

---

# Evaluation of Hashing Methods Performance on Binary Feature Descriptors

Jacek Komorowski<sup>1</sup> and Tomasz Trzcinski<sup>2</sup>

<sup>1</sup> Warsaw University of Technology, Warsaw, Poland

[jacek.komorowski@gmail.com](mailto:jacek.komorowski@gmail.com)

<sup>2</sup> Warsaw University of Technology, Warsaw, Poland [t.trzcinski@ii.pw.edu.pl](mailto:t.trzcinski@ii.pw.edu.pl)

**Summary.** In this paper we evaluate performance of data-dependent hashing methods on binary data. The goal is to find a hashing method that can effectively produce lower dimensional binary representation of 512-bit FREAK descriptors. A representative sample of recent unsupervised, semi-supervised and supervised hashing methods was experimentally evaluated on large datasets of labelled binary FREAK feature descriptors.

**Key words:** data-dependent hashing methods, binary feature descriptors

## 1 Introduction

This paper presents results of an experimental evaluation of recent data-dependent hashing methods applied to binary feature descriptors. The work was motivated by challenges in development of a real-time structure-from-motion solutions for mobile platforms with limited hardware resources. One of the key elements in a typical structure-from-motion processing pipeline is feature matching step, where correspondences between features detected on a new image and features found on previously processed images are being sought. Such correspondences are used to compute camera orientation and build a 3D model of an observed scene.

Binary feature descriptors, such as FREAK [1], are good choice for mobile solutions. They can be efficiently computed and are very compact (512 bits for FREAK versus 512 bytes for real-valued SIFT descriptor). Hamming distance between two binary feature descriptors can be quickly computed using few machine code instructions. But even if comparison of two binary descriptors is very fast, finding correspondence between thousands of features detected in a new image and millions of features on previously processed images requires significant processing power. For real-valued descriptors efficient approximate nearest neighbours search methods can be applied, such as FLANN [13]. Un-

fortunately methods based on clustering do not work well with binary data [14].

In this paper we investigate if hashing methods can be used to reduce dimensionality of 512-bit binary FREAK descriptors to improve feature matching performance and lower storage requirements. Additionally we want to verify if additional training information, in the form of landmark id linked with each descriptor, can be used to improve accuracy of searching for matching descriptors.

Authors have not encountered any results of evaluation of hashing methods on binary data. In the context of image-based information retrieval, hashing algorithms were evaluated on real-valued descriptors such as GIST or SIFT. Lack of such results motivated the research described in this paper.

The paper is structured as follows. Section 2 briefly describes hashing methods. Section 3 presents results of the experimental evaluation of representative hashing methods on large datasets of binary FREAK feature descriptors. The last section concludes the article and presents ideas for future research.

## 2 Overview of hashing methods

Hashing for similarity search is a very active area of development. Many new hashing methods were proposed in the last few years. A number of surveys documenting current state-of-the-art was published recently [16] [19].

Two major categories of hashing methods can be distinguished: *data-independent* methods and *data-dependent* methods. *Data-independent* methods, such as Locality Sensitive Hashing (LSH) [5], do not take into account characteristics of the input data. As such, they have inferior performance in real-life applications, where input data usually has some intrinsic characteristic which can be exploited. We focus our attention on data-dependent on *data-dependent* methods, also known as *learning to hash*. These methods exploit properties of the input data to produce more discriminative and compact binary codes. Data-dependent approach can be further categorized by the level of an external supervision. Unsupervised methods use techniques as spectral analysis or kernelized random projections to compute affinity-preserving binary codes. They exploit the structure among a sample of unlabelled data to learn appropriate embeddings. Semi-supervised or supervised methods exploit additional information from annotated training data. Additional information is usually given as the similarity matrix or list of pairs of similar and dissimilar items. Semi-supervised methods assume that explicit similarity information is provided for only a fraction of an input dataset. Affinity between other elements is inferred from the distance in the input space.

*Learning to hash* is defined [15] as learning a **compound hash function**,  $\mathbf{y} = H(\mathbf{x})$ , mapping an input item  $\mathbf{x}$  to a compact binary code  $\mathbf{y}$ , such that nearest neighbour search in the coding space is efficient and the

result is a good approximation of the true nearest search result in the input space.  $K$ -bit binary code  $\mathbf{y} \in \mathbb{B}^K$  for a sample point  $\mathbf{x} \in \mathbb{R}^D$  is computed as  $\mathbf{y} = [y_1, y_2, \dots, y_K,] = [h_1(\mathbf{x}), h_2(\mathbf{x}), \dots, h_K(\mathbf{x})]$ . Each  $h_k$  is a **binary hash function**, mapping elements from  $\mathbb{R}^D$  to  $\mathbb{B} = \{0, 1\}$ . A compound hash function  $H = [h_1, h_2, \dots, h_K]$  is an ordered set of binary hash functions computing  $K$ -bit binary code.

