13] Nicholas Carlevaris-Bianco, Arash K Ushani, and Ryan M Eustice. University of michigan north campus long-term vision and lidar dataset. *The International Journal of Robotics Research*, 35(9):1023–1035, 2016. 3
- [14] Tommaso Cavallari, Luca Bertinetto, Jishnu Mukhoti, Philip Torr, and Stuart Golodetz. Let’s take this online: Adapting scene coordinate regression network predictions for online rgb-d camera relocalisation. In *3DV*, 2019. 2, 4
- [15] Tommaso Cavallari, Stuart Golodetz, Nicholas A. Lord, Julien Valentin, Luigi Di Stefano, and Philip H. S. Torr. On-The-Fly Adaptation of Regression Forests for Online Camera Relocalisation. In *CVPR*, 2017. 2, 4
- [16] David M Chen, Georges Baatz, Kevin Köser, Sam S Tsai, Ramakrishna Vedantham, Timo Pylvänäinen, Kimmo Roimela, Xin Chen, Jeff Bach, Marc Pollefeys, et al. City-scale landmark identification on mobile devices. In *CVPR 2011*, pages 737–744. IEEE, 2011. 2, 3
- [17] Zetao Chen, Adam Jacobson, Niko Sünderhauf, Ben Upcroft, Lingqiao Liu, Chunhua Shen, Ian D. Reid, and Michael Milford. Deep Learning Features at Scale for Visual Place Recognition. *ICRA*, 2017. 1
- [18] D. DeTone, T. Malisiewicz, and A. Rabinovich. SuperPoint: Self-Supervised Interest Point Detection and Description. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018. 6
- [19] Anh-Dzung Doan, Yasir Latif, Tat-Jun Chin, Yu Liu, Shin-Fang Ch’ng, Thanh-Toan Do, and Ian Reid. Visual localization under appearance change: filtering approaches. *Neural Computing and Applications*, pages 1–14, 2020. 2
- [20] Anh-Dzung Doan, Yasir Latif, Tat-Jun Chin, Yu Liu, Thanh-Toan Do, and Ian Reid. Scalable place recognition under appearance change for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9319–9328, 2019. 2
- [21] Siyan Dong, Qingnan Fan, He Wang, Ji Shi, Li Yi, Thomas Funkhouser, Baoquan Chen, and Leonidas J. Guibas. 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)*, 2021. 2
- [22] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-net: A trainable cnn for joint detection and description of local features. *arXiv preprint arXiv:1905.03561*, 2019. 2, 6
- [23] M. Fischler and R. Bolles. Random Sampling Consensus: A Paradigm for Model Fitting with Application to Image Analysis and Automated Cartography. *Communications of the ACM*, 24:381–395, 1981. 2
- [24] Wolfgang Förstner and Bernhard P. Wrobel. *Photogrammetric Computer Vision*. Springer International Publishing, 2016. 6
- [25] Jan-Michael Frahm, Pierre Fite-Georgel, David Gallup, Tim Johnson, Rahul RaguramC, hangchang Wu, Yi-Hung Jen, Enrique Dunn, Brian Clipp, Svetlana Lazebnik, and Marc Pollefeys. Building Rome on a Cloudless Day. In *European Conference on Computer Vision*, 2010. 2, 4, 6
- [26] Sourav Garg, Tobias Fischer, and Michael Milford. Where is your place, visual place recognition? *arXiv preprint arXiv:2103.06443*, 2021. 2
- [27] Sourav Garg, Niko Sünderhauf, and Michael Milford. Semantic-geometric visual place recognition: a new perspective for reconciling opposing views. *The International Journal of Robotics Research*, page 0278364919839761, 2019. 2
- [28] Sourav Garg, Niko Sünderhauf, Feras Dayoub, Douglas Morrison, Akansel Cosgun, Gustavo Carneiro, Qi Wu, Tat-Jun Chin, Ian Reid, Stephen Gould, et al. Semantics for robotic mapping, perception and interaction: A survey. *arXiv preprint arXiv:2101.00443*, 2021. 2
- [29] Yixiao Ge, Haibo Wang, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-supervising fine-grained region similarities

- for large-scale image localization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1
- [30] Hugo Germain, Guillaume Bourmaud, and Vincent Lepetit. Sparse-to-dense hypercolumn matching for long-term visual localization. In *2019 International Conference on 3D Vision (3DV)*, pages 513–523. IEEE, 2019. 2, 6
- [31] Ben Glocker, Shahram Izadi, Jamie Shotton, and Antonio Criminisi. Real-time rgb-d camera relocalization. In *International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, October 2013. 3
- [32] Bert M. Haralick, Chung-Nan Lee, Karsten Ottenberg, and Michael Nölle. Review and analysis of solutions of the three point perspective pose estimation problem. *IJCV*, 13(3):331–356, 1994. 2
- [33] Richard Hartley, Jochen Trumpp, Yuchao Dai, and Hongdong Li. Rotation averaging. *International journal of computer vision*, 103(3):267–305, 2013. 7
- [34] Caner Hazirbas, Sebastian G. Soyler, Maximilian C. Staab, Laura Leal-Taixé, and Daniel Cremers. Deep depth from focus. In *Asian Conference on Computer Vision (ACCV)*, December 2018. 3
