

Guided Similarity Separation for Image Retrieval

NEURAL INFORMATION PROCESSING SYSTEMS

NeurIPS 2019

Chundi Liu, Guangwei Yu, Maksims Volkovs, Cheng Chang, Himanshu Rai, Junwei Ma, Satya Krishna Gorti

Introduction

> Given a query image, retrieve all the relevant images.



Figure 1: Example query image (left) and four retrieved results (right). Green indicate relevant match while red indicate irrelevant match.

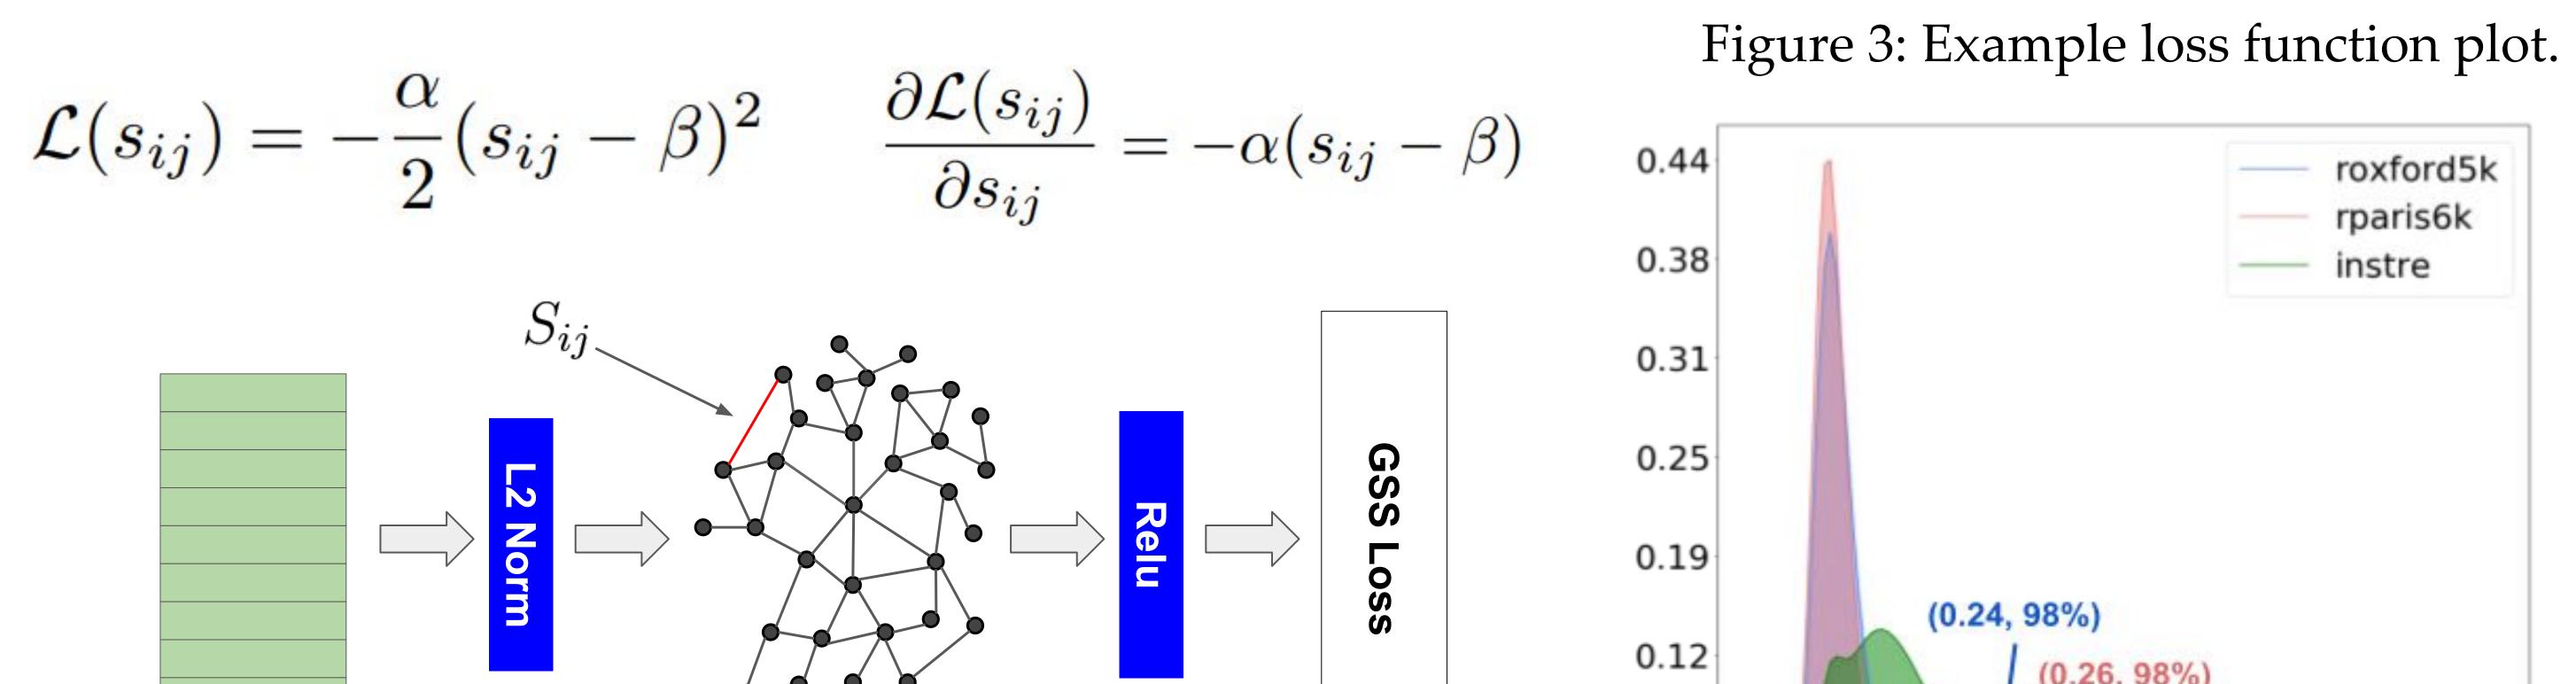
- > We propose a novel model leveraging graph convolutional networks to directly encode neighbor information into image descriptors.
- ➤ A novel guided similarity separation (GSS) loss function is introduced to optimize the proposed model in a fully unsupervised fashion.
- ➤ Refined spatial verification graph can be learned in our model. An approximate inference procedure enables us to leverage the information to achieve better results with a small constant overhead during the inference time.
- > Experiments on five public benchmarks show highly competitive performance with up to 25% relative improvement over leading baselines.