Two most popular choices of a **hash function** are linear projection and kernel-based. Linear projection hash functions are in the form:  $y = h(\mathbf{x}) = \text{sgn}(\mathbf{w}^T \mathbf{x} + b)$ , where  $\mathbf{w} \in \mathbb{R}^D$  is the projection vector and  $b$  is the bias. Kernel-based hash functions are in the form:  $y = h(\mathbf{x}) = \text{sgn}\left(\sum_{t=1}^T w_t K(\mathbf{s}_t, \mathbf{x}) + b\right)$ , where  $K$  is a kernel function,  $\mathbf{s}$  is a set of representative samples that are randomly chosen from the dataset or cluster centres of the dataset and  $w_t$  are weights. Other choices of hash function include spherical functions, Laplacian eigenfunctions, neural networks, decision trees-based and non-parametric functions.

### 3 Evaluation results

#### *Experiment setup*

The experimental evaluation was conducted on datasets consisting of hundred thousands or more labelled 512-bit FREAK feature descriptors. Datasets were created by structure-from-motion application developed in Google Tango project<sup>3</sup>. Each descriptor is labelled with a corresponding landmark id. Descriptors with the same landmark id are projections of the same scene point (landmark) on different images.

Table 1 lists hashing methods evaluated in this paper. Each method was first trained on the training dataset. Learned hash functions were applied to the test dataset to generate hash codes of a different length: 32, 64, 128 and 256-bits. Search precision (*Precision@1*) using the resulting hash codes was evaluated and reported. *Precision@1* for a dataset is calculated as a mean *precision@1* when searching for nearest neighbours using a linear scan for 20 thousand elements randomly sampled from the dataset. *Precision@1* for a sampled element is 1, if its nearest neighbour (based on Hamming distance) in the entire dataset is labelled with the same landmark id. Otherwise, precision is 0.

#### *Unsupervised hashing methods*

This section contains results of an experimental evaluation of unsupervised hashing methods. The evaluation was conducted using publicly available MATLAB implementation [2]. For brief description of methods evaluated in this section refer to [2].

<sup>3</sup> See: <https://get.google.com/tango/>

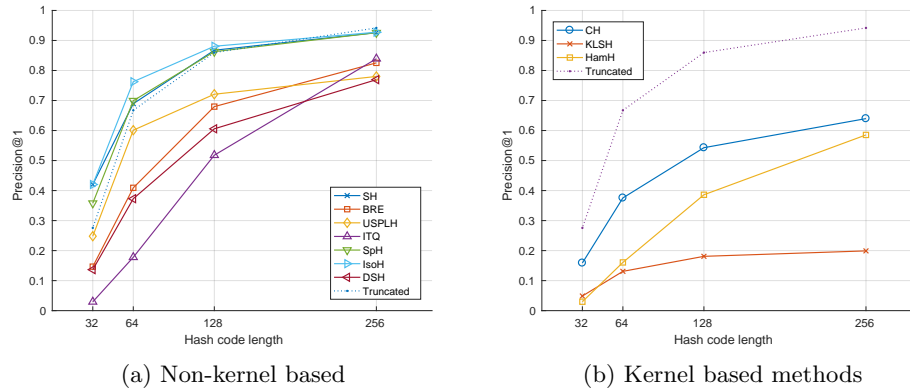
Method	Class	Hash function
Spectral Hashing (SH) [20]	U	Laplacian eigenfunction
Binary Reconstructive Embeddings (BRE) [7]	U	linear projection
Unsupervised Sequential Projection Learning Hashing (USPLH) [18]	U	linear projection
Iterative Quantization (ITQ) [3]	U	linear projection
Isotropic Hashing (IsoH) [6]	U	linear projection
Density Sensitive Hashing (DSH) [11]	U	linear projection
Spherical Hashing (SpH) [4]	U	spherical function
Compressed Hashing [12]	U	kernel-based
Harmonious Hashing (HamH) [22]	U	kernel-based
Kernelized Locality Sensitive Hashing (KLSH) [8]	U	kernel-based
Sequential Projection Learning (SPLH) [17]	SS	linear projection
Bootstrap Sequential Projection Learning – linear version (BTSPLH) [21]	SS	linear projection
Fast supervised hashing (FastHash) [9]	S	boosted decision trees

**Table 1.** List of evaluated hashing methods. Class: U = unsupervised, SS = semi-supervised, S = supervised.

Hashing methods were trained using a dataset consisting of 1 267 346 FREAK descriptors. The evaluation was done on a separated test dataset consisting of 342 602 descriptors associated with 40 704 landmarks.

