

NYU-VPR: Long-Term Visual Place Recognition Benchmark with View Direction and Data Anonymization Influences

Diwei Sheng*, Yuxiang Chai*, Xinru Li, Chen Feng†, Jianzhe Lin, Claudio Silva, John-Ross Rizzo
<https://ai4ce.github.io/NYU-VPR/>

Abstract—Visual place recognition (VPR) is critical in not only localization and mapping for autonomous driving vehicles, but also assistive navigation for the visually impaired population. To enable a long-term VPR system on a large scale, several challenges need to be addressed. First, different applications could require different image view directions, such as front views for self-driving cars while side views for the low vision people. Second, VPR in metropolitan scenes can often cause privacy concerns due to the imaging of pedestrian and vehicle identity information, calling for the need for data anonymization before VPR queries and database construction. Both factors could lead to VPR performance variations that are not well understood yet. To study their influences, we present the NYU-VPR dataset that contains more than 200,000 images over a 2km×2km area near the New York University campus, taken within the whole year of 2016. We present benchmark results on several popular VPR algorithms showing that side views are significantly more challenging for current VPR methods while the influence of data anonymization is almost negligible, together with our hypothetical explanations and in-depth analysis.

I. INTRODUCTION

Visual place recognition (VPR) is the process of retrieving the most similar images for a query one from a database of images with known camera poses, which is often used for loop closing in mapping, localization, and navigation. It relies on representing an image as a global feature vector which describes the portion of the image appearance that is most relevant to its capturing pose. Its applications range from autonomous driving for vehicles, to assistive navigation for the visually impaired people, especially in busy and crowded metropolitan areas where GPS could suffer from the “urban canyon” problem when satellite signals are blocked or multi-reflected to cause large localization errors.

A reliable large-scale and long-term VPR system has to address several challenges. The first one is to choose a proper image view direction, *assuming a non-omnidirectional camera which is commonly available on smartphones unlike panoramic cameras*. In a self-driving car scenario, using front-view images from a dash-mounted camera whose view direction is parallel to the driving/street direction is almost a default choice. In fact, most VPR methods have been investigated and evaluated under such front-view conditions, where features of roads, shapes of skylines, and textures of the roadside buildings can all contribute to describing and discriminating various image locations.

New York University, Brooklyn, NY 11201, USA

* indicates equal contributions.

†Chen Feng is the corresponding author. cfeng@nyu.edu




Fig. 1. Top-1 VPR retrieval accuracy of 4 baseline algorithms under different view directions and data anonymization.

Yet not all downstream applications prefer front-view images. For example, a person with low vision might need image-based wearable navigation assistance to find the entrance of a particular shop on the street. The aforementioned front-view images typically contain more than half of the pixels on the road and the sky while the remaining pixels on street sides that are far away from the image capturing locations. Contrarily, the side view offers fronto-parallel images of buildings along streets, which stores enough information required for such an application scenario.

However, currently, there is a lack of datasets that could evaluate VPR methods specifically on side-view images in comparison with front-view ones. Many datasets contain no side-view images or mix them with front-view images without explicit labels (see Table I). Moreover, as far as we know, there is no systematic comparison of VPR performance between images from the two view directions. Thus, the following questions remain unclear: *are side-view images more challenging for VPR than front-view ones? And if so, why? And how much is the performance difference?*

Another challenge for large-scale VPR systems is data privacy, receiving an increasing attention from the community [12]. Building such systems in large metropolitan areas requires collecting images for a long term, inevitably creating concerns of both violating the privacy of the identity information of individual pedestrians and vehicles, and potentially even leaking their spatial-temporal trajectories. Unlike the privacy-preserving technique in [12] that still operates on raw images, another way could be directly anonymizing the

TABLE I
COMPARISON OF MAJOR PUBLIC OUTDOOR VPR DATASETS WITH NYU-VPR.

