

Self-Supervised Localisation between Range Sensors and Overhead Imagery

Tim Y. Tang*, Daniele De Martini*, Shangzhe Wu[†], and Paul Newman*

*Mobile Robotics Group, University of Oxford [†]Visual Geometry Group, University of Oxford
{ttang, daniele, szwu, pnewman}@robots.ox.ac.uk

Abstract—Publicly available satellite imagery can be an ubiquitous, cheap, and powerful tool for vehicle localisation when a prior sensor map is unavailable. However, satellite images are not directly comparable to data from ground range sensors because of their starkly different modalities. We present a learned metric localisation method that not only handles the modality difference, but is cheap to train, learning in a self-supervised fashion without metrically accurate ground truth. By evaluating across multiple real-world datasets, we demonstrate the robustness and versatility of our method for various sensor configurations. We pay particular attention to the use of millimetre wave radar, which, owing to its complex interaction with the scene and its immunity to weather and lighting, makes for a compelling and valuable use case.

I. INTRODUCTION

The ability to localise relative to an operating environment is central to robot autonomy. Localisation using range sensors, such as lidars [25, 44] and, more recently, scanning millimetre wave radars [33, 36], is an established proposition. Both are immune to changing lighting conditions and directly measure scale, while the latter adds resilience to weather conditions.

Current approaches to robot localisation typically rely on a prior map built using a sensor configuration that will also be equipped on-board, for example a laser map for laser localisation. This paper looks at an alternative method. Public overhead imagery such as satellite images can be a reliable map source, as they are readily available, and often captures information also observable, albeit perhaps in some complex and incomplete way, by sensors on the ground. We can pose the localisation problem in a natural way: find the pixel location of a sensor in an overhead (satellite) image given range data taken from the ground. The task is, however, non-trivial because of the drastic modality difference between satellite image and sparse, ground-based radar or lidar.

Recent work on learning to localise a ground scanning radar against satellite images [39] provides a promising direction which addresses the modality difference by first generating a synthetic radar image from a satellite image. The synthetic image is then “compared” against a live radar image for pose estimation. Such an approach learns metric, cross-modality localisation in an end-to-end fashion, and therefore does not require hand-crafted features limited to a specific environment.

The method in [39] trains a multi-stage network, and needs pixel-wise aligned radar and satellite image pairs for supervision at all stages. This in turn requires sub-metre accurate ground truth position and sub-degree accurate ground truth

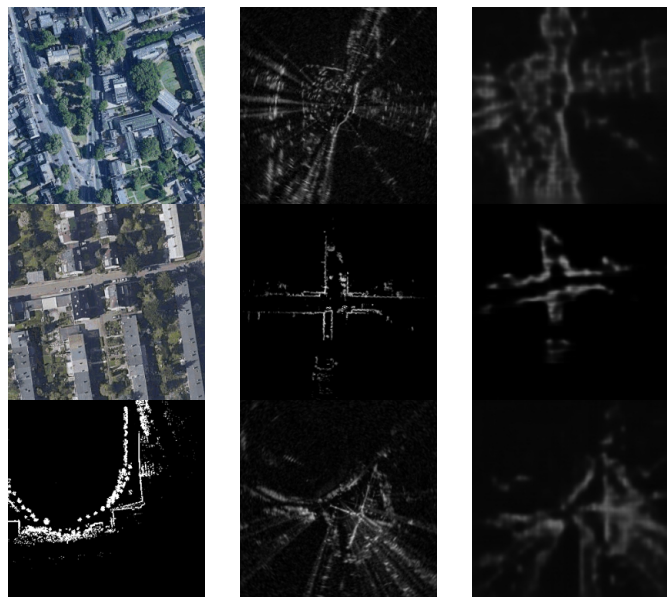


Fig. 1: Given a map image of modality \mathcal{A} (left) and a live data image of modality \mathcal{B} (middle), we wish to find the unknown $SE(2)$ offset between them. To do so, our method generates a synthetic image of modality \mathcal{B} (right) that is pixel-wise aligned with the map image, but contains the same appearance and observed scenes as the live data image. Top: localising radar data against satellite imagery. Middle: localising lidar data against satellite imagery. Bottom: localising radar data against prior lidar map.

heading. In practice, collecting accurate ground truth requires high-end GPS/Inertial Navigation System (INS), and possibly bundle adjustment along with other on-board sensor solutions, bringing in burdens in terms of cost and time consumption.

Building on [39], we propose a method for localising against satellite imagery that is self-supervised. The core idea of both approaches is to generate a synthetic image with the appearance and observed scenes of a live range sensor image, but pixel-wise aligned with the satellite image. We assume a coarse initial pose estimate is available from place recognition, such that there is reasonable overlap between the live ground sensor field-of-view and a queried satellite image.

Vitaly, here we make no use of metrically accurate ground truth for training. Note also that although designed for localising against satellite imagery, our method can naturally handle other forms of cross-modality registration, such as localising a radar against a prior lidar map. Figure 1 shows synthetic images generated by our method used for pose estimation.

To the best of our knowledge, this paper presents the first method to learn the cross-modality, metric localisation of

a range sensor in a self-supervised fashion. Our method is validated experimentally on multiple datasets and achieves performances on-par with a state-of-the-art supervised approach.

II. RELATED WORK

A. Localisation Using Overhead Images

Localisation using aerial or overhead images has been an interest for the community for over a decade. The methods in [24, 26, 34] localise a ground camera using aerial images, by detecting Canny edges from aerial imagery, and matching against lines detected by a ground camera. Several other vision-based approaches project the ground camera images to a top-down perspective via a homography, and compare against the aerial imagery by detecting lane markings [35], SURF features [32], or dense matching [37]. Recent work [13] localises a ground robot in a crop field by matching camera features against landmarks from an aerial map, and incorporated semantics of crops to reduce ambiguity.

Metric localisation of range sensors or point-clouds against overhead imagery requires further pre-processing due to the modality difference. Kaminsky et al. [19] projected point-clouds into images and matched against binary edge images from overhead imagery. The method in [19] also constructs a ray image by ray-tracing each point, and introduces a free-space cost to aid the image registration. The work by Veronese et al. [15] accumulates several lidar scans to produce dense lidar intensity images, which are then matched against satellite images utilising Normalised Mutual Information. Similar as [19], several other methods also pre-process the aerial image before matching against ground laser observations, for example using edge detection [21] or semantic segmentation [16]. Our method directly learns the metric localisation of a range sensor end-to-end, without the need for careful pre-processing.

B. Cross-Modality Localisation

Other forms of cross-modality localisation are also heavily studied by the community. A number of works are proposed to localise a forward facing camera against a prior 3D point-cloud map [43, 11, 47]. Carle and Barfoot [10] localised a ground laser scanner against an orbital elevation map. The works in [40, 8, 41, 31] localise an indoor lidar or stereo camera against architectural floor plans. Brubaker et al. [9] and Floros et al. [17] concurrently proposed matching visual odometry paths to road layouts from OpenStreetMap for localisation.

C. Learning-based State Estimation for Range Sensors

A number of recent works were proposed for learning odometry or localisation of lidars. Barsan et al. [7] represented lidar data as intensity images, and learned a deep embedding for metric localisation by comparing embeddings of live and map lidar intensity images. Methods such as [14, 27] learn deep lidar odometry by projecting lidar point-clouds into other representations before passing through the network. Lu et al. [30] used point-clouds as input to learn descriptors, and utilised 3D Convolutional Neural Networks (CNNs) for solving $SE(2)$ metric localisation by searching in a 3D cost

volume. In their later work, Lu et al. proposed a method to learn $SE(3)$ lidar point-cloud registration end-to-end [29].

