

# U-BEV: Height-aware Bird’s-Eye-View Segmentation and Neural Map-based Relocalization

Andrea Boscolo Camiletto<sup>1,2</sup>, Alfredo Bochicchio<sup>3</sup>, Alexander Liniger<sup>3</sup>, Dengxin Dai<sup>3</sup>, Abel Gawel<sup>3</sup>

**Abstract**—Efficient relocalization is essential for intelligent vehicles when GPS reception is insufficient or sensor-based localization fails. Recent advances in Bird’s-Eye-View (BEV) segmentation allow for accurate estimation of local scene appearance and in turn, can benefit the relocalization of the vehicle. However, one downside of BEV methods is the heavy computation required to leverage the geometric constraints. This paper presents U-BEV, a U-Net inspired architecture that extends the current state-of-the-art by allowing the BEV to reason about the scene on multiple height layers before flattening the BEV features. We show that this extension boosts the performance of the U-BEV by up to 4.11 IoU. Additionally, we combine the encoded neural BEV with a differentiable template matcher to perform relocalization on neural SD-map data. The model is fully end-to-end trainable and outperforms transformer-based BEV methods of similar computational complexity by 1.7 to 2.8 mIoU and BEV-based relocalization by over 26% Recall Accuracy on the nuScenes dataset.

## I. INTRODUCTION

Image-based relocalization and place recognition are among the fundamental challenges for autonomous agents and have been active research topics in the robotics community for over a decade [1], [2]. Efficient and accurate relocalization is especially relevant in Autonomous Driving, where agents travel tens of meters per second in constantly changing environments. The need for relocalization arises when GPS signals are unreliable, such as in densely populated cities where Urban Canyons are common. While the relocalization problem is well studied with high-fidelity map data and onboard sensing, the need for using camera-only systems and lightweight map representations drives contemporary research. Additionally, algorithms need fast inference times with strong generalization capabilities and cover degenerate localization cases where only weak perceptual cues are present. Current image-based methods however are either accurate but prohibitively expensive [3]–[6], suffer under degenerate cases [7], or rely on highly curated [8] or heavy representations [9].

Recent advances in Bird’s-Eye-View (BEV)-based scene understanding boosted the perception capabilities of autonomous agents [4], [10]–[14]. These representations lend themselves well to localization cases as they are trained on common autonomous driving map data. Hence, new BEV-based localization methods [8], [15], [16] elegantly extend cross-view localization by predicting from camera images a local BEV-representation on the 2D ground plane and

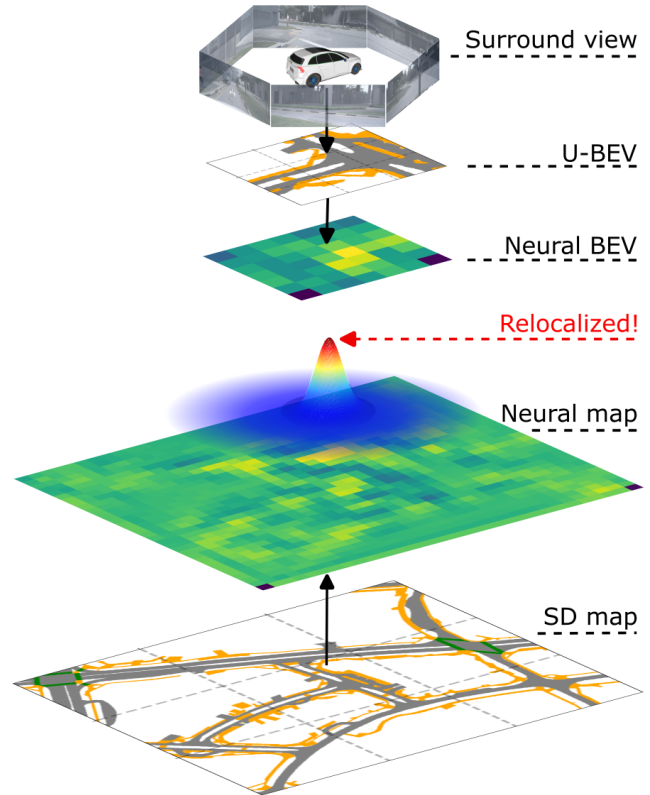


Fig. 1: U-BEV proposes a novel BEV representation from surround-view images for efficient neural relocalization in SD map data.

matching it to a semantic map. However, the flat world assumption reduces BEV accuracy in many 3D real-world cases and lowers localization accuracy.

