Sparse Projections for High-Dimensional Binary Codes

Yan Xia[1], Kaiming He[2], Pushmeet Kohli[2], Jian Sun[2]
[1]University of Science and Technology of China. [2]Microsoft Research.

Recent work on representation learning using deep neural networks has shown that features of thousands of dimensions or even more are useful for various recognition tasks. On the other hand, it has been noticed that for input signals of dimensionality d, the code-length b of the binary code required to achieve reasonable accuracy (compared with no encoding) is usually $O(d)$. Based on these facts, we consider the problem of learning long binary codes for high-dimensional arbitrary input signals.

We observe two key challenges arise while learning and using long binary codes: (1) lack of an effective regularizer for learned high-dimensional mapping and (2) high computational cost for computing long codes. In this paper, we overcome both these problems by introducing a sparsity encouraging regularizer that reduces the effective number of parameters involved in the learned projection operator. This regularizer can not only reduces overfitting but, due to the sparse nature of the projection matrix, also leads to a dramatic reduction in the computational cost.

To be specific, we use $\mathrm{X} \in \mathbb{R}^{d\times n}$ to denote a matrix whose each column is a d-dimensional datum. Our target is to learn a projection matrix $\mathrm{R} \in \mathbb{R}^{b\times d}$ for producing b-bit binary codes by $\mathrm{B} = \mathrm{sign}(\mathrm{RX}) \in \{-1,1\}^{b\times n}$. A widely considered objective function involves minimizing the distortion, as adopted in ITQ [2] and its variants [3, 6]. In ITQ, the objective function is only regularized by an orthogonality constraint, which is a weak regularizer. To further regularize the projection of high-dimensional data, we introduce a sparsity constraint. Our objective function is:

$$\min_{\mathrm{R},\mathrm{B}} \|\mathrm{RX}-\mathrm{B}\|_\mathrm{F}^2 \qquad (1)$$
$$\text{s.t.} \quad \mathrm{R}^\mathrm{T}\mathrm{R} = \mathrm{I}, \quad \text{and} \quad |\mathrm{R}|_0 \le m.$$

Here $\|\cdot\|_\mathrm{F}$ denotes Frobenius norm, and $|\cdot|_0$ denotes the number of non-zero elements in matrix. Parameter m directly controls the sparsity of R.

However, it is challenging to optimize a matrix with both a orthogonality constraint and a sparsity constraint. To find a feasible solution, we adopt the variable-splitting and penalty techniques in optimization [1]. We introduce an auxiliary variable $\bar{\mathrm{R}}$ and relax the problem as:

$$\min_{\mathrm{R},\bar{\mathrm{R}},\mathrm{B}} \|\bar{\mathrm{R}}\mathrm{X}-\mathrm{B}\|_\mathrm{F}^2 + \beta\|\bar{\mathrm{R}}\mathrm{X}-\mathrm{R}\mathrm{X}\|_\mathrm{F}^2 \qquad (2)$$
$$\text{s.t.} \quad \bar{\mathrm{R}}^\mathrm{T}\bar{\mathrm{R}} = \mathrm{I}, \quad \text{and} \quad |\mathrm{R}|_0 \le m,$$

where β is a penalty weight and we fix it to 1 in experiments. We solve (2) in an alternating manner: updating one variable with others fixed.

Step 1. Fix B **and** $\bar{\mathrm{R}}$**, update** R**.** This sub-problem is:

$$\min_{\mathrm{R}} \|\mathrm{RX}-\mathrm{Z}\|_\mathrm{F}^2 \qquad (3)$$
$$\text{s.t.} \quad |\mathrm{R}|_0 \le m,$$

where $\mathrm{Z} = \bar{\mathrm{R}}\mathrm{X}$ is fixed. This sub-problem looks like a sparse coding problem under the m-sparsity constraint. But in commonly studied sparse coding problems, the variable R is a vector and RX is a vector-matrix multiplication. Here we develop a solution to the matrix variable R.