As an emerging sensor for outdoor state estimation, learning-based methods were proposed for scanning Frequency-Modulated Continuous-Wave (FMCW) radars. Aldera et al. [3] utilised an encoder-decoder on polar image representation of radar scans to learn key-points for fast classical radar odometry [12]. Barnes et al. learned image-based radar odometry [5] by masking out regions distracting for pose estimation, and point-based radar odometry [4] by detecting key-points from radar images. Saftescu et al. [36] encoded images of radar polar scans through a rotation-invariant architecture to perform topological localisation (place recognition). These methods, however, are designed to compare data of the same sensor type, and does not address modality difference. Our approach is similar to [7, 14, 3, 5, 42, 36, 4] in that we also represent lidar and radar data as 2D images prior to passing through the network.

D. Unsupervised Image Generation

We seek to generate a synthetic image prior to pose computation, where there is no pixel-wise aligned target image for supervision. CycleGAN [48] achieves unsupervised image-to-image transfer between two domains \mathcal{X} and \mathcal{Y} , by learning two pairs of generators and discriminators, and enforcing cycle-consistency when an image is mapped from \mathcal{X} to \mathcal{Y} and back from \mathcal{Y} to \mathcal{X} , and vice versa. A number of other methods [28, 22] also utilise cycle-consistency, but make different assumptions on how the latent spaces of the two domains are treated. These methods are concerned with generating photo-realistic images. For the problem of metric localisation, however, we need to explicitly encourage the synthetic image to contain information appropriate for pose estimation.

Several prior works are also geometry-aware. The methods in [38, 45, 46] use separate encoders and/or decoders to disentangle geometry and appearance. The results are networks that can interpolate the geometry and appearance of the output images separately. Similarly, our method separately encodes information about appearance and the relative pose offset, resulting in an architecture where the two are disentangled.

III. OVERVIEW AND MOTIVATION

We seek to solve for the $SE(2)$ pose between a map image of modality \mathcal{A} and a live data image of modality \mathcal{B} . Our main focus is when modality \mathcal{A} is satellite imagery, while modality \mathcal{B} are range sensor data represented as an image.

Previously, RSL-Net [39] was proposed to solve for the metric localisation between matched pairs of radar and satellite images. In particular, a synthetic image is generated such that it preserves the appearance and observed scenes of the live radar image, and is pixel-wise aligned with the paired satellite image. The synthetic image and the live radar image are then projected onto deep embeddings, where their pose offset is found by maximising a correlation surface. We follow the same general approach, but, unlike RSL-Net, our method learns in a self-supervised fashion.

A. Hand-crafting Features vs. Learning

A number of works listed in Sections II-B and II-A can achieve decent accuracy on localising a ground range sensor against aerial imagery. However, they typically rely on pre-processing the aerial images using hand-crafted features or transforms designed for a specific set-up and may not generalise to other sensors or different, more complex environments. For example, [21] focuses on detecting edges from a campus dominated by buildings. While [15] directly matches accumulated lidar intensity images against aerial imagery without pre-processing, the same method is unlikely to work for radars.

Our data-driven approach instead learns to directly infer the geometric relationship across modalities, remaining free of hand-crafted features. We show in Section V when localising against satellite imagery, our method works for various types of scenes including urban, residential, campus, and highway.

B. Generating Images vs. Direct Regression

A naive approach would be to take a satellite image and a live data image as inputs, and directly regress the pose. As originally shown in [39], this led to poor results even for the supervised case. Our hypothesis is that when the two images are starkly different in appearance and observed scenes, the problem becomes too complex for direct regression to succeed.

Generating synthetic images first prior to pose estimation brings two advantages over directly regressing the pose. First, generating synthetic images is a simpler and less ill-posed problem than directly regressing the pose, particularly because we can utilise the live data image to condition the generation. Moreover, to generate images, the network loss is distributed over an entire image of $H \times W$ pixels, where H and W are height and width, instead of on just three parameters (x , y , and θ). This introduces greater constraint during optimisation.

C. Conditional Image Generation

Our method of conditional image generation takes in both a map (e.g., satellite) image and a live data image as inputs. An alternative approach is to learn a domain adaptation from the map modality \mathcal{A} to the live data modality \mathcal{B} , without conditioning on the live data image (e.g., standard image-to-image transfer such as CycleGAN [48]).

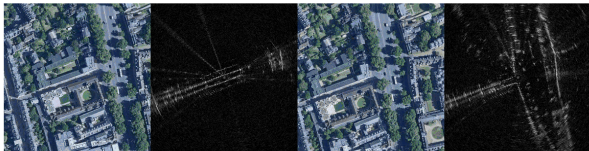


Fig. 2: Two radar images captured 15 seconds apart from each other (2 & 4), pixel-wise aligned with satellite images (1 & 3). Though the overlapping scenes in the satellite images are identical, the radar scans appear significantly different, as they capture different regions in their field-of-view.

In practise, the map (e.g., satellite) image is a denser representation of the environment than a frame of data captured by a range sensor. Only a fraction of the scenes captured in a satellite map are present in a ground sensor field-of-view, resulting in the scan to appear drastically different depending

on the sensor pose. Shown in Figure 2, the overlapping regions of the two satellite images are identical, while the two radar images observe different regions of the scene.

By using a naive image-to-image transfer approach, there is no guarantee the generated image will contain regions of the scene that are useful for pose comparison against the live data image. Figure 11 shows examples of images generated using CycleGAN [48], where the synthetic image highlights different scenes than what are observed by the live data image. The issue with observability or occlusion can potentially be handled by ray-tracing such as in [19]. However, not only is this computationally expensive, it does not apply to FMCW radars which have multiple range returns per azimuth (in which we are particularly interested). This problem is inherently addressed by our approach: by conditioning the image generation with the live data image, we can encourage the synthetic image to capture regions of the scene also observed by the live data image, as shown in Sections IV and V.

IV. SELF-SUPERVISED CROSS-MODALITY LOCALISATION

A. Rotation Inference

Given a paired map (e.g., satellite) image $A \in \mathcal{A}$ and live data image $B \in \mathcal{B}$ with an unknown $SE(2)$ offset, we seek to generate a synthetic image that contains the same appearance and observed scenes as B , but is pixel-wise aligned with A .

Let the $SE(2)$ pose difference between A and B be parametrised as $[x \ y \ \theta]^T$, such that by rotating B by θ , followed by a translation of $[x \ y]^T$, one can pixel-wise align B onto A . The image generation can be formulated as:

$$f(A, B) \rightarrow \tilde{B}_{\theta, \alpha} \quad (1)$$

where $\alpha = [x \ y]^T$. $\tilde{B}_{\theta, \alpha}$ is a generated image of modality \mathcal{B} that synthesises the input live sensor image B applied with a rotation of θ , followed by a translation of $\alpha = [x \ y]^T$. Thus, $\tilde{B}_{\theta, \alpha}$ is pixel-wise aligned with the input map image A , but contains the same observed scenes as B .

However, as originally noted in [39], the mapping in (1) is difficult to learn as the inputs A and B are offset by both a translation and a rotation. CNNs are inherently *equivariant* to translation, but not to rotation [23]. As a result, the CNNs in the generator cannot automatically utilise their mutual information and thereby capture their geometric relationship.

