

## Fast and accurate text classification via multiple linear discriminant projections

Soumen Chakrabarti, Shourya Roy, Mahesh V. Soundalgekar

IIT Bombay; e-mail: soumen@cse.iitb.ac.in

Edited by Y. Joannidis. Received: September 9, 2002 / Accepted: March 3, 2003

Published online: July 16, 2003 – © Springer-Verlag 2003

**Abstract.** Support vector machines (SVMs) have shown superb performance for text classification tasks. They are accurate, robust, and quick to apply to test instances. Their only potential drawback is their training time and memory requirement. For  $n$  training instances held in memory, the best-known SVM implementations take time proportional to  $n^a$ , where  $a$  is typically between 1.8 and 2.1. SVMs have been trained on data sets with several thousand instances, but Web directories today contain millions of instances that are valuable for mapping billions of Web pages into Yahoo!-like directories. We present SIMPL, a nearly linear-time classification algorithm that mimics the strengths of SVMs while avoiding the training bottleneck. It uses Fisher’s linear discriminant, a classical tool from statistical pattern recognition, to project training instances to a carefully selected low-dimensional subspace before inducing a decision tree on the projected instances. SIMPL uses efficient sequential scans and sorts and is comparable in speed and memory scalability to widely used naive Bayes (NB) classifiers, but it beats NB accuracy decisively. It not only approaches and sometimes exceeds SVM accuracy, but also beats the running time of a popular SVM implementation by orders of magnitude. While describing SIMPL, we make a detailed experimental comparison of SVM-generated discriminants with Fisher’s discriminants, and we also report on an analysis of the cache performance of a popular SVM implementation. Our analysis shows that SIMPL has the potential to be the method of choice for practitioners who want the accuracy of SVMs and the simplicity and speed of naive Bayes classifiers.

**Keywords:** Text classification – Discriminative learning – Linear discriminants

### 1 Introduction

Text classification is a well-studied problem in document management. A *classifier* or *learner* is first presented with *training* documents  $d$ , each assigned a label,  $c$ , drawn from two possible labels:  $+1$  or  $-1$ . Depending on the application, the label may indicate some property of the document, e.g., whether a

news article is about sustainable energy or whether an email is a “spam.”

The learner processes the training documents, generally collecting term statistics and estimating various model parameters. Later, *test* instances are presented without the label, and the learner has to choose one of the two labels for each test document.

If an application demands more than two labels (e.g., a Web directory with 15 broad topics at the top level), it is common to train one learner for each topic  $\gamma$ ; documents marked with  $\gamma$  are labeled  $+1$  and all other documents are labeled  $-1$  [8, 14]. This is called “one-vs.-rest” classification.

Text classification has numerous potential applications including the automatic maintenance of topic directories such as the Open Directory (also called Dmoz, see <http://dmoz.org>) and Yahoo! (<http://www.yahoo.com>), filtering email for spam [31], and collaborative filtering [2]. Naive Bayes (NB) [3], maximum entropy (maxent) [28] and support vector machines (SVMs) [8,14,37] are some of the best-known classifiers employed to date on text data.

Not surprisingly, there is a trade-off between simplicity and accuracy. NB classifiers are simple to understand and easy to implement, access disk-resident data efficiently, and run fast, but they may show mediocre accuracy. SVMs are among the most accurate classifiers known for text applications: they beat NB accuracy by a decisive margin and are generally better than maxent classifiers. NB classifiers also tend to score lower than maxent classifiers in terms of accuracy. Such accuracy differences are intriguing because SVM, maxent, and NB classifiers all learn a hyperplane that separates the positive examples ( $c = 1$ ) from the negative ones ( $c = -1$ ), documents being represented as vectors in a high-dimensional term space. SVM and maxent undertake complex nonlinear numeric optimizations (which are highly nontrivial to understand and implement) to search for a high-quality separator, whereas NB makes a quick but generally inferior choice.

NB takes time essentially linear in the number  $n$  of training documents [25,26], whereas SVMs take time proportional to  $n^a$ , where  $a$  is typically between 1.8 and 2.1. Thanks to some clever implementations [15,30], SVMs have been trained on several thousand instances despite their near-quadratic com-

plexity. However, scaling up to hundreds of thousands of instances appears infeasible at this time. Memory footprint is another issue; several popular SVM packages store training vectors in memory. (Some exceptions are noted in Sect. 1.2.)

Scalability and memory footprint can become critical issues as enormous training sets become increasingly available. Web directories such as the Open Directory and Yahoo! contain millions of training instances that occupy tens of gigabytes, whereas even high-end servers are mostly limited to 1–2 GB of RAM. Given the sparsity of data in the text domain, sampling the training data is dangerous because the sample may exclude thousands of useful features. Joachims [14] shows that a large fraction of terms reveal at least some useful class information, and therefore every additional training document could potentially be a source of features. We confirm this in Sect. 4.6, where we see that the ability to scale better also translates to a better accuracy-time trade-off.

In summary, despite the theoretical elegance and superiority of SVMs, their IO behavior and CPU scaling are important concerns. There is a need for easy-to-implement text classifiers that match the simplicity and efficiency of NB classifiers while giving an accuracy comparable to SVMs.

### 1.1 Our contribution

We design, implement, and evaluate a new, simple text classification algorithm that requires very little RAM, deals gracefully with out-of-RAM training data (which it accesses strictly linearly), beats NB accuracy decisively, and even matches SVM accuracy. Our main idea is to:

1. Find a series of projections of the training data by using Fisher’s classical linear discriminant [17, Sect. 11.5] as a subroutine
2. Project all training instances to the low-dimensional subspace found in the previous step
3. Induce a decision tree on the projected low-dimensional data

We call this general framework SIMPL (Simple Iterative Multiple Projection on Lines). SIMPL has several important features: it has very small footprint, linear in the number of terms (dimensions)  $m$  plus the number of documents  $n$ ; it makes only fast sequential scans over the input; its CPU time is almost linear in the total size of the training data; it can be expressed simply in terms of joins, sorts, and GROUP BY operations; and it can be parallelized easily.

To give a quick impression, SIMPL has been implemented using only 600 lines of C++ code and trained on a 65524-document collection in 250 seconds, for which SVM took 3850 seconds.<sup>1</sup> We undertake a careful comparison between SIMPL and SVM with regard to accuracy and performance. We find that, in spite of its simplicity and efficiency, SIMPL is comparable (and sometimes superior) to SVM in terms of accuracy. Although there is no theoretical bound on the number of linear projections SIMPL may need, only two to three projections are usually enough to achieve high accuracy.

The ability to scale to training sets much larger than main memory is a key concern for the data mining community,

which has resulted in excellent out-of-core implementations for traditional classifiers such as decision trees [34]. In the last few years, the machine learning and text mining communities have evolved other powerful classifiers, such as SVMs and maxent classifiers. The scaling and IO behavior of the new and important class of SVM learners are not clearly understood. To this end, we carefully study the performance of a popular SVM implementation accessing documents from a LRU cache having limited size. If the SVM implementation is given a cache of size comparable to the RAM required by SIMPL, it spends a significant portion of its time servicing cache misses, and the performance gap between SIMPL and the SVM implementation grows further.

### 1.2 Related work

Although we are not aware of a hybrid learning strategy similar to our proposal, a few ideas that we discuss here were early hints that a projection-based approach could be promising. A 1988 theorem by Frankl and Maehara [9] showed that a projection of a set of  $n$  points in  $\mathbb{R}^m$  to a *random* subspace of dimension about  $(9/\epsilon^2) \log n$  preserves (to within a  $1 \pm \epsilon$  factor) all relative interpoint distances with high probability. Later work has established random projection as a valuable general technique for dealing with high-dimensional data. Kleinberg projected these points to  $\Theta(m \log^2 m)$  randomly directed lines to answer approximate nearest-neighbor queries efficiently [18]. Dasgupta [5,6] used random projection to learn a mixture of Gaussians, showing en route that well-separated Gaussians remain well separated upon projection.

For a classification task, we need not preserve *all* distances carefully. We simply need a subspace that separates the positive and negative instances well (a special case of “projection pursuits” [10,36]). In an early study by Schütze, Hull, and Pedersen [33], even *single* linear discriminants compared favorably with neural networks for the document routing problem. Lewis et al. [22] reported accurate prediction using a variety of regression strategies for good (single) linear predictors. The recent success of linear SVMs adds further evidence that very few projections could be adequate in the text domain.

In 1999, Shashua established that the decision surface found by a linear SVM is the *same* as the Fisher discriminant for only the “support vectors” (see Sect. 2.4) found by a SVM [35]. Although this result does not directly yield a better SVM algorithm, it gave us the basic intuition behind our idea. Our work is most closely related to linear discriminants [7] and SVMs, which we discuss in detail in Sect. 2 and Sect. 3. Independently, Cooke [4] has suggested discarding well-separated training points before finding Fisher’s linear discriminant but has not used multiple projections to generate a surrogate representation to be used by a more powerful learning algorithm such as a decision tree.

Many researchers have worked on reducing the memory footprint and running time of SVM optimizations. One strategy, pursued by Mangasarian and coworkers, is to change the objective function slightly, which enables use of more efficient mathematical programming machinery without affecting the utility of the solution in practice. Lagrangian SVM [24], proximal SVM [11], and incremental SVM [12] are examples of this paradigm; an incremental SVM can in fact retire and add new

<sup>1</sup> SIMPL is available at <http://www.cse.iitb.ac.in/~soumen>.

training data efficiently. These SVM variants involve inverting an  $m \times m$  matrix ( $m$  is the number of dimensions), which is readily done for moderately large values of  $m$  (hundreds to thousands) but demands too much main memory ( $O(m^2)$ ) and too much time ( $O(m^3)$ ) in the text domain, where  $m > 70000$  is not uncommon. The inverted matrix is generally not sparse. However, two other techniques from this family, successive over-relaxation SVM [23] and reduced SVM (a sampling technique) [20], may compare favorably with SIMPL (apart from being theoretically more elegant). Pavlov, Mao, and Dom have developed a different sampling technique [29].

SIMPL may be interpreted as an approximation to *boosting* [32]. Boosting seeks to improve a “weak” learner (which makes decisions only slightly better than random guessing) by running it many times on successively altered training distributions. The first training distribution is generated from the training data by assigning equal probability  $1/n$  to each instance. Subsequent distributions are generated by boosting the probability of instances that the weak learner labeled incorrectly. Each learner in the sequence also gets a score based on its error rate. The overall learner is a weighted majority of the set of weak learners, where the weights depend on the scores of the weak learners. SIMPL simply throws away correctly learned points but uses a more complex combination of the weak learners.

Our approach is also related to oblique decision trees (ODTs) [27], which try to find nonorthogonal hyperplane cuts in the decision-tree setting. Inducing an ordinary decision tree over the raw term space of a large document collection is already extremely time consuming. ODTs draw on an even more complex hypothesis space than decision trees (an arrangement of simplicial polytopes) and involve a regression over potentially all  $m$  dimensions at *each* node of the decision tree. Consequently, SIMPL is much faster than ODT induction. It is also somewhat faster to apply on test instances than ODTs because we only need to compute a small, fixed number of projections (usually two). We also found SIMPL to be more accurate than orthogonal decision trees. A comparison with ODTs may be worthwhile if ODTs can be trained within reasonable time on high-dimensional data.

