

Local Orthogonal-Group Testing

Ahmet Iscen Ondřej Chum

Visual Recognition Group, Faculty of EE, Czech Technical University in Prague

Abstract. This work addresses approximate nearest neighbor search applied in the domain of large-scale image retrieval. Within the group testing framework we propose an efficient off-line construction of the search structures. The linear-time complexity orthogonal grouping increases the probability that at most one element from each group is matching to a given query. Non-maxima suppression with each group efficiently reduces the number of false positive results at no extra cost. Unlike in other well-performing approaches, all processing is local, fast, and suitable to process data in batches and in parallel. We experimentally show that the proposed method achieves search accuracy of the exhaustive search with significant reduction in the search complexity. The method can be naturally combined with existing embedding methods.

Keywords: approximate nearest neighbours, group testing, image retrieval

1 Introduction

In this paper, we are interested in approximate nearest neighbor search, specifically in large-scale image search. First, since the seminal paper of Sivic and Zisserman [1], the image similarity was based on the bag-of-words approach [2–6]. Efficient image retrieval was performed via inverted file structure [7]. Later, high dimensional non-sparse descriptors were introduced by VLAD [8] and Fisher vectors [9, 10]. Nowadays, image search is dominated by CNN descriptors [11–15], which also use high dimensional non-sparse vectors to represent images. Image similarity is typically measured by cosine similarity of the descriptors, or equivalently by an Euclidean distance of ℓ_2 normalized vectors. Efficient search in this case is performed by (approximate) nearest neighbor search.

A number of methods exist for efficient search in high dimensional data, using a variety of approaches such as partitioning and embedding. In partitioning, the descriptor space is subdivided and only a small fraction of the data is actually considered for possible nearest neighbors, majority of the data is filtered out. These approaches include kd-tree and forests [16], k-means tree [17], LSH [18–20]. Embedding approaches find a mapping from the original descriptor space to some other (typically of lower dimension or binary) space, where the distance, or the ordering by the distance, can be very efficiently approximated. As an example of the embedding methods we mention LSH [18–20] and other compact binary coding methods [21–24], product quantization [25] and methods derived

from it [26, 27]. Commonly, a combination of partitioning and embedding, such as PQ-IVFADC [25], is adopted. More recent works use neural networks to learn embeddings in a supervised manner [28, 29].

An alternative to these approaches are methods based on group testing. Group testing was first used in World War II by the US army [30]. Due to limited resources, the US army did not want to test each individual soldier for an STD; instead they combined several blood samples into a single mixture, and tested the mixtures. If the test was negative, the associated individuals were deemed sound. If the test was positive, then all the recruits that contributed to this mixture were tested individually. Since the percentage of infected soldiers was low, this procedure dramatically reduced the number of tests needed to screen the population of soldiers.

In approximate nearest neighbor search, the individual blood samples are replaced with vectors, that are grouped into memory units. Each memory unit is represented by a vector, that is constructed from dataset vectors in the unit – different means of construction can be used [31], the simplest being a sum. In adaptive group testing [32, 31, 33], memory vectors are used for efficient pre-filtering of candidates. For the candidates, the exact similarity is then computed from the original vectors. To avoid storing the original vectors, methods of non-adaptive group testing were developed. Roughly, each input vector is stored in multiple memory units. The similarity of input vectors to the query is then estimated from a number of relevant memory units.

In this paper, we are interested in non-adaptive group testing. The contribution of the paper is twofold: local processing of the dataset when constructing the search structure and novel grouping of the vectors that introduces additional constraints used in the scoring stage. The proposed method matches the quality of the dictionary learning methods [34] that benefit from timely offline learning on the whole dataset, while preserving indexing efficiency of the basic group testing methods [32]. The process of encoding and search is outlined in Fig. 1.

Local processing. The construction of the memory units and related decoding structures is linear in the size of the dataset. In order to encode the input data, the method (unlike [34]) does not need to see the whole dataset at once. This property makes it efficient in a streaming scenario, where the dataset is gradually increased in batches.

Vector grouping. The group testing methods are efficient when the positive elements to be retrieved (syphilis infected individuals) are sparse in the dataset (population). By a simple and efficient greedy algorithm of selecting the vectors into groups, we try to minimize the chances that more than one element encoded in a single memory vector is positive. In the decoding stage, such a construction allows to suppress false positive responses. We call this process *correction*.

The rest of the paper is structured as follows. The problem is mathematically formulated and current approaches are reviewed in Section 2. The proposed method is described in Section 3. Thorough experimental evaluation is presented in Section 4.

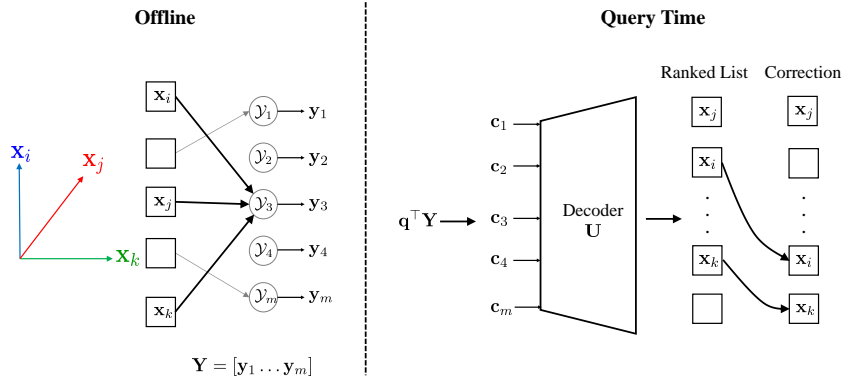


Fig. 1. The input data are represented by vectors $\mathbf{x}_1 \dots \mathbf{x}_k$. By a greedy linear algorithm, vectors that are close to orthogonal are grouped into memory units \mathcal{Y}_i represented by memory vectors \mathbf{y}_i . Sparse decoder matrix \mathbf{U} captures the structure of the group assignment, and is also constructed in the off-line stage. At query time, responses of the memory vectors to the query q are evaluated and used to decode an estimate of the individual input vector similarities. A non-maxima suppression within each memory unit, called correction, is applied to filter out possible false positives. In this particular example, \mathbf{x}_i and \mathbf{x}_k were suppressed as they share \mathcal{Y}_3 with higher ranked \mathbf{x}_j .