The method in [39] proposes to infer the rotation prior to image generation. Namely, reducing (1) to two steps:

$$f_R(A, B) \rightarrow B_\theta \quad (2)$$

$$f_G(A, B_\theta) \rightarrow \tilde{B}_{\theta, \alpha} \quad (3)$$

Here f_R is a function that infers the rotation offset θ between A and B , and outputs B_θ , which is input image B rotated by θ . Now, B_θ is rotation-aligned with the map frame, and therefore offset with A only by a translation, which CNNs can naturally handle. f_G is an image generation function that produces the synthetic image $\tilde{B}_{\theta, \alpha}$. The experiments in [39] show that learning (2) and (3) sequentially resulted in

better performance than learning (1) directly, as the former is congruous with the equivariance properties of CNNs.

In [39], the rotation inference function f_R is parametrised by a deep network as shown in Figure 3, where satellite imagery and radar images are used as an example. Given a coarse initial heading estimate, the live data image B is rotated a number of times with small increments to form a stack of rotated images $\{B\} = \{B_{\theta_0}, B_{\theta_1}, \dots, B_{\theta_n}\}$, where the number of rotations n and the increment are design parameters. Each rotated image is further concatenated with the map image to form a stacked tensor input of n pairs of map and live data images. The output of the network is a softmaxed image from $\{B\}$ that corresponds to B rotated to be rotation-aligned with A , namely B_θ . The core idea is that the network f_R will assign a large softmax weight to the image from $\{B\}$ whose heading most closely aligns with the map image A , and small weights to all other images in $\{B\}$.

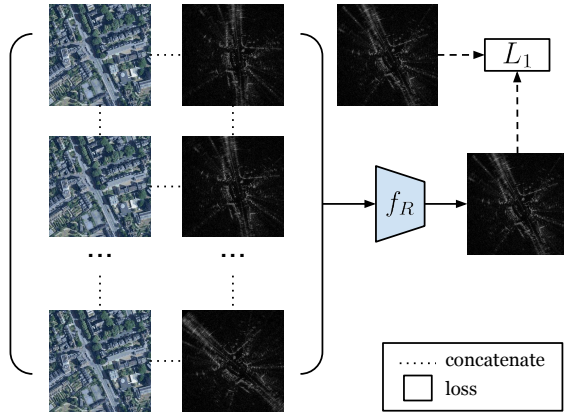


Fig. 3: Prior work in [39] proposes a network to infer the rotation offset. The rotation offset is found by softmaxing a stack of rotated radar images to produce a radar image with the same heading as the satellite image.

If metrically accurate heading ground truth θ is available, then one can rotate B to form an image target to B_θ used for supervising the rotation inference, as in Figure 3. In this work we assume this is never the case, thus the network for f_R must learn to infer the rotation offset self-supervised.

For this reason, while following the same architecture as [39], our method for inferring rotation uses a different training strategy that enables self-supervised learning. In order for the network f_R to produce the correct output, it must be able to infer the rotation from the solution space $\{B\}$, despite there being a modality difference between map image A and live data image B . We make the observation that if the network can infer the rotation offset from a stack of rotated live data images $\{B\}$, then, given a live data image B_{θ_i} , f_R should also be able to output A_{θ_i} from a stack of rotated map images $\{A\}$, where A_{θ_i} is rotation-aligned with B_{θ_i} . Specifically, if we have $B_{\theta_i} = B_\theta$, then the softmaxed map image from $\{A\}$ should be A , as A and B_θ are rotation-aligned.

As such, to learn rotation inference self-supervised, we need to pass through the network f_R twice. The first pass is identical as in the supervised approach in Figure 3, where we denote the output softmaxed image as B_{θ_i} . B_{θ_i} is then used as input

to the second pass through network f_R , together with a stack of map images $\{A\} = \{A, A_{\phi_0}, A_{\phi_1}, \dots, A_{\phi_m}\}$. The rotation angles $[\phi_0 \ \phi_1 \ \dots \ \phi_m]$ can be chosen randomly, and the order of $\{A\}$ is shuffled such that the original non-rotated map image A can be at any index within $\{A\}$. Each image is concatenated with B_{θ_i} to form the input stack for passing through f_R the second time. The network is supervised with an L_1 loss that enforces the output of the second pass to be the non-rotated map image A , which in turn enforces the output of the first pass B_{θ_i} to be B_θ , as B_θ is rotation-aligned with A . Our approach is shown in Figure 4. We use an increment of 2° when forming the rotation stack $\{B\}$.

The estimate for the rotation offset, $\hat{\theta}$, can then be found from the arg-softmax for the rotation stack $\{B\}$.

B. Image Generation

Given A and B_θ we seek to generate a synthetic image $\tilde{B}_{\theta, \alpha}$ as in (3), where $\tilde{B}_{\theta, \alpha}$ is pixel-wise aligned with A . [39] learns the image generation function by a supervised approach, concatenating A and B_θ , and applying an encoder-decoder architecture, as shown in Figure 5. This is possible since a target for the synthetic image $\tilde{B}_{\theta, \alpha}$ can be obtained by applying the ground truth transform.

To generate synthetic images self-supervised, we propose an architecture we call Pose-Aware Separable Encoder Decoder (PASED), shown in Figure 6. PASED is trained in two steps: the first is a pre-training, intra-modality process that can be supervised (top half of Figure 6), while the second handles cross-modality comparison (bottom half of Figure 6).

Taking two random images B^1 and B^2 in the live data modality \mathcal{B} from the training set, where B^1 and B^2 can be at arbitrary heading, we apply a known translation offset $\gamma \in \mathbb{R}^2$ to B^2 . This forms an image B_γ^2 that is a shifted version of B^2 . We pass B^1 through an appearance encoder E_a that encodes its appearance and observed scenes. B_γ^2 and B^2 are passed as inputs to a pose encoder E_p that encodes the translation offset between the input images. The latent spaces from E_a and E_p are combined before passing through a decoder D , which outputs a synthetic image \tilde{B}_γ^1 that is B^1 shifted by a translation γ . In other words, PASED discovers the translation offset between the two images passed as input to E_p , and applies the latent translation encoding to the input image of E_a . The pre-training can be supervised as γ is known, thus we can shift B^1 by γ to produce the target B_γ^1 . The fact that we use different images B^1 and B^2 for inputs to E_a and E_p ensures appearance and pose are disentangled from each other. As shown later, this allows modules of PASED to be separated and re-combined with newly learned modules.