## 2 Preliminaries

A host of linear classifiers have been used for text classification. A document  $d$  is represented as a “feature vector”  $\mathbf{d}$  with a component  $d_t$  for each term  $t$  in the vocabulary. We will overload  $d$  to mean  $\mathbf{d}$  where there is no confusion. Generally, the more often  $t$  occurs in  $d$ , the larger the value of  $d_t$ . A linear classifier is characterized by a vector  $\alpha$  and a scalar constant  $b$ , and it predicts the class of  $d$  as

$$c_{\text{guess}} = \text{sign}(\alpha \cdot \mathbf{d} + b), \quad (1)$$

where  $\cdot$  indicates a dot-product and  $\text{sign}(x) = 0$  if  $x = 0$ ,  $\text{sign}(x) = 1$  if  $x > 0$ , and  $\text{sign}(x) = -1$  if  $x < 0$ .

Linear classifiers work well for document classification along the lines of broad topics. A document about (the game of) cricket will tend to use terms like wicket, run, ball, and pitch frequently and will also tend to use them together, compounding the evidence that the document is about cricket.  $\alpha$  and  $b$

together represent a hyperplane, which cuts across each axis (corresponding to a term) at some offset. The offset acts as a threshold: if the term occurs more frequently, the document is assigned the label  $+1$ . In general, if a linear combination of the frequencies of important terms exceeds a threshold, the document will lie on the “positive” ( $c = 1$ ) side of the hyperplane. This discussion may help explain why linear discriminants are good at text classification.

### 2.1 Naive Bayes (NB) classifiers

Bayesian classifiers estimate a class-conditional document distribution  $\Pr(d|c)$  from the training documents and use Bayes’ rule to estimate  $\Pr(c|d)$  for test documents. The documents are modeled using their terms. The multinomial naive Bayes model assumes that a document is a bag or multiset of terms and the term counts are generated from a multinomial distribution after fixing the document length  $\ell_d$ , which, being fixed for a given document, lets us write

$$\Pr(d|c, \ell_d) = \binom{\ell_d}{\{n(d, t)\}} \prod_{t \in d} \theta_{c,t}^{n(d,t)} \quad (2)$$

where  $n(d, t)$  is the number of times  $t$  occurs in  $d$  and  $\theta_{c,t}$  are suitably estimated [1,26] multinomial probability parameters with  $\sum_t \theta_{c,t} = 1$  for all  $c$ . For the two-class scenario throughout this paper, we only need to compare  $\Pr(c = -1|d)$  against  $\Pr(c = 1|d)$ , or equivalently,  $\log \Pr(c = -1|d)$  against  $\log \Pr(c = 1|d)$ , which simplifies to a comparison between

$$\log \Pr(c = 1) + \sum_{t \in d} n(d, t) \log \theta_{1,t} \quad \text{and} \quad (3)$$

$$\log \Pr(c = -1) + \sum_{t \in d} n(d, t) \log \theta_{-1,t}$$

where  $\Pr(c = \dots)$ , called the class *priors*, are the fractions of training instances in the respective classes. Simplifying Eq. 3, we see that NB is a linear classifier: it makes a decision between  $c = 1$  and  $c = -1$  by thresholding the value of  $\alpha_{\text{NB}} \cdot \mathbf{d} + b$  for a suitable vector  $\alpha_{\text{NB}}$  (which depends on the parameters  $\theta_{c,t}$ ) and constant  $b$ . Here  $\mathbf{d}$  is overloaded to denote a vector of term frequencies (see Sect. 4) and “ $\cdot$ ” denotes a dot-product.

### 2.2 Maximum entropy classifiers

Whereas Bayesian classifiers estimate class-conditional distributions  $\Pr(d|c)$  for each class  $c$ , a maxent classifier directly estimates a parametric model for  $\Pr(c|d)$ . There is a model parameter  $\mu_{c,t}$  for every class  $c$  and term  $t$  (as in NB). A commonly used parametric form of  $\Pr(c|d)$  is

$$\Pr(c|d) \propto \prod_{t \in d} \mu_{c,t}^{n(d,t) / \sum_{\tau} n(d,\tau)} \quad (4)$$

Introducing a normalizing constant to make  $\Pr(c|d)$  add up to 1 over all  $c$ , we get

$$\Pr(c|d) = \frac{1}{Z(d)} \prod_{t \in d} \mu_{c,t}^{n(d,t) / \sum_{\tau} n(d,\tau)} \quad (5)$$

Nigam et al. [28] discuss in detail how to optimize the parameters using training data. In the two-class case, using a maxent

classifier to classify a test document  $d$  amounts to comparing (after taking logs)

$$\sum_{t \in d} \frac{n(d, t)}{\sum_{\tau} n(d, \tau)} \log \mu_{1, t} \quad \text{and} \quad \sum_{t \in d} \frac{n(d, t)}{\sum_{\tau} n(d, \tau)} \log \mu_{-1, t}$$

or  $\sum_{t \in d} n(d, t) \log \mu_{1, t} \quad \text{and} \quad \sum_{t \in d} n(d, t) \log \mu_{-1, t} \quad (6)$

again, clearly a linear discriminant.

### 2.3 Regression techniques

We can regard the classification problem as inducing a linear regression from  $d$  to  $c$  of the form  $c = \alpha \cdot d + b$ , where  $\alpha$  and  $b$  are estimated from the data  $\{(d_i, c_i), i = 1, \dots, n\}$ . This view has been common in a variety of information retrieval (IR) applications. A common objective is to minimize the square error between the observed and predicted class variable:  $\sum_d (\alpha \cdot d + b - c)^2$ . The least-square optimization frequently uses gradient-descent methods, such as the Widrow-Hoff (WH) update rule. Let each vector  $d$  be augmented by one extra element, always set to 1, and a corresponding extra dimension added to  $\alpha$ , to simplify notation and get rid of  $b$ . The WH approach starts with some initial estimate  $\alpha^{(0)}$  (with the extra dimension representing  $b$ ), considers  $d_i, c_i$  one by one, and updates  $\alpha^{(i-1)}$  to  $\alpha^{(i)}$  as follows:

$$\alpha^{(i)} = \alpha^{(i-1)} - 2\eta(\alpha^{(i-1)} \cdot d_i - c_i)d_i. \quad (7)$$

The final  $\alpha$  used for classification is usually the average of all  $\alpha$ s found along the way. Schütze, Lewis, et al. [22,33] have applied WH and other update methods (such as the exponentiated gradient method) to design high-accuracy linear classifiers for text, improving upon traditional Rocchio-style relevance feedback. We will follow the WH approach, but we will not minimize the square error because we are not dependent on a single linear predictor. Instead, our goal is to maximize separation between the classes in the projected subspace, for which we will optimize Fisher’s linear discriminant.

### 2.4 Linear support vector machines

Like NB, linear SVMs (LSVMs) also make a decision by thresholding  $\alpha_{\text{SVM}} \cdot d + b$  (the estimated class is +1 or -1 depending on whether the quantity is greater or less than 0) for a suitable vector  $\alpha_{\text{SVM}}$  and constant  $b$ .  $\alpha_{\text{SVM}}$  is chosen far more carefully than NB. Initially, let us assume that the  $n$  training points in  $\mathbb{R}^m$  from the two classes are linearly separable by a hyperplane perpendicular to a suitable  $\alpha$ . SVM seeks an  $\alpha$  that maximizes the distance of any training point from the hyperplane; this can be written as:

$$\text{Minimize } \frac{1}{2} \alpha \cdot \alpha \quad (= \frac{1}{2} \|\alpha\|^2) \quad (8)$$

subject to  $c_i(\alpha \cdot d_i + b) \geq 1 \quad \forall i = 1, \dots, n$

where  $\{d_1, \dots, d_n\}$  are the training document vectors and  $\{c_1, \dots, c_n\}$  their corresponding classes. (We want an  $\alpha$  such that  $\text{sign}(\alpha \cdot d_i + b) = c_i$ , so that their product is always

positive.) The distance of any training point from the optimized hyperplane (called the *margin*) will be at least  $1/\|\alpha\|$ .

To handle the general case where a single hyperplane may not be able to correctly separate *all* training points, *fudge* variables  $\{\xi_1, \dots, \xi_n\}$  are introduced, and Eq. 8 is enhanced as:

$$\text{Minimize } \frac{1}{2} \alpha \cdot \alpha + C \sum_i \xi_i \quad (9)$$

subject to  $c_i(\alpha \cdot d_i + b) \geq 1 - \xi_i \quad \forall i = 1, \dots, n$   
and  $\xi_i \geq 0 \quad \forall i = 1, \dots, n$

$\sum_i \xi_i$  is the “sum of violations” of the misclassified training points, which is traded off against the margin width  $\frac{1}{2} \alpha \cdot \alpha$  using the tuned constant  $C$ .

SVM packages solve the *dual* of Eq. 9, involving scalars  $\lambda_1, \dots, \lambda_n$ , given by:

$$\text{Maximize } \sum_i \lambda_i - \frac{1}{2} \sum_{i,j} \lambda_i \lambda_j c_i c_j (d_i \cdot d_j) \quad (10)$$

subject to  $\sum_i c_i \lambda_i = 0$   
and  $0 \leq \lambda_i \leq C \quad \forall i = 1, \dots, n$

Having optimized the  $\lambda$ s,  $\alpha$  is recovered as

$$\alpha_{\text{SVM}} = \sum_i \lambda_i c_i d_i \quad (11)$$

If  $0 < \lambda_i \leq C$ , then  $d_i$  is a “support vector.”  $b$  can be estimated as  $c_j - \alpha_{\text{SVM}} \cdot d_j$ , where  $d_j$  is some document for which  $0 < \lambda_j < C$ .

It is common to set  $C$  to  $1/r^2$ , where  $r$  is the radius of the smallest ball containing all training vectors, or the average Euclidean norm of the  $d$ -vectors. A fixed choice of  $C$  saves time but is rarely the best possible value in practice. In principle, the best value of  $b$  should fall out of the training process, but in practice, tuning  $b$  helps as well. Practitioners generally use a held-out validation data set to search over a large range of values for  $C$  and a smaller range of values for  $b$ , which means they have to run the (time-consuming) SVM induction program many, many times. It is natural to want to avoid this extra work, so we compare SIMPL with *two* versions of SVM: one in which  $C = 1/r^2$  and  $b$  is not tuned, and another in which we search over  $C$  and  $b$ , which we call **SVM-best**.

Equation 10 represents a quadratic optimization problem. SVM packages iteratively refine a few  $\lambda$ s at a time (called the *working set*), holding the others fixed. For all but very small training sets, we cannot precompute and store all the inner products  $d_i \cdot d_j$ . As a scaled-down example, if an average document costs 400 bytes in RAM, and there are only  $n = 1000$  documents, the corpus size is 400000 bytes, and the inner products, stored as 4-byte floats, occupy  $4 \times 1000 \times 1000$  bytes, ten times the corpus size. Therefore, the inner products are computed on demand, with an LRU cache of recent values, to reduce recomputation. In all SVM implementations that we know of, all the document vectors are kept in memory so that the inner products can be quickly recomputed when necessary. (This observation excludes SVM variants that do not need quadratic programming, such as the ones discussed in Sect. 1.2.)

### 2.5 Observations leading to our approach

It is natural to question the observed difference between the accuracy of NB classifiers and linear SVMs, given that they

use the same hypothesis space (half-spaces). In assuming attribute independence, NB starts with a large *inductive bias* (loosely speaking, a constraint not guided by the training data) on the space of separating hyperplanes that it will draw from. SVMs do not propose any generative probability distribution for the data points and do not suffer from this form of bias. Another weakness of the NB classifier is that its parameters are based only on sample means; it takes no cognizance of variance. Fisher’s discriminant does take variance into account (see Fig. 2 below).

