Efficient Nearest Neighbors Search for Large-Scale Landmark Recognition

Federico Magliani

IMP lab - Department of Engineering and Architecture
University of Parma

Contact Information:
Parco Area delle Scienze, 181/A
43124 Parma, Italy
Phone: +39 0521/905785
Email: federico.magliani@studenti.unipr.it
Website: implab.ce.unipr.it

Introduction

Recent researches on landmark recognition achieved excellent results in small-scale datasets, but when dealing with large-scale retrieval, issues that were irrelevant with small data (e.g., computational time), quickly become fundamental for an efficient retrieval.

Bag of Indexes (BoI)

BoI is a form of multi-index hashing method [4] for the resolution of Approximate Nearest Neighbors (ANN) search problem and consists in the following steps:

1. creation of L hash tables composed by 2^δ buckets;
2. projection of descriptors using hashing functions (L times);
3. saving each index of descriptors in the corresponding bucket.

During the retrieval, different steps are executed for each query:

1. a BoI structure is created, i.e. a vector of n weights (each corresponding to one image of the database) initiliazed to zero;
2. for each of the L hash tables, weights are assigned to the retrieved indexes in the BoI structure;
3. the elements of the BoI structure with the highest weights are re-ranked according to their Euclidean distance from the query.

Weighing metric

Through the application of Locality Sensitive Hashing (LSH), the indexes are searched only in query buckets. Instead, using Multi-probe LSH, also buckets neighbors to the query buckets are checked. For this reason, the weights for the retrieved indexes are chosen as follows:

$$w(i,q,l) = \begin{cases} \frac{1}{2^{H(i,q)}} & \text{if } H(i,q) \leq l \\ 0 & \text{otherwise} \end{cases}$$

where i is a generic bucket, q is the query bucket and $H(i,q)$ is the Hamming distance between i and q.

