
New Approaches to Support Vector Ordinal Regression

Wei Chu

CHUWEI@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, University College London, London, WC1N 3AR, UK

S. Sathiya Keerthi

SELVARAK@YAHOO-INC.COM

Yahoo! Research Labs, 210 S. DeLacey Avenue, Pasadena, CA-91105, USA

Abstract

In this paper, we propose two new support vector approaches for ordinal regression, which optimize multiple thresholds to define parallel discriminant hyperplanes for the ordinal scales. *Both approaches guarantee that the thresholds are properly ordered at the optimal solution.* The size of these optimization problems is linear in the number of training samples. The SMO algorithm is adapted for the resulting optimization problems; it is extremely easy to implement and scales efficiently as a quadratic function of the number of examples. The results of numerical experiments on benchmark datasets verify the usefulness of these approaches.

1. Introduction

We consider the supervised learning problem of predicting variables of ordinal scale, a setting that bridges metric regression and classification, and referred to as *ranking learning* or *ordinal regression*. Ordinal regression arises frequently in social science and information retrieval where human preferences play a major role. The training samples are labelled by a set of ranks, which exhibits an ordering among the different categories. In contrast to metric regression problems, these ranks are of finite types and the metric distances between the ranks are not defined. These ranks are also different from the labels of multiple classes in classification problems due to the existence of the ordering information.

There are several approaches to tackle ordinal regression problems in the domain of machine learning. The naive idea is to transform the ordinal scales into numeric values, and then solve the problem as a stan-

dard regression problem. Kramer et al. (2001) investigated the use of a regression tree learner in this way. A problem with this approach is that there might be no principled way of devising an appropriate mapping function since the true metric distances between the ordinal scales are unknown in most of the tasks. Another idea is to decompose the original ordinal regression problem into a set of binary classification tasks. Frank and Hall (2001) converted an ordinal regression problem into nested binary classification problems that encode the ordering of the original ranks and then organized the results of these binary classifiers in some *ad hoc* way for prediction. It is also possible to formulate the original problem as a large augmented binary classification problem. Har-Peled et al. (2002) proposed a constraint classification approach that provides a unified framework for solving ranking and multi-classification problems. Herbrich et al. (2000) applied the principle of Structural Risk Minimization (Vapnik, 1995) to ordinal regression leading to a new distribution-independent learning algorithm based on a loss function between pairs of ranks. The main difficulty with these two algorithms (Har-Peled et al., 2002; Herbrich et al., 2000) is that the problem size of these formulations is a quadratic function of the training data size. As for sequential learning, Crammer and Singer (2002) proposed a perceptron-based online algorithm for rank prediction, known as the PRank algorithm.

Shashua and Levin (2003) generalized the support vector formulation for ordinal regression by finding $r - 1$ thresholds that divide the real line into r consecutive intervals for the r ordered categories. However there is a problem with their approach: the ordinal inequalities on the thresholds, $b_1 \leq b_2 \leq \dots \leq b_{r-1}$, are not included in their formulation. This omission may result in disordered thresholds at the solution on some unfortunate cases (see section 4.1 for an example).

In this paper, we propose two new approaches for support vector ordinal regression. The first one takes only the adjacent ranks into account in determining

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

the thresholds, exactly as Shashua and Levin (2003) proposed, but we introduce explicit constraints in the problem formulation that enforce the inequalities on the thresholds. The second approach is entirely new; it considers the training samples from all the ranks to determine each threshold. Interestingly, we show that, in this second approach, the ordinal inequality constraints on the thresholds are automatically satisfied at the optimal solution though there are no explicit constraints on these thresholds. For both approaches the size of the optimization problems is linear in the number of training samples. We show that the popular SMO algorithm (Platt, 1999; Keerthi et al., 2001) for SVMs can be easily adapted for the two approaches. The resulting algorithms scale efficiently; empirical analysis shows that the cost is roughly a quadratic function of the problem size. Using several benchmark datasets we demonstrate that the generalization capabilities of the two approaches are much better than that of the naive approach of doing standard regression on the ordinal labels.

The paper is organized as follows. In section 2 we present the first approach with explicit inequality constraints on the thresholds, derive the optimality conditions for the dual problem, and adapt the SMO algorithm for the solution. In section 3 we present the second approach with implicit constraints. In section 4 we do an empirically study to show the scaling properties of the two algorithms and their generalization performance. We conclude in section 5.

Notations Throughout this paper we will use x to denote the input vector of the ordinal regression problem and $\phi(x)$ to denote the feature vector in a high dimensional reproducing kernel Hilbert space (RKHS) related to x by transformation. All computations will be done using the reproducing kernel function only, which is defined as

$$\mathcal{K}(x, x') = \langle \phi(x) \cdot \phi(x') \rangle \quad (1)$$

where $\langle \cdot \rangle$ denotes inner product in the RKHS. Without loss of generality, we consider an ordinal regression problem with r ordered categories and denote these categories as consecutive integers $\mathbf{Y} = \{1, 2, \dots, r\}$ to keep the known ordering information. In the j -th category, where $j \in \mathbf{Y}$, the number of training samples is denoted as n^j , and the i -th training sample is denoted as x_i^j where $x_i^j \in \mathbb{R}^d$. The total number of training samples $\sum_{j=1}^r n^j$ is denoted as n . $b_j, j = 1, \dots, r-1$ denote the $(r-1)$ thresholds.

2. Explicit Constraints on Thresholds

As a powerful computational tool for supervised learning, support vector machines (SVMs) map the input vectors into feature vectors in a high dimensional

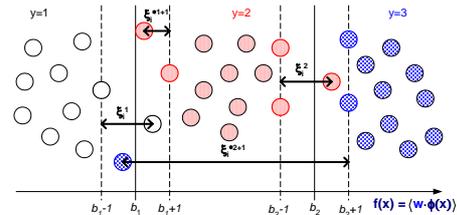


Figure 1. An illustration of the definition of slack variables ξ and ξ^* for the thresholds. The samples from different ranks, represented as circles filled with different patterns, are mapped by $\langle \mathbf{w} \cdot \phi(x) \rangle$ onto the axis of function value. Note that a sample from rank $j+1$ could be counted twice for errors if it is sandwiched by $b_{j+1} - 1$ and $b_j + 1$ where $b_{j+1} - 1 < b_j + 1$, and the samples from rank $j+2, j-1$ etc. never give contributions to the threshold b_j .

