



Effective transductive learning via objective model selection

Ran El-Yaniv ^{a,*}, Leonid Gerzon ^b

^a Computer Science Department, Technion—Israel Institute of Technology, Haifa 32000, Israel

^b Department of Mathematics, Technion—Israel Institute of Technology, Haifa 32000, Israel

Received 21 June 2004; received in revised form 16 March 2005

Available online 3 May 2005

Communicated by L. Bottou

Abstract

This paper is concerned with transductive learning. We study a recent transductive learning approach based on clustering. In this approach one constructs a diversity of unsupervised models of the unlabeled data using clustering algorithms. These models are then exploited to construct a number of hypotheses using the labeled data and the learner selects an hypothesis that minimizes a transductive error bound, which holds with high probability. Empirical examination of this approach, implemented with ‘spectral clustering’, on a suite of benchmark datasets from the UCI repository, indicates that the new approach is effective and comparable with one of the best known transductive learning algorithms to-date.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Transduction; Transductive learning; Clustering; Error bounds; Semi-supervised learning

1. Introduction

Recently, there has been a growing need for and interest in learning algorithms capable of utilizing

both labeled and unlabeled data (see, for example, the survey by [Seeger, 2002](#)). Despite widespread confusion, a distinction between *inductive* ‘semi-supervised’ learning and *transductive* learning should be made. While both approaches attempt to utilize unlabeled data (in addition to the standard labeled training set) to accelerate the learning process, in transduction one is only interested in classifying the given unlabeled points. There is no need to construct a general hypothesis capable of classifying unseen data points. Therefore, at the

* Corresponding author. Tel.: +972 544 227 787; fax: +972 4 829 3379.

E-mail addresses: rani@cs.technion.ac.il (R. El-Yaniv), gleonid@tx.technion.ac.il (L. Gerzon).

URL: <http://www.cs.technion.ac.il/~rani> (R. El-Yaniv).

outset transduction appears to be an easier problem than (semi-supervised) induction.¹

Transductive inference was introduced by Vapnik more than 20 years ago (Vapnik, 1982), and through the years he has been its greatest advocate.² Although the interest in this learning framework is rising, there is still lack of *provably useful* learning principles and algorithms that are both theoretically founded as well as good performers on real world data.

In this paper we implement and study an idea for transductive learning that was recently proposed by Derbeko et al. (2004a,b).³ In this scheme the learner applies a clustering algorithm over the unlabeled data to generate several (unsupervised) models. The learner then utilizes the labeled data to guess labels for entire clusters (so that all points in the same cluster have the same label). This way the algorithm forms a number of hypotheses, one of which is then selected based on an error bound for transduction.

The natural idea of first clustering the unlabeled data and then assigning labels to clusters has been around for a long time and there are plenty of heuristic procedures that attempt to learn using this approach within a semi-supervised or transductive settings (see, for example, Section 2.1 in Seeger, 2002 and references therein). However, to the best of our knowledge, none of the existing procedures was ever theoretically justified in terms of provable reduction of the true risk.

The new approach, given here the acronym CLUST, is both simple and rests on solid theoretical ground. As we show here, CLUST is competitive with one of the best transductive learning methods. In particular, we show in Section 5 a rather

extensive set of experiments where we study the performance of CLUST and compare it to the recent ‘Spectral Graph Transducer’ (SGT) algorithm proposed by Joachims (2003). Our results indicate that CLUST fares well in this comparison. Note that the SGT algorithm is shown in (Joachims, 2003) to outperform other well known transductive learning methods such as the transductive support vector machine (TSVM) of Joachims (1999).

2. The transduction setting

Throughout the paper we consider a classification setting. For simplicity (and initial evaluation of the basic ideas), we focus on *binary* classification.⁴ We rely on the transduction formulation proposed in (Vapnik, 1998, Chapt. 8). In this formulation the learner is given a *full sample* X_{m+u} of $m+u$ unlabeled points. After observing these points, a *training set* X_m of m points is chosen uniformly at random among all m -subsets of the full sample. The labels of all the points in the training set are provided to the learner. Based on the labeled and unlabeled points, the learner’s goal is to predict, as accurately as possible, the labels of the remaining unlabeled points, which constitute the *test set*, $X_u = X_{m+u} \setminus X_m$.

Remark 1. The random choice of the training set from the full sample as described above is equivalent to choosing uniformly at random m points from the full sample, *without replacement*. Therefore, unlike the standard inductive setting, points in the training (and test) sets are *dependent*.

Remark 2. It is not assumed that points in the full sample are drawn i.i.d. from some (unknown) distribution. The only distributional requirement here is that the training set is drawn uniformly at random from all possible m -subsets. Vapnik considered a second transduction setting in which the learner receives training and test sets, which are assumed to be drawn i.i.d. from some unknown

¹ We can make analogies between transduction and a “take-home exam” (where the student gets to see the questions before studying), and between semi-supervised learning and a standard “classroom exam” (where the student gets to see only exam questions from previous years before studying).

² In his 1995 book (Vapnik, 1995) Vapnik writes: “*This model [transduction] can provide the strongest contribution to the 2000 years of discussions about the essence of human reason*”.

³ These papers derived this transductive learning scheme from a general PAC-Bayesian error bound for transduction but did not consider any implementation details or evaluate its practical significance.

⁴ The CLUST algorithm can be easily extended to handle multi-class classifications in a natural manner.

distribution. This second setting is more suitable for some applications (and may appear more natural). However, Theorem 8.1 in (Vapnik, 1998) states that error bounds for the first setting imply the same bounds for the second setting. We therefore restrict our attention, throughout this paper, to the first setting and note that the error bounds and algorithms considered in this paper are applicable in the second setting as well.

