

6-2014

Hierarchical Feature Hashing for Fast Dimensionality Reduction

Bin Zhao

Carnegie Mellon University, binzhao@andrew.cmu.edu

Eric P. Xing

Carnegie Mellon University, epxing@cs.cmu.edu

Follow this and additional works at: http://repository.cmu.edu/machine_learning



Part of the [Theory and Algorithms Commons](#)

Published In

Proceedings of Computer Vision and Pattern Recognition (CVPR 2014), 2051-2058.

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Machine Learning Department by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Hierarchical Feature Hashing for Fast Dimensionality Reduction

Bin Zhao

Eric P. Xing

School of Computer Science, Carnegie Mellon University

{binzhao, epxing}@cs.cmu.edu

Abstract

Curse of dimensionality is a practical and challenging problem in image categorization, especially in cases with a large number of classes. Multi-class classification encounters severe computational and storage problems when dealing with these large scale tasks. In this paper, we propose hierarchical feature hashing to effectively reduce dimensionality of parameter space without sacrificing classification accuracy, and at the same time exploit information in semantic taxonomy among categories. We provide detailed theoretical analysis on our proposed hashing method. Moreover, experimental results on object recognition and scene classification further demonstrate the effectiveness of hierarchical feature hashing.

1. Introduction

Over the past few years, there have been growing interests in building general purpose image classifiers for real world visual recognition [9, 25, 2, 17, 20], with hundreds or thousands of classes. Accompanying the growth of image number and concept space is the huge dimensionality of image feature vectors. For example, [20] builds more than 130K features to represent each image for a 1000-class problem, and for the same problem, [17] represents each image with more than 600K features. Moreover, results in [17] imply that higher dimensional feature representation results in better image classification.

Worse still, a k -class classification problem defines a parameter vector for each class, and d -dimensional feature representation for each image would result in kd parameters to learn. Take linear classification as an example: given $\mathbf{x} \in \mathbb{R}^d$ and $y \in \mathcal{Y}$ with $|\mathcal{Y}| = k$, the joint input-label mapping [26] could be defined as $\Gamma(\mathbf{x}, y) = \mathbf{e}_y \otimes \mathbf{x}$, where $\Gamma(\mathbf{x}, y) \in \mathbb{R}^{kd}$ is obtained by tensor product of \mathbf{x} and vector $\mathbf{e}_y \in \mathbb{R}^k$, where \mathbf{e}_y has every element as 0, except the y -th element being 1. Linear classifier is obtained by learning parameter vector $\mathbf{w} \in \mathbb{R}^{kd}$, such that the label for image \mathbf{x} is given by $y^* = \arg \max_y \mathbf{w}^\top \Gamma(\mathbf{x}, y)$. For large scale problems, kd could be in the scale of billions. Storing such




Figure 1. Left: class taxonomy structured as a tree, with a given data \mathbf{x} from class *Lily*; Right: the hierarchically hashed feature representation for $(\mathbf{x}, Lily)$, where Φ^{Flower} is the hash function defined uniquely for class *Flower*.

huge number of parameters in memory would be a challenge, let alone the intermediate variables generated during learning. Therefore, how to train multi-class classifiers with high-dimensional feature representation and large concept space, is a practical and challenging research problem.

One fruitful line of research in large scale image categorization tries to exploit the inherent structure in the image categories. Specifically, image classes in large scale problems are rarely organized in a flat fashion, but rather with a hierarchical taxonomy [9, 28]. For example, Figure 1 shows an example of tree hierarchy, where leaf nodes are individual image categories¹, and each internal node denotes the cluster of categories in the subtree rooted by the given node. Having access to this taxonomy structure has been shown to benefit the accuracy of the learning algorithms [19, 11, 10, 29]. However, hierarchical classifiers that define a parameter vector for each node in the taxonomy structure [8, 22] result in even larger number of parameters, than flat multi-class classifiers.

In this paper, we propose *Hierarchical Feature Hashing* (*HESHING*), to aggressively reduce dimensionality, when the taxonomy structure in the image categories is known. Specifically, *HESHING* hashes all input-label pairs (\mathbf{x}, y) into a common lower dimensional space with $\Psi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^{d'}$. A linear classifier parameterized by $\mathbf{w} \in \mathbb{R}^{d'}$, trained in the space defined by Ψ , determines the label y^* for a new image \mathbf{x} as $y^* = \arg \max_y \mathbf{w}^\top \Psi(\mathbf{x}, y)$. The parameter vector for a multi-class classifier therefore lives in $\mathbb{R}^{d'}$ instead of in \mathbb{R}^{kd} , where $d' \ll kd$. More precisely, a unique hash function Φ is defined for each node in the category

¹An image can only be classified to one of the leaf categories.




Figure 2. From left to right: (a) tree with leaves at same depth; (b) tree with leaves at different depths; (c) extended tree for (b).

taxonomy, and the hierarchical feature hash for class y is the vector concatenation of all feature hashes along the path from root node to the leaf node corresponding to y . For example, Figure 1 shows the hierarchical feature hash for input vector \mathbf{x} associated with class *Lily*, where the tree path is $\{\text{Plant}, \text{Flower}, \text{Lily}\}$. In theoretical analysis, we show that *HESHING* preserves inner products between two vectors with high probability, following an exponential tail bound. Moreover, the expected number of features lost due to collision in *HESHING* is minimal. Finally, we show that the classifier learned directly in the hashed space achieves similar effects in exploiting class taxonomy, as much more expensively learned hierarchical classifier.

Notations. We assume class labels are structured as a tree, where all leaf nodes are at the same depth, as shown in Figure 2(a). For class hierarchies where leaf labels appear on different level of the tree, such as the one shown in Figure 2(b), we could always extend the shorter path by replicating its leaf node, as shown in Figure 2(c). Given class taxonomy with depth T , define the set of class labels in level t as \mathcal{Y}_t . Specifically, \mathcal{Y}_1 will only contain the root node of the tree, and \mathcal{Y}_T contains all leaf class labels. For example, for the tree in Figure 2(a), we have that $\mathcal{Y}_1 = \{1\}$, $\mathcal{Y}_2 = \{2, 3\}$ and $\mathcal{Y}_3 = \{4, 5, 6, 7, 8\}$. Moreover, define the tree path for class y from root node to leaf node y as $\mathcal{P}(y)$. For example, class 6 in Figure 2(a) corresponds to tree path $\mathcal{P}(6) = \{1, 2, 6\}$. Finally, for any function f indexed by u , we will use $f^u(x) = f(x, u)$ interchangeably throughout the paper and they mean exactly the same.

1.1. Previous Work

While there has been little work in visual recognition that exploits feature hashing for dimensionality reduction, the following three lines of research are related to our work.

Dimensionality reduction: For image categorization with ultra high-dimensional feature representation, conventional dimensionality reduction [13] or feature selection [12] methods, such as PCA, ICA and manifold-based methods, and L_1 -regularization based feature selection methods, such as Lasso [24] and structured Lasso [14], run out of favor due to high computational cost. Moreover, these methods only reduce the dimensionality of the input feature vectors, but have no mechanism to cope with the high dimensionality resulted from large concept space. On the other hand, randomized projection [21] does not need to learn projection matrix, but loses sparsity of the original feature vectors, and introduces additional overhead to store the projection matrices. However, *HESHING* preserves in-

formation as well as randomized projection, but also preserves sparsity with no need for projection matrix.