Dataset	side-view	side-view-label	dynamic-object	crowded-area	anonymization	seasonal-changes	#images
StreetLearn [1]	✓	-	✓	✓	✗	✗	143,000
StreetView [2]	✓	-	✓	✓	✗	✗	62,058
Nordland [3]	✗	-	✗	✗	✗	✓	28,865
VPRiCE 2015 [4]	✗	-	✓	✗	✗	✗	7,778
Tokyo 24/7 [5]	✓	✗	✓	✓	face-only	✗	76,000
Pittsburgh [6]	✓	✗	✓	✗	✗	✓	254,064
KITTI raw [7]	✗	-	✓	✗	✗	✗	12,919
KAIST [8]	✗	-	✓	✗	✗	✗	105,000
Oxford RobotCar [9]	✗	-	✓	✗	✗	✓	19,556,490
Mapillary [10]	✓	✓	✓	✓	✗	✓	1,681,000
NCLT [11]	✗	-	✓	✗	✗	✓	100,000
NYU-VPR(ours)	✓	✓	✓	✓	✓	✓	201,790

images by wiping out all the identity-related pixels (see Figure 3). However, this again brings some unanswered questions for VPR: *would such data anonymization significantly affects the performance of existing VPR algorithms? If so, does it increase or decrease the VPR accuracy/robustness?*

This paper aims at filling the gaps by a new VPR dataset and benchmark. This is a year-long dataset captured outdoors by vehicles with front-view and side-view cameras traveling in a metropolitan area recording more than 200,000 GPS-tagged images. This allows us to answer the above questions by comparing the results under various conditions.

The contributions of this paper are as follows:

- We present NYU-VPR, a unique large-scale, year-long, outdoor VPR benchmark dataset containing both front-view and side-view GPS-tagged images taken at different lighting conditions with seasonal and appearance changes in a busy and crowded urban area of New York City. This dataset and our benchmark code will be released for educational and research purposes.
- We benchmark the performance of several popular VPR algorithms with a focus on the influence of image view directions. As far as we know, this is the first work to systematically demonstrate and analyze the causes of the significant challenge of VPR with side-view images.
- We anonymize the identity information in this dataset by removing pixels of both pedestrians and vehicles to address the privacy concerns of large-scale VPR in urban scenes. This is also the first result to show that all the benchmarked VPR algorithms are only marginally affected by this anonymization.

II. RELATED WORK

Because NYU-VPR contains only outdoor images used for visual place recognition, we review publicly available datasets that have similar characteristics. The main differences between those datasets and our proposed dataset are summarized in Table I.

Side-view and side-view label: In recent years there has been substantial growth in the number of visual place recognition datasets in the urban areas. However, most datasets only contain front-view images, gathered by cameras on the front and back of cars [3, 4, 7–9, 11]. The side-view images featuring the storefront and sidewalk are not included. Few datasets include the side-view images in addition to the

front-view images. For example, the images in Tokyo 24/7 dataset were gathered by pedestrians’ phones and featured both front-view and side-view images [5]. The images in Pittsburgh dataset were perspective images generated from Google Street View panoramas of the Pittsburgh area [6]. But those datasets do not label the images as front-view or side-view. Thus no work can use those datasets to compare the visual place recognition results on the side-view images versus those on the front-view images. NYU-VPR contains images labeled as front-view and images labeled as side-view. We focus on evaluating the visual place recognition algorithms in both categories. We compare the results of algorithms on the side-view images versus the results on the front-view images in order to analyze the influence of view direction on the long-term visual place recognition.

Dynamic objects, crowded-area and anonymization: Dynamic objects such as pedestrians and vehicles in images may affect the performance of visual place recognition due to the changing appearance at the same place or the existence of obstructing the street view. Besides, the presence of pedestrians and vehicles in publicly available datasets may raise privacy issues if images are not anonymized. There are few appearances of dynamic objects in datasets of images gathered in suburban areas. For example, Nordland is a dataset of images taken on the train on a railway line between the cities of Trondheim and Bodø [3]. In contrast, dynamic objects appears much more frequently in the datasets featuring urban areas [4–10]. We define crowded areas as metropolitan areas such as New York City and Tokyo that has a high population density and is crowded with pedestrians and vehicles. In Table I, images from Tokyo 24/7, Mapillary, and NYU-VPR are gathered in crowded areas. Anonymization is needed on those datasets for privacy protection. Tokyo 24/7 only applied face redaction on pedestrians [5], while Mapillary did not apply any anonymization [10]. We use MSeg [13] to replace pedestrians and vehicles with white pixels.

