

Feature relevance determination for ordinal regression in the context of feature redundancies and privileged information

Pfannschmidt, Lukas; Jacob, Jonathan; Hinder, Fabian; Biehl, Michael ; Tino, Peter; Hammer, Barbara

DOI:

[10.1016/j.neucom.2019.12.133](https://doi.org/10.1016/j.neucom.2019.12.133)

License:

Creative Commons: Attribution-NonCommercial-NoDerivs (CC BY-NC-ND)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Pfannschmidt, L, Jacob, J, Hinder, F, Biehl, M, Tino, P & Hammer, B 2020, 'Feature relevance determination for ordinal regression in the context of feature redundancies and privileged information', *Neurocomputing*, vol. 416, pp. 266-279. <https://doi.org/10.1016/j.neucom.2019.12.133>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Feature Relevance Determination for Ordinal Regression in the Context of Feature Redundancies and Privileged Information[☆]

Lukas Pfannschmidt^{a,*}, Jonathan Jakob^a, Fabian Hinder^a, Michael Biehl^b,
Peter Tino^c, Barbara Hammer^a

^a*Machine Learning Group, Bielefeld University, DE*

^b*Intelligent Systems Group, University of Groningen, NL*

^c*Computer Science, University of Birmingham, UK*

Abstract

Advances in machine learning technologies have led to increasingly powerful models in particular in the context of big data. Yet, many application scenarios demand for robustly interpretable models rather than optimum model accuracy; as an example, this is the case if potential biomarkers or causal factors should be discovered based on a set of given measurements. In this contribution, we focus on feature selection paradigms, which enable us to uncover relevant factors of a given regularity based on a sparse model. We focus on the important specific setting of linear ordinal regression, i.e. data have to be ranked into one of a finite number of ordered categories by a linear projection. Unlike previous work, we consider the case that features are potentially redundant, such that no unique minimum set of relevant features exists. We aim for an identification of all strongly and all weakly relevant features as well as their type of relevance (strong or weak); we achieve this goal by determining feature relevance bounds, which correspond to the minimum and maximum feature relevance, respectively, if searched over all equivalent models. In addition, we discuss how this setting enables us to substitute some of the features, e.g. due to their semantics, and how to extend the framework of feature relevance intervals to the setting of privileged information, i.e. potentially relevant information is available for training purposes only, but cannot be used for the prediction itself.

Keywords: Global Feature Relevance, Feature Selection, Interpretability, Ordinal Regression, Privileged Information

[☆]Funding by the DFG in the frame of the graduate school DiDy (1906/3) and by the BMBF (grant number 01S18041A) is gratefully acknowledged.

*Corresponding author

Email addresses: lukas@lpfann.me (Lukas Pfannschmidt),
jjakob@techfak.uni-bielefeld.de (Jonathan Jakob), fhinder@techfak.uni-bielefeld.de
(Fabian Hinder), m.biehl@rug.nl (Michael Biehl), p.tino@cs.bham.ac.uk (Peter Tino),
bhammer@techfak.uni-bielefeld.de (Barbara Hammer)

1. Introduction

Ordinal regression refers to the task to assign data to a finite number of classes or bins, which are ordered qualitatively along a preference scale. Ordinal data often occur in sociodemographic, financial or medical contexts where it is difficult to give absolute quantitative measurements but easily possible to compare samples and assign those to different bins, which are qualitatively ordered, such as the severity of a disease or the risk of a financial transaction. Another popular example for a ranking on ordinal scales takes place in customer feedback or product ranking by humans [1]. Here, the quality is often represented by a five-star rating scale, where five stars correspond to the best rating and one star to the worst. Indeed, many human ratings are represented in an ordinal scale rather than absolute values.

The *ordinal regression problem* (ORP) is the task to embed given data in the real numbers such that they are ordered according to their label, i.e. the target bin. An error is encountered whenever an ordering of two data points assigned to different bins is violated. Although the problem can be treated as a regular regression or classification method, dedicated techniques are often preferred, since they can account for the fact that the distance between ordinal classes in the data is unknown and not necessarily evenly distributed. Examples of ordinal regression include treatments such as the multiclass classification problem [2], and extensions of standard models such as the support vector machine (SVM) or learning vector quantization (LVQ) to ordinal regression tasks [3, 4, 5, 6]. A recent work proposed an incremental and sparse Bayesian approach with favourable scaling properties [7]. Often, ordinal regression is treated as a pairwise ranking problem [8]. Further, there does exist recent theoretical work which establishes consistency of some surrogate losses for ordinal regression, which have better numeric properties [9].

In this work, we will rely on SVM-like treatments of the ORP due to the mathematical elegance and flexibility of this formulation [3, 4, 5].

Recently, methods which enable the interpretability of machine learning models have extensively been discussed [10]. One common way to enhance model interpretability is by means of a determination of the most relevant input dimensions or features, i.e. the relevance of ordinal explanatory variables for the given task. This is particularly relevant when the objective exceeds mere diagnostics, such as safety-critical decision-making, or the design of repair strategies. There do exist a few approaches which address such feature selection for ordinal regression: The approach [11] uses a minimal redundancy formulation based on a feature importance score to find the subset of relevant features. The work in [12] focuses on multiple filter methods which are adapted to ranking data. These models deliver sparse ordinal regression models which enable some insight into the underlying classification prescription. Yet, their result is arbitrary in the case of correlated or redundant features: if there does not exist a unique minimum relevant feature set, it often depends on arbitrary initialization or algorithmic design choices, which feature from a set of redundant features is chosen. Hence, possibly relevant features, so-called weakly relevant features, can

easily be overlooked, albeit they might have a substantial contribution or even causal influence to a model.

The so-called *all* relevant feature selection problem deals with the challenge to determine all features, which are potentially relevant for a given task – a problem which is particularly important for diagnostics purposes if it is not priorly clear which one of a set of relevant, but redundant features to choose. Finding this subset is generally computationally intractable. For standard classification and regression schemes, a few efficient heuristics have been proposed: one possibility is to quantify not only the relevance but also the redundancy of features [13]. Another popular model extends predictive models with statistical tests to discriminate between relevance and irrelevance [14]. Recently, the problem of feature relevance has been investigated in the special case of linear mappings; here, the problem can be phrased in terms of relevance intervals, leading to a convex problem and superior performance in benchmarks [15]. In the presented work, the goal is to extend this approach to the specific relevant setting of ordinal regression tasks, and to demonstrate the benefit of this model in comparison to alternative popular feature selection models such as lasso or ElasticNet.

Besides a formal mathematical modelling by means of linear optimization tasks, we will also demonstrate the suitability of the model to investigate the role of critical features for an ORP. As an example, the integration of criteria such as age, gender, or ethnicity might improve the prediction accuracy of a given model as measured by an appropriate cost function - yet, it might be debatable if these features can have any relevance for the given task as regards a causal relationship on the one hand; on the other hand, it might be unethical or impossible to actually gather such features for a prediction model in its daily use. Examples for a questionable impact of such characteristics on a formal model have recently been debated under the umbrella of model fairness [16]. We will discuss how feature relevance profiles, in particular the identification of weakly relevant features, enable further insight into such settings, by explicitly quantifying the possible impact of such features.