- [35] Jared Heinly, Johannes L. Schönberger, Enrique Dunn, and Jan-Michael Frahm. Reconstructing the world* in six days. In *CVPR*, 2015. 1
- [36] Xinyu Huang, Xinjing Cheng, Qichuan Geng, Binbin Cao, Dingfu Zhou, Peng Wang, Yuanqing Lin, and Ruigang Yang. The apolloscape dataset for autonomous driving. *arXiv:1803.06184*, 2018. 2
- [37] Arnold Irschara, Christopher Zach, Jan-Michael Frahm, and Horst Bischof. From structure-from-motion point clouds to fast location recognition. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2599–2606. IEEE, 2009. 2, 3
- [38] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *CVPR*, 2017. 2, 4
- [39] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A Convolutional Network for Real-Time 6-DOF Camera Relocalization. In *ICCV*, 2015. 2, 3, 4, 6
- [40] R. Kesten, M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinsky, W. Jiang, and V. Shet. Lyft level 5 av dataset. <https://level5.lyft.com/dataset/>, 2019. 2
- [41] Zuzana Kukelova, Jan Heller, and Andrew Fitzgibbon. Efficient Intersection of Three Quadrics and Applications in Computer Vision. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 6
- [42] Gim Hee Lee, Bo Li, Marc Pollefeys, and Friedrich Fraundorfer. Minimal solutions for the multi-camera pose estimation problem. *IJRR*, 34(7):837–848, 2015. 6
- [43] Yunpeng Li, Noah Snavely, and Daniel P Huttenlocher. Location recognition using prioritized feature matching. In *European conference on computer vision*, pages 791–804. Springer, 2010. 2, 3
- [44] Yunpeng Li, Noag Snavely, Dan P. Huttenlocher, and Pascal Fua. Worldwide Pose Estimation Using 3D Point Clouds. In *ECCV*, 2012. 2, 3, 4
- [45] Manuel Lopez-Antequera, Pau Gargallo, Markus Hofinger, Samuel Rota Bulò, Yubin Kuang, and Peter Kotschieder. Mapillary planet-scale depth dataset. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 6
- [46] David G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *IJCV*, 60(2), 2004. 4
- [47] Stephanie Lowry, Niko Sünderhauf, Paul Newman, John J Leonard, David Cox, Peter Corke, and Michael J Milford. Visual place recognition: A survey. *IEEE Transactions on Robotics*, 32(1):1–19, 2015. 2
- [48] Will Maddern, Geoff Pascoe, Chris Linegar, and Paul Newman. 1 Year, 1000km: The Oxford RobotCar Dataset. *IJRR*, 36(1):3–15, 2017. 2, 3
- [49] Michael J Milford and Gordon F Wyeth. SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights. In *ICRA*, 2012. 2
- [50] Tayyab Naseer, Wolfram Burgard, and Cyrill Stachniss. Robust visual localization across seasons. *IEEE Transactions on Robotics*, 34(2):289–302, 2018. 2
- [51] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12716–12725, 2019. 2, 4, 6
- [52] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *CVPR*, 2020. 2, 4, 6
- [53] Paul-Edouard Sarlin, Ajaykumar Unagar, Mans Larsson, Hugo Germain, Carl Toft, Viktor Larsson, Marc Pollefeys, Vincent Lepetit, Lars Hammarstrand, Fredrik Kahl, and Torsten Sattler. Back to the feature: Learning robust camera localization from pixels to pose. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 4, 6
- [54] Torsten Sattler, Michal Havlena, Konrad Schindler, and Marc Pollefeys. Large-Scale Location Recognition and the Geometric Burstiness Problem. In *CVPR*, 2016. 1
- [55] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Fast Image-Based Localization using Direct 2D-to-3D Matching. In *ICCV*, 2011. 2
- [56] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Improving Image-Based Localization by Active Correspondence Search. In *ECCV*, 2012. 2
- [57] T. Sattler, B. Leibe, and L. Kobbelt. Efficient & Effective Prioritized Matching for Large-Scale Image-Based Localization. *PAMI*, 39(9):1744–1756, 2017. 2
- [58] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, and Tomas Pajdla. Benchmarking 6DOF Urban Visual Localization in Changing Conditions. In *CVPR*, 2017. 1, 2, 3, 4, 5, 6, 7
- [59] Torsten Sattler, Akihiko Torii, Josef Sivic, Marc Pollefeys, Hajime Taira, Masatoshi Okutomi, and Tomas Pajdla. Are Large-Scale 3D Models Really Necessary for Accurate Visual Localization? In *CVPR*, 2017. 1

- [60] Torsten Sattler, Tobias Weyand, Bastian Leibe, and Leif Kobbelt. Image Retrieval for Image-Based Localization Revisited. In *British Machine Vision Conference*, 2012. 2, 3, 5, 6
- [61] Torsten Sattler, Qunjie Zhou, Marc Pollefeys, and Laura Leal-Taixe. Understanding the limitations of cnn-based absolute camera pose regression. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 2, 6