Seasonal changes: Matching images that are taken at the same location in different seasons is crucial for long-term visual place recognition. This is because objects on images change with the seasons: new storefront, trees withering, constructions finished, etc. In Table I, Pittsburgh [6] includes images in different seasons but few image locations are visited in all four seasons. Nordland, Oxford RobotCar [9],

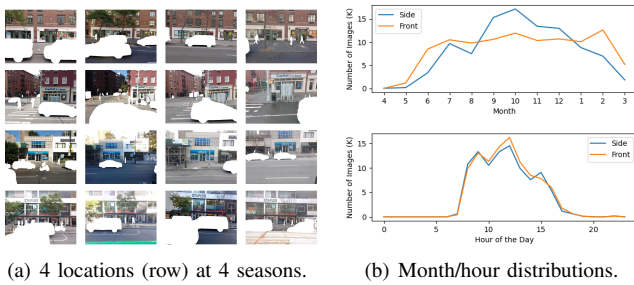


Fig. 2. Dataset visualization of NYU-VPR w.r.t. the image capturing time.

NCLT [11], and Mapillary [10] all include images in four seasons for most locations. NYU-VPR is similar but temporally denser: most locations are visited more than once every month, as shown by the time (Fig. 2(b)) and space (Fig. 4(f) and Fig. 4(e)) distributions of images. So NYU-VPR can be used to evaluate long-term visual place recognition for the influence of seasonal changes.

Baseline methods: Current VPR methods can be roughly grouped into three categories: deep-learning methods, non-deep-learning methods, and methods that only use deep learning descriptors. We select methods in each category. Deep learning methods [14–17] use a convolutional neural networks (CNN) and train CNN in an end-to-end manner directly. We select NetVLAD [14] and PoseNet [15] in this group. For non-deep-learning methods [5, 18–20], two classical ones are bag-of-words (BoW) model and Vector of Locally Aggregated Descriptors (VLAD). In this group, we choose DBoW+ORB [18, 21], which was used in the popular ORB-SLAM for loop closing [22], and VLAD+SURF [19, 23], which was used in [24]. Methods that only use deep learning descriptors take advantage of deep nets’ ability to detect a richer set of key points, such as SuperPoint [25] which we also adopted.

III. THE NYU-VPR DATASET

Our dataset is named NYU-VPR. It contains images recorded in Manhattan, New York from May 2016 to March 2017. The images were recorded with GPS tags by smartphone cameras installed on the front, back, and side parts of (undisclosed) fleet cars with auto-exposure¹. The dataset contains both side-view images and front-view images. There are 100,500 side-view and 101,290 front-view raw images, each with a 640×480 resolution. On the basis of raw images, we use MSeg [13], a semantic segmentation method, to replace moving objects such as people and cars with white pixels. Fig. 3 compares anonymized and raw images.



Fig. 3. Raw images vs. Anonymized images.

The images were recorded on streets around Washington Square Park. The trajectories of the locations where the images were recorded are shown in Fig. 4(d). Since the

¹The raw data was sampled from a larger dataset provided by Carmera.

cameras were placed on fleet cars, and their routes were random, the frequencies of locations where the images were taken are different. The frequencies of the locations where the side-view and front-view images were recorded are shown in Fig. 4(f) and Fig. 4(e) respectively.

Fig. 2(b) shows the time distribution. Since it contains images captured from May 2016 to March 2017, our dataset includes all four seasons. Therefore, it contains various changes of weather, illumination, vegetation, and road construction. As shown in Fig. 2(a), we can see image changes at the same location as the season changes.

Difficulty Level: We assign each side-view query image a difficulty level of easy, medium, or hard. First, we extract SIFT [26] features for each image. Then for each query image, we find the top-8 closest side-view training images by GPS coordinates. The query image and its top-8 closest images form eight image pairs. We use RANSAC to compute a fundamental matrix and the number of inliers for each pair of images. We use three intervals to measure the difficulty level of matching each pair based on each pair’s number of inliers points: 0-19 (hard), 20-80 (medium), >80 (easy). The interval values are determined by artificially viewing the image pairs and checking the similarity of the image pairs. The difficulty level of each side-view query image is the most common difficulty level of its eight pairs.

Uniqueness: Our dataset is unique in two ways. First, comparing to front-view images where sky and road surfaces occupy large areas, side-view images focus on street views such as shop signs and metro entrances. Second, we include image anonymization to protect the privacy of pedestrians and cars. In the meantime, anonymized images provide VPR algorithms static and environment-only information, getting rid of moving objects and pedestrians.