2 Problem statement

In this section we detail the group testing setup and its applications in similarity search and image retrieval.

2.1 Adaptive group testing

Group testing was introduced in similarity search by Shi et al. [32]. Assume that the dataset has N d -dimensional vectors $\{\mathbf{x}_i\}_{i=1}^N$. The entire dataset is denoted by the $d \times N$ matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]$. Each vector \mathbf{x}_i is ℓ_2 -normalized such that $\|\mathbf{x}_i\| = 1$, $1 \leq i \leq N$. The similarity measure between a given query \mathbf{q} and a dataset vector \mathbf{x}_i is computed by the scalar product $s_i = \mathbf{x}_i^\top \mathbf{q}$.

The goal of group testing is to infer the similarities s_i efficiently through M ($M \ll N$) group measurements and a decoder. It has three stages. The encoding stage first assigns vectors to groups. Shi et al. [32] define the encoding matrix \mathbf{G} as a $M \times N$ matrix which keeps group assignments, such that $\mathbf{G}_{ji} = 1$ if \mathbf{x}_i belongs to j th group. \mathbf{G} is populated such that each vector is randomly assigned to m groups and each group has exactly n vectors:

$$M = \frac{mN}{n}. \quad (1)$$

Then the group vectors are created based on their assignments:

$$\mathbf{Y} = \mathbf{X}\mathbf{G}^\top. \quad (2)$$

This is equivalent to summing all image vectors assigned to a group:

$$\mathbf{y}_j = \sum_{\mathbf{x} \in \mathcal{Y}_j} \mathbf{x}, \quad (3)$$

where $\mathcal{Y}_j = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ is the set of n vectors from \mathbf{X} assigned to j th group, for $j = 1, \dots, M$.

The group measurement and decoding stages are performed during the query time. For a given query, the M group measurements are computed

$$\mathbf{c} = \mathbf{q}^\top \mathbf{Y}, \quad (4)$$

and pass them to a decoder to approximate N image similarities:

$$\hat{\mathbf{s}} = \mathbf{c}\mathbf{U} = \mathbf{q}^\top \mathbf{Y}\mathbf{U}. \quad (5)$$

The authors consider $\mathbf{U} = \mathbf{G}^\top$ in their work with an extra back-propagation step. In summary, when a new query is given, they i) compute its similarities with group vectors (4), ii) estimate the image vector similarities (5), iii) perform a back propagation where the exact similarities with top ranked image vectors are computed, iv) rank the images according to their similarity.

Isцен et al. [31] use a different setup but a similar idea. They assign each image vector to a single group (called memory units), thereby reducing the number of groups. Furthermore, the properties of randomly assigned memory units are theoretically analyzed. Assume that memory unit $\mathcal{Y}_j = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ stores n dataset vectors. Query \mathbf{q} is related to \mathbf{x}_1 , such that $\mathbf{q} = \alpha\mathbf{x}_1 + \beta\mathbf{Z}$ where α is the similarity between \mathbf{q} and \mathbf{x}_1 , and \mathbf{Z} is a random vector orthogonal to \mathbf{x}_1 and $\|\mathbf{Z}\| = 1$. When sum aggregation (3) is used to create a group representative vector (called *memory vector*) \mathbf{y}_j from \mathcal{Y}_j , the inner product between \mathbf{q} and \mathbf{y}_j becomes

$$\mathbf{y}_j^\top \mathbf{q} = \mathbf{x}_1^\top \mathbf{q} + \alpha \sum_{\mathbf{x}_i \in \mathcal{Y}_j \setminus \mathbf{x}_1} \mathbf{x}_i^\top \mathbf{x}_1 + \beta \sum_{\mathbf{x}_i \in \mathcal{Y}_j \setminus \mathbf{x}_1} \mathbf{x}_i^\top \mathbf{Z}, \quad (6)$$

where $\mathcal{Y}_j \setminus \mathbf{x}_1$ denotes all vectors in \mathcal{Y}_j except \mathbf{x}_1 .

The main source of noise for (6) comes from the middle term, which is basically the interference between \mathbf{x}_1 and the other \mathbf{x}_i in \mathcal{Y}_j . In an attempt to eliminate this noise, a construction of the memory vectors by computing the pseudo-inverse of all vectors assigned to the group is proposed in [31]:

$$\mathbf{y}_j = ([\mathbf{x} \in \mathcal{Y}_j]^+)^{\top} \mathbf{1}_n, \quad (7)$$

where $+$ denotes *Moore-Penrose pseudo-inverse* [35]. This construction is shown to perform better than the sum construction both theoretically and empirically under some mild conditions. At the final stage, they also re-rank all vectors of the highest scoring groups by computing their true similarity with the query vector.

Both methods indeed perform what is called *adaptive group testing* composed of two steps. The first step computes similarities with group representatives.

These indicate which vector similarities are worth being investigated. The second step is a verification process computing the true similarities for these candidate vectors. Even though this strategy gives very accurate results with high efficiency, it requires group representatives as well as all database vectors to be kept in memory for the second adaptive step.

2.2 Non-adaptive group testing

A more modern view on group testing, called *non-adaptive group testing*, was adopted in other computer science related fields thanks to the advancements made in compressed sensing [36–40]. This approach argues that the identification of the infected individuals is possible just from the results of the group tests realized in the first step. There is no need of a second verification step.

