




ICML  **2005**
BONN - GERMANY

Proceedings of the 22nd International Conference on Machine Learning

Bonn, Germany, 7-11 August, 2005

Edited by
Luc De Raedt
Stefan Wrobel

GOOGLE EXHIBIT 1010
Google v. K.Mizra
IPR2026-00304

**Proceedings of the 22nd International
Conference on Machine Learning**

ICML 2005
F 1 5 6 9 5 1

Luc De Raedt
Stefan Wrobel

ENGR&SCI
Q
325.5
.158
2005





Proceedings of the 22nd International Conference on Machine Learning

This proceedings contains a collection of original papers presented at the Twenty-second International Conference on Machine Learning (ICML-2005) held in Bonn, Germany, August 7-11, 2005. The annual ICML conference gathers scientists interested in the latest advances in the science of machine learning. The papers in this volume have been carefully selected by the program committee. They address theoretical issues, experimental investigation as well as applications, and cover a wide range of topics from all subfields of machine learning, including support vector machines, bayesian methods, evaluation of learning algorithms, reinforcement learning, statistical relational learning, unsupervised learning, boosting, etc. The selected papers are relevant to those in academia and industry keen on practicing and extending the state of the art in machine learning.

ISBN 1-59593-180-5



ICML 2005

Twenty-Second International
Conference on Machine Learning

Edited by Luc De Raedt and Stefan Wrobel

Proceedings, Twenty-Second International Conference on Machine Learning

Edited by Luc De Raedt and Stefan Wrobel

Q
325.5
•I58
2005

per Copyright © 2005

These materials are owned by their respective copyright owners.
All rights reserved.

ISBN 1-59593-180-5

Printed in the United States of America

Table of Contents

Preface.....	xiii
ICML 2005 Organization.....	xv
Program Committee.....	xvi
External Reviewers.....	xvii
Board of the International Machine Learning Society 2005.....	xviii
Sponsors.....	xix
Invited Talks.....	xxiii
Workshop and Tutorial Summaries.....	xxix

ICML 2005 Papers

Exploration and Apprenticeship Learning in Reinforcement Learning..... <i>Pieter Abbeel, Andrew Y. Ng</i>	1
Active Learning for Hidden Markov Models: Objective Functions and Algorithms..... <i>Brigham Anderson, Andrew W. Moore</i>	9
Tempering for Bayesian C&RT..... <i>Nicos Angelopoulos, James Cussens</i>	17
Fast Condensed Nearest Neighbor Rule..... <i>Fabrizio Angiulli</i>	25
Predictive low-rank decomposition for kernel methods..... <i>Francis R. Bach, Michael I. Jordan</i>	33
Multi-Way Distributional Clustering via Pairwise Interactions..... <i>Ron Bekkerman, Ran El-Yaniv, Andrew McCallum</i>	41
Error Limiting Reductions Between Classification Tasks..... <i>Alina Beygelzimer, Varsha Dani, Tom Hayes, John Langford, Bianca Zadrozny</i>	49
Multi-Instance Tree Learning..... <i>Hendrik Blockeel, David Page, Ashwin Srinivasan</i>	57
Action Respecting Embedding..... <i>Michael Bowling, Ali Ghodsi, Dana Wilkinson</i>	65
Clustering Through Ranking On Manifolds..... <i>Markus Breitenbach, Gregory Z. Grudic</i>	73
Reducing Overfitting in Process Model Induction..... <i>Will Bridewell, Narges Bani Asani, Pat Langley, Ljupco Todorovski</i>	81
Learning to Rank using Gradient Descent..... <i>Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, Greg Hullender</i>	89
Learning Class-Discriminative Dynamic Bayesian Networks..... <i>John Burge, Terran Lane</i>	97

Recognition and Reproduction of Gestures using a Probabilistic Framework combining PCA, ICA and HMM	105
<i>Sylvain Calinon, Aude Billard</i>	
Predicting Probability Distributions for Surf Height Using an Ensemble of Mixture Density Networks.....	113
<i>Michael Carney, Padraig Cunningham, Jim Dowling, Ciaran Lee</i>	
Hedged learning: Regret minimization with learning experts	121
<i>Yu-Han Chang, Leslie Pack Kaelbling</i>	
Variational Bayesian Image Modelling	129
<i>Li Cheng, Feng Jiao, Dale Schuurmans, Shaojun Wang</i>	
Preference Learning with Gaussian Processes	137
<i>Wei Chu, Zoubin Ghahramani</i>	
New Approaches to Support Vector Ordinal Regression.....	145
<i>Wei Chu, S. Sathya Keerthi</i>	
A General Regression Technique for Learning Transductions	153
<i>Corinna Cortes, Mehryar Mohri, Jason Weston</i>	
Learning to Compete, Compromise, and Cooperate in Repeated General-Sum Games	161
<i>Jacob W. Crandall, Michael A. Goodrich</i>	
Learning as Search Optimization: Approximate Large Margin Methods for Structured Prediction	169
<i>Hal Daume III, Daniel Marcu</i>	
Multimodal Oriented Discriminant Analysis	177
<i>Fernando De la Torre, Takeo Kanade</i>	
A Practical Generalization of Fourier-based Learning	185
<i>Adam Drake, Dan Ventura</i>	
Combining Model-Based and Instance-Based Learning for First Order Regression	193
<i>Kurt Driessens, Saso Dzeroski</i>	
Reinforcement learning with Gaussian processes	201
<i>Yaakov Engel, Shie Mannor, Ron Meir</i>	
Experimental Comparison between Bagging and Monte Carlo Ensemble Classification	209
<i>Roberto Esposito, Lorenza Saitta</i>	
Supervised Clustering with Support Vector Machines.....	217
<i>Thomas Finley, Thorsten Joachims</i>	
Optimal Assignment Kernels For Attributed Molecular Graphs	225
<i>Holger Fröhlich, Jörg K. Wegner, Florian Sieker, Andreas Zell</i>	
Closed-form dual perturb and combine for tree-based models	233
<i>Pierre Geurts, Louis Wehenkel</i>	
Hierarchical Bayesian Models for Kernel Learning.....	241
<i>Mark Girolami, Simon Rogers</i>	