There exists another popular setting where not all features can or should be used for daily use, hence feature relevances are of particular importance: the scenario of so-called privileged information phrases the situation that some features are available during the training phase only, but not during the test phase, e.g. due to the costs, computational load, or any other restrictions. In classical machine learning, it is commonly assumed, that training and test set have an identical statistical distribution and utilize the same predictive features. In contrast, the *learning using privileged information* paradigm (LUPI) [17] considers additional privileged information only available at training time. This paradigm can be understood as an intelligent teacher feeding the learner extra information to improve the learning process [18]. Additional information could be the output of another model (‘machines-teaching-machines’) or input from a human expert itself, who intuitively knows which examples in the data are hard to discriminate. Examples are medical measurements which require invasive techniques or measurements which require too much time in daily use, but

would be affordable for training. The approach [17] proposed a variant of SVMs that incorporates privileged information for training. The modelling replaces or enriches slack variables, which are required by soft-margin SVMs to correct for hard training samples. This specific approach is known as *similarity control* [18]. The approach [17] introduces the SVM+ in which a smooth function based on the privileged information (PI) is used at training time to improve learning in non-separable classification settings. The method [19] refrained from fully replacing the slack variables and combined them with a smooth function based on PI. It achieved better generalization ability and lower complexity models. Furthermore, this approach also extends the SVM+ to ordinal regression problems.

While approaches to incorporate privileged information exist, and it has been shown that LUPI has the potential to speed up learning [20], the analysis of feature relevances in the context of redundant feature information is still widely open in this setting. In this article, we also introduce an extension of the feature-relevance-interval-computation scheme as proposed in [15] to the LUPI setting; this addresses the question of which features are potentially relevant to facilitate training, i.e. they carry important information to improve the learnability of a task. Irrelevant features in the LUPI framework, on the other hand, do not contribute to the learnability. Unlike standard feature relevances for regression or classification, features relevances for privileged information answer the question whether feature information is beneficial for the learning process itself.

In the following, we will introduce and extend feature relevance learning in the context of redundant features for ordinal regression and privileged information. For this purpose, we recapture two large margin ordinal regression formalizations in section 2, which differ in the type of constraints they enforce on ordinal classes, namely *implicit* and *explicit* constraints. We extend them to an optimization scheme to determine feature relevance bounds in section 3, which can be transferred to several linear optimization problems (Section 3.2). Further we also define the *explicit* formulation to be used in context of learning using privileged information in Section 4. In Section 6 we do several benchmarks to highlight the accuracy and feature selection performance in the classical machine learning case. In Section 6.2 we repeat this in the LUPI setting where we focus on performance measures split by the regular and privileged feature set.

2. Large Margin Ordinal Regression

We consider the following ordinal regression learning task: We assume class labels $L = \{1, 2, \dots, l\}$, which are ordered; w.l.o.g. we represent those as natural numbers. We assume training data are given, $X = \{\mathbf{x}_i^j \in \mathbb{R}^n \mid i = 1, \dots, m_j, j \in L\}$ where data point x_i^j is assigned the class label $j \in L$, i.e. x_i^j is contained in bin number j . The full data set has size $m := m_1 + \dots + m_l$. Here the index j refers to the ordinal target variable the data point x_i^j belongs to. The ORP can be phrased as the search for a mapping $f : \mathbb{R}^n \rightarrow \mathbb{R}$, which preserves the ordering of bins as indicated by the label information. That means the inequality $f(\mathbf{x}_{i_1}^{j_1}) < f(\mathbf{x}_{i_2}^{j_2})$ should hold for all pairs of class labels $j_1 < j_2$ and data indices i_1 and i_2 in these bins.

In the following, we will restrict to the case of a linear function, i.e. $f(\mathbf{x}) = \mathbf{w}^\top \mathbf{x}$ with parameter $\mathbf{w} \in \mathbb{R}^n$. In particular in the case of high dimensional data such a linear prescription is often sufficient to model the underlying regularity. Further, it enables a particularly strong link of feature relevances and underlying model, as already elaborated in popular sparse models such as lasso [21]. There do exist different possibilities to model the ORP learning problem. Here, we will introduce two existing optimization problems, which rely on large margins, and which treat the inequality constraints in two different ways.

Explicit Order Constraints. One way to model ordinal regression is by an embedding of data in the real numbers via f , whereby the bins are separated by adaptive thresholds b_j , which are learned accordingly. A popular formulation which is inspired by support vector machines imposes a margin around all thresholds b_j for this embedding [4]:

$$\min_{\mathbf{w}, \mathbf{b}, \chi, \xi} \frac{1}{2} \|\mathbf{w}\|_1 + C \sum_{i,j} (\chi_i^j + \xi_i^j) \quad (1)$$

s.t. for all i, j

$$\begin{aligned} \mathbf{w}^\top \mathbf{x}_i^j - b_j &\leq -1 + \chi_i^j \\ \mathbf{w}^\top \mathbf{x}_i^{j+1} - b_j &\geq +1 - \xi_i^{j+1} \\ b_j &\leq b_{j+1} \\ \chi_i^j &\geq 0, \xi_i^j \geq 0 \end{aligned} \quad (2)$$

where χ_i^j and ξ_i^j are slack variables, and the thresholds b_j for $j = 1, \dots, l-1$ determine the boundaries which separate the classes, b_j referring to the boundary in between bin j and bin $j+1$. The hyper-parameter $C > 0$ controls the trade-off of the margin and number of errors and it can be chosen through cross validation. We adapt the problem from [4], which uses L_2 regularization, and use L_1 regularization in (Eq. 1), aiming for sparse solutions. In this definition the linear ordering of classes is enforced *explicitly* through constraint $b_j \leq b_{j+1}$. When we refer to (2) in the future, we specifically refer to the constraints of the problem.

Implicit Order Constraints. Another definition first highlighted in [22] enforces the ordering implicitly, by requiring that all data of bin 1 to j are embedded below the threshold b_j , all data from bins $j+1$ to l are above the threshold.

This leads to the implicitly constrained problem:

$$\begin{aligned} \min_{\mathbf{w}, \mathbf{b}, \boldsymbol{\chi}, \boldsymbol{\xi}} \quad & \frac{1}{2} \|\mathbf{w}\|_1 + C \sum_{j=1}^{l-1} \left(\sum_{k=1}^j \sum_{i=1}^{n^k} \chi_{ki}^j + \sum_{k=j+1}^l \sum_{i=1}^{n^k} \xi_{ki}^j \right) \\ \text{subject to} \quad & \mathbf{w}^\top \mathbf{x}_i^k - b_j \leq -1 + \chi_{ki}^j, \quad \chi_{ki}^j \geq 0, \\ & \text{for } k = 1, \dots, j \text{ and } i = 1, \dots, m_k; \\ & \mathbf{w}^\top \mathbf{x}_i^k - b_j \geq +1 - \xi_{ki}^j, \quad \xi_{ki}^j \geq 0 \\ & \text{for } k = j+1, \dots, l \text{ and } i = 1, \dots, m_k. \end{aligned} \tag{3}$$

Again, we adapt the existing problem from [22] and replace the existing regularization $\|\mathbf{w}\|_2$ with $\|\mathbf{w}\|_1$ to induce sparsity. In this definition, not only neighbouring classes are contributing to the overall loss of in between boundaries, but all other classes, as well. This can lead to more robust results in particular in the case of outliers, as shown in [22], but higher computational demand.