A linear SVM carefully finds a single discriminative hyperplane. Consequently, instances projected on the direction  $\alpha_{\text{SVM}}$  normal to this hyperplane show large to perfect inter-class separation. Intuitively, our hope is that we can be slightly sloppy (compared to SVMs) with finding discriminative directions, provided we can quickly find a number of such directions that can *collectively* help a decision tree learner separate the classes in the projected space. To achieve speed and scalability, it must be possible to cast our computation in terms of efficient sequential scans over the data with a small number of accumulators collecting sufficient statistics [13].

### 3 The proposed algorithm

Our proposed training algorithm has the broad outline shown in Figure 1. The outer while-loop of SIMPL finds several linear discriminants  $\alpha^{(0)}, \alpha^{(1)}, \dots$  one by one. To compute each  $\alpha$ , we need a **hill-climbing** step. To prepare the stage for the next  $\alpha$ , we need to **remove** some instances from the training set. We will explain and rationalize these steps in this section.

Testing is simple, and essentially as fast as NB or SVM. We preserve the linear discriminants  $\alpha^{(0)}, \dots, \alpha^{(k-1)}$  as well as the decision tree (in practice we found  $k = 2 \dots 3$  to be adequate). Given a test document  $d$ , we find its  $k$ -dimensional representation  $(d \cdot \alpha^{(0)}, \dots, d \cdot \alpha^{(k-1)})$  and submit this vector to the decision tree classifier, which outputs a class.

#### 3.1 The hill-climbing step

Let the points with  $c = -1$  ( $c = 1$ ) be  $X$  ( $Y$ ). Fisher’s linear discriminant is a (unit) vector  $\alpha$  such that the positive and negative training instances, projected on the direction  $\alpha$ , are as “well-separated” as possible. The separation  $J(\alpha)$ , shown in equation (13), is quantified as the ratio of the square of the difference between the projected means to the sum of the projected variances.

In matrix notation, if  $\mu_X$  and  $\mu_Y$  are the means (centroids) and  $\Sigma_X = (1/|X|) \sum_X (x - \mu_X)(x - \mu_X)^T$  and  $\Sigma_Y = (1/|Y|) \sum_Y (y - \mu_Y)(y - \mu_Y)^T$  are the covariance matrices for point sets  $X$  and  $Y$ , the best linear discriminant can be found in closed form, and has been used in pattern recognition:

$$\arg \max_{\alpha} J(\alpha) = \left( \frac{\Sigma_X + \Sigma_Y}{2} \right)^{-1} (\mu_X - \mu_Y) \quad (12)$$

provided the matrix inverse exists.

However, inverting the covariance matrix is impractical in the document classification domain because it is too large and very likely ill-conditioned. Moreover, inversion will discard sparsity. Instead, we use a gradient ascent or

hill-climbing approach: we start from a reasonable starting value of  $\alpha = (\alpha_1, \dots, \alpha_m)$  and repeatedly find  $\nabla J(\alpha) = (\partial J / \partial \alpha_1, \dots, \partial J / \partial \alpha_m)$ . Denoting the numerator (respectively, denominator) in the rhs of Eq. 13 as  $N(\alpha)$  (respectively,  $D(\alpha) = D_X(\alpha) + D_Y(\alpha)$ ), we can easily write down  $\partial N / \partial \alpha_k$ ,  $\partial D_X / \partial \alpha_k$ , and  $\partial D_Y / \partial \alpha_k$ , as shown in Fig. 2.

From these values we can easily derive the value of  $\partial J / \partial \alpha_i$  for all  $i$  in the term vocabulary. Once we find  $\nabla J(\alpha)$ , we use the standard WH update rule with a learning rate  $\eta$ :

$$\alpha_{\text{next}} \leftarrow \alpha_{\text{current}} + \eta \nabla J(\alpha) \quad (17)$$

As with  $C$  in SVM, setting the learning rate  $\eta$  is a practiced science and art. A small  $\eta$  slows down learning, and a large  $\eta$  may lead to instability and divergence. In neural networks, researchers adapt the rate online as the training progresses [19]. Based on the experience of many WH users published on the Web, we tried values between 0.05 and 0.2 and settled for 0.1. This single, fixed value gave us results that were essentially as good as we could get by tuning  $\eta$  separately for each data set.

We must also fix our “convergence policy.” In SIMPL, we repeat hill-climbing until the increase in  $J(\alpha)$  is less than 5% over three successive iterations. This policy protects us from mild oscillations near the local optimum.

The gist is that we need to maintain the following set of accumulators as we scan the documents sequentially:

- $\sum_X x \cdot \alpha$  (scalar)
- $\sum_Y y \cdot \alpha$  (scalar)
- $\sum_X x_i$  for each  $i$  ( $m$  numbers)
- $\sum_Y y_i$  for each  $i$  ( $m$  numbers)
- $\sum_X x_i (x \cdot \alpha)$  for each  $i$  ( $m$  numbers)
- $\sum_Y y_i (y \cdot \alpha)$  for each  $i$  ( $m$  numbers)

together with the current  $\alpha$ , which is another  $m$  numbers. The total memory footprint is only  $5m + O(1)$  numbers ( $20m$  bytes), where  $m$  is the size of the vocabulary. For our Dmoz data set (see Sect. 4),  $m \approx 1.2 \times 10^6$ , which means we need only about 24MB of RAM. All the vectors have dense array representations, so the time for one hill-climbing step is exactly **linear** in the size of the input data. In the unlikely situation that available RAM is smaller than  $20m$  bytes, the expressions in Fig. 2 can be expressed as GROUP BY and aggregate operations, which can be executed efficiently with limited memory.

#### 3.2 Pruning the training set

After a suitable number of hill-climbing steps, we discard points in  $D$  that are correctly classified by the current  $\alpha$ . This is achieved by projecting all points in  $D$  along  $\alpha$ , so that they are now points on the line (each marked with  $c = 1$  and  $c = -1$ ) and sweeping the line for a minimum-error position where

- most points on one side have  $c = 1$  and most points on the other side have  $c = -1$ , and
- the number of points on the “wrong” side is the minimum possible.

It is easy to do this in one sort of an  $n$ -element array and one sweep with  $O(1)$  extra memory, so the total time for identifying correctly classified documents is  $O(n \log n)$  (and the total space needed is  $O(m + n)$ ).

```

i ← 0
initialize D to the set of all training documents
while D has at least one positive and one negative instance do
  initialize  $\alpha^{(i)}$  to the vector direction joining the positive class centroid to the negative class centroid
  do hill-climbing to find a good linear discriminant  $\alpha^{(i)}$  for D
  remove from D those instances that are correctly classified by  $\alpha^{(i)}$ 
  orthogonalize  $\alpha^{(i)}$  w.r.t.  $\alpha^{(0)}, \dots, \alpha^{(i-1)}$  and scale it so that its  $L_2$  norm  $\|\alpha^{(i)}\| = 1$ 
  i ← i + 1
end while
let  $\alpha^{(0)}, \dots, \alpha^{(k-1)}$  be the k linear discriminant vectors found
for each document vector d in the original training set do
  represent d as a vector of its projections  $(d \cdot \alpha^{(0)}, \dots, d \cdot \alpha^{(k-1)})$ 
  train a decision tree classifier with the k-dimensional training vector
end for

```

**Fig. 1.** The proposed algorithm SIMPL finds several linear discriminants  $\alpha^{(0)}, \dots, \alpha^{(k-1)}$  by using a hill-climbing procedure

$$J(\alpha) = \frac{\overbrace{\left( \frac{1}{|X|} \sum_X x \cdot \alpha - \frac{1}{|Y|} \sum_Y y \cdot \alpha \right)^2}^N}{\underbrace{\frac{1}{|X|} \sum_X (x \cdot \alpha)^2 - \left( \frac{1}{|X|} \sum_X x \cdot \alpha \right)^2}_{D_X} + \underbrace{\frac{1}{|Y|} \sum_Y (y \cdot \alpha)^2 - \left( \frac{1}{|Y|} \sum_Y y \cdot \alpha \right)^2}_{D_Y}} \quad (13)$$

$$\frac{\partial J}{\partial \alpha_i} = \frac{(D_X + D_Y) \frac{\partial N}{\partial \alpha_i} - N \frac{\partial (D_X + D_Y)}{\partial \alpha_i}}{(D_X + D_Y)^2} \quad (14)$$

$$\frac{\partial N}{\partial \alpha_i} = 2 \left( \frac{1}{|X|} \sum_X x \cdot \alpha - \frac{1}{|Y|} \sum_Y y \cdot \alpha \right) \left( \frac{1}{|X|} \sum_X x_i - \frac{1}{|Y|} \sum_Y y_i \right) \quad (15)$$

$$\frac{\partial D_X}{\partial \alpha_i} = \frac{2}{|X|} \left( \sum_X x_i (x \cdot \alpha) - \frac{1}{|X|} (\sum_X x_i) (\sum_X x \cdot \alpha) \right) \quad (16)$$

**Fig. 2.** The main equations involved in the hill-climbing step to find the Fisher’s linear discriminant

The intuition behind this algorithm is quite simple: having found a discriminant  $\alpha$  we should retain only those points that fail to be separated in that direction. Although the hill-climbing steps will always converge to local optima, there is no bound on the number of  $\alpha$ s we will need to extract. In practice, each new  $\alpha$  helps us discard over 80% of *D*. To avoid scanning through the original *D* for every hill-climbing pass, we write out the surviving documents in a new file, which then becomes our *D* for finding the next  $\alpha$ .

Orthogonalizing the set of  $\alpha$ s reduces the correlation between the components of the *k*-dimensional representation of documents. Because decision trees implement orthogonal cuts, we sometimes found a mild improvement in accuracy from the orthogonalization step; it never hurt accuracy. If we replace the decision tree by some other kind of learner, this step may not be needed.

### 3.3 Inducing a decision tree

We used two roughly equivalent decision-tree packages using Quinlan’s C4.5 algorithm: C4.5 itself (<http://www.cse.unsw.edu.au/~quinlan/>) and the Decision-tree package in WEKA [38] (which we simply call WEKA; also see <http://www.cs.waikato.ac.nz/~ml/weka/>). In our context, a decision tree seeks

to partition a set of labeled points  $\{d\}$  in a geometric space, where each *d* is a vector  $(d_0, \dots, d_{k-1})$ .

The decision-tree induction algorithm uses a series of guillotine cuts on the space, each of which is expressed as a comparison of some component  $d_i$  against a constant such that each final rectangular region has only positive or only negative points. The hierarchy of comparisons induces the decision tree, whose leaves correspond to final rectangular regions. To achieve a recall-precision trade-off, just as we can tune the offset *b* for a SVM, we can assign different weights to positive ( $c = 1$ ) and negative ( $c = -1$ ) instances in WEKA.

A decision tree with “pure” (single-class) leaves usually *overfits* the training data and does not generalize well to held-out test data. Better generalization is achieved by *pruning* the tree, trading off the complexity of the tree with the impurity of the leaf rectangles in terms of mixing points belonging to different classes. This does not work very well for large *m* (number of dimensions), which is why decision trees induced on raw text show poor accuracy. This is also why our dimensionality reduction via projection pays off well.

C4.5 and WEKA hold the entire training data in memory, which is usually not a problem because by this stage we have transformed the training data to points with only two to three dimensions. For example, with 100000 documents and three projections, we will need only 1.2MB. Should space become

an issue, we can always use efficient, out-of-core decision-tree implementations like SPRINT [34].

SIMPL induces fairly small decision trees. The number of projected coordinates is already small (two to three) to start with. It is rare to see more than two to three cuts on any one projected dimension. In our experiments, we observe that typically, the depth of the decision tree is less than five, and there are at most a total of 10–15 decision nodes in the tree.