A more recent dictionary learning based approach [34] applies non-adaptive group testing to similarity search. Group testing is defined as an optimization problem, where group vectors \mathbf{Y} and the decoder \mathbf{U} are optimized jointly. Unlike previous approaches [32], there are no assignment or construction constraints. The only constraint is to have a sparse \mathbf{U} to reduce the number of vector multiplications during the query time. The solution is found by a dictionary learning optimization algorithm which yields a sparse decoding matrix \mathbf{U} :

$$\begin{aligned} \underset{\mathbf{Y}, \mathbf{U}}{\text{minimize}} \quad & \frac{1}{2} \|\mathbf{X} - \mathbf{Y}\mathbf{U}\|_F^2 + \lambda \|\mathbf{U}\|_1 \\ \text{subject to} \quad & \|\mathbf{y}_k\|_2 \leq 1 \text{ for all } 0 \leq k < M. \end{aligned} \tag{8}$$

This method has a good search efficiency in terms of complexity and memory footprint, without compromising the search accuracy. Original database vectors need not be stored in the memory anymore since there is no need for re-ranking with original similarities. Nevertheless, its main weakness is its offline complexity. The solution to (8) requires the entire dataset to be available, which means that it needs to be recomputed as new data becomes available. Additionally, complexity of solving (8) grows dramatically as N and M increases, limiting its scalability for very large-scale scenarios.

3 Our method

This section describes our contributions to group testing in this paper. Our framework groups orthogonal vectors together, which allows to have an efficient correction step without computing the true similarities of dataset vectors. Both, the encoder and decoder are learned locally from a subset of the dataset. We use the terminology introduced by [31] throughout the paper. Groups are called memory units, and group representatives are called memory vectors.

3.1 Orthogonal memory units

Random assignment [32] of the input vector to memory units is a basic assignment strategy: m permutations of $\{1, \dots, N\}$ denoted by $\{\boldsymbol{\pi}_k\}_{k=1}^m$ are drawn at

random. For the k -th permutation, the vectors whose indices are $\pi_k((\ell - 1)n + 1), \dots, \pi_k(\ell n)$, $\ell \in \{1, \dots, N/n\}$, are grouped into one memory unit, assuming that n divides N . Random assignment is a convenient choice for large-scale datasets or streaming data due to its low complexity and locality. Other alternatives, such as assigning data based on k -means clustering [31] or kd -tree partitioning [33] show that grouping similar vectors together improve efficiency for adaptive group testing. However, these methods have a potentially expensive extra processing step and a verification step with re-ranking with true similarities is needed.

Our method heads the opposite direction. We propose to create a memory unit so that it contains mutually orthogonal vectors. The reason for such a construction is twofold, first minimizing the interference [31] between the vectors stored in a memory unit (6) and increasing the chances that there is only a single matching vector to a query in each memory unit. As we will show later, this property allows us to correct false positives and significantly improve the retrieval accuracy.

Instead of grouping m random vectors (selected via random permutation) as in [32], a random chunk of km vectors (again via random permutation) is selected. Within a chunk (considering only vectors from that chunk), k memory units of m vectors are constructed by a greedy approach. The memory units are initialized with m randomly selected vectors. In each iteration of the algorithm, each of the k memory units is greedily extended by one vector, that is "the most orthogonal" to the vectors already assigned to that particular memory unit ¹.

Due to the greedy nature of the algorithm, the assignment is not globally optimal. Obtaining the globally optimal assignment is intractable. Our experiments show that in practice the group assignments are sufficiently independent: sum [32] and pseudo-inverse memory vector construction for orthogonal memory units are equally good.

The time complexity of the orthogonal assignment is the same as the time complexity of the random assignment, and the algorithm is fast in practice. This makes the algorithm an efficient option for large-scale scenarios. Since it works on small chunks of the dataset independently, it easily handles additional (streaming) data and is easily parallelized (unlike [34]).

3.2 Local decoder

After the assignment proposed in Section 3.1, M memory vectors are constructed with pseudo-inverse (7) and stored in a $d \times M$ matrix $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_M]$. During the query time, the goal is to approximate similarity of the query vector \mathbf{q} and each individual input vector \mathbf{x}_i . This is achieved through a decoder matrix \mathbf{U} , so that

$$\mathbf{q}^\top \mathbf{X} \approx (\mathbf{q}^\top \mathbf{Y}) \mathbf{U}.$$

Shi et al. [32], propose multiple decoding schemes for a given query. The first proposal is to take the pseudo-inverse of the sparse encoder \mathbf{G} (see Section 2.1):

¹ Matlab source code provided in supplementary material

$\mathbf{U} = \mathbf{G}^+$. This is a costly operation which involves computing and storing the dense pseudo-inverse of \mathbf{G} in the query time. Thus, they use a simpler sparse decoder

$$\mathbf{U} = \mathbf{G}^\top \quad (9)$$

and an extra back-propagation step.

In the dictionary learning approach of Iscen et al. [34], group vectors \mathbf{Y} and decoding matrix \mathbf{U} are estimated by a joint optimization (8), which is extremely time demanding. Nevertheless, once the group vectors are constructed and fixed, the decoder matrix \mathbf{U} is estimated by each column \mathbf{u}_i independently by solving the system of linear equations

$$\mathbf{x}_i = \mathbf{Y}\mathbf{u}_i. \quad (10)$$

For efficiency reasons, it is important that \mathbf{u}_i are sparse vectors. Let $\mathcal{S}(i)$ be the set of indices of the non-zero elements in \mathbf{u}_i . For a given $\mathcal{S}(i)$, the solution is found by solving a system of $|\mathcal{S}(i)|$ linear equations

$$\mathbf{x}_i = \mathbf{Y}_{\mathcal{S}(i)}\mathbf{u}_{i,\mathcal{S}(i)}, \quad (11)$$

where $\mathbf{Y}_{\mathcal{S}(i)}$ are the columns of \mathbf{Y} whose indices are stored in $\mathcal{S}(i)$ and $\mathbf{u}_{i,\mathcal{S}(i)}$ is a vector composed of elements of \mathbf{u}_i with indices $\mathcal{S}(i)$.