Fig. 1 presents search precision in datasets created from the test dataset by applying seven unsupervised, data-dependent hashing methods. To baseline the results, precision of a linear search on the test dataset truncated to first  $k$  bits was evaluated. It must be noted, that discarding last 256 bits of original FREAK descriptor has little effect on the nearest neighbour search precision. *Precision@1* decreases by 2.1% from 0.962 to 0.942. The reason is likely the construction of FREAK descriptor itself, where more discriminative binary tests are used to generate first bits of the descriptor. Interestingly, for 256-bit hash codes, all evaluated hashing methods perform worse compared to naive bit truncation. For shorter codes (128-bits and below) only three hashing methods: Isotropic Hashing [6], Spherical Hashing [4] and Spectral Hashing [20] yield better accuracy. For 128-bit codes, the advantage of best methods over a naive bit truncation is minimal (1-2%). For the shorter codes best hashing methods perform noticeably better. Best performing method is Isotropic Hashing (IsoH) [6].

Performance of kernel-based hashing methods on the binary data is very poor. Results are much worse than naive approach of truncating an original dataset to the first  $k$  bits. This is expected as kernel-based hash functions are based on the set of anchor points, that are representative samples or cluster



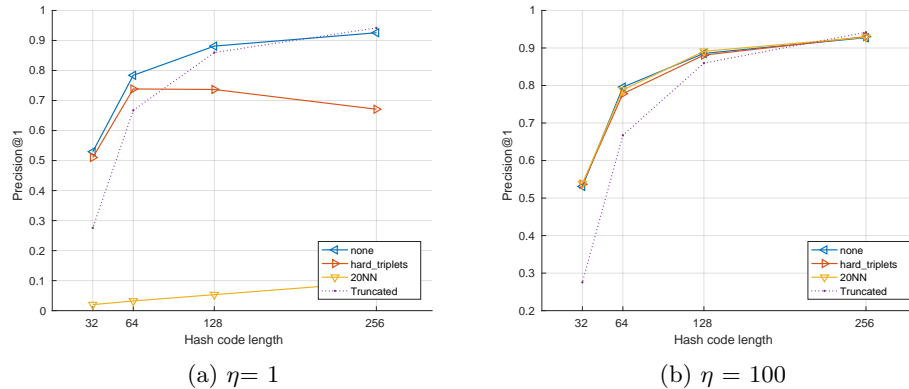
**Fig. 1.** Search precision using hash codes generated from the test dataset by applying unsupervised hashing methods. Search precision in the test dataset truncated to first  $k$  bits is plot with a dotted line for comparison.

centres for the training dataset. Clustering methods perform poorly on data from binary spaces. One of the reasons are decision boundaries in binary spaces [14], as large proportion of points in the Hamming space is equidistant from two randomly chosen anchor points.

#### *Semi-supervised and supervised hashing methods*

Semi-supervised and supervised data-dependent hashing methods evaluated in this section were trained using the dataset consisting of 340 063 descriptors associated with 37 148 landmarks. The evaluation was done on the separated test dataset consisting of 342 602 descriptors associated with 40 704 landmarks.

Figure 2 presents search precision in datasets created from the training dataset by applying semi-supervised Sequential Projection Learning (SPLH) [17] method. SPLH objective function consists of two components: supervised empirical fitness and unsupervised information theoretic regularization. A supervised term tries to minimize empirical error on the labelled data. That is, for each bit minimize a number of instances where elements with the same label are mapped to different values and elements with different labels are mapped to the same value. An unsupervised term provides regularization by maximizing desirable properties like variance and independence of individual bits. Different lines on the plot correspond to different similarity encoding schemes. In *hard triplets* encoding, for each element from the training dataset, the closest element linked with the same landmark id is encoded as similar and the closest element with a different landmark id is encoded as dissimilar. In *20NN* encoding, for each element from the training dataset, the similarity with its 20 nearest neighbours is encoded. For comparison, SPLH method was



