

Bayes Merging of Multiple Vocabularies for Scalable Image Retrieval

Liang Zheng¹, Shengjin Wang¹, Wengang Zhou², and Qi Tian³

¹State Key Laboratory of Intelligent Technology and Systems;

¹Tsinghua National Laboratory for Information Science and Technology;

¹Department of Electronic Engineering, Tsinghua University, Beijing 100084, China

²EEIS Department, University of Science and Technology of China

³University of Texas at San Antonio, TX, 78249, USA

zheng-106@mails.tsinghua.edu.cn wsgsj@tsinghua.edu.cn zhwg@ustc.edu.cn qitian@cs.utsa.edu

Abstract

In the Bag-of-Words (BoW) model, the vocabulary is of key importance. Typically, multiple vocabularies are generated to correct quantization artifacts and improve recall. However, this routine is corrupted by vocabulary correlation, i.e., overlapping among different vocabularies. Vocabulary correlation leads to an over-counting of the indexed features in the overlapped area, or the intersection set, thus compromising the retrieval accuracy. In order to address the correlation problem while preserve the benefit of high recall, this paper proposes a Bayes merging approach to down-weight the indexed features in the intersection set. Through explicitly modeling the correlation problem in a probabilistic view, a joint similarity on both image and feature-level is estimated for the indexed features in the intersection set.

We evaluate our method on three benchmark datasets. Albeit simple, Bayes merging can be well applied in various merging tasks, and consistently improves the baselines on multi-vocabulary merging. Moreover, Bayes merging is efficient in terms of both time and memory cost, and yields competitive performance with the state-of-the-art methods.

1. Introduction

This paper considers the task of Bag-of-Words (BoW) based image retrieval, especially on multi-vocabulary merging. We aim at improving the retrieval accuracy while maintaining affordable memory and time cost.

The vocabulary (also called the codebook or quantizer) lies at the core of the BoW based image retrieval system. It functions by quantizing SIFT descriptors [9] to discrete visual words. The quantized visual words are the nearest

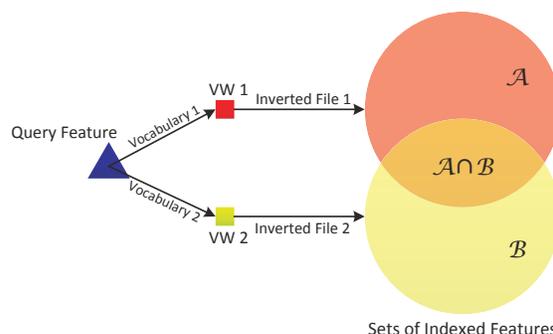


Figure 1. An Illustration of vocabulary correlation. Given a query feature, it is quantized to two visual words in two vocabularies. Then, two sets of indexed features, \mathcal{A} and \mathcal{B} , are identified from the two inverted files, respectively. The area of the intersection set $\mathcal{A} \cap \mathcal{B}$ (denoted as $Card(\mathcal{A} \cap \mathcal{B})$) encodes the extent of correlation between the two sets. In this paper, we focus on the indexed features in $\mathcal{A} \cap \mathcal{B}$.

centers to the feature vectors in the feature space. In order to reduce quantization error and improve recall, multiple vocabularies are often generated, and each feature is quantized to different visual words from multiple vocabularies. The primary benefit of using multiple vocabularies is that more candidate features are recalled, which corrects quantization artifacts to some extent.

However, the routine of multi-vocabulary merging is affected by a crucial problem, i.e., vocabulary correlation [3] (see Fig. 1). Given a query feature, based on the inverted files with two individual vocabularies, two sets of indexed features \mathcal{A} and \mathcal{B} are identified, sharing an intersection set $\mathcal{A} \cap \mathcal{B}$. In this paper, the area of $\mathcal{A} \cap \mathcal{B}$ is approximated by $Card(\mathcal{A} \cap \mathcal{B})$. The larger $Card(\mathcal{A} \cap \mathcal{B})$ is, the larger the correlation will be. In an extreme case, total correlation occurs if $Card(\mathcal{A} \cap \mathcal{B}) = Card(\mathcal{A} \cup \mathcal{B})$, and merging \mathcal{A} and \mathcal{B} brings no benefit.

A straightforward method for multi-vocabulary merging consists in concatenating the BoW histograms of different vocabularies [11]. In a microscopic view of this method, the indexed features in $\mathcal{A} \cap \mathcal{B}$ are counted twice in Fig. 1. Nevertheless, since images in this area are mostly irrelevant ones (the number of relevant images is always very small), the over-counting may actually compromise the retrieval accuracy [12].

In this paper, we consider the situation in which the given vocabularies are correlated, and we aim to reduce the impact of correlation. To address this problem, this paper proposes to model the vocabulary correlation problem from a probabilistic view. In a nutshell, we jointly estimate an image- and feature-level similarity for the indexed features in the intersection set (or overlapping area). Given a query feature, lists of indexed features are extracted from multiple inverted files. Then, we identify the intersection and union sets of the lists, from which the cardinality ratio is calculated. This ratio thus encodes the extent of correlation (see Fig. 1). For the indexed images in the intersection set, its similarity with the query is estimated as a function of the cardinality ratio, and subsequently added to the matching score. Experiments on several benchmark datasets demonstrate that Bayes merging is effective, and yields competitive results with the state-of-the-art methods.

2. Related Work

Vocabulary Generation The vocabulary provides a discrete partitioning of the feature space by visual words. Typically, either flat kmeans [12, 4] or hierarchical kmeans [11] is employed to train a vocabulary in an unsupervised manner. Improved methods include incorporating contextual information into the vocabulary [19], building super-sized vocabulary [16, 20, 10], making use of the active points [15], etc.