Throughout the paper we use the following notations and definitions. We denote by \mathcal{H} a set of binary hypotheses consisting of functions from the input space \mathcal{X} to $\mathcal{Y} = \{\pm 1\}$. The learner's goal is to choose a good hypothesis from \mathcal{H} . For each $h \in \mathcal{H}$ and a set $Z = x_1, \dots, x_N$ of samples define

$$R_h(Z) \triangleq \frac{1}{N} \sum_{i=1}^N \ell(h(x_i), \phi(x_i)), \quad (1)$$

where $\phi(x_i) \in \mathcal{Y}$ is the true label of x_i and $\ell(\cdot, \cdot)$ is the 0/1-loss. We make use of the following quantities, which are all instances of (1). The quantity $R_h(X_{m+u})$ is called the *full sample risk* of the hypothesis h , $R_h(X_u)$ is referred to as the *transduction risk* or *test error* (of h), and $R_h(X_m)$ is the *training error* (of h). It is important to observe that in our transduction setting $R_h(X_{m+u})$ is *not* a random variable, but both $R_h(X_m)$ and $R_h(X_u)$ are random variables, due to the random selection of the samples X_m from X_{m+u} . The goal of the learner is to choose $h \in \mathcal{H}$ with minimal transduction risk $R_h(X_u)$.

3. Vapnik's bounds for transduction

The transductive error bounds we present here are based on a slight adaptation of Vapnik's arguments presented in (Vapnik, 1982, 1998). The idea is to bound the deviation between the two random variables $R_h(X_u)$ and $R_h(X_m)$, which are both concentrated around their mean $R_h(X_{m+u})$. This deviation can be explicitly written as a tail of an hypergeometric distribution, which converges exponentially fast, and in fact, has faster convergence rates than the tail of the binomial distribution (which governs the learning rates in standard inductive learning).

Fix m and u and consider some hypothesis $h \in \mathcal{H}$. Suppose that h makes k_h errors on the full sample (i.e. $k_h = (m+u)R_h(X_{m+u})$). Consider a random choice of the training set X_m from the full sample, and let $B(r, k_h)$ be the probability that h makes exactly r errors over the training set X_m . This probability is by definition the hypergeometric distribution, given by

$$B(r, k_h) \triangleq \frac{\binom{k_h}{r} \binom{m+u-k_h}{m-r}}{\binom{m+u}{m}}.$$

Define

$$\begin{aligned} C(\varepsilon, k_h) &\triangleq \Pr\{R_h(X_u) - R_h(X_m) > \varepsilon\} \\ &= \Pr\left\{\frac{k_h - r}{u} - \frac{r}{m} > \varepsilon\right\} \\ &= \sum_r B(r, k_h), \end{aligned}$$

where the summation is over all values of r such that $\max\{k_h - u, 0\} \leq r \leq \min\{m, k_h\}$ and

$$\frac{k_h - r}{u} - \frac{r}{m} > \varepsilon. \quad (2)$$

Define

$$\Gamma_1(\varepsilon) \triangleq \max_k C(\varepsilon, k), \quad (3)$$

$$\Gamma_2(\varepsilon) \triangleq \max_k C\left(\sqrt{\frac{k}{m+u}} \cdot \varepsilon, k\right). \quad (4)$$

We now state Vapnik's implicit bounds for transduction. The bounds are adapted slightly in order to incorporate a prior probability over \mathcal{H} (the original bounds deal with a uniform prior). This change makes them similar to McAllester's PAC-Bayesian bounds in (McAllester, 1999). Also, we note that the original bounds in (Vapnik, 1982) are two-sided and the following theorem considers one-sided bounds.⁵

Theorem 3 (Vapnik). *Let δ be given, let \mathbf{p} be a prior distribution over \mathcal{H} that may depend on the*

⁵ Specifically, in the original bound (Vapnik, 1982), the two-sided condition $|\frac{k_h - r}{u} - \frac{r}{m}| > \varepsilon$ is used instead of condition (2).

full sample X_{m+u} . Let $\varepsilon_1^*(h)$ (resp. $\varepsilon_2^*(h)$) be the minimal value of ε that satisfies $\Gamma_1(\varepsilon) \leq \mathbf{p}(h)\delta$ (resp. $\Gamma_2(\varepsilon) \leq \mathbf{p}(h)\delta$). Then, with probability at least $1 - \delta$, for all $h \in \mathcal{H}$,

$$R_h(X_u) < R_h(X_m) + \varepsilon_1^*(h) \tag{5}$$

and respectively,

$$R_h(X_u) < R_h(X_m) + \sqrt{R_h(X_{m+u})} \cdot \varepsilon_2^*(h). \tag{6}$$

Proof. We sketch the proof of (6).

$$\begin{aligned} & \Pr \left\{ \exists h \in \mathcal{H} : \frac{R_h(X_u) - R_h(X_m)}{\sqrt{R_h(X_{m+u})}} > \varepsilon_2^*(h) \right\} \\ & \leq \sum_{h \in \mathcal{H}} C \left(\sqrt{\frac{k_h}{m+u}} \varepsilon_2^*(h), k_h \right) \leq \sum_{h \in \mathcal{H}} \Gamma(\varepsilon_2^*(h)) \\ & \leq \sum_{h \in \mathcal{H}} \mathbf{p}(h)\delta = \delta, \end{aligned} \tag{7}$$

where (7) follows from the union bound. \square

It is easy to eliminate the (unknown) quantity $R_h(X_{m+u})$ from the right hand side of (6). Subtracting $R_h(X_m)$ from both sides of (6), then squaring both sides and substituting $\frac{m}{m+u}R_h(X_m) + \frac{u}{m+u}R_h(X_u)$ for $R_h(X_{m+u})$, we get a quadratic inequality where the “unknown” variable is $R_h(X_u)$. Solving for $R_h(X_u)$ yields the following result, as in Eq. (8.15) in (Vapnik, 1998).