Architecture Output descriptors Non-linearity Adjacency matrix Output descriptors Non-linearity Adjacency matrix Output descriptors Non-linearity Adjacency matrix Input descriptors Mon-linearity Adjacency descriptors matrix Output Adjacency descriptors matrix Output Adjacency descriptors New Index Descriptors Non-linearity Non-linearity Adjacency descriptors matrix Output Adjacency descriptors New Index Descriptors Non-linearity Non-linearity Non-linearity Adjacency descriptors Output Adjacency descriptors New Index New Index New Index Non-linearity Non-linea

Figure 2: Overall architecture

Guided Similarity Separation Loss

- The main idea behind guided similarity separation is to increase s_{ij} if it is above a given threshold and lower it otherwise.
- ➤ This has a clustering effect where images with higher similarity scores move closer together, and those with lower scores get pushed further apart. 0.15
- > Loss function and gradient:



Descriptors

New kNN Graph

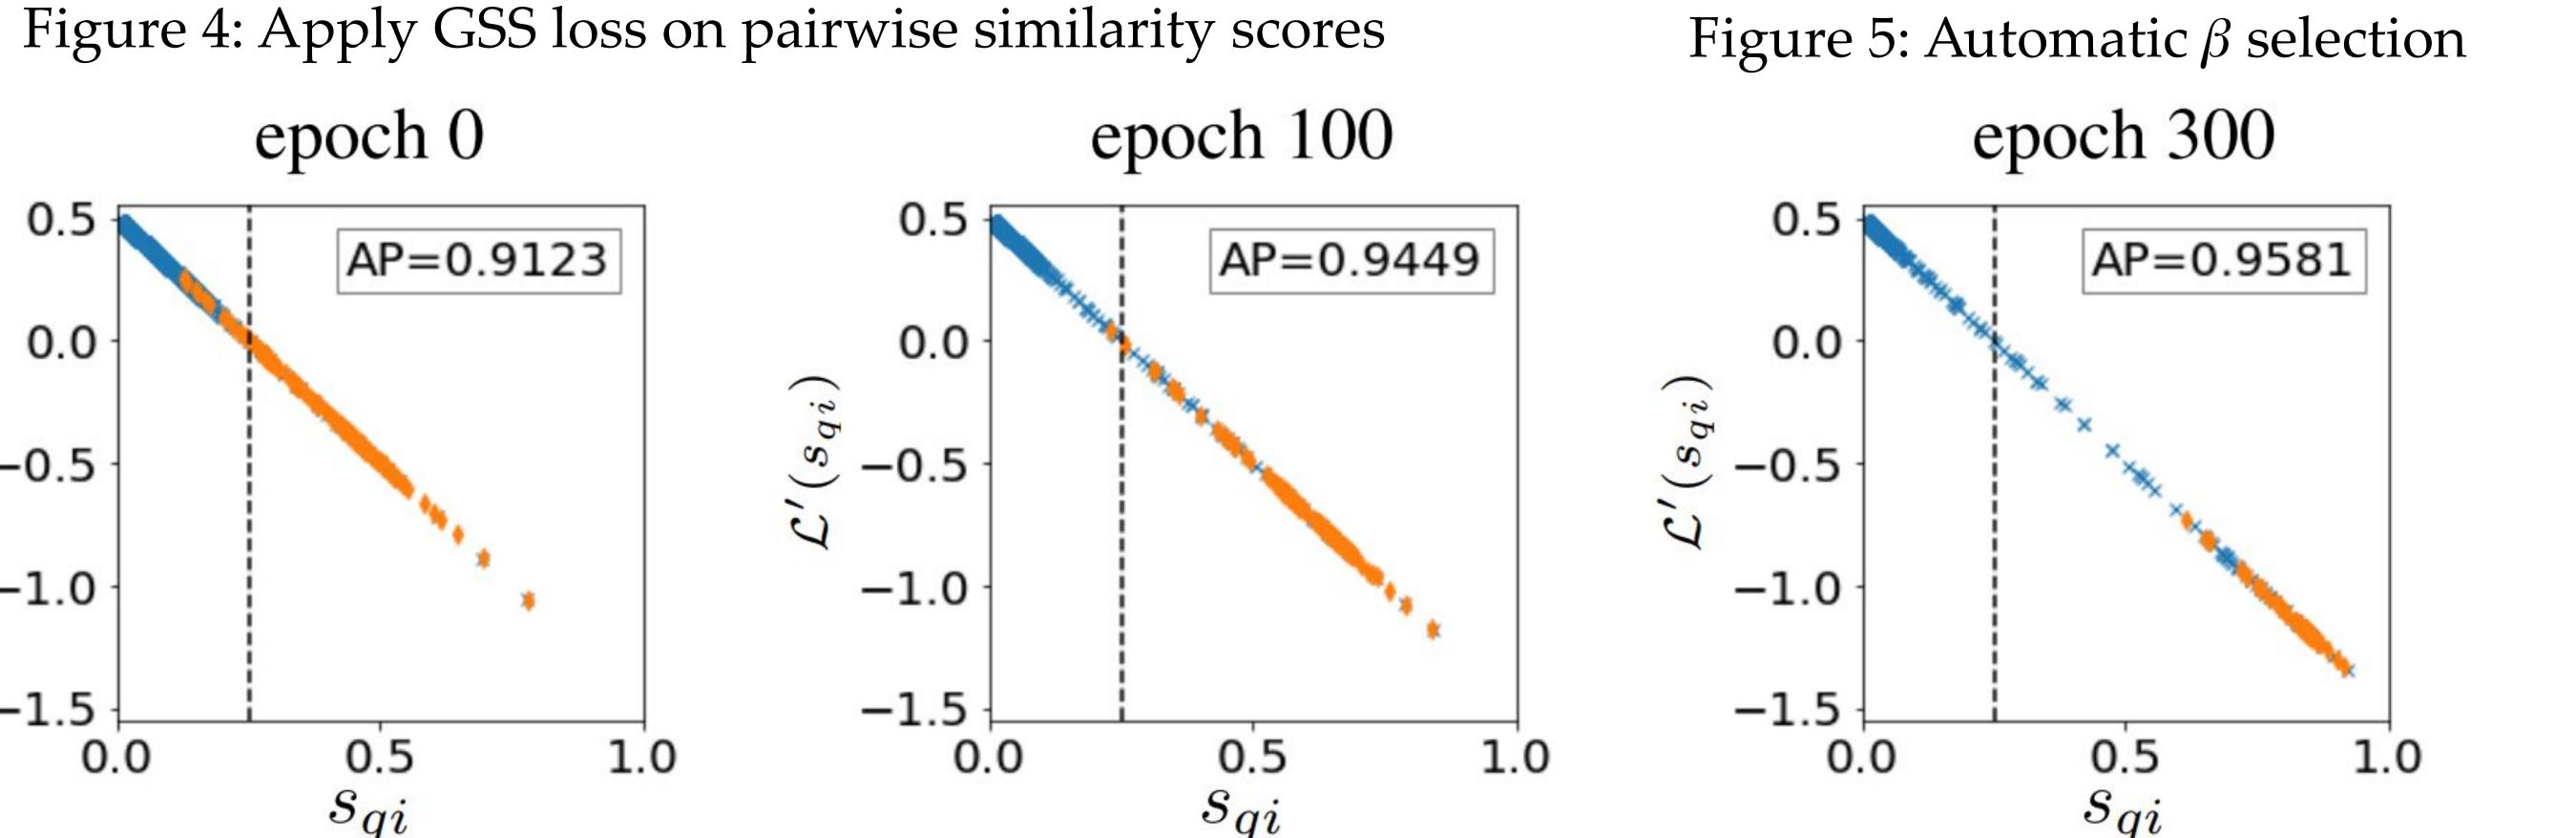


Figure 6: GSS gradients against the similarity scores between a query image and index images.

Spatial Verification

- ➤ The edges in the adjacency matrix are verified by spatial verification.
- Once our model encodes information from the refined adjacency matrix into new index descriptors, the benefit of spatial verification can be effectively preserved without explicitly applying it to query during inference.

