

Transfer Hashing with Privileged Information

Joey Tianyi Zhou[†], Xinxing Xu[†], Sinno Jialin Pan[‡], Ivor W. Tsang[§],
Zheng Qin[†], and Rick Siow Mong Goh[†]

[†]Institute of High Performance Computing, Singapore

[‡]Nanyang Technological University, Singapore

[§]QCIS, University of Technology Sydney, Australia

[†]{zhouty,xuwx,qinz,gohsm}@ihpc.a-star.edu.sg, [‡]sinnopan@ntu.edu.sg, [§]ivor.tsang@uts.edu.au

Abstract

Most existing learning to hash methods assume that there are sufficient data, either labeled or unlabeled, on the domain of interest (i.e., the target domain) for training. However, this assumption cannot be satisfied in some real-world applications. To address this data sparsity issue in hashing, inspired by transfer learning, we propose a new framework named Transfer Hashing with Privileged Information (THPI). Specifically, we extend the standard learning to hash method, Iterative Quantization (ITQ), in a transfer learning manner, namely ITQ+. In ITQ+, a new slack function is learned from auxiliary data to approximate the quantization error in ITQ. We developed an alternating optimization approach to solve the resultant optimization problem for ITQ+. We further extend ITQ+ to LapITQ+ by utilizing the geometry structure among the auxiliary data for learning more precise binary codes in the target domain. Extensive experiments on several benchmark datasets verify the effectiveness of our proposed approaches through comparisons with several state-of-the-art baselines.

1 Introduction

Hashing methods have been applied for efficient similarity search in many application areas, especially for information retrieval [Wang *et al.*, 2016]. The goal of hashing is to design or learn a compact binary code, with each bit taking values of either -1/1 or 0/1, in a low-dimensional space for each data instance such that similar instances in the original space are mapped to similar binary codes. As a result, data instances can be stored in a low cost, and the similarity between instances can be efficiently computed with the Hamming distance using binary operation (XOR).

Most existing learning to hash methods require a large amount of data instances to learn a set of hash functions to construct binary codes [Wang *et al.*, 2016]. However in some real-world applications, for a domain of interest, i.e., the target domain, the data instances may not be sufficient enough to learn a precise hashing model. For example, Taobao.com provides a platform for small businesses and individual entrepreneurs to open online stores. Suppose an individual en-

trepreneur wants to build a hashing system for all the images of the products sold in his/her online store. Unfortunately, the number of images are not sufficient enough to build a precise hashing system. A straightforward solution is to download images of the same or related products from other e-commerce websites, such as Amazon.com or eBay.com, to help learning a hashing system. However, in many individual online stores in Taobao.com, the images used to demonstrate the products are usually amateur, which are taken by the individual entrepreneur himself/herself, while the images of products demonstrated in Amazon.com are professional. A simple aggregation of these two different kinds of images may not help to learn a precise hashing system for the target images. Motivated by transfer learning [Pan and Yang, 2010], instead of using the images downloaded from Amazon.com or other e-commerce websites directly, one can extract knowledge from the auxiliary images, and then transfer the knowledge to help learning a more precise hashing system for the target images.

Specifically, we propose a novel framework named “Transfer Hashing with Privileged Information” (THPI), which is a marriage of transfer learning and a new learning paradigm, namely Learning Using Privileged Information (LUPI), proposed by [Vapnik and Vashist, 2009], where privileged information is assumed to be available for each training instance in the training phrase, and missing in the testing phrase. Our aim is to construct precise hash codes for a target domain instances, e.g., the images of product in an individual online store, by encoding the privileged information from a source domain, e.g., images on products downloaded from Amazon.com, into a learning to hash model. In this example, the target images can be considered as training instances, while auxiliary images of the same or similar products as the target images can be considered as their privileged information in the training phrase. As shown in [Vapnik and Izmailov, 2015], the amount of training data required for training can be dramatically reduced with the help of privileged information. Therefore, we expect that with privileged information, we are able to learn precise hashing functions for the target domain where data instances are not sufficient.

Note that the proposed THPI framework is different from cross-modal hashing which assumes that training instances of different modalities are rich, and are sufficient to learn reliable hash codes, respectively [Kumar and Udupa, 2011].

In most cross-modal hashing methods, full correspondences across different modalities are required as input [Kumar and Udupa, 2011; Wu *et al.*, 2015]. Moreover, the goal of cross-modal hashing is to retrieve relevant data across modalities. In THPI, the goal is to learn a set of good hash functions with insufficient training instances in a target domain, and retrieve relevant information in the target domain.

The novelties of our work are summarized as follows,

- We propose a novel framework named “Transfer Hashing with Privileged Information” (THPI) to alleviate data sparsity in a target domain by transferring knowledge from a source domain.
- A new algorithm named ITQ+ is proposed, where a novel slack function for incorporating privileged information is introduced to regularize the learning of hash codes for the target domain.
- We further extend ITQ+ to LapITQ+, where underlying graph structure extracted from the source domain is encoded as a prior for learning more precise hash codes for the target domain.

2 Related Work

2.1 Learning to Hash

Most hashing methods focus on how to quantize data with minimal information loss. For example, Locality-sensitive Hashing (LSH) [Raginsky and Lazebnik, 2009] uses a set of random projections followed by thresholding. Spectral Hashing (SH) formulates the quantization as spectral graph partitioning [Weiss *et al.*, 2008], where the graph geometry on original feature space is preserved. Iterative Quantization (ITQ) [Gong and Lazebnik, 2011] is proposed to refine the initial projections, e.g., Principal Component Analysis (PCA) [Tipping and Bishop, 1999], Canonical Correlation Analysis (CCA) [Hardoon *et al.*, 2004], such that quantization error can be reduced. All these methods require sufficient data for learning hash functions for a domain of interest. Different from these methods, THPI aims at alleviating data sparsity and further improving hashing performance on the target domain by exploiting knowledge from other domains, which may be of heterogeneous features. Recently, cross-modal hashing [Kumar and Udupa, 2011; Zhang *et al.*, 2011] that aims to learn hash codes with data from different modalities draws much attention. However, cross-modal hashing assumes training instances of different modalities be sufficient to learn reliable hashing functions respectively, which is different from THPI.