Matching Refinement Feature-to-feature matching is a key issue in the BoW model. The baseline approach employs a coarse word-to-word matching, resulting in undesirable low precision. To improve precision, some works analyze the spatial contexts [16, 21, 24] of SIFT features, and use the spatial constraints as solution to refining matching. Another line of works extracts binary signatures from SIFT descriptors [4] or its contexts [23, 8]. The feature matching is thus refined by a further check of the Hamming distance between binary signatures. In this paper, however, we argue that even if two features are adjacent in the feature space, the corresponding images are probably very different. Therefore, we are supposed to look one step further by estimating a joint similarity on both image- and feature-level from clues in multiple vocabularies.

Multiple Vocabularies It is well known that multi-vocabulary merging is effective in improving recall [3, 18]. Typically, multi-vocabulary merging can be performed ei-

ther at score level, *e.g.*, by concatenating the BoW histograms [11], or at rank level, *e.g.*, by rank aggregation [7]. On the other hand, some works also provide clues that multiple vocabularies also improve precision [1, 17]. To address the problem of vocabulary correlation, Xia *et al.* [18] propose to create the vocabularies jointly and reduce correlation from the view of vocabulary generation. A more relevant work includes [3], which uses PCA to implicitly remove correlation of given vocabularies, resulting in a low dimensional image representation. Our work departs from previous works in two aspects. First, we explicitly model the vocabulary correlation problem from a probabilistic view. Second, our work is proposed for the BoW based image retrieval task, which differs from NN search problems.

3. Background

3.1. Notations

Assume that the K vocabularies are denoted as $\mathcal{V}^{(k)} = \{v_1^{(k)}, v_2^{(k)}, \dots, v_{s_k}^{(k)}\}$, $k = 1, \dots, K$, where $v_i^{(k)}$ represents a visual word and s_k is the vocabulary size. Correspondingly, built on $\mathcal{V}^{(k)}$, K inverted files are organized as $\mathcal{W}^{(k)} = \{W_1^{(k)}, W_2^{(k)}, \dots, W_{s_k}^{(k)}\}$, $k = 1, \dots, K$, where each entry $W_i^{(k)}$ contains a list of indexed features.

Given a query SIFT feature x , it is quantized to a visual word tuple $(v^{(1)}, v^{(2)}, \dots, v^{(K)})$, where $v^{(k)}$, $k = 1, \dots, K$ is the nearest centroid in $\mathcal{V}^{(k)}$ to x . With the K visual words we can identify K sets of indexed features in entries $\{W_{i_k}^{(k)}\}_{k=1}^K$. From the K sets, we can define three types of sets to be used in this paper.

Definition 1 (k^{th} -order intersection set) *The intersection set of k , and only k sets, denoted as $\cap^{(k)}$, $k \geq 2$.*

Definition 2 (k^{th} -order union set) *The union set of k , and only k sets, denoted as $\cup^{(k)}$, $k \geq 2$.*

Definition 3 (difference set) *The set in which no overlapping exists, i.e., $\cup^{(K)} - \sum_{k=2}^K \cap^{(k)}$, $K \geq 2$.*

3.2. Baselines

Single vocabulary baseline (B₀) For a single vocabulary, we adopt the baseline introduced in [12, 4]. Specifically, vocabularies are trained by AKM on the independent Flickr60K data [4], and average IDF [22] weighting scheme is used. We replace the original SIFT descriptor with root-SIFT [13]. In this scenario, we denote the matching function between two features x and y as,

$$f_0(x, y) = \delta_{v_x, v_y} \quad (1)$$

where v_x and v_y are visual words of x and y in the vocabulary, respectively, and $\delta(\cdot)$ is the Kronecker delta response.

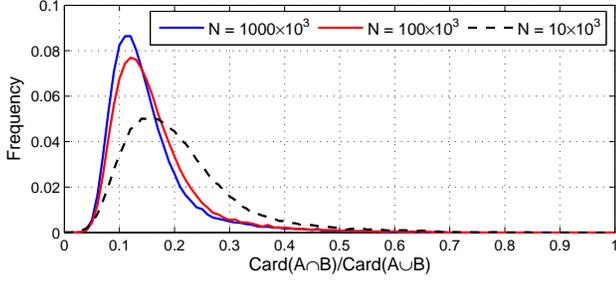


Figure 2. The distribution of the cardinality ratio on 10K, 100K, 1000K images, respectively. We use two vocabularies of size 20K.

Conventional vocabulary merging (B₁) Given K vocabularies, B_1 simply concatenates multiple BoW histograms [11]. It is equivalent to a simple score-level addition of the outputs of multiple vocabularies. The matching function between features x and y can be defined as

$$f_1(x, y) = \sum_{k=1}^K \delta_{v_x^{(k)}, v_y^{(k)}} \quad (2)$$

where $v_x^{(k)}$ and $v_y^{(k)}$ are visual words in vocabulary $\mathcal{V}^{(k)}$ for x and y , respectively. Eq. 2 shows that in baseline B_1 , an indexed feature is counted k times if it is in the k^{th} -order intersection set $\cap^{(k)}$, and only once if in the difference set (since there is no overlapping).

Multi-index based vocabulary merging (B₂) In [1], a multi-index is organized as a multi-dimensional structure. In its nature, given K vocabularies, two features are considered as a match *iff* they are in the K^{th} -order intersection set $\cap^{(K)}$ of the indexed feature lists. Therefore, in baseline B_2 , the matching function is defined as

$$f_2(x, y) = \prod_{k=1}^K \delta_{v_x^{(k)}, v_y^{(k)}} \quad (3)$$

Eq. 3 only counts the indexed features in $\cap^{(K)}$, discarding the rest. Therefore, the recall is low for B_2 .

