Spectral Geometric Verification

Re-Ranking Point Cloud Retrieval for Metric Localization

Kavisha Vidanapathirana, Peyman Moghadam, Sridha Sridharan, Clinton Fookes
IEEE Robotics and Automation Letters (RA-L)

Registration-free point cloud re-ranking for LiDAR-based 6-DoF metric localization.

Motivation

LiDAR-based Metric Localization in large scale environments is typically formulated as a hierarchical process of first estimating a coarse-level place candidate (using retrieval), and then estimating the 6-DoF transformation.

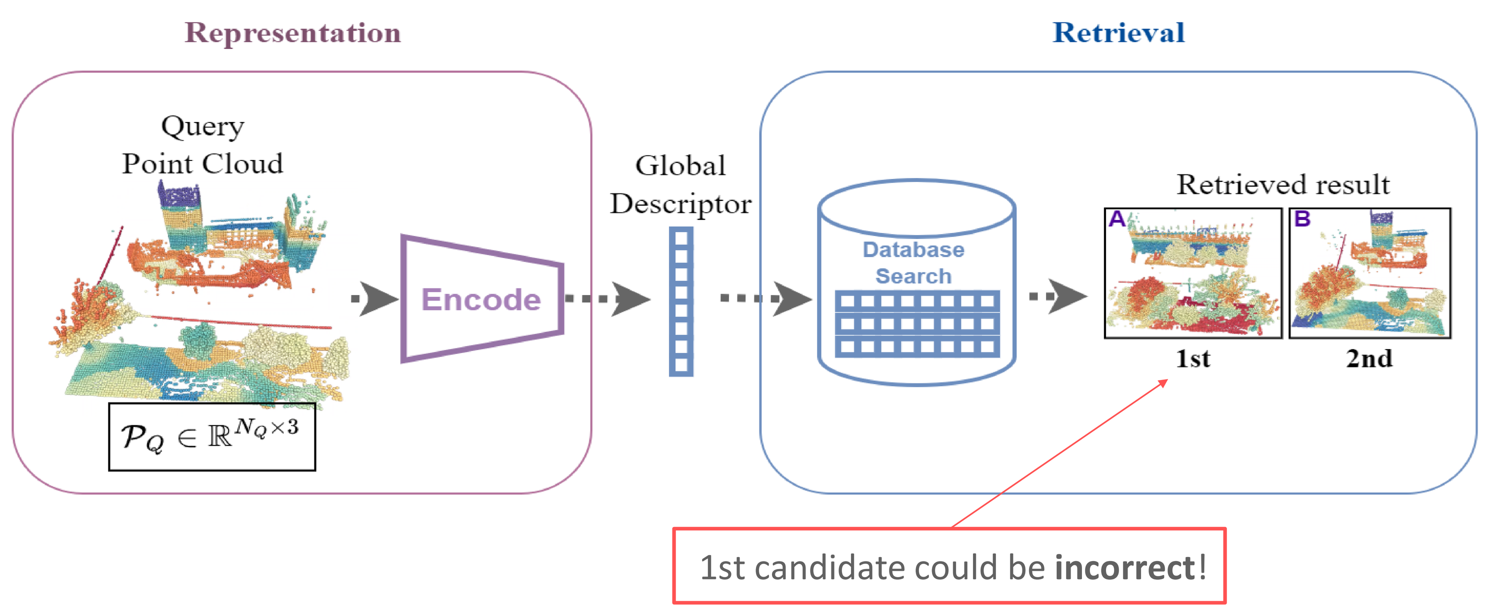

Figure 1: The first step in metric localization, known as place recognition, is typically formulated as a retrieval problem, by encoding the query point cloud into a compact global descriptor, which is used for retrieval over a large database.

The best retrieved candidate can sometimes be incorrect, and an incorrect retrieval will hinder the registration and prevent accurate localization.