2.2 Transfer Learning

Transfer learning (TL) [Pan and Yang, 2010] aims to transfer knowledge across different domains so that rich source domain knowledge can be used to build better classifiers on a target domain where the transferred knowledge can be labels [Zhou *et al.*, 2014b], features [Pan *et al.*, 2011], cross-domain correspondences [Zhou *et al.*, 2014a; 2016]. TL has shown promising results in many machine learning tasks, such as classification and regression. To the best of our

knowledge, there is only one work on studying transfer learning for hashing [Ou *et al.*, 2014]. Different from their work, we focus on how to transfer knowledge across heterogeneous feature spaces in an unsupervised manner.

2.3 Learning using Privileged Information

Recently, Vapnik and Vashist [2009] introduced a new learning paradigm namely learning using privileged information (LUPI). In LUPI, auxiliary privileged features are assumed to be available in the training phrase but not available in the testing phrase. A new model SVM+ is proposed by exploiting the privileged features to construct a correcting function in traditional Support Vector Machines (SVMs) to control the loss such that the learned classifiers can embrace stronger generalization ability. Specifically, given a set of training data $\{\mathbf{x}_i, \tilde{\mathbf{x}}_i\}_{i=1}^l$, where $\tilde{\mathbf{x}}$ is the corresponding privileged features for the original features \mathbf{x} , SVM+ aims to learn a target classifier $f(\mathbf{x}) = \mathbf{w}^\top \mathbf{x}$ from the original feature vectors and a slack approximation function $\epsilon(\tilde{\mathbf{x}}) = \tilde{\mathbf{w}}^\top \tilde{\mathbf{x}}$ from the privileged feature vectors, simultaneously. The goal of the slack approximation function is to control the loss of the target classifier by incorporating privileged features. The objective function of SVM+ is written as follows:

$$\begin{aligned} \min_{\mathbf{w}, \tilde{\mathbf{w}}, b} \quad & \frac{1}{2} (\|\mathbf{w}\|^2 + \lambda \|\tilde{\mathbf{w}}\|^2) + C \sum_{i=1}^l \epsilon(\tilde{\mathbf{x}}) \\ \text{s.t.} \quad & y_i \mathbf{w}^\top \mathbf{x}_i \geq 1 - \epsilon(\tilde{\mathbf{x}}_i), \forall i = 1, 2, \dots, l, \\ & \epsilon(\tilde{\mathbf{x}}) \geq 0, \forall i = 1, 2, \dots, l, \end{aligned}$$

Inspired by the formulation of SVM+ and recent advances on LUPI [Niu *et al.*, 2016; Xu *et al.*, 2015; Sharmanska *et al.*, 2013], in this work, we aim to construct a slack function for hashing by incorporating privileged information to learn a more precise hash model in the target domain where data is sparse.

3 Iterative Quantization with Privileged Information (ITQ+)

Suppose that we are given n data points $\{\mathbf{x}_{T_1}, \mathbf{x}_{T_2}, \dots, \mathbf{x}_{T_n}\}$ with $\mathbf{x}_{T_i} \in \mathbb{R}^{d_T}$. Denote by $\mathbf{X}_T = [\mathbf{x}_{T_1}, \dots, \mathbf{x}_{T_n}]^\top \in \mathbb{R}^{n \times d_T}$ the data matrix. Without loss in generality, we assume that all the n points have been zero-centered, i.e., $\sum_{i=1}^n \mathbf{x}_{T_i} = \mathbf{0}$. The goal of learning to hash is to learn a binary code matrix $\mathbf{B}_T \in \mathbb{R}^{n \times c}$ with its elements in $\{-1, 1\}$, where c is the length of each hash code. For each bit $k \in \{1, \dots, c\}$, a binary function $h_k(\mathbf{x}_T) = \text{sgn}(\mathbf{r}_k^\top \mathbf{x}_T)$ is learned, where $\mathbf{r}_k \in \mathbb{R}^{d_T}$ is the hyperplane for the k -th bit. Denote by $\mathbf{R} \in \mathbb{R}^{d_T \times c}$ the projection matrix for all the c bits, the binary code matrix \mathbf{B}_T can be obtained by setting $\mathbf{B}_T = \text{sgn}(\mathbf{X}_T \mathbf{R})$.

When n is small, i.e., the available target training data is limited, the hash codes learned by existing methods may not perform well. How to learn a precise hash model from sparse data is a crucial issue for most existing learning-to-hash algorithms. Inspired by the exciting results of LUPI, which prove that the amount of training data can be significantly reduced with privileged information for training a precise predictive model [Lapin *et al.*, 2014; Vapnik and Vashist, 2009; Vapnik and Izmailov, 2015], we propose a new framework