Corollary 4 (Vapnik). *Under the conditions of Theorem 3,*

$$\begin{aligned} R_h(X_u) & \leq R_h(X_m) + \frac{\varepsilon^*(h)^2 u}{2(m+u)} \\ & + \varepsilon_2^*(h) \sqrt{R_h(X_m) + \left(\frac{\varepsilon_2^*(h)u}{2(m+u)} \right)^2}. \end{aligned} \tag{8}$$

We thus have two general bounds: The “absolute” bound (5) and the “relative” bound (8).⁶

It is possible to show that the relative bound is tighter than the absolute bound for smaller values of the empirical error $R_h(X_m)$ and that the absolute bound is tighter for larger values of $R_h(X_m)$ (see,

⁶ In his book Vapnik only focuses on the relative bound. Both bounds are mentioned in (Bottou et al., 1994).

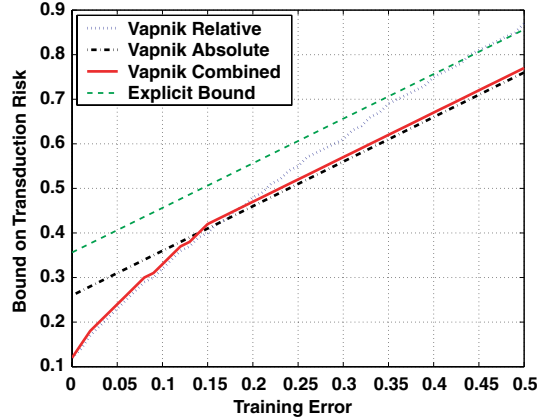


Fig. 1. The Vapnik absolute, relative and combined bounds and the explicit approximation bound (9); all bounds are computed with $m = u = 100$ and $\mathbf{p}(h) = 1/608$. This prior corresponds to $C = 20$, $T = 1$ and $\tau = 5$ (see Section 4). The Vapnik relative and absolute bounds are computed with $\delta = 0.05$, and therefore the Vapnik combined bound corresponds to $\delta = 0.1$. The curve of the explicit bound was also calculated with $\delta = 0.1$.

for example Fig. 1). We therefore use the union bound to combine the bounds into one bound by applying each of the bounds with confidence $\delta/2$ and taking the minimum of the two at each point. Throughout the paper we refer to the combined bound as “the Vapnik bound”. Note that a related result has recently been presented in (Blum and Langford, 2003).⁷

The Vapnik bound is rather tight and therefore, it is intriguing to see if the bound can be effectively utilized in a practical model selection, as we attempt in this paper. Possible sources of slackness in this bound are only introduced through the utilization of the union bound in (7) and the definitions of $\Gamma_i(\varepsilon)$ in Eqs. (3) and (4). However, note that $\varepsilon_i^*(h)$ ($i = 1, 2$) is a complicated *implicit* function of m , u , $\mathbf{p}(h)$ and δ leading to a bound that is difficult to interpret and (as noted also by Vapnik) must be computed in order to be used. A number of *explicit* (but looser) bounds for trans-

⁷ In particular, Theorem 6 in (Blum and Langford, 2003) presents a similar bound for transduction based on a direct calculation of the hypergeometric tail. Note however that the bound in (Blum and Langford, 2003) was the first to consider a general prior \mathbf{p} ; Vapnik’s bound was stated with a uniform prior.

duction are presented in (Derbeko et al., 2004b). For example, under the same conditions of Theorem 3, the following bound holds with probability at least $1 - \delta$, uniformly for all $h \in \mathcal{H}$,

$$R_h(X_u) \leq R_h(X_m) + \sqrt{\left(\frac{m+u}{u}\right) \left(\frac{\ln \frac{1}{p(h)} + \ln \frac{1}{\delta}}{m}\right)}. \quad (9)$$

Fig. 1 shows the above bounds as a function of the empirical error for realistic assignments of the other problem parameters.

4. Transduction via clustering

In order to utilize bounds, such as the Vapnik bound or (9), one has to specify a prior distribution p over \mathcal{H} . As observed in (Derbeko et al., 2004a), within the transduction setting it is often convenient to apply such bounds because the prior distribution can be constructed *after* observing the (unlabeled) full sample. Derbeko et al. also proposed the following learning scheme, based on clustering.

Let $\mathcal{A}(\tau, \theta)$ be a parametric clustering algorithm that can partition the full sample into any desirable number τ of clusters. The hyper-parameter $\theta \in \Theta$ can specify different clustering algorithms and/or their parameters (e.g., the variance in spectral clustering; see below). For $\tau = 2, \dots, C$ ($C \leq m$) and $\theta \in \Theta$, we apply $\mathcal{A}(\tau, \theta)$ on the full sample and obtain a collection of partitions $\mathcal{C}_{\tau, \theta}$, $\tau = 2, \dots, C$, $\theta \in \Theta$. Setting $T = |\Theta|$ we thus have $(C - 1)T$ partitions of X_{m+u} . We now receive the training set X_m and would like to utilize it in order to construct $(C - 1)T$ hypotheses, one for each partition. For each partition $\mathcal{C}_{\tau, \theta}$, we consider each of its τ clusters and label all the points in the same cluster according to a majority vote among all the labeled points (from X_m) falling within that cluster. If none of the labeled training points is in the cluster (or in case of a tie), we arbitrarily give the entire cluster a single label from $\{\pm 1\}$. Hence, we end up with a collection of $(C - 1)T$ hypotheses. We consider each hypothesis h , compute its training error $R_h(X_m)$ and then compute its Vapnik bound based on appropriate prior

$p(h)$ (see below). We select the hypothesis with the minimal error bound.

