

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

5-2016

HDIdx: High-dimensional indexing for efficient approximate nearest neighbor search

Ji WAN

Chinese Academy of Sciences

Sheng TANG

Chinese Academy of Sciences

Yongdong ZHANG

Chinese Academy of Sciences

Jintao LI

Chinese Academy of Sciences

Pengcheng WU

Singapore Management University, pcwu@smu.edu.sg

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

WAN, Ji; TANG, Sheng; ZHANG, Yongdong; LI, Jintao; WU, Pengcheng; and HOI, Steven C. H.. HDIdx: High-dimensional indexing for efficient approximate nearest neighbor search. (2016). *Neurocomputing*. 237, 401-404.

Available at: https://ink.library.smu.edu.sg/sis_research/3413

This Journal Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Ji WAN, Sheng TANG, Yongdong ZHANG, Jintao LI, Pengcheng WU, and Steven C. H. HOI

HDIIdx: High-Dimensional Indexing for Efficient Approximate Nearest Neighbor Search

Ji Wan, Sheng Tang, Yongdong Zhang, Jintao Li
† *Key Lab of Intelligent Information Processing,
ICT, CAS, Beijing 100190, China*
‡ *University of the CAS, Beijing 100049, China*
{wanji,ts,zhyd,jtli}@ict.ac.cn

Pengcheng Wu, Steven C.H. Hoi
* *School of Information Systems,
Singapore Management University,
Singapore 178902*
{pcwu,chhoi}@smu.edu.sg

Abstract

Fast Nearest Neighbor (NN) search is a fundamental challenge in large-scale data processing and analytics, particularly for analyzing multimedia contents which are often of high dimensionality. Instead of using exact NN search, extensive research efforts have been focusing on approximate NN search algorithms. In this work, we present “HDIIdx”, an efficient high-dimensional indexing library for fast approximate NN search, which is open-source and written in Python. It offers a family of state-of-the-art algorithms that convert input high-dimensional vectors into compact binary codes, making them very efficient and scalable for NN search with very low space complexity.

Keywords: High-Dimensional Indexing, Approximate Nearest Neighbor Search, Product Quantization, Spectral Hashing

1. Introduction

Nearest neighbor (NN) search, also known as proximity search or similarity search, aims to find closest or most similar data points/items from a collection of data points/items. It is a fundamental technique for many application domains. In many real-world applications, like content-based image retrieval [1] and multimedia data mining [2], data is often represented in high-dimensional spaces. This makes *NN* search on large-scale high-dimensional data very challenging given limited storage and computational resources. Instead of achieving the *exact* NN search, recent research efforts have been devoted to *approximate* NN (*ANN*) search. Among varied solutions, hashing based techniques have been one of the most powerful techniques to index large-scale high-dimensional data. There are two kinds of hashing algorithms for *ANN* search: 1) mapping data to the Hamming space while preserving similarities in the original space, and 2) compressing data into short codes and approximating the distance with the one computed over these codes. Comprehensive studies of hashing algorithms for *ANN* search can be found in [3, 4].

In this work, we present a new library of high-dimensional indexing solutions for fast *ANN* search, which is open-source and written in Python. Python is a powerful and successful programming language for scientific computing, and becomes very popular for big data analysis re-

This work was done when Ji Wan visited Dr Hoi’s group at SMU.




Figure 1: System architecture of the proposed HDIdx solution.

cently [5, 6]. Several *ANN* search libraries have already existed for Python, including the well-known *scikit-learn* [5], *Annoy*¹ and *NearPy*². All these libraries are based on Locality Sensitive Hashing (LSH) [7], a scalable *ANN* search solution in high-dimensional space. However, LSH is a data-independent algorithm which maps the input vectors to hash codes by random projections. It often requires long codes to guarantee a high precision and many hash tables to ensure a high recall. Moreover, original vectors are also needed for ranking the candidates returned by LSH. As a result, LSH based *ANN* solutions usually require large amounts of memory.

In contrast to the above mentioned libraries, we use *unsupervised learning to hash* algorithms in HDIdx. By exploring the underlying distribution of data using machine learning techniques, we can map the high-dimensional input vectors to compact hash codes and apply highly efficient search algorithms on these binary codes. Original vectors can be discarded once they are mapped to the compact codes, thus leading to very small memory cost required. Empirical studies show that HDIdx can search a million-scale high-dimensional database in less than one millisecond.

2. Software Framework

Fig. 1 shows the proposed system architecture of HDIdx, which consists of three main modules: 1) *Encoder*: it compresses the input vectors into compact hash codes, 2) *Indexer*: it indexes the base items and performs *NN* search, and 3) *Storage*: it encapsulates the underlying data storage and provides unified interface for the *Indexer*.