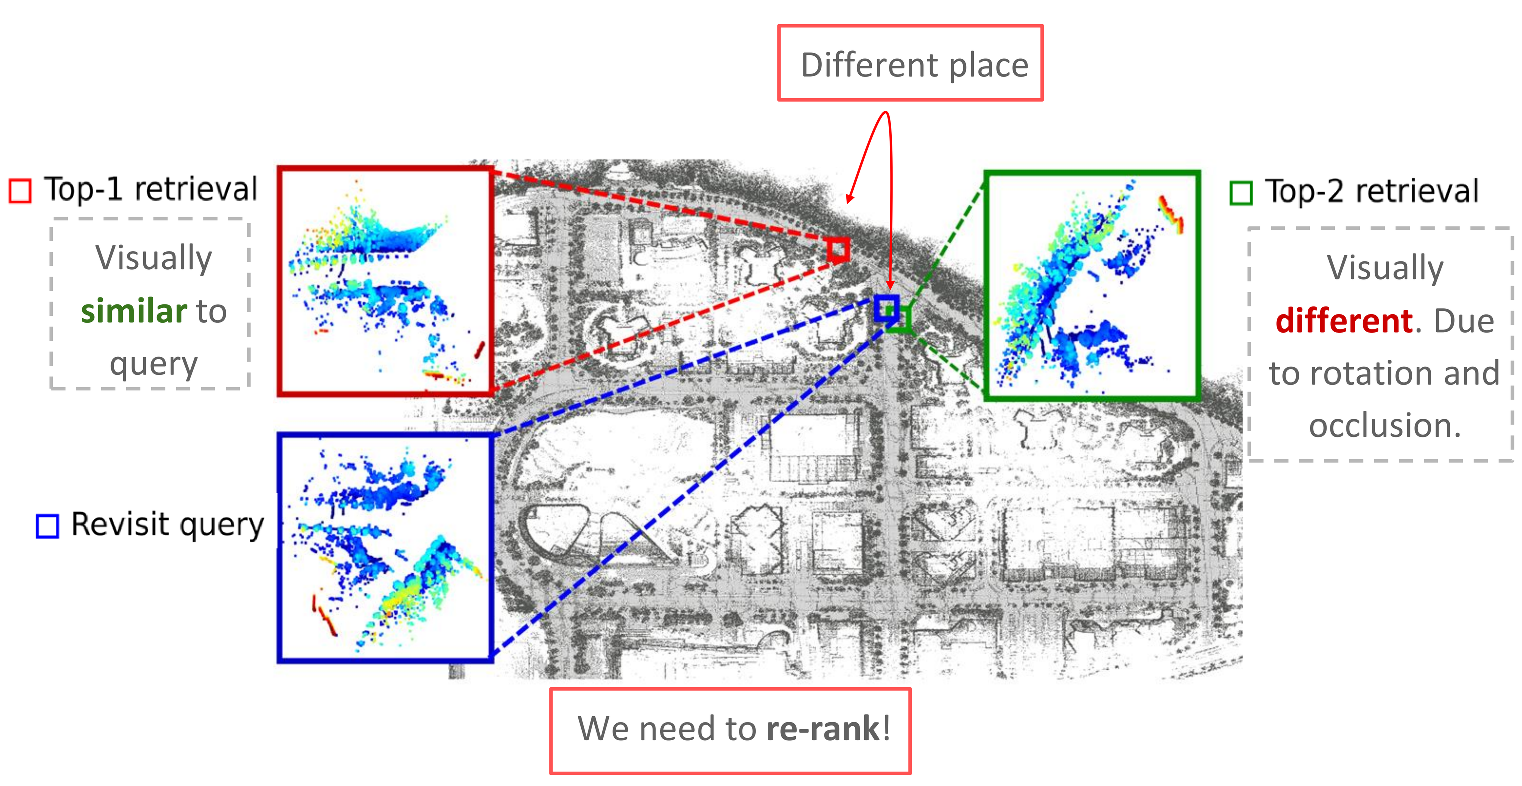

Figure 2: For a query point cloud (shown in the blue box), the top-1 candidate (shown in the red box) appears visually similar. However, we see that top-1 candidate is actually from a different place. Instead, the correct candidate (shown in green) is retrieved as the second-best option. This happens when there are multiple structurally similar places in the database and the network struggles to discriminate between them. Re-ranking aims to address this failure case by re-ordering this set of candidates such the correct candidate will be ranked first.