We now discuss the construction of the prior. Initially \mathcal{H} consists of all 2^{m+u} possible dichotomies of the full sample. Let $\mathcal{H}_{\tau, \theta} \subset \mathcal{H}$ be the set of all binary hypotheses that assign an identical label to all points in the same cluster in partition $\mathcal{C}_{\tau, \theta}$. Clearly, $|\mathcal{H}_{\tau, \theta}| = 2^\tau$ for each θ , because $\mathcal{H}_{\tau, \theta}$ contains all possible dichotomies over τ clusters. Without any prior information about the quality of the clustering as a function of the parameter θ , each hypothesis in $\mathcal{H}_\tau \triangleq \cup_\theta \mathcal{H}_{\tau, \theta}$ is assigned the same prior, which will be proportional to $1/(T \cdot 2^\tau)$. Also, without any additional information we do not know the “correct” number of clusters. Therefore, each of the $C - 1$ subsets \mathcal{H}_τ , $\tau = 2, \dots, C$, is assigned the same probability mass, $1/(C - 1)$. The final prior distribution is, thus,

$$p(h) = \begin{cases} \frac{1}{(C - 1)T2^\tau}, & \text{if } h \in \mathcal{H}_\tau, \text{ for some } 2 \leq \tau \leq C; \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

To summarize, each of the partitions $\mathcal{C}_{\tau, \theta}$ induces one hypothesis $h_{\tau, \theta}$, through the use of the training set. Clearly, $\sum_{\tau, \theta} p(h_{\tau, \theta}) = 1$, and therefore, the Vapnik bound applies. Similarly, using the prior (10) in (9) we have, that with probability at least $1 - \delta$ (over random choices of the training set from the full sample), for all $h \in \mathcal{H}_\tau$,

$$R_h(X_u) \leq R_h(X_m) + \sqrt{\left(\frac{m+u}{u}\right) \left(\frac{\tau + \log((C - 1)T) + \ln \frac{1}{\delta}}{m}\right)}. \quad (11)$$

It can be shown that the right hand side of (11) upper bounds the Vapnik bound. From (11) it follows that if we are lucky enough to compute a “quality clustering” of the data, which gives rise to an hypothesis with a small training error, based on a small number of clusters, then we have a guarantee that the transductive test error of the corresponding hypothesis will be small. A nice property of this bound is that the divergence between the true and empirical risks increases only

logarithmically with the number of models we consider (i.e., $(C - 1)T$). The pseudo-code in Fig. 2 summarizes the CLUST algorithm.

4.1. Which clustering algorithm?

A successful implementation of the CLUST algorithm may critically depend on the choice of a clustering algorithm $\mathcal{A}(\tau, \theta)$. Although the area of unsupervised learning is very active and offers a great many types of algorithms, the task of selecting a “suitable” clustering algorithm for a given data is ill-defined and cannot have a principled solution without considering the objective function. However, the strength of the CLUST algorithm should be exactly in its ability to consider *many* clusterings simultaneously and automatically

choose a relatively good one! Therefore, we seek an algorithm that is reasonably fast, and flexible in the sense that it can generate a diversity of models based on a diversity of metrics. Particularly convenient candidates are the *spectral clustering* (or kernel PCA) algorithms, which are reasonably fast and well-motivated, and flexible in the sense that they use a kernel to assign the initial pairwise dissimilarities. For our implementation we (rather arbitrarily) selected the spectral algorithm of (Ng et al., 2002). This algorithm uses an RBF kernel to assign pairwise dissimilarities and has one parameter σ (which here is interchangeably referred to as θ). It can also cluster the data into any desired number τ of clusters. Due to lack of space we refer the reader to (Ng et al., 2002) for the details of this algorithm.