In the second step, we fix the weights of E_a , E_p , and D which are optimised from the pre-training step. This narrows down the self-supervision problem to learning a cross-modality pose encoder E_p^* that discovers the translation offset between an image of modality \mathcal{A} and another of \mathcal{B} . Taking A and B_θ as inputs, E_p^* should encode the unknown translation offset α between them. Concurrently, B_θ is encoded by E_a , where the latent space is combined with the latent space produced

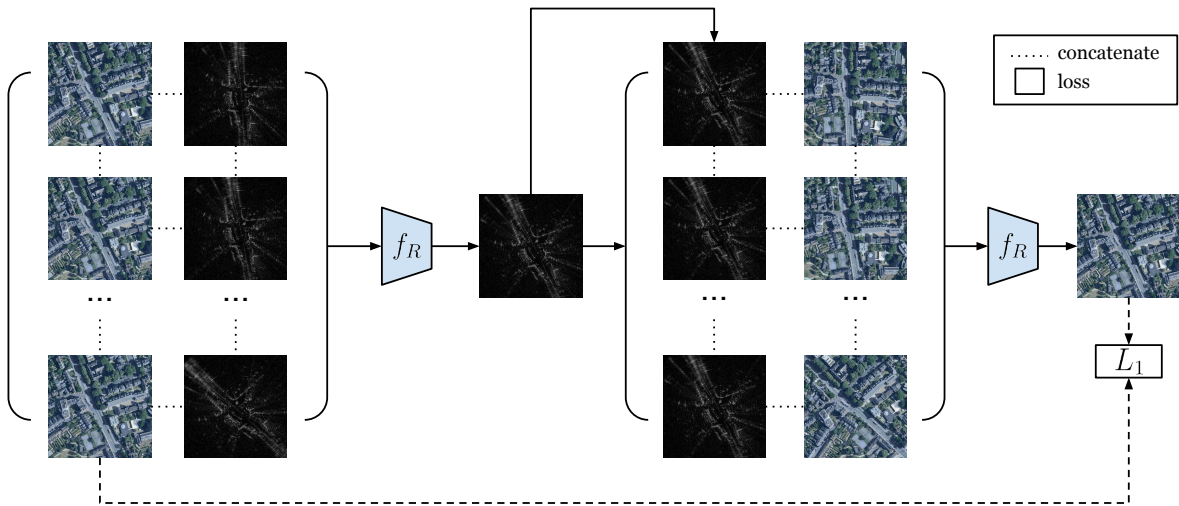


Fig. 4: Given A and a rotation stack $\{B\}$ the network f_R finds B_{θ_i} by taking softmax. Then, given B_{θ_i} and a rotation stack $\{A\}$, the network outputs a softmaxed map image from $\{A\}$. A loss is applied to enforce the output of the second pass to be A , which in turn enforces the output of the first pass to be B_{θ} . Here both symbols for f_R in the figure refer to the same network, but at different forward passes.

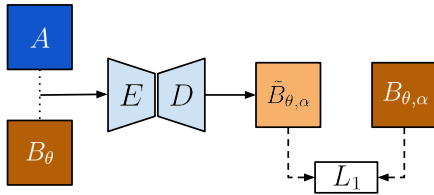


Fig. 5: Architecture for image generation in prior supervised approach [39].

by $E_p^*(A, B_{\theta})$, before decoded by D . This encoder-decoder combination will generate a synthetic image $\tilde{B}_{\theta, \alpha}$, which we do not have a target for.

We can apply a known shift to the centre position of A to query another map image A' , where A' is offset with B_{θ} by an unknown translation β . Using the same encoder-decoder combination as before, we can take A' and B_{θ} to generate a synthetic image $\tilde{B}_{\theta, \beta}$. Furthermore, given B_{θ} and the networks learned from pre-training, we can easily generate \tilde{B}_{θ} by encoding a zero shift. If we pass \tilde{B}_{θ} and $\tilde{B}_{\theta, \alpha}$ to the pre-trained pose encoder E_p , then the latent space will encode a shift of $-\alpha$. Combining this latent space with $E_a(\tilde{B}_{\theta, \beta})$, we can decode a synthetic image $\tilde{B}_{\theta, \beta - \alpha}$. Here $\beta - \alpha$ is a known value as it is the translation offset applied to A to obtain A' .

We can shift B_{θ} by $\beta - \alpha$ to get $B_{\theta, \beta - \alpha}$. Using $B_{\theta, \beta - \alpha}$ and B_{θ} , we can generate $\tilde{B}_{\theta, \beta - \alpha}$ with E_a , E_p , and D , shown on the bottom right of Figure 6. A loss can then be established between the two synthetic images $\tilde{B}_{\theta, \beta - \alpha}$, where the latter one is a target image created by modules with weights fixed. By back-propagation the loss optimises the network E_p^* . Alternatively we can use $B_{\theta, \beta - \alpha}$ as the target, but using $\tilde{B}_{\theta, \beta - \alpha}$ led to faster convergence.

For the loss to be minimised, two conditions must hold true. First, $\tilde{B}_{\theta, \beta}$ must have correctly encoded the appearance and observed scenes in B_{θ} . Second, $\tilde{B}_{\theta, \alpha}$ and $\tilde{B}_{\theta, \beta}$ must have the correct translations α and β , respectively. By satisfying these two constraints we can ensure E_p^* is able to discover the translation offset across modalities, and is compatible with

pre-trained networks E_a and D for image generation.

C. Pose Estimation

Taking $\tilde{B}_{\theta, \alpha}$ and B_{θ} , we embed them to a joint space, where their translation offset is found by maximising correlation on the learned embeddings. This can be performed efficiently in the Fourier domain, as is done in prior works that use a similar approach [7, 5, 39]. In this step, we can infer $\hat{\alpha} = [\hat{x} \ \hat{y}]^T$, which is our posterior estimate to the translation.

The embeddings are thus learned to further ensure the synthetic image and the live image can be correctly correlated. Without ground truth α , we can self-supervise using a similar approach as in learning PASED, by applying a known shift. The architecture for learning the embeddings is shown in Figure 7, where we denote the embedding network for real and synthetic images to be H_B and $H_{\tilde{B}}$, respectively. Given learned deep embeddings $\tilde{B}_{\theta, \beta}^{\dagger}$ and B_{θ}^{\dagger} , the translation offset by correlation maximisation is found to be $\hat{\beta}$. If we replace $\tilde{B}_{\theta, \beta}^{\dagger}$ with $\tilde{B}_{\theta, \alpha}^{\dagger}$ and reverse the order, the offset found will be $-\hat{\alpha}$. The sum of the two offsets $\beta - \alpha$ is known, and can be used to establish a loss term. Similar as in Section IV-B, $\tilde{B}_{\theta, \beta}$ is obtained by shifting the map image A to get A' .

V. EXPERIMENTAL VALIDATION

The overall pipeline for data flow at inference time is shown in Figure 8. The inference runs at about 10 Hz on a single 1080 Ti GPU. We evaluate on a large number of public, real-world datasets collected with vehicles equipped with on-board range sensors. The datasets we use come with metric ground truths that are decently accurate, though we noticed the GPS/INS solutions in certain places can drift up to a few metres.