In this work, we propose a new approach for estimating BEVs that leverages height-aware feature embedding, allowing the network to reason on the depth dimension without the need for heavy computations. The base architecture leverages an encoder-decoder structure with an additional geometrical projection layer to reason both in the image (during the encoding) and BEV space (during the decoding). We advocate for a 2-step localization approach, where autonomous agents first globally estimate their location within several meters and then rely on local methods to obtain application-required accuracies. Therefore, we utilize lightweight Standard Definition (SD) map data and aim for accuracies of one-shot relocalization below 10m. In this work, we integrate our BEV representation with a differentiable template matcher,

<sup>1</sup>Work done while at BayerLab, Huawei Zurich Research Center.

<sup>2</sup>MPI for Informatics. aboscolo@mpi-inf.mpg.de

<sup>3</sup>BayerLab, Huawei. firstname.lastname@huawei.com

which allows for end-to-end training and real-time relocalization. The localization architecture can work on arbitrary BEV methods by encoding them and the corresponding map data into a Neural representation. This work, combining U-BEV with the relocalization module outperforms both other BEV methods and contemporary BEV-based localization on the nuScenes dataset by over 26.4% Recall Accuracy at 10m. In summary, this paper presents the following contributions:

- 1) A new lightweight U-BEV architecture that is geometrically constrained and leverages the height of the points from the ground instead of their depth from the camera.
- 2) An end-to-end trainable real-time global localization algorithm between Neural BEV and Neural encoded SD maps.
- 3) Improved BEV (1.7 to 2.8 higher IoU) and localization performance (+26.4% Recall Accuracy at 10m) on the nuScenes dataset.

## II. RELATED WORK

This publication contributes to the state-of-the-art in BEV estimation and image-based relocalization. We review related work in the following.

### A. Bird's-Eye-View Segmentation

Semantic Segmentation of the car's surroundings in BEV is a common task in the autonomous driving domain. The goal is to flatten and discretize the environment around the car into tiles of fixed size and segment each tile into labels e.g., road, sidewalk, building, etc.

This segmentation is often addressed as a side task to the goal of 3D object detection and leverages both camera data and LiDAR input [4], [17], [18]. Recently, the community is shifting towards addressing the problem with cameras only [3], [10], [12], [19]–[21].

In this shift, handling depth is the main differentiator. To reason in BEV space, pixels require projection from the image to the world coordinates. This in turn requires spatial understanding. One line of works leverages a discretized explicit depth estimation [3], [4], [11], which is used to weight the features along the ray accordingly. Another option is to learn a mapping from image space to BEV space via a transformer or attention layer like [10], [22]. Finally, features can also be warped to the BEV space using homographies explicitly [13] or learned transformations [14].

While predicting the depth and projecting accordingly is undeniably a desirable geometric constraint, it can be resource-intensive. A common trend is to discretize the depth in the range of [1, 60] meters with a step size of 0.5 meter [4], [23] which leads to the creation of a very dense point cloud representation. In this work, we propose a different direction: instead of predicting dense depth, we estimate a sparse set of height layers.

### B. Image-based Relocalization

Image-based relocalization is closely related to Visual Place Recognition which defines the problem of retrieving

a query image from a database [2], [24] and estimating the (relative) pose. The most popular approach is to estimate local feature embeddings per image and either query them in a database (map) [24] or regress the map localization by spatially rolling out image features. When viewpoints are vastly different as is the case in this work, the problem is coined as cross-view localization. Recent works bridge the viewpoint difference using a sparse spatial-semantic representation [7], [25], [26], leverage generative networks that generate ground-view images from aerial views and perform database retrieval [27], [28], or learn a joint feature embedding for different views [29]–[35]. Another line of work lifts this challenge to the semantic domain to further bridge the perceptual gap between ground-view imagery (e.g., surround-view images from a car) to top-down map data. They generate BEV representations from ground-view images and match them to a learned or off-the-shelf semantic map [7], [8], [15], [16]. This work follows a similar approach, but in contrast to [15], [16] employs the more powerful U-BEV representation that allows localization in lower detailed maps. The proposed work also bridges large-scale localization errors of hundreds of meters, larger than [8] which addresses  $2m$  errors, and covers more degenerate cases than [7], e.g., our method is able to localize on open roads without distinct landmarks such as poles or signs.