We propose to construct the set of indices $\mathcal{S}(i)$ based on the groupings of the input vectors into the memory units. Consider the bipartite graph B illustrated in Fig. 2. In this graph, one type of nodes corresponds to the input vectors, the other type corresponds to memory units. There is an edge between two nodes if the corresponding input vector is a member of the memory unit. It is natural to attempt to decode input vector \mathbf{x}_i with memory units that vector is assigned to. We call this 0-order local decoder. For this decoder, the $\mathcal{S}_0(i)$ contains indices of memory units connected by a single edge to \mathbf{x}_i in graph B , as illustrated in Fig. 2 left. Note that this construction has the same set of non-zero elements as the construction of Shi et al. [32], eqn. (9), however the weights estimated by (11) will differ from \mathbf{G}^\top .

Memory vectors with indices $\mathcal{S}_0(i)$, contain other input vectors than \mathbf{x}_i . These vectors influence the estimate of $\mathbf{q}^\top \mathbf{x}_i$. Using the same reasoning as before, we propose to select indices $\mathcal{S}_1(i)$ of the memory units that contain any input vector \mathbf{x}_j assigned to any memory unit with indices $\mathcal{S}_0(i)$. We call this selection 1-order local decoder, set $\mathcal{S}_1(i)$ contains indices of memory units connected by up to three edges to \mathbf{x}_i in graph B , as illustrated in Fig. 2 right.

The proposed selection of the non-zero entries of the decoder \mathbf{U} efficiently picks relevant memory vectors to be used during the estimation of $\mathbf{q}^\top \mathbf{x}_i$ for each \mathbf{x}_i in time complexity that is independent of the size of the input data. This construction is local, does not require the presence of all memory vectors in the memory and thus is suitable for batch processing and parallelization.

Sparse decoder. To sparsify the decoder matrix \mathbf{U} further, we propose to add ℓ_1 -norm regularization into (11):

$$\underset{\mathbf{u}_i}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{x}_i - \mathbf{Y}_{\mathcal{S}(i)}\mathbf{u}_{i,\mathcal{S}(i)}\|_2^2 + \lambda \|\mathbf{u}_{i,\mathcal{S}(i)}\|_1 \quad (12)$$




Fig. 2. Left: An illustration of 0-order local decoder for \mathbf{x}_i . $\mathcal{S}_0(i)$ contains indices of memory units that \mathbf{x}_i is assigned to. **Right:** An illustration of 1-order local decoder for \mathbf{x}_i . $\mathcal{S}_1(i)$ contains indices of the memory units that contain any input vector \mathbf{x}_j assigned to any memory unit with indices $\mathcal{S}_0(i)$.

In practice, we use a greedy algorithm called Orthogonal Matching Pursuit (OMP) [41, 42]. OMP allows to choose the exact number of non-zero elements (instead of setting the parameter λ) in the solution so that the complexity of the decoder \mathbf{U} is directly adjusted. All the local properties of the decoder construction are preserved.

Cascade decoder. The decoder \mathbf{U} can be decomposed into two matrices $\mathbf{U} = \mathbf{U}^0 + \mathbf{U}^1$. Let $\mathbf{c} = \mathbf{q}^\top \mathbf{Y}$, the decoder is then written as:

$$\mathbf{q}^\top \mathbf{X} \approx \mathbf{c} \mathbf{U}^0 + \mathbf{c} \mathbf{U}^1. \quad (13)$$

Since the estimated columns \mathbf{u}_i of \mathbf{U} contain a few significant elements (with high absolute value), some of them typically corresponding to $\mathcal{S}_0(i)$, and a larger number of less significant elements, the columns \mathbf{u}_i^0 of \mathbf{U}^0 are very sparse and contain the most significant entries, while the columns \mathbf{u}_i^1 of \mathbf{U}^1 contain the remaining elements of \mathbf{u}_i . The decomposition is then used in a cascade, first a short list of elements is efficiently obtained by using a rough approximation

$$\mathbf{q}^\top \mathbf{X} \approx \hat{\mathbf{s}}^0 = \mathbf{c} \mathbf{U}^0. \quad (14)$$

Only for the shortlist \mathcal{R} , the finer estimate of the similarity (13) is performed as

$$\mathbf{q}^\top \mathbf{x}_i \approx \hat{\mathbf{s}}^0 + \mathbf{c} \mathbf{u}_i^1 \quad i \in \mathcal{R}.$$

This process reduces the number of operations since \mathbf{U}^0 is much sparser than \mathbf{U} and $|\mathcal{R}| \ll N$. The trade-off between the search accuracy and efficiency is controlled by the sparsity of \mathbf{U}^0 . The memory requirements are increased only marginally, depending on the used representation of sparse matrices.

Synergy with embedding methods. In this section, we discuss the combination of the proposed method with embedding methods, in particular product quantization (PQ) [25]. In order to further reduce the memory footprint and the number of operations, the memory vectors \mathbf{y}_i stored in \mathbf{y} can be compressed by PQ. Approximating $\mathbf{q}^\top \mathbf{y}_i$ by asymmetric product quantizer is equivalent to actually evaluating $\mathbf{q}^\top \hat{\mathbf{y}}_i$, where $\hat{\mathbf{y}}_i$ is a quantized version of \mathbf{y}_i . Due to the local properties of the proposed algorithm, the decoder \mathbf{U} can be construed from $\hat{\mathbf{Y}} = [\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_M]$ without any additional cost.