We add large artificial pose offsets to the ground truth when querying for a satellite image, thereby simulating a realistic robot navigation scenario where the initial pose estimate can solve place recognition, but is too coarse for the robot's metric pose. Using a map (e.g., satellite) image queried at this coarse

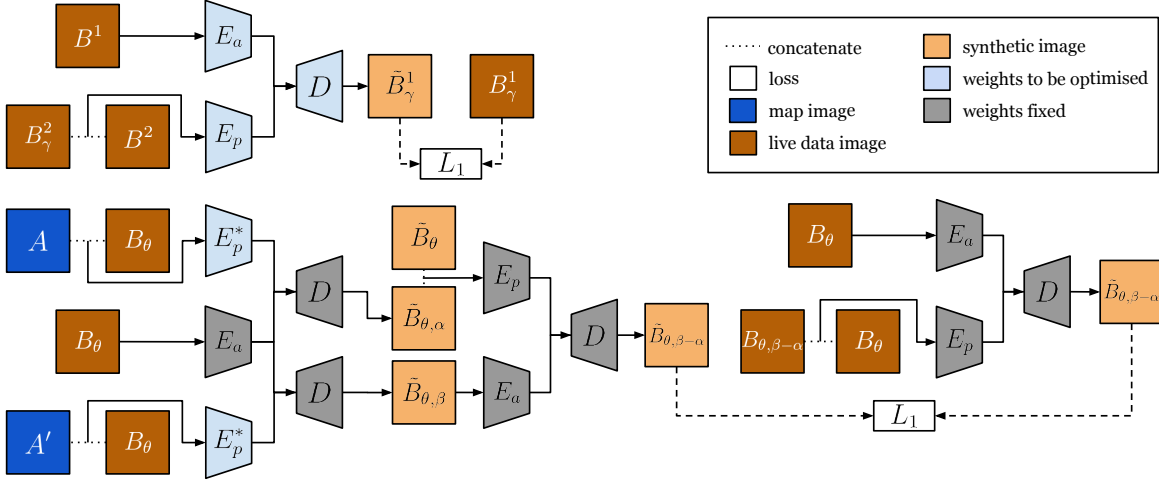


Fig. 6: Top: during pre-training, we can learn an appearance encoder E_a , and a pose encoder E_p that discovers the translation offset between an image of B and a shifted version of itself. Bottom: Taking E_a , E_p , and D and fixing their weights, we seek to learn E_p^* which discovers the translation offset between two images from different modalities. E_a , E_p , and D can provide the necessary geometric and appearance relationships used for learning E_p^* self-supervised.

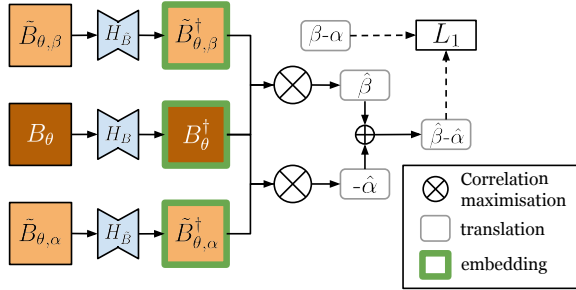


Fig. 7: The networks H_B and $H_{\tilde{B}}$ are learned to project real live images and synthetic images to a joint embedding, where their translation offset can be found by maximising correlation.

initial pose estimate, our method solves metric localisation by comparing against the live sensor data. The true pose offsets are hidden during training as our method is self-supervised, and are only revealed at test time for evaluation purposes.

The artificial offset is chosen such that the initial estimate has an unknown heading error in the range $[-\frac{\pi}{8}, \frac{\pi}{8}]$, therefore given the initial estimate θ_0 , the rotation inference must choose a solution space of at least $[\theta_0 - \frac{\pi}{8}, \theta_0 + \frac{\pi}{8}]$ to guarantee the correct solution can be found. We use a pixel-wise translation error in the range $[-25, 25]$ pixels. Depending on the resolution for a specific experiment, this corresponds to an error of at least $[-10\text{m}, 10\text{m}]$ and up to more than $[-20\text{m}, 20\text{m}]$.

A. Radar Localisation Against Satellite Imagery

We evaluate on two datasets with FMCW radar and GPS: the Oxford Radar RobotCar Dataset [6] and the MulRan Dataset [20]. The satellite images for RobotCar are queried using Google Maps Platform [2]. For MulRan they are queried using Bing Maps Platform [1], as high-definition Google satellite imagery is unavailable at the place of interest.

We benchmark against the prior supervised method RSL-Net [39] in our experiments, which is evaluated only on the RobotCar Dataset. Both datasets contain repeated traversals of the same routes. We separately train, validate, and test for every dataset, splitting the data as in Figure 9. For the

RobotCar Dataset, we split the trajectories the same way as in [39] for a fair comparison. For the RobotCar Dataset, the training set consists of training data from sequences no. 2, no. 5, and no. 6, while we test on the test data from sequence no. 2. For the MulRan Dataset, we used sequences KAIST 01 and Sejong 01. The RobotCar test set features an urban environment, while KAIST 01 is in a campus and Sejong 01 is primarily a highway.

	Mean Error (metric)			(pixel)	
	x (m)	y (m)	θ ($^\circ$)	x	y
RobotCar (ours)	3.44	5.40	3.03	3.97	6.23
MulRan (ours)	6.02	7.02	2.92	7.64	8.91
RobotCar (RSL-Net [39], supervised)	2.74	4.26	3.12	3.16	4.92
MulRan (RSL-Net)	5.85	7.11	1.88	7.42	9.03

TABLE I: Mean error for radar localisation against satellite imagery.

We test on every fifth frame, resulting in 201 frames from the RobotCar Dataset and 358 from the MulRan Dataset, spanning a total distance of near 4 km. The resolution used is 0.8665 m/pixel for RobotCar and 0.7876 m/pixel for MulRan. The mean errors are reported in Table I.

B. Lidar Localisation Against Satellite Imagery

For this experiment, we evaluate on the RobotCar Dataset [6] which also has two Velodyne HDL-32E lidars mounted in a tilted configuration, and KITTI (raw dataset) [18] which has a Velodyne HDL-64E lidar and GPS data.

For the RobotCar Dataset, the trajectories are split into training, validation, and test sets approximately the same way as in Section V-A. For the KITTI Dataset, the training set includes sequences 20110929_drive0071, 20110930_drive0028, and 20111003_drive0027. Sequence 20110926_drive0117 is used for validation. Finally, data in 20111003_drive0034 are split into training and test, as shown in Figure 9. To turn 3D lidar point-clouds to lidar images, the point-clouds are projected to the $x - y$

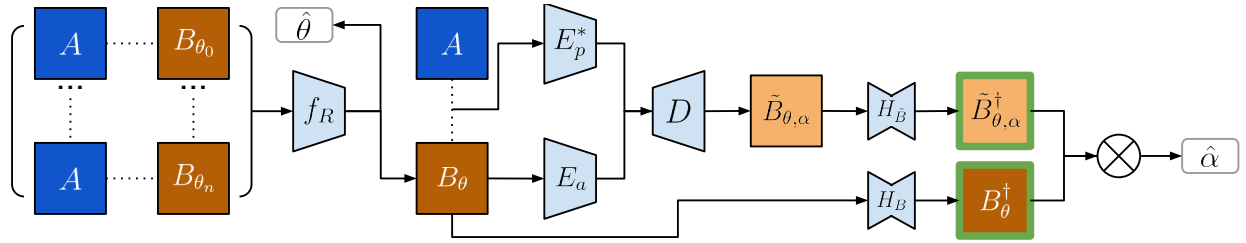


Fig. 8: Overall data flow of our method at inference: given map image A and live data image B , based on the initial heading estimate, we form a stack of rotated images $\{B_{\theta_0}, \dots, B_{\theta_n}\}$, from which f_R discovers B_{θ} that is B rotated to be rotation-aligned with A . This process also infers the heading estimate $\hat{\theta}$. A and B_{θ} are used to generate a synthetic image $\tilde{B}_{\theta, \alpha}$ that has the same appearance and observed scene as B_{θ} and is pose-aligned with A . $\tilde{B}_{\theta, \alpha}$ and B_{θ} are projected to deep embeddings $\tilde{B}_{\theta, \alpha}^{\dagger}$ and B_{θ}^{\dagger} , where the estimate for the translation offset $\hat{\alpha}$ is found by correlation maximisation.