In the following we introduce feature relevance bounds for the explicit variant which is an extension from existing work for simple linear classification in [15]. The definition for the implicit variant is very similar and can be found in Appendix B.

3. Feature Relevance Bounds for Ordinal Regression with Explicit Order

Assume a training set X is given. We denote an optimum solution of problem (1) as $(\tilde{\mathbf{w}}, \tilde{\mathbf{b}}, \tilde{\boldsymbol{\xi}}, \tilde{\boldsymbol{\chi}})$. This solution induces the value

$$\mu_X := \frac{1}{2} \|\tilde{\mathbf{w}}\|_1 + C \cdot \sum_{i,j} (\tilde{\chi}_i^j + \tilde{\xi}_i^j)$$

which is uniquely determined by X . The quantity μ_X is unique by definition, albeit the solution $(\tilde{\mathbf{w}}, \tilde{\mathbf{b}}, \tilde{\boldsymbol{\xi}}, \tilde{\boldsymbol{\chi}})$ is not.

We are interested in the class of equivalent good hypotheses, i.e. all weight vectors \mathbf{w} which yield (almost) the same quality as regards the regression error and generalization ability as the function induced by $\tilde{\mathbf{w}}$. This class might contain an infinite number of alternative hypothesis: in the context of correlated features, for example, we can trade one feature for the other. However, the function class cannot explicitly be computed, since the generalization ability is unknown for future data. We use the following surrogate induced by μ_X

$$\begin{aligned} F_\delta(X) \quad & := \{ \mathbf{w} \in \mathbb{R}^n \mid \exists \mathbf{b}, \boldsymbol{\xi}, \boldsymbol{\chi} \text{ such that constraints (2) hold,} \\ & \frac{1}{2} \|\mathbf{w}\|_1 + C \cdot \sum_{i,j} (\xi_i^j + \chi_i^j) \leq (1 + \delta) \cdot \mu_X \} \end{aligned} \tag{4}$$

These constraints ensure the following properties:

1. The empirical error of equivalent functions in $F_\delta(X)$ is minimum, as measured by the slack variables.
2. The loss of the generalization ability is limited, as guaranteed by a small L_1 -norm of the weight vector and learning theoretical guarantees as provided, e.g. by Theorem 7 in [23] and Corollary 5 in [24].

The parameter $\delta \geq 0$ quantifies the tolerated deviation to accept a function as yet good enough, C is determined by Problem (1).

Solutions \mathbf{w} in $F_\delta(X)$ are sparse in the sense that irrelevant features are uniformly weighted as 0 for all solutions in $F_\delta(X)$. Relevant but potentially redundant features can be weighted arbitrarily, disregarding sparsity, similar in spirit to the ElasticNet; yet the latter weights mutually redundant features equally and can therefore hide the relevance in the case of many redundant features [25]. In this contribution we are interested in the relevance of features for forming good hypotheses; more precisely, we are interested in the following more specific characteristics:

- **Strong relevance** of feature \mathcal{I} for $F_\delta(X)$: Is feature \mathcal{I} relevant for all hypotheses in $F_\delta(X)$, i.e. all weight vectors $\mathbf{w} \in F_\delta(X)$ yield $w_{\mathcal{I}} \neq 0$?
- **Weak relevance** of feature \mathcal{I} for $F_\delta(X)$: Is feature \mathcal{I} relevant for at least one hypothesis in $F_\delta(X)$ in the sense that one weight vector $\mathbf{w} \in F_\delta(X)$ exists with $w_{\mathcal{I}} \neq 0$, but this does not hold for all weight vectors in $F_\delta(X)$?
- **Irrelevance** of feature \mathcal{I} for $F_\delta(X)$: Is feature \mathcal{I} irrelevant for every hypothesis in $F_\delta(X)$, i.e. all weight vectors $\mathbf{w} \in F_\delta(X)$ yield $w_{\mathcal{I}} = 0$?

A feature is irrelevant for $F_\delta(X)$ if it is neither strongly nor weakly relevant. The questions of strong and weak relevance can be answered via the following optimization problems:

Problem minrel(\mathcal{I}):

$$\begin{aligned} \min_{\mathbf{w}, \mathbf{b}, \boldsymbol{\chi}, \boldsymbol{\xi}} \quad & |w_{\mathcal{I}}| & (5) \\ \text{s.t. for all } i, j \quad & \text{conditions (2) hold and} \\ & \frac{1}{2} \|\mathbf{w}\|_1 + C \cdot \sum_{k,l} (\chi_k^l + \xi_k^l) \leq (1 + \delta) \cdot \mu_X & (6) \end{aligned}$$

Here $|w_{\mathcal{I}}|$ denotes the absolute value of feature \mathcal{I} in w . Feature \mathcal{I} is strongly relevant for $F_\delta(X)$ iff minrel(\mathcal{I}) yields an optimum larger than 0.

Problem maxrel(\mathcal{I}):

$$\begin{aligned} \max_{\mathbf{w}, \mathbf{b}, \boldsymbol{\chi}, \boldsymbol{\xi}} \quad & |w_{\mathcal{I}}| & (7) \\ \text{s.t. for all } i, j \quad & \text{conditions (2) and (6) hold} \end{aligned}$$

Feature \mathcal{I} is weakly relevant for $F_\delta(X)$ iff minrel(\mathcal{I}) yields an optimum at 0 and maxrel(\mathcal{I}) yields an optimum larger than 0.

These two optimization problems span a real-valued interval for every feature \mathcal{I} with the result of $\text{minrel}(\mathcal{I})$ as lower and $\text{maxrel}(\mathcal{I})$ as upper bound. This interval characterizes the range of weights for \mathcal{I} occupied by good solutions in $F_\delta(X)$. Hence, besides information about a feature’s relevance, some indication about the degree up to which a feature is relevant or can be substituted by others, is given. Note, however, that the solutions are in general not consistent estimators of an underlying ‘true’ weight vector as regards its exact value, as has been discussed, e.g. for lasso [26]. For consistency, it is advisable to use L_2 regularization after the selection of a set of relevant features.