## III. METHOD

The proposed full algorithm localizes a set of surround-view images in an SD map. It generates a local BEV representation from the surround images and a neural map encoding from an SD map tile given a coarse 3D location prior  $\xi_{init} = (x_{init}, y_{init}, \phi_{init})$  from onboard sensors (e.g. a noisy GPS signal and a compass). A deep template matcher then slides the local Neural BEV over the global Neural map resulting in a similarity map. The localization finishes by returning the Soft-Argmax of the similarity map. An overview of our method is illustrated in Fig. 2.

### A. Bird's Eye View Reconstruction

We propose a novel lightweight and accurate BEV architecture to reconstruct the car's surroundings from a set of surround-view images. Our model, called U-BEV, is structured as an Encoder-Decoder with an in-between layer that geometrically projects the extracted features of the bottleneck and skip connections from the image to the BEV space. An overview can be seen in Fig. 4.

Given a set of 6 images and their intrinsic and extrinsic calibrations, we predict a BEV  $\mathbf{B} \in \mathbb{R}^{S \times S \times N}$ , where  $S$  is the size of the BEV in pixel and  $N$  is the number of labels available in the map. We use the center of the back wheel axis as our origin point following the convention in the nuScenes dataset [36].

**Feature Extraction:** We extract features at different resolutions from all 6 images with a light-weight pretrained EfficientNet [37] backbone, a common approach in smaller models [10], [11]. Specifically, we extract features at strides  $\times 2, \times 4, \times 8, \times 16$ , intentionally omitting the last stride for



the ratio of the resolutions and channels constant, we can plug them into a classic decoder style with skip connections, yielding the final BEV output (Yellow part in Fig. 4).

### B. Map encoding

Maps are passed into our system as rasterized boolean  $N$ -channels discretized surfaces, where  $N$  is the number of classes, i.e. each semantic class is assigned a separate channel. In the case of polygon representations, as common in SD maps for autonomous driving, we preprocess the map by rasterizing the polygons of each class into channels.

### C. Localization

To perform localization, we leverage both a local BEV  $\mathbf{B}$  of the surroundings produced by the U-BEV model and a map tile of the global map  $\mathbf{M}_{loc} \in \mathbb{R}^{L \times L \times N}$  cropped given a coarse location prior  $\xi_{init}$ . Localization is achieved by differentiable template-matching in the latent space. To compensate imperfection of the local BEV reconstruction, the localization module extracts a neural representation from both the map tile and local BEV and builds a probability map  $\mathbf{M}_{prob} \in \mathbb{R}^{L \times L}$  over the map tile following [39].

Matching the Neural BEV prediction and the Neural map in the feature space robustifies the localization module against errors and imperfections in the local BEV that may be due to occlusion, or perceptual degradation (e.g., lack of illumination, or bad weather) at the cost of resolution. To extract those features, we use the EfficientNet architecture as in the U-BEV model.

We apply a softmax with  $\tilde{\mathbf{M}} = \text{softmax}_{2D}(\mathbf{M}_{prob})$  to use as our probability distribution. We extract a prediction by performing a soft-argmax [40] in both  $x$  and  $y$  direction with

$$\Phi(\tilde{\mathbf{M}}) = \sum_{i,j}^S (\mathbf{W} \otimes \tilde{\mathbf{M}})_{ij} \quad \text{with} \quad \begin{matrix} \mathbf{W} \in \mathbb{R}^{S \times S \times 2} \\ \mathbf{W}_{ij} = (i, j) \end{matrix} . \quad (3)$$

## IV. EXPERIMENTS

We evaluate the proposed U-BEV algorithm individually and the full relocalization system including deep template matching on the public nuScenes driving dataset [36] with its map expansion. In our evaluation, we use surround camera data, ground truth poses, and ground truth SD map information. We use three main categories of the map data, i.e. drivable surface, crosswalks, and walkways. Since one goal is to investigate degenerate cases like feature-poor scenarios, we avoid the use of distinct objects, such as traffic lights in the localization scenario to force the network to learn about road network shape over distinct landmarks.