## 4 Experiments

The core of SIMPL (excluding document scanning and pre-processing) was implemented in only 600 lines of C++ code, making generous use of ANSI templates. The core of C4.5 is roughly another 1000 lines of C code. (In contrast, SVM-LIGHT, a very popular SVM package, is over 6000 lines.) “g++ -O3” was used for compilation. Programs were run on Pentium III machines with 500–700MHz CPUs and 512–2048MB of RAM.

Several implementations of SVM are widely used and publicly available: Sequential Minimum Optimization (SMO) by John Platt [8,30], NodeLib by Gary Flake (<http://www.neci.nec.com/homepages/flake/nodelib/html/>), and SVM-LIGHT by Thorsten Joachims [15]. We found SVM-LIGHT to be comparable or better in accuracy compared to published SMO numbers. NodeLib has no built-in support for massive, sparse input data at this time. Our experiments are based on SVM-LIGHT, evaluating it for  $C$  between 1 and 60 (SVM-LIGHT’s default is the reciprocal of the average value of  $d \cdot d$ , which range between 20 and 30 for our TFIDF representation). We also evaluate several values of  $b$  in a suitable band around the separator. Other settings and flags are left at default values except where noted.

We use a few standard accuracy measures. For the following contingency table

(Number of documents)	Estimated class	
	$\bar{c}$	$c$
Actual class $\bar{c}$	$n_{00}$	$n_{01}$
$c$	$n_{10}$	$n_{11}$

recall and precision are defined as  $R = n_{11}/(n_{11} + n_{10})$  and  $P = n_{11}/(n_{11} + n_{01})$ , respectively.  $F_1 = 2RP/(R + P)$  is also a commonly used measure. A classifier may have parameters using which one can trade off  $R$  for  $P$  or vice versa. When these parameters are adjusted to get  $R = P$ , this value is called the “break-even point.”

### 4.1 Data sets

We use one synthetic data set and several standard real-life benchmark data sets. We use the synthetic data to study properties of the discriminants found by SVM and SIMPL, and we report precision and recall numbers on the real-life data sets.

#### 4.1.1 Synthetic data

Our synthetic data generator is based on the “TCAT concept” defined by Joachims [16]. A TCAT concept is specified via a set

$p_i$	$n_i$	$f_i$	Feature type
77	29	98	High-frequency positive (HFP)
4	21	52	High-frequency negative (HFN)
16	2	431	Medium-frequency positive (MFP)
1	12	341	Medium-frequency negative (MFN)
9	1	5045	Low-frequency positive (LFP)
1	21	24276	Low-frequency negative (LFN)
169	191	8116	Rest

$p_i$	$n_i$	$f_i$	Feature type
33	2	65	High frequency positive (HFP)
32	65	152	High frequency negative (HFN)
2	1	171	Medium frequency positive (MFP)
3	21	974	Medium frequency negative (MFN)
3	1	3455	Low frequency positive (LFP)
1	10	17020	Low frequency negative (LFN)
78	52	5821	Rest

Fig. 3. Parameters for the TCAT-based synthetic data generator

of 3-tuples  $\{(p_i, n_i, f_i), i = 1, 2, \dots\}$ . The  $i$ th tuple specifies a vocabulary subset having  $f_i$  terms. The overall vocabulary is the union of all these subsets. Positive documents ( $c = 1$ ) use  $p_i$  out of the  $f_i$  terms from subset  $i$ , whereas negative documents ( $c = -1$ ) use  $n_i$  of the  $f_i$  terms from subset  $i$ . Note that the same term may be picked for both positive and negative documents.

Joachims argued through examples that TCAT concepts closely model real-life text classification tasks. He estimated TCAT parameters for some well-known classification benchmarks (Sect. 4.1.2), two of which are shown in Fig. 3 and are used in our experiments.

We use a simple TCAT-based data generator, which chooses the  $p_i$  (respectively,  $n_i$ ) terms uniformly at random, with replacement, from the  $f_i$  available features in set  $i$ . This implies all the synthetic documents have the same length (number of words), which is not very realistic. Note that our TCAT-based synthetic data are used purely to compare the linear discriminants computed by SVM and SIMPL under controlled circumstances, not to judge overall accuracy. Joachims used the TCAT characterization in a more powerful manner: he proved that TCAT concepts are learnable by LSVMs with small testing error, without making any assumptions about how the words are picked and how often they are repeated. He also derived results involving noise term distributions added to the basic TCAT model.

#### 4.1.2 Real-life data

We use the following real-life data sets. The first three are well known in recent IR literature, small in size, and suitable for controlled experiments on accuracy and CPU scaling. The last two data sets are large; they approach the scale we envisage for real applications. They were mainly used to compare run-time performance.

*Reuters [21]*: About 7700 training and 3000 test documents (“MOD-APTE” split), 30000 terms, 135 categories. The raw text takes about 21 MB.

*20NG*: About 18800 total documents organized in a directory structure with 20 topics. For each topic the files are listed alphabetically and the first 75% chosen as training documents. There are 94000 terms. The raw concatenated text takes up 25 MB and can be downloaded from <http://kdd.ics.uci.edu/databases/20newsgroups/20newsgroups.html>

*WebKB*: About 8300 documents in 7 categories. About 4300 pages on 7 categories (faculty, project, etc.) were collected from 4 universities, and about 4000 miscellaneous pages were collected from other universities. For each classification task, any one of the four university pages are selected as test documents and rest as training documents. The raw text is about 26 MB and can be downloaded from <http://www.cs.cmu.edu/afs/cs.cmu.edu/project/theo-20/www/data/>.

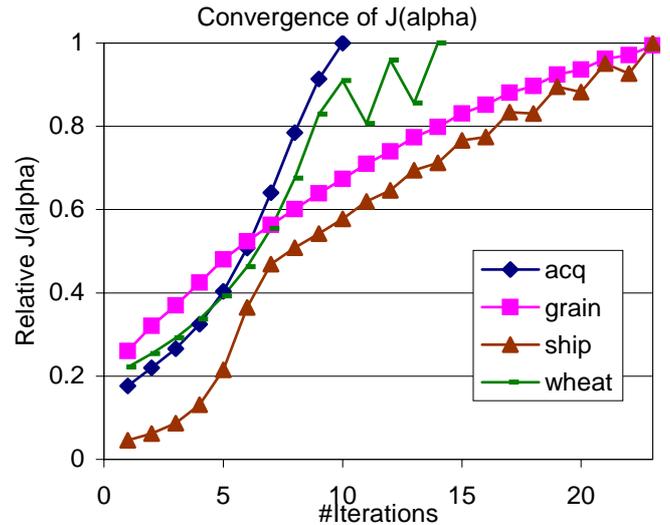
*OHSUMED*: 348566 abstracts from medical journals, having around 230000 terms and 308511 topics and available at <http://ftp.ics.uci.edu/pub/machine-learning-databases/ohsumed/>. The raw text is 400 MB. The first 75% are selected as training documents and the rest are test documents.

*Dmoz*: From the RDF file published at <http://dmoz.org/>, we picked a sample of 140000 URLs, successfully crawled some 120000 of them, and stripped HTML tags, leaving plain text behind. The number of distinct tokens was over 500000. Twelve top-level topics of Dmoz were used; these were populated with 5000–21000 labeled documents each. The raw text (available on request) occupied 271 MB.

#### 4.2 Document representation

We use the standard multinomial model (Eq. 2) for NB and the standard “TFIDF” document representation from IR for SVM and SIMPL. In keeping with some of the best systems at TREC (<http://trec.nist.gov/>), our IDF for term  $t$  is  $\log(|D|/|D_t|)$ , where  $D$  is the document collection and  $D_t \subseteq D$  is the set of documents containing  $t$ . The term frequency  $\text{TF}(d, t) = 1 + \ln(1 + \ln(n(d, t)))$ , where  $n(d, t) > 0$  is the raw frequency of  $t$  in document  $d$  (TF is zero if  $n(d, t) = 0$ ).  $d$  is represented as a sparse vector with the  $t$ th component being  $\text{IDF}(t) \text{TF}(d, t)$ . The  $L_2$  norm of each document vector is scaled to 1 before submitting to the classifiers.

All our data sets sport more than two labels. For each label (e.g., “cocoa”), a two-class problem (“cocoa” vs. “not cocoa”) is formulated. All tokens are turned to lowercase and standard SMART stopwords



**Fig. 4.** Hill-climbing to maximize  $J(\alpha)$  is fast. Relative values of  $J(\alpha)$  (scaled to 1 at convergence) are plotted against the number of hill-climbing steps

(<ftp://ftp.cs.cornell.edu/pub/smart/>) are removed, but no stemming is performed. No feature selection is used prior to running any of our classification algorithms: the naive Bayes classifier in the Rainbow library [25] (with Laplace and Lidstone’s methods evaluated for parameter smoothing), SVMLIGHT, and SIMPL. Alternatively, one may preprocess the collection through a common feature selector and then submit them to each classifier, which adds a fixed time to each classifier.

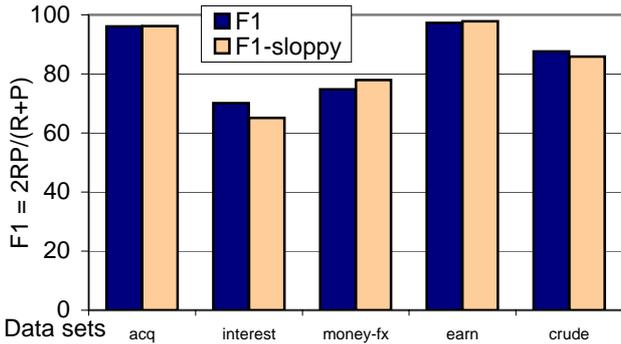
#### 4.3 Hill climbing

The hill-climbing approach is fast and practical. We usually settle at a maximum within 15–25 iterations: Fig. 4 shows that  $J(\alpha)$  quickly grows and stabilizes with successive iterations. Our convergence policy (see Sect. 3.1) guards well against mild problems of local maxima and overshoots in case the learning rate  $\eta$  in Eq. 17, which is set to 0.1 throughout, is not the best possible choice. This condition usually manifests itself in small oscillations in  $J(\alpha)$ , e.g., for topics *ship* and *wheat*. Our results are insensitive, within a wide range, to the specific choices of all these parameters.

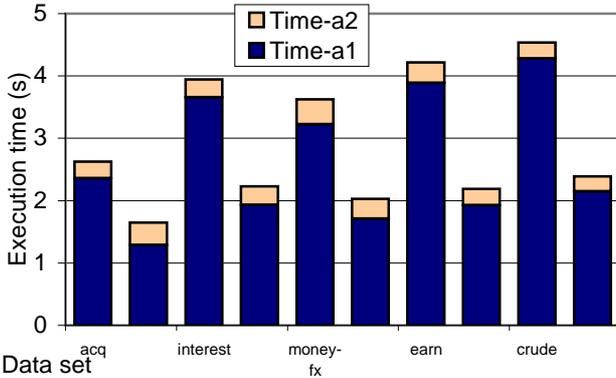
An important concern for us was the lack of a guarantee of global optimality in the hill-climbing step. Suppose the hill-climbing process for  $\alpha^{(0)}$  gets trapped in a local maximum, and we accept this value, continuing with  $\alpha^{(1)}$  etc. Might this completely upset the optimization of subsequent  $\alpha$ s and lead us arbitrarily astray from the best set of  $\alpha$ s?