3.3 Decoder correction

The search results are affected by inaccuracy of the similarity $\hat{\mathbf{s}} \approx \mathbf{q}^\top \mathbf{X}$ estimate (5). The estimate $\hat{\mathbf{s}}$ can be used as a pre-filtering, true similarities can be computed for highly ranked vectors and these scores can be back-propagated to update other scores, as proposed by Shi et al. [32]. This was shown to improve the search accuracy, but requires all dataset vectors to be accessible in the memory during query time. While such an approach reduces the complexity of the search (compared to the exhaustive search), it increases the memory footprint requirements.

Based on the orthogonal memory unit grouping, we propose a simple and efficient correction scheme, which completely avoids the utilization of the input vectors and their true similarities. Due to the orthogonality assignment, we make an assumption that there is at most one matching vector per memory unit to a query. In other words, if multiple vectors from the same memory unit score high in the ranked list, it is highly likely that only one is matching and the rest are false positives. We propose non-maxima suppression per memory unit. In practice, in a single pass through the ranked results, each top ranked (so far non-suppressed) result \mathbf{x}_i suppresses all other vectors that appear in any of the memory vectors together with \mathbf{x}_i .

Experimentally, we show that the orthogonality of the vectors is essential for the correction to work. This correction scheme improves the search accuracy, especially for memory units composed of smaller number n of vectors.

4 Experiments

In this section, we experimentally verify the proposed method (all its components) in a large-scale image retrieval scenario. We first show the benefit of the orthogonal grouping and correction compared to random grouping. We define the following variants for our experiments. **LO-GT** is memory vectors with orthogonal assignment (sec. 3.1) and local decoder (sec. 3.2). **LO-GT*** additionally has the correction step (sec. 3.3). Their random counterparts are **RND** and **RND***. Memory vectors from memory units are created with the pseudo-inverse construction (7) for all cases.

4.1 Experimental setup

Datasets. We use two well-known large-scale image retrieval benchmarks in our experiments: Oxford105k and Paris106k. They contain about 105k and 106k images respectively. They are formed by adding 100k distractor images from Flickr [2] to Oxford Buildings [2] and Paris [3] datasets. We also perform a larger-scale experiment in revisited \mathcal{R} Oxford + \mathcal{R} 1M [43], which consists of new 1M challenging distractor set. We evaluate using the Medium setup. Following the standard evaluation, the search performance is measured by mean average precision (mAP).

Image Representation. We use state-of-the-art image descriptors extracted from a ResNet101 network fine-tuned for image retrieval [15]. Each descriptor is extracted from 3 different image scales using GeM pooling, and combined into a single descriptor as in [43]. Each descriptor has $d = 2048$ dimensionality.

Complexity Analysis. Following the existing work [34], efficiency is reported by measuring the complexity ratio. This metric is based on the total number of scalar operations during the search. It is computed as $\rho = (Md + s)/dN$, where M is the total number of memory vectors, d is the dimensionality, s is the number of non-zero elements in the decoder \mathbf{U} , and N is the number of images in the dataset. Smaller ρ means more efficient search. Since our method does not require any image vectors to be loaded in the memory, ρ also relates to the gain in memory footprint compared to the exhaustive search. The only exception is the cascade decoder, where some of the columns of \mathbf{U}^1 are not touched.

4.2 Retrieval performance

We now compare our orthogonal assignment with random assignment in a retrieval scenario. Various components of group testing, such as group size n , decoder order l and sparsity of \mathbf{U} are analyzed. All the experiments in this subsection are performed on the Oxford105k dataset.

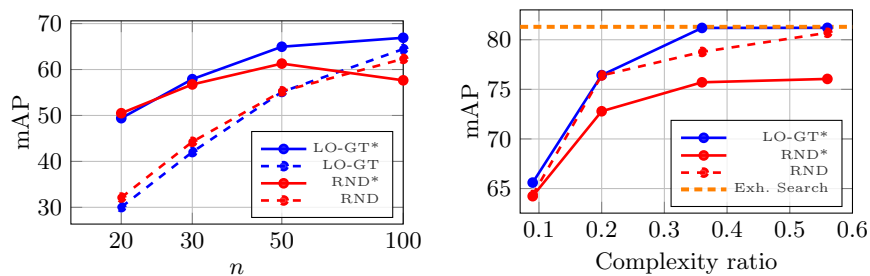


Fig. 3. **Left:** Impact of different n for $m = n/10$ and 0-order decoder on Oxford105k dataset. **Right:** Complexity ration vs. retrieval quality for $l = 1$ -order decoding, $n = 50$, complexity ratio varies by changing $m = 2 \dots 5$.

Group size n . To keep the overall complexity fixed in this experiment, that is keeping sizes of \mathbf{Y} and \mathbf{U} constant, we set $m/n = M/N = 1/10$. Increasing the group size n also increases the number m of memory units each database vector is assigned to. Figure 3 shows the mAP with the 0-order decoder for different values of n , comparing the proposed methods with the random grouping. For this settings, LO-GT without the decoder correction does not bring any significant improvement compared to random grouping. Using the decoder correction (LO-GT*) performs significantly better than the random grouping. As shown in the same figure, the correction actually degrades the performance of random grouping after certain complexity. This shows the benefit of our orthogonal grouping approach, which allows the correction process by assuming that there is only a single matching vector in the dataset.

Performance of the 1-order decoder. The 1-order decoder, non-sparse in this experiment, gives significantly better performance compared to 0-order decoder with higher complexity and memory footprint. In this experiment, we fix $n = 50$ and show the mAP and complexity (ρ) for varying m in Figure 3. Our proposed method LO-GT* achieves the same mAP as exhaustive search with at the complexity ratio of $\rho = 0.36$ (corresponding to $m = 4$), outperforming the random variants.