Front-view vs. Side-view: We hypothesize that side-view images are more challenging than front-view images for existing VPR methods, based on several theoretical reasons and observations. First, as illustrated in Fig. 5(a), since the images were taken on a moving vehicle at a low frequency, the overlap of two sequential side-view images would only occupy a limited part of the whole image. In contrast, the overlap of two sequential front-view images occupies more space. Furthermore, if the road is narrow and stores are very close to the vehicle, the spatial area covered in the side-view image is less than that covered in the front-view image at the same location. Besides, trees may impact the success rate for side-view images because trees will block features such as banners and entrances. Most trees look the same and thus may cause the uncertainty of localization.

Another important reason is that images are taken on a moving vehicle, suffering from motion blur. When the spatial area coverage of side-view images is less than front-view images, the motion blur effect is more serious on side-view images. We set a threshold for the variance of Laplacian to detect the blurriness of an image². The result in Fig. 5(b)

²<https://www.pyimagesearch.com/2015/09/07/blur-detection-with-opencv/>

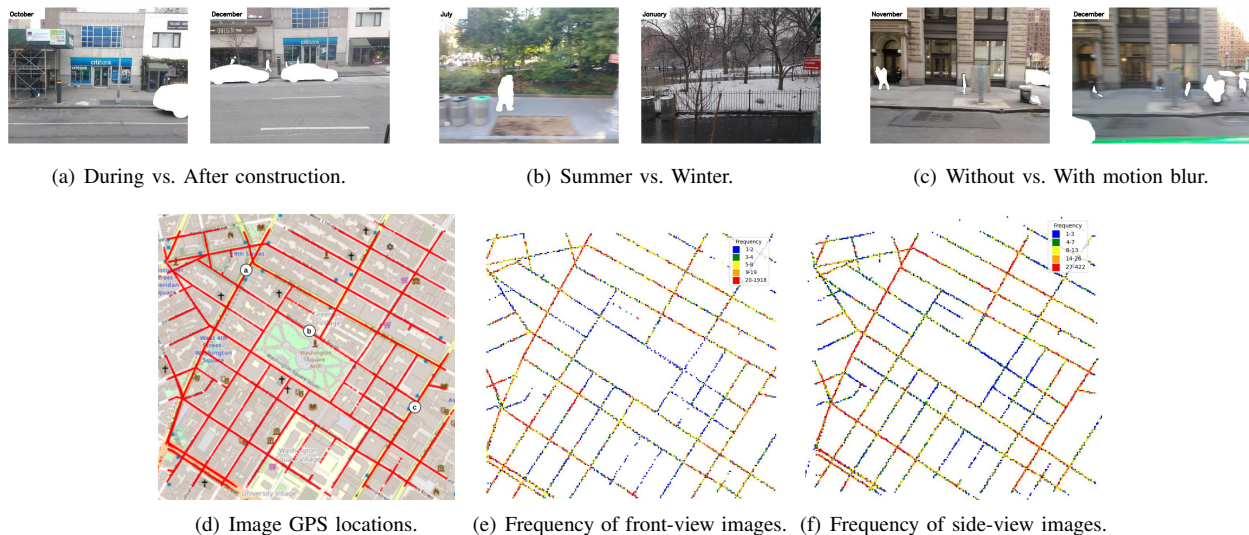


Fig. 4. Dataset visualization of NYU-VPR with respect to the image capturing location. The locations of (a)-(c) are highlighted in (d).

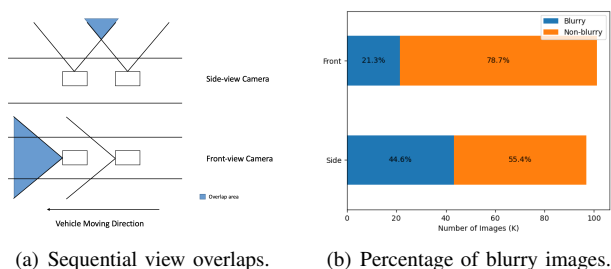


Fig. 5. Further explanation and illustration of our dataset.

right shows that side-view images have a larger percentage of blurry images than front-view images. We think blurry images are more challenging for VPR and we will analyze the relationship further in the experiment section.

Other Challenges: In addition to the side-view images, there are several other challenges in our dataset. Because our dataset is one-year long, the images taken at the same location have artificial or natural differences. First, Fig. 4(a) left was taken in October 2016 with sideways constructions and the right was taken in December 2016 after the construction. At some locations, the construction may cover the whole image. Second, different seasons cause different appearances at the same location. Fig. 4(b) left was taken in summer, July 2016, and the right was taken in winter, January 2017. In this case, the vegetation in Washington Square Park had changed a lot and snow was covering the ground in winter. Furthermore, if the vehicle was moving fast, the images taken by the vehicle will be blurry (Fig. 4(c)). Although two images were taken at the same location, the blurry one will cause more difficulty during VPR.