It is difficult to find a perfect answer to this question because an exhaustive search for all the optimal  $\alpha$ s is infeasible, i.e., we have no access to “ground truth.” Instead, to get some heuristic evidence of the robustness of SIMPL, we calculated two different values of  $\alpha^{(0)}$ :

- Perform hill climbing for  $\alpha^{(0)}$  until convergence (this may give a global or a local optimum). Suppose this takes  $k$  iterations.
- Degrade the value of  $\alpha^{(0)}$  deliberately by running the hill-climbing process for only  $k/2$  iterations.



**Fig. 5.** Quitting hill climbing early has little or no impact on  $F_1$  accuracy; in fact, in some cases, accuracy *improves*



**Fig. 6.** For each topic, the first bar shows time to compute  $\alpha^{(0)}$  to convergence, then the time to compute  $\alpha^{(1)}$  to convergence. The second bar shows the time to compute  $\alpha^{(0)}$  sloppily, followed by computing  $\alpha^{(1)}$  to convergence. Because earlier  $\alpha$ s take more time to compute, we may even save overall time

After finding  $\alpha^{(0)}$ , subsequent  $\alpha$ s are computed as usual. We wish to study the implications of these two choices on both the accuracy and the speed of SIMPL.

For several data sets, early termination of hill climbing for  $\alpha^{(0)}$  had negligible effect on the overall accuracy of SIMPL (see Fig. 5). We determined that later  $\alpha$ s can compensate for sloppiness in finding an earlier  $\alpha$ . In fact, in some cases, because one of precision and recall falls less than the other, we may even *gain* accuracy in  $F_1$  terms. Such resilience to variation in policy and parameters is very desirable.

Figure 6 shows the effect of incomplete hill climbing on running time. For a given  $\alpha$ , one iteration of hill climbing takes a fixed time. Therefore, halving the number of iterations also halves the time taken to find  $\alpha^{(0)}$ . Finding  $\alpha^{(0)}$  sloppily has two potential related effects: fewer training points may be eliminated and a larger number of subsequent  $\alpha$ s may be required. We focus on the run time for finding  $\alpha^{(1)}$  in Fig. 6. Although the time to find  $\alpha^{(1)}$  generally increases if  $\alpha^{(0)}$  is found sloppily, the extra time required for  $\alpha^{(1)}$  (and subsequent  $\alpha$ s) is very small compared to the time saved while finding  $\alpha^{(0)}$ .

We save overall time because we can lop off over 80% of the current training set for every additional  $\alpha^{(i)}$  that we find. For example, for the Dmoz data the numbers of surviving documents were 116528, 7074, 230, and 2 before finding  $\alpha^{(0)}$ ,  $\alpha^{(1)}$ ,  $\alpha^{(2)}$ , and  $\alpha^{(3)}$  respectively. Consequently, for all

our data sets, SIMPL generates only two to four linear projections before running out of training documents.

The preceding discussion raises the following question: how will SIMPL behave if we push sloppiness to the limit, i.e., we perform no hill climbing *at all* after initializing  $\alpha$ ? We might need to find more  $\alpha$ s, but each will be computed trivially. Our main experimental observations, from a suitably modified version of SIMPL, are:

- The number of surviving training points after each  $\alpha$  tends to be larger than in the hill-climbing case:

Hill climbing	No hill climbing
7776 $\rightarrow$ 203 $\rightarrow$ 0	7776 $\rightarrow$ 250 $\rightarrow$ 5
5836 $\rightarrow$ 159 $\rightarrow$ 0	5836 $\rightarrow$ 520 $\rightarrow$ 2
5836 $\rightarrow$ 149 $\rightarrow$ 0	5836 $\rightarrow$ 406 $\rightarrow$ 1

- However, the number of extra  $\alpha$ s required is almost always one or two compared to the previous number of  $\alpha$ s obtained through hill-climbing.
- The time saved on  $\alpha^{(0)}$  by avoiding any hill climbing overwhelms any cost for finding extra  $\alpha$ s; therefore, the total time is smaller if no hill climbing is performed. For example, in some runs with real-life data, we cut down total time from 21 s to 6 s, from 35 s to 6 s, and 125 s to 28 s. Both sets of times are substantially smaller than the time taken by SVM.
- However, there is generally a drop in accuracy. For example, in some runs,  $F_1$  was reduced from 0.96 to 0.945, and from 0.73 to 0.6. In a few cases, hill climbing made the difference between beating vs. not beating SVM.

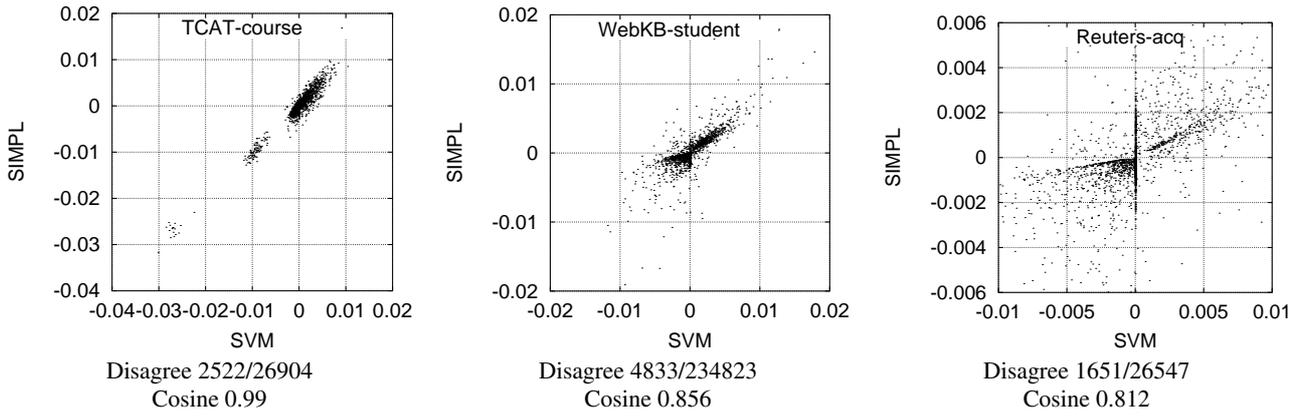
#### 4.4 Discriminants and projections

The next set of measurements shows the nature of discriminants found by SIMPL and other algorithms and the quality of separation of the training (and test) data on these discriminants.

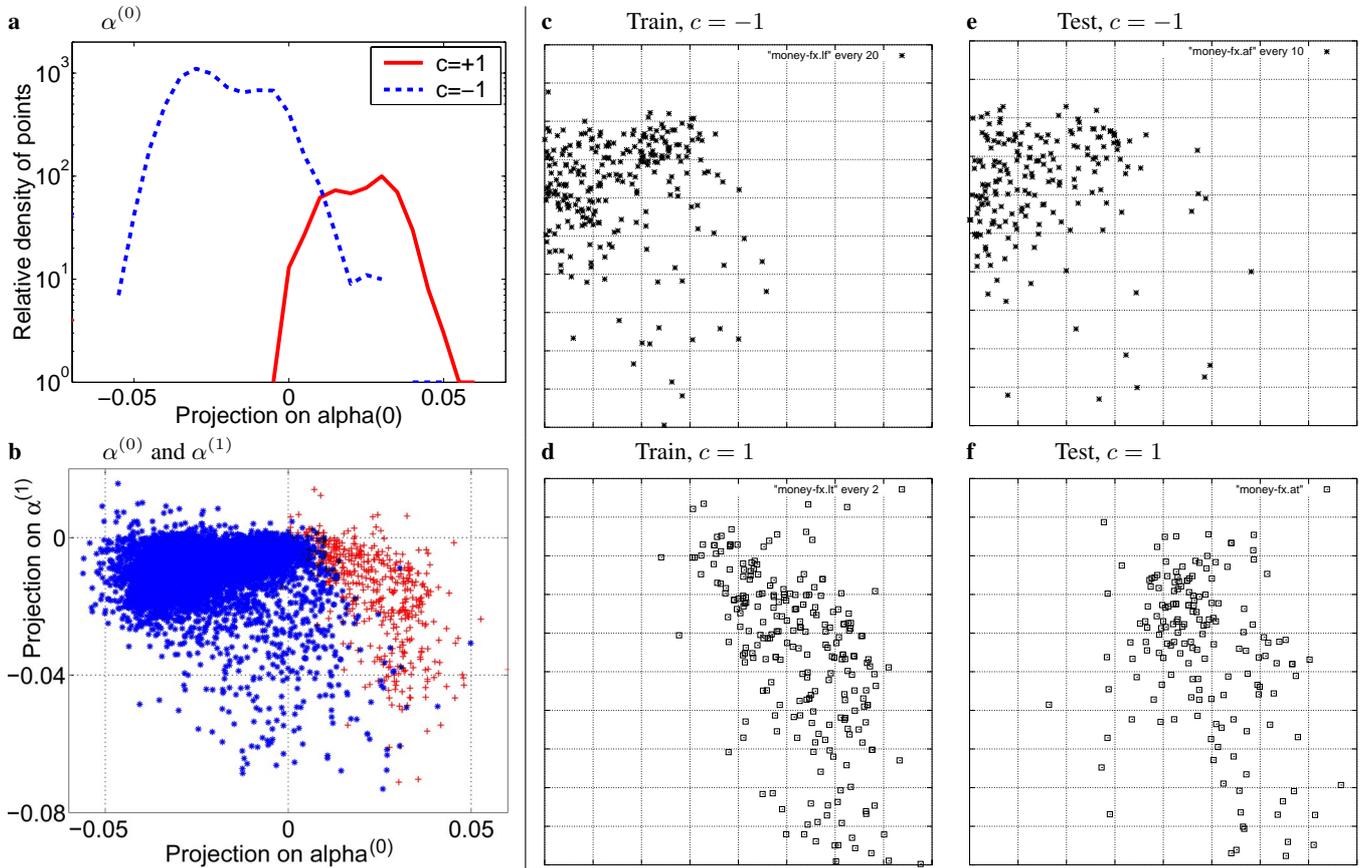
One indicator of the promise of SIMPL is the general agreement between the directions  $\alpha_{\text{SIMPL}}^{(0)}$  and  $\alpha_{\text{SVM}}$ . Assuming both vectors are scaled to unit  $L_2$  norm, the agreement can be expressed as their dot-product, i.e., the cosine of the angle between them, a single number between  $-1$  (perfect anticorrelation) and  $+1$  (perfect correlation). We can also visualize a scatterplot of points  $(\alpha_{\text{SVM}}^{(0)}(t), \alpha_{\text{SIMPL}}^{(0)}(t))$  for each term  $t$ .

Note that a low correlation between  $\alpha_{\text{SIMPL}}^{(0)}$  and  $\alpha_{\text{SVM}}$  does not necessarily mean that SIMPL is in trouble because SIMPL derives more directions for projection.

In Fig. 7, we show the scatterplots between the discriminants found by SVM and SIMPL. For synthetic data, the discriminants found by SVM and SIMPL are virtually identical, with typical cosines around 0.99. Real data sets contain noise, and this is clear from scatterplots that appear to be more spread out. However, the apparent density of points in the print is misleading: the  $(+, +)$  and  $(-, -)$  quadrants have an overwhelming majority of points (e.g., 234823) compared to the  $(+, -)$  and  $(-, +)$  quadrants (e.g., 4833). The cosine remains quite large, typically 0.78 to over 0.85. These may appear small compared to 0.99, but keep in mind that the  $\alpha$  vectors have tens of thousands of dimensions, where slight noise in each dimension is capable of turning the vectors away from



**Fig. 7.** Comparison between  $\alpha_{\text{SVM}}$  and  $\alpha^{(0)}$  found by SIMPL. “Disagree  $x/y$ ” means that  $x$  coordinates differed in sign and  $y$  coordinates did not, i.e., there are in all  $x$  points in quadrants 2 and 4 and  $y$  points in quadrants 1 and 3. Cosine is as described in the text



**Fig. 8a-f.** Although  $\alpha^{(0)}$  by itself may fail to separate the labels in the training data well (a), assistance from  $\alpha^{(1)}$  (and subsequent  $\alpha$ s in general) manages to separate the labels sufficiently well (b). For clarity, we separate training and test points for each  $c$  in parts c-f. We see that the projections cluster training and test data in very similar ways, hinting that high-accuracy decision trees can be induced from the reduced-dimension data

each other. Cosine values above 0.78 therefore indicate substantial agreement between the discriminants found by SVM and SIMPL.