The work flow of HDIdx is detailed as follows. First of all, an *Encoder* is learned from some vectors. Then, the base vectors are mapped to hash codes using the learned *Encoder*. After that, an *Indexer* builds indexes over these hash codes, and writes the indexes to memory or database. When a query vector arrives, it is first mapped to some hash codes by the same *Encoder*, and then the *Indexer* searches for similar items from hashing code database to match this query according to its hash codes.

¹<https://github.com/spotify/annoy>

²<https://github.com/pixelogik/NearPy>

3. Implementation and Empirical Results

3.1. Implementation Details

HDIdx is mainly developed and implemented in the Python language, with some computationally intensive part implemented in C language. The *Storage* module consists of wrappers for various storage media, like memory and databases, and is easy to understand. Here we only present details of the *Encoder* and the *Indexer* modules.

3.1.1. Encoder

The *Encoder* module compresses the input vectors into compact hash codes, which requires little storage resource and can be searched very fast. Two unsupervised learning to hash algorithms are used for ANN search in HDIdx: Spectral Hashing (SH) [8] and Product Quantization (PQ) [9].

SH maps the input vector to binary codes, and measures the similarities by Hamming distance. HDIdx uses the *compiler intrinsics* to compute the Hamming distance between these codes, which is extremely fast. Even with ordinary computational resources (e.g. a laptop or a PC), it can compute one million distances for 64-bit binary codes in 0.01 seconds. Additionally, since the number of distinct Hamming distances is very limited, which is bounded by the number of bits, the NNs can be found using efficient partial counting sort algorithm with a time complexity of $O(N)$, where N is the number of base vectors.

PQ divides the vector space into m parts and maps an input vector to m precomputed centroids in each sub-space. Then the distance in the original space can be approximated by the one between these centroids. HDIdx uses *Asymmetric Distance Computation (ADC)* for distance approximation. ADC only maps the base vectors to the precomputed centroids, then approximates the distance by comparing the centroids and the original query vector, instead of quantized query vector. The distance calculation between PQ codes is accelerated by look-up tables. For exhaustive search over PQ codes, the time complexity is $O(N \log R)$, where R is the number of returned base vectors.

The main parameter in HDIdx is the number of bits b . For the PQ algorithm, we fix the codebook size of each sub-quantizer to 256, so each index is encoded by 8 bits and the number of sub-quantizer $m = b/8$.

3.1.2. Indexer

The *Indexer* module indexes the base items and performs fast NN search, and provides non-exhaustive searching schemes to accelerate the search on large scale database.

One way to perform non-exhaustive search on the SH codes is to build a hash table with the SH codes as indexes, then the NNs can be found by exploring a sequence of Hamming balls with progressively increasing radius. However, the number of buckets will increase exponentially with respect to the code length, and most computation will be wasted on checking empty buckets when the codes get longer. Multi-index hashing (MIH) [10] is used to address this issue. The basic idea is to split the b -bit code into t substrings, then build t hash tables using these t substrings as keys respectively. We can get candidate NNs by exploring much smaller Hamming balls in each subspace, and final results can be obtained by ranking these candidates with full-length codes.

To accelerate the search on PQ codes, we implements an inverted file system (IVF) to perform non-exhaustive search [9]. The base vectors are first separated into k' groups by a coarse quantizer q_c , then the residual vectors in each group are compressed to PQ codes. Here residual vector is the offset with respect to the centroid in each group: $r(x) = x - q_c(x)$. When a query comes, only w groups correspond to the w nearest neighbors of the query in the coarse codebook will be checked. The final results is obtained by searching the PQ codes in these groups exhaustively.




Figure 2: Recall rate of different indexing schemes.

Table 1: Search time for 100- NN search (ms).

Number of bits	8	16	32	64	128
Spectral Hashing	3.923	5.843	7.221	11.965	20.478
Product Quantization	15.078	15.340	16.932	20.628	28.296

3.2. Empirical results

We evaluate HDIdx on the SIFT1M dataset [9], which consists of 3 subsets: 1) 1-million base vectors on which the search is performed, 2) 10,000 query vectors, and 3) 100,000 training vectors for learning the *Encoder*, where each vector is of 128 dimensions. The time costs of 100- NN search is used to measure the search efficiency, and $recall@R$ is used to measure the search quality. $recall@R$ is the proportion of query vectors for which the NN is ranked in the first R positions. The experiments were conducted on a laptop with 2.40GHz i7 processor and 16GB RAM.