Online Feature Selection for Pixel Classification	249
<i>Karen Glocer, Damian Eads, James Theiler</i>	
Learning Strategies for Story Comprehension: A Reinforcement Learning Approach	257
<i>Eugene Grois, David C. Wilkins</i>	
Near-Optimal Sensor Placements in Gaussian Processes	265
<i>Carlos Guestrin, Andreas Krause, Ajit Paul Singh</i>	
Robust One-Class Clustering Using Hybrid Global and Local Search	273
<i>Gunjan Gupta, Joydeep Ghosh</i>	
Statistical and Computational Analysis of Locality Preserving Projection	281
<i>Xiaofei He, Deng Cai, Wanli Min</i>	
Intrinsic Dimensionality Estimation of Submanifolds in $\Re^d$	289
<i>Matthias Hein, Jean-Yves Audibert</i>	
Bayesian Hierarchical Clustering	297
<i>Katherine A. Heller, Zoubin Ghahramani</i>	
Online Learning over Graphs	305
<i>Mark Herbster, Massimiliano Pontil, Lisa Wainer</i>	
Adapting Two-Class Classification Methods to Many Class Problems	313
<i>Simon I. Hill, Arnaud Doucet</i>	
A Martingale Framework for Concept Change Detection in Time-Varying Data Streams	321
<i>Shen-Shyang Ho</i>	
Multi-Class protein fold detection using adaptive codes	329
<i>Eugene Ie, Jason Weston, William Stafford Noble, Christina Leslie</i>	
Learning Approximate Preconditions for Methods in Hierarchical Plans	337
<i>Okhtay Ilghami, Héctor Muñoz-Avila, Dana S. Nau, David W. Aha</i>	
Evaluating Machine Learning for Information Extraction	345
<i>Neil Ireson, Fabio Ciravegna, Mary Elaine Califf, Dayne Freitag, Nicholas Kushmerick, Alberto Lavelli</i>	
Learn to Weight Terms in Information Retrieval Using Category Information	353
<i>Rong Jin, Joyce Y. Chai, Luo Si</i>	
A Smoothed Boosting Algorithm Using Probabilistic Output Codes	361
<i>Rong Jin, Jian Zhang</i>	
Efficient discriminative learning of Bayesian network classifier via Boosted Augmented Naive Bayes	369
<i>Yushi Jing, Vladimir Pavlovic, James M. Rehg</i>	
A Support Vector Method for Multivariate Performance Measures	377
<i>Thorsten Joachims</i>	
Error Bounds for Correlation Clustering	385
<i>Thorsten Joachims, John Hopcroft</i>	

Interactive Learning of Mappings from Visual Percepts to Actions	393
<i>Sébastien Jodogne, Justus H. Piater</i>	
A Causal Approach to Hierarchical Decomposition of Factored MDPs.....	401
<i>Anders Jonsson, Andrew Barto</i>	
A Comparison of Tight Generalization Error Bounds	409
<i>Matti Kääriäinen, John Langford</i>	
Generalized LARS as an Effective Feature Selection Tool for Text Classification With SVMs	417
<i>S. Sathiya Keerthi</i>	
Ensembles of Biased Classifiers.....	425
<i>Rinat Khossainov, Andreas Heß, Nicholas Kushmerick</i>	
Computational Aspects of Bayesian Partition Models	433
<i>Mikko Koivisto, Kismat Sood</i>	
Learning the Structure of Markov Logic Networks.....	441
<i>Stanley Kok, Pedro Domingos</i>	
Using Additive Expert Ensembles to Cope with Concept Drift	449
<i>Jeremy Kolter, Marcus Maloof</i>	
Semi-supervised Graph Clustering: A Kernel Approach	457
<i>Brian Kulis, Sugato Basu, Inderjit S. Dhillon, Raymond J. Mooney</i>	
A Brain Computer Interface with Online Feedback based on Magnetoencephalography.....	465
<i>Thomas Navin Lal, Michael Schröder, N. Jeremy Hill, Hubert Preissl, Thilo Hinterberger, Jürgen Mellinger, Martin Bogdan, Wolfgang Rosenstiel, Thomas Hofmann, Niels Birbaumer, Bernhard Schölkopf</i>	
Relating Reinforcement Learning Performance to Classification Performance	473
<i>John Langford, Bianca Zadrozny</i>	
PAC-Bayes Risk Bounds for Sample-Compressed Gibbs Classifiers	481
<i>François Laviolette, Mario Marchand</i>	
Heteroscedastic Gaussian Process Regression.....	489
<i>Quoc V. Le, Alex J. Smola, Stéphane Canu</i>	
Predicting Relative Performance of Classifiers from Samples.....	497
<i>Rui Leite, Pavel Brazdil</i>	
Logistic Regression with an Auxiliary Data Source	505
<i>Xuejun Liao, Ya Xue, Lawrence Carin</i>	
Predicting Protein Folds with Structural Repeats Using a Chain Graph Model.....	513
<i>Yan Liu, Eric P. Xing, Jaime Carbonell</i>	
Unsupervised Evidence Integration	521
<i>Philip M. Long, Vinay Varadan, Sarah Gilman, Mark Treshock, Rocco A. Servedio</i>	
Naive Bayes Models for Probability Estimation.....	529
<i>Daniel Lowd, Pedro Domingos</i>	

ROC Confidence Bands : An Empirical Evaluation	537
<i>Sofus A. Macskassy, Foster Provost, Saharon Rosset</i>	
Modeling Word Burstiness Using the Dirichlet Distribution	545
<i>Rasmus E. Madsen, David Kauchak, Charles Elkan</i>	
Proto-Value Functions: Developmental Reinforcement Learning	553
<i>Sridhar Mahadevan</i>	
The cross entropy method for classification	561
<i>Shie Mannor, Dori Peleg, Reuven Y. Rubinstein</i>	
Bounded Real-Time Dynamic Programming: RTDP with monotone upper bounds and performance guarantees.....	569
<i>H. Brendan McMahan, Maxim Likhachev, Geoffrey J. Gordon</i>	
Comparing Clusterings - An Axiomatic View.....	577
<i>Marina Meila</i>	
Weighted Decomposition Kernels	585
<i>Sauro Menchetti, Fabrizio Costa, Paolo Frasconi</i>	
High Speed Obstacle Avoidance using Monocular Vision and Reinforcement learning	593
<i>Jeff Michels, Ashutosh Saxena, Andrew Y. Ng</i>	
Dynamic Preferences in Multi-Criteria Reinforcement Learning	601
<i>Sriaram Natarajan, Prasad Tadepalli</i>	
Learning First-Order Probabilistic Models with Combining Rules	609
<i>Sriaram Natarajan, Prasad Tadepalli, Eric Altendorf, Thomas G. Dietterich, Alan Fern, Angelo Restificar</i>	
An Efficient Method for Simplifying Support Vector Machines	617
<i>DucDung Nguyen, Tu Bao Ho</i>	
Predicting Good Probabilities With Supervised Learning	625
<i>Alexandru Niculescu-Mizil, Rich Caruana</i>	
Recycling Data for Multi-Agent Learning.....	633
<i>Santiago Ontañón, Enric Plaza</i>	
A Graphical Model for Chord Progressions Embedded in a Psychoacoustic Space	641
<i>Jean-François Paiement, Douglas Eck, Samy Bengio, David Barber</i>	
Q-Learning of Sequential Attention for Visual Object Recognition from Informative Local Descriptors	649
<i>Lucas Paletta, Gerald Fritz, Christin Seifert</i>	
Discriminative versus Generative Parameter and Structure Learning of Bayesian Network Classifiers.....	657
<i>Franz Pernkopf, Jeff Bilmes</i>	
Optimizing Abstaining Classifiers using ROC Analysis	665
<i>Tadeusz Pietraszek</i>	
Independent Subspace Analysis Using Geodesic Spanning Trees	673
<i>Barnabás Póczos, András Lörincz</i>	