Hashing: Feature hashing has been previously studied in binary classification problems [23, 27], and successfully applied to applications such as email spam filtering [27]. To our knowledge, *HESHING* is the first attempt to exploit category taxonomy in hash design for aggressive dimensionality reduction. Moreover, we show in Theorem 2 that the redundancy introduced by *HESHING* greatly reduces information loss due to hash collision. Furthermore, *hierarchical feature hashing* should not be confused with locality-sensitive hashing [1, 18], one of the approximate-nearest-neighbor algorithms, designed for efficient data indexing. In *hierarchical feature hashing*, hashing technique is adopted to reduce dimensionality for multi-class classification, instead of approximate nearest neighbor search.

Hierarchical classification: Given class taxonomy, hierarchical classifiers exploiting category taxonomy have been studied in the literature [15, 8, 6, 2, 30, 4, 3]. These approaches could be roughly separated into two groups: (1) learn each parameter vector separately, or (2) learn all parameter vectors together. However, the first group of methods face the well-known error propagation problem, where errors made close to the root node are propagated through the tree and yield misclassification. On the other hand, learning all parameter vectors together suffers severely from curse of dimensionality, especially on problems with large concept space. However, as shown in Theorem 3, *HESHING* achieves similar effect as hierarchical classifier in exploiting class taxonomy, but with a much smaller memory footprint, and consequently much more favorable on large scale visual recognition problems.

1.2. Summary of Contributions

To conclude the introduction, we summarize the main contributions of this paper as follows. (1). We propose an approach to aggressively reduce the dimensionality of parameter space in multi-class visual recognition, with no learning, preserves feature sparsity and introduces no additional overhead to store any projection matrix. (2). We provide detailed theoretical analysis on the proposed method. (3). Empirical results on object recognition and scene classification demonstrate its effectiveness: better classification accuracy with much smaller memory footprint.

2. Hierarchical Feature Hashing

Given taxonomy structure of image categories (e.g., Figure 1), we impose similarity requirements between feature representations for adjacent labels in the hierarchy. Specifically, assume the joint input-output feature representation for image \mathbf{x} with class y is $\Psi(\mathbf{x}, y)$, with the optimal label predicted as $y^* = \arg \max_y \mathbf{w}^\top \Psi(\mathbf{x}, y)$. In the spirit of Bayesian methods, we require $\Psi(\mathbf{x}, \textit{Lily})$ to be more simi-

lar to $\Psi(\mathbf{x}, \text{Rose})$, than $\Psi(\mathbf{x}, \text{Maple})$ in Figure 1. Following such feature representation, the resulting classifier may tolerate minor mistakes, such as predicting a sibling of the correct label, but avoids gross errors, such as predicting a node in a completely different part of the taxonomy.

Definition 1. Define $\bar{\mathcal{Y}} = \cup_{t=1}^T \mathcal{Y}_t$ as the set of all labels in the taxonomy tree with depth T . For any $u \in \bar{\mathcal{Y}}$, denote by h^u a hash function $h^u : \mathbb{N} \rightarrow \{1, \dots, m\}$, and ξ^u another hash function $\xi^u : \mathbb{N} \rightarrow \{\pm 1\}$, such that $\forall u' \neq u: h^{u'} \neq h^u$ and $\xi^{u'} \neq \xi^u$. For any feature-label pair (\mathbf{x}, y) with $y \in \mathcal{Y}_T$, hierarchical feature hash $\Psi(\mathbf{x}, y)$ is defined as

$$\Psi(\mathbf{x}, y) = \frac{1}{\sqrt{|\mathcal{P}(y)|}} \begin{pmatrix} \Phi(\mathbf{x}, u_1) \\ \vdots \\ \Phi(\mathbf{x}, u_{|\mathcal{P}(y)|}) \end{pmatrix} \quad (1)$$

where $\mathcal{P}(y) = \{u_1, \dots, u_{|\mathcal{P}(y)|}\}$ with $|\mathcal{P}(y)| = T$ is the path from root node to leaf y , and $\Phi(\mathbf{x}, u) \in \mathbb{R}^m$ is defined as $\forall i \in \{1, \dots, m\} : \Phi_i(\mathbf{x}, u) = \sum_{j: h^u(j)=i} \xi^u(j) x_j$.

For hierarchical feature hashing, any hash function could be used, as long as it has uniform distribution among the m hash brackets. We will discuss more details in Section 4. *HESHING* maps any input-output pair $(\mathbf{x}, y) \in \mathbb{R}^d \times \mathcal{Y}$ into a much lower dimensional space, where a single mT -dimensional parameter vector \mathbf{w} is learned for multi-class classification. Since we usually have $mT \ll d|\mathcal{Y}|$, learning multi-class classifier in the hashed space takes much less memory and computational time than in the original feature space. This is especially important for large scale image categorization with high-dimensional feature representation and large concept space. Moreover, for closely related classes sharing common nodes on their tree paths, such as *Lily* and *Rose* in Figure 1, their hierarchical feature hashes are similar, in the sense that parts of their feature vectors are the same. For example, the hierarchical feature hash for $(\mathbf{x}, \text{Lily})$ is the concatenation of $\Phi(\mathbf{x}, \text{Plant})$, $\Phi(\mathbf{x}, \text{Flower})$ and $\Phi(\mathbf{x}, \text{Lily})$, and the hierarchical feature hash for $(\mathbf{x}, \text{Rose})$ only differs in the last part $\Phi(\mathbf{x}, \text{Rose})$. If we denote the learned parameter vector in hashed space as $\mathbf{w} = [\mathbf{w}_1^\top, \dots, \mathbf{w}_T^\top]^\top$, where $\forall t \in \{1, \dots, T\}$, $\mathbf{w}_t \in \mathbb{R}^m$ discriminates between classes at level t . Therefore, *HESHING* congregates discriminative information from all levels of the class taxonomy, and at the same time maintains a very small memory footprint.

HESHING could be used with any classification algorithm. In this paper, we adopt the max-margin classifier [7] and formulate multi-class classification as follows:

$$\begin{aligned} \min_{\mathbf{w}, \xi} \quad & \frac{\lambda}{2} \|\mathbf{w}\|^2 + \frac{1}{n} \sum_{i=1}^n \xi_i \\ \text{s.t.} \quad & \forall i = 1, \dots, n, r = 1, \dots, k \\ & \mathbf{w}^\top \Psi(\mathbf{x}_i, y_i) + \delta_{y_i r} - \mathbf{w}^\top \Psi(\mathbf{x}_i, r) \geq 1 - \xi_i \end{aligned} \quad (2)$$

where $\delta_{y_i r} = 1$ if $y_i = r$ and 0 otherwise. Optimization problem (2) is solved using *stochastic gradient descent* [5].

3. Analysis

In this section, we provide theoretical analysis on *hierarchical feature hashing*.