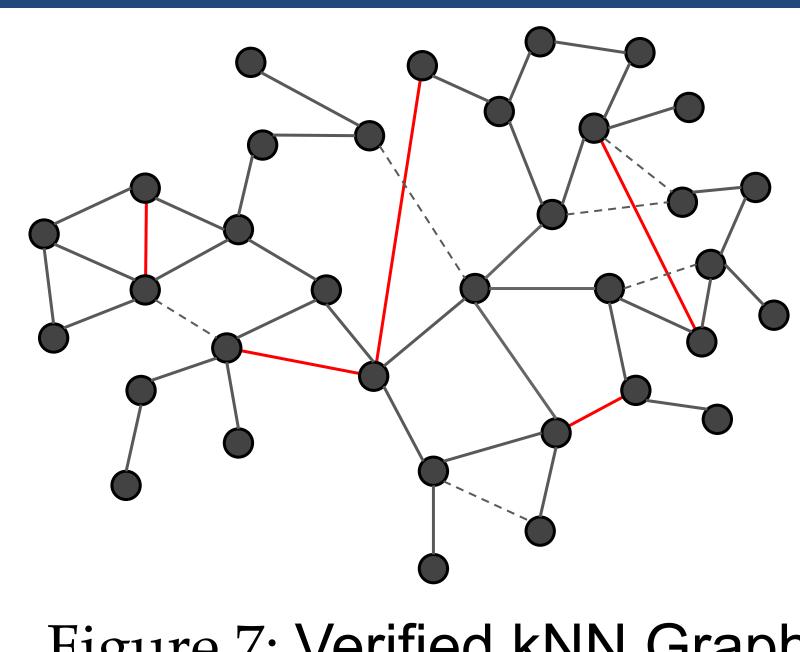


Figure 7: Verified kNN Graph

Results

	mAP				
Method	INSTRE	ROxford		\mathcal{R} Paris	
		Medium	Hard	Medium	Hard
GeM [33]	69.1	64.7	38.5	77.2	56.3
GeM+aQE [33]	74.6	67.2	40.8	80.7	61.8
GeM+DFS [19]	81.1	69.8	40.5	88.9	78.5
GeM+FSR [17]	78.2	70.7	42.2	88.7	78.0
GeM+DFS-FSR [18]	77.9	70.5	40.3	88.7	78.1
GeM+IME [45]	82.3	70.4	45.6	85.0	68.7
GeM+DSM [36]	_	65.3	39.2	77.4	56.2
GeM+DSM [36]+DFS	_	75.0	46.2	89.3	79.3
GeM+GSS	89.2	77.8	57.5	92.4	83.5
Spatial Verification					
GeM+aQE+DELF [28]-SV	87.4	77.2	54.9	88.9	74.8
GeM+DFS+HessAff-ASMK [40]-SV	_	79.1	52.7	91.0	81.0
HessAffNet-HardNet++ [27]+ HQE [20]-SV	1 -	75.2	53.3	73.1	48.9
$GeM+GSS_{\mathcal{V}}$	90.5	79.1	62.2	93.4	85.3
GeM+GSS $_{V}$ -SV	92.4	80.6	64.7	93.4	85.3

Table 1: Retrieval results on INSTRE, ROxford and RParis (Medium and Hard).