RKHS (Vapnik, 1995; Schölkopf & Smola, 2002), where a linear machine is constructed by minimizing a regularized functional. For binary classification (a special case of ordinal regression with $r = 2$), SVMs find an optimal direction that maps the feature vectors into function values on the real line, and a single optimized threshold is used to divide the real line into two regions for the two classes respectively. In the setting of ordinal regression, the support vector formulation could attempt to find an optimal mapping direction \mathbf{w} , and $r-1$ thresholds, which define $r-1$ parallel discriminant hyperplanes for the r ranks accordingly. For each threshold b_j , Shashua and Levin (2003) suggested considering the samples from the two adjacent categories, j and $j+1$, for empirical errors (see Figure 1 for an illustration). More exactly, each sample in the j -th category should have a function value that is less than the lower margin $b_j - 1$, otherwise $\langle \mathbf{w} \cdot \phi(x_i^j) \rangle - (b_j - 1)$ is the error (denoted as ξ_i^j); similarly, each sample from the $(j+1)$ -th category should have a function value that is greater than the upper margin $b_j + 1$, otherwise $(b_j + 1) - \langle \mathbf{w} \cdot \phi(x_i^{j+1}) \rangle$ is the error (denoted as ξ_i^{*j+1}).¹ Shashua and Levin (2003) generalized the *primal* problem of SVMs to ordinal regression as follows:

$$\min_{\mathbf{w}, b, \xi, \xi^*} \frac{1}{2} \langle \mathbf{w} \cdot \mathbf{w} \rangle + C \sum_{j=1}^{r-1} \left(\sum_{i=1}^{n^j} \xi_i^j + \sum_{i=1}^{n^{j+1}} \xi_i^{*j+1} \right) \quad (2)$$

subject to

$$\begin{aligned} \langle \mathbf{w} \cdot \phi(x_i^j) \rangle - b_j &\leq -1 + \xi_i^j, \\ \xi_i^j &\geq 0, \quad \text{for } i = 1, \dots, n^j; \\ \langle \mathbf{w} \cdot \phi(x_i^{j+1}) \rangle - b_j &\geq +1 - \xi_i^{*j+1}, \\ \xi_i^{*j+1} &\geq 0, \quad \text{for } i = 1, \dots, n^{j+1}; \end{aligned} \quad (3)$$

where j runs over $1, \dots, r-1$ and $C > 0$.

¹The superscript $*$ in ξ_i^{*j+1} denotes that the error is associated with a sample in the adjacent upper category of the j -th threshold.

A problem with the above formulation is that the natural ordinal inequalities on the thresholds, i.e., $b_1 \leq b_2 \leq \dots \leq b_{r-1}$ cannot be guaranteed to hold at the solution. To tackle this problem, we explicitly include the following constraints in (3):

$$b_{j-1} \leq b_j, \quad \text{for } j = 2, \dots, r-1. \quad (4)$$

2.1. Primal and Dual Problems

By introducing two auxiliary variables $b_0 = -\infty$ and $b_r = +\infty$, the modified *primal* problem in (2)–(4) can be equivalently written as follows:

$$\min_{\mathbf{w}, \mathbf{b}, \xi, \xi^*} \frac{1}{2} \langle \mathbf{w} \cdot \mathbf{w} \rangle + C \sum_{j=1}^r \sum_{i=1}^{n^j} (\xi_i^j + \xi_i^{*j}) \quad (5)$$

subject to

$$\begin{aligned} \langle \mathbf{w} \cdot \phi(x_i^j) \rangle - b_j &\leq -1 + \xi_i^j, & \xi_i^j &\geq 0, \forall i, j; \\ \langle \mathbf{w} \cdot \phi(x_i^j) \rangle - b_{j-1} &\geq +1 - \xi_i^{*j}, & \xi_i^{*j} &\geq 0, \forall i, j; \\ b_{j-1} &\leq b_j, \forall j. \end{aligned} \quad (6)$$

The *dual* problem can be derived by standard Lagrangian techniques. Let $\alpha_i^j \geq 0$, $\gamma_i^j \geq 0$, $\alpha_i^{*j} \geq 0$, $\gamma_i^{*j} \geq 0$ and $\mu^j \geq 0$ be the Lagrangian multipliers for the inequalities in (6). The Lagrangian for the *primal* problem is:

$$\begin{aligned} \mathcal{L}_e &= \frac{1}{2} \langle \mathbf{w} \cdot \mathbf{w} \rangle + C \sum_{j=1}^r \sum_{i=1}^{n^j} (\xi_i^j + \xi_i^{*j}) \\ &- \sum_{j=1}^r \sum_{i=1}^{n^j} \alpha_i^j (-1 + \xi_i^j - \langle \mathbf{w} \cdot \phi(x_i^j) \rangle + b_j) \\ &- \sum_{j=1}^r \sum_{i=1}^{n^j} \alpha_i^{*j} (-1 + \xi_i^{*j} + \langle \mathbf{w} \cdot \phi(x_i^{*j}) \rangle - b_{j-1}) \\ &- \sum_{j=1}^r \gamma_i^j \xi_i^j - \sum_{j=1}^r \gamma_i^{*j} \xi_i^{*j} - \sum_{j=1}^r \mu^j (b_j - b_{j-1}). \end{aligned} \quad (7)$$

The KKT conditions for the *primal* problem require the following to hold:

$$\frac{\partial \mathcal{L}_e}{\partial \mathbf{w}} = \mathbf{w} - \sum_{j=1}^r \sum_{i=1}^{n^j} (\alpha_i^{*j} - \alpha_i^j) \phi(x_i^j) = 0; \quad (8)$$

$$\frac{\partial \mathcal{L}_e}{\partial \xi_i^j} = C - \alpha_i^j - \gamma_i^j = 0, \quad \forall i, \forall j; \quad (9)$$

$$\frac{\partial \mathcal{L}_e}{\partial \xi_i^{*j}} = C - \alpha_i^{*j} - \gamma_i^{*j} = 0, \quad \forall i, \forall j; \quad (10)$$

$$\frac{\partial \mathcal{L}_e}{\partial b_j} = \sum_{i=1}^{n^j} (\alpha_i^j + \mu^j) - \sum_{i=1}^{n^{j+1}} (\alpha_i^{*j+1} + \mu^{j+1}) = 0, \forall j.$$