for learning to hash, namely Transfer Hashing with Privileged Information (THPI). In THPI, the data sparsity issue on the target domain is alleviated by using privileged information from an auxiliary domain which is referred to as the source domain. Apart from a target feature vector \mathbf{x}_T , in THPI, we assume that corresponding privileged information $\mathbf{x}_S \in \mathbb{R}^{d_S}$ from the source domain is available for training as well, which means that there are n corresponding data pairs $\{(\mathbf{x}_{S_1}, \mathbf{x}_{T_1}), (\mathbf{x}_{S_2}, \mathbf{x}_{T_2}), \dots, (\mathbf{x}_{S_n}, \mathbf{x}_{T_n})\}$ for training. Furthermore, we denote by $\mathbf{X}_{SC} = [\mathbf{x}_{S_1}, \dots, \mathbf{x}_{S_n}]^\top \in \mathbb{R}^{n \times d_S}$ the matrix of the n corresponding instances on the source domain, and $\mathbf{X}_{SU} = [\mathbf{x}_{S_{n+1}}, \mathbf{x}_{S_{n+2}}, \dots, \mathbf{x}_{S_{n+n_S}}]^\top \in \mathbb{R}^{n_S \times d_S}$ the matrix of the additional n_S source domain instances. Note that the privileged information is only available for training but not for testing.

3.1 Iterative Quantization

The ITQ algorithm [Gong and Lazebnik, 2011] aims to construct hashing functions using an iterative quantization method to learn a rotation matrix by minimizing the quantization error. Specifically, an orthogonal projection matrix $\mathbf{R}_T \in \mathbb{R}^{d_T \times c}$ is learned with the code matrix $\mathbf{B}_T \in \{-1, 1\}^{n \times c}$ by optimizing the following quantization loss:

$$\begin{aligned} \min_{\mathbf{B}_T, \mathbf{R}_T} \quad & \|\mathbf{B}_T - \mathbf{X}_T \mathbf{R}_T\|_F^2 \\ \text{s.t.} \quad & \mathbf{R}_T^\top \mathbf{R}_T = \mathbf{I}. \end{aligned} \quad (1)$$

3.2 The objective function for ITQ+

Here, we assume that for \mathbf{B}_T , each bit is balanced, i.e., $\mathbf{B}_T \in \mathcal{B} = \{\mathbf{B} \in \{-1, 1\}^{n \times c} \text{ and } \sum_{i=1}^n \mathbf{b}_i = \mathbf{0}\}$, where \mathbf{b}_i is the i -th column of \mathbf{B} . Define $\mathbf{E} = \mathbf{B}_T - \mathbf{X}_T \mathbf{R}_T$, where \mathbf{E} is the error matrix induced by the quantization process. With the privileged data $\mathbf{X}_{SC} \in \mathbb{R}^{n \times c}$, we aim to approximate the quantization error matrix \mathbf{E} by using a slack function $\mathbf{g}(\mathbf{X}_{SC}) = \mathbf{X}_{SC} \mathbf{P}$, where $\mathbf{P} \in \mathbb{R}^{d_S \times c}$ is another orthogonal projection matrix to be learned. Therefore, we formulate Iterative Quantization with privileged information (ITQ+) as

$$\begin{aligned} \min_{\mathbf{B}_T \in \mathcal{B}, \mathbf{R}_T, \mathbf{P}} \quad & \frac{1}{2} \|\mathbf{E}\|_F^2 + \lambda_1 \|\mathbf{E} - \mathbf{g}(\mathbf{X}_{SC})\|_F^2 \\ \text{s.t.} \quad & \mathbf{R}_T^\top \mathbf{R}_T = \mathbf{I}, \text{ and } \mathbf{P}^\top \mathbf{P} = \mathbf{I} \end{aligned} \quad (2)$$

where $\lambda_1 > 0$ is a tradeoff parameter. Note that in SVM+, the privileged information is used to approximate the slack variables, which can be considered as tolerance functions that allow the margin constraints to be violated. Here, in ITQ+, we borrow the high-level idea of SVM+ to use source-domain information to approximate the target-domain quantization error \mathbf{E} . On one hand, the constructed slack function models the difficulty in quantizing the target domain data with privileged information from the source domain. On the other hand, the constructed slack function can provide a way to regularize the quantization error to avoid overfitting when the size of target domain training data is small.

3.3 Optimization

The solution for the optimization problem (2) can be obtained by alternately updating the binary code matrix \mathbf{B}_T and the rotation matrices \mathbf{R}_T and \mathbf{P} . The procedure is summarized in Algorithm 1, and the details are described in this section.

Update \mathbf{B}_T by fixing \mathbf{R}_T and \mathbf{P}

By fixing \mathbf{R}_T and \mathbf{P} , the binary code matrix \mathbf{B}_T can be obtained by solving the following optimization problem,

$$\begin{aligned} \min_{\mathbf{B}_T \in \mathcal{B}} \quad & \|\mathbf{B}_T - \mathbf{X}_T \mathbf{R}_T\|_F^2 \\ & + \lambda_1 \|(\mathbf{B}_T - \mathbf{X}_T \mathbf{R}_T) - \mathbf{X}_{SC} \mathbf{P}\|_F^2. \end{aligned} \quad (3)$$

As \mathbf{R}_T and \mathbf{P} are constant when optimizing \mathbf{B}_T , (3) can be reformulated as

$$\max_{\mathbf{B}_T \in \mathcal{B}} \text{tr}(\mathbf{B}_T (\lambda_1 \mathbf{P}^\top \mathbf{X}_{SC}^\top + (\lambda_1 + 1) \mathbf{R}_T^\top \mathbf{X}_T^\top)), \quad (4)$$

where $\text{tr}(\cdot)$ denotes the trace of a matrix. The solution for (3) can be obtained by sorting the matrix \mathbf{M} , column-wisely, and then projecting the sorted matrix onto the constraint \mathcal{B} , where

$$\mathbf{M} = \lambda_1 \mathbf{P}^\top \mathbf{X}_{SC}^\top + (\lambda_1 + 1) \mathbf{R}_T^\top \mathbf{X}_T^\top.$$