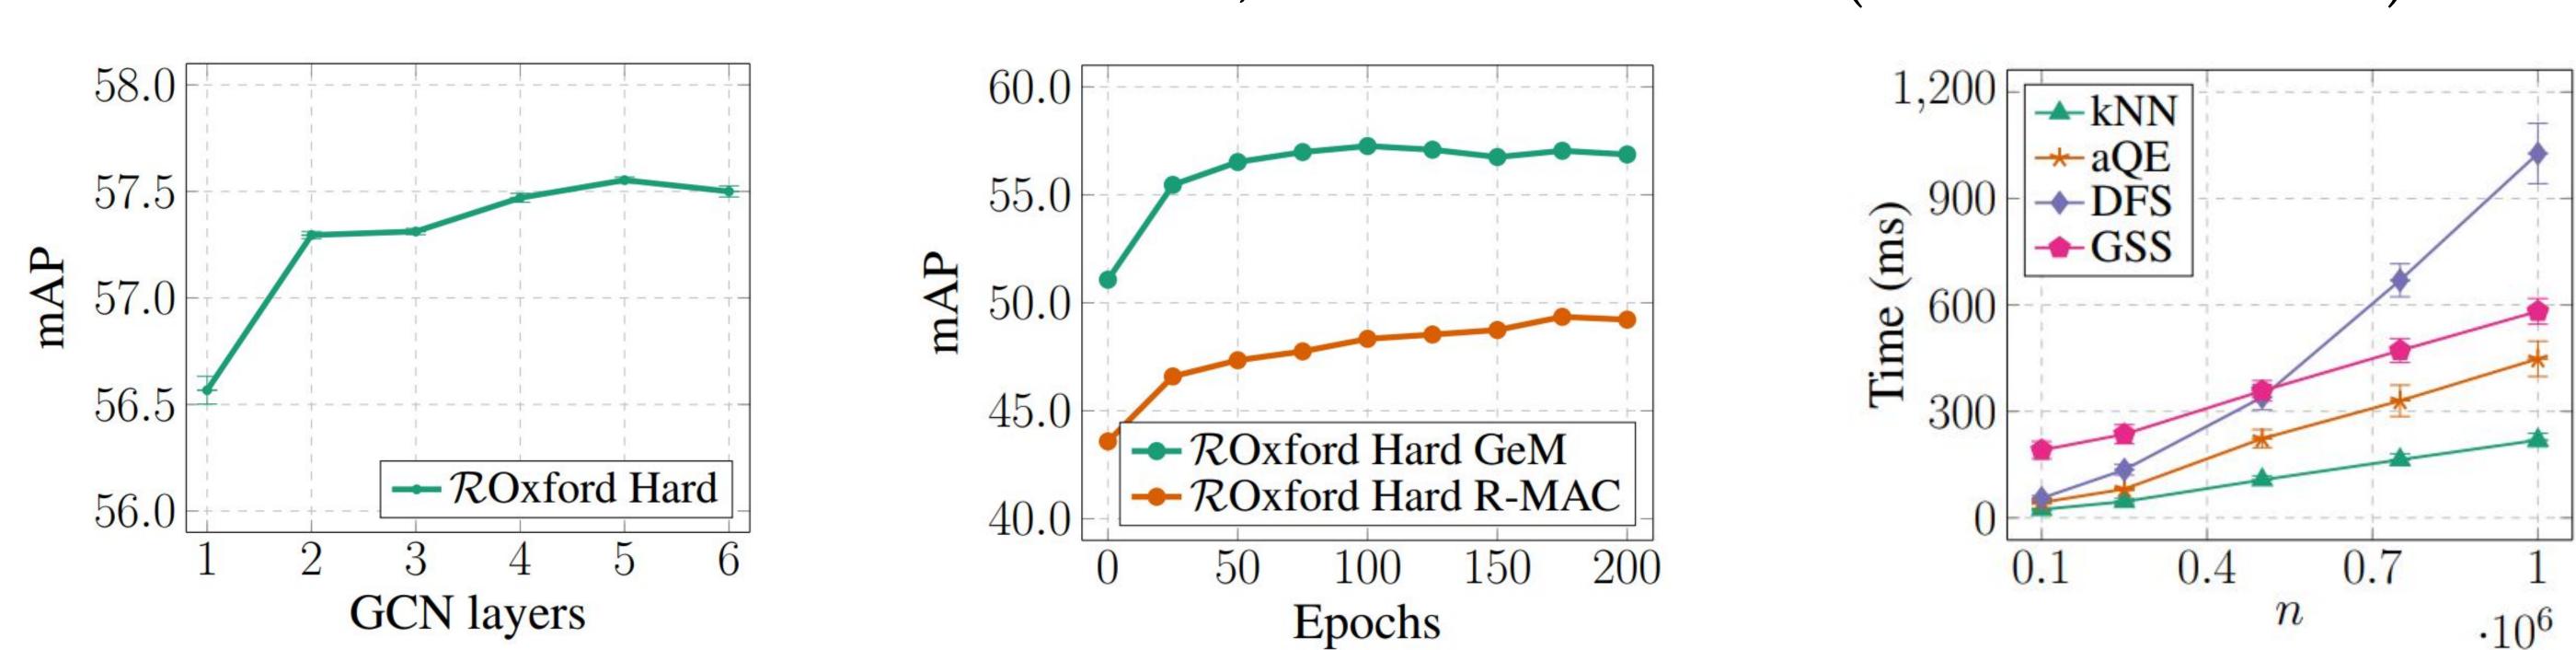


Figure 8: Multiple layers, training