

TSCM: A Teacher-Student Model for Vision Place Recognition Using Cross-Metric Knowledge Distillation

Yehui Shen

Mingmin Liu

Huimin Lu

Xieyuanli Chen

Abstract—Visual place recognition (VPR) plays a pivotal role in autonomous exploration and navigation of mobile robots within complex outdoor environments. While cost-effective and easily deployed, camera sensors are sensitive to lighting and weather changes, and even slight image alterations can greatly affect VPR efficiency and precision. Existing methods overcome this by exploiting powerful yet large networks, leading to significant consumption of computational resources. In this paper, we propose a high-performance teacher and lightweight student distillation framework called TSCM. It exploits our devised cross-metric knowledge distillation to narrow the performance gap between the teacher and student models, maintaining superior performance while enabling minimal computational load during deployment. We conduct comprehensive evaluations on large-scale datasets, namely Pittsburgh30k and Pittsburgh250k. Experimental results demonstrate the superiority of our method over baseline models in terms of recognition accuracy and model parameter efficiency. Moreover, our ablation studies show that the proposed knowledge distillation technique surpasses other counterparts. The code of our method has been released at <https://github.com/nubot-nudt/TSCM>.

I. INTRODUCTION

Place recognition plays a significant role in the autonomous navigation of mobile robots, allowing them to locate the global place within previously mapped environments [1], [2]. Existing approaches can be categorized into visual place recognition (VPR) and LiDAR place recognition (LPR). Though LPR methods [3], [4], [5] have demonstrated good performance in outdoor environments, 3D LiDAR sensors are comparatively more costly and power-intensive. In contrast, cameras offer the advantages of cost-efficiency and abundant perceptual data, leading to growing research interest in outdoor place recognition [6], [7].

Current VPR methods [8] typically compress the query image into a global descriptor that is then used to determine the current location by comparing it with stored data in the database. However, visual sensors are sensitive to environmental changes, such as variations in lighting, weather conditions, and occlusions, which can negatively impact the robustness of these descriptors, resulting in poor VPR performance. Therefore, the main focus in VPR research is

Y. Shen is with Faculty of Robot Science and Engineering, Northeastern University, China. X. Chen and H. Lu are with the College of Intelligence Science and Technology, National University of Defense Technology, China. M. Liu is with SIASUN Robot& Automation Co., Ltd, China.

Corresponding author: Xieyuanli Chen (xieyuanli.chen@nudt.edu.cn)

This work has partially been funded by the National Science Foundation of China under Grant U1913202, U22A2059, U20A20197, and 62203460, Fund for key Laboratory of Space Flight Dynamics Technology (Num 2022-JYAPAF-F1028), and Young Elite Scientists Sponsorship Program by CAST (No. 2023QNRC001).

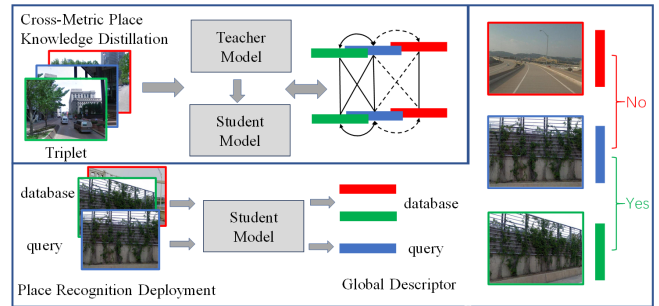


Fig. 1: Our model consists of cross-metric knowledge distillation and place recognition deployment. It uses cross-metric learning to transfer knowledge from teacher to student offline. During online inference, it uses the lightweight student model to generate descriptors from the input image and identify potential places by comparing them against those stored in the database.

to develop a robust network capable of generating stable descriptors. Unfortunately, such networks tend to be complex, with a high number of parameters, and they often operate at computational speeds lower than the sensor's frame rate, which limits their practical applicability [9].

To fill this gap, Hui et al. [10] employed knowledge distillation to enable knowledge transfer from the powerful teacher network to the lightweight student network. However, the performance of the student network fell significantly short of that of the teacher network. Cai et al. [11] introduced a self-learning approach to improving the student network in acquiring the teacher network's knowledge. While this approach enhanced the student network's performance, it did not reduce the size of the student network, making it unsuitable for mobile robots with limited onboard resources.

In this paper, we introduce a novel concept called cross-metric knowledge distillation (KD) and apply it to VPR within our teacher-student network. Our network comprises a teacher and a student model. The teacher model integrates components from ResNet [12], vision transformer [13], and inter-transformer [14], resulting in increased complexity and superior performance compared to state-of-the-art baselines in VPR. In contrast, the student network is designed to be lightweight, retaining only essential components. We transfer place knowledge from the powerful teacher to the lightweight student through our proposed cross-metric distillation. This approach empowers the lightweight student network to match or exceed the teacher's performance while reducing inference time, enabling fast and accurate VPR in real-world applications, as illustrated in Fig. 1. We evaluated our model on two large outdoor datasets, Pittsburgh30k and Pittsburgh250k.