3.1. Preserving Inner Product

One of the most important requirements for dimensionality reduction is to preserve inner product [21], such that, for any two input data samples, the inner product between their resulting low-dimensional representations is close to that of their corresponding high-dimensional features. The following theorem shows that under hierarchical feature hashing, the inner product between any two vectors is preserved with high probability.

Theorem 1. Given (\mathbf{x}, y) and (\mathbf{x}', y') , where $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$, and $\|\mathbf{x}\|_2 = \|\mathbf{x}'\|_2 = 1$. For hierarchical feature hashing defined in (1), $\forall \epsilon < 1$, if $m \geq \frac{288 \log(6/\delta)}{\epsilon^2 |\mathcal{P}(y)|}$ and $\|\mathbf{x}\|_\infty \leq \sqrt{\frac{3|\mathcal{P}(y)|}{8m \log(6m/\delta)}}$, we have that

$$\mathbb{P}(|\langle \Psi(\mathbf{x}, y), \Psi(\mathbf{x}', y') \rangle - \langle \mathbf{x}, \mathbf{x}' \rangle| \geq \epsilon) \leq \delta \quad (3)$$

Proof. For inner product, we have that

$$\langle \mathbf{x} - \mathbf{x}', \mathbf{x} - \mathbf{x}' \rangle = \|\mathbf{x}\|_2^2 + \|\mathbf{x}'\|_2^2 - 2\langle \mathbf{x}, \mathbf{x}' \rangle \quad (4)$$

$$\Rightarrow \langle \mathbf{x}, \mathbf{x}' \rangle = \frac{1}{2} (\|\mathbf{x}\|_2^2 + \|\mathbf{x}'\|_2^2 - \|\mathbf{x} - \mathbf{x}'\|_2^2) \quad (5)$$

Similarly, for inner product in the hashed space

$$\langle \Psi(\mathbf{x}, y), \Psi(\mathbf{x}', y') \rangle = \frac{1}{2} (\|\mathbf{x}\|_\Psi^2 + \|\mathbf{x}'\|_\Psi^2 - \|\mathbf{x} - \mathbf{x}'\|_\Psi^2) \quad (6)$$

where $\|\mathbf{x}\|_\Psi^2 = \Psi(\mathbf{x}, y)^\top \Psi(\mathbf{x}, y)$. Using (5) and (6), we get

$$\begin{aligned} 2|\langle \Psi(\mathbf{x}, y), \Psi(\mathbf{x}', y') \rangle - \langle \mathbf{x}, \mathbf{x}' \rangle| &= |\|\mathbf{x}\|_\Psi^2 - \|\mathbf{x}\|_2^2 \\ &+ \|\mathbf{x}'\|_\Psi^2 - \|\mathbf{x}'\|_2^2 + \|\mathbf{x} - \mathbf{x}'\|_\Psi^2 - \|\mathbf{x} - \mathbf{x}'\|_2^2| \\ &\leq |\|\mathbf{x}\|_\Psi^2 - \|\mathbf{x}\|_2^2| + |\|\mathbf{x}'\|_\Psi^2 - \|\mathbf{x}'\|_2^2| + |\|\mathbf{x} - \mathbf{x}'\|_\Psi^2 - \|\mathbf{x} - \mathbf{x}'\|_2^2| \end{aligned} \quad (7)$$

We need to bound each of the three terms in the right hand side of (7). Specifically, we will show that

$$\mathbb{P}(2|\|\mathbf{x}\|_\Psi^2 - \|\mathbf{x}\|_2^2| > \epsilon \|\mathbf{x}\|_2^2) \leq \frac{1}{3} \delta \quad (8)$$

$$\mathbb{P}(2|\|\mathbf{x}'\|_\Psi^2 - \|\mathbf{x}'\|_2^2| > \epsilon \|\mathbf{x}'\|_2^2) \leq \frac{1}{3} \delta \quad (9)$$

$$\mathbb{P}(2|\|\mathbf{x} - \mathbf{x}'\|_\Psi^2 - \|\mathbf{x} - \mathbf{x}'\|_2^2| > \epsilon \|\mathbf{x} - \mathbf{x}'\|_2^2) \leq \frac{1}{3} \delta \quad (10)$$

Without loss of generality, we will prove the first inequality, and the other two could be shown similarly. Since $\|\mathbf{x}\|_2 = 1$, we need to show $\mathbb{P}(|\|\mathbf{x}\|_\Psi^2 - 1| > \frac{1}{2}\epsilon) \leq \frac{1}{3}\delta$. Before getting into details of the proof, we first introduce the following two lemmas from [16] and [27]:

Lemma 1. $\forall \mathbf{x} \in \mathbb{R}^d$ with $\|\mathbf{x}\|_2 = 1$, define diagonal matrix $\mathbf{D}_\mathbf{x} \in \mathbb{R}^{d \times d}$ as $(\mathbf{D}_\mathbf{x})_{jj} = x_j$. For any column-normalized matrix $\mathbf{A} \in \mathbb{R}^{m \times d}$, define $\|\mathbf{x}\|_\mathbf{A} = \max_{\mathbf{y}: \|\mathbf{y}\|_2=1} \|\mathbf{A} \mathbf{D}_\mathbf{x} \mathbf{y}\|_2$. For any i.i.d. random ± 1 diagonal matrix \mathbf{D}_s , the following holds: $\forall \mathbf{x}$, if $\|\mathbf{x}\|_\mathbf{A} \leq \frac{\epsilon}{6\sqrt{\log(1/\delta)}}$, then $\mathbb{P}(|\|\mathbf{A} \mathbf{D}_s \mathbf{x}\|_2 - 1| > \epsilon) \leq \delta$.

Lemma 2. Define hash function $h^u(\mathbf{x}) : \mathbb{N} \rightarrow \{1, \dots, m\}$, if \mathbf{x} is such that $\|\mathbf{x}\|_2 = 1$ and $\|\mathbf{x}\|_\infty \leq \sqrt{\frac{3}{8m \log(m/\delta)}}$, then define $\sigma_*^2 = \max_i \sum_{j=1}^d x_j^2 \delta_{ih^u(j)}$ where i ranges over all hash buckets. We have that $\mathbb{P}(\sigma_*^2 > \frac{2}{m}) \leq \delta$.