IV. BENCHMARK EXPERIMENTS

A. Settings

We selected five classical as well as state-of-the-art descriptors and methods for evaluation of visual place recognition performance on the NYU-VPR dataset.

Dataset: We randomly split both front-view and side-view images into training, validation and testing sections by 80%, 5%, and 15% respectively. For each view direction, both anonymized and raw images share the same split result. All images are resized to 640×480 . We also use Python module `utm` to convert GPS coordinates to UTM coordinates for more precise distance calculation between two locations.

VLAD+SURF: We use Vector of Locally Aggregated Descriptors (VLAD) [19] to aggregate speeded up robust features (SURF) [23] descriptors for image retrieval. Through experiments, we find the optimal cluster number is 32 within 8, 16, 32, and 64, by using MiniBatchKmeans with batch size at 5000. This cluster number gives high accuracy and acceptable training time. The training of 77608 images took about 8 hours on CPU with 64 GB available memory.

VLAD+SuperPoint: We use SuperPoint model pre-trained on MS-COCO generic image dataset [25]. We use nVidia RTX 2080S to extract SuperPoint features. Then we use VLAD to aggregate SuperPoint descriptors for image retrieval. We set the cluster number at 32, just as we do in VLAD+SURF, by using MiniBatchKmeans with batch size at 100. Notice the dimension of SuperPoint descriptors is larger than the dimension of SURF descriptors. The training of 77608 images took about 20 hours on CPU and GPU with 64 GB available memory.

NetVLAD: We directly use the pre-trained model weight for 30 epochs on Pittsburgh-250k datasets [6] to complete our testing. For the hardware, the CPU we adapt is Intel® Core™ i7-8700k, and the GPU we use is NVIDIA GEFORCE GTX 1080 TI. We first complete an initial clustering on training data to find out the centroids used for the testing process. The input testing data with the extracted deep feature are assigned to different clusters afterward. The batch size during testing is 24.

PoseNet: We use PoseNet model with ResNet34 as the base architecture [15]. For training, PoseNet requires the Cartesian coordinates of images as input besides images

themselves. So we gather latitude and longitude information of training images from the camera. We convert latitude and longitude to universal transverse mercator (UTM) coordinates to improve the accuracy of PoseNet’s estimation of images’ relative position. We use images with normalized UTM coordinates as the input to PoseNet. The GPU we use is NVIDIA GEFORCE RTX 2080S. The batch size during training is 32. For testing, PoseNet outputs the estimation of normalized UTM coordinates of the query images, which are used for evaluation.

DBoW: We use Distributed Bag of Words (DBoW) model³. We choose Oriented FAST and Rotated BRIEF (ORB) [21] descriptors for representing features. We use DBoW to generate a vocabulary constructed by ORB descriptors of training and test images. For testing, We generate the top-5 retrieval images by using DBoW3 to generate a score between each training and test images and selecting the top-5 scores for each test image. We run the testing process with multi-thread for efficiency.

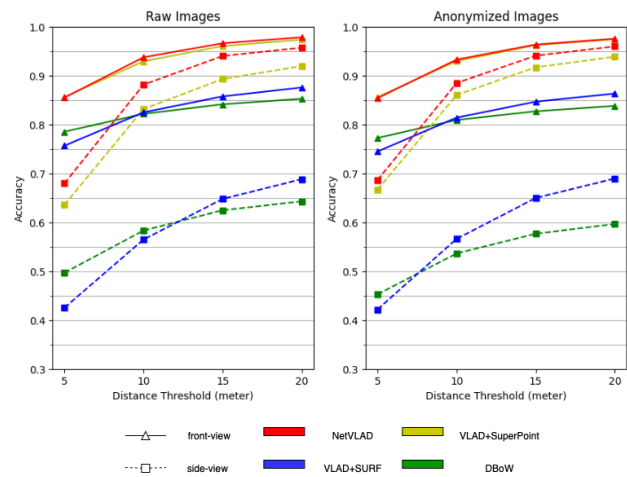
Evaluation: Following [24], we measure both top-1 and top-5 retrieval accuracy under four distance thresholds (5, 10, 15, 20 meters). If any of the top-k retrieval images are within the range of the distance from the query image, we count it as a successful retrieval. The top-k ranking is based on the similarity between image features calculated by VPR algorithms. In essence, this evaluation metric is similar to the more common precision-recall curve.