The experimental results demonstrate that our model outperforms baseline models in terms of both recognition accuracy and parameter efficiency. Furthermore, we provide multiple ablation studies showing that our proposed distillation method surpasses other distillation techniques.

In summary, the contributions of this paper are threefold: (i) We introduce TSCM, a novel cross-metric knowledge distillation approach applied to VPR, which enables the student model to even outperform the teacher model. (ii) Both our devised new teacher and student networks achieve better performance in terms of VPR accuracy than the state-of-the-art baseline methods; (iii) While achieving superior performance, our method is more lightweight than the baseline, displaying outstanding computational efficiency. Our student model generates one descriptor and finds a matching within a 10k-image database in less than 2ms, achieving real-time performance.

II. RELATED WORK

Localization is a crucial capability for robots, commonly achieved through Simultaneous Localization and Mapping (SLAM). However, over time, inevitable drift adversely impacts localization accuracy. Responding to this challenge, researchers have increasingly focused on place recognition. Place recognition aids in loop-closure detection and re-localization by identifying previously visited locations and mitigating localization drift [3], [15], [16], [17]. This paper primarily centers on visual place recognition (VPR).

Early work by Milford et al. [18] introduced SeqSLAM for outdoor VPR. They employed handcrafted features to calculate the most suitable candidate among sequences. Vysotska and Stachniss [19] exploit multi-sequence maps as well as sequential image information to achieve effective and robust VPR. With the development of neural network technology, there has been a surge in the development of learning-based methods. Among these, NetVLAD [20], a Convolutional Neural Network (CNN)-based architecture, stands out as a pivotal advancement. It facilitates end-to-end training for place recognition, thereby achieving considerable success. Subsequently, many NetVLAD-based approaches have emerged [21], [22], [8], [23], [24]. With the advent of the Transformer [25], an increasing number of Transformer-based methods have come to the forefront. Wang et al. [26] introduced a hybrid CNN-Transformer feature extraction network, addressing the inherent locality limitation of Convolutional Neural Networks (CNNs) by incorporating the Vision Transformer to model image contextual information and dynamically combining task-related features. Hou et al. [27] used Transformer-based networks for feature extraction, presenting a hierarchical transformer for VPR.

Knowledge distillation was initially introduced for image classification by Hinton et al. [28]. This technique utilizes a powerful, well-trained teacher model to guide the training process of a lighter-weight student model. The goal is to simplify the model while preserving the original teacher model’s accuracy. Due to the success of knowledge distillation, many related works have emerged. [29], [30], [31]. Oki et al. [32]

introduced the concept of metric learning into knowledge distillation, leveraging pairs or triplets of training samples to bring the student model closer to the teacher model, which is closely aligned with the concept of cross-metric knowledge distillation proposed in this paper. In recent years, researchers have also applied knowledge distillation to place recognition. Shen et al. [33] proposed StructVPR, a novel training architecture for VPR, aiming to enhance structural knowledge within RGB global features, thereby enhancing feature stability in dynamic environments. Cui et al. [34] introduced a continual contrastive learning method named CCL to address the issue of catastrophic forgetting and enhance the overall robustness of LPR approaches. This method employs knowledge distillation to update the model, enabling it to sustain effective place recognition over extended periods and across diverse environments. Cai et al. [11] introduced STUN, a self-teaching framework that predicts place and estimates prediction uncertainty simultaneously using knowledge distillation without reducing the size.

In contrast to the abovementioned approaches, we introduce a novel knowledge distillation technique called cross-scale knowledge distillation, which helps us achieve a good balance between performance and model simplicity for VPR.

III. OUR APPROACH

In this section, we present the detailed network architecture of our teacher-student models, TSCM, and how we apply the devised cross-metric KD to VPR.

A. Teacher-Student Model

The architecture of our designed teacher network is presented in Fig. 2. It fuses powerful components from ResNet [12], vision transformer [13], inter-transformer [14] to increase complexity, enabling better learning of place knowledge in complex outdoor environments and generating distinctive place descriptors. It first exploits ResNet [12] to extract high-level and fine-grained features from the input image. While deeper network architectures have the potential to extract more intricate features, they often encounter training complexities. ResNet effectively addresses these challenges through the use of residual connections. Therefore, we employ several ResNet blocks as the backbone network, as illustrated in Fig. 2. We exclude its final stage to enable the teacher model to allocate more general features relevant to places, which is verified to be more beneficial than the features of the last layer in our ablation studies. Previous research [11] did not utilize the middle features directly for the final descriptor, but our findings show that using these features improves outcomes. As a result, we preserve the (1×1024) features from the ResNet branch and incorporate them into the final global descriptor aggregation process. This insight is also verified in Fig. 7.