In hierarchical feature hashing, each node u in the class taxonomy corresponds to a unique pair of hash functions (h^u, ξ^u) . For any $u \in \mathcal{P}(y)$, define $\mathbf{A}^u \in \mathbb{R}^{m \times d}$ as $A_{ij}^u = \delta_{ih^u(j)}$ and \mathbf{D}_s^u as $(\mathbf{D}_s^u)_{jj} = \xi^u(j)$. Define the tree path from root to leaf for label y as $\mathcal{P}(y) = \{u_1, \dots, u_{|\mathcal{P}(y)|}\}$ where u_1 is the root node and $u_{|\mathcal{P}(y)|} = y$ is the leaf. Consequently, we will have a sequence of matrices $\{\mathbf{A}^u\}$ and $\{\mathbf{D}_s^u\}$, where u runs through $\{u_1, \dots, u_{|\mathcal{P}(y)|}\}$. Define

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}^{u_1} & & \\ & \ddots & \\ & & \mathbf{A}^{u_{|\mathcal{P}(y)|}} \end{pmatrix}, \tilde{\mathbf{x}} = \frac{1}{\sqrt{|\mathcal{P}(y)|}} \begin{pmatrix} \mathbf{x} \\ \vdots \\ \mathbf{x} \end{pmatrix} \quad (11)$$

$$\mathbf{D}_s = \begin{pmatrix} \mathbf{D}_s^{u_1} & & \\ & \ddots & \\ & & \mathbf{D}_s^{u_{|\mathcal{P}(y)|}} \end{pmatrix}, \mathbf{y} = \begin{pmatrix} \mathbf{y}^{u_1} \\ \vdots \\ \mathbf{y}^{u_{|\mathcal{P}(y)|}} \end{pmatrix} \quad (12)$$

where $\tilde{\mathbf{x}} \in \mathbb{R}^{d|\mathcal{P}(y)|}$ is obtained by stacking up $|\mathcal{P}(y)|$ replicates of \mathbf{x} , and each $\mathbf{y}^u \in \mathbb{R}^d$ with $u \in \{u_1, \dots, u_{|\mathcal{P}(y)|}\}$, such that $\mathbf{y} \in \mathbb{R}^{d|\mathcal{P}(y)|}$ and $\|\mathbf{y}\|_2 = 1$. Moreover, define diagonal matrix $\mathbf{D}_x \in \mathbb{R}^{d \times d}$ as $(\mathbf{D}_x)_{jj} = x_j$, and $\mathbf{D}_{\tilde{\mathbf{x}}} \in \mathbb{R}^{d|\mathcal{P}(y)| \times d|\mathcal{P}(y)|}$ as $(\mathbf{D}_{\tilde{\mathbf{x}}})_{jj} = \tilde{x}_j$. Clearly, we have that $\mathbf{D}_{\tilde{\mathbf{x}}} = \frac{1}{\sqrt{|\mathcal{P}(y)|}} \text{diag}(\mathbf{D}_x, \dots, \mathbf{D}_x)$, with $|\mathcal{P}(y)|$ replicates of \mathbf{D}_x aligned on its diagonal. Next we will check the conditions of Lemma 1 for $\{\mathbf{A}, \tilde{\mathbf{x}}, \mathbf{D}_s, \mathbf{D}_{\tilde{\mathbf{x}}}\}$. First, each matrix \mathbf{A}^u is column normalized, as exactly one element in each column of \mathbf{A}^u is 1 and all other elements are 0, thus, \mathbf{A}^u is column normalized. Consequently, \mathbf{A} is also column normalized. For $\tilde{\mathbf{x}}$, we have that $\|\tilde{\mathbf{x}}\|_2 = \frac{1}{\sqrt{|\mathcal{P}(y)|}} |\mathcal{P}(y)| \|\mathbf{x}\|_2 = \|\mathbf{x}\|_2 = 1$, and

$$\begin{aligned} \|\mathbf{A} \mathbf{D}_{\tilde{\mathbf{x}}} \mathbf{y}\|_2^2 &= \sum_{u \in \mathcal{P}(y)} \frac{1}{|\mathcal{P}(y)|} \|\mathbf{A}^u \mathbf{D}_x \mathbf{y}^u\|_2^2 \\ &= \sum_{u \in \mathcal{P}(y)} \frac{1}{|\mathcal{P}(y)|} \left[\sum_{i=1}^m \left(\sum_{j=1}^d \delta_{ih^u(j)} x_j y_j^u \right)^2 \right] \\ &= \sum_{u \in \mathcal{P}(y)} \frac{1}{|\mathcal{P}(y)|} \left[\sum_{i=1}^m \left(\sum_{j=1}^d \sqrt{\delta_{ih^u(j)} x_j} \sqrt{\delta_{ih^u(j)} y_j^u} \right)^2 \right] \\ &\leq \sum_{u \in \mathcal{P}(y)} \frac{1}{|\mathcal{P}(y)|} \left[\sum_{i=1}^m \left(\sum_{j=1}^d \delta_{ih^u(j)} x_j^2 \sum_{j=1}^d \delta_{ih^u(j)} y_j^{u2} \right) \right] \quad (13) \end{aligned}$$

$$\leq \sum_{u \in \mathcal{P}(y)} \frac{1}{|\mathcal{P}(y)|} \left[\sum_{i=1}^m \left(\sigma_*^2 \sum_{j=1}^d \delta_{ih^u(j)} y_j^{u2} \right) \right] \quad (14)$$

$$\begin{aligned} &= \sum_{u \in \mathcal{P}(y)} \frac{\sigma_*^2}{|\mathcal{P}(y)|} \left[\sum_{i=1}^m \sum_{j=1}^d \delta_{ih^u(j)} y_j^{u2} \right] \\ &= \sum_{u \in \mathcal{P}(y)} \frac{\sigma_*^2}{|\mathcal{P}(y)|} \|\mathbf{y}^u\|_2^2 = \frac{\sigma_*^2}{|\mathcal{P}(y)|} \quad (15) \end{aligned}$$

where (13) applies Cauchy-Schwarz inequality, (14) is true according to the definition of σ_*^2 and (15) uses the fact that $\|\mathbf{y}\|_2 = 1$. Moreover, $\|\tilde{\mathbf{x}}\|_\infty = \frac{1}{\sqrt{|\mathcal{P}(y)|}} \|\mathbf{x}\|_\infty \leq \sqrt{\frac{3}{8m \log(6m/\delta)}}$. Using Lemma 2, we have

$$\mathbb{P} \left(\|\mathbf{A} \mathbf{D}_{\tilde{\mathbf{x}}} \mathbf{y}\|_2^2 > \frac{2}{m|\mathcal{P}(y)|} \right) \leq \frac{1}{6} \delta \quad (16)$$

or equivalently, $\|\mathbf{A} \mathbf{D}_{\tilde{\mathbf{x}}} \mathbf{y}\|_2^2 \leq \frac{2}{m|\mathcal{P}(y)|}$ with probability at least $1 - \delta/6$. Since this holds for any \mathbf{y} with $\|\mathbf{y}\|_2 = 1$ and $m \geq \frac{288 \log(6/\delta)}{\epsilon^2 |\mathcal{P}(y)|}$, we have with probability at least $1 - \delta/6$,

$$\|\tilde{\mathbf{x}}\|_A \leq \sqrt{\frac{2}{m|\mathcal{P}(y)|}} \leq \frac{\epsilon/2}{6\sqrt{\log(6/\delta)}} \quad (17)$$

Applying Lemma 1 and union bound, we have

$$\mathbb{P} \left(\|\mathbf{A} \mathbf{D}_s \tilde{\mathbf{x}}\|_2 - 1 > \frac{1}{2} \epsilon \right) \leq \frac{1}{3} \delta \quad (18)$$