B. Results


Figure 6(a) shows our main results focused on the performance of non-/anonymized front-/side-view datasets, showing the influence of anonymization and view directions.

Performance: Fig. 1 and Fig. 6(a) shows the result of our experiments. Obviously, the result of the top 5 retrieval image accuracy is higher than the top 1 retrieval image accuracy, with an average 10% higher. And when using VLAD to aggregate descriptors, the result of SuperPoint descriptors is much more accurate than the result of SURF descriptors. In both two figures, the accuracy of NetVLAD is the highest, followed by VLAD+SuperPoint and VLAD+SURF. The last is the DBoW method. We may attribute the low accuracy of DBoW to the unsustainability of ORB features. PoseNet outputs a GPS coordinate and using that coordinate we can find the closest top 1 retrieval image, and through experiments, the accuracy of PoseNet is 15.3% and 37.5% when the distance threshold is 5 and 10 meters respectively. *Due to the low performance of PoseNet, we omit it in other experiments and do not plot the results.* We also calculate the accuracy in terms of difficulty level as mentioned before. Fig. 6(b) shows the result of top 5 retrieval in different difficulty level. The visual result is shown in Fig 7. Clearly, we can see VLAD+SuperPoint has better performance than VLAD+SURF and DBoW.

Anonymization: From Fig. 1, we can draw the conclusion that the anonymization does not have a large impact on the



(a) Top-5 retrieval results.



(b) Top-5 retrieval results in terms of difficulty level.

Fig. 6. Main results of the benchmark.

visual place recognition result, either of front-view dataset or side-view dataset. The anonymization, however, has little influence on the accuracy of some methods. For example, VLAD+SuperPoint gets 1.1% increase on average, while DBoW and VLAD+SURF have around 2.1% and 3.4% decrease on average respectively. Therefore, when doing experiments on visual place recognition, we can anonymize raw images to protect the privacy of people and cars.

View Direction: View direction does have a conspicuous influence on retrieval accuracy. In Fig. 1, the accuracies calculated from front-view images are higher than those from side-view images. This phenomenon happens in every method, indicating that front-view images is less challenging than side-view images in VPR, as we expected previously.

Next, we provide some in-depth analysis according to the reasons we hypothesized in Fig. 5(a). We first define that *a query image is successfully retrieved if any of the top 5 retrieval images is within 5 meters of the query image.*

Dynamic Objects: Fig. 8 shows the success rate of query images vs. different anonymization rates. We calculate the anonymization rate using the percentage of white pixels in the image. As the anonymization rate increases, the success rate for the front view increases while the success rate for the

³<https://github.com/rmsalinas/DBow3>

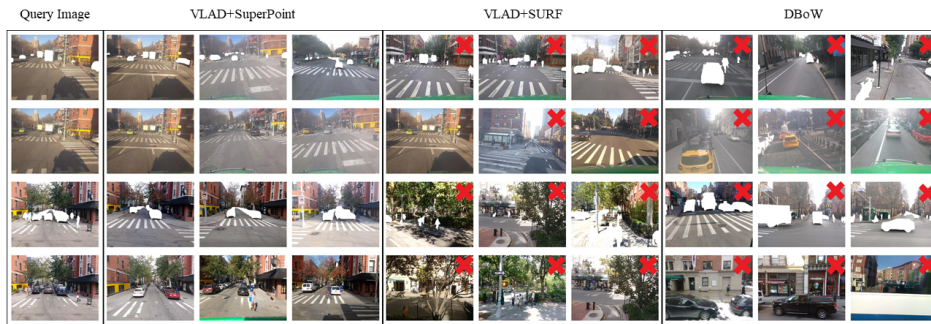


Fig. 7. We visualize VPR results only for VLAD+SuperPoint, VLAD+SURF, and DBoW, because the result of NetVLAD is very similar to VLAD+SuperPoint. We randomly picked one location (more locations results in the supplementary). The first row is for the anonymized front-view query. The second row is for the corresponding raw front-view query. The third row is for the side-view query taken at the same location as the front-view query. The last row is for the corresponding raw side-view query. The red cross means the location of the retrieval image is not in the distance threshold (10 meters). We show the top 3 retrievals for each method.