While ResNet has demonstrated good performance, its feature extraction tends to overlook interrelationships among distinct regions within an image. In contrast, Vision Transformer (ViT) [13] extracts both global and local features from the image while simultaneously capturing correlation

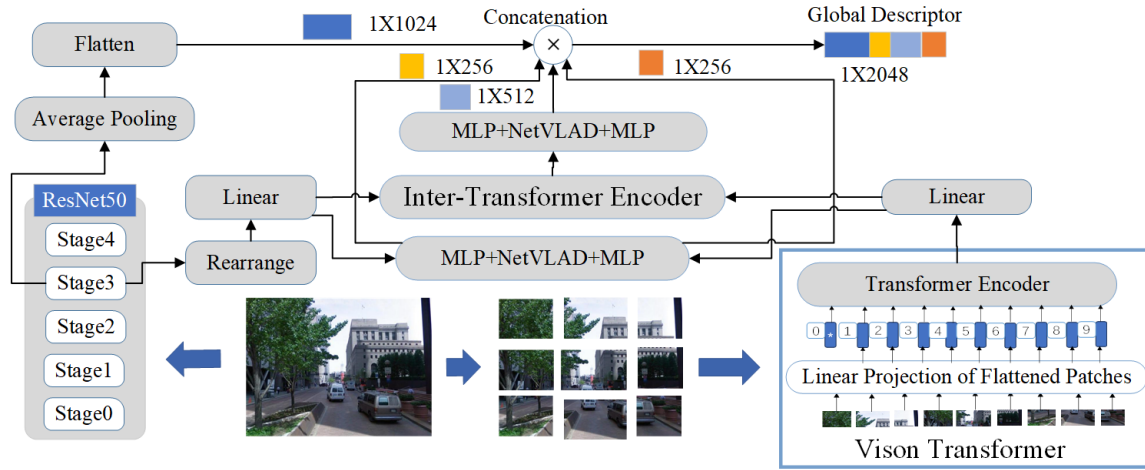


Fig. 2: The pipeline overview of our proposed teacher network. It processes input images through a multi-stage feature extraction process using ResNet and Vision Transformer (ViT). The extracted features are further processed by NetVLAD and MLPs to create initial descriptors. The final global descriptor is the concatenation of features from all branches.




Fig. 3: The structure of the Inter-Transformer Encoder

information among various image regions. It establishes connections among these disparate regions, enhancing its capacity to capture global and contextual information. Therefore, we also employ ViT to generate global features from the input image. After acquiring fine-grained features of (1×256) from ResNet and global features of the same size from ViT, we combine them to create more representative image features, generating distinctive global descriptors.

To this end, we introduce an inter-transformer encoder to fuse the features extracted by ResNet and ViT as shown in Fig. 3. It has two branches, each with a multi-head attention(MHSA) and a feed-forward network(FFN). The outputs of each branch are finally fused by a NetVLAD and MLP module to obtain the initial descriptor of (1×512) . The MHSA of Inter-Transformer in Fig. 3 can be mathematically formulated as follows:

$$\begin{aligned} F_{Att}^{Res} &= \text{Attention}(Q^{Res}, K^{ViT}, V^{ViT}) \\ &= \text{softmax}\left(\frac{Q^{Res} K^{ViT}}{\sqrt{d_k}}\right) V^{ViT} \end{aligned} \quad (1)$$

where Q^{Res} is substantially the query split of the feature

from ResNet. K^{ViT}, V^{ViT} are the key and value splits of the feature from ViT F^{ViT} . $\sqrt{d_k}$ represents the dimension of splits. F_{Att}^{Res} denotes the output feature from the MHSA in the ResNet branch. This feature is subsequently fed into the FFN to generate the attention-enhanced feature, facilitated by the following mathematical formulation:

$$F_{ITF}^{Res} = \text{FFN}(F_{Att}^{Res} + F^{Res}) + F_{Att}^{Res} + F^{Res} \quad (2)$$

Likewise, we perform the same operations on the ViT branch.:

$$\begin{aligned} F_{Att}^{ViT} &= \text{Attention}(Q^{ViT}, K^{Res}, V^{Res}) \\ &= \text{softmax}\left(\frac{Q^{ViT} K^{Res}}{\sqrt{d_k}}\right) V^{Res} \end{aligned} \quad (3)$$

$$F_{ITF}^{ViT} = \text{FFN}(F_{Att}^{ViT} + F^{ViT}) + F_{Att}^{ViT} + F^{ViT} \quad (4)$$

Features F_{ITF}^{Res} and F_{ITF}^{ViT} are concatenated to form the composite feature F_{ITF} . Subsequently, this feature F_{Inter} is fed into the NetVLAD-MLPs combos [20] to generate a descriptor. The final global descriptor is created by integrating this descriptor with those from other branches. While the teacher network gathers maximum information using multiple complex components, including potentially redundant knowledge, this does not affect deployment, as a lighter student network is deployed into practical applications.