Furthermore, we have

$$\mathbf{A} \mathbf{D}_s \tilde{\mathbf{x}} = \frac{1}{\sqrt{|\mathcal{P}(y)|}} \begin{pmatrix} \mathbf{A}^{u_1} \mathbf{D}_s^{u_1} \mathbf{x} \\ \vdots \\ \mathbf{A}^{u_{|\mathcal{P}(y)|}} \mathbf{D}_s^{u_{|\mathcal{P}(y)|}} \mathbf{x} \end{pmatrix} \quad (19)$$

Since $\forall u \in \{u_1, \dots, u_{|\mathcal{P}(y)|}\}$ and $\forall i \in \{1, \dots, m\}$, $(\mathbf{A}^u \mathbf{D}_s^u \mathbf{x})_i = \sum_{j=1}^d \delta_{ih^u(j)} \xi^u(j) x_j = \Phi_i(\mathbf{x}, u)$, we have that $\|\mathbf{A} \mathbf{D}_s \tilde{\mathbf{x}}\|_2 = \|\Psi(\mathbf{x}, y)\|_2 = \|\mathbf{x}\|_\Psi$. Therefore, $\mathbb{P}(\|\mathbf{x}\|_\Psi - 1 > \frac{1}{2} \epsilon) \leq \frac{1}{3} \delta$. Similarly, we could prove (9) and (10). Combine the three inequalities together using union bound, we get (3), where we used the fact that $\|\mathbf{x} - \mathbf{x}'\|_2^2 \leq 2$. \square

3.2. Feature Loss due to Collision

Hierarchical feature hashing reduces memory footprint of multi-class classification by potentially projecting multiple features into the same hash bracket, which we call a collision. One of the key fears of using *HESHING* is that information might be lost due to collision, which unavoidably restrains the best-possible prediction. However, *hierarchical feature hashing* utilizes multiple hash functions, each defined for a node in the corresponding tree path. As long as one of the multiple hash functions avoids collision on x_j , the information residing in x_j could be exploited in learning. We first formally define feature loss as follows.

Definition 2. Feature lost due to collision: For any (\mathbf{x}, y) with $\mathbf{x} \in \mathbb{R}^d$, we say feature x_j is lost due to collision, if $\forall u \in \mathcal{P}(y), \exists j' \neq j$, such that $h(j, u) = h(j', u)$.

According to the above definition, information of the j -th feature of \mathbf{x} is considered lost if for all $u \in \mathcal{P}(y)$, x_j is projected to a hash bracket shared by another feature. We study the expected number of features lost due to collision in the following theorem.

Theorem 2. Given a tree-structured class hierarchy, define $\tau = \min_y |\mathcal{P}(y)|$ as length of the shortest tree path. For hierarchical feature hashing defined in (1), the expected number of features lost due to collision is at most

$$\mathcal{E}_c = d \left(\frac{d-1}{m} \right)^\tau \quad (20)$$

Proof. We will first consider one feature, say x_j . For x_j , $\forall u \in \mathcal{P}(y)$, since the hash function h^u has a uniform distribution over all hash brackets, the probability of another feature $j' \neq j$ colliding with j in hash function h^u is $\frac{1}{m}$. Mathematically, this means for $j' \neq j$:

$$\mathbb{P}(h^u(j') = h^u(j)) = \frac{1}{m} \quad (21)$$

Define \mathbb{P}_{ju} as the probability of the j -th feature lost due to collision in h^u , then we have that

$$\begin{aligned} \mathbb{P}_{ju} &= \mathbb{P}(\exists j' \neq j : h^u(j') = h^u(j)) \\ &= \mathbb{P}(\cup_{j' \neq j} h^u(j') = h^u(j)) \\ &\leq \sum_{j' \neq j} \mathbb{P}(h^u(j') = h^u(j)) = \frac{d-1}{m} \end{aligned} \quad (22)$$

where we used union bound. The probability \mathbb{P}_j that j -th feature is lost due to collision, which means that it collides with other features in all $|\mathcal{P}(y)|$ hashes is therefore upper bounded by $\mathbb{P}_{ju}^{|\mathcal{P}(y)|}$. Finally, the expected number of features lost due to collision, \mathcal{E}_c satisfies

$$\mathcal{E}_c = d \cdot \mathbb{P}_j \leq d \left(\frac{d-1}{m} \right)^{|\mathcal{P}(y)|} \leq d \left(\frac{d-1}{m} \right)^\tau \quad (23)$$

since $\tau = \min_y |\mathcal{P}(y)|$ and we usually have $m > d$. \square
Let's take the feature representation in [20] as an example, where $d = 130K$ and $\tau = 5$. Assume we use $m = 2^{20}$, which is only 0.8% of kd , the number of parameters if we learn in the original feature space. According to Theorem 2, HESHING loses less than 4 features due to collision.

3.3. Comparison with Hierarchical Classifier

Given taxonomy structure for image categories, one could also build hierarchical classifier to exploit the taxonomy. One group of approaches [22, 8] formulate as follows:

$$\begin{aligned} \min_{\{\mathbf{w}_u\}, \xi} & \frac{\lambda}{2} \sum_{t=1}^T \sum_{u \in \mathcal{Y}_t} \|\mathbf{w}_u\|^2 + \frac{1}{n} \sum_{i=1}^n \xi_i \\ \text{s.t.} & \forall i = 1, \dots, n, r = 1, \dots, k \\ & \frac{1}{\sqrt{T}} \sum_{u \in \mathcal{P}(y_i)} \mathbf{w}_u^\top \mathbf{x}_i + \delta_{y_i r} - \frac{1}{\sqrt{T}} \sum_{v \in \mathcal{P}(r)} \mathbf{w}_v^\top \mathbf{x}_i \geq 1 - \xi_i \end{aligned} \quad (24)$$

where each node u in the class taxonomy corresponds to a unique parameter vector \mathbf{w}_u , and the prototype for each image category is the sum of parameter vectors along the tree path. For example in Figure 2(a), the prototype for class 6 is $\mathbf{W}_6 = \mathbf{w}_1 + \mathbf{w}_2 + \mathbf{w}_6$. The following theorem compares (24) with HESHING based multi-class classification.

Theorem 3. $\forall \{\mathbf{w}_u\}_{u \in \mathcal{Y}_t, t \in \{1, \dots, T\}}$ with $\mathbf{w}_u \in \mathbb{R}^d$, define

$$\mathbf{w}_h = \begin{pmatrix} \sum_{u \in \mathcal{Y}_1} \Phi(\mathbf{w}_u, u) \\ \vdots \\ \sum_{u \in \mathcal{Y}_T} \Phi(\mathbf{w}_u, u) \end{pmatrix} \quad (25)$$

Denote objective value of problem (24) for $\{\mathbf{w}_u\}$ with the smallest feasible $\sum_{i=1}^n \xi_i$ as $O_{HieSVM}(\{\mathbf{w}_u\})$, and objective value of problem (2) for \mathbf{w}_h with the smallest feasible $\sum_{i=1}^n \xi_i$ as $O_{HieHash}(\mathbf{w}_h)$. Then $|O_{HieSVM}(\{\mathbf{w}_u\}) - O_{HieHash}(\mathbf{w}_h)|$ is minimal. Moreover, the decision values according to $\{\mathbf{w}_u\}$ and \mathbf{w}_h are close.