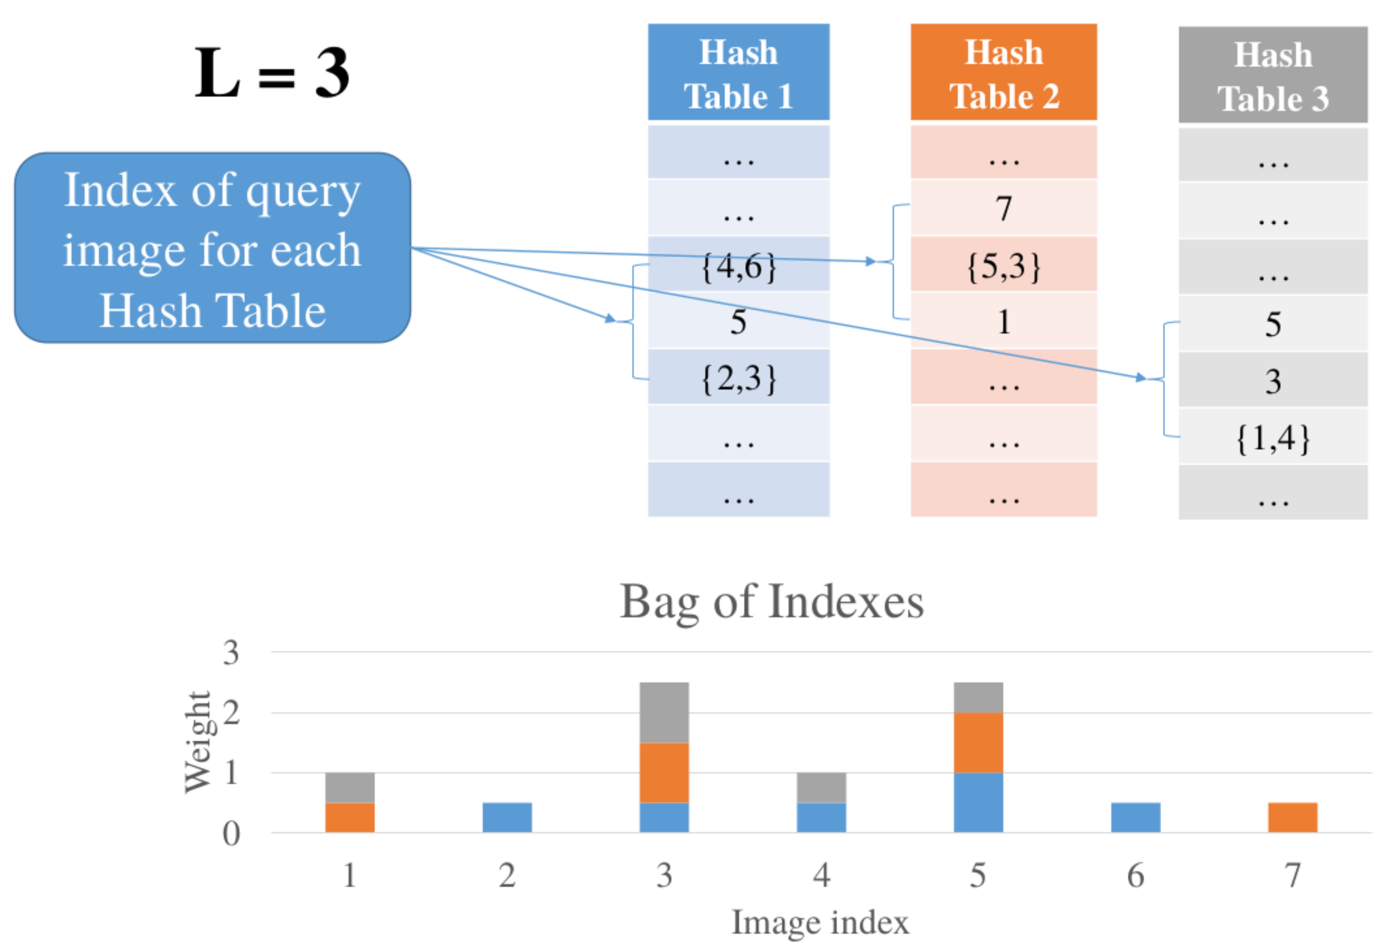

Figure 1: Overview of the retrieval through BoI multi-probe LSH.

BoI adaptive multi-probe LSH

BoI multi-probe LSH approach has the drawback of increasing the computational time since it also needs to search in neighboring buckets. To mitigate this drawback, we introduce BoI adaptive multi-probe LSH, that iteratively refine the search bucket space, by starting with a large number of neighboring buckets γ_0 and slowly reduce it when the number of hash tables increases.

Symbol	Definition	Chosen value	
		Holidays+Flickr1M	SIFT1M
n	number of images	1M	1M
δ	hash dimension	$2^8 = 256$	$2^8 = 256$
L	number of hash tables	100	50
γ_0	initial gap	68	68
l	neighbors bucket	3-neighbors	3-neighbors
ε	size of re-ranking list	250	500
-	reduction	sublinear	sublinear

Table 1: Summary of notation and values used in the experiments.

Results on Holidays+Flickr1M and SIFT1M

Method	Holidays+Flickr1M	
	mAP	avg retrieval time
PP-index [1] *	82.70%	2 844 msec
LOPQ [2]	36.37%	4 msec
FLANN [3]	83.97%	995 msec
BoI adaptive multi-probe LSH	**86.69%**	**8 msec**

Table 2: Results on Holidays+Flickr1M. * indicates our re-implementation.

Method	R=1	R=10	R=100	avg retrieval time
PP-index [1] *	94.32%	94.98%	94.98%	16 999 msec
LOPQ [2]	36.34%	80.11%	96.18%	104 msec
FLANN [3]	95.06%	95.86%	95.86%	31 msec
BoI adaptive multi-probe LSH	**96.80%**	**97.44%**	**97.44%**	**28 msec**

Table 3: Results on SIFT1M. * indicates our re-implementation.

Conclusions

The advantages of the BoI approach are *i)* the possibility to use it in combination with different projection functions, and *ii)* that BoI adaptive multi-probe LSH allows to drastically reduce the query time, outperforming the accuracy results compared to the state-of-the-art methods for large-scale landmark recognition.

References

[1] Chavez, Figueroa, and Navarro. Effective proximity retrieval by ordering permutations. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 2008.

[2] Kalantidis and Avrithis. Locally optimized product quantization for approximate nearest neighbor search. *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.

[3] Muja and Lowe. Scalable nearest neighbor algorithms for high dimensional data. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 2014.

[4] Norouzi, Punjani, and Fleet. Fast search in hamming space with multi-index hashing. *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.