### A. Dataset

The nuScenes dataset is the only publicly available one that provides 6-views surrounding images and SD maps aligned with the recorded data. It consists of 33989 data samples from 2 different cities. The standard data split of nuScenes separates the sequences temporally (i.e. samples recorded next to each other are in the same split). This is

well suited for most tasks that deal with dynamic objects as object detection, but suboptimal when static objects are involved as in the relocalization scenario. This creates data contamination: samples from different sequences on the same road are distributed into training and validation sets while observing the same part of the scene. We notice when training on this split that the models can learn unreasonably well and can predict parts of the maps that are not observable due to occlusion. Therefore, we use the split proposed by [41]. This split separates samples both temporally and spatially and avoids overlaps in training and validation sets. We thus have 28008 data samples for training and 5981 for testing.

### B. Experimental Setup

**BEV:** For evaluating the U-BEV performance, we build BEV segmentation from 6 surround images of  $100m \times 100m$  at a resolution of 50cm/px. The sets of 6 images of resolution  $1600 \times 900$  are first cropped on the top by 27% and down-sampled to a resolution of  $544 \times 224$ . We compare U-BEV with Cross-View-Transformer (CVT) [10] a state-of-the-art BEV network which we retrained in our split for a baseline comparison under similar conditions as U-BEV. Retraining heavy BEV models with large GPU memory footprints is beyond the scope of our comparison as we aim to find a lightweight solution suitable for online relocalization.

**Relocalization:** For relocalization, we perturb the ground-truth location of each set of surround images uniformly within  $[0m, 100m]$  and crop tiles from the map data around the perturbed location of size  $300m \times 300m$ . Images and map tiles are given to the end-to-end localizer which estimates the position. We compare our localizer with the variant that includes CVT as BEV network and OrienterNet [15]. For the latter, we use their pretrained model and finetune it on a set of different map resolutions and BEV ranges on the nuScenes training set. Furthermore, we disable the rotational degree of freedom in OrienterNet and exchange the OpenStreetMap tiles with nuScenes map tiles in training and validation. We evaluate relocalization with Recall Accuracy (RA) at distance thresholds 1, 2, 5, 10m

### C. Training

Our training pipeline is divided into two steps. First, we pre-train the U-BEV module to predict the surrounding BEV. Only then, do we train the end-to-end model. All models are trained on a single NVIDIA GV100 GPU with 32GB of VRAM.

**BEV:** To have a fair comparison with CVT we train one model per class in our experiments, but we use a single multi-purpose model in the end-to-end pipeline. We use a weighted focal loss function [42] on the BEV prediction and a sparse Huber loss to supervise the height estimated using Equation 1.

For data augmentation, we rotate the camera poses around the car by a random amount. This does not affect the feature extraction and height prediction but helps the BEV decoder avoid overfitting by constantly changing the camera view directions. We train for 12 epochs, using 1-cycle learning




Fig. 4: Architecture of the U-BEV Model. (a) The pretrained backbone (in blue) extracts features from all 6 cameras around the car. The first decoder (in orange) predicts the height of each pixel on each input image. This height is used to project features from different cameras into a single BEV (in green). Deeper features get projected to lower-resolution BEVs and are then upsampled in an encoder-decoder fashion with skip connections (in yellow). (b) Illustrates the projection operation from surround view images and heights to different BEV layers.