Update \mathbf{R}_T by fixing \mathbf{P} and \mathbf{B}_T

By fixing \mathbf{B}_T and \mathbf{P} , the optimization problem with respect to \mathbf{R}_T can be written as follows,

$$\begin{aligned} \min_{\mathbf{R}_T} \quad & \|\mathbf{B}_T - \mathbf{X}_T \mathbf{R}_T\|_F^2 \\ & + \lambda_1 \|(\mathbf{B}_T - \mathbf{X}_T \mathbf{R}_T) - \mathbf{X}_{SC} \mathbf{P}\|_F^2 \\ \text{s.t.} \quad & \mathbf{R}_T^\top \mathbf{R}_T = \mathbf{I}. \end{aligned} \quad (5)$$

As \mathbf{P} is fixed, and only \mathbf{R}_T is to be optimized, we further rewrite the above optimization problem as

$$\begin{aligned} \min_{\mathbf{R}_T} \quad & \left\| \mathbf{X}_T \mathbf{R}_T - \left(\mathbf{B}_T - \frac{\lambda_1}{\lambda_1 + 1} (\mathbf{X}_{SC} \mathbf{P}) \right) \right\|_F^2 \\ \text{s.t.} \quad & \mathbf{R}_T^\top \mathbf{R}_T = \mathbf{I}, \end{aligned} \quad (6)$$

which is an orthogonal procrustes problem [Schönemann, 1966], and can be solved analytically. To be specific, by performing the Singular Value Decomposition (SVD) on

$$\left(\mathbf{B}_T - \frac{\lambda_1}{\lambda_1 + 1} (\mathbf{X}_{SC} \mathbf{P}) \right)^\top \mathbf{X}_T, \text{ i.e.,}$$

$$\left(\mathbf{B}_T - \frac{\lambda_1}{\lambda_1 + 1} (\mathbf{X}_{SC} \mathbf{P}) \right)^\top \mathbf{X}_T = \hat{\mathbf{S}} \mathbf{\Sigma} \mathbf{S}^\top,$$

we obtain

$$\mathbf{R}_T = \hat{\mathbf{S}} \mathbf{S}^\top. \quad (7)$$

Update \mathbf{P} by fixing \mathbf{R}_T and \mathbf{B}_T

With \mathbf{R}_T and \mathbf{B}_T fixed, we obtain an optimization problem with respect to \mathbf{P} as

$$\begin{aligned} \min_{\mathbf{P}} \quad & \|(\mathbf{B}_T - \mathbf{X}_T \mathbf{R}_T) - \mathbf{X}_{SC} \mathbf{P}\|_F^2 \\ \text{s.t.} \quad & \mathbf{P}^\top \mathbf{P} = \mathbf{I}, \end{aligned} \quad (8)$$

which again is a standard orthogonal procrustes problem, and can be solved analytically as follows,

$$\mathbf{P} = \hat{\mathbf{Q}} \mathbf{Q}^\top, \quad (9)$$

where $(\mathbf{B}_T - \mathbf{X}_T \mathbf{R}_T)^\top \mathbf{X}_{SC} = \hat{\mathbf{Q}} \mathbf{\Lambda} \mathbf{Q}^\top$ is obtained by performing the SVD on $(\mathbf{B}_T - \mathbf{X}_T \mathbf{R}_T)^\top \mathbf{X}_{SC}$.

Algorithm 1 : Alternating optimization procedure for ITQ+ (or LapITQ+)

- 1: Initialize $\mathbf{R}_T^0, \mathbf{P}^0$ to be the random orthogonal matrices, and set $\tau = 0$.
 - 2: **while** not converge **do**
 - 3: Update $\mathbf{B}_T^{\tau+1}$ by solving (4) (or (11) for LapITQ+).
 - 4: Update $\mathbf{R}_T^{\tau+1}$ according to (7).
 - 5: Update $\mathbf{P}^{\tau+1}$ according to (9).
 - 6: $\tau = \tau + 1$.
 - 7: **end while**
-

4 Extension for ITQ+ (LapITQ+)

In ITQ+, only \mathbf{X}_{SC} is used for learning a hashing model for the target domain. In practice, besides \mathbf{X}_{SC} , we may have a large amount of training instances \mathbf{X}_{SU} on the source domain, whose corresponding feature vectors in the target domain are unknown. To fully exploit all the source domain data to learn a more precise hashing model for the target domain, we proposed an extension of ITQ+ in this section, namely LapITQ+.

Our motivation is from multi-view learning, where the underlying graph structures in different views are assumed to be similar [He and Lawrence, 2011]. Intuitively, as we have a large amount of training instances on the source domain, we can learn a precise graph structure for the source domain, and encode the structure as a regularization term for learning the hash codes for the target domain. Specifically, we can first apply ITQ to learn the hash codes \mathbf{B}_S using all the available source domain data $\mathbf{X}_S = [\mathbf{X}_{SC}^\top, \mathbf{X}_{SU}^\top]^\top$ by optimizing the following quantization loss,

$$\begin{aligned} \min_{\mathbf{B}_S, \mathbf{R}_S} \quad & \|\mathbf{B}_S - \mathbf{X}_S \mathbf{R}_S\|_F^2 \\ \text{s.t.} \quad & \mathbf{R}_S^\top \mathbf{R}_S = \mathbf{I}, \end{aligned} \quad (10)$$

where $\mathbf{X}_S \in \mathbb{R}^{(ns+n) \times ds}$, $\mathbf{B}_S \in \{-1, 1\}^{(ns+n) \times c}$ and $\mathbf{R}_S \in \mathbb{R}^{ds \times c}$. This can be done offline in advance.