3.1. Generalization Bounds

In the beginning of Section 3 we introduced the set $F_\delta(X)$ of all equivalent good hypotheses which yield (almost) the same quality regarding regression error and generalization ability. However, the impact of the norm of \mathbf{w} and the high loss $\sum_{i,j} (\tilde{\chi}_i^j + \tilde{\xi}_i^j)$ are not considered separately, i.e. a low norm of \mathbf{w} allows a high loss, and vice versa. We would like to control the generalization error by means of l_1 -regularization. To do so, we consider both quantities separately, i.e. we define

$$\mathcal{H}_\delta(\tilde{\mathbf{w}}) := \left\{ \mathbf{w} \in \mathbb{R}^n \mid \exists \mathbf{b}, \boldsymbol{\xi}, \boldsymbol{\chi} \text{ such that constraints (2) hold,} \right. \\ \left. \|\mathbf{w}\|_1 \leq (1 + \delta)\|\tilde{\mathbf{w}}\|_1 \text{ and} \right. \quad (8)$$

$$\left. \sum_{i,j} (\xi_i^j + \chi_i^j) \leq \sum_{i,j} (\tilde{\xi}_i^j + \tilde{\chi}_i^j) \right\}. \quad (9)$$

This allows us to extend the results from [15] to our scenario, i.e. show that the generalization error of all hypothesis with the same or a lower high loss is bounded by means of the l_1 -regularization. Recall Theorem 26.15 from Understanding Machine Learning [27]:

Theorem 1. *Suppose that \mathcal{D} is a distribution on $X \times Y$ such that with probability 1 we have $\|x\|_\infty \leq R$. Let $\mathcal{H} = \{\mathbf{w} \in \mathbb{R}^d \mid \|\mathbf{w}\|_1 \leq B\}$ and let $l : \mathcal{H} \times X \times Y \rightarrow \mathbb{R}$ be of the form $l(\mathbf{w}, (x, y)) = \phi(\langle \mathbf{w}, x \rangle, y)$ where $\phi : \mathbb{R} \times Y \rightarrow \mathbb{R}$ is such that for all $y \in Y$, the function $a \mapsto \phi(a, y)$ is η -Lipschitz and such that $\max_{a \in [-RB, RB]} |\phi(a, y)| \leq c$. Then, for any $\tau \in (0, 1)$ with probability of at least $1 - \tau$ over the choice of i.i.d. sample of size n , for all $\mathbf{w} \in \mathcal{H}$,*

$$\mathbb{E}_{(x,y) \sim \mathcal{D}}[l(\mathbf{w}, x, y)] \leq \frac{1}{n} \sum_{i=1}^n l(\mathbf{w}, x_i, y_i) + 2\eta RB \sqrt{\frac{2 \log(2d)}{n}} + c \sqrt{\frac{2 \ln(2/\tau)}{n}}.$$

To apply this theorem we have to reformulate our classifier as a collection of binary classifiers. Since all classes use the same subspace spanned by \mathbf{w} it is enough to distinguish neighbouring classes, i.e. every b_j gives rise to a classifier that allows us to decide whenever x belongs to one of $0, \dots, j$ or $j + 1, \dots, |L|$.

Consider the ramp loss

$$\begin{aligned} l_{\prec j}(\mathbf{w}, \mathbf{b}, x, y) &= \min\{1, \max\{0, 1 - \mathbf{1}_{y \prec j}(\mathbf{w}^\top x - b_j)\}\}, \\ l_j(\mathbf{w}, \mathbf{b}, x, y) &= l_{\leq j}(\mathbf{w}, \mathbf{b}, x, y) + l_{\geq j}(\mathbf{w}, \mathbf{b}, x, y), \\ l(\mathbf{w}, \mathbf{b}, x, y) &= l_y(\mathbf{w}, \mathbf{b}, x, y) \end{aligned}$$

where $\mathbf{1}_{y \prec j} = 1$ if $y \prec j$ and -1 otherwise for some comparison operation $\cdot \prec \cdot$. Notice that l corresponds to the implicit order constraints, which is an upper bound for the explicit loss where only neighbouring classes are considered, rather than all classes. By using this loss function it is clear that the loss of the original classifier is bounded by the sum of all those binary classifiers. Since the ramp loss is 1-Lipschitz and maps to the interval $[0, 1]$ we may apply Theorem 1 to obtain

$$\begin{aligned} \mathbb{E}_{(x,y) \sim \mathcal{D}}[l(\mathbf{w}, x, y)] &\leq \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\sum_{j=1}^{|L|} (l_{\leq j}(\mathbf{w}, x, y) + l_{\geq j}(\mathbf{w}, x, y)) \right] \\ &= \sum_{j=1}^{|L|} (\mathbb{E}_{(x,y) \sim \mathcal{D}} [l_{\leq j}(\mathbf{w}, x, y)] + \mathbb{E}_{(x,y) \sim \mathcal{D}} [l_{\geq j}(\mathbf{w}, x, y)]) \\ &\leq \sum_{j=1}^{|L|} \left(\frac{1}{n} \sum_{i=1}^n (l_{\leq j}(\mathbf{w}, x_i, y_i) + l_{\geq j}(\mathbf{w}, x_i, y_i)) \right. \\ &\quad \left. + 4RB \sqrt{\frac{2 \log(2d)}{n}} + 2 \sqrt{\frac{2 \ln(2/\tau)}{n}} \right) \end{aligned}$$

for all \mathbf{w} such that $\|\mathbf{w}\|_1 \leq B$ with probability $1 - \tau$ over the choice of sample. In particular, setting $\rho_j = \sum_i \tilde{\xi}_i^j + \tilde{\chi}_i^j$ and $\rho = \sum_j \rho_j$ to the hinge loss of the baseline classifier and using the fact that the hinge loss upper bounds ramp loss, this gives rise to

$$L_{\mathcal{D}}(\tilde{\mathbf{w}}, \tilde{\mathbf{b}}) \leq |L| \left(\frac{\rho}{n} + 4\|\tilde{\mathbf{w}}\|_1 R \sqrt{\frac{2 \log(2d)}{n}} + 2 \sqrt{\frac{2 \ln(2/\tau)}{n}} \right)$$

for the generalization error of the baseline linear classifier $(\tilde{\mathbf{w}}, \tilde{\mathbf{b}})$ and

$$L_{\mathcal{D}}(h) \leq |L| \left(\frac{\rho}{n} + 4(1 + \delta)\|\tilde{\mathbf{w}}\|_1 R \sqrt{\frac{2 \log(2d)}{n}} + 2 \sqrt{\frac{2 \ln(2/\tau)}{n}} \right)$$

for all $h \in \mathcal{H}_\delta(\tilde{\mathbf{w}})$, with probability at least $1 - \tau$ over the choice of training sample, i.e. our choice of constraints allow the generalization error upper bound to increase by $4\delta\|\tilde{\mathbf{w}}\|_1 |L| R \sqrt{\frac{2 \log(2d)}{n}}$.

3.2. Feature Relevance Bounds as Linear Problem

The problems from Section 3 are not yet linear problems, but they can be transferred to linear optimization problems, for which particularly efficient solvers are available.