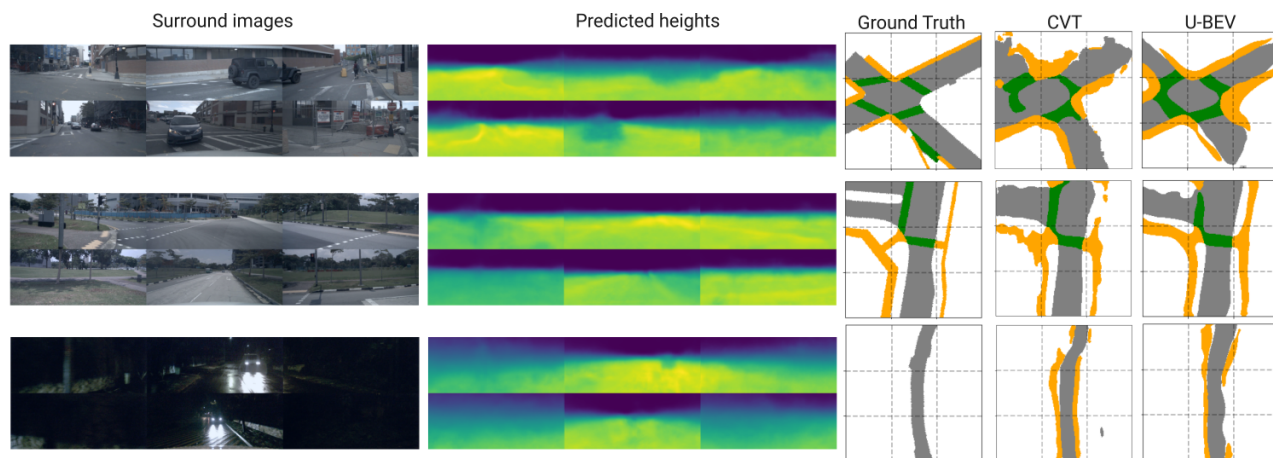


Fig. 5: Sample inputs and outputs of the proposed U-BEV, including surround images, predicted heights, and predicted and ground truth BEV. Compared to CVT, U-BEV more truthfully reconstructs drivable surfaces and sidewalks.

rate scheduler [43] with a maximum learning rate of  $1 \cdot 10^{-3}$  and AdamW as optimizer [44].

**Relocalization:** Before training the relocalization part, we freeze the U-BEV image encoder and its height decoder. This allows us to avoid overfitting to the limited nuScenes dataset. As before, we find benefits in augmenting the dataset by rotating both the local BEV (by modifying the extrinsic parameters) and the global map (by rotating the map) by random amounts during training until the last two epochs, where we disable this process.

We train this model similar to U-BEV, for 12 epochs and with the same LR and optimizer.

#### D. Results

**BEV:** Table II summarizes the BEV results of U-BEV and CVT. U-BEV outperforms CVT in all categories in terms of IoU. While the performance on drivable surfaces improves by 1.7 IoU, the improvement is much higher on

Walkways and Crossing with an improvement of 2.8 and 2.3 respectively. Notably, U-BEV achieves this performance boost at a reduced computational complexity, requiring 5% less GFLOPs during inference. Single-class inference consistently improves U-BEV by up to 2 IoU. These results indicate that the representational power of the proposed U-BEV is considerably larger than the Transformer-based CVT baseline. We attribute this to firstly, the multi-height encoding giving a boost of up to 1.7 IoU compared to the single-height case. Secondly, U-BEV retains full resolution through skip connections while CVT upsamples from a bottleneck of size  $25 \times 25$ . Sample results in Fig. 5 demonstrate outputs from our system, including height predictions, and BEVs from U-BEV and CVT. Notably, U-BEV shows superior performance in drivable surface and sidewalk segmentation.