Agreement between these discriminants gives us hope that if a linear SVM manages to separate the (projected) training data by a wide margin, Fisher’s discriminant should achieve separation that is almost as good. This, however, is not often the case. As Fig. 8a shows, the “confusion zone” after project-

ing the training points to  $\alpha^{(0)}$  is still significant. If only one Fisher’s discriminant were used for classification, we would expect SIMPL’s accuracy to be poorer than LSVM’s.

Luckily, SIMPL can make use of multiple projections. Figure 8b shows the result of projecting the training points on to the first *two* discriminants,  $\alpha^{(0)}$  and  $\alpha^{(1)}$ . The confusion zone in Fig. 8a is now stretched out over the  $y$ -axis, better separating training points with different labels. One or two more projec-

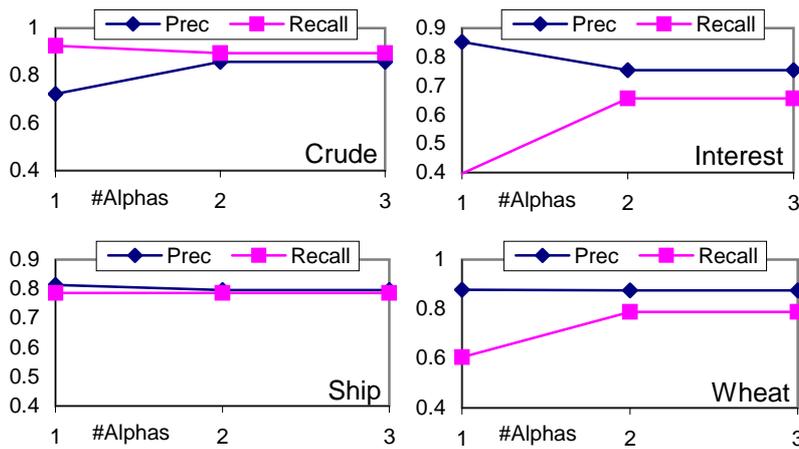


Fig. 9. Variation of precision and recall with the number of linear projections used by C4.5. Two projections are almost always adequate

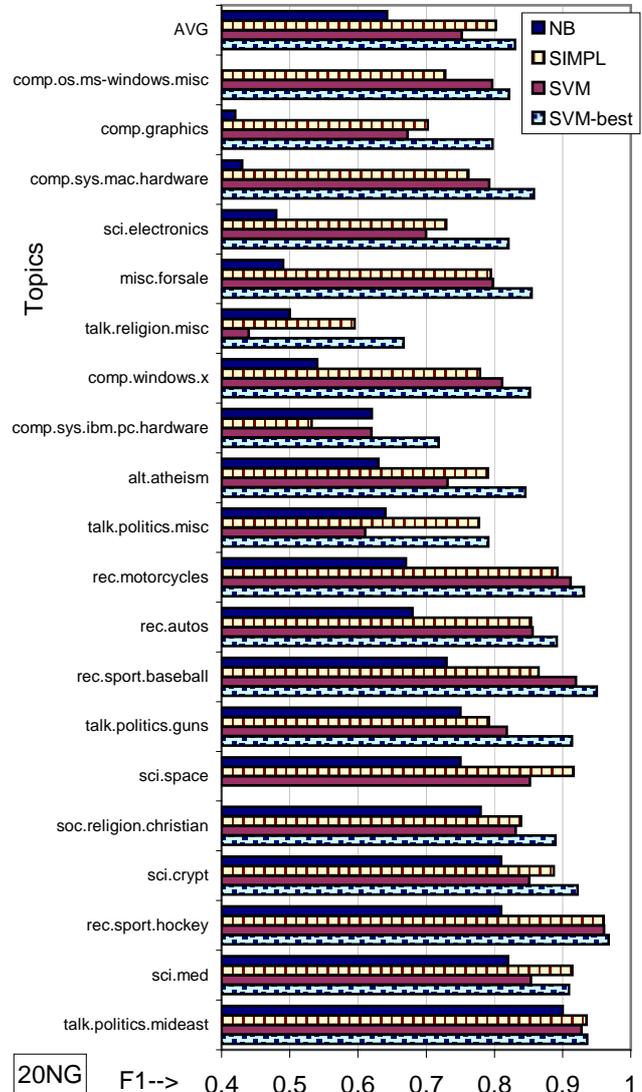
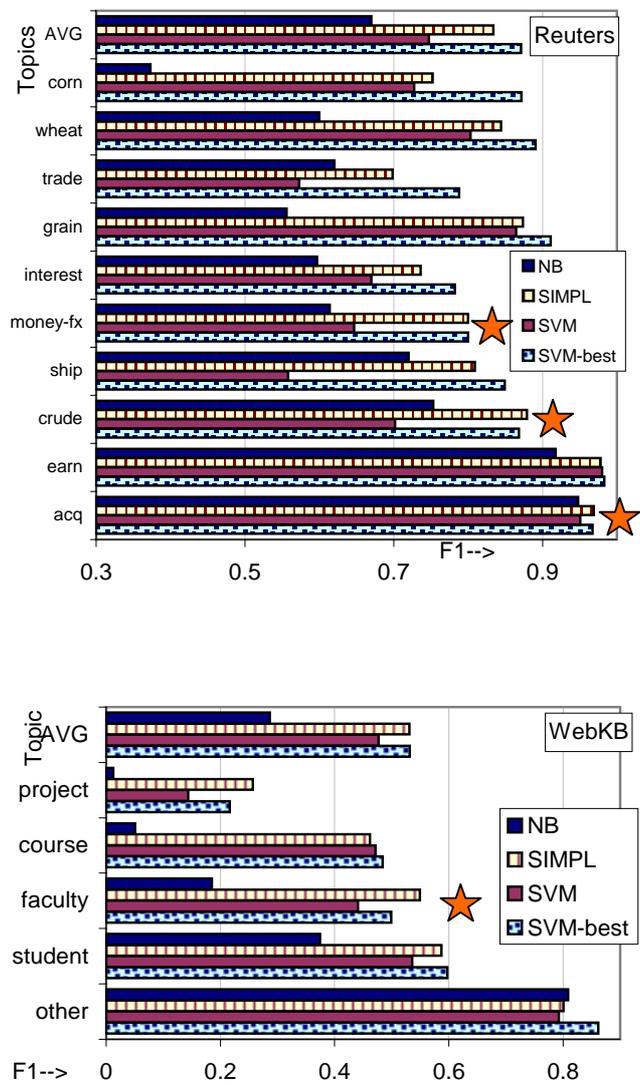


Fig. 10. Aggregate  $F_1$  scores for NB, SIMPL, and SVM for some of the most populated topics of the Reuters, 20NG, and WebKB data sets. We decisively beat NB (by a 15–20% margin) in most cases. We also frequently beat SVM (average 6% margin) and lose to SVM-best by a narrow margin (average 3%). In some cases we win over even SVM-best

tions generally suffice to separate all training collections we have seen.

Obviously, it would be reassuring to know that the projected test points are also clustered similar to the projected training points (otherwise a decision tree may perform poorly with the projected representations). Figure 8c–f show projections of training and test data, separated into different plots for the two classes. Clearly, projected points from training and test data appear to be distributed similarly in the plane.

How many projections are needed to retain enough information for C4.5 to achieve high accuracy? Figure 9 shows the effect of including up to the first three  $\alpha$ s. In all the cases, the first *two*  $\alpha$ s are sufficient for peak accuracy. The general trend is that precision and recall approach each other as we include more projections, improving  $F_1$ . The loss in one is more than compensated for by the gain in the other. This result also shows that we can improve beyond linear regression (Sect. 2.3) by using additional projections. It is also reassuring that including more  $\alpha$ s (up to five) than necessary never hurt accuracy.

#### 4.5 Accuracy

How does SIMPL measure up against SVM in terms of overall accuracy? Figure 10 shows a bird’s-eye view of many data sets and all algorithms: it compares the  $F_1$  score for naive Bayes (NB), SIMPL, SVM, and SVM-best (see Sect. 2.4). SIMPL beats NB in 33 out of 35 cases, and usually by large margins.

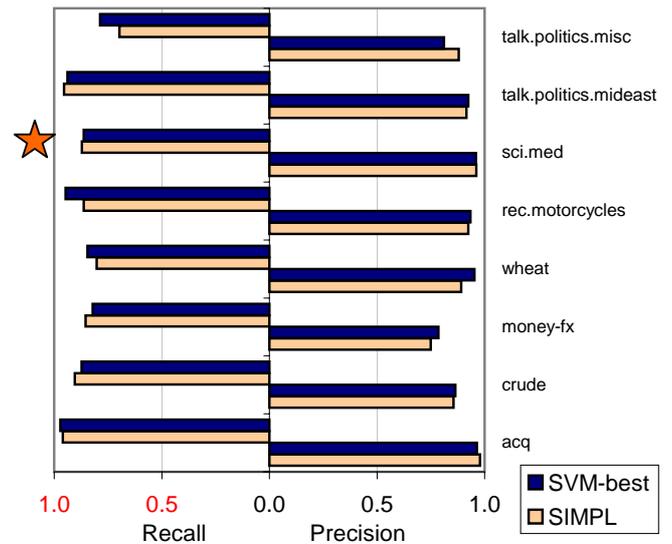
SIMPL beats SVM in 23 out of 35 cases. On average we are a few percent short of SVM-best, but there are several cases (marked by stars) where we beat even SVM-best. (The extent to which tuning the offset parameter  $b$  improved SVM-best beyond SVM surprised us and is not reported elsewhere. Tuning  $C$  had less effect.) We note that this is a comprehensive study over many diverse, standard data sets, and the high accuracy of SIMPL is very stable across the board.

$F_1$ , being the harmonic mean, favors algorithms whose recall is close to precision, which is the case with SIMPL. To be fair, a closer look shows that SIMPL usually loses to SVM-best by a small margin in either recall or precision but beats it in the other (Fig. 11).

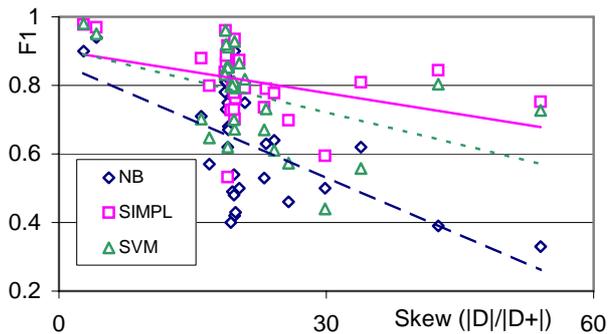
In a few cases, we beat SVM in *both* recall and precision. This is possible because we are not limited to a single planar separator. Because we use a decision tree in the projected space, we can learn, say, a function like EXOR (but we don’t expect a text classification task to pose such a challenge), which a linear SVM cannot. Although SVMs with more complex kernels may be used, they are slower to train than linear SVMs.