Next, we construct an adjacency graph \mathbf{G} from the hash codes \mathbf{B}_S as follow: for each code \mathbf{B}_{S_i} (each row of \mathbf{B}_S), connect \mathbf{B}_{S_i} to its k nearest neighbors with a weight of value 1, where Hamming distance is applied. After constructing the adjacency graph, we use it to define the graph Laplacian $\mathbf{L}_C \in \mathbb{R}^{n \times n}$ for the target domain data. Finally, the proposed LapITQ+ method is formulated as

$$\begin{aligned} \min_{\substack{\mathbf{B}_T \in \mathcal{B}, \\ \mathbf{R}_T, \mathbf{P}_T}} \quad & \|\mathbf{E}\|_F^2 + \lambda_1 \|\mathbf{E} - \mathbf{g}(\mathbf{X}_{SC})\|_F^2 + \lambda_2 \text{tr}(\mathbf{B}_T^\top \mathbf{L}_C \mathbf{B}_T) \\ \text{s.t.} \quad & \mathbf{R}_T^\top \mathbf{R}_T = \mathbf{I}, \text{ and } \mathbf{P}^\top \mathbf{P} = \mathbf{I}. \end{aligned}$$

where $\lambda_1, \lambda_2 \geq 0$ are parameters. Compared to (2), the third term in the above objective is to transfer the graph structure from the source domain to the target domain. Note that instead of constructing the graph Laplacian matrix on the original space [Weiss *et al.*, 2008], we construct the graph Laplacian on the binary code space. In this way, we quantify the local properties of the manifold on the source domain and thus naturally transfer the graph structure across domains for learning hash codes. Though learning the hash codes on the

source domain and the constructing adjacency graph can be very expensive, they can be done offline. The procedure to solve the above optimization is the same as that used for ITQ+ except for the update on \mathbf{B}_T .

4.1 Updating \mathbf{B}_T by fixing \mathbf{P} and \mathbf{R}_T

We first relax the constraint $\mathbf{B}_T \in \{-1, 1\}^{n \times c}$ to $\mathbf{B}_T \in [-1, 1]^{n \times c}$ on the feasible domain \mathcal{B} , and obtain the following constrained quadratic programming optimization,

$$\min_{\mathbf{B}_T \in \mathcal{B}} \quad -2\text{tr}(\mathbf{B}_T \mathbf{K}) + \lambda_2 \|\mathbf{B}_T \mathbf{L}\|_F^2. \quad (11)$$

where $\mathbf{K} = ((1 + \lambda_1) \mathbf{R}^\top \mathbf{X}_T^\top) + \lambda_1 \mathbf{P}^\top \mathbf{X}_{SC}^\top$, and $\mathbf{L}_C = \mathbf{L}^\top \mathbf{L}$. Finally, we binarize the codes by $\mathbf{B}_T = \text{sgn}(\mathbf{B}_T)$.

5 Complexity Analysis

The computational cost for proposed algorithms mainly depends on two parts: 1) to optimize the binary codes \mathbf{B}_T and 2) to optimize the orthogonal rotation matrices \mathbf{R}_T and \mathbf{P} . For updating \mathbf{B}_T , the time complexity for ITQ+ which involves sorting is $O(n_T \log(n_T)c)$, and for LapITQ+, which involves QP programming, the time complexity is $O(n_T^3)$. For updating the orthogonal rotation matrices \mathbf{R}_T and \mathbf{P} , the time complexities are bounded by $O(c^2 d_T + d_T^3)$ and $O(c^2 d_S + d_S^3)$ respectively. In transfer learning, n_T is supposed to be not large. Moreover, in learning to hash, the code length c is supposed to be small. The dimensions d_S and d_T can be pre-processed to be small through dimensionality reduction techniques, such as CCA or PCA. Therefore, the overall complexities for ITQ+ and LapITQ+ are reasonably small.

6 Experiments

6.1 Datasets and Experimental Setup

To verify the effectiveness of our proposed approaches, ITQ+ and LapITQ+, we conduct a series of experiments on three benchmark datasets: BBC Collection [Greene and Cunningham, 2006], multilingual Reuters [Amini *et al.*, 2009], and NUS-WIDE [Chua *et al.*, 2009].

BBC Collection was collected for multi-view learning, where each instance is represented by three views. Specifically, this dataset was constructed from a single-view BBC corpora by splitting news article into related ‘‘views’’ of text. On this dataset, we consider View 1 as the source domain, and View 2 as the target domain.

Multilingual Reuters Collection is a text dataset with over 11,000 news articles from 6 categories in 5 languages, e.g., English, French, etc., which are represented by a bag-of-words weighted by TF-IDF. Each document was also translated into the other four languages to construct correspondences. In the experiments, we use the English as the source domain and French as the target domain. Note that the original data is of very high dimensionality, we first perform PCA with 60% energy preserved on the TF-IDF features. After that, we obtain 1131- and 1230-dimensional features for the English and French documents respectively.

NUS-WIDE dataset consists of 269,648 images from 81 concepts with a total number of 5,018 unique tags downloaded from Flickr. Following [Song *et al.*, 2013], we use

150-D color moment for each image. For the corresponding text documents, we use bag-of-word features based on the 5,018 tags provided by NUS-WIDE, and further reduce its dimensionality by LDA to obtain a 60-D textual feature vector for each document. On this dataset, we treat image features as the target data and the text features as the source data.