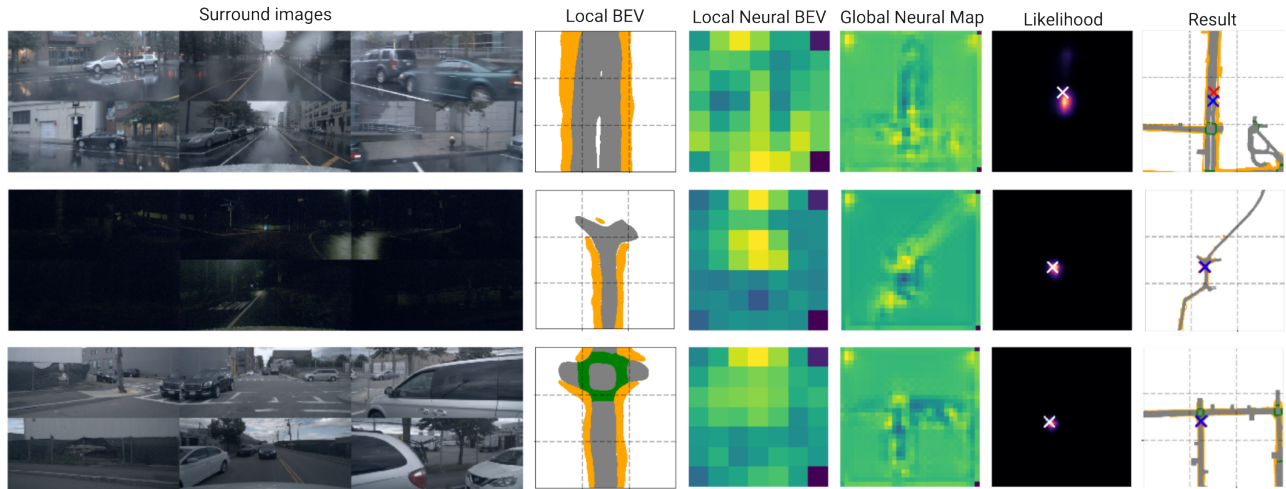


Fig. 6: Example of inputs and outputs of the localization pipeline, including surround images, local BEV, neural encoding of local BEV and map tile, likelihood of the prediction, and visualization of the result. In the visualization, the blue cross is the ground truth pose while the red cross is the prediction.

Type	BEV (m)	map (m <sup>2</sup> )	cm/px	Layers	R@1m	R@2m	R@5m	R@10m
OrienterNet	60	300×300	75	3	17.29	38.22	57.21	65.11
OrienterNet	60	300×300	30	3	<b>31.23</b>	<b>46.90</b>	58.89	66.04
OrienterNet only drivable	60	300×300	30	1	<u>24.08</u>	38.42	54.63	62.20
CVT Neural Relocalization	60	300×300	30	3	7.691	23.61	58.44	77.03
U-BEV only drivable (Ours)	60	300×300	30	1	13.69	37.67	<u>69.57</u>	<u>82.78</u>
U-BEV Neural Relocalization (Ours)	60	300×300	30	3	16.89	<u>41.60</u>	<b>71.33</b>	<b>83.46</b>

TABLE I: Localization results in terms of Recall Accuracy at 1m, 2m, 5m, 10m.

Model	GFLOPs ↓	Drivable ↑	Walkways ↑	Crossings ↑
CVT	18.7	65.3	25.4	26.1
CVT mc	18.7	62.0	29.7	27.4
U-BEV, w/o H	<b>17.16</b>	66.23	30.8	<b>28.0</b>
U-BEV mc light	<u>17.81</u>	66.08	30.96	26.63
U-BEV mc, w/o H	<b>17.16</b>	62.69	27.93	24.5
U-BEV mc	21.46	<u>66.8</u>	<u>32.0</u>	27.7
U-BEV	<u>17.81</u>	<b>67.0</b>	<b>32.5</b>	<b>29.7</b>

TABLE II: IoU performance on the [45] split with 100 × 100m and 50cm per pixel. mc refers to multi-class models, w/o H refers to models without height.

**Relocalization:** Table I summarizes the relocalization results. OrienterNet performs best in the R@1m and R@2m bins which we attribute to the more sophisticated matching pipeline at full resolution. OrienterNet is good at picking up distinct features of the environment but suffers under semantically poor regions. U-BEV relocalization on the other side largely outperforms the baselines in the R@5m and R@10m bins which we attribute to its ability to reason about the shape of roads over the presence of distinct features. We argue that the good performance in the R@5m and R@10m bins is highly relevant for a lightweight relocalization system in diverse scenarios, including feature-poor regions. Task-specific local perception systems can bridge the remaining error to a desired accuracy as elaborated in Section I. CVT-based relocalization with the proposed template matcher has the lowest performance in the lower bins but improves on the R@10m bin. U-BEV increases performance by 21.1%

R@5m and 26.4% R@10m over the best baselines. We also perform experiments with only drivable surfaces in the map data. Here, U-BEV suffers only a minor performance decrease, while the baseline degrades especially in the R@1m and R@2m bins. We attribute this to the ability of U-BEV to rely more on the scene structure than on distinct semantic landmarks as the baseline.

## V. CONCLUSIONS

This paper proposes a novel U-Net inspired BEV architecture “U-BEV” that leverages intermediate Segmentation on multiple heights. This architecture outperforms contemporary transformer-based BEV architectures in Segmentation performance by 1.7 to 2.8 IoU. In addition, we propose a novel relocalization approach that leverages the proposed U-BEV for localization against neurally encoded SD map data together with Deep Template Matching. The relocalization extension largely outperforms related methods by over 26% in Recall Accuracy at 10m. Notably, the need for only a few classes in map data, especially the road surface, paves the way to relocalization in feature-poor environments.