Instead of directly optimizing (3), we solve the following problem:

$$\min_{\mathrm{R},\mathrm{S}} \phi(\mathrm{R},\mathrm{S}) = \min_{\mathrm{R},\mathrm{S}} \|\mathrm{RX}-\mathrm{Z}\|_\mathrm{F}^2 + \|\mathrm{X}\|_\mathrm{F}^2\|\mathrm{R}-\mathrm{S}\|_\mathrm{F}^2 - \|\mathrm{RX}-\mathrm{SX}\|_\mathrm{F}^2, \qquad (4)$$

where ϕ is a surrogate objective function [5] and S is a matrix of the same size as R. The solution to (4) can be proven to give the solution to (3). We optimize (4) by an alternating algorithm as in [5]: fixing S and solving for R, and vice versa. In this way, the iterative solution to (3) is given as

$$\mathrm{R}_{t+1} = thr_m\left(\mathrm{R}_t + \frac{1}{\|\mathrm{X}\|_\mathrm{F}^2}(\bar{\mathrm{R}}-\mathrm{R}_t)\mathrm{X}\mathrm{X}^\mathrm{T}\right), \qquad (5)$$

This is an extended abstract. The full paper is available at the Computer Vision Foundation webpage.

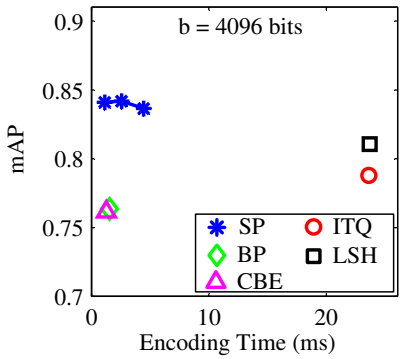

Figure 1: ANN search accuracy *vs.* encoding time on a dataset of one million 4096-d DNN features. Our sparse projection (SP) method, with 5%, 10%, and 15% non-zero elements, is more accurate and about 10× faster than ITQ [2] and LSH, and is more accurate than BP [3] and CBE [6].

where thr_m is an operator that only keeps the largest m entries. A natural initialization to (5) is to let $\mathrm{R}_0 = \bar{\mathrm{R}}$. Then the first iteration of (5) gives:

$$\mathrm{R} = thr_m\left(\bar{\mathrm{R}}\right). \qquad (6)$$

In experiments, we find that the one-step solution gives comparable accuracy, but is much faster for training.

Step 2. Fix B **and** R**, update** $\bar{\mathrm{R}}$**.** This sub-problem is an orthogonal procrustes problem. The solution is $\bar{\mathrm{R}} = \mathrm{V}\mathrm{U}^\mathrm{T}$, where U and V are from the SVD of $\mathrm{X}\mathrm{Y}^\mathrm{T}$. It is worth noting that according to modern studies [4], the above solution is valid for all $b \ge d$. Thus our method and ITQ are both naturally applicable for generating *higher*-dimensional binary codes.

Step 3. Fix R **and** $\bar{\mathrm{R}}$**, update** B**.** This sub-problem is: $\min_B \|\bar{\mathrm{R}}\mathrm{X}-\mathrm{B}\|_\mathrm{F}^2 = \max_B \sum_{i,j}(\bar{\mathrm{R}}\mathrm{X})_{ij}\mathrm{B}_{ij}$. Since $\mathrm{B}_{ij} \in \{-1,1\}$, the solution is $\mathrm{B} = \mathrm{sign}(\bar{\mathrm{R}}\mathrm{X})$.

We iteratively run in Step 1-3 (start from Step3) for 50 iterations, and finally a sparse matrix R is produced for projecting data.

We run comprehensive experiments of ANN search, image retrieval and image classification. Results show that our method not only leads to better accuracy than competitive dense projection methods (ITQ [2] and LSH) with the same code lengths, but is also over one order of magnitude faster. Furthermore, our method is more accurate than two recent methods (BP [3] and CBE [6]) that are designed for fast high-dimensional binary encoding. As an example, Figure 1 shows a comparison on one million 4096-d DNN features.

[1] Richard Courant. Variational methods for the solution of problems of equilibrium and vibrations. *Bulletin of the American Mathematical Society*, 49(1):1–23, 1943.

[2] Yunchao Gong and Svetlana Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *CVPR*, 2011.

[3] Yunchao Gong, Sanjiv Kumar, Henry A Rowley, and Svetlana Lazebnik. Learning binary codes for high-dimensional data using bilinear projections. In *CVPR*, 2013.

[4] John C Gower and Garmt B Dijksterhuis. *Procrustes problems*, volume 3. Oxford University Press Oxford, 2004.

[5] Kenneth Lange, David R Hunter, and Ilsoon Yang. Optimization transfer using surrogate objective functions. *Journal of computational and graphical statistics*, 2000.

[6] Felix X Yu, Sanjiv Kumar, Yunchao Gong, and Shih-Fu Chang. Circulant binary embedding. In *ICML*, 2014.