While different features may lead to different retrieval performances, the evaluation of different features is not the focus of this paper. To simulate the partial cross-domain correspondence setting, we randomly select a fraction of training examples that are of both modalities as correspondences and denote the correspondences ratio by α , i.e., $\alpha = n/(n + n_S)$, and from the remaining data, we randomly selected 10% as the test samples. The parameters λ_1 and λ_2 for the proposed methods are tuned by cross validation in the range of $[0, 0.001, 0.005, \dots, 1, 2]$. We set the maximum number of iterations to be 150. To remove any randomness caused by random selection of training set, the results are averaged over 10 training-testing splits. To assess the performance of different algorithms, following the evaluation protocols in [Gong and Lazebnik, 2011; Raginsky and Lazebnik, 2009], a nominal threshold of the average distance to the 50th nearest neighbor is used to determine whether a database point returned for a given query is considered a true positive. Finally, we adopt the widely used criterion Mean Average Precision (MAP) [Gong and Lazebnik, 2011; Kumar and Udupa, 2011] for evaluation.

6.2 Compared Methods and Evaluation

We first evaluate the performance of different methods by varying the number of hashing bits in the range of $\{8, 16, 32, 64\}$, with fixed $\alpha = 0.5$. The proposed transfer hashing approach is compared with four state-of-the-art hashing methods, i.e., LSH [Andoni and Indyk, 2006], DSH [Jin *et al.*, 2014], CCA-ITQ [Gong and Lazebnik, 2011] and one cross-modal hashing method CVH [Kumar and Udupa, 2011].

- **LSH:** Local Sensitive Hashing (LSH) [Andoni and Indyk, 2006] is based on the a series of random projection to preserve pairwise distances for data points.
- **DSH:** Density Sensitive Hashing (DSH) [Jin *et al.*, 2014] is to exploit the clustering results to generate a set of candidate hash functions and to select the hash functions which can split the data most equally.
- **CCA-ITQ:** For a fair comparison, we utilize the data from two domains by using CCA instead of PCA to generate initializations for ITQ [Gong and Lazebnik, 2011].
- **CVH:** Cross-View Hashing (CVH) [Kumar and Udupa, 2011] extends Spectral Hashing [Weiss *et al.*, 2008] in a cross-modal manner, which maps similar data to similar codes across views based on a similarity graph.

We report MAP over all the test data for different methods in the Table 1. From the results, we can see that the proposed methods ITQ+ and LapITQ+ perform much better than the baselines. The reason is that ITQ+ and LapITQ+ introduce the new slack function to regularize the quantization loss by using privileged information, which is very important for

Table 1: MAP (%) over 10 runs with $\alpha = 0.5$.

BBC						
Bit	LSH	DSH	CCA-ITQ	CVH	ITQ+	LapITQ+
8	14.00	20.28	21.58	17.76	24.32	26.62
16	16.79	23.25	24.69	24.61	27.69	28.06
32	21.71	25.00	27.09	25.36	29.00	30.30
64	24.61	28.36	26.51	27.62	29.58	31.33
Reuters						
Bit	LSH	DSH	CCA-ITQ	CVH	ITQ+	LapITQ+
8	6.37	7.10	8.75	8.46	9.67	10.12
16	6.55	8.12	10.00	9.49	10.96	11.32
32	7.17	8.15	11.93	9.91	12.85	13.51
64	7.24	7.81	13.06	12.97	13.95	14.70
NUS-wide						
Bit	LSH	DSH	CCA-ITQ	CVH	ITQ+	LapITQ+
8	14.86	23.82	33.18	21.15	35.07	35.42
16	20.49	25.69	36.99	24.49	38.51	39.13
32	25.58	33.10	40.68	27.30	42.16	42.85
64	28.50	35.42	42.45	30.57	44.01	45.40

loss generalization, and insensitive to noise or outliers, especially when the target data is sparse. Compared to LSH and DSH, both ITQ and CVH show better performance as they can make use of knowledge of both the source domain and the target domain by projecting them onto a common space. In this way, the source domain data can be utilized to slightly alleviate the data sparsity issue on the target domain. However these two methods still show inferior performance compared to ITQ+ and LapITQ+. As CCA-ITQ simply performs CCA as a preprocessing step, CCA does not explicitly affect the quantization loss during learning of hashing codes. The CVH method extends spectral hashing in a cross-modal manner and learns hashing codes by performing an eigenvalue decomposition, which usually requires a large number of training data.

From the experimental results, we conclude that the proposed slack function is a better way to transfer source domain knowledge for hashing. Most cross-modal hashing methods require a lot of cross-domain data correspondences, and learn hashing functions only on the correspondences. In contrast, LapITQ+ utilizes all source domain data including unparallelled data to learn source-domain hash codes offline, and use the structure underlying these hash codes to regularize the learning of hash codes on the target domain. Finally, we also observe that LapITQ+ outperforms ITQ+ by incorporating data geometry structure, and thus consistently obtain improvement by 1-2% in MAP.

6.3 Training Data Size and Retrieved Sample Size

We randomly select [10%, 30%, 50%, 70%] of the data from the target domain as the training set to evaluate the influence of training size on all the methods. Furthermore, for these data, we are also given the corresponding privileged data during training. Correspondences ratio is set to be $\alpha = [0.1, 0.3, 0.5, 0.7]$ accordingly. Results are reported in the Figure 1. From the figure, we observe that ITQ+ and LapITQ+ outperform other baselines by large margins, especially when the target training data size is small. Although CCA-ITA shows promising results compared with other baselines, it still performs worse than our proposed methods.