4. Proposed Method

For multi-vocabulary merging, the major problem is the over-counting of the intersection sets $\cap^{(k)}$, $k = 1, \dots, K$. On the other hand, the major benefit is a high recall, which is encoded in the difference set. Taking both issues into consideration, we propose to exert a likelihood on the intersection sets and preserve the difference set (scored as B_1). Without loss of generality, we start from the case of two vocabularies and then generalize it to multiple vocabularies.

4.1. Model Formulation

Given a query feature x in image Q , two sets of indexed features \mathcal{A} and \mathcal{B} are identified in two inverted files, respectively. Here, we want to evaluate the likelihood that a SIFT

feature y is a true neighbor of x given that y belongs to the intersection set of \mathcal{A} and \mathcal{B} . This likelihood can be modeled as the following conditional probability,

$$w(x, y) = p(y \in T_x \mid y \in \mathcal{A} \cap \mathcal{B}). \quad (4)$$

In Eq. 4, we define T_x as the set of features which are visually similar to x (locally) *and* belong to the ground truth images of Q (globally). On the other hand, F_x is defined as the features which violate any of the two criteria. Therefore, T_x and F_x satisfy the follows

$$p(y \in T_x) + p(y \in F_x) = 1. \quad (5)$$

For simplicity, we denote $y \in \mathcal{A} \cap \mathcal{B}$ as $\mathcal{A} \cap \mathcal{B}$, $y \in T_x$ as T_x , and $y \in F_x$ as F_x . Then, using the formula of Bayes' theorem as well as Eq. 5, we get

$$\begin{aligned} p(T_x \mid \mathcal{A} \cap \mathcal{B}) &= \frac{p(\mathcal{A} \cap \mathcal{B} \mid T_x) \cdot p(T_x)}{p(\mathcal{A} \cap \mathcal{B})} \\ &= \frac{p(\mathcal{A} \cap \mathcal{B} \mid T_x) \cdot p(T_x)}{p(\mathcal{A} \cap \mathcal{B} \mid T_x) \cdot p(T_x) + p(\mathcal{A} \cap \mathcal{B} \mid F_x) \cdot p(F_x)}. \end{aligned} \quad (6)$$

Then, re-formulating Eq. 6, we have

$$p(T_x \mid \mathcal{A} \cap \mathcal{B}) = \left(1 + \frac{p(\mathcal{A} \cap \mathcal{B} \mid F_x)}{p(\mathcal{A} \cap \mathcal{B} \mid T_x)} \cdot \frac{p(F_x)}{p(T_x)} \right)^{-1}. \quad (7)$$

In Eq. 7, there are actually three random variables to estimate, *i.e.*, $p(\mathcal{A} \cap \mathcal{B} \mid F_x)$ (**term 1**), $p(\mathcal{A} \cap \mathcal{B} \mid T_x)$ (**term 2**), and $p(F_x)/p(T_x)$ (**term 3**). In Section 4.2, we will exploit the estimation of these probabilities.

4.2. Probability Estimation

Estimation of term 1 In Eq. 7, the term $p(\mathcal{A} \cap \mathcal{B} \mid F_x)$ encodes the probability that feature y lies in the set $\mathcal{A} \cap \mathcal{B}$ given that y is a false match of query feature x . In this case, we should consider the distribution of the x ' false matches in sets \mathcal{A} and \mathcal{B} . In large databases, the number of true matches (both locally and globally) is limited. In other words, false matches dominate the space covered by \mathcal{A} and \mathcal{B} . Therefore, we assume that false matches are uniformly distributed in \mathcal{A} and \mathcal{B} , and term 1 can be estimated as

$$p(\mathcal{A} \cap \mathcal{B} \mid F_x) = \frac{\text{Card}(\mathcal{A} \cap \mathcal{B})}{\text{Card}(\mathcal{A} \cup \mathcal{B})}, \quad (8)$$

where $\text{Card}(\cdot)$ represents the cardinality of a set. Eq. 8 implies that, the probability that a false match falls into $\mathcal{A} \cap \mathcal{B}$ is proportional to the cardinality ratio $\frac{\text{Card}(\mathcal{A} \cap \mathcal{B})}{\text{Card}(\mathcal{A} \cup \mathcal{B})}$. Intuitively, the larger the intersection set is, the more probable that a false match will fall into it. Fig. 2 depicts the distribution of this cardinality ratio on different database scales.

Estimation of term 2 In contrast to term 1, the probability encoded in term 2 reflects the likelihood that y , a true neighbor of query x , falls into the intersection set $\mathcal{A} \cap \mathcal{B}$.

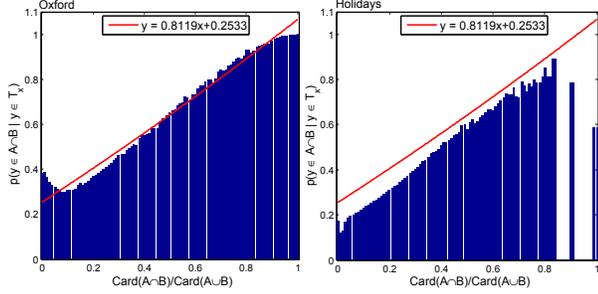


Figure 3. Distribution of term 2 as a function of $\frac{Card(\mathcal{A} \cap \mathcal{B})}{Card(\mathcal{A} \cup \mathcal{B})}$ on Oxford (left) and Holidays (right) datasets. Least Square Fitting of degree 1 is performed on Oxford, plotted as the red line. We find that the same line also fits the trends of Holidays dataset.