Note that the dummy variables associated with b_0 and b_r , i.e. μ^1 , μ^r , α_i^{*1} 's and α_i^r 's, are always zero. The conditions (9) and (10) give rise to the constraints $0 \leq \alpha_i^j \leq C$ and $0 \leq \alpha_i^{*j} \leq C$ respectively. Let us now apply Wolfe duality theory to the *primal* problem. By introducing the KKT conditions (8)–(10) into the Lagrangian (7) and applying the kernel trick (1), the *dual* problem becomes a maximization problem involving the Lagrangian multipliers α , α^* and μ :

$$\max \sum_{j,i} (\alpha_i^j + \alpha_i^{*j}) - \frac{1}{2} \sum_{j,i} \sum_{j',i'} (\alpha_i^{*j} - \alpha_i^j) (\alpha_{i'}^{*j'} - \alpha_{i'}^{j'}) \mathcal{K}(x_i^j, x_{i'}^{j'}) \quad (11)$$

subject to

$$\begin{aligned} 0 &\leq \alpha_i^j \leq C, & \forall i, \forall j, \\ 0 &\leq \alpha_i^{*j+1} \leq C, & \forall i, \forall j, \\ \sum_{i=1}^{n^j} \alpha_i^j + \mu^j &= \sum_{i=1}^{n^{j+1}} \alpha_i^{*j+1} + \mu^{j+1}, & \forall j, \\ \mu^j &\geq 0, & \forall j, \end{aligned} \quad (12)$$

where j runs over $1, \dots, r-1$. Leaving the dummy variables out of account, the size of the optimization problem is $2n - n^1 - n^r$ (α and α^*) plus $r-2$ (for μ).

The *dual* problem (11)–(12) is a convex quadratic programming problem. Once the α , α^* and μ are obtained by solving this problem, \mathbf{w} is obtained from (8). The determination of the b_j 's will be addressed in the next section. The discriminant function value for a new input vector x is

$$f(x) = \langle \mathbf{w} \cdot x \rangle = \sum_{j,i} (\alpha_i^{*j} - \alpha_i^j) \mathcal{K}(x_i^j, x). \quad (13)$$

The predictive ordinal decision function is given by $\arg \min_i \{i : f(x) < b_i\}$.

2.2. Optimality Conditions for the Dual

To derive proper stopping conditions for algorithms that solve the *dual* problem and also determine the thresholds b_j 's, it is important to write down the optimality conditions for the *dual*. Though the resulting conditions that are derived below look a bit clumsy because of the notations, the ideas behind them are very much similar to those for the binary SVM classifier case. The Lagrangian for the *dual* can be written down as follows:

$$\begin{aligned} \mathcal{L}_d &= \frac{1}{2} \sum_{j,i} \sum_{j',i'} (\alpha_i^{*j} - \alpha_i^j) (\alpha_{i'}^{*j'} - \alpha_{i'}^{j'}) \mathcal{K}(x_i^j, x_{i'}^{j'}) \\ &+ \sum_{j=1}^{r-1} \beta_j (\sum_{i=1}^{n^j} \alpha_i^j - \sum_{i=1}^{n^{j+1}} \alpha_i^{*j+1} + \mu^j - \mu^{j+1}) \\ &- \sum_{j,i} \eta_i^j (\alpha_i^j + \alpha_i^{*j}) - \sum_{j,i} \pi_i^j (C - \alpha_i^j) + \\ &\pi_i^{*j} (C - \alpha_i^{*j}) - \sum_{j=2}^{r-1} \lambda^j \mu^j - \sum_{j,i} (\alpha_i^j + \alpha_i^{*j}) \end{aligned}$$

where the Lagrangian multipliers η_i^j , η_i^{*j} , π_i^j , π_i^{*j} and λ^j are non-negative, while β_j can take any value.

The KKT conditions associated with β_j can be given as follows:

$$\begin{aligned} \frac{\partial \mathcal{L}_d}{\partial \alpha_i^j} &= -f(x_i^j) - 1 - \eta_i^j + \pi_i^j + \beta_j = 0, \pi_i^j \geq 0, \\ \eta_i^j &\geq 0, \pi_i^j (C - \alpha_i^j) = 0, \eta_i^j \alpha_i^j = 0, \text{ for } i = 1, \dots, n^j; \\ \frac{\partial \mathcal{L}_d}{\partial \alpha_i^{*j+1}} &= f(x_i^{j+1}) - 1 - \eta_i^{*j+1} + \pi_i^{*j+1} - \beta_j = 0, \\ \pi_i^{*j+1} &\geq 0, \eta_i^{*j+1} \geq 0, \pi_i^{*j+1} (C - \alpha_i^{*j+1}) = 0, \\ \eta_i^{*j+1} \alpha_i^{*j+1} &= 0, \text{ for } i = 1, \dots, n^{j+1}; \end{aligned} \quad (14)$$

where $f(x)$ is defined as in (13), while the KKT conditions associated with the μ_j are

$$\beta_j - \beta_{j-1} - \lambda^j = 0, \lambda^j \mu^j = 0, \lambda^j \geq 0, \quad (15)$$

where $j = 2, \dots, r-1$. The conditions in (14) can be re-grouped into the following six cases:

$$\begin{aligned}
 \text{case 1 : } & \alpha_i^j = 0 & f(x_i^j) + 1 & \leq \beta_j \\
 \text{case 2 : } & 0 < \alpha_i^j < C & f(x_i^j) + 1 & = \beta_j \\
 \text{case 3 : } & \alpha_i^j = C & f(x_i^j) + 1 & \geq \beta_j \\
 \text{case 4 : } & \alpha_i^{*j+1} = 0 & f(x_i^{*j+1}) - 1 & \geq \beta_j \\
 \text{case 5 : } & 0 < \alpha_i^{*j+1} < C & f(x_i^{*j+1}) - 1 & = \beta_j \\
 \text{case 6 : } & \alpha_i^{*j+1} = C & f(x_i^{*j+1}) - 1 & \leq \beta_j
 \end{aligned}$$