## REFERENCES

- [1] S. Lowry, N. Sünderhauf, P. Newman, *et al.*, “Visual place recognition: A survey,” *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 1–19, 2015.

- [2] S. Garg, T. Fischer, and M. Milford, “Where is your place, visual place recognition?” *International Joint Conference on Artificial Intelligence (IJCAI2021)*, 2021.
- [3] E. Xie, Z. Yu, D. Zhou, *et al.*, “M<sup>2</sup>BEV: Multi-camera joint 3D detection and segmentation with unified birds-eye view representation,” in *ArXiv Preprint*, 2022. arXiv: 2204.05088.
- [4] Z. Liu, H. Tang, A. Amini, *et al.*, “Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [5] Y. Zhang, Z. Zhu, and D. Du, “Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction,” *arXiv preprint arXiv:2304.05316*, 2023.
- [6] Y. Wei, L. Zhao, W. Zheng, Z. Zhu, J. Zhou, and J. Lu, “Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving,” *arXiv preprint arXiv:2303.09551*, 2023.
- [7] Y. Zhou, X. Li, S. Li, and X. Wang, “Visual mapping and localization system based on compact instance-level road markings with spatial uncertainty,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 802–10 809, 2022.
- [8] Z. Zhang, M. Xu, W. Zhou, T. Peng, L. Li, and S. Poslad, “Bev-locator: An end-to-end visual semantic localization network using multi-view images,” *arXiv preprint arXiv:2211.14927*, 2022.
- [9] X. Zhu, X. Cao, Z. Dong, *et al.*, “Nemo: Neural map growing system for spatiotemporal fusion in bird’s-eye-view and bdd-map benchmark,” *arXiv preprint arXiv:2306.04540*, 2023.
- [10] B. Zhou and P. Krähenbühl, “Cross-view transformers for real-time map-view semantic segmentation,” in *CVPR*, 2022.
- [11] J. Phillion and S. Fidler, “Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d,” in *Proceedings of the European Conference on Computer Vision*, 2020.
- [12] T. Roddick and R. Cipolla, “Predicting semantic map representations from images using pyramid occupancy networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 138–11 147.
- [13] Y. B. Can, A. Liniger, O. Unal, D. Paudel, and L. Van Gool, “Understanding bird’s-eye view of road semantics using an onboard camera,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3302–3309, 2022.
- [14] T. Roddick and R. Cipolla, “Predicting semantic map representations from images using pyramid occupancy networks,” *CoRR*, vol. abs/2003.13402, 2020.
- [15] P.-E. Sarlin, D. DeTone, T.-Y. Yang, *et al.*, “OrionNet: Visual Localization in 2D Public Maps with Neural Matching,” in *CVPR*, 2023.
- [16] P.-E. Sarlin, E. Trulls, M. Pollefeys, J. Hosang, and S. Lynen, “Snap: Self-supervised neural maps for visual positioning and semantic understanding,” *arXiv preprint arXiv:2306.05407*, 2023.
- [17] S. Vora, A. H. Lang, B. Helou, and O. Beijbom, “Pointpainting: Sequential fusion for 3d object detection,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4604–4612.
- [18] T. Yin, X. Zhou, and P. Krähenbühl, “Multimodal virtual point 3d detection,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 16 494–16 507, 2021.
- [19] B. Pan, J. Sun, H. Y. T. Leung, A. Andonian, and B. Zhou, “Cross-view semantic segmentation for sensing surroundings,” *IEEE Robotics and Automation Letters*, vol. 5, pp. 4867–4873, 2019.
- [20] A. Hu, Z. Murez, N. C. Mohan, *et al.*, “Fiery: Future instance prediction in bird’s-eye view from surround monocular cameras,” *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 15 253–15 262, 2021.
- [21] A. Saha, O. A. M. Maldonado, C. Russell, and R. Bowden, “Enabling spatio-temporal aggregation in birds-eye-view vehicle estimation,” *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5133–5139, 2021.
- [22] N. B. Gosala and A. Valada, “Bird’s-eye-view panoptic segmentation using monocular frontal view images,” *IEEE Robotics and Automation Letters*, vol. PP, pp. 1–1, 2021.
- [23] J. Huang, G. Huang, Z. Zhu, and D. Du, “Bevdet: High-performance multi-camera 3d object detection in bird-eye-view,” *ArXiv*, vol. abs/2112.11790, 2021.
- [24] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, “Netvlad: Cnn architecture for weakly supervised place recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 5297–5307.
- [25] H. Wang, C. Xue, Y. Tang, W. Li, F. Wen, and H. Zhang, “Ltsr: Long-term semantic relocalization based on hd map for autonomous vehicles,” in *2022 International Conference on Robotics and Automation (ICRA)*, IEEE, 2022, pp. 2171–2178.
- [26] A. Gawel, C. Del Don, R. Siegwart, J. Nieto, and C. Cadena, “X-view: Graph-based semantic multi-view localization,” *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1687–1694, 2018.
- [27] A. Toker, Q. Zhou, M. Maximov, and L. Leal-Taixé, “Coming down to earth: Satellite-to-street view synthesis for geo-localization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6488–6497.
- [28] X. Lu, Z. Li, Z. Cui, M. R. Oswald, M. Pollefeys, and R. Qin, “Geometry-aware satellite-to-ground image synthesis for urban areas,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 859–867.