Still, we estimate this probability as a function of the cardinality ratio $\frac{Card(\mathcal{A} \cap \mathcal{B})}{Card(\mathcal{A} \cup \mathcal{B})}$. However, since the number of true matches is very small compared to false ones, we do not adopt the method in estimating term 1. Instead, image data with ground truth is used to analyze the distribution.

Specifically, empirical analysis is performed on Oxford and Holidays datasets. Given a feature x in the query image Q , true matches are defined as the features which have a Hamming distance [4] smaller than 20 to x and which appear in the ground truth images of Q . Then we calculate the ratio of the number of true matches in $\mathcal{A} \cap \mathcal{B}$ to the number of true matches in $\mathcal{A} \cup \mathcal{B}$. Finally, the relationship between the ratio and $\frac{Card(\mathcal{A} \cap \mathcal{B})}{Card(\mathcal{A} \cup \mathcal{B})}$ is depicted in Fig. 3.

A surprising fact from Fig. 3 is that $p(\mathcal{A} \cap \mathcal{B} | F_x)$ increases linearly with $\frac{Card(\mathcal{A} \cap \mathcal{B})}{Card(\mathcal{A} \cup \mathcal{B})}$. Contrary to our expectation, true matches do not aggregate around the query point. Instead, they tend to scatter in the high-dimensional feature space. Otherwise, the curves in Fig. 3 would take on a $\log(\cdot)$ -like profile. On the other hand, Fig. 3 also implies that the indexed features in $\mathcal{A} \cap \mathcal{B}$ are mostly false matches. This explains why the over-counting compromises the retrieval accuracy. Moreover, we also find that the trend in Fig. 3 seems to be database-independent.

Estimation of term 3 Term 3, *i.e.*, $p(F_x)/p(T_x)$, can be interpreted as the ratio of the probability of y being a false match to y being a true match. Typically, as the database grows, the number of false images will become larger, and the value of term 3 will increase. To model this property, and thus making our system adjustable to large scale settings, we set term 3 as

$$\frac{p(F_x)}{p(T_x)} = \log(N \cdot c), \quad (9)$$

where N is the number of images in the database, and c is a weighting parameter. Note that we add a $\log(\cdot)$ operator due to numerical considerations.

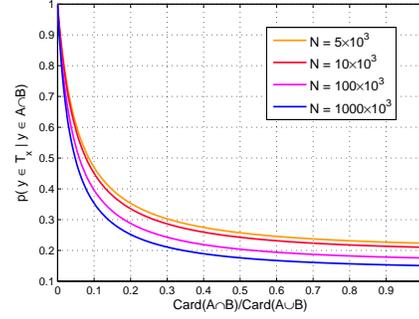


Figure 4. The estimation of Eq. 4 as a function of $\frac{Card(\mathcal{A} \cap \mathcal{B})}{Card(\mathcal{A} \cup \mathcal{B})}$. Four curves are presented, corresponding to $N = 5\text{K}$, 10K , 100K , and 1M , respectively. The vocabulary sizes are both 20K .

4.3. Similarity Interpretation

Using the estimation methods introduced in Section. 4.2, we are able to provide an explicit implementation of the probability model (Eq. 4). Specifically, we assume four database sizes are involved, *i.e.*, 5K , 10K , 100K , 1M , and we set the parameter c to 1 for better illustration. The derived probability function is plotted against $\frac{Card(\mathcal{A} \cap \mathcal{B})}{Card(\mathcal{A} \cup \mathcal{B})}$ in Fig. 4. From the curves in Fig. 4, we can get several implications in terms of physical interpretation.

First, when the intersection area is very small (the cardinality ratio is close to zero), it is very likely that y is a true match if it falls into this area. In this scenario, the discriminative power of the intersection set is high, and can be trusted when merging vocabularies.

Second, when the cardinality ratio approaches 1, *i.e.*, sets \mathcal{A} and \mathcal{B} share a large overlap, the probability of y being a true match is small. This makes more sense if we take into consideration the fact that false images dominate the entire feature space. Moreover, a larger intersection means a larger dependency (or correlation) between two vocabularies, in which situation our method exerts a punishment (low weight) and overcomes this problem to some extent.

Third, as the database becomes larger, the curves lean towards the origin. In fact, for large databases, the chances that y is a true match will be more remote under each cardinality ratio. Nevertheless, the cardinality ratio tends to get smaller (see Fig. 2) as the database grows, so the estimated probability will be compensated to some extent.

As a summary, Fig. 4 reveals some interesting properties of $\mathcal{A} \cap \mathcal{B}$. The formula Eq. 4 will be adopted into the BoW-based image retrieval framework in Section. 4.5.

4.4. Generalization to Multiple Vocabularies

In this section, we generalize our method to the case of multiple vocabularies ($K \geq 2$).

Given K vocabularies, a query feature x is quantized to K visual words, and subsequently K sets of indexed features are identified, *i.e.*, $\{\mathcal{A}_i\}_{i=1}^K$. If a database feature

Algorithm 1 Bayes merging for image retrieval

Input:

The query image Q with L descriptors x_1, x_2, \dots, x_L ;
The K vocabularies $\mathcal{V}^{(1)}, \mathcal{V}^{(2)}, \dots, \mathcal{V}^{(K)}$;
The K inverted files $\mathcal{W}^{(1)}, \mathcal{W}^{(2)}, \dots, \mathcal{W}^{(K)}$;

Iteration:

- 1: **for** $n = 1 : L$ **do**
 - 2: Quantize x_n into K visual words $v^{(1)}, \dots, v^{(K)}$;
 - 3: Identify K lists of indexed features $\mathcal{A}^{(1)}, \dots, \mathcal{A}^{(K)}$;
 - 4: Find all k^{th} -order intersection sets, $k = 2, \dots, K$;
 - 5: Find all k^{th} -order union sets, $k = 2, \dots, K$;
 - 6: **for** each indexed feature in \cup^K **do**
 - 7: Find the k^{th} -order intersection set it falls in;
 - 8: Find the k^{th} -order union set it falls in;
 - 9: Calculate $\frac{Card(\cap^{(k)})}{Card(\cup^{(k)})}$;
 - 10: Calculate matching strength using Eq. 10;
 - 11: Vote for the candidate image using Eq. 11;
 - 12: **end for**
 - 13: **end for**
-

y falls into the k^{th} -order intersection set of $\{\mathcal{A}_i\}_{i=1}^K$, the probability of it being a true match to x is defined as

$$w(x, y) = p(y \in T_x \mid y \in \mathcal{A}_1 \cap \mathcal{A}_2 \cap \dots \cap \mathcal{A}_k). \quad (10)$$

Using the similarity function derived in Section 4.3, we can estimate Eq. 10 as a function of the cardinality ratio of the k^{th} -order intersection and union sets $\frac{Card(\mathcal{A}_1 \cap \mathcal{A}_2 \cap \dots \cap \mathcal{A}_k)}{Card(\mathcal{A}_1 \cup \mathcal{A}_2 \cup \dots \cup \mathcal{A}_k)}$.

4.5. Proposed Image Retrieval Pipeline

In this section, the matching function of the Bayes merging method is defined as follows,

$$f(x, y) = \begin{cases} kw(x, y), & \text{if } y \in \cap^k, k \geq 2 \\ \sum_{i=1}^K \delta_{v_x^{(k)}, v_y^{(k)}}, & \text{otherwise} \end{cases} \quad (11)$$

where $w(x, y)$ is the similarity function defined in Eq. 10. If $w(x, y) = 1$, Bayes merging reduces to the baseline B_1 .

The pipeline of Bayes merging is summarized in Algorithm 1. In the offline steps, K vocabularies are trained and the corresponding K inverted files are organized. During online retrieval, given a query image Q with L descriptors, for each feature x_n , we quantize it to K visual words (step 2). Then, K lists of indexed features are identified (step 3), from which all k^{th} -order intersection and union sets are identified (step 4, 5). For each indexed feature in \cup^K , we find the k^{th} -order intersection and union sets it falls in (step 7, 8), and calculate the cardinality ratio (step 9). Finally, matching strength is calculated according to Eq. 10 and used in the matching function as Eq. 11 (steps 10 and 11).

For one query feature, we have to traverse \cup^K twice in Algorithm 1, which doubles the query time. However, in

Dataset	# images	# queries	# descriptors	Evaluation
<i>Holidays</i>	1491	500	4,455,091	mAP
<i>Oxford</i>	5063	55	13,516,660	mAP
<i>Ukbench</i>	10200	10200	19,415,079	N-S score

Table 1. Details of the datasets in the experiments.

Value of c in Eq. 9	10	20	30	40	50
<i>Oxford</i> , mAP(%)	46.72	46.80	46.77	46.73	46.73
<i>Holidays</i> , mAP(%)	58.35	58.47	58.51	58.50	58.36

Table 2. The impact of parameter c on image retrieval accuracy. Results (mAP in percent) on Oxford 5K and Holidays datasets are presented. We set $c = 30$ from these results.

the **supplementary material**, we demonstrate that we can accomplish this process by traversing \cup^K only once, thus solving the efficiency problem of Bayes merging.

5. Experiments

In this section, the proposed Bayes merging is evaluated on three benchmark datasets, *i.e.*, Holidays [4], Oxford [12], and Ukbench [11]. The details of the datasets are summarized in Table 1. We also add the Flickr 1M dataset [4] of one million images to test the scalability of our method. All the vocabularies are trained independently on the Flickr60K dataset [4] using AKM [12] with different initial seeds.

5.1. Parameter Analysis

One parameter, *i.e.*, the weighting parameter c in Eq. 9 is involved in the probabilistic model. We evaluate c on the Holidays and Oxford datasets, and record in Table 2 the mAP results against different values of c . We can see that the mAP results remain stable when c ranges from 10 to 50, probably due to the effect of the *log* operator in Eq. 9. We therefore set c to 30 in the following experiments.

5.2. Evaluation

Comparison with the baselines We first compare Bayes merging with the baselines, *i.e.*, B_0, B_1, B_2 defined in Section 3.2. The results are demonstrated in Fig. 5 and Fig. 6. From these results we find that baseline B_2 does not benefit from introducing multiple vocabularies, and that its performance drops when merging more vocabularies, because the recall further decreases. We speculate that Multiple Assignment will bring benefit [1, 17] to B_2 . Moreover, baseline B_1 brings limited improvements over B_0 . In fact, B_1 has a higher recall than B_0 , but this benefit is impaired by vocabulary correlation in which many irrelevant images are over-counted.