Proof. We first introduce following Lemma 3 to upper bound the interaction between different hash mappings.

Lemma 3. Let $\Phi^u = (h^u, \xi^u)$ and $\Phi^v = (h^v, \xi^v)$, $u \neq v$, be two different hash functions. Then for any $\mathbf{w}, \mathbf{x} \in \mathbb{R}^d$, the inner product $\langle \Phi^u(\mathbf{w}), \Phi^v(\mathbf{x}) \rangle$ is bounded by

$$\mathbb{P}(|\langle \Phi^u(\mathbf{w}), \Phi^v(\mathbf{x}) \rangle| > \epsilon) \leq 2e^{-\frac{m\epsilon^2/2}{\|\mathbf{w}\|_2^2 \|\mathbf{x}\|_2^2 + m\epsilon \|\mathbf{w}\|_\infty \|\mathbf{x}\|_1/3}} \quad (26)$$

Due to space limit, we move proof of Lemma 3 to the supplementary material. Let's start with Problem (24), given $\{\mathbf{w}_u\}$, each ξ_i could be optimized individually,

$$\xi_i^* = \max_{r=1, \dots, k} \left\{ 1 - \delta_{y_i r} - \frac{1}{\sqrt{T}} \sum_{u \in \mathcal{P}(y)} \mathbf{w}_u^\top \mathbf{x}_i + \frac{1}{\sqrt{T}} \sum_{v \in \mathcal{P}(r)} \mathbf{w}_v^\top \mathbf{x}_i \right\} \quad (27)$$

Consequently, the optimal objective value corresponding to $\{\mathbf{w}_u\}$ could be computed as $O_{HieSVM}(\{\mathbf{w}_u\}) = \frac{\lambda}{2} \sum_{t=1}^T \sum_{u \in \mathcal{Y}_t} \|\mathbf{w}_u\|_2^2 + \frac{1}{n} \sum_{i=1}^n \xi_i^*$. Similarly, for Problem (2), $O_{HieHash}(\mathbf{w}_h) = \frac{\lambda}{2} \|\mathbf{w}_h\|_2^2 + \frac{1}{n} \sum_{i=1}^n \xi_i^{**}$, with each ξ_i also optimized individually,

$$\xi_i^{**} = \max_{r=1, \dots, k} \left\{ 1 - \delta_{y_i r} - \mathbf{w}_h^\top \Psi(\mathbf{x}_i, y_i) + \mathbf{w}_h^\top \Psi(\mathbf{x}_i, r) \right\} \quad (28)$$

Next we will show that both $|\sum_{t=1}^T \sum_{u \in \mathcal{Y}_t} \|\mathbf{w}_u\|_2^2 - \|\mathbf{w}_h\|_2^2|$ and $|\xi_i^* - \xi_i^{**}|$ are minimal. Specifically,

$$\begin{aligned} & \left| \sum_{t=1}^T \sum_{u \in \mathcal{Y}_t} \|\mathbf{w}_u\|_2^2 - \|\mathbf{w}_h\|_2^2 \right| \\ &= \left| \sum_{t=1}^T \sum_{u \in \mathcal{Y}_t} \|\mathbf{w}_u\|_2^2 - \sum_{t=1}^T \left\| \sum_{u \in \mathcal{Y}_t} \Phi(\mathbf{w}_u, u) \right\|_2^2 \right| \end{aligned} \quad (29)$$

$$\begin{aligned} &= \left| \sum_{t=1}^T \sum_{u \in \mathcal{Y}_t} (\|\mathbf{w}_u\|_2^2 - \|\Phi(\mathbf{w}_u, u)\|_2^2) \right. \\ &\quad \left. - 2 \sum_{t=1}^T \sum_{u \neq v: u, v \in \mathcal{Y}_t} \Phi(\mathbf{w}_u, u)^\top \Phi(\mathbf{w}_v, v) \right| \end{aligned} \quad (30)$$

$$\approx \left| \sum_{t=1}^T \sum_{u \in \mathcal{Y}_t} (\|\mathbf{w}_u\|_2^2 - \|\Phi(\mathbf{w}_u, u)\|_2^2) \right| \quad (31)$$

$$\leq \sum_{t=1}^T \sum_{u \in \mathcal{Y}_t} \|\mathbf{w}_u\|_2^2 - \|\Phi(\mathbf{w}_u, u)\|_2^2 \approx 0 \quad (32)$$

where (31) is true according to Lemma 3, and (32) is true by applying Theorem 1 to a single node tree with only root node u . Similarly, $\forall r \in \{1, \dots, k\}$:

$$\left| \frac{1}{\sqrt{T}} \sum_{u \in \mathcal{P}(r)} \mathbf{w}_u^\top \mathbf{x}_i - \mathbf{w}_h^\top \Psi(\mathbf{x}_i, r) \right| = \left| \frac{1}{\sqrt{T}} \sum_{u \in \mathcal{P}(r)} \mathbf{w}_u^\top \mathbf{x}_i - \frac{1}{\sqrt{T}} \sum_{t=1}^T \sum_{u \in \mathcal{Y}_t} \Phi(\mathbf{w}_u, u)^\top \Phi(\mathbf{x}_i, y_i^t) \right| \quad (33)$$

$$\approx \frac{1}{\sqrt{T}} \left| \sum_{u \in \mathcal{P}(r)} \mathbf{w}_u^\top \mathbf{x}_i - \sum_{t=1}^T \Phi(\mathbf{w}_{y_i^t}, y_i^t)^\top \Phi(\mathbf{x}_i, y_i^t) \right| \quad (34)$$

$$\leq \frac{1}{\sqrt{T}} \sum_{u \in \mathcal{P}(r)} \left| \mathbf{w}_u^\top \mathbf{x}_i - \Phi(\mathbf{w}_u, u)^\top \Phi(\mathbf{x}_i, u) \right| \approx 0 \quad (35)$$

where (34) is true since $\forall u \neq y_i^t$: $|\Phi(\mathbf{w}_u, u)^\top \Phi(\mathbf{x}_i, y_i^t)| \approx 0$ according to Lemma 3, and (35) is true according to Theorem 1. Consequently, we have that $|\xi_i^* - \xi_i^{**}| \approx 0$. Therefore, we have that for any $\{\mathbf{w}_u\}_{u \in \mathcal{Y}_t, t \in \{1, \dots, T\}}$, and corresponding \mathbf{w}_h defined in (25),

$$|O_{HierSVM}(\{\mathbf{w}_u\}) - O_{HierHash}(\mathbf{w}_h)| \approx 0 \quad (36)$$

Therefore, parameter vector \mathbf{w}_h constructed in (25) is approximately the optimal solution to Problem (2), i.e., the parameter vector learned in hierarchical feature hashing. Moreover, given a new data point $\mathbf{x} \in \mathbb{R}^d$, the prediction task now consists of calculating $\mathbf{w}_h^\top \Psi(\mathbf{x}, r)$ for any $r \in \mathcal{Y}_T$. The results in (32) shows that the decision value computed in hashed feature space is very close to the decision value, $\frac{1}{\sqrt{T}} \sum_{u \in \mathcal{P}(r)} \mathbf{w}_u^\top \mathbf{x}_i$, computed in the original feature space using hierarchical classifier. \square

Moreover, since the weight vector actually learned in the hashed space has information about the collision of features, it could be trained to handle such collisions by avoiding putting heavy weights on those ambiguous features, hence resulting in even better classification.