We also compared our accuracy with that of C4.5 run directly on the raw text, reported in earlier work [8, 14]. We see that although we too use C4.5, our accuracy is substantially better, thanks to our novel feature combination and transformation steps. In addition, SIMPL runs much faster than C4.5 on the raw data.

Finally, we show a scatterplot of  $F_1$  scores against the positive class *skew* (ratio of the number of documents to the number of documents with  $c = 1$ ) in Fig. 12. While all methods suffer somewhat from skew, clearly NB suffers most and SIMPL suffers least.



**Fig. 11.** Most often we lose to SVM-best by a small margin in one of recall and precision and beat it in the other. Stars mark where we win in both



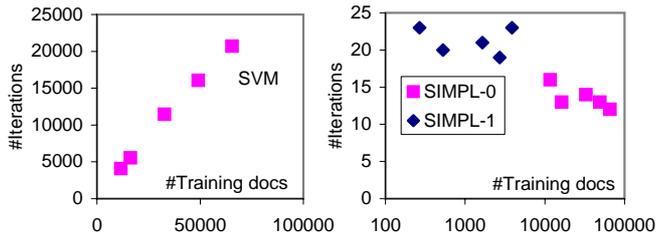
**Fig. 12.** SIMPL shows the least adverse effect of class skew on  $F_1$  accuracy among naive Bayes (NB), SIMPL, and SVM

#### 4.6 Performance

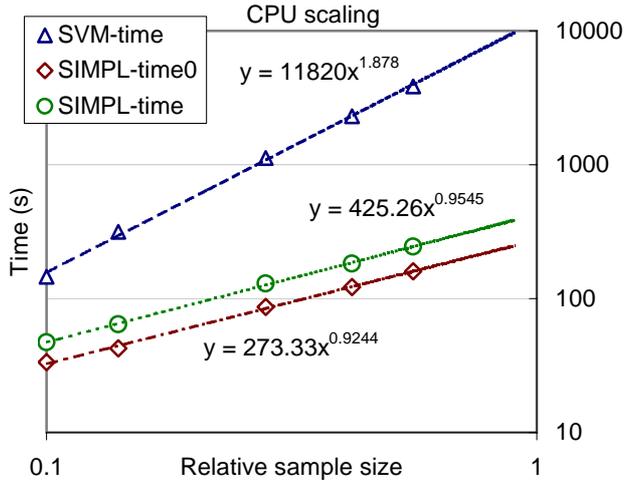
Having established that our accuracy is comparable to SVM, we turn to a detailed investigation of the scalability and IO behavior of the two algorithms. Throughout this part of the paper, “SVM” means SVMLIGHT. Comparison with the SVM variants discussed in Sect. 1.2 could be a fruitful area of future work.

We did not include the initial time required to turn the raw training documents into the compact representation required for sequential scans in the case of SIMPL and the LRU cache required by SVM. We started timing the algorithms only after these initial disk structures were ready. The time consumed for preprocessing depends on a host of nonstandard factors such as the raw representation and system policies: single vs. many files, stopword detection and word stemming and truncation, term weighting, etc. We used the OHSUMED and Dmoz data for performance measurements. We report on Dmoz, OHSUMED being broadly similar.

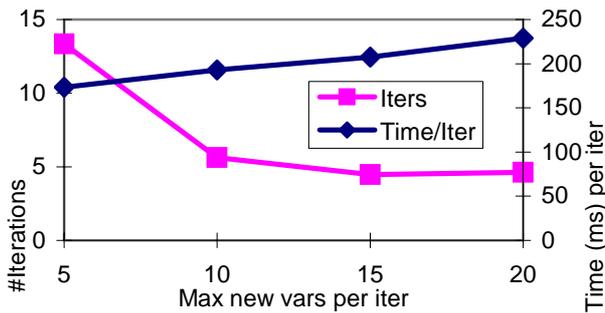
We will first compare the scaling of CPU time with training set size, assuming enough RAM is available. We observe in Fig. 13 that, unlike SVM, the number of iterations needed for our hill-climbing step is largely independent of the number



**Fig. 13.** Number of iterations needed by one run of SVM and two runs of SIMPL (one finding  $\alpha^{(0)}$ , the other,  $\alpha^{(1)}$ ). The number of SVM iterations seems to scale up with the problem size, unlike SIMPL



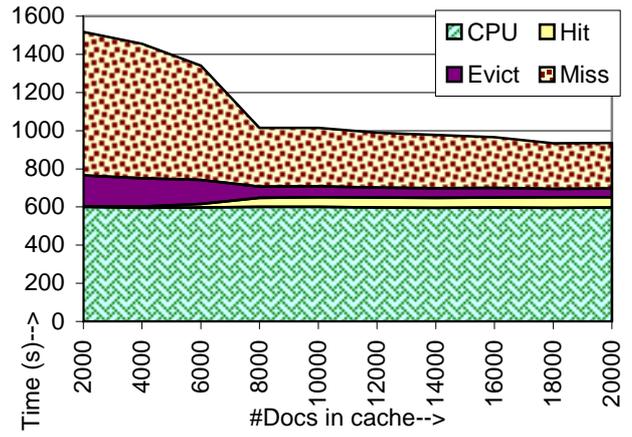
**Fig. 14.** Scaling of overall CPU time, excluding preprocessing time, with training set size for SVM and SIMPL, keeping all training data in memory. The line marked “SIMPL-time0” shows the time for finding just the first  $\alpha$ , and the line marked “SIMPL” shows the total time for SIMPL. A sample of 65000 documents were chosen from Dmoz



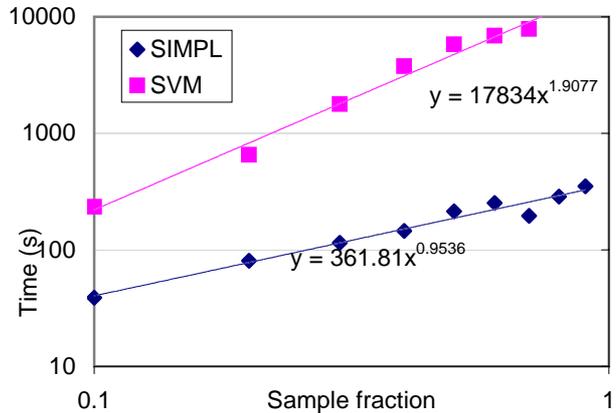
**Fig. 15.** Although cache misses and time per iteration increase if the number of new variables optimized per iteration is capped, the drastic savings in the number of iterations reduces overall time. The corpus and cache sizes were 24000 and 2000 documents and “-q 20” was used

of training documents. The time taken by single hill-climbing iteration is linear in the total input size, defined as  $\sum_d |t \in d|$ , plus  $O(n \log n)$  for  $n$  documents. Because  $\log n$  is small and the time for sorting is very small compared to the  $\alpha$  update step, the total time for SIMPL is expected to be essentially linear.

This is confirmed in the log-log plot shown in Fig. 14, where the least-square fit for SIMPL is roughly  $t \propto n^{0.955}$ ,



**Fig. 16.** CPU time, cache hit time, miss time, and eviction time plotted against the relative size of cache available to SVM

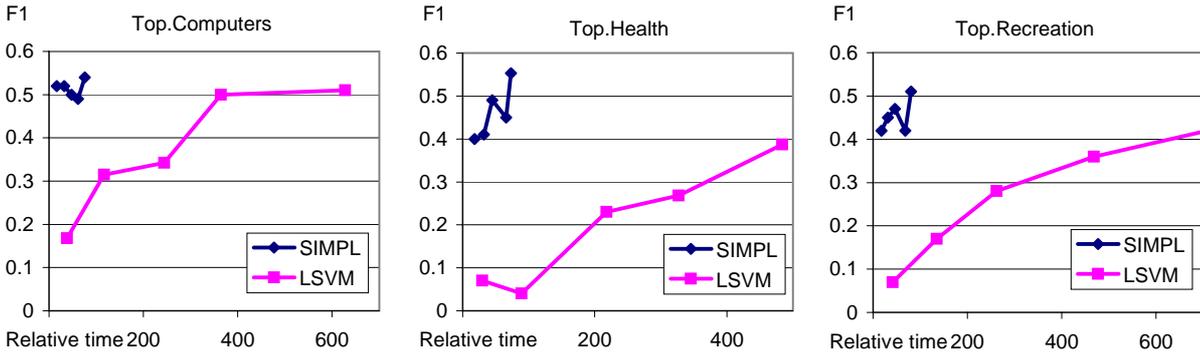


**Fig. 17.** When SVM and SIMPL are given the same amount of RAM as we scale up the size of the training set, SIMPL performs far better

where  $t$  is the execution time. In contrast, the regression for the running time of SVM is  $t \propto n^{1.88}$ , which corroborates earlier measurements by Platt, Joachims, and others. This difference translates to a running time ratio (SVM to SIMPL) of almost two orders of magnitude for  $n$  as small as 65000. For collections at the scale of Yahoo! or the full Dmoz data set (millions of documents) the ratio will reach several orders of magnitude.

All public (quadratic programming) SVM implementations that we know of, including SVMLIGHT and SMO, load all the training document vectors into memory. With limited memory, we expect the performance gap between SVM and SIMPL to be even larger. In our final set of experiments, we study the behavior of SVM with limited memory.

As mentioned above (Sect. 2.4), SVM optimizes the  $\lambda$ s corresponding to a small working set of documents at a time. In SVMLIGHT, the size of this working set (typically 10–50) is set by the “-q” option. There are standard heuristics for picking the next working set to speed up convergence, but applying these may replace the entire working set, reducing cache locality. SVMLIGHT provides another option, “-n,” which limits the number of new  $\lambda$ s that are permitted to enter the working set in each iteration. Reducing this increases cache locality but can lead to more iterations. A sample of this trade-off is shown



**Fig. 18.** Because SIMPL scales better than SVMLIGHT, we can achieve  $F_1$  comparable to SVMLIGHT in much shorter time. To match our time, SVMLIGHT must sample the training data, which reduces accuracy. For three Dmoz topics, we plot  $F_1$  vs. relative execution time

in Fig. 15. Even though a heavy replacement rate decreases cache hits and increases the time per iteration, the number of iterations is cut so drastically that a large value for  $-n$  is almost always a better choice.

How do cache hits and misses translate into running time overheads? This can be measured either by limiting physical memory available to the computer (in Linux, using a line of the form `append="mem=512M"` in `/etc/lilo.conf`), or by letting SVM do its own document cache management and instrumenting this cache. The former option is appealing to the end user. However, system processes, OS buffer cache, device interrupts, and paging make measurements unreliable. Moreover, the OS cache is physical and cannot exploit the structure of the SVM computation. Therefore, we expect large-scale SVM packages to implement their own caching. We used a disk without any file system as a raw block device (an ATA/66 7200RPM Seagate-ST330630A drive with 2 MB on-board cache on `/dev/hdc1`), which precluded interference owing to OS buffering. We built an LRU cache on it with a preconfigured main-memory quota. Servicing a miss usually involved exactly one seek on disk unless the data were in the disk's on-board cache.

Figure 16 shows a break-up of times spent in computation (CPU), hit, miss, and eviction processing. This sample from Dmoz had 23428 documents (23 MB of sparse term vectors) with 360000 features, so SIMPL needs only 6 MB of RAM and runs for only 80 s with close-to 100% CPU utilization. In contrast, if SVM is given 6 MB of cache (about 6000 documents), it takes over 1300 seconds, of which 60% is spent in servicing evictions and misses.