Sparse decoder is obtained by adding a sparsity constraint on the solution of \mathbf{U} (12). We show the impact of such sparsity constraint in Figure 4. Orthogonal Matching Pursuit algorithm is used to adjust the number of non-zeros (L) on each column of \mathbf{U} . Smaller values L leads to a sparser solution, hence lower complexity ratio. It is shown that we achieve better accuracy than the random variants for all values of L . Setting $L = 300$ gives us the same accuracy as the exhaustive search. This corresponds to the complexity ratio of 0.23. As the complexity of the decoder increases, the benefit of the correction is less pronounced. At the same time LO-GT without correction, *i.e.* due to better estimates of the similarity, significantly outperforms both random grouping variants.

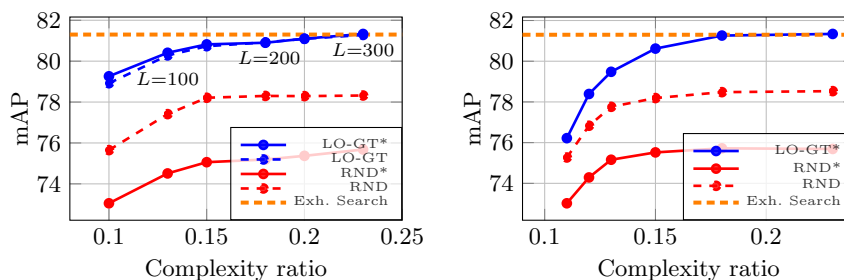


Fig. 4. Left: Search accuracy for sparse \mathbf{U} of 1-order decoder. L is the number of non-zeros in each column of \mathbf{U} . **Right:** Impact of cascade decoder. Complexity ratio is varied by changing $p = 50\% \dots 100\%$. \mathbf{U}^0 is populated with the most significant entries carrying p percent of the column energy. Rest of non-zero entries are assigned to \mathbf{U}_1 .

Cascade decoder is the decomposition of the decoder \mathbf{U} in two matrices $\mathbf{U} = \mathbf{U}^0 + \mathbf{U}^1$ (sec. 3.2). In this experiment, $n = 50$, $m = 4$ and the sparse decoder \mathbf{U} with $L = 300$ from the previous experiment is decomposed into two matrices. Each column of \mathbf{U}^0 is populated with the most significant entries carrying p percent of the column energy. Remaining non-zero entries are assigned to \mathbf{U}^1 . Figure 4 shows the outcome of this approach. Different complexity ratio is obtained by changing $p = 50\% \dots 100\%$. It is shown that the cascade decoder reduces the complexity ratio even further without harming the search accuracy. **Actual search time** is measured as seconds instead of complexity ratio. In a single-thread CPU, exhaustive search in Oxford105k dataset takes 0.198 s per query on average. Our search time (when the complexity ratio is 0.18) in the same environment is 0.054 s per query. That’s a ratio of 0.27, but we would like to note that we use a simple Matlab implementation which is not optimized for this task.

4.3 Comparison with other methods

We compare our method against the existing group testing techniques in the literature [32, 34]. Two of our comparisons are made against the two variants of the group testing framework proposed by Shi et al. [32]. The first variant involves computing the full pseudo-inverse of the assignment graph \mathbf{G} . Since the resulting matrix is a large dense matrix, this variant is not efficient and is only included to serve as a baseline. We also compare against the back-propagation variant proposed by Shi et al. [32], where the scores are updated based on the true distance computations with dataset vectors. This variant is more efficient but requires higher memory footprint than the exhaustive search. All dataset vectors, in addition to memory vectors, need to be available in the memory.

We also compare against the dictionary learning solution proposed by Iscen et al. [34]. This method also has two variants. The first case involves using the entire the dataset to learn group vectors. Offline processing of this variant is not efficient for practical large-scale applications. Optimization problem (8) takes a long time to learn the group vectors. Alternatively, coresets [44, 45] are proposed to reduce the indexing time. Coresets are representative data points sampled from the dataset. Number of coresets is set to $N/5$.

	Complexity	Memory	Oxford105k		Paris106k	
			Index Time	mAP	Index Time	mAP
Baseline	1.00	1.00	-	81.3	-	83.4
GT [32] pinv	4.20	4.20	<1	63.9	<1	56.5
GT [32] w/ bp.	0.36	1.36	<1	73.4	<1	73.6
DL [34] w/ cset	0.18	0.18	273	81.4	288	85.2
DL [34]	0.11	0.11	435	86.8	492	86.2
Ours, $l = 1$	0.18	0.23	4	81.3	4	83.7

Table 1. Comparison of our method and existing group testing methods. Index time (in minutes) is the time it takes to create group vectors and the decoder. We set $M = 8408$ for all methods. We set the size of coresets $N/5$ for DL [34].

Table 1 shows the comparison between our method and prior art in group testing algorithms. Timings are reported on a server with 32 cores. For every variant we set $M = 8408$, which corresponds to $n = 50$ and $m = 4$ for [32] and the proposed method. Note that DL has higher mAP but this comes at a significant offline cost. Finding group vectors with DL [34] involves solving an expensive optimization problem. It also requires entire dataset to be loaded in the memory at once, limiting its scalability for very large N . Furthermore, it is shown that the offline complexity of DL increases exponentially as N increases [34]. Therefore, this method is scalable only if the size of the dataset is reduced with coresets. Search accuracy is lower in that case, but the offline indexing time is still significantly higher than our method.

Batch processing. We evaluate the search performance in a scenario where the data becomes available in batches over time. We divide the Oxford1M dataset randomly into b batches of equal size and create memory vectors \mathbf{Y}_b and the decoding matrix \mathbf{U}_b separately for each batch. After processing all the batches, we concatenate all \mathbf{Y}_b and \mathbf{U}_b and perform the search. Figure 5 shows the mAP for different number of batches. It can be observed that the search accuracy of DL [34] degrades significantly as the data are divided in more batches. This can be explained by the nature of this method, the fewer data are used in the global optimization, the less efficient search. The indexing time of DL is extremely high. Indexing the whole dataset in one go is not tractable, and for $b = 10$ it takes about 50 hours. On the contrary, our method can handle any scenario where matching vectors become available over time. This clearly shows in Figure 5, where the performance of our method is stable regardless of b .