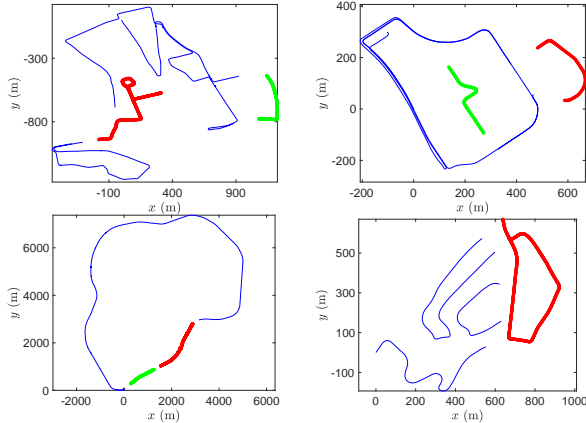


Fig. 9: Training (blue), validation (green), and test (red) trajectories for RobotCar (top left), KAIST (top right), Sejong (bottom left) and 20111003_drive0034 (bottom right). Certain data are removed to avoid overlap between the splits.

plane. We discard points with z values smaller than zero to remove ground points when creating the lidar images.

Since lidars have a shorter range than radars, we use satellite images of a greater zoom level, with resolution 0.4332 m/pixel for RobotCar and 0.4592 m/pixel for KITTI. The test set consists of 200 frames for RobotCar and 253 for KITTI, spanning a total distance of near 3 km. The test set for KITTI features a residential area. The results are reported in Table II.

	Mean Error (metric)			(pixel)	
	x (m)	y (m)	θ ($^{\circ}$)	x	y
RobotCar (ours)	1.54	1.85	2.34	3.55	4.27
KITTI (ours)	3.05	3.13	1.67	6.64	6.82
RobotCar (RSL-Net)	2.31	2.55	2.08	5.33	5.89
KITTI (RSL-Net)	2.45	2.79	1.59	5.34	6.08

TABLE II: Mean error for lidar localisation against satellite imagery.

C. Radar Localisation Against Prior Lidar Map

	Mean Error (metric)			(pixel)	
	x (m)	y (m)	θ ($^{\circ}$)	x	y
RobotCar (ours)	2.21	2.57	2.65	2.55	2.97
MulRan (ours)	3.57	3.26	2.15	4.53	4.13
RobotCar (RSL-Net)	2.66	3.41	2.45	3.07	3.93
MulRan (RSL-Net)	3.37	2.61	1.40	4.28	3.32
RobotCar (CycleGAN)	6.41	9.05	2.65	7.40	10.44
MulRan (CycleGAN)	4.84	4.39	2.15	6.14	5.58

TABLE III: Mean error for radar localisation against prior lidar map.

Though our method is designed for localising against satellite imagery, we show it can also handle more standard forms

of cross-modality localisation. Here we build a lidar map using a prior traversal, and localise using radar from a later traversal.

We demonstrate on the RobotCar and MulRan datasets, where we use the same resolution as in Section V-A. For RobotCar, we use ground truth to build a lidar map from sequence no. 2. Radar data in the training sections from no. 5 and no. 6 as in Figure 9 form the training set, while the test section from sequence no. 5 forms the test set. For MulRan, lidar maps are built from KAIST 01 and Sejong 01, and we localised using radar data from KAIST 02 and Sejong 02, which are split into training, validation, and test sets. This resulted in a test set consisting of 201 frames from RobotCar and 272 frames from MulRan, spanning a total distance of near 4 km. The localisation results are shown in Table III.

This experiment is more suitable for naive image generation methods such as CycleGAN [48] than previous experiments, as the field-of-view is considerably more compatible when both modalities are from range sensors. In Table III, we show results where we replaced the image generation stage of our method by CycleGAN, and kept other modules. The localisation results are however much worse when modality A is satellite imagery, as shown qualitatively in Figure 11.

D. Online Pose-Tracking System

In prior experiments we assumed place recognition is always available, providing a coarse initial estimate for every frame. Here we present a stand-alone pose-tracking system by continuously localising against satellite imagery. Given a coarse initial estimate (e.g., from GPS) for the first frame, the vehicle localises and computes its pose within the satellite map. The initial estimate for every frame onward is then set to be the computed pose of the previous frame. We only need place recognition once at the very beginning; the vehicle then tracks its pose onward without relying on any other measurements.

1) *Introspection*: As localising using satellite imagery is challenging, the result will not always be accurate. Our method, however, naturally allows for introspection. A synthetic image $\tilde{B}_{\theta, \alpha}$ was generated from A and B_{θ} . We can apply a known small translation offset δ to A to form A_{δ} . Taking A_{δ} and B_{θ} we can generate $\tilde{B}_{\theta, \alpha + \delta}$. Finally, we can compute a translation offset $\hat{\delta}$ by passing $\tilde{B}_{\theta, \alpha + \delta}$ and $\tilde{B}_{\theta, \alpha}$ through the learned embeddings and maximising correlation.

Let $d_{\text{intro}} = \|\delta - \hat{\delta}\|$. A large value of d_{intro} indicates the generated images are erroneous. This allows us to examine

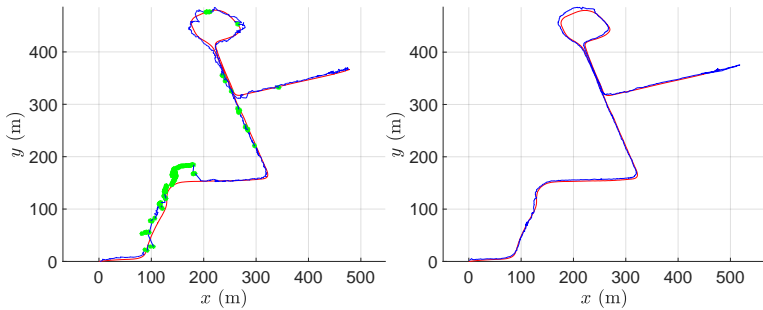


Fig. 10: Estimated pose (blue) vs. ground truth pose (red) for localising a radar (left) and a lidar (right) against satellite imagery. Our system continuously tracks the vehicle’s pose over 1 km, where we occasionally fall back to odometry for the radar experiment (green). Our system is stand-alone and requires GPS only for the first frame.

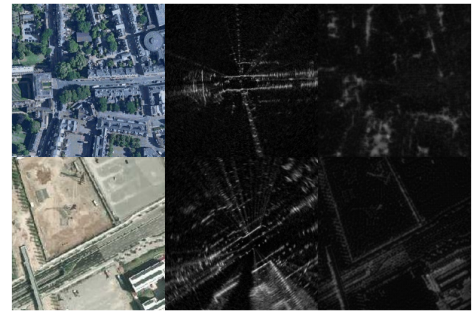


Fig. 11: Results of CycleGAN: satellite image A (left), ground truth radar image $B_{\theta, \alpha}$ (middle), synthetic radar image $\tilde{B}_{\theta, \alpha}$ (right). This led to large localisation error as $\tilde{B}_{\theta, \alpha}$ does not contain scenes observed by B .