In the application of information retrieval, users are usually more interested in precision at the first K returned results. We

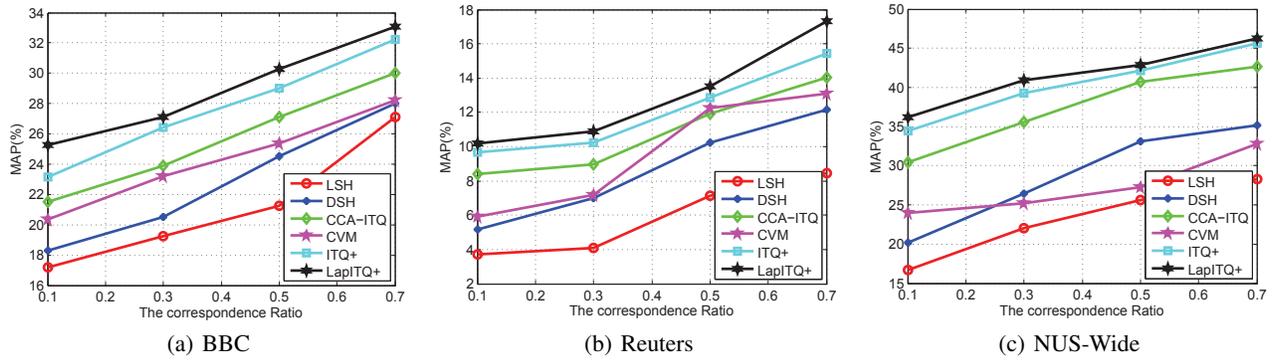


Figure 1: MAP v.s. Privileged Data Size.

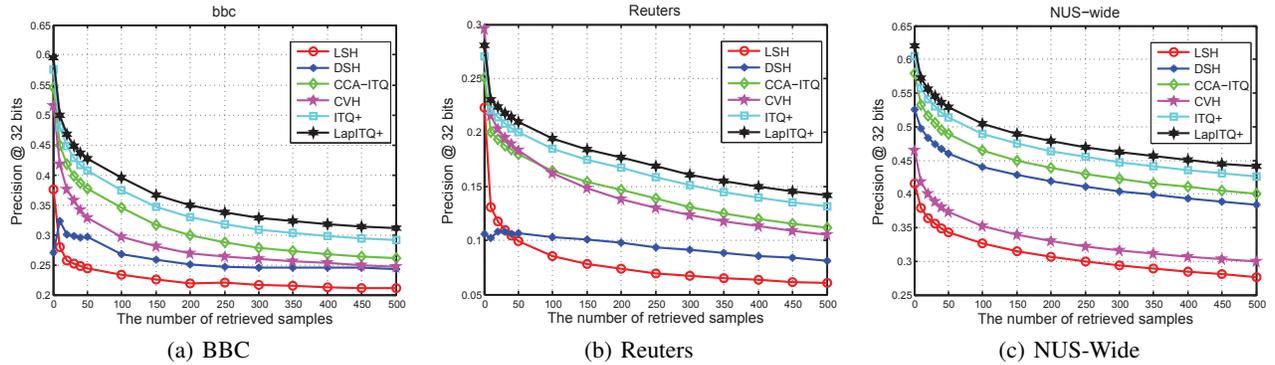


Figure 2: Precision v.s. The Number of Retrieved Samples.

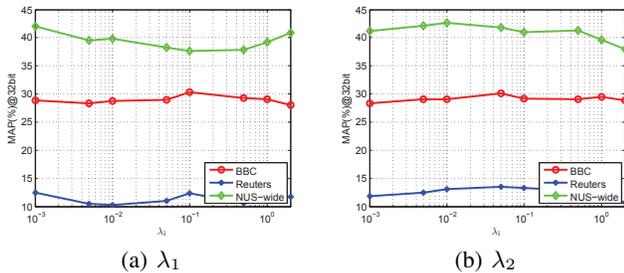


Figure 3: Parameter Analysis.

also report the Top-K precision [Liu *et al.*, 2011] with varying numbers of K retrieved samples on the three datasets with 32 bits in Figure 2. As can be seen from the figure, our proposed methods achieve the best precisions for different values of K . For different numbers of bits, we have similar observations. Thus, we do not report the results here due to space limitation.

6.4 Parameter Analysis

In ITQ+, there is one parameter λ_1 , and in LapITQ+, there are two parameters λ_1 and λ_2 . As LapITQ+ is an extension of ITQ+, we only analyze the parameter sensitivity of LapITQ+ in the range of $[0, 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1, 2]$. We first fix $\lambda_2 = 0.01$ and vary λ_1 . The results of LapITQ+ with 32 bits on the three datasets are shown in the Figure 3(a) with

x -log scale. In the second experiment, we fix $\lambda_1 = 0.01$ and vary λ_2 . The results are reported in the Figure 3(b).¹ We observe that LapITQ+ is not sensitive to λ_1 and λ_2 .

7 Conclusion

In this paper, we propose a new learning framework for hashing named Transfer Hashing with Privileged Information (THPI), where privileged information is used to approximate a slack function to regularize the learning of hashing functions with insufficient data instances in the target domain. Based on the framework, we develop two particular transfer learning methods named ITQ+ and LapITQ+. We conduct extensive experiments on three benchmark datasets. Experimental results verify the superiority of the proposed methods ITQ+ and LapITQ+.

8 Acknowledgement

Sinno Jialin Pan is supported by the NTU Singapore Nanyang Assistant Professorship (NAP) grant M4081532.020. Ivor W. Tsang is grateful for the support from the ARC Future Fellowship FT130100746 and ARC grant LP150100671.

¹The results on the three datasets are 28.49, 11.52, 41.78 with $\lambda_1 = 0$, and 27.56, 13.37, 42.41 with $\lambda_2 = 0$, respectively.