Theorem 2. *Problem $\text{minrel}(\mathcal{I})$ is equivalent to the following linear optimization problem:*

$$\begin{aligned}
\mathbf{minrel}^*(\mathcal{I}) : \quad & \min_{\mathbf{w}, \mathbf{w}, \mathbf{b}, \boldsymbol{\chi}, \boldsymbol{\xi}} \quad \hat{w}_{\mathcal{I}} \\
& \text{s.t. for all } i, j \quad \text{conditions (2) hold} \\
& \frac{1}{2} \sum_k \hat{w}_k + C \cdot \sum_{k,l} (\chi_k^l + \xi_k^l) \leq (1 + \delta) \cdot \mu_X \quad (10) \\
& w_i \leq \hat{w}_i, -w_i \leq \hat{w}_i \quad (11)
\end{aligned}$$

Problem $\text{maxrel}(\mathcal{I})$ can be solved by taking the optimum of the following two linear optimization problems:

$$\begin{aligned}
\mathbf{maxrel}_{\text{pos}}^*(\mathcal{I}) : \quad & \max_{\mathbf{w}, \mathbf{w}, \mathbf{b}, \boldsymbol{\chi}, \boldsymbol{\xi}} \quad \hat{w}_{\mathcal{I}} \\
& \text{s.t. for all } i, j \quad \text{conditions (2) hold} \\
& \frac{1}{2} \sum_k \hat{w}_k + C \cdot \sum_{k,l} (\chi_k^l + \xi_k^l) \leq (1 + \delta) \cdot \mu_X \\
& w_i \leq \hat{w}_i, -w_i \leq \hat{w}_i \\
& \hat{w}_{\mathcal{I}} \leq w_{\mathcal{I}} \quad (12)
\end{aligned}$$

and the problem

$$\begin{aligned}
\mathbf{maxrel}_{\text{neg}}^*(\mathcal{I}) : \quad & \max_{\mathbf{w}, \mathbf{w}, \mathbf{b}, \boldsymbol{\chi}, \boldsymbol{\xi}} \quad \hat{w}_{\mathcal{I}} \\
& \text{s.t. for all } i, j \quad \text{conditions (2) hold} \\
& \frac{1}{2} \sum_k \hat{w}_k + C \cdot \sum_{k,l} (\chi_k^l + \xi_k^l) \leq (1 + \delta) \cdot \mu_X \\
& w_i \leq \hat{w}_i, -w_i \leq \hat{w}_i \\
& \hat{w}_{\mathcal{I}} \leq -w_{\mathcal{I}} \quad (13)
\end{aligned}$$

The proof can be found in the appendix.

In practice, it might be a good strategy to split constraint (5) into two, separately limiting the weight vector

$$\frac{1}{2} \sum_k \hat{w}_k \leq (1 + \delta) \cdot \|\tilde{\mathbf{w}}\|_1$$

and error term

$$\sum_{k,l} (\chi_k^l + \xi_k^l) \leq \sum_{k,l} (\tilde{\chi}_k^l + \tilde{\xi}_k^l)$$

where the symbols marked $\tilde{\cdot}$ refer to the optimum solution of the original margin-based ordinal regression problem. This split enables us to better control the loss of generalization ability and error terms, and it also mediates the dependency on the hyper-parameter C of the space of equivalent good functions. At a small down-side, this split depends on the found solution and it is no longer uniquely defined by the given training data, albeit we did not observe large variation in practical applications.

4. Learning using Privileged Information

Let us shortly recall the classical setting considered so far: Given ordered class labels $L = \{1, 2, \dots, l\}$ and training data $X = \{\mathbf{x}_i^j \in \mathbb{R}^n \mid i = 1, \dots, m_j, j \in L\}$ where data point x_i^j is assigned the class label $j \in L$. The full data set has size $m := m_1 + \dots + m_l$. Here the index j refers to the ordinal target variable (represented by b_j) the data point x_i^j belongs to.

In the LUPI setting, we work with two types of information X and $X^* = \{\mathbf{x}_i^{*j} \in \mathbb{R}^{n^*} \mid i = 1, \dots, m_j, j \in L\}$ which is a set of additional information commonly called privileged information (PI) where p is the amount of privileged features we have available. The information is privileged in the sense that it is not available in the testing and prediction phase, and it is only present when training the model. This fact does not necessarily imply that the privileged information is of higher quality or exhibits correlation with the label y at all. Rather, there are reasons why it cannot be gathered at prediction time: examples are too costly computations (such as extensive feature preprocessing), unavailability of sensors, unavailability of the information (such as information which is available only in retrospective, or privacy issues which prevent gathering the data (such as personal information)). X and X^* , in general, do not have to share the same space or modality. As an example, X could cover numerical features, and X^* could be textual input from an expert.

4.1. Modelling Slacks in Ordinal Regression

There are several ways to integrate privileged information into the learning model [28]. In the following we only consider *similarity control* where privileged information is interpreted as the teacher giving hints about the difficulty for each training example. These hints can be incorporated into an SVM by means of slack variables which was shown in [19] already. In the following we will extend our *explicit* definition of ordinal regression to handle privileged information by adapting similarity control as used in [19].

We recall that in the explicit variant two types of slacks are used. Each slack value represents a deviation from the classification rule. In the LUPI case, we replace χ_i^j by

$$p_\chi^j(\mathbf{x}_i^*) := \left(\mathbf{w}_\chi^* \cdot \mathbf{x}_i^{*j} + d_\chi \right)$$

and ξ_i^j by the function

$$p_\xi^j(\mathbf{x}_i^*) := \left(\mathbf{w}_\xi^* \cdot \mathbf{x}_i^{*j} + d_\xi \right).$$

$$\begin{aligned}
\min_{\mathbf{w}, \mathbf{b}, \mathbf{w}^*, \mathbf{d}} \quad & \frac{1}{2} \|\mathbf{w}\|_1 + \frac{\gamma}{2} (\|\mathbf{w}_X^*\|_1 + \|\mathbf{w}_\xi^*\|_1) + C \sum_{j=1}^l \sum_{i=1}^{n^k} (p_\chi^j(\mathbf{x}_i^*) + p_\xi^j(\mathbf{x}_i^*)) \\
\text{s.t. for every } j = 1, \dots, l-1, \quad & \mathbf{w}^\top \mathbf{x}_i^j - b_j \leq -1 + p_\chi^j(\mathbf{x}_i^*) \\
& \mathbf{w}^\top \mathbf{x}_i^{j+1} - b_j \geq +1 - p_\xi^{j+1}(\mathbf{x}_i^*) \\
& b_j \leq b_{j+1} \\
& p_\chi^j(\mathbf{x}_i^*) \geq 0, p_\xi^j(\mathbf{x}_i^*) \geq 0
\end{aligned} \tag{14}$$

γ is an additional hyperparameter to scale the influence of privileged information. This allows us to reject nonsense PI by simplifying the model and relying solely on X when considering a cross validation scheme where we expect better generalization ability by a simpler model. The adaption of [19] now enables us to define relevance bounds as in Section 3.

4.2. Feature Relevance Bounds for Ordinal Regression with Privileged Information

We now consider two sets of features. In the following we define bounds for both regarding their relevance to the machine learning procedure when both sets are present. Because PI are not present while predicting they are always irrelevant for that phase. They are relevant to speed up learning by mediating the distribution of slack variables.