The overview of our student network is depicted in Fig. 4. Unlike the complex architecture of the teacher network, the student network has a much simpler structure, as illustrated in Fig. 7. The ResNet branch is pivotal in place recognition, so we keep it in the student network. To simplify the student network, we replace both the ViT and the Inter-Transformer with a layer of convolutional layers. Our student network is relatively smaller but still includes multiple branches, allowing it to learn various types of place knowledge from the teacher network through cross-metric knowledge distillation, explained in the next section.




Fig. 4: The overview of the student network

B. Cross-Metric Knowledge Distillation

Knowledge distillation (KD) transfers complex knowledge from a larger teacher model to a smaller student model while keeping the student compact. Existing KD methods for VPR typically use triplet loss [35], as shown in the left figure of Fig. 5. $T(\cdot)$ represents the teacher model, and $S(\cdot)$ represents the student model. a , p , and n correspond to the anchor, positive, and negative samples. Solid lines indicate distances to be reduced, while dashed lines indicate distances to be increased. Triplet loss is well-suited for supervising VPR tasks, as it helps find similar places and distinguish different places. During KD, the teacher model remains fixed, and the triplet loss is applied to the student model as follows:

$$L_{\text{Tri}} = \max \{d(S(a) - S(p)) - d(S(a) - S(n)) + m, 0\}, \quad (5)$$

where m is the margin.

Traditional KD aligns outputs of the student model with those of the teacher, keeping distances between $T(a)$, $T(n)$, and $T(p)$ unchanged, known as soft targets (red lines in Fig. 5), and the corresponding loss function is:

$$L_{\text{soft}} = \sum_i^N d(S(a_i) - T(a_i)) + d(S(p_i) - T(p_i)) + d(S(n_i) - T(n_i)). \quad (6)$$

In VPR, our objective is to ensure that positive sample outputs are close to the anchor while negative sample outputs are far from it, referred to as hard targets (illustrated by the green lines in Fig. 5). This goal should be consistent whether dealing with teacher samples, student samples, or interactions between them. Traditionally, the triplet loss focuses only on the relationship between the outputs of the student model, overlooking the relationship between the outputs of the student and teacher models (e.g., the relationship between $S(a)$ and $T(p)$). One can apply the triplet loss within a single model to achieve the desired hard targets ($L_{\text{Tri}} = L_{\text{hard}}$). However, there is limited research on achieving these hard targets between the teacher and student models.

To fill in this gap, we investigated the relationships between samples in both the teacher and student models during the distillation process. We aimed to bring a positive student sample, $S(p)$, closer to the teacher anchor, $T(a)$, while




Fig. 5: Comparison of different knowledge distillation (KD) strategies on triplet loss. $T(\cdot)$ represents the output of the teacher model, and $S(\cdot)$ represents the output of the student model. a , p , n represent anchor, positive, and negative samples. Solid lines indicate distances to be reduced, while dashed lines indicate distances to be increased. The left figure shows the traditional KD method. The middle one is a fully connected knowledge distillation. The right figure illustrates our proposed strategy, which exploits constraints within each model and facilitates cross-model interactions, enhancing supervision for knowledge distillation in place recognition.

pushing negative student samples, $S(n)$, away from $T(a)$, and vice versa. To capture these relationships, we introduced black lines in addition to the red and green lines, as depicted in the second figure of Fig. 5, to represent the distances between samples across teacher and student. This can be formulated as:

$$L_{\text{cm-full}} = \max \{d(S(a) - T(p)) + d(S(p) - T(a)) - d(S(a) - T(n)) - d(S(n) - T(a)) + m, 0\}. \quad (7)$$

Ideally, using Eq. (7) as the loss for VPR would maximize information utilization during KD. However, including all sample distances presents a challenge: We aim to separate $S(a)$ from $T(n)$ and $S(n)$ from $T(a)$, but this may inadvertently cause $S(a)$ to diverge from $T(a)$ and $S(n)$ to diverge from $T(n)$, contradicting our soft targets (red solid line). To address this, we propose our cross-metric KD, which discards the dashed black lines and maintains the solid black lines as cross-model constraints, formulated as:

$$L_{\text{cm}} = \sum_i^N d(S(a_i) - T(p_i)) + d(S(p_i) - T(a_i)) \quad (8)$$

As $S(a)$ approaches $T(p)$ and $S(p)$ approaches $T(a)$, the distances between $S(a)$ and $T(a)$, $S(p)$ and $T(p)$, and $S(a)$ and $S(p)$ tend to decrease. As illustrated in the right figure of Fig. 5, our design benefits both soft and hard targets, thus increasing the KD performance specifically in VPR tasks.