In comparison, it is clear that Bayes merging yields great improvements. Take Holidays for example, when merging

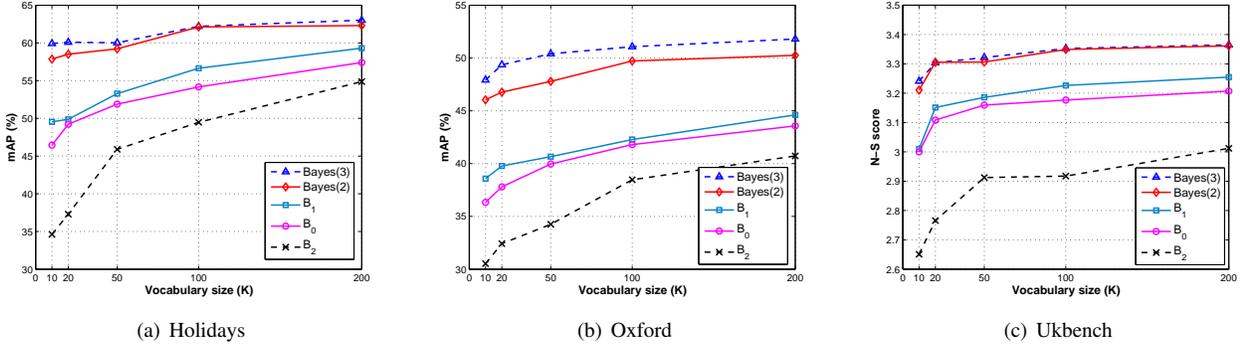


Figure 6. Image retrieval performance as a function of the vocabulary size. Methods include the three baselines, *i.e.*, B_0 , B_1 , B_2 , as well as the proposed Bayes merging of two and three vocabularies, *i.e.*, **Bayes(2)** and **Bayes(3)**, respectively. mAP on (a) Holidays and (b) Oxford, and N-S score on (c) Ukbench are presented.

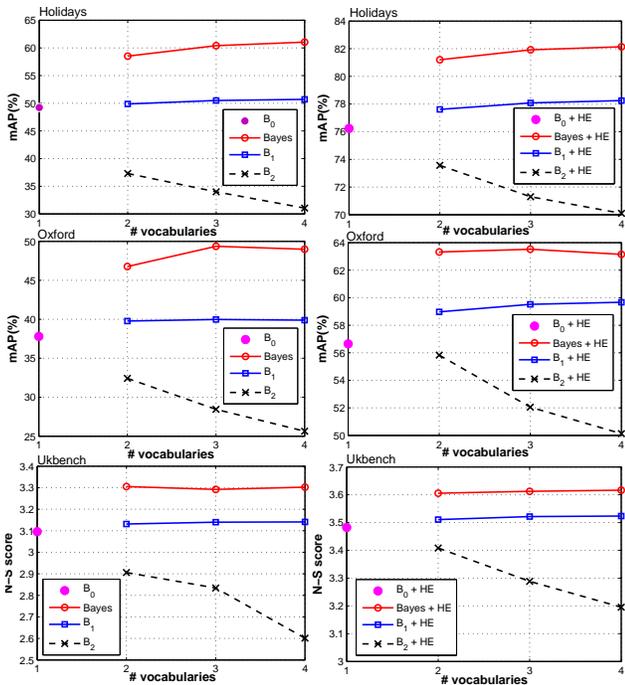


Figure 5. Image retrieval accuracy as a function of the number of merged vocabularies. Results of Holidays (**Top**), Oxford (**Middle**), and Ukbench (**Bottom**) are presented. We compare three baselines, *i.e.*, B_0 , B_1 , and B_2 (see Section. 3.2), with our method (Bayes) (**Left**). We also show the results combined with Hamming Embedding (HE) (**Right**). The vocabulary size is 20K.

two vocabularies of size 20K, the gains in mAP over the three baselines are 9.28%, 8.64%, and 21.21%, respectively. The improvement is even higher for three vocabularies. Nevertheless, we favor two vocabularies due to the fact that the marginal improvement is prominent, while introducing little computational complexity.

Impact of vocabulary sizes The vocabulary size may have an impact on the effectiveness of Bayes merging. To

this end, we generate vocabularies of size 10K, 20K, 50K, 100K, and 200K on the independent Flickr60K data. In Fig. 6, we demonstrate the results obtained from various vocabulary sizes on the three datasets. Except for the three baselines, we also report results obtained by Bayes merging of two or three vocabularies.

From Fig. 6, we can see that B_1 still yields limited improvement over B_0 . Moreover, B_1 and B_2 perform better under those larger vocabularies. This is due to the fact that larger vocabularies reduce correlation. But for large databases, vocabularies are never large enough, so the correlation problem would be more severe in the large-scale case. Moreover, it is clear that the Bayes merging method exceeds the baselines consistently under different vocabulary sizes. Meanwhile, Bayes merging of three vocabularies has a slightly higher performance than two vocabularies.

Merging vocabularies of different sizes Bayes merging can also be generalized to merging vocabularies of different sizes, and the procedure is essentially the same with Algorithm 1. As with the contribution of each vocabulary, we adopt the same unit weight for all vocabularies, as it is shown to yield satisfying performance in [3]. In this paper, we report the merging results on Oxford dataset in Table 4.

Table 4 demonstrates that merging vocabularies of different sizes marginally improves mAP on Oxford. For example, Bayes merging of two vocabularies of size 10K and 20K improves over the $2 \times 10K$ and $2 \times 20K$ Bayes methods by 1.07% and 0.34%, respectively. We speculate that vocabularies of different sizes provide extra complementary information, which can be captured by our method. However, since the smaller vocabulary introduces more noise, the benefit is limited.

Combination with Hamming Embedding To test whether Bayes merging is complementary to some prior arts, we combine it with Hamming Embedding (HE) [4] and burstiness weighting [5] using the default parameters. HE effectively improves the precision of feature matching. In

Methods	Holidays, mAP(%)			Oxford, mAP(%)			Ukbench, N-S		
	2×20K	3×20K	2×50K	2×20K	3×20K	2×100K	2×20K	3×20K	2×100K
B ₀	49.23	49.23	54.17	37.80	37.80	38.80	3.11	3.11	3.17
B ₁	49.87	50.49	56.64	39.78	39.97	41.28	3.15	3.16	3.22
Bayes	58.51	60.40	59.22	46.77	49.36	49.72	3.31	3.30	3.35
B ₁ + HE	77.61	78.08	77.48	58.97	59.52	60.08	3.51	3.51	3.48
Bayes + HE	81.20	81.56	80.60	63.32	63.53	63.96	3.61	3.62	3.57
Bayes + HE + Burst	81.53	81.92	81.08	65.01	64.82	64.73	3.62	3.62	3.59

Table 3. Results on three benchmark datasets for different methods: baselines B₀ and B₁, the proposed method (Bayes), Hamming Embedding (HE) [4], and burstiness weighting (Burst) [5]. We consider the merging of 2×20K, 3×20K, and 2×50K vocabularies, respectively.