curve and run-time analysis

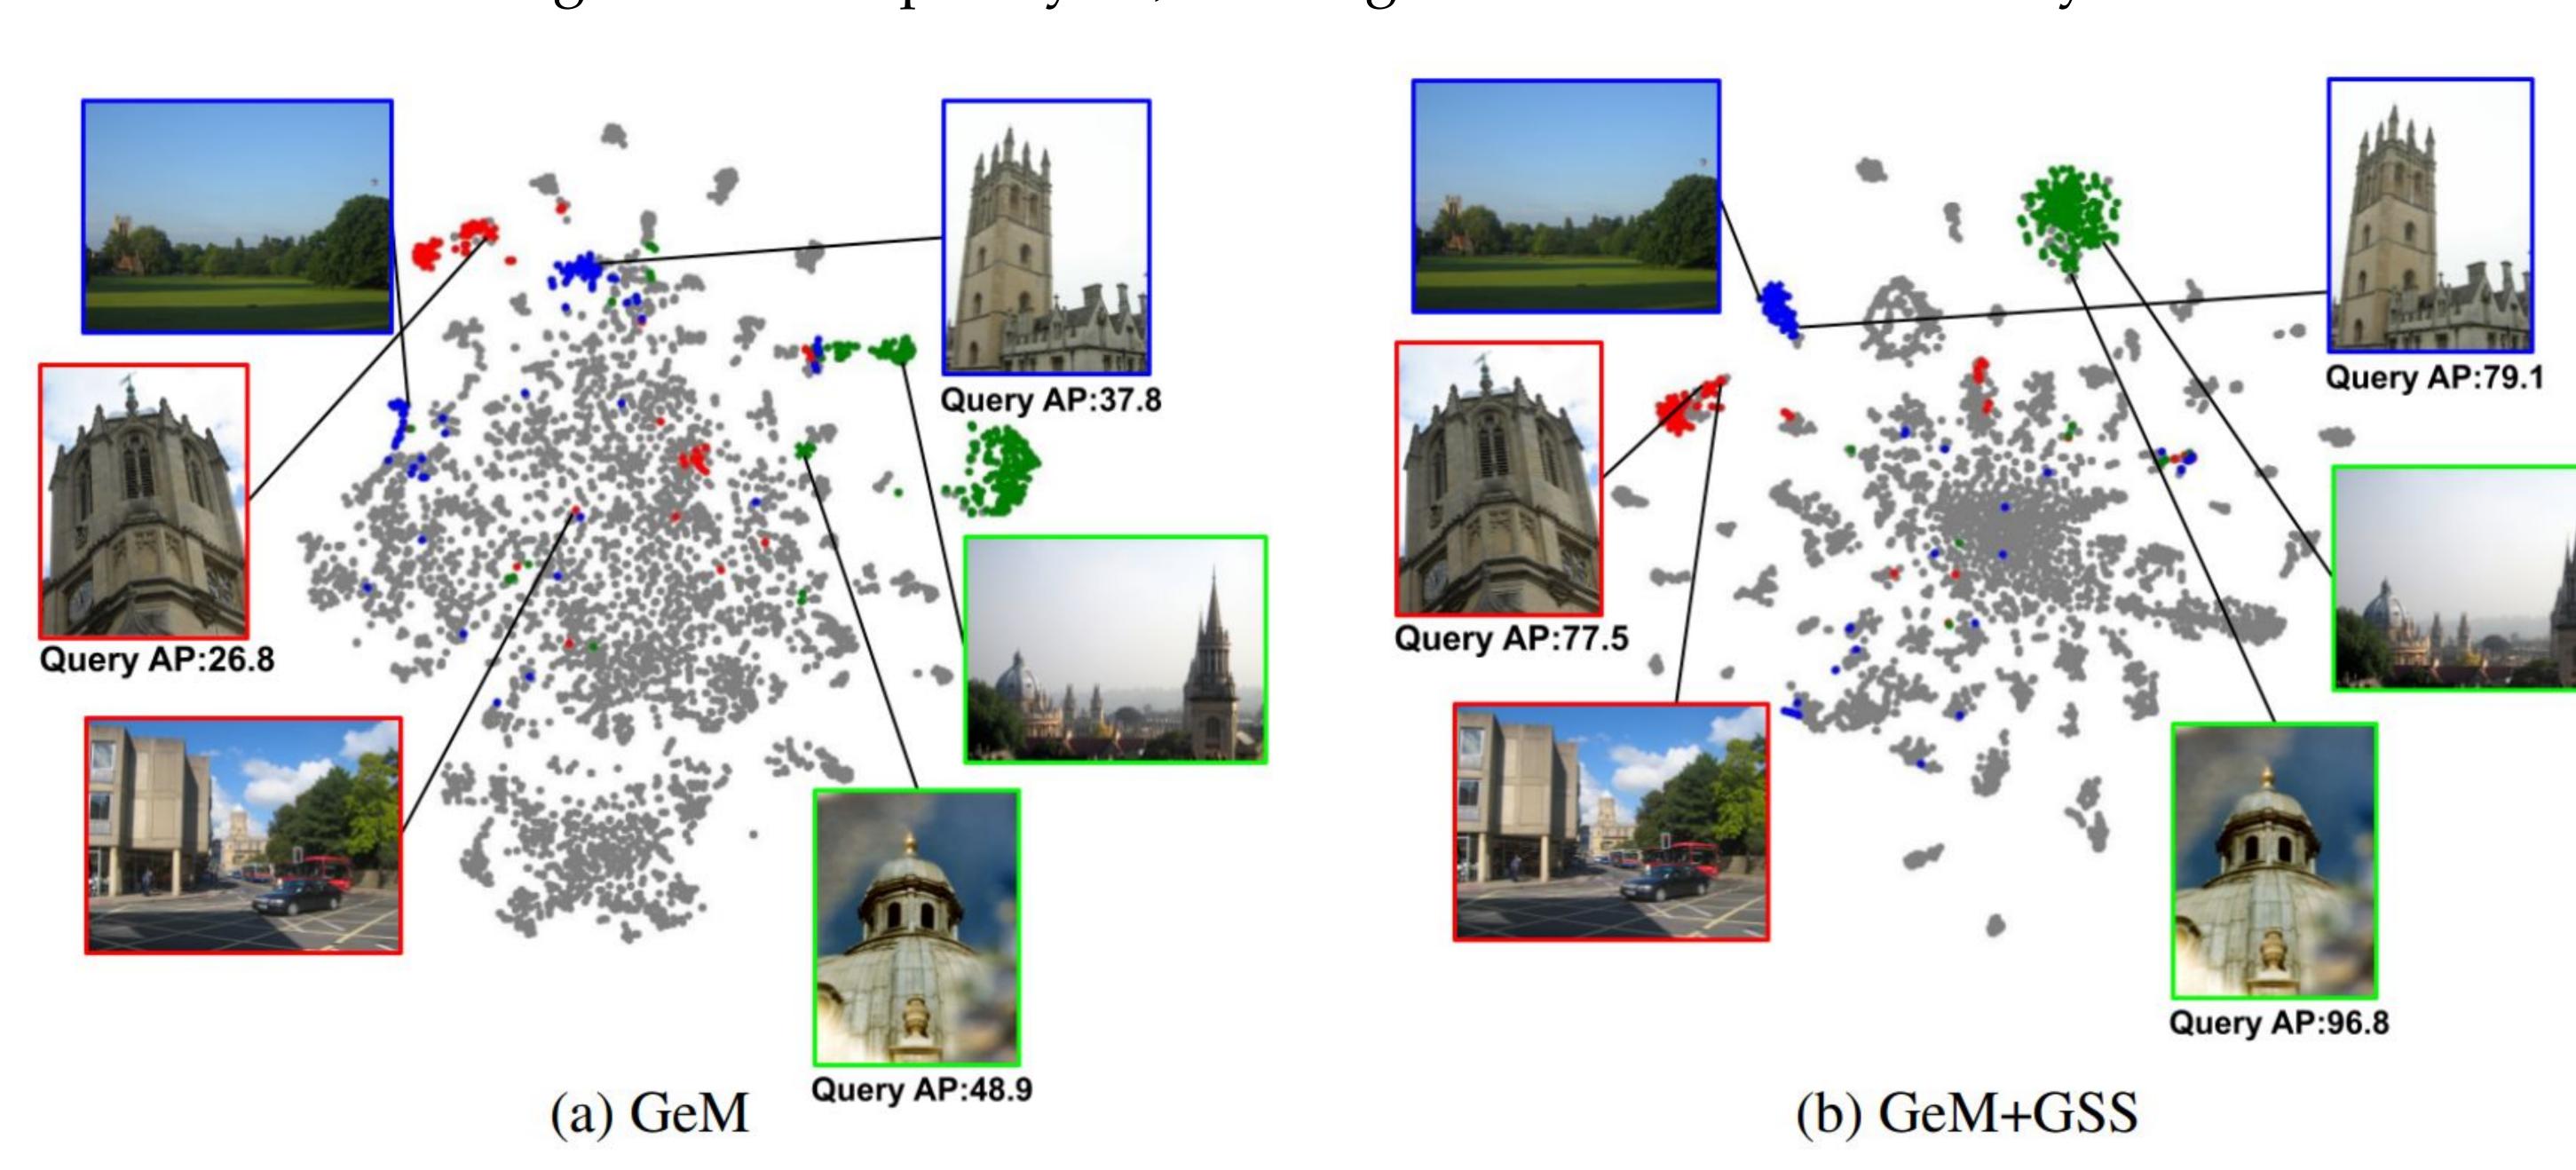


Figure 9: Qualitative analysis on ROxford. GeM and GeM+GSS descriptors are plotted using PCA followed by t-SNE projection to two dimensions.