We can classify any variable into one of the following six sets:

$$\begin{aligned}
 I_{0a}^j &= \{i \in \{1, \dots, n^j\} : 0 < \alpha_i^j < C\} \\
 I_{0b}^j &= \{i \in \{1, \dots, n^{j+1}\} : 0 < \alpha_i^{*j+1} < C\} \\
 I_1^j &= \{i \in \{1, \dots, n^{j+1}\} : \alpha_i^{*j+1} = 0\} \\
 I_2^j &= \{i \in \{1, \dots, n^j\} : \alpha_i^j = 0\} \\
 I_3^j &= \{i \in \{1, \dots, n^j\} : \alpha_i^j = C\} \\
 I_4^j &= \{i \in \{1, \dots, n^{j+1}\} : \alpha_i^{*j+1} = C\}
 \end{aligned}$$

Let us denote $I_0^j = I_{0a}^j \cup I_{0b}^j$, $I_{up}^j = I_0^j \cup I_1^j \cup I_3^j$ and $I_{low}^j = I_2^j \cup I_4^j$. We further define $F_{up}^i(\beta_j)$ on the set I_{up}^j as

$$F_{up}^i(\beta_j) = \begin{cases} f(x_i^j) + 1 & \text{if } i \in I_{0a}^j \cup I_3^j \\ f(x_i^{*j+1}) - 1 & \text{if } i \in I_{0b}^j \cup I_1^j \end{cases}$$

and $F_{low}^i(\beta_j)$ on the set I_{low}^j as

$$F_{low}^i(\beta_j) = \begin{cases} f(x_i^j) + 1 & \text{if } i \in I_{0a}^j \cup I_2^j \\ f(x_i^{*j+1}) - 1 & \text{if } i \in I_{0b}^j \cup I_4^j \end{cases}$$

Then the conditions can be simplified as

$$\beta_j \leq F_{up}^i(\beta_j) \forall i \in I_{up}^j \text{ and } \beta_j \geq F_{low}^i(\beta_j) \forall i \in I_{low}^j,$$

which can be compactly written as:

$$b_{low}^j \leq \beta_j \leq b_{up}^j \quad (16)$$

where $b_{up}^j = \min\{F_{up}^i(\beta_j) : i \in I_{up}^j\}$ and $b_{low}^j = \max\{F_{low}^i(\beta_j) : i \in I_{low}^j\}$.

The KKT conditions in (15) indicate that the condition, $\beta_{j-1} \leq \beta_j$ always holds, and that $\beta_{j-1} = \beta_j$ if $\mu^j > 0$. To merge the conditions (15) and (16), let us define

$$\tilde{B}_{low}^j = \max\{b_{low}^k : k = 1, \dots, j\}$$

and

$$\tilde{B}_{up}^j = \min\{b_{up}^k : k = j, \dots, r-1\},$$

where $j = 1, \dots, r-1$. The overall optimality conditions can be simply written as

$$B_{low}^j \leq \beta_j \leq B_{up}^j \quad \forall j$$

where

$$B_{low}^j = \begin{cases} \tilde{B}_{low}^{j+1} & \text{if } \mu^{j+1} > 0 \\ \tilde{B}_{low}^j & \text{otherwise} \end{cases}$$

and

$$B_{up}^j = \begin{cases} \tilde{B}_{up}^{j-1} & \text{if } \mu^j > 0 \\ \tilde{B}_{up}^j & \text{otherwise.} \end{cases}$$

Table 1. The basic framework of the SMO algorithm for support vector ordinal regression using explicit threshold constraints.

SMO	start at a valid point, α , α^* and μ , that satisfy (12), find the current B_{up}^j and $B_{low}^j \forall j$
Loop	do <ol style="list-style-type: none"> 1. determine the <i>active threshold</i> J 2. optimize the pair of active variables and the set μ_a 3. compute B_{up}^j and $B_{low}^j \forall j$ at the new point while the optimality condition (17) has not been satisfied
Exit	return α , α^* and b

We introduce a tolerance parameter $\tau > 0$, usually 0.001, to define approximate optimality conditions. The overall stopping condition becomes

$$\max\{B_{low}^j - B_{up}^j : j = 1, \dots, r-1\} \leq \tau. \quad (17)$$

From the conditions in (14) and (3), it is easy to see the close relationship between the b_j 's in the *primal* problem and the multipliers β_j 's. In particular, at the optimal solution, β_j and b_j are identical. Thus b_j can be taken to be any value from the interval, $[B_{low}^j, B_{up}^j]$. We can resolve any non-uniqueness by simply taking $b_j = \frac{1}{2}(B_{low}^j + B_{up}^j)$. Note that the KKT conditions in (15), coming from the additive constraints in (4) we introduced in Shashua and Levin's formulation, enforce $B_{low}^{j-1} \leq B_{low}^j$ and $B_{up}^{j-1} \leq B_{up}^j$ at the solution, which guarantee that the thresholds specified in these feasible regions will satisfy the inequality constraints $b_{j-1} \leq b_j$; without the constraints in (4), the thresholds might be disordered at the solution!

2.3. SMO Algorithm

In this section we adapt the SMO algorithm (Platt, 1999; Keerthi et al., 2001) for the solution of (11)–(12). The key idea of SMO consists of starting with a valid initial point and optimizing only one pair of variables at a time while fixing all the other variables. The suboptimization problem of the two active variables can be solved analytically. Table 1 presents an outline of the SMO implementation for our optimization problem.

In order to determine the pair of *active variables* to optimize, we select the *active threshold* first. The index of the *active threshold* is defined as $J = \arg \max_j \{j : B_{low}^j - B_{up}^j > \tau\}$. Let us assume that B_{low}^J and B_{up}^J are actually defined by $b_{low}^{j_o}$ and $b_{up}^{j_u}$ respectively, and that the two multipliers associated with $b_{low}^{j_o}$ and $b_{up}^{j_u}$ are α_o and α_u . The pair of multipliers (α_o, α_u) is optimized from the current point $(\alpha_o^{old}, \alpha_u^{old})$ to reach the new point, $(\alpha_o^{new}, \alpha_u^{new})$.