Assume a training set $X = \{\mathbf{x}_i^j \in \mathbb{R}^n\}$ and $X^* = \{\mathbf{x}_i^{*j} \in \mathbb{R}^{n^*}\}$. Further we define

$$\mathcal{L} := C \sum_{j=1}^l \sum_{i=1}^{n^k} (p_\chi^j(\mathbf{x}_i^*) + p_\xi^j(\mathbf{x}_i^*))$$

as the total slack loss of problem (14). Denote an optimum solution of the problem as $(\tilde{\mathbf{w}}, \tilde{\mathbf{b}}, \tilde{\mathbf{w}}_\chi^*, \tilde{\mathbf{w}}_\xi^*, \tilde{d}_\chi, \tilde{d}_\xi)$ and its total loss as $\tilde{\mathcal{L}}$. Analogous to Section 3, this solution induces the value

$$\mu_{X, X^*} := \frac{1}{2} \|\tilde{\mathbf{w}}\|_1 + \frac{\gamma}{2} (\|\tilde{\mathbf{w}}_X\|_1 + \|\tilde{\mathbf{w}}_\xi\|_1) + \tilde{\mathcal{L}}.$$

Furthermore, we use the following proxy induced by μ_{X, X^*}

$$\begin{aligned}
F_\delta(X, X^*) \quad & := \{ \mathbf{w} \in \mathbb{R}^n, \mathbf{w}_\chi^*, \mathbf{w}_\xi^* \in \mathbb{R}^{n^*} \mid \exists \mathbf{b}, d_\chi, d_\xi \\
& \text{such that constraints (14) hold and} \\
& \frac{1}{2} \|\mathbf{w}\|_1 + \frac{\gamma}{2} (\|\mathbf{w}_\chi^*\|_1 + \|\mathbf{w}_\xi^*\|_1) + \mathcal{L} \leq (1 + \delta) \cdot \mu_{X, X^*} \}
\end{aligned} \tag{15}$$

This proxy allows us to define similar feature relevances as found in Section 3 for non-privileged feature \mathcal{I} in X :

- **Strong relevance** of feature \mathcal{I} for $F_\delta(X, X^*)$: Is feature \mathcal{I} relevant for all hypotheses in $F_\delta(X, X^*)$, i.e. all weight vectors $\mathbf{w} \in F_\delta(X, X^*)$ yield $w_{\mathcal{I}} \neq 0$?
- **Weak relevance** of feature \mathcal{I} for $F_\delta(X, X^*)$: Is feature \mathcal{I} relevant for at least one hypothesis in $F_\delta(X, X^*)$ in the sense that one weight vector $\mathbf{w} \in F_\delta(X, X^*)$ exists with $w_{\mathcal{I}} \neq 0$, but this does not hold for all weight vectors in $F_\delta(X, X^*)$?
- **Irrelevance** of feature \mathcal{I} for $F_\delta(X, X^*)$: Is feature \mathcal{I} irrelevant for every hypothesis in $F_\delta(X, X^*)$, i.e. all weight vectors $\mathbf{w} \in F_\delta(X, X^*)$ yield $w_{\mathcal{I}} = 0$?

and similarly for feature \mathcal{P} in X^* with $\mathbf{w}_\bullet^* := \{\mathbf{w}_\chi^*, \mathbf{w}_\xi^* \mid (\mathbf{w}^*, \mathbf{w}_\chi^*, \mathbf{w}_\xi^*) \in F_\delta(X, X^*)\}$:

- **Strong relevance** of feature \mathcal{P} for $F_\delta(X, X^*)$: Is feature \mathcal{P} relevant for all hypotheses in $F_\delta(X, X^*)$, i.e. for all \mathbf{w}_\bullet^* in $F_\delta(X, X^*)$ at least one weight vector in \mathbf{w}_\bullet^* for one bin of the ordered classes yields $w_{\bullet, \mathcal{P}}^* \neq 0$?
- **Weak relevance** of feature \mathcal{P} for $F_\delta(X, X^*)$: Is feature \mathcal{P} relevant for at least one hypothesis in $F_\delta(X, X^*)$ in the sense that one weight vector \mathbf{w}_\bullet^* exists with $w_{\bullet, \mathcal{P}}^* \neq 0$, but this does not hold for all \mathbf{w}_\bullet^* in $F_\delta(X, X^*)$?
- **Irrelevance** of feature \mathcal{P} for $F_\delta(X, X^*)$: Is feature \mathcal{P} irrelevant for every hypothesis in $F_\delta(X, X^*)$, i.e. all weight vectors \mathbf{w}_\bullet^* yield $w_{\bullet, \mathcal{P}}^* = 0$?

A feature is irrelevant for $F_\delta(X, X^*)$ if it is neither strongly nor weakly relevant.

The questions of strong and weak relevance can be answered via the following optimization problems:

Problem minrel(\mathcal{P}):

$$\begin{aligned}
& \max_{\bullet \in \{\chi, \xi\}} \min_{\mathbf{w}, \mathbf{w}_\bullet^*, \mathbf{b}, d_\bullet} |w_{\bullet, \mathcal{P}}^*| & (16) \\
& \text{s.t. for all } i, j & \text{conditions (14) hold and} \\
& & \frac{1}{2} \|\mathbf{w}\|_1 + \frac{\gamma}{2} (\|\mathbf{w}_\chi^*\|_1 + \|\mathbf{w}_\xi^*\|_1) + \mathcal{L} \leq (1 + \delta) \cdot \mu_{X, X^*}
\end{aligned}$$

Because of two slack functions and the corresponding weights \mathbf{w}_χ^* and \mathbf{w}_ξ^* we need to optimize two inner feature relevancies $|w_{\bullet, \mathcal{P}}^*|$. To aggregate them to a global feature relevance we take the maximum to express that a feature could be used only in one of both functions, i.e. it is not relevant for all slack functions but at least in one. One could define an additional relevance classification by taking into account cases where the $\min \min > 0$, i.e. the feature is relevant for all slack functions. In the following we limit ourselves to the former case.

Feature \mathcal{P} is strongly relevant for $F_\delta(X, X^*)$ iff minrel(\mathcal{P}) yields an optimum larger than 0.

Problem maxrel(\mathcal{P}):

$$\begin{aligned}
& \max_{\bullet \in \{\chi, \xi\}} \max_{\mathbf{w}, \mathbf{w}_\bullet^*, \mathbf{b}, \chi, \xi} |w_{\bullet \mathcal{P}}^*| & (17) \\
& \text{s.t. for all } i, j & \text{conditions (14) hold and} \\
& & \frac{1}{2} \|\mathbf{w}\|_1 + \frac{\gamma}{2} (\|\mathbf{w}_\chi^*\|_1 + \|\mathbf{w}_\xi^*\|_1) + \mathcal{L} \leq (1 + \delta) \cdot \mu_{X, X^*}
\end{aligned}$$

Similar to the first problem we consider the maximum inner feature relevance to express the global feature relevance.