Taking into account the aforementioned considerations, we define the final loss function for our cross-metric KD for VPR as follows:

$$L_{\text{total}} = L_{\text{hard}} + L_{\text{soft}} + L_{\text{cm}}. \quad (9)$$

IV. EXPERIMENTAL EVALUATION

The main focus of this work is a teacher-student network exploiting our devised cross-metric knowledge distillation (KD) for fast and accurate visual place recognition (VPR). We present our experiments to show the capabilities of our method and support our key claims, which are: (i) TSCM introduces the concept of cross-metric KD to VPR, allowing

our smaller student model to perform similarly, and sometimes even better, than the larger teacher model; (ii) TSCM attains superior recognition accuracy while maintaining a more lightweight model in comparison to the state-of-the-art baseline methods; (iii) TSCM demonstrates exceptional computational efficiency, compressing images into descriptors in 1.3 ms and finding a matching in under 0.6 ms per query using a 10k-image database.

A. Experimental Setup

1) Dataset: Our evaluation of TSCM and other methods is conducted on the Pittsburgh dataset [36], a comprehensive image repository sourced from Google Street View. The original dataset encompasses 250k training samples and 24k test query samples. For the training procedure, we employ the same partitioning scheme as NetVLAD [20], using a reduced subset of 10k samples for training and validation while employing Pittsburgh30k with 10k test data and Pittsburgh250k with 84k test data for testing. Notably, these images have been captured at various times of the day and span multiple years.

2) Configuration: In TSCM, the Vision Transformer (ViT) is pre-trained at a resolution of 224x224, with a patch size of 16. For the Inter-Transformer, our setup aligns with that of CVTNet [14], featuring the following specifications: an embedding dimension of $d_{\text{model}} = 512$, several heads of $n_{\text{heads}} = 4$, and an intermediate dimension for the feed-forward layer of $d_{\text{ff}} = 1024$. In the case of NetVLAD following the Inter-Transformer, we configure it with an input feature vector dimension of $\text{feature_size} = 512$, an output feature dimension of $d_{\text{output}} = 512$, and a total of $d_K = 64$ clusters. For other instances of NetVLAD, we adjust the input feature vector dimension to $\text{feature_size} = 256$, with an output feature dimension of $d_{\text{output}} = 256$, while keeping the remaining parameters consistent with NetVLAD following the Inter-Transformer. During the network training phase, we adhere to the STUN setup [11], employing a batch size of 8, utilizing the Adam optimizer with an initial learning rate of 1×10^{-5} , which undergoes decay by a factor of 0.99 after each epoch, and applying a weight decay of 0.001. Our training infrastructure is hosted on a cloud server equipped with the following specifications: a single V100-SXM2-32GB (32GB) GPU, complemented by a 12-core Intel(R) Xeon(R) Platinum 8255C CPU @ 2.50GHz. The server environment operates on the PyTorch 1.10.0 image, utilizing Python 3.8 (Ubuntu 20.04) and CUDA 11.3.

3) Evaluation Metrics: Our experiments are designed to assess the recognition performance of our model and the impact of cross-metric knowledge distillation. We employ recall@N, mAP@N, and AP as metrics to measure recognition performance and recall@N as a metric to measure the impact of cross-metric knowledge distillation.

B. VPR Performance

The first experiment aims to assess the VPR performance of our approach. The experimental results support that TSCM achieves superior recognition accuracy compared to

TABLE I: Recognition performance on the Pittsburgh30k dataset

	r@1/5/10 \uparrow	mAP@1/5/10 \uparrow	AP \uparrow
MC Dropout [37]	0.302/0.523/0.611	0.302/0.108/0.061	0.463
PFE w/o MLS [38]	0.550/0.805/0.876	0.550/0.266/0.167	0.690
PFE w/ MLS [38]	0.444/0.680/0.764	0.444/0.199/0.120	0.655
BTL [39]	0.515/0.766/0.840	0.515/0.252/0.158	0.591
STUN(teacher) [11]	0.548/0.802/0.877	0.548/0.268/0.167	0.678
STUN(student) [11]	0.613/0.840/0.898	0.613 /0.280/0.171	0.739
TSCM(our teacher)	<u>0.730/0.909/0.942</u>	<u>0.730 / 0.334/0.204</u>	<u>0.862</u>
TSCM(our student)	0.735/0.914/0.945	0.735 / 0.329/0.198	0.864

bold denotes the best and underline denotes the second best results.

TABLE II: Recognition performance on the Pittsburgh250k dataset

	r@1/5/10 \uparrow	mAP@1/5/10 \uparrow	AP \uparrow
STUN(teacher) [11]	0.519/0.757/0.829	0.519/0.235/0.143	0.687
STUN(student) [11]	0.559/0.788/0.852	0.559/0.249/0.150	0.710
TSCM(our teacher)	<u>0.680/0.862/0.900</u>	<u>0.680/0.299/0.177</u>	<u>0.846</u>
TSCM(our student)	0.684/0.867/0.906	0.684/0.294/0.173	0.847

bold denotes the best and underline denotes the second best results.

the baseline and our proposed KD strategy, allowing the smaller student model to perform similarly, and sometimes even better, than the larger teacher model. We follow the experimental setup the same as STUN [11].