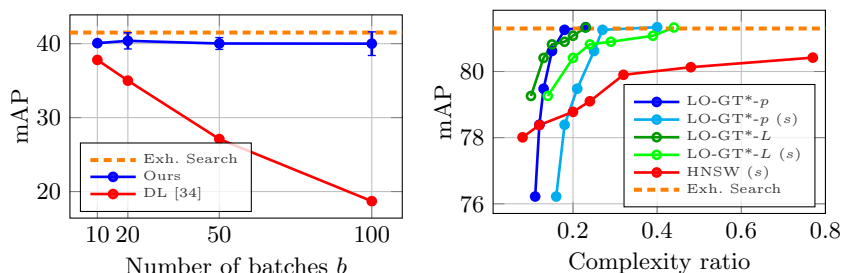


Fig. 5. Left: The mAP of the indexing structure constructed sequentially. \mathcal{R} Oxford + \mathcal{R} 1M is divided into b batches and \mathbf{Y}_b and \mathbf{U}_b is created separately for each batch. We run our method multiple times and show the standard variation in vertical bars. **Right:** Comparison between our methods and HNSW [46] in Oxford105k. Curves with the complexity ratio based on measured run-time are denoted by (s), otherwise measured by the number of scalar operations.

Comparison with partitioning. Finally, we compare our framework against the well-known FLANN toolbox [16]. We use the “autotuned” setting of FLANN, setting the target precision to 0.95. The average speed-up ratio after 5 runs

is 1.45, which corresponds to a complexity ratio of 0.69. Compared to that, our method is able to achieve the baseline mAP performance with only 0.2 complexity ratio and memory footprint.

Comparison with graph-based indexing. Finally, we compare our method against a popular graph-based method by Malkov and Yashunin [46]. We use the implementation provided in FAISS toolbox [47]. To report the complexity ratio, a ratio of approximate-search time using the HNSW index to an exhaustive-search time, both on a single-thread CPU, is computed. A comparison of mAP on Oxford105k for different complexity ratios is shown in Fig 5(right plot). Two variants of the proposed method are shown: sparse 1-order decoder and cascade decoder $\mathbf{U}_0 + \mathbf{U}_1$ (as in Fig. 4), where the complexity is controlled by parameter L and p respectively. Additionally, we also report complexity ratio measured as a ratio of approximate-search and exhaustive-search times (as in HNSW). The plot shows that our method clearly outperforms HNSW. Furthermore, our framework requires significantly smaller memory footprint.

Combination with PQ. As described in Section 3.2, our method is compatible with existing embedding techniques, such as product quantization (PQ) [25]. We compress 2048D memory vectors into c bytes. Table 2 shows the mAP for different c . Note that learning \mathbf{U} from compressed $\hat{\mathbf{Y}}$ (denoted by \mathbf{U}^{Pq} in the table) is important in this case. If \mathbf{U} is learned from non-compressed \mathbf{Y} , then mAP is significantly degraded.

Table 2. Combination of PQ and our method. Each memory vector is compressed into c bytes.

Decoder \ c	256	32	16
\mathbf{U} from \mathbf{Y}	77.7	52.2	44.8
\mathbf{U}^{Pq} from $\hat{\mathbf{Y}}$	79.2	70.5	62.3

5 Conclusions

We have proposed two contributions to the group testing framework. First, the linear-time complexity orthogonal grouping increases the probability that at most one element from each group is matching to a given query. Non-maxima suppression with each group efficiently reduces the number of false positive results at no extra cost. Second, unlike in other similarly performing approaches, such as dictionary learning [34], all processing is local, orders of magnitude faster, and suitable to process data in batches and in parallel. We experimentally show, that for any choice of the efficiency, the proposed method significantly outperforms previously used random grouping. Finally, the proposed method achieves search accuracy of the exhaustive search with significant reduction in the search complexity.

Acknowledgments The authors were supported by MSMT LL1303 ERC-CZ grant and the OP VVV funded project CZ.02.1.01/0.0/0.0/16_019/0000765 “Research Center for Informatics”.