Method	B ₁	B ₂	Bayes
10K + 20K	40.89	32.85	47.11
20K + 50K	41.20	34.70	48.85
10K + 20K + 50K	42.31	35.82	49.05

Table 4. The mAP of Bayes merging of vocabularies of different sizes on Oxford dataset. In comparison, Bayes merging of two vocabularies of the same size yields an mAP of 46.04%, 46.77%, and 47.79% for the 10K, 20K, 50K vocabularies, respectively.

our experiment, HE with a single vocabulary achieves an mAP of 76.24% and 56.65% on Holidays and Oxford, and an N-S score of 3.49 on Ukbench, respectively.

The results in Fig. 5 and Table 3 indicate that Bayes merging yields consistent improvements of the B₀ + HE method. Specifically, when merging two vocabularies of 20K, the mAP is improved from 76.24% to 81.20% and from 56.65% to 63.32% on Holidays and Oxford, respectively. Similar trend can be observed on Ukbench: N-S score rises from 3.49 to 3.61. In its nature, HE results in refined matching in the feature space (locally). Complementary, the Bayes merging jointly considers the image- and feature-level similarity. Therefore, while good matching in the feature space can be guaranteed by HE, our method punishes those of a false match in the image space. In this scenario, we actually raise an interesting question: *can we simply trust feature-to-feature similarity in image retrieval?*

In addition, combining burstiness weights brings about extra, though limited improvement (see Table 3). Our implementation differs from [5] in that we do not apply the weights on images in the intersection set, but instead on the difference set ($\mathcal{A} \cup \mathcal{B} - \mathcal{A} \cap \mathcal{B}$) only. A performance summary of various methods is presented in Table 3.

Large-scale experiments To test the scalability of our method, we add the Flickr1M distractor images [4] to the Holidays and Oxford datasets. For comparison, we report the results of baselines B₀ and B₁. From Fig. 7, it is clear that Bayes merging outperforms the two baselines significantly. On Holidays dataset mixed with one million images, Bayes merging achieves mAP of 39.60%, compared with 28.19% and 29.26% of baseline B₀ and B₁, respectively.

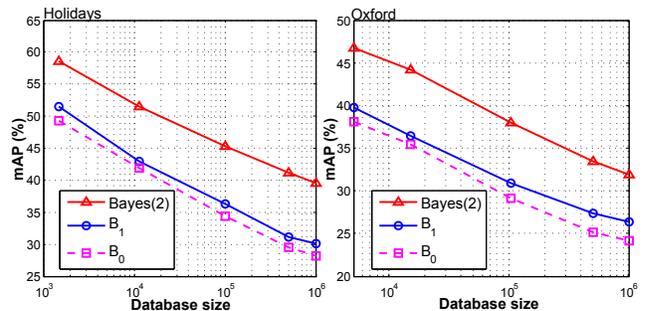


Figure 7. The mAP results as a function of the database size on Holidays and Oxford datasets. Three methods are compared, *i.e.*, baselines B₀, B₁, and Bayes merging of two vocabularies. The vocabulary size is 20K for all methods.

In terms of efficiency, the baseline method B₀ consumes 4 bytes per feature, and 1.9 GB for indexing one million images. The Bayes merging of two vocabularies doubles the memory cost to about 3.8 GB on Flickr1M.

On the other hand, it takes 2.52s and 4.87s for B₀ and B₁ to perform one query on 1 million image size, respectively, using a server with 3.46 GHz and 64GB memory. Bayes merging involves identifying the intersection set and calculate the cardinality ratio. In fact, the cardinality ratio can be computed and stored offline. Moreover, as shown in the **supplementary material**, we are able to perform both the identification and the voting tasks by traversing the two lists of indexed features only once. Therefore, our method only marginally increases the query time to 5.12s.

Comparison with state-of-the-arts We first compare our method with [3] which employs PCA to addresses the correlation problem implicitly. In [3], merging four 16K vocabularies and eight 8K vocabularies yield an mAP of 55.8% and 56.7%, respectively. Moreover, merging vocabularies of multiple sizes obtains a best mAP of 58.8% on Holidays. In comparison, the result obtained by Bayes merging is 58.5% and 60.4% for two and three vocabularies of size 20K, respectively.

Second, we compare the Bayes merging with the Rank Aggregation (RA) method [2, 7] in Table 5. Following [7],

Method	Bayes		Rank aggregation			
	2	3	3	5	7	9
<i>Holidays</i> , mAP(%)	58.5	60.4	49.9	50.7	51.0	51.0
<i>Oxford</i> , mAP(%)	46.8	49.4	39.5	40.8	41.2	41.4

Table 5. Comparisons with rank aggregation (RA). Different numbers of vocabularies are trained to test RA. Vocabulary size is 20K.

Method	Bayes					
	[20]	[6]	[5]	[16]	[15]	
<i>Holidays</i> , mAP(%)	81.9	80.9	81.3	83.9	78.1	-
<i>Oxford</i> , mAP(%)	65.0	68.7	61.5	64.7	-	66.4
<i>Ukbench</i> , N-S	3.62	3.60	3.42	3.54	3.56	3.50

Table 6. Comparisons with the state-of-the-art methods.

we take the median of multiple ranks as the final rank. Since RA works on the rank level, it does not address the correlation problem, so its performance is limited. The results demonstrate the superiority of Bayes merging.