- [29] F. Fervers, S. Bullinger, C. Bodensteiner, M. Arens, and R. Stiefelhagen, "Uncertainty-aware vision-based metric cross-view geolocalization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 621–21 631.
- [30] S. Wang, Y. Zhang, A. Vora, A. Perincherri, and H. Li, "Satellite image based cross-view localization for autonomous vehicle," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2023, pp. 3592–3599.
- [31] Z. Xia, O. Booiij, M. Manfredi, and J. F. Kooij, "Visual cross-view metric localization with dense uncertainty estimates," in *European Conference on Computer Vision*, Springer, 2022, pp. 90–106.
- [32] T. Lentsch, Z. Xia, H. Caesar, and J. F. Kooij, "Slice-match: Geometry-guided aggregation for cross-view pose estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17 225–17 234.
- [33] J. Zhao, Q. Zhai, P. Zhao, R. Huang, and H. Cheng, "Co-visual pattern-augmented generative transformer learning for automobile geo-localization," *Remote Sensing*, vol. 15, no. 9, p. 2221, 2023.
- [34] J. Lin, Z. Zheng, Z. Zhong, *et al.*, "Joint representation learning and keypoint detection for cross-view geolocalization," *IEEE Transactions on Image Processing*, vol. 31, pp. 3780–3792, 2022.
- [35] H. Yang, X. Lu, and Y. Zhu, "Cross-view geolocalization with layer-to-layer transformer," *Advances in Neural Information Processing Systems*, vol. 34, pp. 29 009–29 020, 2021.
- [36] H. Caesar, V. Bankiti, A. H. Lang, *et al.*, "Nuscenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [37] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," *ICML 2019*, vol. abs/1905.11946, 2019.
- [38] H. Li, P. Xiong, J. An, and L. Wang, "Pyramid attention network for semantic segmentation," *ArXiv*, vol. abs/1805.10180, 2018.
- [39] J. Cheng, Y. Wu, W. Abd-Almageed, and P. Natarajan, "Qatm: Quality-aware template matching for deep learning," *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11 545–11 554, 2019.
- [40] D. C. Luvizon, H. Tabia, and D. Picard, "Human pose regression by combining indirect part detection and contextual information," *Computers & Graphics*, vol. 85, pp. 15–22, 2019.
- [41] T. Roddick and R. Cipolla, "Predicting semantic map representations from images using pyramid occupancy networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 138–11 147.
- [42] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 2980–2988.
- [43] L. N. Smith and N. Topin, "Super-convergence: Very fast training of neural networks using large learning rates," in *Artificial Intelligence and Machine Learning for Multi-Domain Operations applications*, SPIE, vol. 11006, 2019, pp. 369–386.
- [44] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2017.
- [45] T. Roddick and R. Cipolla, "Predicting semantic map representations from images using pyramid occupancy networks," *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11 135–11 144, 2020.