- [62] Johannes L. Schönberger and Jan-Michael Frahm. Structure-From-Motion Revisited. In *CVPR*, June 2016. 1, 4, 6
- [63] Johannes L Schönberger, Marc Pollefeys, Andreas Geiger, and Torsten Sattler. Semantic visual localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6896–6906, 2018. 2, 4
- [64] Stephen Se, David Lowe, and Jim Little. Mobile robot localization and mapping with uncertainty using scale-invariant visual landmarks. *The international Journal of robotics Research*, 21(8):735–758, 2002. 2
- [65] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene Coordinate Regression Forests for Camera Relocalization in RGB-D Images. In *CVPR*, 2013. 2, 4
- [66] Noah Snavely, Steven M. Seitz, and Richard Szeliski. Modeling the world from internet photo collections. *IJCV*, 80(2):189–210, 2008. 1
- [67] Niko Sünderhauf, Peer Neubert, and Peter Protzel. Are we there yet? challenging seqslam on a 3000 km journey across all four seasons. 2, 3
- [68] Linus Svärm, Olof Enqvist, Fredrik Kahl, and Magnus Oskarsson. City-Scale Localization for Cameras with Known Vertical Direction. *PAMI*, 39(7):1455–1461, 2017. 2
- [69] Chris Sweeney, Victor Fragoso, Tobias Höllerer, and Matthew Turk. gDLS: A Scalable Solution to the Generalized Pose and Scale Problem. In *ECCV*. Springer, 2014. 6
- [70] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. InLoc: Indoor visual localization with dense matching and view synthesis. In *CVPR*, 2018. 2, 3, 4, 5, 6
- [71] Carl Toft, Daniyar Turmukhambetov, Torsten Sattler, Fredrik Kahl, and Gabriel J Brostow. Single-image depth prediction makes feature matching easier. In *European Conference on Computer Vision*, pages 473–492. Springer, 2020. 6, 8
- [72] Akihiko Torii, Relja Arandjelović, Josef Sivic, Masatoshi Okutomi, and Tomas Pajdla. 24/7 Place Recognition by View Synthesis. In *CVPR*, 2015. 1, 2, 3
- [73] Akihiko Torii, Josef Sivic, and Tomas Pajdla. Visual localization by linear combination of image descriptors. In *Proceedings of the 2nd IEEE Workshop on Mobile Vision, with ICCV*, 2011. 1
- [74] Akihiko Torii, Josef Sivic, Tomas Pajdla, and Masatoshi Okutomi. Visual place recognition with repetitive structures. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 883–890, 2013. 2, 3
- [75] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment—a modern synthesis. In *International workshop on vision algorithms*, pages 298–372. Springer, 1999. 4
- [76] Julien Valentin, Angela Dai, Matthias Nießner, Pushmeet Kohli, Philip Torr, Shahram Izadi, and Cem Keskin. Learning to Navigate the Energy Landscape. In *3DV*, 2016. 2
- [77] Jonathan Ventura, Clemens Arth, Gerhard Reitmayr, and Dieter Schmalstieg. A Minimal Solution to the Generalized Pose-and-Scale Problem. In *CVPR*, 2014. 6
- [78] Olga Vysotska and Cyrill Stachniss. Lazy data association for image sequences matching under substantial appearance changes. *IEEE Robotics and Automation Letters*, 1(1):213–220, 2015. 2
- [79] Florian Walch, Caner Hazirbas, Laura Leal-Taixé, Torsten Sattler, Sebastian Hilsenbeck, and Daniel Cremers. Image-Based Localization Using LSTMs for Structured Feature Correlation. In *ICCV*, 2017. 2, 4
- [80] Johanna Wald, Torsten Sattler, Stuart Golodetz, Tommaso Cavallari, and Federico Tombari. Beyond controlled environments: 3D camera re-localization in changing indoor scenes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2, 3, 4
- [81] Frederik Warburg, Soren Hauberg, Manuel López-Antequera, Pau Gargallo, Yubin Kuang, and Javier Civera. Mapillary street-level sequences: A dataset for lifelong place recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2626–2635, 2020. 2
- [82] Erik Wijmans and Yasutaka Furukawa. Exploiting 2d floorplan for building-scale panorama rgbd alignment. In *Computer Vision and Pattern Recognition, CVPR*, 2017. 3
- [83] Amir R. Zamir and Mubarak Shah. Accurate Image Localization Based on Google Maps Street View. In *ECCV*, 2010. 1
- [84] Bernhard Zeisl, Torsten Sattler, and Marc Pollefeys. Camera pose voting for large-scale image-based localization. In *ICCV*, 2015. 2
- [85] Wei Zhang and Jana Kosecka. Image based Localization in Urban Environments. In *3DPVT*, 2006. 1
- [86] Zichao Zhang, Torsten Sattler, and Davide Scaramuzza. Reference Pose Generation for Long-term Visual Localization via Learned Features and View Synthesis. *International Journal of Computer Vision*, pages 1–1, 2020. 5