Tab. I presents the recognition performance of various methods on the Pittsburgh30k dataset. In terms of VPR performance, both our student and teacher method outperforms all other baseline methods. Notably, our student model outperforms our large teacher model in recall metrics while performing similarly in mAP and AP metrics. Our student model surpasses other baseline models in all metrics, including larger teacher models of baseline methods.

To validate our method’s generalization ability on a larger database, we compared it with STUN on Pittsburgh250k without retraining or fine-tuning. As depicted in Tab. II, our method consistently achieves significantly better performance compared to the baseline methods. A more detailed comparison is also provided in the top-left figure of Fig. 7. The recall and precision curves clearly illustrate that our method outperforms STUN on two datasets, affirming its superior performance. We additionally provide more visualization results in Fig. 6. As can be seen, our TSCM consistently identifies true positives in VPR with smaller distances to the query compared to STUN, which produces incorrect matches.

C. Ablation Studies and Insights

In this section, we assess the impact of our proposed cross-metric KD and validate our designs on the architecture and our KD strategies. In Tab. IV, we denote D_{SaTp} as the distances between student anchors and teacher positive samples, and so on. As can be seen, our cross-metric knowledge distillation method enhances place recognition within TSCM. Moreover, the performance of our proposed $L_{\text{cm}} = xx$ is notably superior using D_{SaTp} and D_{SpTa} , compared to scenarios where D_{SpTa} , D_{SpTa} , D_{SnTa} , and D_{SaTn} are all included, or when either D_{SpTa} , D_{SpTa} , D_{SnTa} or D_{SpTa} ,

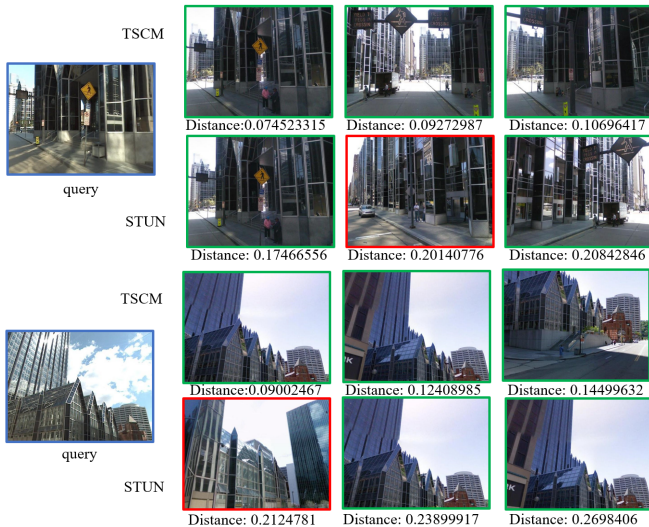


Fig. 6: Comparing qualitative VPR results between ours and STUN

TABLE III: Comparisons in params and runtime at database of 10k

model	Total params	FLOPS (10^9)	Descriptor generation	Matching time	Total time
Our teacher	194 M	81.626	4.5 ms	0.7 ms	5.2 ms
STUN's student	27 M	4.136	1.3 ms	0.6 ms	1.9 ms
Our student	13 M	3.440	1.3 ms	0.6 ms	1.9 ms

D_{SpTa} , D_{SaTn} are added. This substantiates our hypotheses presented in Fig. 5.

We conducted ablation experiments to assess the effectiveness of specific architectural components within our teacher network. This included investigating the impact of retaining the ResNet branch, removing the last ResNet stage, and evaluating the contribution of the Inter-Transformer Encoder. The results of these experiments, depicted in Fig. 7, revealed the following key findings: i) Retaining the ResNet branch significantly improved model performance; ii) Removing the final ResNet stage did not lead to performance degradation but rather resulted in improvement; iii) The inclusion of the Inter-Transformer Encoder had a positive impact on model performance. In summary, our modified architecture outperforms the original one.

We conducted ablation experiments on our teacher network, examining the effects of specific architectural components. We analyze the impact of retaining the ResNet branch, removing the last ResNet stage, and evaluating the contribution of the Inter-Transformer encoder. As shown in Fig. 7, the results reveal the following insights: i) Keeping the ResNet branch improved model performance significantly; ii) Eliminating the final ResNet stage leads to performance improvement; iii) The inclusion of the Inter-Transformer Encoder positively affected model performance. In sum, all of our design choices contribute positively to VPR tasks.