Finally, we compare the results of Bayes merging with state-of-the-arts in Table 6. On the three datasets, we achieve mAP = **81.9%** on *Holidays*, mAP = **65.0%** on *Oxford*, and N-S = **3.62** on *Ukbench*. We have also tested on the data provided by [14], where the codebook size is 65K. On *Oxford* dataset, the mAP is 77.3%. Note that some sophisticated techniques are absent in our system, such as spatial constraints [6, 16], semantic consistency [20], etc. Still, the results demonstrate that the performance of Bayes merging is very competitive. We also provide some sample retrieval results in the **supplementary material**.

6. Conclusion

Multi-vocabulary merging is an effective method to improve the recall of visual matching. However, this process is impaired by vocabulary correlation. To address the problem, this paper proposes a Bayes merging approach to explicitly estimate the matching strength of the indexed features in the intersection sets, while preserving those in the difference set. In a probabilistic view, Bayes merging is capable of jointly modeling an image- and feature-level similarity from multiple sets of indexed features. Specifically, we exploit the probability that an indexed feature is a true match (both locally and globally) if it is located in the intersection sets of multiple inverted files. Extensive experiments demonstrate that Bayes merging effectively reduces the impact of vocabulary correlation, thus improving the retrieval accuracy significantly. Further, our method is efficient, and yields competitive results with state-of-the-arts.

Acknowledgement This work was supported by the National High Technology Research and Development Program of China (863 program) under Grant No. 2012AA011004 and the National Science and Technolo-

gy Support Program under Grant No. 2013BAK02B04. This work also was supported in part to Dr. Qi Tian by ARO grant W911NF-12-1-0057, Faculty Research Awards by NEC Laboratories of America, and 2012 UTSA STARTER Research Award respectively. This work was supported in part by National Science Foundation of China (NSFC) 61128007.

References

- [1] A. Babenko and V. Lempitsky. The inverted multi-index. In *CVPR*, 2012.
- [2] R. Fagin, R. Kumar, and D. Sivakumar. Efficient similarity search and classification via rank aggregation. In *ACM SIGMOD*, 2003.
- [3] H. Jégou and O. Chum. Negative evidences and co-occurrences in image retrieval: The benefit of pca and whitening. In *ECCV*, 2012.
- [4] H. Jégou, M. Douze, and C. Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *ECCV*, 2008.
- [5] H. Jégou, M. Douze, and C. Schmid. On the burstiness of visual elements. In *CVPR*, 2009.
- [6] H. Jégou, M. Douze, and C. Schmid. Improving bag-of-features for large scale image search. *IJCV*, 2010.
- [7] H. Jégou, C. Schmid, H. Harzallah, and J. Verbeek. Accurate image search using the contextual dissimilarity measure. *PAMI*, 32(1):2–11, 2010.
- [8] Z. Liu, H. Li, W. Zhou, R. Zhao, and Q. Tian. Contextual hashing for large-scale image search. *TIP*, 23(4):1606–1614, 2014.
- [9] D. G. Lowe. Distinctive image features from scale invariant keypoints. *IJCV*, 2004.
- [10] A. Mikulík, M. Perdoch, O. Chum, and J. Matas. Learning a fine vocabulary. In *ECCV*, 2010.
- [11] D. Niester and H. Stewenius. Scalable recognition with a vocabulary tree. In *CVPR*, 2006.
- [12] J. Philbin, O. Chum, M. Isard, and A. Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *CVPR*, 2007.
- [13] A. Relja and A. Zisserman. Three things everyone should know to improve object retrieval. In *CVPR*, 2012.
- [14] G. Tolias, Y. Avrithis, and H. Jégou. To aggregate or not to aggregate: Selective match kernels for image search. In *ICCV*, 2013.
- [15] J. Wang, J. Wang, Q. Ke, G. Zeng, and S. Li. Fast approximate k-means via cluster closures. In *CVPR*, 2012.
- [16] X. Wang, M. Yang, T. Cour, S. Zhu, K. Yu, and T. X. Han. Contextual weighting for vocabulary tree based image retrieval. In *ICCV*, 2011.
- [17] Z. Wu, Q. Ke, J. Sun, and H.-Y. Shum. A multi-sample, multi-tree approach to bag-of-words image representation for image retrieval. In *CVPR*, 2009.
- [18] Y. Xia, K. He, F. Wen, and J. Sun. Joint inverted index. In *ICCV*, 2013.
- [19] S. Zhang, Q. Huang, G. Hua, S. Jiang, W. Gao, and Q. Tian. Building contextual visual vocabulary for large-scale image applications. In *ACM MM*, 2010.
- [20] S. Zhang, M. Yang, X. Wang, Y. Lin, and Q. Tian. Semantic-aware co-indexing for near-duplicate image retrieval. In *ICCV*, 2013.
- [21] L. Zheng and S. Wang. Visual phraselet: Refining spatial constraints for large scale image search. *Signal Processing Letters, IEEE*, 20(4):391–394, 2013.
- [22] L. Zheng, S. Wang, Z. Liu, and Q. Tian. Lp-norm idf for large scale image search. In *CVPR*, 2013.
- [23] L. Zheng, S. Wang, Z. Liu, and Q. Tian. Packing and padding: Coupled multi-index for accurate image retrieval. In *CVPR*, 2014.
- [24] W. Zhou, H. Li, Y. Lu, and Q. Tian. Principal visual word discovery for automatic license plate detection. *TIP*, 21(9):4269–4279, 2012.