First of all, we evaluate the exhaustive search performance of SH and PQ codes with respect to b . Fig. 2 shows that $recall@R$ increases with b , and PQ is always better than SH with the same b . On the other hand, Table 1 shows that SH is much faster than PQ, especially for small b .

Moreover, we compare HDIdx with *Annoy* and *NearPy* in a 100- NN search task. We evaluate 64-bit SH and PQ codes, together with multi-index hashing (MIH) on SH codes for $t = 4$, inverted file (IVF) on PQ codes for $k' = 1024$, $w \in \{5, 10\}$. *Annoy* is evaluated with $nt \in \{16, 32\}$, where nt is the number of trees. *NearPy* is evaluated with $nb \in \{16, 32\}$ and $L = 8$, where nb is the number of binary projections and L is the number of hash tables. Server observations can be obtained from Table 2. First, both MIH and IVF significantly speed up the search on SH and PQ codes without sacrificing the search quality. Second, both HDIdx and *Annoy* are much faster and more accurate than *NearPy*. Third, the search time and quality of IVF is comparable with *Annoy*, but HDIdx is more efficient in terms of storage costs since it does not store the original vectors while *Annoy* does. It requires 512MB for 1 million 128-D vectors, but only 8MB for 64-bit hash codes.

In the above experiments, IVF beats MIH in terms of both speed and quality. Optimization for MIH according to [11] is under development, there are also research on optimization of PQ [12]. These improvements will be put into the future release of HDIdx.

Table 2: Performance for 100-*NN* search.

Method	SH	PQ	MIH	IVF	IVF	Annoy	Annoy	NearPy	NearPy
Parameters	-	-	$t = 4$	$w = 5$	$w = 10$	$nt = 16$	$nt = 32$	$nb = 16$	$nb = 32$
Time (ms)	11.965	20.628	4.978	0.540	0.915	0.516	0.920	138.076	2.412
recall@R	0.532	0.915	0.543	0.826	0.904	0.753	0.892	0.654	0.205

4. Illustrative Examples

Illustrative examples of HDIdx can be found at: <http://goo.gl/eMyh3Z>.

5. Conclusions

This work presented the HDIdx software that offers cutting-edge solutions for indexing large-scale high-dimensional data. The current library implemented two state-of-the-art high-dimensional feature compressing algorithms, SH and PQ, and their respective indexing solutions, MIH and IVF. Our empirical results showed that our solution is very efficient in both time and storage costs. In the future work, we wish to keep improving HDIdx by exploring some more state-of-the-art research.

References

- [1] J. Wan, D. Wang, S. C. Hoi, P. Wu, J. Zhu, Y. Zhang, J. Li, Deep learning for content-based image retrieval: A comprehensive study, in: Proceedings of the ACM MM, 2014, pp. 157–166.
- [2] S. Wang, Q. Huang, S. Jiang, Q. Tian, L. Qin, Nearest-neighbor method using multiple neighborhood similarities for social media data mining, *Neurocomputing* 95 (2012) 105–116.
- [3] J. Wang, H. T. Shen, J. Song, J. Ji, Hashing for similarity search: A survey, CoRR abs/1408.2927.
- [4] K. Grauman, R. Fergus, Learning Binary Hash Codes for Large-Scale Image Search, Springer Berlin Heidelberg, 2013.
- [5] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al., Scikit-learn: Machine learning in python, *The Journal of Machine Learning Research* 12 (2011) 2825–2830.
- [6] I. Idris, Python Data Analysis, Packt Publishing Ltd, 2014.
- [7] P. Indyk, R. Motwani, Approximate nearest neighbors: Towards removing the curse of dimensionality, in: Proceedings of the ACM STOC, 1998, pp. 604–613.
- [8] Y. Weiss, A. Torralba, R. Fergus, Spectral hashing, in: Proceedings of the NIPS, 2008, pp. 1753–1760.
- [9] H. Jégou, M. Douze, C. Schmid, Product quantization for nearest neighbor search, *IEEE Trans. Pattern Anal. Mach. Intell.* 33 (1) (2011) 117–128.
- [10] M. Norouzi, A. Punjani, D. J. Fleet, Fast search in hamming space with multi-index hashing, in: Proceedings of the IEEE CVPR, 2012, pp. 3108–3115.
- [11] J. Wan, S. Tang, Y. Zhang, L. Huang, J. Li, Data driven multi-index hashing, in: Proceedings of the IEEE ICIP, 2013, pp. 2670–2673.
- [12] T. Ge, K. He, Q. Ke, J. Sun, Optimized product quantization, *IEEE Trans. Pattern Anal. Mach. Intell.* 36 (4) (2014) 744–755.