A Model for Handling Approximate, Noisy or Incomplete Labeling in Text Classification.....	681
<i>Ganesh Ramakrishnan, Krishna Prasad Chitrapura, Raghu Krishnapuram, Pushpak Bhattacharyya</i>	
Healing the Relevance Vector Machine through Augmentation.....	689
<i>Carl Edward Rasmussen, Joaquin Quiñero-Candela</i>	
Supervised versus Multiple Instance Learning: An Empirical Comparison	697
<i>Soumya Ray, Mark Craven</i>	
Generalized Skewing for Functions with Continuous and Nominal Attributes.....	705
<i>Soumya Ray, David Page</i>	
Fast Maximum Margin Matrix Factorization for Collaborative Prediction	713
<i>Jason D. M. Rennie, Nati Srebro</i>	
Coarticulation: An Approach for Generating Concurrent Plans in Markov Decision Processes.....	721
<i>Khashayar Rohanimanesh, Sridhar Mahadevan</i>	
Why Skewing Works: Learning Difficult Boolean Functions with Greedy Tree Learners.....	729
<i>Bernard Rosell, Lisa Hellerstein, Soumya Ray, David Page</i>	
Integer Linear Programming Inference for Conditional Random Fields	737
<i>Dan Roth, Wen-Tau Yih</i>	
Learning Hierarchical Multi-Category Text Classification Models	745
<i>Juho Rousu, Craig Saunders, Sandor Szedmak, John Shawe-Taylor</i>	
Expectation Maximization Algorithms for Conditional Likelihoods	753
<i>Jarkko Salojärvi, Kai Puolamäki, Samuel Kaski</i>	
Estimating and computing density based distance metrics.....	761
<i>Sajama, Alon Orlitsky</i>	
Supervised dimensionality reduction using mixture models	769
<i>Sajama, Alon Orlitsky</i>	
Object Correspondence as a Machine Learning Problem	777
<i>Bernhard Schölkopf, Florian Steinke, Volker Blanz</i>	
Analysis and Extension of Spectral Methods for Nonlinear Dimensionality Reduction	785
<i>Fei Sha, Lawrence K. Saul</i>	
Non-Negative Tensor Factorization with Applications to Statistics and Computer Vision	793
<i>Amnon Shashua, Tamir Hazan</i>	
Fast Inference and Learning in Large-State-Space HMMs.....	801
<i>Sajid M. Siddiqi, Andrew W. Moore</i>	
New D-Separation Identification Results for Learning Continuous Latent Variable Models.....	809
<i>Ricardo Silva, Richard Scheines</i>	
Identifying Useful Subgoals in Reinforcement Learning by Local Graph Partitioning	817
<i>Özgür Simsek, Alicia P. Wolfe, Andrew G. Barto</i>	

Beyond the Point Cloud: from Transductive to Semi-supervised Learning	825
<i>Vikas Sindhwani, Partha Niyogi, Mikhail Belkin</i>	
Active Learning for Sampling in Time-Series Experiments With Application to Gene Expression Analysis....	833
<i>Rohit Singh, Nathan Palmer, David Gifford, Bonnie Berger, Ziv Bar-Joseph</i>	
Compact approximations to Bayesian predictive distributions	841
<i>Edward Snelson, Zoubin Ghahramani</i>	
Large Scale Genomic Sequence SVM Classifiers	849
<i>Sören Sonnenburg, Gunnar Rätsch, Bernhard Schölkopf</i>	
A Theoretical Analysis of Model-Based Interval Estimation	857
<i>Alexander L. Strehl, Michael L. Littman</i>	
Explanation-Augmented SVM: an Approach to Incorporating Domain Knowledge into SVM Learning	865
<i>Qiang Sun, Gerald DeJong</i>	
Unifying the Error-Correcting and Output-Code AdaBoost within the Margin Framework	873
<i>Yijun Sun, Sinisa Todorovic, Jian Li, Dapeng Wu</i>	
Finite Time Bounds for Sampling Based Fitted Value Iteration	881
<i>Csaba Szepesvári, Rémi Munos</i>	
TD(λ) Networks: Temporal-Difference Networks with Eligibility Traces	889
<i>Brian Tanner, Richard S. Sutton</i>	
Learning Structured Prediction Models: A Large Margin Approach	897
<i>Ben Taskar, Vassil Chatalbashev, Daphne Koller, Carlos Guestrin</i>	
Learning Discontinuities with Products-of-Sigmoids for Switching between Local Models	905
<i>Marc Toussaint, Sethu Vijayakumar</i>	
Core Vector Regression for Very Large Regression Problems	913
<i>Ivor W. Tsang, James T. Kwok, Kimo T. Lai</i>	
Propagating Distributions on a Hypergraph by Dual Information Regularization	921
<i>Koji Tsuda</i>	
Hierarchical Dirichlet Model for Document Classification	929
<i>Sriharsha Veeramachaneni, Diego Sona, Paolo Avesani</i>	
Implicit Surface Modelling as an Eigenvalue Problem	937
<i>Christian Walder, Olivier Chapelle, Bernhard Schölkopf</i>	
New Kernels for Protein Structural Motif Discovery and Function Classification	945
<i>Chang Wang, Stephen D. Scott</i>	
Exploiting Syntactic, Semantic and Lexical Regularities in Language Modeling via Directed Markov Random Fields	953
<i>Shaojun Wang, Shaomin Wang, Russell Greiner, Dale Schuurmans, Li Cheng</i>	
Bayesian Sparse Sampling for On-line Reward Optimization	961
<i>Tao Wang, Daniel J. Lizotte, Michael Bowling, Dale Schuurmans</i>	

Learning Predictive Representations from a History	969
<i>Eric Wiewiora</i>	
Incomplete-Data Classification using Logistic Regression	977
<i>David Williams, Xuejun Liao, Ya Xue, Lawrence Carin</i>	
Learning Predictive State Representations in Dynamical Systems Without Reset.....	985
<i>Britton Wolfe, Michael R. James, Satinder Singh</i>	
Linear Asymmetric Classifier for Cascade Detectors	993
<i>Jianxin Wu, Matthew D. Mullin, James M. Rehg</i>	
Building Sparse Large Margin Classifiers.....	1001
<i>Mingrui Wu, Bernhard Schölkopf, Gökhan Bakir</i>	
Dirichlet Enhanced Relational Learning	1009
<i>Zhao Xu, Volker Tresp, Kai Yu, Shipeng Yu, Hans-Peter Kriegel</i>	
Learning Gaussian Processes from Multiple Tasks.....	1017
<i>Kai Yu, Volker Tresp, Anton Schwaighofer</i>	
Augmenting Naive Bayes for Ranking	1025
<i>Harry Zhang, Liangxiao Jiang, Jiang Su</i>	
A New Mallows Distance Based Metric For Comparing Clusterings	1033
<i>Ding Zhou, Jia Li, Hongyuan Zha</i>	
Learning from Labeled and Unlabeled Data on a Directed Graph	1041
<i>Dengyong Zhou, Jiayuan Huang, Bernhard Schölkopf</i>	
2D Conditional Random Fields for Web Information Extraction	1049
<i>Jun Zhu, Zaiqing Nie, Ji-Rong Wen, Bo Zhang, Wei-Ying Ma</i>	
Harmonic mixtures: combining mixture models and graph-based methods for inductive and scalable semi-supervised learning.....	1057
<i>Xiaojin Zhu, John Lafferty</i>	
Large Margin Non-Linear Embedding	1065
<i>Alexander Zien, Joaquin Quiñero-Candela</i>	

The cross entropy method for classification

Shie Mannor

Dept. of Electrical and Computer Engineering, McGill University, Montreal, Canada

SHIE@ECE.MCGILL.CA

Dori Peleg

Dept. of Electrical Engineering, Technion Institute of Technology, Haifa 32000, Israel

DORIP@TECHUNIX.TECHNION.AC.IL

Reuven Rubinstein

Faculty of Industrial Engineering and Management, Technion Institute of Technology, Haifa 32000, Israel

IERRR01@IE.TECHNION.AC.IL

Abstract

We consider support vector machines for binary classification. As opposed to most approaches we use the number of support vectors (the " L_0 norm") as a regularizing term instead of the L_1 or L_2 norms. In order to solve the optimization problem we use the cross entropy method to search over the possible sets of support vectors. The algorithm consists of solving a sequence of efficient linear programs. We report experiments where our method produces generalization errors that are similar to support vector machines, while using a considerably smaller number of support vectors.