Geometric Verification in Real-Time

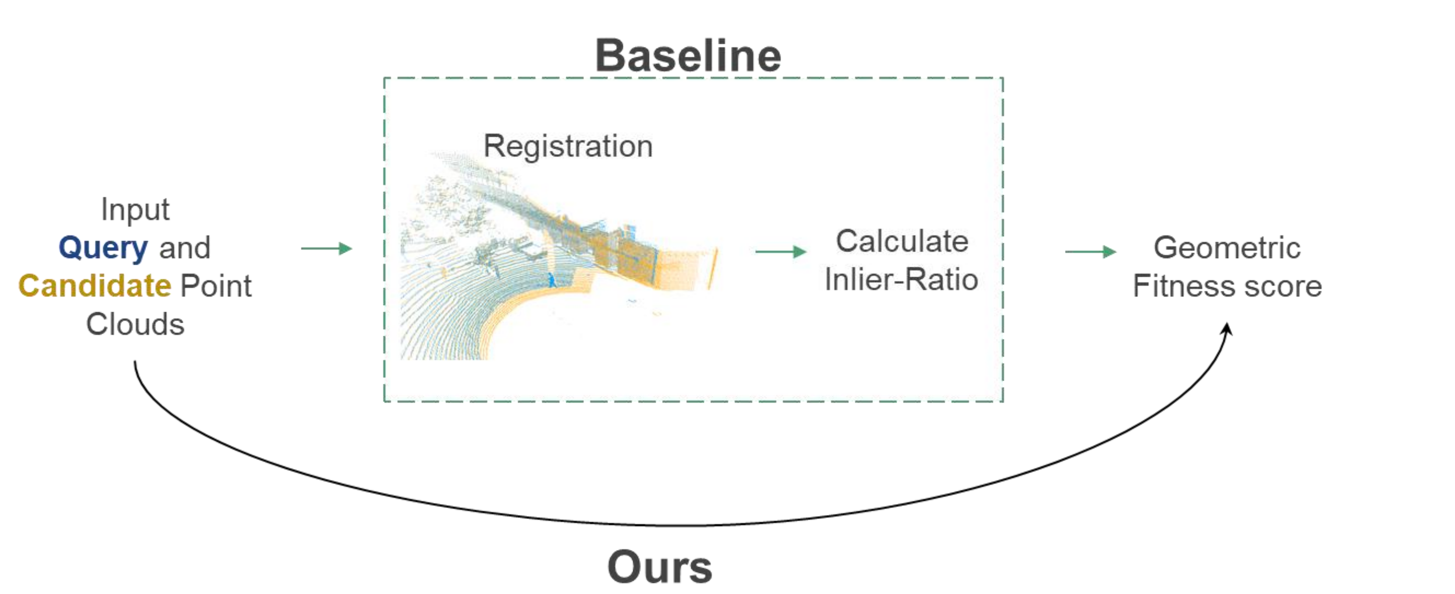

Figure 3: The baseline method for checking geometric consistency is to register the query and candidate point clouds and calculate the inlier ratio after registration. This process of pairwise registration for all candidates prohibits processing a large number of candidates during real-time operation. To address this, we propose a geometric verification method that does not require registration and directly calculates a geometric fitness score.

Method

We can avoid the calculation of a registration-based inlier score by instead using a 'correspondence compatibility' score. We compute our 'correspondence compatibility' score by constructing a graph where the edges between correspondence pairs are weighted by the degree of preservation of pair-wise distances between their points.