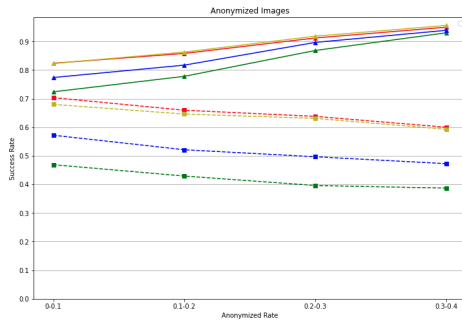


Fig. 8. VPR success rate vs. anonymized rate. The same legend as Fig. 6.

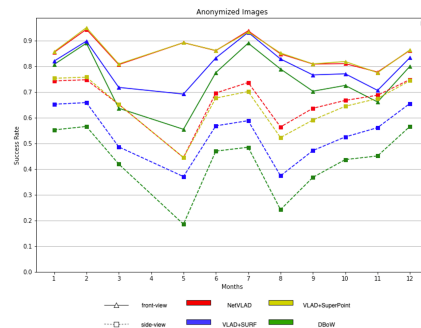


Fig. 10. VPR success rate vs. query image month.




Fig. 9. VPR success rate vs. image blurriness.

side view decreases. We propose a hypothesis for why this happens: in front view, dynamic objects such as cars and pedestrians are noisy signals for VPR. After anonymizing dynamic objects, VPR algorithms focus on the street features, which increases the success rate. However, in side view, those dynamic objects may block many street features. A higher anonymization rate means more street features blocked, which decreases the success rate.

Motion Blur: Fig. 9 shows the relationship among success rate, VPR method, view directions, and blurriness. It clearly shows that side-view images are more challenging than front-view images, and blurry images are more challenging than non-blurry ones. Moreover, given the same blurriness condition, side-view is still more challenging than front-view. We believe this is also due to the reason we hypothesized before: front-view contains more view overlaps among neighboring images, thus better VPR feature matching.

Seasonal Change: In addition, we analyzed the success rate of query images in different months as shown in Fig. 10. This result further confirmed in each month our observa-

tion that the side view is more challenging than the front view. The fluctuations between months also reflect some season/weather related influences on VPR performance.

V. CONCLUSIONS

After the large-scale experiments and analysis, we can finally answer our questions with confidence. Are side-view images more challenging for VPR than front-view ones? *Yes, and the performance drops of all VPR methods are significant*, although the dataset has no significant spatial/temporal differences on the distribution of images captured from the two view directions. Would our data anonymization significantly affect the performance of existing VPR algorithms? *No, and for some methods, the anonymization could even bring marginal improvements*, potentially due to the removal of those VPR noises. Our future work includes benchmarking more VPR methods and with geometric verification.

ACKNOWLEDGMENTS

We would like to thank Carmera for providing the raw NYC image data set that we used for creating the NYU-VPR. This research is funded by the Connected Cities for Smart Mobility towards Accessible and Resilient Transportation (C2SMART), a Tier 1 University Center awarded by U.S. Department of Transportation under contract 69A3351747124.

REFERENCES

- [1] P. Mirowski, A. Banki-Horvath, K. Anderson, D. Teplyashin, K. M. Hermann, M. Malinowski, M. K.