1. Introduction

Many classification algorithms are focused on finding a classifier which minimizes the sum of the training error (or some related cost functional) and a complexity penalty. The intuition is that a classifier which is too complex might be tailored to the data and result in overfitting. The exact choice of a complexity penalty is often motivated by generalization abilities of the resulting classifier. However, this choice is also restricted by guaranteeing that the optimization problem is tractable. In Support Vector Machines (SVMs) one typically considers an L_2 complexity term (e.g., Cristianini & Shawe-Taylor, 2000; Vapnik, 1998), which can be justified in terms of minimizing the margin. This yields a quadratic optimization problem which can be solved remarkably efficiently up to a few thousand samples (see Platt, 1998; Keerthi et al., 2001). The main problem with using the L_2

complexity term is that it does not encourage sparse solutions. That is, the resulting classifier involves a significant fraction of the training points. Moreover, it was observed (e.g., Downs et al., 2001) that the solution produced by SVMs can even be made more sparse without changing the resulting classifier. Other kernel-based methods (e.g., Schölkopf & Smola, 2002) attempt to guarantee that the obtained classifier is sparse using different ideas. Most notably, the ν -SVM formulation (Chang & Lin, 2002) attempts to control the fraction of Support Vectors (SVs) by introducing a parameter used to tradeoff the number of SVs with the size of errors. Still, regularization is done using L_2 norm. In the context of classification, using the L_1 norm for regularization was used in Bradley and Mangasarian (1998) and Zhu et al. (2003).

A classifier with a small number of SVs is termed "sparse". Such classifiers enjoy several merits. First, sparse classifiers typically have better generalization error (e.g., Floyd & Warmuth, 1995 and more recently within the luckiness framework of Herbrich and Williamson (2002) and via compression bounds of Herbrich, 2002). Second, such classifiers are considered more robust to noise in the data (in the context of L_1 classification see Zhu et al., 2003, for a discussion). Third, the resulting classifier can be computed more quickly on test data, as it involves less terms. Improving the complexity of classification (over the test data) was considered Burges (1996); Burges and Schölkopf (1997); Downs et al. (2001), who considered simplifying a classifier obtained from solving an SVM. Typically, some SVs are discarded with little (or no, as in Downs et al., 2001) effect on the resulting classifier. Our approach is different. Instead of trying to simplify a solution obtained from solving an SVM, we directly solve an optimization problem aimed at producing a sparse classifier.

In this paper we consider the classical soft SVM prob-

Appearing in *Proceedings of the 22nd International Conference on Machine Learning*, Bonn, Germany, 2005. Copyright 2005 by the author(s)/owner(s).

lem, but regularize using the number of SVs (the so-called L_0 norm). As a result we obtain a discontinuous and non-convex optimization problem. We formulate the problem as a search problem where one looks for the set of SVs. We apply the Cross Entropy (CE) method for efficiently searching for the set of SVs. The CE method provides a sequence of candidate sets of SVs. For each such candidate set, finding the weights of the SVs amounts to solving a reasonably sized linear program provided that the size of the candidate set is not too large. The CE method is a state-of-the-art heuristic method for solving combinatorial optimization problems and continuous multi-extremal problems (Rubinstein & Kroese, 2004). The basic idea behind the CE method is to transform a combinatorial optimization problem into a stochastic optimization problem, and then solve the stochastic optimization method adaptively. In many problems of practical interest, such as the travelling salesperson problem and the Max-cut problem, the CE method was observed to quickly reach a solution close to optimal (for a discussion and examples, see Rubinstein & Kroese, 2004). We note that it is not claimed that the CE method is capable of solving every instance of some NP-hard problem efficiently, but rather that it typically performs well. For a convergence proof of the CE method see Rubinstein and Kroese. (2004). We will use the CE-method to search over possible subsets of SVs.

We are concerned with the classical classification setup, i.e., given a training set of pairs $\{\mathbf{x}^{(i)}, y^{(i)}\}_{i=1}^n$, where $\mathbf{x}^{(i)} \in \mathbb{R}^d$ are input patterns and $y^{(i)} \in \{-1, 1\}$ are the corresponding labels, find a classifier $f(\mathbf{x})$ which will classify new patterns with as few errors as possible. We are also given a kernel $K: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$. For a vector $\alpha \in \mathbb{R}^n$ let $\|\alpha\|_0 = \sum_{i=1}^n 1_{\{\alpha_i \neq 0\}}$ denote the " L_0 norm" that counts how many elements are not zero. The objective of this paper is to solve the problem

$$\begin{aligned} & \text{minimize} \quad \|\alpha\|_0 + C \sum_{i=1}^n \phi(y_i f(\mathbf{x}_i)) \\ & \text{subject to} \quad 0 \leq \alpha \leq C, \end{aligned} \quad (1)$$

with variables $\alpha \in \mathbb{R}^n$ and $b \in \mathbb{R}$, and where the hinge loss $\phi: \mathbb{R} \rightarrow \mathbb{R}_+$ is defined as $\phi(z) = 1 - z$ if $z \leq 1$, and zero otherwise; C is a hyperparameter; and the classifier f is defined by:

$$f(\mathbf{x}) = \sum_{i=1}^n \alpha_i y_i K(\mathbf{x}_i, \mathbf{x}) + b.$$

The remainder of the paper is organized as follows. In Section 2 we review the CE method and provide the

essential details needed here. In Section 3 we describe the classification algorithm based on the CE method. In Section 4 we report experimental results with an artificial data set and with several common data sets. In Section 5 we discuss the results and offer some conclusions.

2. The Cross Entropy Method

In this section we recapitulate the principles of the CE method. The reader is referred to Rubinstein and Kroese. (2004); de-Boer et al. (2005) and references therein for context, extensions, and applications. As mentioned, the basic idea behind the CE method is to transform the original (combinatorial) optimization problem to an associated stochastic optimization problem, and then to tackle the stochastic problem efficiently by an adaptive sampling algorithm. By doing so one constructs a random sequence of solutions which converges (probabilistically) to the optimal or at least a reasonable solution. Once the associated stochastic optimization is defined, the CE method alternates the following two phases:

1. Generation of a sample of random data (trajectories, vectors, etc.) according to a specified random mechanism.
2. Update of the parameters of the random mechanism, on the basis of the data, in order to produce a "better" sample in the next iteration.