Feature \mathcal{P} is weakly relevant for $F_\delta(X, X^*)$ iff $\text{minrel}(\mathcal{P})$ yields an optimum 0 and $\text{maxrel}(\mathcal{P})$ yields an optimum larger than 0

4.3. Privileged Feature Relevance Bounds as Linear Problem

Both problems can be transferred to linear optimization problems:

Theorem 3. *Problem $\text{minrel}(\mathcal{P})$ is equivalent to taking the maximum over following two linear optimization problems:*

$$\begin{aligned}
& \mathbf{minrel}_\chi^*(\mathcal{P}) : & & \hat{w}_{\chi \mathcal{P}}^* & (18) \\
& \min_{\mathbf{w}, \hat{\mathbf{w}}, \mathbf{w}_\chi^*, \hat{\mathbf{w}}_\chi^*, \mathbf{w}_\xi^*, \hat{\mathbf{w}}_\xi^*, \mathbf{b}, d_\chi, d_\xi} & & & \\
& \text{s.t. for all } i, j & & \text{conditions (14) hold and} \\
& & & \frac{1}{2} \sum_k \hat{w}_k + \frac{\gamma}{2} \sum_k \hat{w}_{\chi k} + \frac{\gamma}{2} \sum_k \hat{w}_{\xi k} + \mathcal{L} \leq (1 + \delta) \cdot \mu_X \\
& & & w_i \leq \hat{w}_i, -w_i \leq \hat{w}_i \\
& & & \chi_i \leq \hat{\chi}_i, -\chi_i \leq \hat{\chi}_i \\
& & & \xi_i \leq \hat{\xi}_i, -\xi_i \leq \hat{\xi}_i
\end{aligned}$$

and

$$\begin{aligned}
& \mathbf{minrel}_\xi^*(\mathcal{P}) : & & \hat{w}_{\xi \mathcal{P}}^* & (19) \\
& \min_{\mathbf{w}, \hat{\mathbf{w}}, \mathbf{w}_\chi^*, \hat{\mathbf{w}}_\chi^*, \mathbf{w}_\xi^*, \hat{\mathbf{w}}_\xi^*, \mathbf{b}, d_\chi, d_\xi} & & & \\
& \text{s.t. for all } i, j & & \text{conditions (14) hold and} \\
& & & \frac{1}{2} \sum_k \hat{w}_k + \frac{\gamma}{2} \sum_k \hat{w}_{\chi k} + \frac{\gamma}{2} \sum_k \hat{w}_{\xi k} + \mathcal{L} \leq (1 + \delta) \cdot \mu_X \\
& & & w_i \leq \hat{w}_i, -w_i \leq \hat{w}_i \\
& & & \chi_i \leq \hat{\chi}_i, -\chi_i \leq \hat{\chi}_i \\
& & & \xi_i \leq \hat{\xi}_i, -\xi_i \leq \hat{\xi}_i
\end{aligned}$$

For $\text{maxrel}(\mathcal{P})$ we first define the linear optimization problem

$$\begin{aligned}
& \text{maxrel}_{\lambda, \bullet}^*(\mathcal{P}) : && \hat{w}_{\bullet\mathcal{P}}^* && (20) \\
& \max_{\substack{\mathbf{w}, \hat{\mathbf{w}}, \mathbf{w}_{\chi}^*, \hat{\mathbf{w}}_{\chi}^*, \mathbf{w}_{\xi}^*, \hat{\mathbf{w}}_{\xi}^*, \\ \mathbf{b}, d_{\chi}, d_{\xi}}} && && \\
& \text{s.t. for all } i, j && \text{conditions (14) hold and} && \\
& && \frac{1}{2} \sum_k \hat{w}_k + \frac{\gamma}{2} \sum_k \hat{w}_{\chi k}^* + \frac{\gamma}{2} \sum_k \hat{w}_{\xi k}^* + \mathcal{L} \leq (1 + \delta) \cdot \mu_X && \\
& && w_i \leq \hat{w}_i, -w_i \leq \hat{w}_i && \\
& && \chi_i \leq \hat{\chi}_i, -\chi_i \leq \hat{\chi}_i && \\
& && \xi_i \leq \hat{\xi}_i, -\xi_i \leq \hat{\xi}_i && \\
& && \hat{w}_{\bullet\mathcal{P}}^* \leq \lambda \cdot w_{\bullet\mathcal{P}}^* &&
\end{aligned}$$

such that

$$\text{maxrel}(\mathcal{P}) := \max_{\substack{\lambda \in \{-1, +1\}, \\ \bullet \in \{\chi, \xi\}}} \text{maxrel}_{\lambda, \bullet}^*(\mathcal{P}),$$

i.e. the maximum of four linear problems.

A proof of this theorem is similar to Section 3.2 and is omitted for the sake of brevity.

5. Relevance Bounds for Feature Selection

While the relevance bounds should give truthful indication of feature relevance, in practice the discrimination between relevant and irrelevant features is challenging: variations of the underlying distributions of the features have the implication that thresholds for feature relevance can vary for different features. The use of slack variables in the overall model and thus our relevance bounds allow variation in the contribution of features which improves finding stable solutions but also adds noise. This is exacerbated by the behaviour of linear programming solvers, which often have exhibit loss of precision. For relevance bounds specifically, even if feature \mathcal{I} is independent we often observe $\text{maxRel}(\mathcal{I}) > 0$ and $0 < \text{minRel}(\mathcal{I}) < 10^{-5}$.

We do not aim for a data independent threshold to discriminate between noise and relevant features. Instead, we introduce distribution dependent thresholds: we estimate the distribution of relevances of noise features given the model constraints. We expect for a given model class defined by $F_{\delta}(X)$ the same amount of slackness in the relevances for irrelevant variables. This slackness is introduced by the parameters of the algorithm itself (δ , C) and the LP-solvers internal ones and should be similar for truly non-correlated variables. Therefore, we propose to estimate the parameters of a normal distribution and the corresponding prediction interval Π to obtain a data dependent threshold [29].

An existing work proposes a similar resampling based approach to estimate a stopping threshold for a forward feature selection approach [30].

To estimate this noise distribution we use randomly permuted input features from X to imitate irrelevant features. We define $p(\mathcal{I})$ as the random permutation of values in \mathcal{I} and $X^{p(\mathcal{I})} := \{X \setminus \mathcal{I}\} \cup p(\mathcal{I})$ as the dataset where \mathcal{I} was replaced by its random permutation. With these we define two random sample populations

$$\widehat{\pi}(\text{maxrel}) := \{\text{maxrel}(p(\mathcal{I}), X^{p(\mathcal{I})}) \mid \text{where } \mathcal{I} \text{ randomly chosen from } X\}$$

and

$$\widehat{\pi}(\text{minrel}) := \{\text{minrel}(p(\mathcal{I}), X^{p(\mathcal{I})}) \mid \text{where } \mathcal{I} \text{ randomly chosen from } X\}$$

where a population with n samples is defined as $\widehat{\pi}(\cdot)_n$.