References

- [Amini *et al.*, 2009] Massih-Reza Amini, Nicolas Usunier, and Cyril Goutte. Learning from multiple partially observed views - an application to multilingual text categorization. In *NIPS*, pages 28–36, 2009.
- [Andoni and Indyk, 2006] Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *FOCS*, pages 459–468. IEEE, 2006.
- [Chua *et al.*, 2009] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yan-Tao. Zheng. Nus-wide: A real-world web image database from national university of singapore. In *CIVR*, Santorini, Greece., 2009.
- [Gong and Lazebnik, 2011] Yunchao Gong and Svetlana Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *CVPR*, pages 817–824. IEEE, 2011.
- [Greene and Cunningham, 2006] Derek Greene and Pádraig Cunningham. Practical solutions to the problem of diagonal dominance in kernel document clustering. In *ICML*, pages 377–384. ACM Press, 2006.
- [Hardoon *et al.*, 2004] David R. Hardoon, Sándor Szedmák, and John Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural Computation*, 16(12):2639–2664, 2004.
- [He and Lawrence, 2011] Jingrui He and Rick Lawrence. A graph-based framework for multi-task multi-view learning. In *ICML*, pages 25–32, 2011.
- [Jin *et al.*, 2014] Zhongming Jin, Cheng Li, Yue Lin, and Deng Cai. Density sensitive hashing. *Cybernetics, IEEE Transactions on*, 44(8):1362–1371, 2014.
- [Kumar and Udupa, 2011] Shaishav Kumar and Raghavendra Udupa. Learning hash functions for cross-view similarity search. In *IJCAI*, pages 1360–1365, 2011.
- [Lapin *et al.*, 2014] Maksim Lapin, Matthias Hein, and Bernt Schiele. Learning using privileged information: SVM+ and weighted SVM. *Neural Networks*, 53:95–108, 2014.
- [Liu *et al.*, 2011] Wei Liu, Jun Wang, Sanjiv Kumar, and Shih-Fu Chang. Hashing with graphs. In *ICML*, pages 1–8, 2011.
- [Niu *et al.*, 2016] Li Niu, Xinxing Xu, Lin Chen, Lixin Duan, and Dong Xu. Action and event recognition in videos by learning from heterogeneous web sources. *IEEE Trans. Neural Netw. Learning Syst.*, March 2016.
- [Ou *et al.*, 2014] Xinyu Ou, Lingyu Yan, Hefei Ling, Cong Liu, and Maolin Liu. Inductive transfer deep hashing for image retrieval. In *ACM MM*, pages 969–972, 2014.
- [Pan and Yang, 2010] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.*, 22(10):1345–1359, October 2010.
- [Pan *et al.*, 2011] Sinno Jialin Pan, Ivor W. Tsang, James T. Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *IEEE Trans. Neural Netw.*, 22(2):199–210, 2011.
- [Raginsky and Lazebnik, 2009] Maxim Raginsky and Svetlana Lazebnik. Locality-sensitive binary codes from shift-invariant kernels. In *NIPS*, pages 1509–1517, 2009.
- [Schönemann, 1966] Peter H Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10, 1966.
- [Sharmanska *et al.*, 2013] Viktoriia Sharmanska, Novi Quadrianto, and Christoph H. Lampert. Learning to rank using privileged information. In *ICCV*, pages 825–832, 2013.
- [Song *et al.*, 2013] Jingkuan Song, Yang Yang, Yi Yang, Zi Huang, and Heng Tao Shen. Inter-media hashing for large-scale retrieval from heterogeneous data sources. In *SIGMOD*, pages 785–796. ACM, 2013.
- [Tipping and Bishop, 1999] Michael E Tipping and Christopher M Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- [Vapnik and Izmailov, 2015] Vladimir Vapnik and Rauf Izmailov. Learning using privileged information: Similarity control and knowledge transfer. *J. Mach. Learn. Res.*, 16:2023–2049, 2015.
- [Vapnik and Vashist, 2009] Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural Networks*, 22(5):544–557, 2009.
- [Wang *et al.*, 2016] Jun Wang, Wei Liu, Sanjiv Kumar, and Shih-Fu Chang. Learning to hash for indexing big data - A survey. *Proceedings of the IEEE*, 104(1):34–57, 2016.
- [Weiss *et al.*, 2008] Yair Weiss, Antonio Torralba, and Robert Fergus. Spectral hashing. In *NIPS*, pages 1753–1760, 2008.
- [Wu *et al.*, 2015] Botong Wu, Qiang Yang, Wei-Shi Zheng, Yizhou Wang, and Jingdong Wang. Quantized correlation hashing for fast cross-modal search. In *IJCAI*, pages 3946–3952, 2015.
- [Xu *et al.*, 2015] Xinxing Xu, Wen Li, and Dong Xu. Distance metric learning using privileged information for face verification and person re-identification. *IEEE Trans. Neural Netw. Learning Syst.*, 26(12):3150–3162, 2015.
- [Zhang *et al.*, 2011] Dan Zhang, Fei Wang, and Luo Si. Composite hashing with multiple information sources. In *ACM SIGIR*, pages 225–234. ACM, 2011.
- [Zhou *et al.*, 2014a] Joey Tianyi Zhou, Sinno Jialin Pan, Ivor W. Tsang, and Yan Yan. Hybrid heterogeneous transfer learning through deep learning. In *AAAI*, pages 2213–2220, 2014.
- [Zhou *et al.*, 2014b] Joey Tianyi Zhou, Ivor W. Tsang, Sinno Jialin Pan, and Mingkui Tan. Heterogeneous domain adaptation for multiple classes. In *AISTATS*, pages 1095–1103, 2014.
- [Zhou *et al.*, 2016] Joey Tianyi Zhou, Sinno Jialin Pan, Ivor W. Tsang, and Shen-Shyang Ho. Transfer learning for cross-language text categorization through active correspondences construction. In *AAAI*, 2016.