We now present a proto-typical version of the CE method. Suppose we wish to minimize some cost function $S(\mathbf{z})$ over all $\mathbf{z}$ in some set $\mathcal{Z}$. Let us denote the minimum by γ^* , thus

$$\gamma^* = \min_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{z}). \quad (2)$$

We randomize our deterministic problem by defining a family of auxiliary pdfs $\{f(\cdot; \mathbf{v}), \mathbf{v} \in \mathcal{V}\}$ on $\mathcal{Z}$ and we associate with Eq. (2) the following estimation problem for a given scalar γ :

$$\mathbb{P}_{\mathbf{u}}(S(\mathbf{Z}) \leq \gamma) = \mathbb{E}_{\mathbf{u}} [I_{\{S(\mathbf{Z}) \leq \gamma\}}],$$

where $\mathbf{u}$ is some known (initial) parameter. We consider the event "cost is low" to be the rare event $I_{\{S(\mathbf{Z}) \leq \gamma\}}$ of interest. To estimate this event, the CE method generates a sequence of tuples $\{(\hat{\gamma}_t, \hat{\mathbf{v}}_t)\}$, that converge (with high probability) to a small neighborhood of the optimal tuple $(\gamma^*, \mathbf{v}^*)$, where γ^* is the solution of the program (2) and $\mathbf{v}^*$ is a pdf that emphasizes values in $\mathcal{Z}$ with a low cost. We note that

typically the optimal $\mathbf{v}^*$ is degenerated as it concentrates on the optimal solution (or a small neighborhood thereof). Let ρ denote the fraction of the best samples used to find the threshold γ . The process that is based on sampled data is termed the *stochastic counterpart* since it is based on stochastic samples of data. The number of samples in each stage of the stochastic counterpart is denoted by N , which is a predefined parameter. The following is a standard CE procedure for minimization borrowed from de-Boer et al. (2005). We initialize by setting $\hat{\mathbf{v}}_0 = \mathbf{v}_0 = \mathbf{u}$ and choose a not very small ρ , say $10^{-2} \leq \rho$. We then proceed iteratively as follows:

1. **Adaptive updating of γ_t .** For a fixed $\mathbf{v}_{t-1}$, let γ_t be a $\rho 100\%$ -percentile of $S(\mathbf{Z})$ under $\mathbf{v}_{t-1}$. That is, γ_t satisfies $\mathbb{P}_{\mathbf{v}_{t-1}}(S(\mathbf{Z}) \leq \gamma_t) \geq \rho$ and $\mathbb{P}_{\mathbf{v}_{t-1}}(S(\mathbf{Z}) \geq \gamma_t) \geq 1 - \rho$ where $\mathbf{Z} \sim f(\cdot; \mathbf{v}_{t-1})$. A simple estimator $\hat{\gamma}_t$ of γ_t can be obtained by taking a random sample $\mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(N)}$ from the pdf $f(\cdot; \mathbf{v}_{t-1})$, calculating the performances $S(\mathbf{Z}^{(\ell)})$ for all ℓ , ordering them from smallest to biggest as $S_{(1)} \leq \dots \leq S_{(N)}$ and finally evaluating the $\rho 100\%$ sample percentile as $\hat{\gamma}_t = S_{(\lceil \rho N \rceil)}$.

2. **Adaptive updating of $\mathbf{v}_t$.** For a fixed γ_t and $\mathbf{v}_{t-1}$, derive $\mathbf{v}_t$ from the solution of the program

$$\max_{\mathbf{v}} D(\mathbf{v}) = \max_{\mathbf{v}} \mathbb{E}_{\mathbf{v}_{t-1}} I_{\{S(\mathbf{Z}) \leq \gamma_t\}} \log f(\mathbf{Z}; \mathbf{v}). \quad (3)$$

The stochastic counterpart of (3) is as follows: for fixed $\hat{\gamma}_t$ and $\hat{\mathbf{v}}_{t-1}$, derive $\hat{\mathbf{v}}_t$ from the following program:

$$\max_{\mathbf{v}} \hat{D}(\mathbf{v}) = \max_{\mathbf{v}} \frac{1}{N} \sum_{\ell=1}^N I_{\{S(\mathbf{Z}^{(\ell)}) \leq \hat{\gamma}_t\}} \log f(\mathbf{Z}^{(\ell)}; \mathbf{v}). \quad (4)$$

We note that if f belongs to the Natural Exponential Family (e.g., Gaussian, Bernoulli), then Eq. (4) has a closed form solution (see de-Boer et al., 2005). The CE optimization algorithm is summarized in Algorithm 2.1. In this paper we will assume that f belongs to a Bernoulli family. In our case $\mathcal{Z} = \{0, 1\}^n$ and $\mathbf{v}$ is an n dimensional vector of numbers between 0 and 1. The update formula of the k th element in $\mathbf{v}$ (Eq. (4)) in this case simply becomes:

$$\hat{v}_t(k) = \frac{\sum_{\ell=1}^N I_{\{S(\mathbf{Z}^{(\ell)}) \leq \hat{\gamma}_t\}} I_{\{Z_k^{(\ell)}=1\}}}{\sum_{\ell=1}^N I_{\{S(\mathbf{Z}^{(\ell)}) \leq \hat{\gamma}_t\}}}.$$

This formula has the interpretation that it counts how many times a value of 1 (in $I_{\{Z_k^{(\ell)}=1\}}$) led to a significant result (matches with the indicator $I_{\{S(\mathbf{Z}^{(\ell)}) \leq \hat{\gamma}_t\}}$),

how many times a value of 0 led to a significant result, and normalize the value of the parameter accordingly. Instead of the updating the parameter vector $\mathbf{v}$ directly via the solution of Eq. (4) we use the following *smoothed* version

$$\hat{\mathbf{v}}_t = \beta \hat{\mathbf{v}}_t + (1 - \beta) \hat{\mathbf{v}}_{t-1}, \quad (5)$$

where $\hat{\mathbf{v}}_t$ is the parameter vector obtained from the solution of (4), and β is a smoothing parameter, with $0.7 < \beta < 1$. The reason for using the smoothed (5) instead of the original updating rule is to smooth out the values of $\hat{\mathbf{v}}_t$, and to reduce the probability that some component $\hat{v}_{t,i}$ of $\hat{\mathbf{v}}_t$ will be zeros or unities at an early stage, and the algorithm will get stuck in a local maxima. Note that for $0 < \beta < 1$ we always have that $\hat{v}_{t,i} > 0$, while for $\beta = 1$ one might have (even at the first iterations) that either $\hat{v}_{t,i} = 0$ or $\hat{v}_{t,i} = 1$ for some indices i . As a result, the algorithm may converge to a wrong solution.

Algorithm 2.1 The CE Method for Stochastic Optimization

1. Choose some $\hat{\mathbf{v}}_0$. Set $t = 1$ (level counter).
2. Generate a sample $\mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(N)}$ from the density $f(\cdot; \mathbf{v}_{t-1})$ and compute the sample $\rho 100\%$ -percentile $\hat{\gamma}_t$ of the sample scores.
3. Use the same sample $\mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(N)}$ and solve the stochastic program (4). Denote the solution by $\hat{\mathbf{v}}_t$.
4. Apply (5) to smooth out the vector $\hat{\mathbf{v}}_t$.
5. If for some $t \geq d$, say $d = 3$, $\hat{\gamma}_t = \hat{\gamma}_{t-1} = \dots = \hat{\gamma}_{t-d}$ then stop; otherwise set $t = t + 1$ and reiterate from step 2.