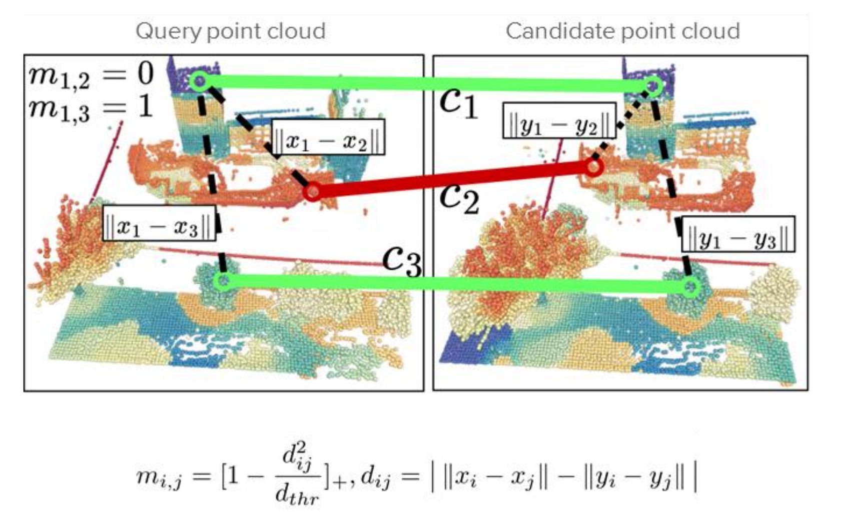

Figure 4: In this figure, we see 3 example correspondences where the green lines indicate correct correspondences, and the red line indicates an incorrect correspondence. The compatibility of two correspondences, i and j, noted as m_{ij}, is computed by checking the change in relative distances between their points. That is the change in length of the dotted black lines.

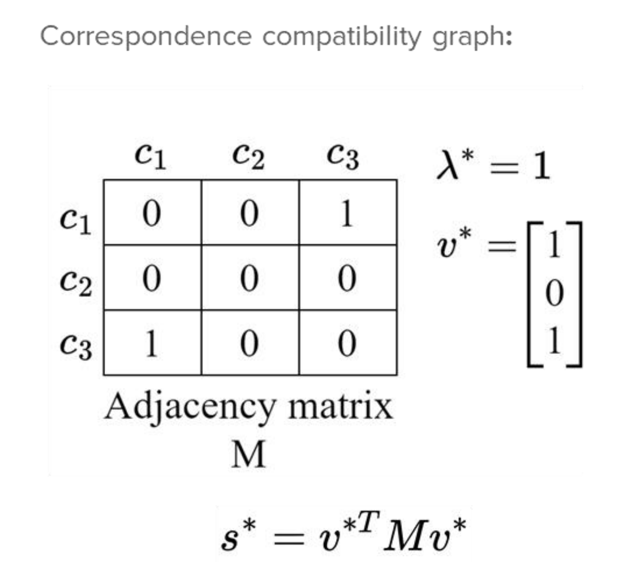

Figure 5: The leading eigenvector of this matrix can be used as an indicator vector that estimates the nodes of the main cluster of this graph. The score s^* is the average compatibility score of the main cluster. s^* can be used as a fitness score for geometric verification based re-ranking.

We then use spectral graph theory to estimate the subset of inlier correspondences and take the average of their compatibility score as the fitness score for our re-ranking.

Results

Place Recognition:

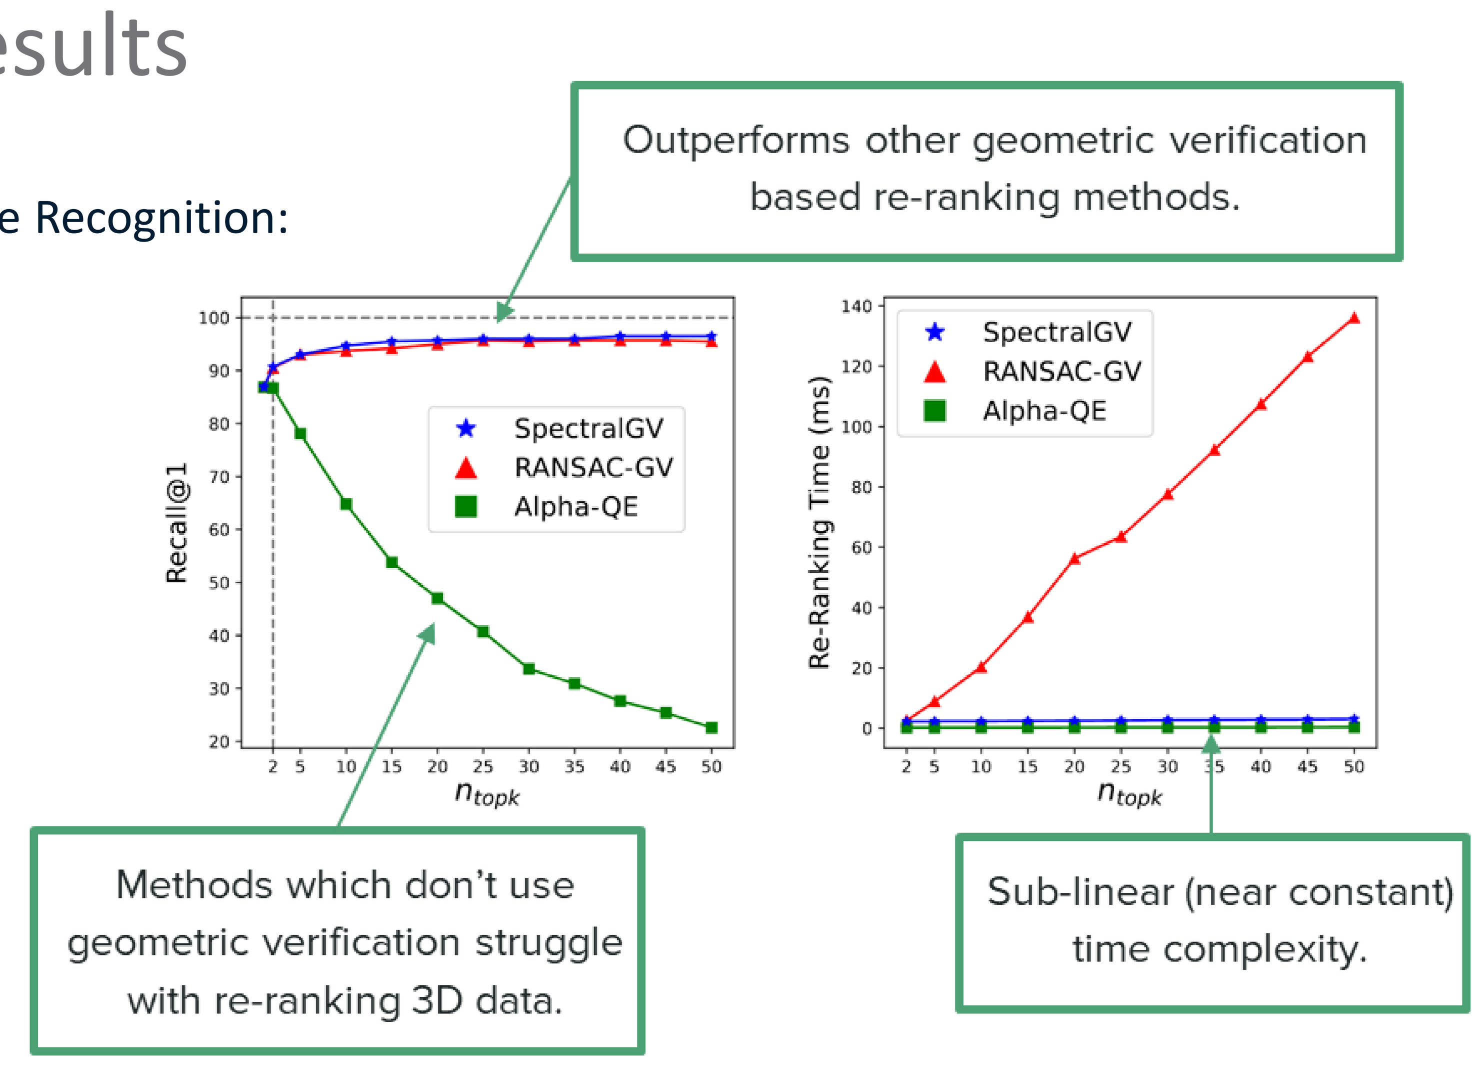

Metric Localization:

Evaluation on unseen dataset

Architecture agnostic

	S (↑)	RTE (↓)	RRE (↓)
EgoNN [5]	88.2	42.45	13.36
LCDNet [10]	94.0	10.40	6.37
LoGG3D-Net [11]	97.0	9.51	3.64
EgoNN + SGV	98.7	4.32	1.44
LCDNet + SGV	98.5	**0.79**	2.35
LoGG3D-Net + SGV	**99.5**	2.88	**1.07**

As Australia's national science agency and innovation catalyst, CSIRO is solving the greatest challenges through innovative science and technology.

CSIRO. Unlocking a better future for everyone.

FOR FURTHER INFORMATION
Kavisha Vidanapathirana
kavisha.vidanapathirana@data61.csiro.au
research.csiro.au/robotics
github.com/csiro-robotics/SpectralGV