It is possible that $j_o \neq j_u$. In this case, named as *cross update*, more than one equality constraint in (12) is involved in the optimization that may update the

variable set $\boldsymbol{\mu}_a = \{\mu^{\min\{j_o, j_u\}+1}, \dots, \mu^{\max\{j_o, j_u\}}\}$, a subset of $\boldsymbol{\mu}$. In the case of $j_o = j_u$, named as *standard update*, only one equality constraint is involved and the variables of $\boldsymbol{\mu}$ are keep intact, i.e. $\boldsymbol{\mu}_a = \emptyset$. These suboptimization problems can be solved analytically, and the detailed formulas for updating can be found in our longer technical report (Chu & Keerthi, 2005).

3. Implicit Constraints on Thresholds

In this section we present a new approach to support vector ordinal regression. Instead of considering only the empirical errors from the samples of adjacent categories to determine a threshold, we allow the samples in all the categories to contribute errors for each threshold. A very nice property of this approach is that the ordinal inequalities on the thresholds are satisfied automatically at the optimal solution in spite of the fact that such constraints on the thresholds are not explicitly included in the new formulation.

Figure 2 explains the new definition of slack variables ξ and ξ^* . For a threshold b_j , the function values of all the samples from all the lower categories, should be less than the lower margin $b_j - 1$; if that does not hold, then $\xi_{ki}^j = \langle \mathbf{w} \cdot \phi(x_i^k) \rangle - (b_j - 1)$ is taken as the error associated with the sample x_i^k for b_j , where $k \leq j$. Similarly, the function values of all the samples from the upper categories should be greater than the upper margin $b_j + 1$; otherwise $\xi_{ki}^{*j} = (b_j + 1) - \langle \mathbf{w} \cdot \phi(x_i^k) \rangle$ is the error associated with the sample x_i^k for b_j , where $k > j$. Here, the subscript ki denotes that the slack variable is associated with the i -th input sample in the k -th category; the superscript j denotes that the slack variable is associated with the lower categories of b_j ; and the superscript $*j$ denotes that the slack variable is associated with the upper categories of b_j .

3.1. Primal Problem

By taking all the errors associated with all $r-1$ thresholds into account, the *primal* problem can be defined as follows:

$$\min_{\mathbf{w}, \mathbf{b}, \boldsymbol{\xi}, \boldsymbol{\xi}^*} \frac{1}{2} \langle \mathbf{w} \cdot \mathbf{w} \rangle + C \sum_{j=1}^{r-1} \left(\sum_{k=1}^j \sum_{i=1}^{n^k} \xi_{ki}^j + \sum_{k=j+1}^r \sum_{i=1}^{n^k} \xi_{ki}^{*j} \right) \quad (18)$$

subject to

$$\begin{aligned} \langle \mathbf{w} \cdot \phi(x_i^k) \rangle - b_j &\leq -1 + \xi_{ki}^j, & \xi_{ki}^j &\geq 0, \\ \text{for } k &= 1, \dots, j \text{ and } i = 1, \dots, n^k; \\ \langle \mathbf{w} \cdot \phi(x_i^k) \rangle - b_j &\geq +1 - \xi_{ki}^{*j}, & \xi_{ki}^{*j} &\geq 0, \\ \text{for } k &= j+1, \dots, r \text{ and } i = 1, \dots, n^k; \end{aligned} \quad (19)$$

where j runs over $1, \dots, r-1$. Note that there are $r-1$ inequality constraints for each sample x_i^k (one for each threshold).

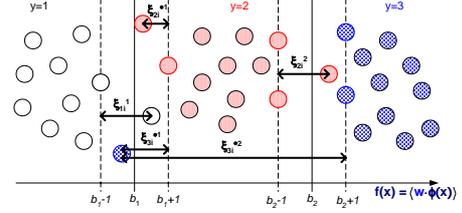


Figure 2. An illustration on the new definition of slack variables ξ and ξ^* that imposes implicit constraints on the thresholds. All the samples are mapped by $\langle \mathbf{w} \cdot \phi(x) \rangle$ onto the axis of function values. Note the term ξ_{3i}^{*1} in this graph.

To prove the inequalities on the thresholds at the optimal solution, let us consider the situation where \mathbf{w} is fixed and only the b_j 's are optimized. Note that the ξ_{ki}^j and ξ_{ki}^{*j} are automatically determined once the b_j are given. To eliminate these variables, let us define, for $1 \leq k \leq r$,

$$\begin{aligned} I_k^{\text{low}}(b) &= \{i \in \{1, \dots, n^k\} : \langle \mathbf{w} \cdot \phi(x_i^k) \rangle - b \geq -1\}, \\ I_k^{\text{up}}(b) &= \{i \in \{1, \dots, n^k\} : \langle \mathbf{w} \cdot \phi(x_i^k) \rangle - b \leq 1\}. \end{aligned}$$

It is easy to see that b_j is optimal iff it minimizes the function

$$e_j(b) = \sum_{k=1}^j \sum_{i \in I_k^{\text{low}}(b)} (\langle \mathbf{w} \cdot \phi(x_i^k) \rangle - b + 1) + \sum_{k=j+1}^r \sum_{i \in I_k^{\text{up}}(b)} (-\langle \mathbf{w} \cdot \phi(x_i^k) \rangle + b + 1) \quad (20)$$

Let B_j^* denote the set of all minimizers of $e_j(b)$. By convexity, B_j^* is a closed interval. Given two intervals $B_1 = [c_1, d_1]$ and $B_2 = [c_2, d_2]$, we say $B_1 \leq B_2$ if $c_1 \leq c_2$ and $d_1 \leq d_2$.