Figure 1. A proto-typical CE algorithm.

It is found empirically (Rubinstein & Kroese, 2004) that the CE method is robust with respect to the choice of its parameters N, ρ and β , as long as ρ is not too small, $\beta < 1$, and N is large enough. Typically those parameters satisfy that $0.01 \leq \rho \leq 0.1$, $0.5 \leq \beta \leq 0.9$, and $N \geq 3n$, where n is the number of parameters. Improvements of Algorithm 2.1 include the Fully Adaptive CE (FACE) variant, where the parameters N and ρ are updated online. See de-Boer et al. (2005) for more details.

3. The CE classifier algorithm

Problem (1) is a difficult non-convex discontinuous optimization problem due to the 'zero norm' of α . In

order to circumvent the discontinuity we introduce a vector of n binary variables $\sigma \in \{0, 1\}^n$ and define $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$. We replace α with $\Sigma\hat{\alpha}$, i.e., the vector σ multiplies the n dimensional vector $\hat{\alpha}$ componentwise. Thus (1) is transformed to the problem

$$\begin{aligned} & \text{minimize} && \|\Sigma\hat{\alpha}\|_0 + C\mathbf{1}^T\xi \\ & \text{subject to} && Y(KY\Sigma\hat{\alpha} + b\mathbf{1}) \geq 1 - \xi \\ & && \xi \geq 0 \\ & && 0 \leq \Sigma\hat{\alpha} \leq C \\ & && \sigma \in \{0, 1\}^n, \end{aligned} \quad (6)$$

with variables $\hat{\alpha}, \sigma, b$ and $\xi \in \mathbb{R}^n$, where $\mathbf{1}$ is a vector of ones, K is the kernel matrix satisfying $K_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)$, and $Y = \text{diag}(y_1, \dots, y_n)$. Since any component of α can be written as $1 \cdot \hat{\alpha}$ if $\alpha > 0$ or $0 \cdot \hat{\alpha}$ if $\alpha = 0$ then problems (1) and (6) are equivalent.

The zero norm counts the number of non-zero elements in a vector. Therefore $\|\Sigma\hat{\alpha}\|_0 = \sum_{i=1}^n \sigma_i$ equals the number of SVs (in practice $\hat{\alpha}$ is never exactly zero). Denote by $\tilde{\alpha}$ the vector of non-zero elements of vector $\Sigma\hat{\alpha}$. Denote $\tilde{K}$ as the kernel matrix K without the columns indicated by the zeros of σ and $\tilde{Y}$ as the diagonal matrix of labels without the labels indicated by the zeros of σ . For a given vector σ , the problem (6) is reduced to:

$$\begin{aligned} & \text{minimize} && \mathbf{1}^T\xi \\ & \text{subject to} && Y(\tilde{K}\tilde{\alpha} + b\mathbf{1}) \geq 1 - \xi \\ & && \xi \geq 0 \\ & && 0 \leq \tilde{\alpha} \leq C, \end{aligned} \quad (7)$$

with variables $\tilde{\alpha}, \xi, b$. Notice that the dimension of $\tilde{\alpha}$ is equal to the number of SVs, i.e., $\sum_{i=1}^n \sigma_i$. This is typically a much smaller number than the number of training patterns n . Problem (7) is a Linear Programming (LP) problem which can be solved with highly efficient optimization algorithms. The number of variables is $(\sum_{i=1}^n \sigma_i) + n + 1$ and the number of constraints is $2n + 2(\sum_{i=1}^n \sigma_i)$. Therefore the algorithm is affected by the number of features only when the kernel matrix is initially calculated. We propose to solve problem (1) by using the CE algorithm for the combinatorial optimization problem of selecting the set of SVs. We use Bernoulli distribution for $f(\cdot; \mathbf{v})$, with the variable σ playing the role of $\mathbf{z}$ in the CE method. The objective function is

$$S(\sigma) = \mathbf{1}^T\sigma + C\mathbf{1}^T\xi,$$

where $\mathbf{1}^T\xi$ was calculated by solving (7). We chose $\rho = 0.03, \beta = 0.7, d = 3, N = 5n$, where n is the number of training patterns. The values of parameters ρ, β, d are standard¹ and the value of N was determined via trial

¹We found that the CE algorithm is not sensitive to

and error to be the lowest number which consistently produces the same results as with a larger N .

Consider the following optimization problem, known as the "one norm SVM":

$$\begin{aligned} & \text{minimize} && \mathbf{1}^T\alpha + C\mathbf{1}^T\xi \\ & \text{subject to} && Y(KY\alpha + b\mathbf{1}) \geq 1 - \xi \\ & && \xi \geq 0 \\ & && 0 \leq \alpha \leq C, \end{aligned} \quad (8)$$

Problem (8) is an upper bound on the "zero norm" problem (1). This approach was introduced by Smola et al., (2000, pp. 141-142) and produces less sparse solutions than those generated by the "zero norm" problem. From a computational point of view, this problem can be solved efficiently. In order to reduce the complexity of our algorithm, we first run a one norm SVM, and then only consider the SVs that were computed by the one norm as potential SVs. It is reasonable to assume that the majority of SVs of the zero norm problem are SVs of the one-norm problem because these SVs are representative in some sense. A similar approach was adopted in Burges and Schölkopf (1997). Thus we suggest to use the solution of problem (8) as an initial solution for problem (6).

Three modifications were made to the standard CE algorithm:

1. The solution $\tilde{\alpha}$ of (7) may have nearly zero elements. To remove redundancy and to make the term of number of SVs in the target function $S(\sigma)$ slightly less rigid, we decided that if the value of $\tilde{\alpha}_j < 10^{-6}$ then the corresponding σ_j was set to zero.
2. In a CE iteration frequently identical σ vectors are generated (especially near convergence). Given the same σ , the solution of the convex problem (7) is identical. Thus computing (7) again is superfluous. To avoid this redundancy we used a cache for the results of (7). We note that caching can be done in an extremely simple way. Each σ (a vector of ones and zeros) is assigned a binary number corresponding to its value. The σ 's are sorted in ascending order. For each σ problem (7) is computed only if its value does not appear in the (sorted) list of already computed values².
3. Before the CE algorithm is executed the non SVs of the solution of Problem (8) are excluded from

these parameters as long as they are chosen from the ranges specified at the end of Section 2.

²After the target function is calculated without redundancy, the σ vectors are sorted according to the value of their target function as explained in (2.1).

consideration as potential SVs. All the training points (including the "removed" SVs) are still taken into account when solving (7). Note that this steps guarantees sparser solution than those generated by Smola et al. (2000).

In summary, the training phase of the CE classifier algorithm consists of solving typically several thousand LP problems. We note that the number of variables used by each LP is independent of the number of features that only affects the computation of the kernel matrix, which is performed once.

4. Experiments

The CE classifier algorithm was compared to the SVM and the one-norm SVM (Smola et al., 2000) algorithms on one synthetic dataset and five real world two class problem datasets taken from the UCI repository (see Table 1).