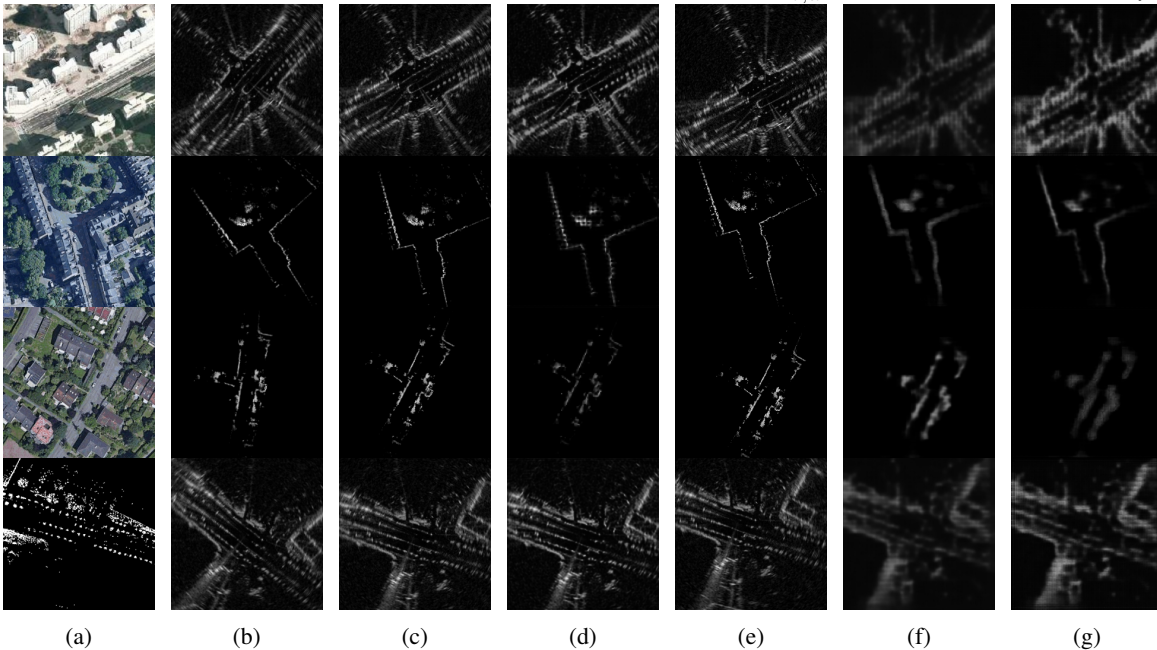


Fig. 12: Images at various stages of our method: map image A (a), live data image B (b), output of rotation inference B_{θ} (c), embedding B_{θ}^{\dagger} (d), pixel-wise aligned ground truth $B_{\theta, \alpha}$ (e), synthetic image $\tilde{B}_{\theta, \alpha}$ (f), embedding $\tilde{B}_{\theta, \alpha}^{\dagger}$ (g). From top to bottom: radar localisation against satellite imagery evaluated on MulRan, lidar localisation against satellite imagery evaluated on RobotCar and KITTI, radar localisation against lidar map evaluated on MulRan.

the solution quality; our system falls back to using odometry for dead-reckoning when d_{intro} exceeds a threshold. We do not require high-quality odometry, but rather only use a naive approach by directly maximising correlation between two consecutive frames without any learned modules. In our experiments, we set δ to be $[10 \ 10]^T$, and d_{intro} to be 5.

2) *Results*: We conduct two experiments on the test set of RobotCar, one where we track a radar using satellite imagery, and one where we track a lidar. For both experiments we run localisation at 4 Hz. The results are shown in Figure 10. If the solution error is too large, then the initial estimate will be too off for a sufficient overlap between the next queried satellite image and live data, resulting in losing track of the vehicle. Though the solution error can be larger than 10 m at times, our system continuously localises the vehicle for over a kilometre without completely losing track. For the lidar experiment, the solutions are sufficiently accurate to not require any odometry. Our experiments are single-frame localisations, and we make

no attempt at windowed/batch optimisation or loop closures.

E. Further Qualitative Results

Additional qualitative results are presented in Figure 12 showing various stages of our methods for different modalities.

VI. CONCLUSION AND FUTURE WORK

We present self-supervised learning to address cross-modality metric localisation between satellite imagery and on-board range sensors, without using metrically accurate ground truth for training. Our method is validated across a large number of experiments for multiple modes of localisation, with results on-par with prior supervised approach. A coarse initial pose estimate is needed for our method to compute metric localisation. An extension would then be to solve place recognition for a range sensor within a large satellite map.

ACKNOWLEDGMENTS

We thank Giseop Kim from IRAP Lab, KAIST for providing GPS data for the MulRan Dataset.

REFERENCES

- [1] Bing Maps. <https://docs.microsoft.com/en-us/bingmaps/>.
- [2] Google Maps Platform. <https://developers.google.com/maps/documentation/maps-static/intro/>.
- [3] Roberto Aldera, Daniele De Martini, Matthew Gadd, and Paul Newman. Fast Radar Motion Estimation with a Learnt Focus of Attention Using Weak Supervision. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 1190–1196. IEEE, 2019.
- [4] Dan Barnes and Ingmar Posner. Under the Radar: Learning to Predict Robust Keypoints for Odometry Estimation and Metric Localisation in Radar. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Paris, 2020. URL <https://arxiv.org/abs/2001.10789>.
- [5] Dan Barnes, Rob Weston, and Ingmar Posner. Masking by Moving: Learning Distraction-Free Radar Odometry from Pose Information. In *Conference on Robot Learning (CoRL)*, 2019. URL <https://arxiv.org/pdf/1909.03752>.
- [6] Dan Barnes, Matthew Gadd, Paul Murcutt, Paul Newman, and Ingmar Posner. The Oxford Radar RobotCar Dataset: A Radar Extension to the Oxford RobotCar Dataset. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Paris, 2020. URL <https://arxiv.org/abs/1909.01300>.
- [7] Ioan Andrei Barsan, Shenlong Wang, Andrei Pokrovsky, and Raquel Urtasun. Learning to Localize Using a LiDAR Intensity Map. In *CoRL*, pages 605–616, 2018.
- [8] Federico Boniardi, Tim Caselitz, Rainer Kümmerle, and Wolfram Burgard. Robust LiDAR-based Localization in Architectural Floor Plans. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3318–3324. IEEE, 2017.
- [9] Marcus A Brubaker, Andreas Geiger, and Raquel Urtasun. Lost! Leveraging the Crowd for Probabilistic Visual Self-localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3057–3064, 2013.
- [10] Patrick JF Carle and Timothy D Barfoot. Global Rover Localization by Matching Lidar and Orbital 3D Maps. In *2010 IEEE International Conference on Robotics and Automation*, pages 881–886. IEEE, 2010.
- [11] Tim Caselitz, Bastian Steder, Michael Ruhnke, and Wolfram Burgard. Monocular Camera Localization in 3D Lidar Maps. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1926–1931. IEEE, 2016.
- [12] Sarah H Cen and Paul Newman. Radar-Only Ego-Motion Estimation in Difficult Settings via Graph Matching. *arXiv preprint arXiv:1904.11476*, 2019.
- [13] Nived Chebrolu, Philipp Lottes, Thomas Läbe, and Cyrill Stachniss. Robot Localization Based on Aerial Images for Precision Agriculture Tasks in Crop Fields. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 1787–1793. IEEE, 2019.
- [14] Younggun Cho, Giseop Kim, and Ayounng Kim. DeepLO: Geometry-Aware Deep LiDAR Odometry. *arXiv preprint arXiv:1902.10562*, 2019.
- [15] Lucas de Paula Veronese, Edilson de Aguiar, Rafael Correia Nascimento, Jose Guivant, Fernando A Auat Cheein, Alberto Ferreira De Souza, and Thiago Oliveira-Santos. Re-emission and Satellite Aerial Maps Applied to Vehicle Localization on Urban Environments. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4285–4290. IEEE, 2015.
- [16] Can Ulas Dogruer, A Bugra Koku, and Melik Dolen. Outdoor Mapping and Localization Using Satellite Images. *Robotica*, 28(7):1001–1012, 2010.
- [17] Georgios Floros, Benito Van Der Zander, and Bastian Leibe. OpenStreetSLAM: Global Vehicle Localization Using OpenStreetMaps. In *2013 IEEE International Conference on Robotics and Automation*, pages 1054–1059. IEEE, 2013.
- [18] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision Meets Robotics: The KITTI Dataset. *International Journal of Robotics Research (IJRR)*, 2013.
- [19] Ryan S Kaminsky, Noah Snavely, Steven M Seitz, and Richard Szeliski. Alignment of 3D Point Clouds to Overhead Images. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 63–70. IEEE, 2009.
- [20] Giseop Kim, Yeong Sang Park, Younghun Cho, Jinyong Jeong, and Ayounng Kim. MulRan: Multimodal Range Dataset for Urban Place Recognition. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2020. Submitted. Under Review.
- [21] Rainer Kümmerle, Bastian Steder, Christian Dornhege, Alexander Kleiner, Giorgio Grisetti, and Wolfram Burgard. Large Scale Graph-based SLAM using Aerial Images as Prior Information. *Autonomous Robots*, 30(1):25–39, 2011.
- [22] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse Image-to-Image Translation via Disentangled Representations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 35–51, 2018.
- [23] Karel Lenc and Andrea Vedaldi. Understanding Image Representations by Measuring Their Equivariance and Equivalence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 991–999, 2015.
- [24] Keith Yu Kit Leung, Christopher M Clark, and Jan P Huissoon. Localization in Urban Environments by Matching Ground Level Video Images with an Aerial Image. In *2008 IEEE International Conference on Robotics and Automation*, pages 551–556. IEEE, 2008.
- [25] Jesse Levinson and Sebastian Thrun. Robust Vehicle Localization in Urban Environments Using Probabilistic Maps. In *2010 IEEE International Conference on Robotics and Automation*, pages 4372–4378. IEEE, 2010.