References

1. Sivic, J., Zisserman, A.: Video Google: A text retrieval approach to object matching in videos. In: ICCV. (2003)
2. Philbin, J., Chum, O., Isard, M., Sivic, J., Zisserman, A.: Object retrieval with large vocabularies and fast spatial matching. In: CVPR. (June 2007)
3. Philbin, J., Chum, O., Isard, M., Sivic, J., Zisserman, A.: Lost in quantization: Improving particular object retrieval in large scale image databases. In: CVPR. (June 2008)
4. Jégou, H., Douze, M., Schmid, C.: Improving bag-of-features for large scale image search. IJCV **87**(3) (February 2010) 316–336
5. Jégou, H., Schmid, C., Harzallah, H., Verbeek, J.: Accurate image search using the contextual dissimilarity measure. IEEE Trans. PAMI **32**(1) (January 2010) 2–11
6. van Gemert, J.C., Veenman, C., Smeulders, A.W., Geusebroek, J.: Visual word ambiguity. IEEE Trans. PAMI **32**(7) (July 2010) 1271–1283
7. Babenko, A., Lempitsky, V.: The inverted multi-index. In: CVPR. (June 2012)
8. Jégou, H., Douze, M., Schmid, C., Pérez, P.: Aggregating local descriptors into a compact image representation. In: CVPR. (June 2010)
9. Perronnin, F., Dance, C.R.: Fisher kernels on visual vocabularies for image categorization. In: CVPR. (June 2007)
10. Perronnin, F., J.Sánchez, Mensink, T.: Improving the fisher kernel for large-scale image classification. In: ECCV. (September 2010)
11. Babenko, A., Slesarev, A., Chigorin, A., Lempitsky, V.: Neural codes for image retrieval. In: ECCV. (2014)
12. Tolias, G., Sicre, R., Jégou, H.: Particular object retrieval with integral max-pooling of cnn activations. ICLR (2016)
13. Gordo, A., Almazan, J., Revaud, J., Larlus, D.: Deep image retrieval: Learning global representations for image search. ECCV (2016)
14. Radenović, F., Tolias, G., Chum, O.: CNN image retrieval learns from bow: Un-supervised fine-tuning with hard examples. ECCV (2016)
15. Radenović, F., Tolias, G., Chum, O.: Fine-tuning cnn image retrieval with no human annotation. arXiv preprint arXiv:1711.02512 (2017)
16. Muja, M., Lowe, D.G.: Scalable nearest neighbor algorithms for high dimensional data. IEEE Trans. PAMI **36** (2014)
17. Nistér, D., Stewénius, H.: Scalable recognition with a vocabulary tree. In: CVPR. (June 2006) 2161–2168
18. Indyk, P., Motwani, R.: Approximate nearest neighbors: towards removing the curse of dimensionality. In: STOC. (1998) 604–613
19. Gionis, A., Indyk, P., Motwani, R.: Similarity search in high dimension via hashing. In: VLDB. (1999) 518–529
20. Datar, M., Immorlica, N., Indyk, P., Mirrokni, V.: Locality-sensitive hashing scheme based on p-stable distributions. In: Proceedings of the Symposium on Computational Geometry. (2004)
21. Lv, Q., Charikar, M., Li, K.: Image similarity search with compact data structures. In: CIKM. (November 2004) 208–217
22. Norouzi, M., Punjani, A., Fleet, D.J.: Fast search in hamming space with multi-index hashing. In: CVPR. (2012)
23. Weiss, Y., Torralba, A., Fergus, R.: Spectral hashing. In: NIPS. (December 2009)
24. Raginsky, M., Lazebnik, S.: Locality-sensitive binary codes from shift-invariant kernels. In: NIPS. (2010)

25. Jégou, H., Douze, M., Schmid, C.: Product quantization for nearest neighbor search. *IEEE Trans. PAMI* **33**(1) (January 2011) 117–128
26. Ge, T., He, K., Ke, Q., Sun, J.: Optimized product quantization for approximate nearest neighbor search. In: *CVPR*. (June 2013)
27. Kalantidis, Y., Avrithis, Y.: Locally optimized product quantization for approximate nearest neighbor search. In: *CVPR*. (2014)
28. Jain, H., Zepeda, J., Pérez, P., Gribonval, R.: Subic: A supervised, structured binary code for image search. In: *ICCV*. (2017)
29. Jain, H., Zepeda, J., Pérez, P., Gribonval, R.: Learning a complete image indexing pipeline. (2018)
30. Dorfman, R.: The detection of defective members of large populations. *The Annals of Mathematical Statistics* **14**(4) (1943) 436–440
31. Iscen, A., Furon, T., Gripon, V., Rabbat, M., Jégou, H.: Memory vectors for similarity search in high-dimensional spaces. *IEEE Trans. Big Data* **4**(1) (2018)
32. Shi, M., Furon, T., Jégou, H.: A group testing framework for similarity search in high-dimensional spaces. In: *ACM Multimedia*. (November 2014)
33. Iscen, A., Amsaleg, L., Furon, T.: Scaling group testing similarity search. In: *ACM ICMR*. (2016)
34. Iscen, A., Rabbat, M., Furon, T.: Efficient large-scale similarity search using matrix factorization. In: *CVPR*. (2016)
35. Rao, C.R., Mitra, S.K.: Generalized inverse of matrices and its applications. Volume 7. (1971)
36. Aldridge, M., Baldassini, L., Johnson, O.: Group testing algorithms: bounds and simulations. *IEEE Trans. Inform. Theory* (2014)
37. Bickson, D., Baron, D., Ihler, A., Avissar, H., Dolev, D.: Fault identification via nonparametric belief propagation. *IEEE Transactions on Signal Processing* **59**(6) (2011) 2602–2613
38. Cheraghchi, M., Hormati, A., Karbasi, A., Vetterli, M.: Compressed sensing with probabilistic measurements: A group testing solution. In: *47th Annual Allerton Conference on Communication, Control, and Computing*. (2009)
39. Gilbert, A., Indyk, P.: Sparse recovery using sparse matrices. *Proceedings of the IEEE* **98**(6) (2010) 937–947
40. Sejdinovic, D., Johnson, O.: Note on noisy group testing: asymptotic bounds and belief propagation reconstruction. In: *48th Annual Allerton Conference on Communication, Control, and Computing*. (2010)
41. Pati, Y.C., Rezaifar, R., Krishnaprasad, P.: Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In: *ASILOMAR*. (1993) 40–44
42. Davis, G.M., Mallat, S.G., Zhang, Z.: Adaptive time-frequency decompositions with matching pursuit. In: *SPIE’s International Symposium on Optical Engineering and Photonics in Aerospace Sensing*. (1994) 402–413
43. Radenović, F., Iscen, A., Toliás, G., Avrithis, Y., Chum, O.: Revisiting oxford and paris: Large-scale image retrieval benchmarking. In: *CVPR*. (2018)
44. Agarwal, P.K., Har-Peled, S., Varadarajan, K.R.: Approximating extent measures of points. *Journal of the ACM* **51**(4) (2004) 606–635
45. Feldman, D., Feigin, M., Sochen, N.: Learning big (image) data via coresets for dictionaries. *Journal of Mathematical Imaging and Vision* **46**(3) (2013) 276–291
46. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734* (2017)

47. Malkov, Y.A., Yashunin, D.: Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. arXiv preprint arXiv:1603.09320 (2016)