4.1. Synthetic data

The synthetic dataset was generated in the following manner. The probability of $y = 1$ or -1 was equal. Features x_1, x_2 were drawn as $x|y \sim \mathcal{N}(y\mathbf{1}, \Lambda)$, with the covariance matrix $\Lambda = VDV^T$ and $\mathbf{1}$ a vector of ones. The matrix D was diagonal with the values 1, 2 and V a randomly generated unitary matrix. The number of training patterns was 300 and the test set consisted of 1000 patterns.

The linear kernel was used for all algorithms. The value of hyperparameter C for each algorithm was set as the minimizer of the errors on the test set. The SVM algorithm achieved 24.5% with 56.7% of the training patterns. The one-norm SVM algorithm achieved 24.5% with 0.667% of the training patterns. The CE algorithm achieved 24.3% with 0.333% of the training patterns. Figures 2-4 illustrate the sparsity of the one-norm SVM and CE algorithm in comparison to the SVM algorithm. The class of each training pattern is marked by a circle or a square. If the shape is full, this means that the training pattern acts as a support vector. Note that the CE algorithm utilizes only one training pattern out of 300 for the best results.

4.2. Real-world data

Each dataset was divided ten times into a training set of 150 patterns and a test set of the remaining patterns. Three kernels were used:

1. Linear: $K_{i,j} = \langle \mathbf{x}_i, \mathbf{x}_j \rangle$.
2. Polynomial of degree 5: $K_{i,j} = (\langle \mathbf{x}_i, \mathbf{x}_j \rangle + 1)^5$.




Figure 2. SVM classifier for synthetic problem (56.7% SVs).




Figure 3. One-norm SVM classifier for synthetic problem (0.667% SVs).




Figure 4. CE classifier for synthetic problem (0.333% SVs).

3. Gaussian, or radial basis functions (RBF): $K_{i,j} = \exp(-(\langle \mathbf{x}_i, \mathbf{x}_j \rangle^2)/(2\sigma_{RBF}^2))$.

For the linear and polynomial kernels the 5-fold cross validation error was calculated for $C = \lambda/(1 - \lambda)$, where λ had values between 0.1 and 0.9 in increments of 0.1. The RBF kernel is strongly dependant on its hyperparameter σ_{RBF} . Therefore we set $C = 1$ and the 5-fold cross validation error was calculated for 9 linearly spaced values of σ_{RBF} between 1 and 5. Nine values were used to match the number of hyperparameter values for the other kernels. This particular range of value of hyperparameter σ_{RBF} was used because we observed that typically for these datasets the kernel matrix shifted from an almost unit matrix ($\sigma_{RBF} = 1$) to a matrix of ones ($\sigma_{RBF} = 5$).

The training features of all the datasets, algorithms and kernels were preprocessed to have zero mean. In the case of the linear kernel only zero mean preprocessing was used. In the case of the polynomial kernel all datasets but the first dataset were preprocessed to have a unit variance (in addition to zero mean). In the case of the Gaussian kernel for all but the first dataset and only the SVM algorithm for the second and third datasets the data were preprocessed to have a unit variance (all datasets were preprocessed to have zero mean). The test examples were processed accordingly.

Table 1. The real-world datasets. The value of hyperparameter C was equal to λ for the Bupa dataset to allow convergence.

No.	Name	Features	Patterns
1	Ionosphere	34	351
2	Pima	8	768
3	Wdbc	30	569
4	Bupa	6	345
5	Sonar	60	208

The algorithms are compared by two criteria: the error rate and the number of Support Vectors (SVs). Results for linear kernel, polynomial kernel, and Gaussian kernel are presented in Tables 2, 3, and 4, respectively. The criteria was averaged over the partitions into training and test sets.

The CE classifier produced error rates comparable to the standard SVM algorithm. Yet it consistently did so by using a significantly smaller percentage of SVs. On average, across the real-world datasets, the CE classifier used 29%, 70%, 40% less SVs than the one-norm algorithm with the linear, polynomial and Gaussian kernels, respectively. It is interesting to note that this behavior was observed for both a "local" kernel (such as the RBF kernel) and a "global" kernel (such as the linear kernel), indicating that this phenomena cannot be attributed to the locality of the kernel.

5. Discussion

In this paper we considered solving SVM type classification problems with zero norm regularization. Since the optimization problem is non-convex and even discontinuous we formulated it as a combinatorial optimization problem and applied the CE method to solve it. The suggested method was shown to be able to cope well with several rather difficult problems. It produced test errors that were statistically similar to SVMs with L_2 regularization. The main benefit of the method is the sparsity of the solution which was significant in all the instances we experienced with.

A natural question at this point would be why did we employ the CE method and not some other heuristic search method such as genetic algorithms or simulated annealing for searching over the set of SVs. There are two broad types of heuristic search methods. The first type is based on having a proximity relation that facilitates efficient local search. Examples of this type include simulated annealing (Aarts & Korst, 1989) and guided local search (Voudouris & Tsang, 1999). The second type includes global methods that are typically based on iteratively creating collections of solutions, improving the quality of the collection from iteration to iteration. The first type of search algorithms is not suitable for the zero-norm minimization problem because the problem has many discontinuities, and algorithms that are based on local search are likely to get stuck. Algorithms that employ some penalty function on traversed areas (such as the guided local search algorithm) are not likely to perform well since there are many discontinuities leading to an overly complicated penalty function. The CE method belongs to the second type of algorithms. As opposed to other algorithms of this type, such as genetic algorithms, the CE method naturally fits the classification problem and is easy to use. The CE method does not require tedious tweaking or definition of unnatural operators as genetic algorithms do. Finally, similarly to simulated annealing the CE method is guaranteed to converge to the optimal solution in the large sample limit. This provides some theoretical guarantees on the performance of the algorithm.

The computational cost of running the CE based classifier is determined by the size of the LP problems and by the available computational resources. Currently, there are excellent and highly optimized LP solvers that are readily available. In the problems we describe above, a typical run of the CE method required solving a few thousand small linear programs, which can be done rather efficiently. It is worth noting that the CE classification algorithm is "embarrassingly paral-

Table 2. The error and SV percentage (in parentheses) for the real-world datasets with a linear kernel. The number after the $\pm$ represents the standard deviation of either the error or the percentage of SVs. The lowest error and sparsest solution for every dataset are in boldface.

No.	SVM	One-norm SVM	CE
1	14.7 $\pm$ 2.0 (36.4 $\pm$ 8.9)	13.0$\pm$1.8 (15.1 $\pm$ 2.5)	14.6 $\pm$ 1.8 (7.7$\pm$2.6)
2	24.3$\pm$1.4 (51.2 $\pm$ 6.2)	24.6 $\pm$ 1.1 (4.9 $\pm$ 0.5)	24.8 $\pm$ 1.5 (3.9$\pm$1.2)
3	5.7$\pm$1.2 (10.1 $\pm$ 2.5)	5.9 $\pm$ 1.4 (4.8 $\pm$ 1.1)	5.9 $\pm$ 0.8 (3.7$\pm$1.4)
4	32.6 $\pm$ 2.1 (71.9 $\pm$ 3.8)	32.5$\pm$1.7 (4.0 $\pm$ 0.0)	33.4 $\pm$ 2.8 (3.1$\pm$0.6)
5	25.9 $\pm$ 3.7 (53.7 $\pm$ 7.9)	25.5$\pm$4.7 (14.7 $\pm$ 2.4)	25.5$\pm$4.7 (10.3 $\pm$ 1.9)

Table 3. The error and SV percentage for the real-world datasets with a polynomial kernel of degree 5.