Lemma 1. $B_1^* \leq B_2^* \leq \dots \leq B_{r-1}^*$

Proof. The “right side derivative” of e_j with respect to b is

$$g_j(b) = -\sum_{k=1}^j |I_k^{\text{low}}(b)| + \sum_{k=j+1}^r |I_k^{\text{up}}(b)| \quad (21)$$

Take any one j and consider $B_j^* = [c_j, d_j]$ and $B_{j+1}^* = [c_{j+1}, d_{j+1}]$. Suppose $c_j > c_{j+1}$. Define $b_j^* = c_j$ and $b_{j+1}^* = c_{j+1}$. Since b_{j+1}^* is strictly to the left of the interval B_j^* that minimizes e_j , we have $g_j(b_{j+1}^*) < 0$. Since b_{j+1}^* is a minimizer of e_{j+1} we also have $g_{j+1}(b_{j+1}^*) \geq 0$. Thus we have $g_{j+1}(b_{j+1}^*) - g_j(b_{j+1}^*) > 0$; also, by (21) we get

$$0 < g_{j+1}(b_{j+1}^*) - g_j(b_{j+1}^*) = -|I_{j+1}^{\text{low}}(b_{j+1}^*)| - |I_{j+1}^{\text{up}}(b_{j+1}^*)|$$

which is impossible. In a similar way, $d_j > d_{j+1}$ is also not possible. This proves the lemma.

If the optimal b_j are all unique,² then Lemma 1 implies that the b_j satisfy the natural ordinal ordering. Even when one or more b_j 's are non-unique, Lemma 1 says that there exist choices for the b_j that obey

²If, in the primal problem, we regularize the b_j also (i.e., include the extra cost term $\sum_j b_j^2/2$) then the b_j are guaranteed to be unique. Lemma 1 still holds in this case.

the natural ordering. The fact that the order preservation comes about automatically is interesting and non-trivial, which differs from the PRank algorithm (Crammer & Singer, 2002) where the order preservation on the thresholds is easily brought in via their update rule.

It is also worth noting that Lemma 1 holds even for an extended problem formulation that allows the use of different costs (different C values) for different misclassifications (class k misclassified as class j can have a C_k^j). In applications such as collaborative filtering such a problem formulation can be very appropriate; for example, an A rated movie that is misrated as C may need to be penalized much more than if a B rated movie is misrated as C. Shashua and Levin’s formulation and its extension given in section 2 of this paper do not precisely support such a differential cost structure. This is another good reason in support of the implicit problem formulation of the current section.

3.2. Dual Problem

Let $\alpha_{ki}^j \geq 0$, $\gamma_{ki}^j \geq 0$, $\alpha_{ki}^{*j} \geq 0$ and $\gamma_{ki}^{*j} \geq 0$ be the Lagrangian multipliers for the inequalities in (19). Using ideas parallel to those in section 2.1 we can show that the *dual* of (18)–(19) is the following maximization problem that involves only the multipliers α and α^* :

$$\begin{aligned} \max_{\alpha, \alpha^*} & -\frac{1}{2} \sum_{k,i} \sum_{k',i'} \left(\sum_{j=1}^{k-1} \alpha_{ki}^{*j} - \sum_{j=k}^{r-1} \alpha_{ki}^j \right) \left(\sum_{j=1}^{k'-1} \alpha_{k'i'}^{*j} \right. \\ & \left. - \sum_{j=k'}^{r-1} \alpha_{k'i'}^j \right) \mathcal{K}(x_i^k, x_{i'}^{k'}) + \sum_{k,i} \left(\sum_{j=1}^{k-1} \alpha_{ki}^{*j} + \sum_{j=k}^{r-1} \alpha_{ki}^j \right) \end{aligned} \quad (22)$$

subject to

$$\begin{aligned} \sum_{k=1}^j \sum_{i=1}^{n^k} \alpha_{ki}^j &= \sum_{k=j+1}^r \sum_{i=1}^{n^k} \alpha_{ki}^{*j} \quad \forall j \\ 0 \leq \alpha_{ki}^j &\leq C \quad \forall j \text{ and } k \leq j \\ 0 \leq \alpha_{ki}^{*j} &\leq C \quad \forall j \text{ and } k > j. \end{aligned} \quad (23)$$

The *dual* problem (22)–(23) is a convex quadratic programming problem. The size of the optimization problem is $(r-1)n$ where $n = \sum_{k=1}^r n^k$ is the total number of training samples. The discriminant function value for a new input vector x is

$$f(x) = \langle \mathbf{w} \cdot x \rangle = \sum_{k,i} \left(\sum_{j=1}^{k-1} \alpha_{ki}^{*j} - \sum_{j=k}^{r-1} \alpha_{ki}^j \right) \mathcal{K}(x_i^k, x).$$

The predictive ordinal decision function is given by $\arg \min_i \{i : f(x) < b_i\}$.

The ideas for adapting SMO to (22)–(23) are similar to those in section 2.3. The resulting suboptimization problem is analogous to the case of *standard update* in

section 2.3 where only one of the equality constraints from (23) is involved. Full details of the derivation of the dual problem as well as the SMO algorithm have been skipped for lack of space. These details are given in our longer technical report (Chu & Keerthi, 2005).

4. Numerical Experiments

We have implemented the two SMO algorithms for the ordinal regression formulations with explicit constraints (EXC) and implicit constraints (IMC),³ along with the algorithm of Shashua and Levin (2003) for comparison purpose. The function caching technique and the double-loop scheme proposed by Keerthi et al. (2001) have been incorporated in the implementation for efficiency. We begin this section with a simple dataset to illustrate the typical behavior of the three algorithms, and then empirically study the scaling properties of our algorithms. Then we compare the generalization performance of our algorithms against standard support vector regression on eight benchmark datasets for ordinal regression. The following Gaussian kernel was used in these experiments:

$$\mathcal{K}(x, x') = \exp \left(-\frac{\kappa}{2} \sum_{\zeta=1}^d (x_{\zeta} - x'_{\zeta})^2 \right) \quad (24)$$

where x_{ζ} denotes the ζ -th element of the input vector x . The tolerance parameter τ was set to 0.001 for all the algorithms. We have utilized two evaluation metrics which quantify the accuracy of predicted ordinal scales $\{\hat{y}_1, \dots, \hat{y}_t\}$ with respect to true targets $\{y_1, \dots, y_t\}$: a) *Mean absolute error* is the average deviation of the prediction from the true target, i.e. $\frac{1}{t} \sum_{i=1}^t |\hat{y}_i - y_i|$, in which we treat the ordinal scales as consecutive integers; b) *Mean zero-one error* is simply the fraction of incorrect predictions.