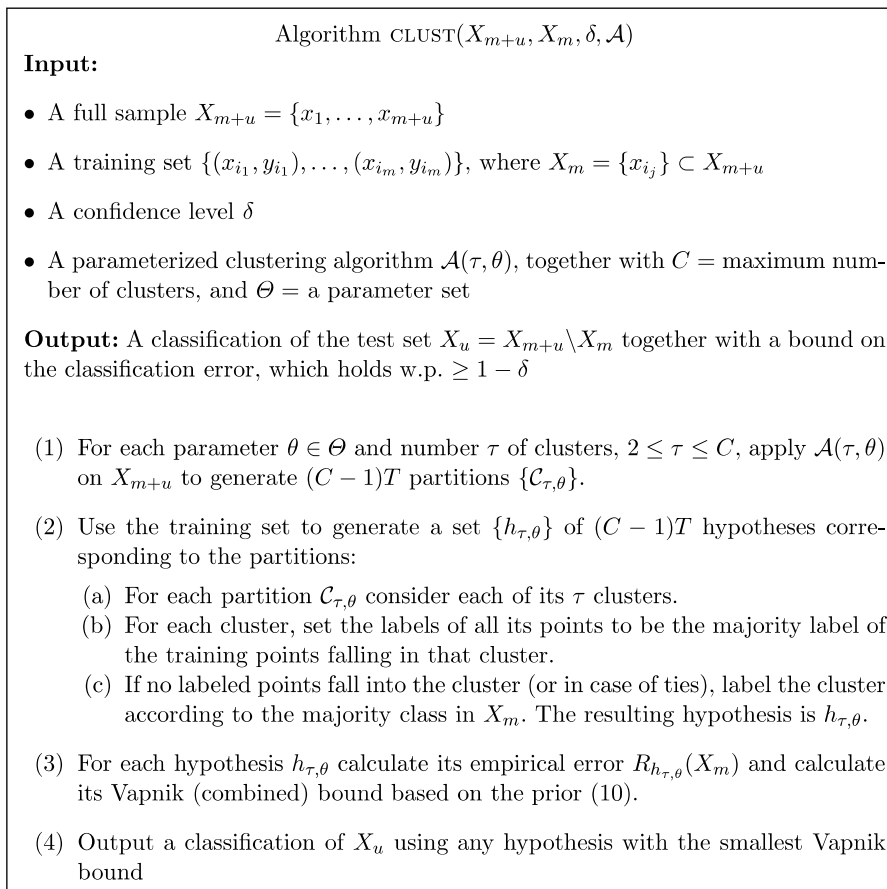


Fig. 2. CLUST pseudo-code.

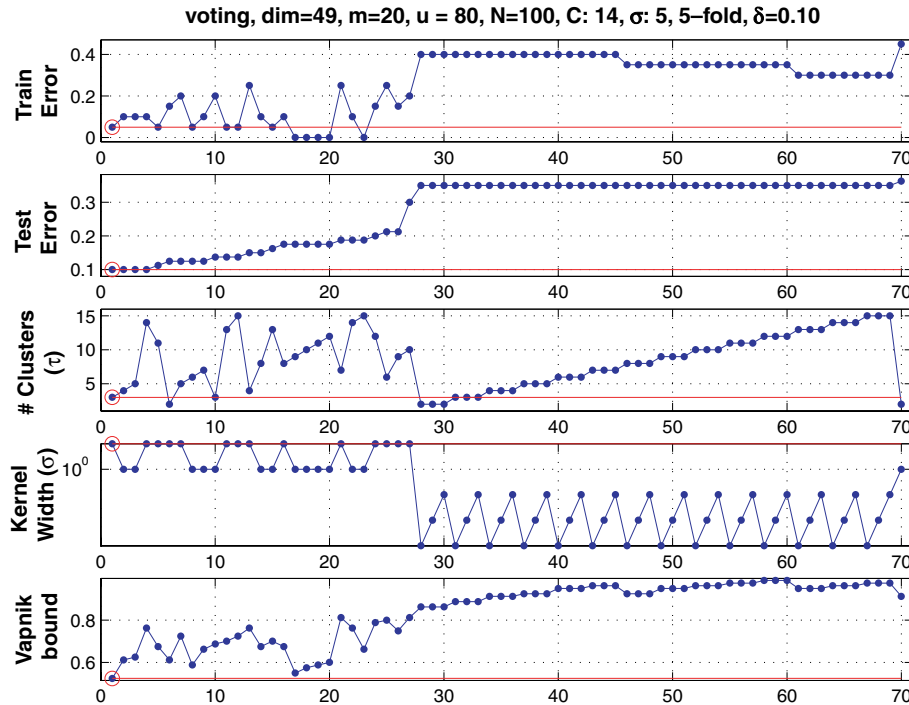


Fig. 3. Application of the CLUST algorithm on the Voting dataset with $m = 20$ and $u = 80$. The five panels show (from top to bottom), for all the $14 \times 5 = 70$ hypotheses generated, their training error, test error, number of clusters (τ) and kernel width (σ) used by the spectral clustering algorithm, and the Vapnik bound. The hypotheses are ordered (on the x -axis of each panel) according to increasing magnitude of the transduction risk (test error). The hypothesis selected by CLUST is circled.

In all the experiments described below we applied the spectral clustering algorithm with $C = 15$ and we fixed $\Theta = \{10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1\}$. That is, we clustered the data to $\tau = 2, \dots, 15$ clusters and for each τ we attempted five different metrics. These parameter sets were chosen and fixed before we applied the algorithm to any of the datasets.⁸ Overall, the CLUST algorithm attempted $5 \times 14 = 70$ unsupervised models (using the spectral clustering algorithm) in each of the applications described below. In all cases most of these unsupervised models were very unsuccessful, but for most datasets this diversity

⁸ However, we clearly used some “prior knowledge” in selecting Θ . Some experience with SVMs indicates that such a logarithmic “grid” (Θ) of possible values for the RBF “width” (σ) is often sufficient to achieve high quality results when applying (supervised) SVMs on datasets of medium sized dimension, as those we consider.

of models often included a small number of high quality models.

4.2. Example

Before we present our empirical study of the CLUST algorithm, we describe in some detail one run of the algorithm. Consider Fig. 3 (see the figure’s caption for a description of the five panels). The figure shows a particularly successful application of the algorithm, where the final selection is of the best hypothesis among all the 70 hypotheses that were generated.⁹ The selected hypothesis (circled) does not have the smallest empirical (train) error and we see that the Vapnik bound did pro-

⁹ This run corresponds to one of the folds (among the five; see below). Another fold was also particularly successful and in the other three the algorithm’s performance was less than perfect.

test the algorithm from selecting a number of hypotheses with a lower (or similar) training error, but with a significantly larger test error. Another striking fact is that most of the clusterings were extremely bad and did not yield useful hypotheses (for example, see the ones corresponding to small σ). The algorithm easily avoided all those “irrelevant” hypotheses. On the down side, we see that the Vapnik bound was not sufficiently tight to provide the final hypothesis with a useful “performance guarantee” (w.h.p.). In particular, the bound corresponding to the best hypothesis is slightly larger than 0.5. Unfortunately, this behavior was typical in all the experiments we conducted and is partly due to the very small sample size we used ($m = 20$ and $u = 80$ in the current example) and the relatively large number of hypotheses, which increases the slack of the bound (due to the use of the union bound in Eq. (7)).