No.	SVM	One-norm SVM	CE
1	15.3 $\pm$ 2.7 (36.1 $\pm$ 3.7)	13.7 $\pm$ 2.6 (20.5 $\pm$ 8.4)	12.5$\pm$1.3 (7.1 $\pm$ 1.1)
2	33.2 $\pm$ 1.5 (48.8 $\pm$ 5.2)	30.6 $\pm$ 1.8 (29.5 $\pm$ 4.6)	30.2$\pm$2.4 (11.2 $\pm$ 6.6)
3	6.0 $\pm$ 2.1 (21.9 $\pm$ 2.7)	8.5 $\pm$ 2.8 (15.1 $\pm$ 3.2)	5.6$\pm$1.3 (2.5 $\pm$ 0.7)
4	33.7$\pm$5.2 (58.0 $\pm$ 6.0)	36.3 $\pm$ 2.2 (33.9 $\pm$ 3.5)	37.9 $\pm$ 4.4 (14.4 $\pm$ 9.9)
5	15.9$\pm$4.7 (70.3 $\pm$ 1.7)	20.3 $\pm$ 7.0 (51.1 $\pm$ 6.8)	23.3 $\pm$ 5.3 (6.9 $\pm$ 1.6)

Table 4. The error and SV percentage (in parentheses) for the real-world datasets with a Gaussian kernel and $C = 1$.

No.	SVM	One-norm SVM	CE
1	9.8 $\pm$ 2.3 (76.3 $\pm$ 2.2)	6.2$\pm$1.5 (19.3 $\pm$ 3.1)	6.6 $\pm$ 2.3 (14.1 $\pm$ 2.6)
2	27.5 $\pm$ 1.7 (67.9 $\pm$ 5.1)	25.2$\pm$3.0 (12.9 $\pm$ 4.9)	25.4 $\pm$ 3.3 (8.5 $\pm$ 1.7)
3	7.5 $\pm$ 0.8 (42.4 $\pm$ 3.4)	4.6$\pm$1.5 (14.4 $\pm$ 1.5)	4.7 $\pm$ 1.4 (9.7 $\pm$ 1.2)
4	34.4$\pm$3.0 (93.4 $\pm$ 1.6)	36.9 $\pm$ 3.9 (28.3 $\pm$ 25.5)	36.9 $\pm$ 4.6 (10.4 $\pm$ 4.3)
5	46.7 $\pm$ 6.3 (100.0 $\pm$ 0.0)	24.3$\pm$3.5 (41.7 $\pm$ 6.2)	24.5 $\pm$ 3.7 (22.5 $\pm$ 2.5)

lel" since the sampling method and calculation of the score (solving the LP) can be efficiently divided across processors. As a result, a speedup which is linear in the number of available CPUs can be obtained.

We observed in experiments that using the CE method as a second phase after running one norm SVM results in a sparser solution with a similar error. The CE method can therefore be considered as a post-processing phase, where classifiers that are obtained using a "standard" method are being made sparser. We expect that even greater gain would be obtained for regression problems. For such problems, the fraction of SVs that is typically obtained by kernel methods is large. The issue of using the CE method to further sparsify regressors is left for future research.

The classification algorithm we suggested operates on a single data set. In order to speed-up the computation for large data sets, it may be of interest to use active set methods (see Mangasarian & Musicant, 2001 for consideration of active set methods in SVMs). The idea is to partition the classification problem to multiple smaller optimization problems. After each of these problems is solved separately a global solution is obtained by merging the solutions of the smaller

problems. This may be done, perhaps, by interleaving activation of one norm SVMs with CE method type iterations on different parts of the data set. This issue is left for future research.

Acknowledgements

This research was partially supported by NSERC.

References

- Aarts, E., & Korst, J. (1989). *Simulated Annealing and Boltzmann Machines*. John Wiley & Sons.
- Bradley, P. S., & Mangasarian, O. L. (1998). Feature selection via concave minimization and support vector machines. *Proceedings of the 15th International Conference on Machine Learning* (pp. 82-90).
- Burges, C. J. C. (1996). Simplified support vector decision rules. *Proceedings of the 13th International Conference on Machine Learning* (pp. 71-77).
- Burges, C. J. C., & Schölkopf, B. B. (1997). Improving speed and accuracy of support vector learning ma-

- chines. *Advances in Neural Information Processing Systems*, 9 (pp. 375 - 381).
- Chang, C.-C., & Lin, C.-J. (2002). Training nu-support vector regression: theory and algorithms. *Neural Computation*, 14, 1959-1977.
- Cristianini, N., & Shawe-Taylor, J. (2000). *An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods*. Cambridge, England: Cambridge University Press.
- de-Boer, P., Kroese, D., Mannor, S., & Rubinstein, R. (2005). A tutorial on the cross-entropy method. *Annals of Operations Research*, 134, 19-67.
- Downs, T., Gates, K., & Masters, A. (2001). Exact simplification of support vector solutions. *Journal of Machine Learning Research*, 2, 293-297.
- Floyd, S., & Warmuth, M. (1995). Sample compression, learnability, and the Vapnik-Chervonenkis dimension. *Machine Learning*, 21, 269-304.
- Herbrich, R. (2002). *Learning kernel classifiers: Theory and algorithms*. Boston: MIT Press.
- Herbrich, R., & Williamson, R. C. (2002). Algorithmic luckiness. *Journal of Machine Learning Research*, 3, 175-212.
- Keerthi, S. S., Shevade, S. K., Bhattacharyya, C., & Murthy, K. R. K. (2001). Improvements to platts SMO algorithm for SVM classifier design. *Neural Computation*, 13, 637 -649.
- Mangasarian, O., & Musicant, D. (2001). Active set support vector machine classification. *Advances in Neural Information Processing Systems*, 13 (pp. 577-583).
- Platt, J. (1998). Fast training of support vector machines using sequential minimal optimization. In B. Schölkopf, C. Burges and A. Smola (Eds.), *Advances in kernel methods - support vector learning*. MIT press.
- Rubinstein, R., & Kroese., D. (2004). *The cross-entropy method: A unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*. Springer-Verlag.
- Schölkopf, B., & Smola, A. J. (2002). *Learning with kernels*. MIT Press.
- Smola, A. J., Bartlett, P., Schölkopf, B., & Schuurmans, D. (2000). *Advances in large margin classifiers*. MIT Press.
- Vapnik, V. N. (1998). *Statistical Learning Theory*. New York: Wiley Interscience.
- Voudouris, C., & Tsang, E. (1999). Guided local search. *European Journal of Operational Research*, 113, 469-499.
- Zhu, J., Rosset, S., Hastie, T., & Tibshirani, R. (2003). 1-norm support vector machines. *Advances in Neural Information Processing Systems*, 16.