Extending from the small-scale experiment above, we present in Fig. 17 a large-scale experiment where we go from a 10% to a 100% sample of our Dmoz data (117920 documents, 110 MB in RAM). For each sample we determine the amount of RAM needed by SIMPL, give that quantity of cache to SVM, and compare the running times. The graph is superficially similar to Fig. 14, but a closer look shows that the ratio of SVM to SIMPL running times is larger owing to cache overheads. Summarizing, SIMPL beats SVM with respect to both CPU and cache performance, but the near-quadratic CPU scaling of SVM makes cache overheads appear less serious than they really are: the *total* time spent by SIMPL is less than 20% of the time spent by SVM on cache management alone.

Finally, we show some examples of the overall trade-off between accuracy and running time of SIMPL and (quadratic

programming) SVMs. In Fig. 18, we consider three of the top-level topics of Dmoz. The plots show  $F_1$  accuracy against training time. Points are generated by taking uniform random samples of different sizes (1, 5, 10%, etc.) of the entire training set of 120000 documents. SVMLIGHT approaches the best  $F_1$  at a large value of execution time, whereas SIMPL achieves peak  $F_1$  with very short execution times because it can handle much larger samples within that time. Obviously, smarter sampling techniques [20, 29] will fare better than uniform sampling and should be compared to SIMPL in future work.

## 5 Conclusion

We have presented SIMPL, a new classifier for high-dimensional data such as text. SIMPL is very simple to understand and easy to implement. SIMPL is very fast, scaling linearly with input size, as opposed to SVM involving quadratic programming, which shows almost quadratic scaling. SIMPL uses efficient sequential scans over the training data, unlike popular SVM implementations, which have less efficient disk and cache access patterns and poorer locality of reference. This performance boost carries little or no penalty in terms of accuracy: we often beat SVM in the  $F_1$  measure and closely match SVM in recall and precision. SIMPL beats naive Bayes and decision-tree classifiers decisively for text learning tasks.

Our work shows that even though SVMs are elegant, powerful, and theoretically appealing, they have not rendered the search for practical and IO-efficient alternatives unnecessary or fruitless. A natural area of future work is to identify properties of data sets that guarantee near-SVM accuracy using SIMPL. Another area of applied work is to test SIMPL vis-a-vis a nonlinear SVM for nontextual training data, which are more difficult to separate than text. It would also be of interest to compare SIMPL with SVM variants as well as use the success of SIMPL to guide the search and optimization process in SVMs.

*Acknowledgements.* Thanks to Pedro Domingos, Gary Flake, and Vladimir Vapnik for helpful discussions, Thorsten Joachims for generous help with SVMLIGHT, Kunal Punera for help with preparing some data sets, and Shantanu Godbole for helpful comments on the manuscript. This research was partially supported by Tata Consultancy Services and IBM Corporation.

## References

1. Agrawal R, Bayardo RJ, Srikant R (2000) Athena: mining-based interactive management of text databases. In: Proceedings of the 7th international conference on extending database technology (EDBT), Konstanz, Germany, March 2000. <http://www.almaden.ibm.com/cs/people/rAgrawal/papers/athena.ps>
2. Basu C, Hirsh H, Cohen WW (1998) Recommendation as classification: using social and content-based information in recommendation. In: Proceedings of the 15th national conference on artificial intelligence, Madison, WI, July 1998, pp 714–720
3. Chakrabarti S, Dom B, Agrawal R, Raghavan P (1998) Scalable feature selection, classification and signature generation for organizing large text databases into hierarchical topic taxonomies. VLDB J [http://www.cs.berkeley.edu/~soumen/VLDB54\\_3.pdf](http://www.cs.berkeley.edu/~soumen/VLDB54_3.pdf)
4. Cooke T (2002) Two variations on Fisher's linear discriminant for pattern recognition. IEEE Trans Patt Analysis Machine Intell (PAMI) 24(2):268–273 <http://www.computer.org/~tpami/tp2002/i0268abs.htm>
5. Dasgupta S (1999) Learning mixtures of Gaussians. In: FOCS, pp 634–644 <http://charlotte.ucsd.edu/users/dasgupta/papers/focs2.ps>
6. Dasgupta S (2000) Experiments with random projection. UAI 16:143–151 <http://charlotte.ucsd.edu/users/dasgupta/papers/random.ps>
7. Duda R, Hart P (1973) Pattern classification and scene analysis. Wiley, New York
8. Dumais S, Platt J, Heckerman D, Sahami M (1998) Inductive learning algorithms and representations for text categorization. In: Proceedings of the 7th conference on information and knowledge management, 1998. <http://www.research.microsoft.com/~jplatt/cikm98.pdf>
9. Frankl P, Maehara H (1988) The Johnson-Lindenstrauss lemma and the sphericity of some graphs. J Combin Theory B 44:355–362
10. Friedman JH (1987) Exploratory projection pursuit. J Am Stat Assoc 82:249–266
11. Fung G, Mangasarian OL (2001) Proximal support vector classifiers. In: Provost F, Srikant R (eds) Proceedings of the 7th ACM SIGKDD international conference on knowledge discovery and data mining, San Francisco, August 2001, pp 77–86 University of Wisconsin Data Mining Institute Technical Report 01-02, <http://www.cs.wisc.edu/~gfung/>
12. Fung G, Mangasarian OL (2002) Incremental support vector machine classification. In: Proceedings of the 2nd SIAM international conference on data mining, Arlington, VA, April 2002, pp 247–260 University of Wisconsin Data Mining Institute Technical Report 01-08, <ftp://ftp.cs.wisc.edu/pub/dmi/tech-reports/01-08.ps>
13. Graefe G, Fayyad UM, Chaudhuri S (1998) On the efficient gathering of sufficient statistics for classification from large SQL databases. In: Knowledge discovery and data mining, vol 4. AAAI Press, New York, pp 204–208
14. Joachims T (1998) Text categorization with support vector machines: learning with many relevant features. In: Nédellec C, Rouveirol C (eds) Proceedings of ECML-98, 10th European conference on machine learning, Lecture notes in computer science, vol 1398. Springer, Berlin Heidelberg New York, pp 137–142
15. Joachims T (1999) Making large-scale SVM learning practical. In: Schölkopf B, Burges C, Smola A (eds) Advances in kernel methods: support vector learning. MIT Press, Cambridge, MA [http://www-ai.cs.uni-dortmund.de/DOKUMENTE/joachims\\_99a.pdf](http://www-ai.cs.uni-dortmund.de/DOKUMENTE/joachims_99a.pdf)
16. Joachims T (2001) A statistical learning model of text classification for support vector machines. In: Croft WB, Harper DJ, Kraft DH, Zobel J (eds) Proceedings of the international conference on research and development in information retrieval, vol 24, New Orleans, September 2001, ACM Press, New York, pp 128–136
17. Johnson RA, Wichern DW (2001) Applied multivariate statistical analysis, 3rd edn. Prentice-Hall, New Delhi
18. Kleinberg JM (1997) Two algorithms for nearest-neighbor search in high dimensions. In: Proceedings of the ACM symposium on theory of computing, pp 599–608
19. LeCun Y, Simard PY, Pearlmetter B (1993) Automatic learning rate maximization by on-line estimation of the Hessian's eigenvectors. In: Hanson SJ, Cowan JD, Lee-Giles C (eds) Advances in neural information processing systems, vol 5. Morgan Kaufmann, San Mateo, CA, pp 156–163
20. Lee YJ, Mangasarian OL (2001) RSVM: reduced support vector machines. In: Proceedings of the 1st SIAM international conference on data mining, Chicago, April 2001. [http://www.siam.org/meetings/sdm01/pdf/sdm01\\_13.pdf](http://www.siam.org/meetings/sdm01/pdf/sdm01_13.pdf)
21. Lewis DD (1997) The Reuters-21578 text categorization test collection, 1997. <http://kdd.ics.uci.edu/databases/reuters21578/reuters21578.html>
22. Lewis DD, Schapire RE, Callan JP, Papka R (1996) Training algorithms for linear text classifiers. In: Frei HP, Harman D, Schäuble P, Wilkinson R (eds) Proceedings of SIGIR-96, 19th ACM international conference on research and development in information retrieval, ACM Press, New York, pp 298–306
23. Mangasarian OL, Musicant DR (1999) Successive over-relaxation for support vector machines. In: IEEE Trans Neural Netw 10:1032–1037 <ftp://ftp.cs.wisc.edu/math-prog/tech-reports/98-18.ps>
24. Mangasarian OL, Musicant DR (2000) Lagrangian support vector machines. Technical Report 00-06, Data Mining Institute, University of Wisconsin, Madison, June 2000. <http://www.cs.wisc.edu/~musicant/>
25. McCallum A (1998) Bow: a toolkit for statistical language modeling, text retrieval, classification and clustering. Software available from <http://www.cs.cmu.edu/~mccallum/bow/>
26. McCallum A, Nigam K (1998) A comparison of event models for naive Bayes text classification. In: AAAI-ICML-98 workshop on learning for text categorization, AAAI Press, pp 41–48 Also technical report WS-98-05, CMU, <http://www.cs.cmu.edu/~knigam/papers/multinomial-aaaiws98.pdf>
27. Murthy SK, Kasif S, Salzberg S (1994) A system for induction of oblique decision trees. J Artif Intell Res 2:1–32
28. Nigam K, Lafferty J, McCallum A (1999) Using maximum entropy for text classification. In: IJCAI-99 workshop on machine learning for information filtering, pp 61–67. <http://www.cs.cmu.edu/~knigam/> and <http://www.cs.cmu.edu/~mccallum/papers/maxent-ijcaiws99.ps.gz>
29. Pavlov D, Mao J, Dom B (2000) Scaling-up support vector machines using boosting algorithm. In: Proceedings of the international conference on pattern recognition (ICPR), Barcelona, September 2000. <http://www.cvc.uab.es/ICPR2000/>

30. Platt J (1998) Sequential minimal optimization: a fast algorithm for training support vector machines. Technical Report MSR-TR-98-14, Microsoft Research. <http://www.research.microsoft.com/users/jplatt/smoTR.pdf>
31. Sahami M, Dumais S, Heckerman D, Horvitz E (1998) A Bayesian approach to filtering junk E-mail. In: Learning for text categorization: papers from the 1998 workshop, Madison, WI, AAAI Technical Report WS-98-05
32. Schapire RE (2001) The boosting approach to machine learning: an overview. In: Proceedings of the MSRI workshop on nonlinear estimation and classification, Berkeley, CA, March 2001. <http://stat.bell-labs.com/who/cocteau/nec/> and <http://www.research.att.com/~schapire/boost.html>
33. Schutze H, Hull DA, Pederson JO (1995) A comparison of classifiers and document representations for the routing problem. In: SIGIR, pp 229–237. <ftp://parcftp.xerox.com/pub/qca/SIGIR95.ps>
34. Shafer JC, Agrawal R, Mehta M (1996) SPRINT: A scalable parallel classifier for data mining. VLDB, pp 544–555
35. Shashua A (1999) On the equivalence between the support vector machine for classification and sparsified Fisher's linear discriminant. Neural Processing Lett 9(2):129–139 <http://www.cs.huji.ac.il/~shashua/papers/fisher-NPL.pdf>
36. Swayne DF, Cook D, Buja A (1998) XGobi: interactive dynamic data visualization in the x window system. J Computat Graph Stat 7(1) <http://lib.stat.cmu.edu/general/XGobi/>
37. Vapnik V, Golowich S, Smola AJ (1996) Support vector method for function approximation, regression estimation, and signal processing. In: Advances in neural information processing systems. MIT Press, Cambridge, MA
38. Witten IH, Frank E (1999) Data mining: practical machine learning tools and techniques with Java implementations. Morgan Kaufmann, San Francisco