D. Computational Efficiency and Runtime

This section provides a comparative analysis of the number of parameters and FLOPS in our teacher model, student model, and STUN's student model, as illustrated in Tab. III. The results show that our student model, when compared to

TABLE IV: Effect of cross-metric knowledge distillation

L_h	L_s	D_1	D_2	D_3	D_4	recall@1/5/10 on 30k	recall@1/5/10 on 250k
Δ						0.730/0.909/0.942	0.680/0.862/0.900
\checkmark	\checkmark					0.727/0.908/0.941	0.679/0.863/0.898
\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	0.722/0.904/0.937	0.671/0.851/0.890
\checkmark	\checkmark	\checkmark	\checkmark	\checkmark		0.729/0.908/0.942	0.680/0.864/0.890
\checkmark	\checkmark	\checkmark	\checkmark		\checkmark	0.725/0.911/0.943	0.676/0.859/0.896
\checkmark	\checkmark	\checkmark	\checkmark		\checkmark	0.735/0.914/0.945	0.684/0.867/0.906

Δ denotes the teacher model. **bold** denotes the best result L_h indicates L_{hard} ; L_s indicates L_{soft} ; D_1 indicates D_{SaTp} ; D_2 indicates D_{SpTa} ; D_3 indicates D_{SaTn} ; D_4 indicates D_{SnTa} .




Fig. 7: (a) Recall and precision curve for STUN and our method. (b) ablation study on ResNet branch. (c) ablation study on last stage. (d) ablation study on Inter-Transformer Encoder

both our teacher model and STUN's student model, has fewer parameters and lower computational requirements, making it a more lightweight option. Tab. III also presents the runtime outcomes for TSCM, affirming our third claim that TSCM exhibits remarkable computational efficiency, capable of compressing captured images into descriptors about 1.3 ms and matching in less than 0.6 ms per query at a database size of 10k within the experimental environment.

V. CONCLUSION

In this paper, we introduce a novel knowledge distillation technique for visual place recognition called cross-metric knowledge distillation. This method empowers the student network to surpass the teacher network's performance when integrated into our teacher-student model for VPR. We conducted extensive experiments using various datasets and compared with existing methods to validate our claims. The results demonstrate that our approach achieves superior recognition accuracy while maintaining a more lightweight model than the baseline method. Furthermore, our proposed cross-metric knowledge distillation method outperforms conventional knowledge distillation techniques. Finally, we evaluated the model's processing speed, confirming its ability to compress captured images into descriptors in 1.3 ms and achieve matching in less than 0.6 ms per query with a database size of 10k, demonstrating real-time performance.