4.1. Grading Dataset

The grading dataset was used in chapter 4 of Johnson and Albert (1999) as an example of the ordinal regression problem.⁴ There are 30 samples of students’ score. The “sat-math score” and “grade in prerequisite probability course” of these students are used as input features, and their final grades are taken as the targets. In our experiments, the six students with final grade A or E were not used, and the feature associated with the “grade in prerequisite probability course” was treated as a continuous variable though it had an ordinal scale. In Figure 3 we present the solution obtained by the

³The source code (written in ANSI C) of our implementation of the two algorithms can be found at <http://www.gatsby.ucl.ac.uk/~chuwei/svor.htm>.

⁴The grading dataset is available at <http://www.mathworks.com/support/books/book1593.jsp>.

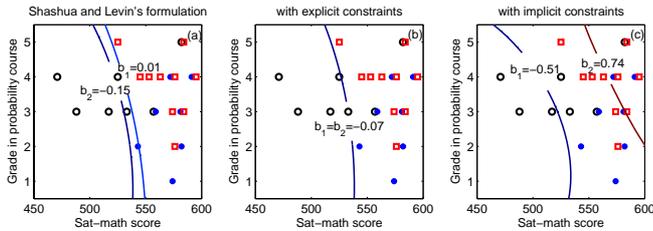


Figure 3. The training results of the three algorithms using a Gaussian kernel on the grading dataset. The discriminant function values are presented as contour graphs indexed by the two thresholds. The circles denote the students with grade D, the dots denote grade C, and the squares denote grade B.

three algorithms using the Gaussian kernel (24) with $\kappa = 0.5$ and the regularization factor value of $C = 1$. In this particular setting, the solution to Shashua and Levin (2003)’s formulation has disordered thresholds $b_2 < b_1$ as shown in Figure 3 (left plot); the formulation with explicit constraints corrects this disorder and yields equal values for the two thresholds as shown in Figure 3 (middle plot).

4.2. Scaling

In this experiment, we empirically studied how the two SMO algorithms scale with respect to training data size and the number of ordinal scales in the target. The California Housing dataset was used in the scaling experiments.⁵ Twenty-eight training datasets with sizes ranging from 100 to 5,000 were generated by random selection from the original dataset. The continuous target variable of the California Housing data was discretized to ordinal scale by using 5 or 10 equal-frequency bins. The standard support vector regression (SVR) was used as a baseline, in which the ordinal targets were treated as continuous values and $\epsilon = 0.1$. These datasets were trained by the two algorithms using a Gaussian kernel with $\kappa = 1$ and a regularization factor value of $C = 100$. Figure 4 gives plots of the computational costs of the three algorithms as functions of the problem size, for the two cases of 5 and 10 target bins. Our algorithms scale well with scaling exponents between 2.13 and 2.33, while the scaling exponent of SVR is about 2.40 in this case. This near-quadratic property in scaling comes from the sparseness property of SVMs, i.e., non-support vectors affect the computational cost only mildly. The EXC and IMC algorithms cost more than the SVR approach due to the larger problem size. For large sizes, the cost of EXC is only about x times that of SVR. As expected, we also noticed that the computational cost of IMC is dependent on r , the number of ordinal scales in the

⁵The California Housing dataset can be found at <http://lib.stat.cmu.edu/datasets/>.

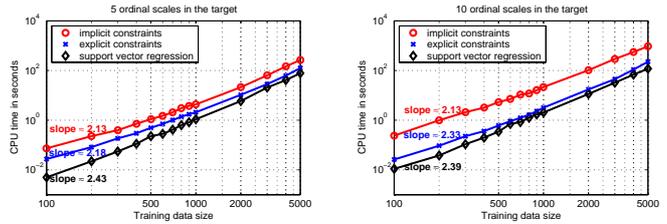


Figure 4. Plots of CPU time versus training data size on log – log scale, indexed by the estimated slopes respectively. We used the Gaussian kernel with $\kappa = 1$ and the regularization factor value of $C = 100$ in the experiment.

target. The cost for 10 ranks is observed to be roughly 5 times that for 5 ranks, whereas the cost of EXC is nearly the same for the two cases. These observations are consistent with the size of the optimization problems. The problem size of IMC is $(r - 1)n$ (which is heavily influenced by r) while the problem size of EXC is about $2n + r$ (which largely depends on n only since we usually have $n \gg r$). This factor of efficiency can be a key advantage for the EXC formulation.

4.3. Benchmark datasets

Next, we compared the generalization performance of the two approaches against the naive approach of using standard support vector regression (SVR) and the method (SLA) of Shashua and Levin (2003). We collected eight benchmark datasets that were used for metric regression problems.⁶ For each dataset, the target values were discretized into ten ordinal quantities using equal-frequency binning. We randomly partitioned each dataset into training/test splits as specified in Table 2. The partitioning was repeated 20 times independently. The input vectors were normalized to zero mean and unit variance, coordinate-wise. The Gaussian kernel (24) was used for all the algorithms. 5-fold cross validation was used to determine the optimal values of model parameters (the Gaussian kernel parameter κ and the regularization factor C) involved in the problem formulations, and the test error was obtained using the optimal model parameters for each formulation. The initial search was done on a 7×7 coarse grid linearly spaced in the region $\{(\log_{10} C, \log_{10} \kappa) \mid -3 \leq \log_{10} C \leq 3, -3 \leq \log_{10} \kappa \leq 3\}$, followed by a fine search on a 9×9 uniform grid linearly spaced by 0.2 in the $(\log_{10} C, \log_{10} \kappa)$ space. The ordinal targets were treated as continuous values in standard SVR, and the predictions for test cases were rounded to the nearest ordinal scale. The insensitive zone parameter, ϵ of SVR was fixed at 0.1. The test results of the four algorithms are recorded in Table 2. It is very clear that the generalization capabilities

⁶These regression datasets are available at <http://www.liacc.up.pt/~ltorgo/Regression/DataSets.html>.