5. Experiments

We tested the CLUST algorithm on standard learning problems from the UCI repository. In order to minimize the risk of over-fitting due to “dataset selection,” we committed ourselves before the start of the experiment to the entire collection of datasets used by Blum and Chawla (2001) in their paper on the transductive ‘graph-mincut’ algorithm.¹⁰ This benchmark consists of eight datasets from the UCI repository. Some properties of these sets are summarized in Table 1. Since the spectral clustering algorithm we employ assumes numerical vectors, we transformed nominal features to numerical ones. This was done in a standard way through the use of binary indicator attributes. We also normalized each attribute independently so that the dynamic range of the attributes was $[0, 1]$. No other preprocessing was used.

Using this collection of problems we conducted a rather extensive experiment, focusing on a very hard setting with extremely small sample size. For each dataset, we examined the performance

Table 1
The datasets (‘MUSH’ is the “mushroom” dataset)

Dataset	Dimension	Bias (%)	Nominal?
MUSH	21 (117)	48	✓
TAE	5 (57)	20	✓
VOTING	16 (48)	38	✓
MUSK	166	39	×
PIMA	8	35	×
IONOSPHERE	34	42	×
BUPA	6	41	×
MI	7 (17)	47	✓

The dimension is the original number of features. In cases where the original features are nominal (indicated in last column) the effective dimension, after transforming to numerical features, appears in parenthesis; ‘Bias’ is the proportion of the majority class as measured on the full sample X_{m+u} .

of the CLUST algorithm over nine $m:u$ partitions that were constructed from a full sample of $m + u = 100$ points selected uniformly at random from the original dataset. The $m:u$ partitions we examined are: 10:90, 20:80, ..., 90:10. For each $m:u$ partition we provided the algorithm with the unlabeled full sample consisting of 100 points together with a labeled training set of size m that was selected randomly and uniformly among all m -subsets. We repeated this random choice of the training set five times and all the results reported here are averages over these fivefolds.

5.1. Benchmark algorithms

As our main benchmark algorithm we selected the Spectral Graph Transducer (SGT) algorithm presented recently in (Joachims, 2003). This sophisticated algorithm appears to be one of the best transductive learning algorithms known today, as judged by the empirical study presented by Joachims. In particular, SGT clearly shows an advantage over inductive learning algorithms and in many cases it beats other well known transductive algorithms such as the transductive support vector machine (TSVM) variant of Joachims (1999), the graph-mincut algorithm of Blum and Chawla (2001) and the co-training algorithm of Blum and Mitchell (1998), which can be viewed as a special case of SGT.

The SGT algorithm has a number of parameters. As indicated by a sensitivity analysis presented in (Joachims, 2003), all but one of these parameters

¹⁰ We have not tested the CLUST algorithm on other datasets and we report below on all our results.

Table 2

Fivefold average errors (\pm standard error of the mean) for $m = 20$ and $u = 80$. The lowest error among the three algorithms (CLUST, SGT and NAIVE) appears in boldface. Also, the best among CLUST* and SGT* appears in boldface

	CLUST	SGT	NAIVE	CLUST*	SGT*
MUSH	.11 \pm .04	.12 \pm .04	.53 \pm .01	.07 \pm .00	.06 \pm .01
TAE	.17 \pm .01	.28 \pm .03	.19 \pm .01	.17 \pm .00	.18 \pm .01
VOTING	.13 \pm .01	.13 \pm .02	.37 \pm .01	.10 \pm .00	.08 \pm .01
MUSK	.45 \pm .03	.42 \pm .01	.45 \pm .05	.32 \pm .03	.31 \pm .02
PIMA	.31 \pm .03	.39 \pm .01	.36 \pm .01	.27 \pm .02	.29 \pm .01
IONO	.24 \pm .03	.33 \pm .04	.48 \pm .04	.16 \pm .04	.25 \pm .02
BUPA	.43 \pm .03	.42 \pm .02	.46 \pm .04	.38 \pm .02	.31 \pm .02
MI	.47 \pm .02	.41 \pm .01	.51 \pm .02	.35 \pm .01	.32 \pm .02

are not critical. However, one parameter (k , which determines the number of nearest neighbors of each point, for the purpose of graph construction) is critical. Joachims discovered a nice heuristic to determine a good value for this parameter, based on a solution for the optimization problem solved by the SGT algorithm. In all the experiments we conducted with SGT we assigned the value of k using the same heuristic as used by Joachims in his paper. We set the other parameters to their recommended values.¹¹ The assignment for k was chosen to be the best one (according to Joachims' heuristic) among all 100 possible choices in $\{1, 2, \dots, 100\}$ (recall that the size of our full sample is always 100).

In order to gain some perspective on the results we also tested and report on the test errors achieved by two extreme algorithms. The first is a trivial algorithm that labels the test set according to the majority class *as observed on the training set*. We denote this algorithm by NAIVE. Clearly any useful transductive learning algorithm must not consistently underperform NAIVE. The second benchmark is the best hypothesis *in hindsight* generated by algorithm CLUST. This benchmark is denoted by CLUST*. We also report on the best hypothesis *in hindsight* considered by SGT (among all 100 models) and denote it by SGT*.