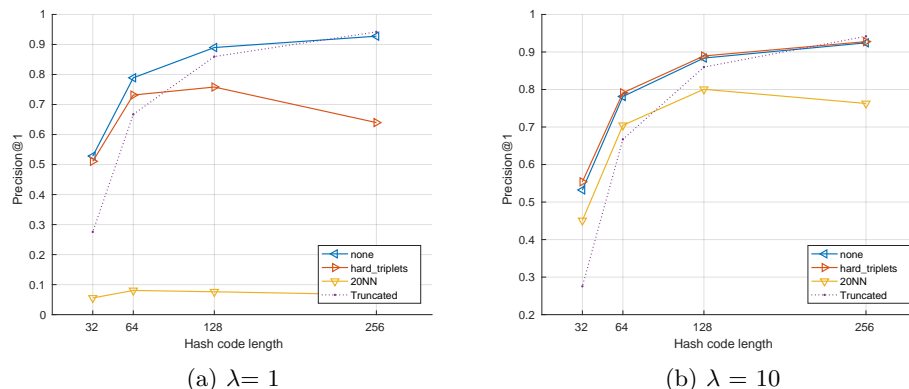
**Fig. 2.** Search precision using hash codes generated from the test dataset by semi-supervised SPLH method. Different subplots correspond to different weight  $\eta$  of an unsupervised term in the objective function. On each plot results obtained using three different similarity encodings schemes are shown. Search precision in the test dataset truncated to first  $k$  bits is plot with a dotted line for comparison.

evaluated without any similarity information (*none*), using only unsupervised term in the optimization function. Surprisingly, using supervised information does not provide any improvement in the search accuracy. In contrary, when supervised term in the objective function has higher weight (lower  $\eta$ , left subplot), the results are noticeably worse, especially for longer codes. The best results are achieved when supervised information is not used at all (*none*). When unsupervised term in the objective function has higher weight (higher  $\eta$ , right subplot), the encoded similarity information has little effect and the performance is the same as without using any supervised information.

Semi-supervised nonlinear hashing (BTSPLH) [21] is an enhancement of SPLH method. Instead of the boosting-like process in SPLH, authors propose a bootstrap-style sequential learning scheme to derive the hash function by correcting the errors incurred by all previously learned bits. The results of evaluation of BTSPLH method are depicted on Figure 3. The results are very similar to previous SPLH method. When supervised term in the objective function has higher weight (lower  $\lambda$ , left subplot), the results using supervised information (*hard\_triplets* and *20NN* encoding) are poor. The best results are achieved when supervised information is not used at all.

The last method evaluated in this paper is supervised FastHash [10] algorithm. It uses a two-step learning strategy: binary code inference is followed by binary classification step using an ensemble of decision trees.

The key decision when using supervised hashing methods is a choice of similarity encoding scheme. For large datasets it's not feasible to encode similarity between all  $N \times N$  pairs of elements. Figure 4a presents results when




**Fig. 3.** Search precision using hash codes generated from the test dataset by semi-supervised BTSPLH method. Different subplots correspond to different weight  $\lambda$  of an unsupervised term in the objective function. On each plot results obtained using three different similarity encodings schemes are shown. Search precision in the test dataset truncated to first  $k$  bits is plot with a dotted line for comparison.

for each element from the training dataset, similarity for its 20 nearest neighbours, all similar and 100 randomly chosen dissimilar elements is encoded. Figure 4b shows results when number of selected dissimilar elements is increased to 300. In both cases search precision is worse compared to naive bit truncation. It can be observed that encoding similarity information between more pairs improves the performance. Unfortunately, for practical reasons, it was not possible to further increase amount of supervised training data. For a training dataset consisting of over 300 thousand elements, encoding similarities between each element and over 300 other elements produces over 100 million pairs. Learning procedure requires over 10 GB of memory to efficiently process such amount of data.

In both cases severe overfitting can be observed. As decision tree depth grows and model complexity increases, the discrepancy between performance on the test and training set grows.

## 4 Conclusions and future work

Evaluation of hashing methods on large datasets of binary FREAK descriptors yield some surprising results. For 256-bit hash codes (half of the length of the original FREAK descriptor) none of the examined methods performed better compared to naive bit truncation approach. In theory, linear projection-based hash methods should be able to produce at least as good result. Yet all of the evaluated methods performed worse. For short codes, hashing methods gain advantage. For 128-bit codes Isotropic Hashing [6], the best unsupervised



**Fig. 4.** Search precision using 256-bit hash codes generated from the test dataset by supervised FastHash method. For each element from the training set, similarity information was encoded for 20 nearest neighbours, all similar and different number of dissimilar elements: 100 in (a) and 300 in (b). For comparison, search precision on the training dataset is shown with a dashed line and search precision in the test dataset truncated to first  $k$  bits is plot with a dotted line.

method, allows achieving 2% better search accuracy compared to naive bit truncation. For 64-bit codes, Isotropic Hashing produces 15% better results.