Table 2. Test results of the four algorithms using a Gaussian kernel. The targets of these benchmark datasets were discretized by 10 equal-frequency bins. The results are the averages over 20 trials, along with the standard deviation. d denotes the input dimension and “training/test” denotes the partition size. We use bold face to indicate the lowest average value among the results of the four algorithms. The symbols \star are used to indicate the cases significantly worse than the winning entry; A p-value threshold of 0.01 in Wilcoxon rank sum test was used to decide this.

Dataset	Partition		Mean zero-one error				Mean absolute error			
	d	training/test	SVR	SLA	EXC	IMC	SVR	SLA	EXC	IMC
Pyrimidines	27	50/24	0.777±0.068 \star	0.756±0.073	0.752±0.063	0.719±0.066	1.404±0.184	1.400±0.255	1.331±0.193	1.294±0.204
MachineCPU	6	150/59	0.693±0.056 \star	0.643±0.057	0.661±0.056	0.655±0.045	1.048±0.141	1.002±0.121	0.986±0.127	0.990±0.115
Boston	13	300/206	0.589±0.025 \star	0.561±0.023	0.569±0.025	0.561±0.026	0.785±0.052	0.765±0.057	0.773±0.049	0.747±0.049
Abalone	8	1000/3177	0.758±0.017 \star	0.739±0.008 \star	0.736±0.011	0.732±0.007	1.407±0.021 \star	1.389±0.027 \star	1.391±0.021 \star	1.361±0.013
Bank	32	3000/5192	0.786±0.004 \star	0.759±0.005 \star	0.744±0.005	0.751±0.005 \star	1.471±0.010 \star	1.414±0.012 \star	1.512±0.017 \star	1.393±0.011
Computer	21	4000/4192	0.494±0.006 \star	0.462±0.006	0.462±0.005	0.473±0.005 \star	0.632±0.011 \star	0.597±0.010	0.602±0.009	0.596±0.008
California	8	5000/15640	0.677±0.003 \star	0.640±0.003	0.640±0.003	0.639±0.003	1.070±0.008 \star	1.068±0.006 \star	1.068±0.005 \star	1.008±0.005
Census	16	6000/16784	0.735±0.004 \star	0.699±0.002	0.699±0.002	0.705±0.002 \star	1.283±0.009 \star	1.271±0.007 \star	1.270±0.007 \star	1.205±0.007

of the three ordinal regression algorithms are better than that of the approach of SVR. The performance of Shashua and Levin’s method is similar to our EXC approach, as expected, since the two formulations are pretty much the same. Our ordinal algorithms are comparable on the mean zero-one error, but the results also show the IMC algorithm yields much more stable results on mean absolute error than the EXC algorithm.⁷ From the view of the formulations, EXC only considers the extremely worst samples between successive ranks, whereas IMC takes all the samples into account. Thus the outliers may affect the results of EXC significantly, while the results of IMC are relatively more stable in both validation and test.

5. Conclusion

In this paper we proposed two new approaches to support vector ordinal regression that determine $r - 1$ parallel discriminant hyperplanes for the r ranks by using $r - 1$ thresholds. The ordinal inequality constraints on the thresholds are imposed explicitly in the first approach and implicitly in the second one. The problem size of the two approaches is linear in the number of training samples. We also designed SMO algorithms that scale only about quadratically with the problem size. The results of numerical experiments verified that the generalization capabilities of these approaches are much better than the naive approach of applying standard regression.

Acknowledgments

A part of the work was carried out at IPAM of UCLA. WC was supported by the National Institutes of Health and its National Institute of General Medical Sciences division

⁷As pointed out by a reviewer, $\xi_i^j + \xi_i^{*j+1}$ in (2) of EXC is an upper bound on the zero-one error of the i -th example, while, in (18) of IMC, $\sum_{k=1}^j \xi_{ki}^j + \sum_{k=j+1}^r \xi_{ki}^{*j}$ is an upper bound on the absolute error. Note that, in all the examples we use consecutive integers to represent the ordinal scales.

under Grant Number 1 P01 GM63208.

References

Chu, W., & Keerthi, S. S. (2005). *New approaches to support vector ordinal regression* (Technical Report). Yahoo! Research Labs.

Crammer, K., & Singer, Y. (2002). Pranking with ranking. *Advances in Neural Information Processing Systems 14* (pp. 641–647). Cambridge, MA: MIT Press.

Frank, E., & Hall, M. (2001). A simple approach to ordinal classification. *Proceedings of the European Conference on Machine Learning* (pp. 145–165).

Har-Peled, S., Roth, D., & Zimak, D. (2002). Constraint classification: A new approach to multiclass classification and ranking. *Advances in Neural Information Processing Systems 15*.

Herbrich, R., Graepel, T., & Obermayer, K. (2000). Large margin rank boundaries for ordinal regression. *Advances in Large Margin Classifiers* (pp. 115–132). MIT Press.

Johnson, V. E., & Albert, J. H. (1999). *Ordinal data modeling (statistics for social science and public policy)*. Springer-Verlag.

Keerthi, S. S., Shevade, S. K., Bhattacharyya, C., & Murthy, K. R. K. (2001). Improvements to Platt’s SMO algorithm for SVM classifier design. *Neural Computation, 13*, 637–649.

Kramer, S., Widmer, G., Pfahringer, B., & DeGroeve, M. (2001). Prediction of ordinal classes using regression trees. *Fundamenta Informaticae, 47*, 1–13.

Platt, J. C. (1999). Fast training of support vector machines using sequential minimal optimization. *Advances in Kernel Methods - Support Vector Learning* (pp. 185–208). MIT Press.

Schölkopf, B., & Smola, A. J. (2002). *Learning with kernels*. The MIT Press.

Shashua, A., & Levin, A. (2003). Ranking with large margin principle: two approaches. *Advances in Neural Information Processing Systems 15* (pp. 937–944).

Vapnik, V. N. (1995). *The nature of statistical learning theory*. New York: Springer-Verlag.