5.2. Empirical observations

As mentioned, overall we conducted, $9 \times 8 = 72$ experiments, with nine different $m:u$ partitions and eight datasets. Each experiment included fivefolds.

Table 2 specifies the average errors (over the fivefolds) suffered by the various algorithms, for train/test partitions with $m = 20$ and $u = 80$. Such partitions, with very small training sets, are often of practical interest. We see that CLUST is significantly better than SGT in three datasets (TAE, PIMA and IONO). SGT is significantly better than CLUST in one dataset (MI). In most cases both CLUST and SGT are significantly better than NAIVE. In one case (MUSK) CLUST is no better than NAIVE but SGT is significantly worse than NAIVE in two cases (TAE and PIMA). In the majority of sets SGT* achieved lower test error than CLUST*.

Table 3 provides a general view on the entire set of experiments. The table specifies the average errors of the various algorithms, where the average is taken with respect to the errors obtained in all nine $m:u$ partitions. Here we see that both CLUST and SGT significantly beat each other exactly three times. We also see that both CLUST and SGT lose twice to NAIVE. A rough conclusion from this table is that no algorithm is better than the other and that each of the algorithms excels in different learning problems. We also see that the performance of SGT* is significantly better than that of CLUST*. Except for the IONO set, where SGT* (and SGT) are particularly bad, on all other sets SGT* achieves smaller error rates. Perhaps this is not so surprising because for the computation of SGT* we considered 100 different models whereas for

¹¹ Specifically, we used $D = 80$ and $C = 3200$. Also, note that we used Joachims' code to run the SGT algorithm. This code can be downloaded from Joachims' web site.

Table 3

“Total” average errors (\pm averaged standard errors of the mean). Each entry is an average of the nine fivefold averages, corresponding to the nine $m:u$ partitions. Best results appear in boldface as described in Table 2

	CLUST	SGT	NAIVE	CLUST*	SGT*
MUSH	.09 \pm .02	.09 \pm .02	.54 \pm .02	.07 \pm .02	.05 \pm .01
TAE	.16 \pm .03	.23 \pm .04	.17 \pm .02	.13 \pm .01	.12 \pm .01
VOTING	.14 \pm .02	.13 \pm .03	.39 \pm .05	.10 \pm .01	.06 \pm .01
MUSK	.44 \pm .04	.38 \pm .03	.42 \pm .05	.29 \pm .04	.24 \pm .02
PIMA	.29 \pm .03	.39 \pm .03	.35 \pm .03	.25 \pm .03	.24 \pm .02
IONO	.16 \pm .03	.33 \pm .03	.43 \pm .03	.08 \pm .02	.21 \pm .02
BUPA	.43 \pm .03	.41 \pm .03	.42 \pm .03	.33 \pm .02	.24 \pm .01
MI	.41 \pm .03	.31 \pm .02	.53 \pm .03	.25 \pm .02	.19 \pm .02

CLUST*, only 70 models. In any case, the relative success of SGT* (and of CLUST*) may indicate that there is still room for significant improvements of both algorithms!

The two scatter plots in Fig. 4 provide somewhat different general views of the experiments. Each point in the main scatter plot corresponds to one of the 72 experiments (and by itself represent an average over fivefolds). The experiments are ordered by datasets as indicated by the labels of the x -axis (and the y -axis). Within each dataset range, the nine experiments are ordered in increasing size of m (each corresponding to one of the $m:u$

partitions). Suppose that in the i th experiment (in this order) the test error of CLUST is a and the error of SGT is b . If $a < b$ then the point corresponding to this experiment is located at the coordinate $(i, i + b/a)$. Otherwise, the point is located at $(i, i - ab)$. Thus, a point below the diagonal represents an advantage for SGT and vice versa. The subplot in Fig. 4 is a standard scatter plot. Here again there is a point for each experiment and the location is simply determined by the error rates. From these plots we see again that each of the algorithms is an overall winner on 3 different datasets and there is a mixed behavior on the other two sets. On the IONO set we observe a major advantage of CLUST over SGT. There is no such major advantage of SGT on other datasets. From the standard scatter (sub)plot we can see that there is rather weak positive correlation between the behaviors of the two algorithms over the set of experiments.

The experiments described above consider settings where the total number of points $m + u = 100$ is rather small. While many applications of interest have limited sample sizes (e.g., medical applications), some others (e.g., text categorization) may have access to large repositories of unlabeled data. It is therefore interesting to ask what could be the resulting learning curves for a fixed m and growing u . Consider the explicit bound (11). The only appearance of u in the bound is within the factor $\sqrt{(m + u)/u}$. For $u \geq m$, this factor is bounded at the interval $(1, \sqrt{2}]$. Therefore, one should not expect any dramatic improvement of this bound by taking larger u . In our algorithm we use the (implicit) Vapnik bound that is only bounded above by the explicit bound (11).

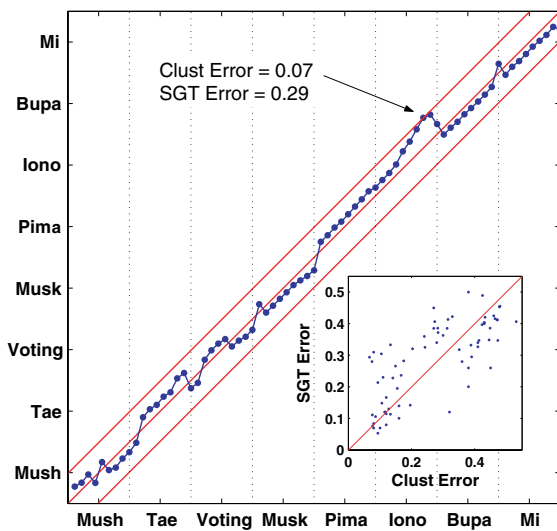


Fig. 4. Two scatter plots. Each point represent one experiment among the 72 experiments conducted. The subplot is a standard scatter plot of the algorithms’ errors. For a description of the main plot, see the text.