Examined semi-supervised and supervised methods were not able to benefit from the additional supervised information in the form of landmark ids linked with feature descriptors in the training set. Surprisingly, adding more supervised information in the examined semi-supervised methods produced worse results. The likely reason is, that due to limited hardware resources, similarity information can be explicitly encoded only for fraction of pairs of points from the training dataset, which leads to suboptimal performance.

Hashing methods can be beneficial when storage is a primary concern and short codes are required. They can be used to generate compact, 64-bit binary codes, from original 512-bit FREAK feature descriptors. This would decrease storage requirements four times, and increase descriptor matching speed at the expense of moderate decrease in search precision (about 15%). Further reduction to 32 bits has a detrimental effect on search precision, reducing it by over 40%.

As a future work it will be beneficial to investigate other approaches to supervised dimensionality reduction, such as Mahalanobis metric learning.

## Acknowledgment

This research was supported by Google’s Sponsor Research Agreement under the project ”Efficient visual localization on mobile devices”.



## References

1. A. Alahi, R. Ortiz, and P. Vandergheynst. Freak: Fast retina keypoint. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, June 2012.
2. D. Cai. A Revisit of Hashing Algorithms for Approximate Nearest Neighbor Search. *ArXiv e-prints*, December 2016.
3. Y. Gong, S. Lazebnik, A. Gordo, and F. Perronnin. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12), Dec 2013.
4. J. P. Heo, Y. Lee, J. He, S. F. Chang, and S. E. Yoon. Spherical hashing. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, June 2012.
5. Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, New York, NY, USA, 1998. ACM.
6. Weihao Kong and Wu jun Li. Isotropic hashing. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*. Curran Associates, Inc., 2012.
7. Brian Kulis and Trevor Darrell. Learning to hash with binary reconstructive embeddings. In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*. Curran Associates, Inc., 2009.
8. Brian Kulis and Kristen Grauman. Kernelized locality-sensitive hashing for scalable image search. In *Computer Vision, 2009 IEEE 12th International Conference on*. IEEE, 2009.
9. Guosheng Lin, Chunhua Shen, and Anton van den Hengel. Supervised hashing using graph cuts and boosted decision trees. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2015.
10. Guosheng Lin, Chunhua Shen, and Anton van den Hengel. Supervised hashing using graph cuts and boosted decision trees. *IEEE transactions on pattern analysis and machine intelligence*, 37(11), 2015.
11. Yue Lin, Deng Cai, and Cheng Li. Density sensitive hashing. *CoRR*, abs/1205.2930, 2012.
12. Yue Lin, Rong Jin, Deng Cai, Shuicheng Yan, and Xuelong Li. Compressed hashing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
13. Marius Muja and David G Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. *VISAPP (1)*, 2(331-340), 2009.
14. Tomasz Trzcinski, Vincent Lepetit, and Pascal Fua. Thick boundaries in binary space and their influence on nearest-neighbor search. *Pattern Recognition Letters*, 33(16), 2012.
15. Jingdong Wang, Heng Tao Shen, Jingkuan Song, and Jianqiu Ji. Hashing for similarity search: A survey. *CoRR*, abs/1408.2927, 2014.
16. Jingdong Wang, Ting Zhang, Jingkuan Song, Nicu Sebe, and Heng Tao Shen. A survey on learning to hash. *CoRR*, abs/1606.00185, 2016.
17. Jun Wang, Shih-Fu Chang, and S. Kumar. Semi-supervised hashing for large-scale search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(undefined), 2012.

18. Jun Wang, Sanjiv Kumar, and Shih fu Chang. Sequential projection learning for hashing with compact codes. In Johannes Frnkranz and Thorsten Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*. Omnipress, 2010.
19. Jun Wang, Wei Liu, Sanjiv Kumar, and Shih-Fu Chang. Learning to hash for indexing big data - A survey. *CoRR*, abs/1509.05472, 2015.
20. Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral hashing. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*. Curran Associates, Inc., 2009.
21. Chenxia Wu, Jianke Zhu, Deng Cai, Chun Chen, and Jiajun Bu. Semi-supervised nonlinear hashing using bootstrap sequential projection learning. *IEEE Trans. Knowl. Data Eng.*, 25(6), 2013.
22. Bin Xu, Jiajun Bu, Yue Lin, Chun Chen, Xiaofei He, and Deng Cai. Harmonious hashing. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*. AAAI Press, 2013.