4. Experiments

In this section, we test *hierarchical feature hashing* on two fundamental computer vision tasks: object recognition and scene classification.

4.1. Basic Setup

We first describe the evaluation metrics in terms of classification accuracy and memory footprint. For the classification accuracy, we report two measures: flat error and hierarchical error. Specifically, given an image, flat error equals 1 if the predicted class is not its true class, and 0 otherwise. On the other hand, hierarchical error reports the minimum height of the lowest common ancestor between true and predicted classes. For example in Figure 2(a), if the predicted class is 4 while true class is 8, hierarchical error

Data set	T	#Class	#Train	#Test
amphibian	4	8	8857	1200
fish	8	11	12503	1650
furniture	4	23	32959	3450
geo	5	12	19804	1800
music	6	25	38013	3750
reptile	7	29	29851	4350
tool	6	26	31087	3900
SUN	3	397	19850	19850

Table 1. Data sets details.

for this image is 2. For both criteria, the overall error score for an algorithm is the average over all test images. For memory footprint, we report the amount of memory taken during training.

We compare our method to three different feature representations. The first setting learns multi-class classifiers in the original feature space: (1). one-vs-rest classification (*OVR*), with each binary classifier trained using SVM; (2). multi-class SVM (*MSVM*) [7]; (3). hierarchical SVM (*HieSVM*) [8]; (4). tree classifier (*TreeSVM*) through a top-down approach, which trains a multi-class SVM at each node in the class taxonomy [15]. The second setting *FlatHash* also solves Problem 2 for multi-class classifier, but applies flat feature hashing using only leaf class, ignoring class taxonomy structure. For example in Figure 1(a), the flat feature hashing for $(\mathbf{x}, Lily)$ is just $\Phi(\mathbf{x}, Lily)$. The third setting applies PCA for dimensionality reduction, and utilizes one-vs-rest classification (*PCA*). All classifiers are trained with *stochastic gradient descent*.

We adopt MurmurHash² as the hash function in both hierarchical feature hashing and flat feature hashing. To obtain a unique hash function Φ^u for each class u , we set the seed in MurmurHash to u and compute the hash. However, it should be noted that our result is not dependent on specific choice of hash implementations.

4.2. Object Recognition on ImageNet

We first experiment on the object recognition task on ImageNet [9], consisting of images collected from the web that are organized according to the WordNet hierarchy. Specifically, we report results on 7 subtrees of ImageNet that are used in the study in [9]: *amphibian*, *fish*, *furniture*, *geological formation (geo)*, *musical instrument (music)*, *reptile* and *tool*. The training, validation and test sets for these subtrees are taken from the 2010 ImageNet Large Scale Visual Recognition Challenge data set. We used two feature representations for each image: (1) low dimensional feature composed using SIFT-based bag-of-words representation, and (2) high dimensional feature computed the same way as in [17]. Specifically, for both feature representations, we start with computing dense SIFT descriptors for each image, and then run *k-means* clustering on a random subset

²<https://sites.google.com/site/murmurhash/>

of 1 million SIFT descriptors to form a visual vocabulary. For the low dimensional features, we generate 1000 visual words and build bag-of-words feature, such that each image is represented as a 1000-dimensional vector. On the other hand, for the high dimensional features, we get a visual vocabulary of 8192 visual words. Using the learned vocabulary, we employ *Locality-constrained Linear Coding (LLC)* [17] for feature coding. Finally, a single feature vector is computed for each image using max pooling on a spatial pyramid, with each image represented as a 170016-dimensional vector. For *HESHING* and *FlatHash*, we fix the hash size $m = 2^{11}$ for low dimensional feature representation, and $m = 2^{18}$ for the high dimensional feature representation. It should be noted that $m = 2^{11}$ still effectively reduces dimensionality of parameter space, as all non-hash based methods will encounter at least kd parameters, where d is original feature dimension and k is the number of classes. For *PCA*, we fix the reduced dimensionality to 500 for the low dimensional feature representation, and to 1000 for the high dimensional one. The reason for picking such dimensionality for *PCA* is to ensure highest possible accuracy without running into out-of-memory problem, as *PCA* needs to store a $d \times d'$ projection matrix, with d being the original feature dimension and d' the reduced one.

As shown in Table 2, *HESHING* achieves comparable classification results with the best competing algorithm across all data sets, and in many cases even better classification accuracy than all other approaches. The fact that *HESHING* provides equally good or better classification results than classifiers trained in the original feature space, with or without knowledge of the class taxonomy, demonstrates that *HESHING* does not lose critical discriminative information after aggressive dimensionality reduction. Moreover, *HESHING* not only outperforms *FlatHash* in hierarchical error, but also generates better flat error. Therefore, exploiting class taxonomy in feature hashing not only benefits in achieving lower hierarchical errors, but also introduces redundancy in the hashed representation, effectively reducing information loss due to hash collision, as shown in Theorem 2, and thus enables lower flat errors.

Moreover, Figure 3 shows the memory footprint of various algorithms on the *furniture* data set. Clearly, *HESHING* requires much less memory than other non-hash based methods, rendering it especially attractive in large-scale visual recognition problems, where hashing is necessary to make training multi-class classifier computationally feasible at all, as competing methods would run out of memory. Finally, the memory footprints for *HESHING* and *FlatHash* are close, while *HESHING* outperforms *FlatHash* on classification accuracy according to Table 2.

Furthermore, we investigate the influence of hash size m on the classification errors. Specifically, we run *HESHING* on the high dimensional feature representation, with




Figure 3. Memory footprint comparison on data set furniture (left) and SUN (right). Please note that the vertical axis is in log scale.

various values of m . As shown in Figure 4, when the hash size m decreases, both flat error and hierarchical error increase. Moreover, take the *furniture* data set as an example, according to Theorem 2, a 21 bit hash table causes almost no collisions. Nonetheless, an 18 bit hash which has almost 20% collisions performs equally well on the problem. This leads to rather memory efficient implementations.




Figure 4. Impact of code length on flat error (left) and hierarchical error (right). Horizontal axis is the number of bits in hash table, i.e., number 18 on the axis means $m = 2^{18}$.

4.3. Scene Recognition on SUN

Next we experiment on the scene classification task on the SUN data set [28], by far the largest scene recognition data set. It captures a full variety of 899 scene categories. We used 397 well-sampled categories to run the experiment as in [28]. For each class, 50 images are used for training and the other 50 for test. For image features, we adopt the same setting as the high dimensional feature representation used in ImageNet experiment, and represent each image as a 170016-dimensional vector. For *HESHING* and *FlatHash*, we fix the hash size $m = 2^{19}$. For *PCA*, we set the reduced dimensionality to 1000. Classification errors and memory footprint are summarized in Table 2 and Figure 3, with similar observations as in the study on ImageNet.