- [26] Ang Li, Vlad I Morariu, and Larry S Davis. Planar Structure Matching Under Projective Uncertainty for Geolocation. In *European Conference on Computer Vision*, pages 265–280. Springer, 2014.
- [27] Qing Li, Shaoyang Chen, Cheng Wang, Xin Li, Chenglu Wen, Ming Cheng, and Jonathan Li. LO-Net: Deep Real-time Lidar Odometry. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8473–8482, 2019.
- [28] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised Image-to-Image Translation Networks. In *Advances in Neural Information Processing Systems*, pages 700–708, 2017.
- [29] Weixin Lu, Guowei Wan, Yao Zhou, Xiangyu Fu, Pengfei Yuan, and Shiyu Song. DeepVCP: An End-to-End Deep Neural Network for Point Cloud Registration. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [30] Weixin Lu, Yao Zhou, Guowei Wan, Shenhua Hou, and Shiyu Song. L3-Net: Towards Learning Based LiDAR Localization for Autonomous Driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6389–6398, 2019.
- [31] Malcolm Mielle, Martin Magnusson, and Achim J Lilienthal. The Auto-Complete Graph: Merging and Mutual Correction of Sensor and Prior Maps for SLAM. *Robotics*, 8(2):40, 2019.
- [32] Masafumi Noda, Tomokazu Takahashi, Daisuke Deguchi, Ichiro Ide, Hiroshi Murase, Yoshiko Kojima, and Takashi Naito. Vehicle Ego-Localization by Matching in-Vehicle Camera Images to an Aerial Image. In *Asian Conference on Computer Vision*, pages 163–173. Springer, 2010.
- [33] Yeong Sang Park, Jinyong Jeong, Youngsik Shin, and Ayoung Kim. Radar Dataset for Robust Localization and Mapping in Urban Environment. In *ICRA Workshop on Dataset Generation and Benchmarking of SLAM Algorithms for Robotics and VR/AR*, Montreal, May. 2019.
- [34] Martin P Parsley and Simon J Julier. Towards the Exploitation of Prior Information in SLAM. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2991–2996. IEEE, 2010.
- [35] Oliver Pink. Visual Map Matching and Localization Using a Global Feature Map. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 1–7. IEEE, 2008.
- [36] Stefan Saftescu, Matthew Gadd, Daniele De Martini, Dan Barnes, and Paul Newman. Kidnapped Radar: Topological Radar Localisation using Rotationally-Invariant Metric Learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Paris, 2020. URL <https://arxiv.org/abs/2001.09438>.
- [37] Turgay Senlet and Ahmed Elgammal. A Framework for Global Vehicle Localization Using Stereo Images and Satellite and Road maps. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 2034–2041. IEEE, 2011.
- [38] Zhixin Shu, Mihir Sahasrabudhe, Riza Alp Guler, Dimitris Samaras, Nikos Paragios, and Iasonas Kokkinos. Deforming Autoencoders: Unsupervised Disentangling of Shape and Appearance. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 650–665, 2018.
- [39] T. Y. Tang, D. De Martini, D. Barnes, and P. Newman. RSL-Net: Localising in Satellite Images From a Radar on the Ground. *IEEE Robotics and Automation Letters*, 5(2):1087–1094, April 2020. ISSN 2377-3774. doi: 10.1109/LRA.2020.2965907.
- [40] Xipeng Wang, Ryan J Marcotte, and Edwin Olson. GLFP: Global Localization from a Floor Plan.
- [41] Xipeng Wang, Steve Vozar, and Edwin Olson. FLAG: Feature-based Localization between Air and Ground. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3178–3184. IEEE, 2017.
- [42] Rob Weston, Sarah Cen, Paul Newman, and Ingmar Posner. Probably Unknown: Deep Inverse Sensor Modelling Radar. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5446–5452. IEEE, 2019.
- [43] Ryan W Wolcott and Ryan M Eustice. Visual Localization within Lidar Maps for Automated Urban Driving. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 176–183. IEEE, 2014.
- [44] Ryan W Wolcott and Ryan M Eustice. Fast LIDAR Localization Using Multiresolution Gaussian Mixture Maps. In *2015 IEEE international Conference on Robotics and Automation (ICRA)*, pages 2814–2821. IEEE, 2015.
- [45] Wayne Wu, Kaidi Cao, Cheng Li, Chen Qian, and Chen Change Loy. Transgaga: Geometry-Aware Unsupervised Image-to-Image Translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8012–8021, 2019.
- [46] Xianglei Xing, Tian Han, Ruiqi Gao, Song-Chun Zhu, and Ying Nian Wu. Unsupervised Disentangling of Appearance and Geometry by Deformable Generator Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10354–10363, 2019.
- [47] Yuquan Xu, Vijay John, Seiichi Mita, Hossein Tehrani, Kazuhisa Ishimaru, and Sakiko Nishino. 3D Point Cloud Map Based Vehicle Localization Using Stereo Camera. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 487–492. IEEE, 2017.
- [48] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2223–2232, 2017.