REFERENCES

- [1] H. Yin, X. Xu, S. Lu, X. Chen, R. Xiong, S. Shen, C. Stachniss, and Y. Wang, "A survey on global lidar localization: Challenges, advances and open problems," 2023.
- [2] P. Yin, S. Zhao, I. Cisneros, A. Abuduweili, G. Huang, M. Milford, C. Liu, H. Choset, and S. Scherer, "General place recognition survey: Towards the real-world autonomy age," 2022.
- [3] X. Chen, T. Läbe, A. Milioto, T. Röhling, O. Vysotska, A. Haag, J. Behley, and C. Stachniss, "OverlapNet: Loop Closing for LiDAR-based SLAM," in *Proc. of Robotics: Science and Systems (RSS)*, 2020. [Online]. Available: <https://www.ipb.uni-bonn.de/wp-content/papercite-data/pdf/chen2020rss.pdf>
- [4] J. Ma, J. Zhang, J. Xu, R. Ai, W. Gu, and X. Chen, "Overlaptransformer: An efficient and yaw-angle-invariant transformer network for lidar-based place recognition," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 3, pp. 6958–6965, 2022.
- [5] J. Ma, X. Chen, J. Xu, and G. Xiong, "Seqot: A spatial-temporal transformer network for place recognition using sequential lidar data," *IEEE Trans. on Industrial Electronics (TIE)*, vol. 70, no. 8, pp. 8225–8234, 2022.
- [6] X. Zhang, L. Wang, and Y. Su, "Visual place recognition: A survey from deep learning perspective," *Pattern Recognition*, vol. 113, p. 107760, 2021.
- [7] C. Masone and B. Caputo, "A survey on deep visual place recognition," *IEEE Access*, vol. 9, pp. 19 516–19 547, 2021.
- [8] S. Hausler, S. Garg, M. Xu, M. Milford, and T. Fischer, "Patch-netvlad: Multi-scale fusion of locally-global descriptors for place recognition," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 14 141–14 152.
- [9] J. Yu, C. Zhu, J. Zhang, Q. Huang, and D. Tao, "Spatial pyramid-enhanced netvlad with weighted triplet loss for place recognition," *IEEE Trans. on neural networks and learning systems(TNNLS)*, vol. 31, no. 2, pp. 661–674, 2019.
- [10] L. Hui, M. Cheng, J. Xie, J. Yang, and M.-M. Cheng, "Efficient 3d point cloud feature learning for large-scale place recognition," *IEEE Trans. on Image Processing*, vol. 31, pp. 1258–1270, 2022.
- [11] K. Cai, C. X. Lu, and X. Huang, "Stun: Self-teaching uncertainty estimation for place recognition," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 6614–6621.
- [12] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016. [Online]. Available: <https://arxiv.org/pdf/1512.03385>
- [13] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," 2021.
- [14] J. Ma, G. Xiong, J. Xu, and X. Chen, "CVTNet: A Cross-View Transformer Network for LiDAR-Based Place Recognition in Autonomous Driving Environments," *IEEE Trans. on Industrial Informatics (TII)*, 2023.
- [15] X. Chen, T. Läbe, A. Milioto, T. Röhling, J. Behley, and C. Stachniss, "OverlapNet: A Siamese Network for Computing LiDAR Scan Similarity with Applications to Loop Closing and Localization," *Autonomous Robots*, 2021. [Online]. Available: <http://www.ipb.uni-bonn.de/pdfs/chen2021auro.pdf>
- [16] Y. Wang, Z. Sun, C.-Z. Xu, S. E. Sarma, J. Yang, and H. Kong, "Lidar iris for loop-closure detection," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5769–5775.
- [17] S. Siam and H. Zhang, "Fast-SeqSLAM: A Fast Appearance Based Place Recognition Algorithm," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2017.
- [18] M. Milford and G. Wyeth, "SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2012.
- [19] O. Vysotska and C. Stachniss, "Effective Visual Place Recognition Using Multi-Sequence Maps," *IEEE Robotics and Automation Letters (RA-L)*, vol. 4, pp. 1730–1736, 2019. [Online]. Available: <http://www.ipb.uni-bonn.de/pdfs/vysotska2019ral.pdf>
- [20] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "Netvlad: Cnn architecture for weakly supervised place recognition," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 5297–5307.
- [21] J. Yu, C. Zhu, J. Zhang, Q. Huang, and D. Tao, "Spatial pyramid-enhanced netvlad with weighted triplet loss for place recognition," *IEEE Trans. on Neural Networks and Learning Systems(TNNLS)*, vol. 31, no. 2, pp. 661–674, 2020.
- [22] A. Khaliq, S. Ehsan, Z. Chen, M. Milford, and K. McDonald-Maier, "A holistic visual place recognition approach using lightweight cnns for significant viewpoint and appearance changes," *IEEE Trans. on Robotics (TRO)*, vol. 36, no. 2, pp. 561–569, 2020.
- [23] Z. Li, C. D. W. Lee, B. X. L. Tung, Z. Huang, D. Rus, and M. H. Ang, "Hot-netvlad: Learning discriminatory key points for visual place recognition," *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 2, pp. 974–980, 2023.
- [24] A. Khaliq, M. Milford, and S. Garg, "Multires-netvlad: Augmenting place recognition training with low-resolution imagery," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 2, pp. 3882–3889, 2022.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Proc. of the Advances in Neural Information Processing Systems (NIPS)*, vol. 30, 2017.
- [26] Y. Wang, Y. Qiu, P. Cheng, and J. Zhang, "Hybrid cnn-transformer features for visual place recognition," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 33, no. 3, pp. 1109–1122, 2023.
- [27] Z. Hou, Y. Yan, C. Xu, and H. Kong, "Hitpr: Hierarchical transformer for place recognition in point cloud," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2022, pp. 2612–2618.
- [28] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint*, 2015.
- [29] J. H. Cho and B. Hariharan, "On the efficacy of knowledge distillation," in *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2019, pp. 4794–4802.
- [30] W. Park, D. Kim, Y. Lu, and M. Cho, "Relational knowledge distillation," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 3967–3976.
- [31] B. Zhao, Q. Cui, R. Song, Y. Qiu, and J. Liang, "Decoupled knowledge distillation," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 11 953–11 962.
- [32] H. Oki, M. Abe, J. Miyao, and T. Kurita, "Triplet loss for knowledge distillation," in *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2020, pp. 1–7.
- [33] Y. Shen, S. Zhou, J. Fu, R. Wang, S. Chen, and N. Zheng, "Structvpr: Distill structural knowledge with weighting samples for visual place recognition," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 11 217–11 226.
- [34] J. Cui and X. Chen, "CCL: Continual contrastive learning for LiDAR place recognition," *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 8, pp. 4433–4440, 2023.
- [35] E. Hoffer and N. Ailon, "Deep metric learning using triplet network," in *International Workshop on Similarity-Based Pattern Recognition*, 2015, pp. 84–92.
- [36] A. Torii, J. Sivic, M. Okutomi, and T. Pajdla, "Visual place recognition with repetitive structures," *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 37, no. 11, pp. 2346–2359, 2015.
- [37] A. Taha, Y.-T. Chen, X. Yang, T. Misu, and L. Davis, "Exploring uncertainty in conditional multi-modal retrieval systems," *arXiv preprint*, 2019.
- [38] Y. Shi and A. K. Jain, "Probabilistic face embeddings," in *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2019, pp. 6902–6911.
- [39] F. Warburg, M. Jørgensen, J. Civera, and S. Hauberg, "Bayesian triplet loss: Uncertainty quantification in image retrieval," in *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2021, pp. 12 158–12 168.