- Grimes, K. Simonyan, K. Kavukcuoglu, A. Zisserman et al., “The streetlearn environment and dataset,” *arXiv preprint arXiv:1903.01292*, 2019. 2
- [2] A. R. Zamir and M. Shah, “Image geo-localization based on multiplenearest neighbor feature matching usinggeneralized graphs,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 8, pp. 1546–1558, 2014. 2
- [3] N. Sünderhauf, P. Neubert, and P. Protzel, “Are we there yet? challenging seqslam on a 3000 km journey across all four seasons,” in *Proc. IEEE Int’l Conf. Robotics and Automation (ICRA)*, 2013, p. 2013. 2
- [4] “The VPRiCE Challenge 2015 – Visual Place Recognition in Changing Environments - Public - Confluence.” [Online]. Available: <https://roboticvision.atlassian.net/wiki/spaces/PUB/pages/14188617/The+VPRiCE+Challenge+2015+Visual+Place+Recognition+in+Changing+Environments> 2
- [5] A. Torii, R. Arandjelovic, J. Sivic, M. Okutomi, and T. Pajdla, “24/7 place recognition by view synthesis,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1808–1817. 2, 3
- [6] A. Torii, J. Sivic, T. Pajdla, and M. Okutomi, “Visual place recognition with repetitive structures,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2013. 2, 4
- [7] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The kitti dataset,” *Int’l J. Robotics Research*, 2013. 2
- [8] Y. Choi, N. Kim, K. Park, S. Hwang, J. S. Yoon, Y. In, and I. Kweon, “All-day visual place recognition: Benchmark dataset and baseline,” in *Workshop on Visual Place Recognition in Changing Environments*, 06 2015. 2
- [9] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, “1 year, 1000 km: The oxford robotcar dataset,” *Int’l J. Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017. 2
- [10] F. Warburg, S. Hauberg, M. Lopez-Antequera, P. Gargallo, Y. Kuang, and J. Civera, “Mapillary street-level sequences: A dataset for lifelong place recognition,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 2626–2635. 2, 3
- [11] N. Carlevaris-Bianco, A. K. Ushani, and R. M. Eustice, “University of michigan north campus long-term vision and lidar dataset,” *The International Journal of Robotics Research*, vol. 35, no. 9, pp. 1023–1035, 2016. 2, 3
- [12] P. Speciale, J. L. Schonberger, S. B. Kang, S. N. Sinha, and M. Pollefeys, “Privacy preserving image-based localization,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 5493–5503. 1
- [13] J. Lambert, Z. Liu, O. Sener, J. Hays, and V. Koltun, “MSeg: A composite dataset for multi-domain semantic segmentation,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 3
- [14] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, “Netvlad: Cnn architecture for weakly supervised place recognition,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 5297–5307. 3
- [15] A. Kendall, M. Grimes, and R. Cipolla, “Posenet: A convolutional network for real-time 6-dof camera relocalization,” in *Proc. IEEE Int’l Conf. Computer Vision (ICCV)*, 2015, pp. 2938–2946. 3, 4
- [16] M. Chancán, L. Hernandez-Nunez, A. Narendra, A. B. Barron, and M. Milford, “A hybrid compact neural architecture for visual place recognition,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 993–1000, April 2020.
- [17] Z. Chen, A. Jacobson, N. Sünderhauf, B. Upcroft, L. Liu, C. Shen, I. Reid, and M. Milford, “Deep learning features at scale for visual place recognition,” in *Proc. IEEE Int’l Conf. Robotics and Automation (ICRA)*. IEEE, 2017, pp. 3223–3230. 3
- [18] D. Gálvez-López and J. D. Tardós, “Bags of binary words for fast place recognition in image sequences,” *IEEE Trans. Robotics*, vol. 28, no. 5, pp. 1188–1197, October 2012. 3
- [19] H. Jégou, M. Douze, C. Schmid, and P. Pérez, “Aggregating local descriptors into a compact image representation,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2010, pp. 3304–3311. 3, 4
- [20] T. Sattler, B. Leibe, and L. Kobbelt, “Improving image-based localization by active correspondence search,” in *Proc. European Conf. Computer Vision (ECCV)*. Springer, 2012, pp. 752–765. 3
- [21] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, “Orb: An efficient alternative to sift or surf,” in *Proc. IEEE Int’l Conf. Computer Vision (ICCV)*, 2011, pp. 2564–2571. 3, 5
- [22] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, “Orbslam: A versatile and accurate monocular slam system,” *IEEE Trans. Robotics*, vol. 31, no. 5, pp. 1147–1163, 2015. 3
- [23] H. Bay, T. Tuytelaars, and L. Van Gool, “Surf: Speeded up robust features,” in *Proc. European Conf. Computer Vision (ECCV)*. Springer, 2006, pp. 404–417. 3, 4
- [24] X. Yu, S. Chaturvedi, C. Feng, Y. Taguchi, T.-Y. Lee, C. Fernandes, and S. Ramalingam, “VLASE: Vehicle localization by aggregating semantic edges,” *Proc. IEEE/RSJ Int’l Conf. Intelligent Robots and Systems (IROS)*, 2018. 3, 5
- [25] D. DeTone, T. Malisiewicz, and A. Rabinovich, “Superpoint: Self-supervised interest point detection and description,” in *CVPR Deep Learning for Visual SLAM Workshop*, 2018. [Online]. Available: <http://arxiv.org/abs/1712.07629> 3, 4
- [26] D. G. Lowe, “Object recognition from local scale-invariant features,” in *Proc. IEEE Int’l Conf. Computer Vision (ICCV)*, vol. 2, 1999, pp. 1150–1157 vol.2. 3