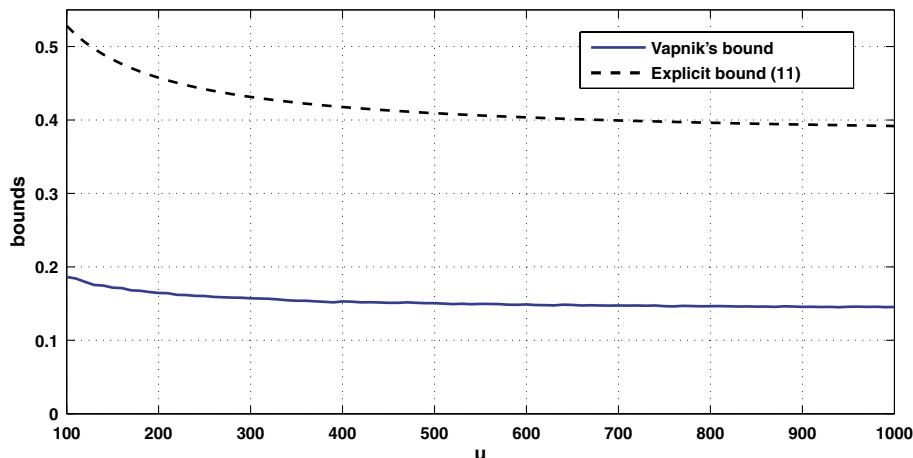


Fig. 5. The explicit bound (11) and the Vapnik bound on the test error, as a function of u . Both graphs are computed with a fixed $m = 100$, $T = 10$, $C = 10$ and $\delta = 0.1$.

However, as can be seen in the graphs in Fig. 5 both bounds are qualitatively similar and we observe that also the Vapnik bound would not benefit from larger values of u .

6. Concluding remarks

We studied a new learning scheme for transduction based on clustering. Unlike many other transductive learning heuristics the new scheme is theoretically well-founded. Our empirical study indicates that the `CLUST` algorithm achieves state-of-the-art performance, and we conclude that it is competitive with the the recent `SGT` algorithm (while being considerably simpler than `SGT`). Nevertheless, we note that the running time of `CLUST`, compared to `SGT` (using the parameters reported above), are considerably longer.

Our study concerned a straightforward implementation of the `CLUST` approach. We made no attempts to improve the basic scheme. It would be interesting to consider more sophisticated prior constructions, perhaps ones which are based on more extensive knowledge that can be derived from the full sample geometry. Also, it would be interesting to see if the `SGT` algorithm can benefit from a principled model selection as we propose here.

Transductive inference is an intriguing type of learning that is not well understood. We believe

that transduction offers new opportunities for machine learning research. In any case, while the precise relation between the inductive and transductive settings remains unclear, the interest in transduction should not be underestimated. If transduction allows for faster (and/or computationally more efficient) learning than the best possible inductive learning, then much more research attention should be devoted to transduction.

References

- Blum, A., Chawla, S., 2001. Learning from labeled and unlabeled data using graph mincuts. In: ICML-2001.
- Blum, A., Mitchell, T., 1998. Combining labeled and unlabeled data with co-training. In: COLT-1998, URL [cite-seer.nj.nec.com/blum98combining.html](http://citeseer.nj.nec.com/blum98combining.html).
- Blum, A., Langford, J., 2003. PAC-MDL Bounds. In: COLT-2003, pp. 344–357.
- Bottou, L., Cortes, C., Vapnik, V., 1994. On the effective vc dimension. Tech. rep., Neuroprose ([ftp://archive.cis.ohio-state.edu/pub/neuroprose](http://archive.cis.ohio-state.edu/pub/neuroprose)), also on <http://leon.bottou.com/publications>.
- Derbeko, P., El-Yaniv, R., Meir, R., 2004a. Error bounds for transductive learning via compression and clustering. In: Advances in Neural Information Processing Systems (NIPS) 16, pp. 1085–1092.
- Derbeko, P., El-Yaniv, R., Meir, R., 2004b. Explicit learning curves for transduction and application to clustering and compression algorithms. J. Artificial Intelligence Res. (JAIR) 22, pp. 117–142.

- Joachims, T., 1999. Transductive inference for text classification using support vector machines. In: ICML-1999.
- Joachims, T., 2003. Transductive learning via spectral graph partitioning. In: ICML-2003.
- McAllester, D., 1999. Some PAC-Bayesian theorems. *Machine Learning* 37 (3), 355–363.
- Ng, A., Jordan, M., Weiss, Y., 2002. On spectral clustering: analysis and an algorithm. In: NIPS 14, pp. 849–856.
- Seeger, M., 2002. Learning with labeled and unlabeled data. Tech. rep. Available at <http://www.dai.ed.ac.uk/homes/seeger/papers/>.
- Vapnik, V.N., 1982. *Estimation of Dependences Based on Empirical Data*. Springer Verlag, New York.
- Vapnik, V.N., 1995. *The Nature of Statistical Learning Theory*. Springer Verlag, New York.
- Vapnik, V.N., 1998. *Statistical Learning Theory*. Wiley Interscience, New York.