5. Conclusions and Discussions

Hierarchical feature hashing aggressively reduces dimensionality of parameter space, and exploits class taxonomy to reduce information loss and avoids high cost errors of labeling an image to a class far away from its true category. Besides the empirical study discussed in this paper, another promising application of *HESHING* is on mobile visual search systems like Google Goggles. Since a mobile device faces more strict constraint on memory usage, the fact that *HESHING* requires no projection matrix makes it

Algorithm		amphibian	fish	furniture	geo	music	reptile	tool	SUN
OVR	flat	0.61 0.50	0.60 0.51	0.52 0.47	0.47 0.40	0.76 0.74	0.85 0.76	0.73 0.68	0.83
	hie	1.89 1.57	3.57 3.12	2.27 2.11	2.09 1.72	4.44 4.31	4.59 4.18	3.87 3.44	1.62
MSVM	flat	0.62 0.58	0.66 0.62	0.58 0.53	0.50 0.41	0.75 0.72	0.83 0.78	0.80 0.74	0.86
	hie	1.94 1.76	3.98 3.81	2.68 2.35	2.25 1.88	4.36 4.20	4.38 4.22	4.93 3.92	1.78
HieSVM	flat	0.62 0.59	0.63 0.61	0.53 0.52	0.50 0.43	0.75 0.75	0.88 0.80	0.74 0.69	0.89
	hie	1.86 1.72	3.28 3.24	2.14 2.10	2.12 1.87	4.33 4.31	4.52 4.49	3.76 3.47	1.74
TreeSVM	flat	0.69 0.65	0.75 0.69	0.64 0.58	0.52 0.44	0.79 0.74	0.87 0.83	0.85 0.77	0.88
	hie	2.42 2.23	4.69 4.08	3.16 2.77	2.36 1.96	4.61 4.26	4.79 4.61	5.51 4.20	1.78
FlatHash	flat	0.66 0.51	0.68 0.57	0.61 0.49	0.51 0.41	0.74 0.73	0.86 0.74	0.82 0.71	0.86
	hie	2.10 1.78	4.10 3.62	2.87 2.23	2.30 1.83	4.29 4.24	4.75 4.16	5.18 3.55	1.75
PCA	flat	0.60 0.49	0.62 0.53	0.53 0.46	0.50 0.44	0.78 0.74	0.86 0.75	0.76 0.66	0.84
	hie	1.83 1.56	3.71 3.37	2.33 2.10	2.28 1.98	4.59 4.33	4.74 4.25	4.09 3.39	1.68
HESHING	flat	0.60 0.49	0.56 0.49	0.53 0.47	0.48 0.42	0.69 0.65	0.81 0.74	0.74 0.65	0.82
	hie	1.82 1.50	3.22 3.01	2.35 2.08	2.16 1.80	3.78 3.49	4.26 4.07	3.95 3.26	1.51

Table 2. Misclassification for various algorithms, where *flat* means flat error, *hie* stands for hierarchical error. For ImageNet data sets, the number before ‘|’ corresponds to the low dimensional feature representation, and the number after is for high dimensional features.

very attractive in such applications.

Acknowledgements

This research is supported by Google, NSF IIS-0713379, ONR N000140910758, and AFOSR FA9550010247.

References

- [1] A. Andoni and P. Indyk. Near-optimal hashing algorithms for near neighbor problem in high dimensions. In *FOCS*, 2006. **2**
- [2] S. Bengio, J. Weston, and D. Grangier. Label embedding trees for large multi-class tasks. In *NIPS*, 2010. **1, 2**
- [3] A. Beygelzimer, J. Langford, Y. Lifshits, G. Sorkin, and A. Strehl. Conditional probability tree estimation analysis and algorithms. In *UAI*, 2009. **2**
- [4] A. Beygelzimer, J. Langford, and P. Ravikumar. Error-correcting tournaments. In *ALT*, 2009. **2**
- [5] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *COMPSTAT*, 2010. **3**
- [6] L. Cai and T. Hofmann. Hierarchical document categorization with support vector machines. In *CIKM*, 2004. **2**
- [7] K. Crammer and Y. Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *JMLR*, 2:265–292, 2001. **3, 6**
- [8] O. Dekel, J. Keshet, and Y. Singer. Large margin hierarchical classification. In *ICML*, 2004. **1, 2, 5, 6**
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. **1, 6**
- [10] R. Fergus, H. Bernal, Y. Weiss, and A. Torralba. Semantic label sharing for learning with many categories. In *ECCV*, 2010. **1**
- [11] G. Griffin and P. Perona. Learning and using taxonomies for fast visual categorization. In *CVPR*, 2008. **1**
- [12] I. Guyon and A. Elisseeff. An introduction to variable and feature selection. *JMLR*, 3:1157–1182, 2003. **2**
- [13] A. Jain, R. Duin, and J. Mao. Statistical pattern recognition: A review. *TPAMI*, 22(1), 2000. **2**
- [14] S. Kim and E. Xing. Tree-guided group lasso for multi-task regression with structured sparsity. In *ICML*, 2010. **2**
- [15] D. Koller and M. Sahami. Hierarchically classifying documents using very few words. In *ICML*, 1997. **2, 6**
- [16] E. Liberty, N. Ailon, and A. Singer. Dense fast random projections and lean Walsh transforms. In *APPROX/RANDOM*, 2008. **3**
- [17] Y. Lin, F. Lv, S. Zhu, M. Yang, T. Cour, K. Yu, L. Cao, and T. Huang. Large-scale image classification: fast feature extraction and svm training. In *CVPR*, 2011. **1, 6, 7**
- [18] W. Liu, J. Wang, S. Kumar, and S. Chang. Hashing with graphs. In *ICML*, 2011. **2**
- [19] M. Marszalek and C. Schmid. Semantic hierarchies for visual object recognition. In *CVPR*, 2007. **1**
- [20] F. Perronnin, Z. Akata, Z. Harchaoui, and C. Schmid. Towards good practice in large-scale learning for image classification. In *CVPR*, 2012. **1, 5**
- [21] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *NIPS*, 2008. **2, 3**
- [22] R. Salakhutdinov, A. Torralba, and J. Tenenbaum. Learning to share visual appearance for multiclass object detection. In *CVPR*, 2011. **1, 5**
- [23] Q. Shi, J. Petterson, J. Langford, A. Smola, and A. Strehl. Hash kernels. In *AISTATS*, 2009. **2**
- [24] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Royal. Statist. Soc B*, 58:267–288, 1996. **2**
- [25] A. Torralba, R. Fergus, and W. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *PAMI*, 30:1958–1970, 2008. **1**
- [26] I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. *JMLR*, 6:1453–1484, 2005. **1**
- [27] K. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, 2009. **2, 3**
- [28] J. Xiao, J. Hays, K. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010. **1, 7**
- [29] B. Zhao, L. Fei-Fei, and E. Xing. Large-scale category structure aware image categorization. In *NIPS*, 2011. **1**
- [30] D. Zhou, L. Xiao, and M. Wu. Hierarchical classification via orthogonal transfer. In *ICML*, 2011. **2**