The prediction interval is then defined as

$$\Pi(\cdot)_n := \overline{\widehat{\pi}(\cdot)_n} \pm T_{n-1}(p) \cdot \sigma(\widehat{\pi}(\cdot)) \sqrt{1 + (1/n)}.$$

Here $\bar{\pi}_n$ denotes the sample mean and $\sigma(x)$ the standard deviation, and T represents Student's t-distribution with $n - 1$ degrees of freedom. The size of Π depends on parameter p , the expected probability that a new value is included in the interval. We propose default values of $p = 0.999$ for a low false positive rate and $n \geq 50$ which yielded robust thresholds for common feature set sizes in our experiments without adding too many computations to the complexity, which we analyse in Section 5.1.

To classify feature \mathcal{I} as irrelevant we check if its relevance bounds are element of our prediction intervals. We therefore replace the theoretical classifications from Section 3 with the following:

- **Strong relevance:** $\text{maxrel}(\mathcal{I}) \notin \Pi(\text{maxrel}) \wedge \text{minrel}(\mathcal{I}) \notin \Pi(\text{minrel})$
- **Weak relevance:** $\text{maxrel}(\mathcal{I}) \notin \Pi(\text{maxrel}) \wedge \text{minrel}(\mathcal{I}) \in \Pi(\text{minrel})$
- **Irrelevance:** $\text{maxrel}(\mathcal{I}) \in \Pi(\text{maxrel}) \wedge \text{minrel}(\mathcal{I}) \in \Pi(\text{minrel})$

5.1. Time complexity

In the following we outline the scaling behaviour of our proposed method for feature selection. Our method can be divided in three separate computational steps which differ in their algorithmic complexity. We consider a problem with n samples and d features.

The initial baseline solution is analogue to a standard ordinal regression SVM solution which can be solved using the sequential minimal optimization (SMO) algorithm [31, 4] which is in $\mathcal{O}(n^3)$. The relevance bounds are given by a set of linear programs for which interior point methods exist [32, 33, 34] which are in $\mathcal{O}(n^{2.5})$. This complexity bound is very general and one could reformulate and adapt these problems using existing outlines [35, 36]. In the normal setting we consider the constant $z = 3$ for the number of linear programs needed (Section 3.2) and $z = 6$ in the LUPI setting (Section 4.3) such that the

relevance interval for each feature is in $\mathcal{O}(zn^{2.5})$. This results in $\mathcal{O}(dzn^{2.5})$ for all relevance bounds. Additionally, we employ a permutation test approach which adds a constant c additional LPs to achieve statistical stability which is overall in $\mathcal{O}(cn^{2.5})$. Overall our method is in $\mathcal{O}(n^3 + (dz + c)n^{2.5})$ when considering $n > d$.

Because the $dz + c$ LPs are a significant factor, we proposed to solve them in parallel [37] which we evaluate in Appendix B.1.

6. Experiments

We evaluate our methodology in two steps. First, we focus on our ordinal regression approach in the classical machine learning setting - using regular data. Then we examine the adaption of our method to the LUPI paradigm - using data that incorporates privileged information.

6.1. Classical Setting of ORP

In this section, we focus on our ordinal regression method for regular data. We show the quality of our feature selection by evaluating the results of both the explicit and the implicit variant of our method, on theoretically generated data with know ground truth. In addition, we compare both variants with regard to their classification accuracy and run time on standard benchmark datasets. The accuracy is measured using the *Macro-averaged Mean Absolute Error (MMAE)* which is specifically designed for ordinal regression data with imbalanced classes:

$$MMAE = \frac{1}{l} \sum_{j=1}^l \frac{\sum_{i=1}^{m_j} |j - f(x_i^j)|}{m_j}, \quad (21)$$

where l is the number of bins, f refers to the bin the sample x_i^j is assigned to by the learned model, and m_j refers to the number of samples in class j .

The section is rounded off by an analysis of a real world data set, showcasing the insights that can be gained from our method.

6.1.1. Artificial Data

We adapt the generation method presented in [15] for ordinal regression. By using equal frequency binning we convert the continuous regression variable into an ordered discrete target variable with five ordinal classes. The data is generated from a suitable set of informative features. From those we form strongly relevant features by simply picking the desired number out of the informative set. Weakly relevant features are created as linear combinations of informative features. Finally, irrelevant features are drawn from random Gaussian noise. All features are normalized to zero mean and unit variance. The exact characteristics of the datasets used in our experiments are shown in Table 1.

Table 1: Artificially created data sets with known ground truth. The model of which the data is drawn from is based on the strongly relevant features. The weakly relevant features are linear combinations of strong ones. Characteristics of the sets are taken from [15] and [38]. All sets have target variables with five ordinal classes.

Dataset	#Instances	#Strong	#Weak	#Irrelevant
Set 1	150	6	0	6
Set 2	150	0	6	6
Set 3	150	3	4	3
Set 4	256	6	6	6
Set 5	512	1	2	11
Set 6	200	1	20	0
Set 7	200	1	20	20
Set 8	1000	10	20	10
Set 9	1000	10	20	200

For evaluation, we use the F-measure to quantify the detection of the all relevant feature set found by our method (dubbed feature relevance interval - FRI)¹ with regard to the true all relevant features of the data.

Because of the lack of other feature selection methods in this context we emulate the behaviour of lasso [21] and the ElasticNet (EN) [25]. For that we utilize a cross-validated recursive feature elimination², using the ordinal regression model given by Equation 1 with an ElasticNet penalty and parameter p . The parameter p , controlling the ratio between the L_1 and L_2 norm of the EN model, is optimized with a search over the values $p \in \{0, 0.01, 0.1, 0.2, 0.5, 0.7, 1\}$. Setting $p = 0$ corresponds to a lasso like sparsity constraint, and we test that scenario explicitly. Our surrogates are called $M_e^{L_1}$ (lasso) and $M_e^{L_1+L_2}$ (EN), both based on the explicit variant.

Hyper parameters are selected according to 5-fold cross validation, and all scores are averaged over 30 independent runs.

The results are given in Table 2, where FRI_e and FRI_i denote the explicit and the implicit variant respectively. Because lasso and EN performed nearly identical we only give the results for the EN.

The results show, that FRI in both variants is superior to $M_e^{L_1+L_2}$ on every data set, especially for clean data where it scores nearly perfect on every measure. It only shows slightly worse precision in Set 9 where the feature space is big. $M_e^{L_1+L_2}$ on the other hand, is very precise in that setting, but selects only 37% of relevant features. Having shown that, we are now interested in which of the two FRI variants is performing better. Since they both score perfectly on clean data, we increase the challenge by adding Gaussian noise with a standard deviation of $\sigma = 0.5$ to all sets. The theory, as given in [22], indicates that the implicit variant should perform better on noisy data, because for every decision boundary

¹ Implementation in Python: <https://github.com/lpfann/fri>

²Implementation in Python: RFECV from scikit-learn

to be determined it has access to more data samples than the explicit variant, thus gaining an advantage with regard to stability. However, our experiments do not support this notion as both variants of FRI perform equally well on noisy data. Interestingly, the $M_e^{L_1+L_2}$ improved its performance on those sets with a lot of weakly relevant features. This could be explained by assuming that the model has to rely on more of the weak, thus inter-correlated features, to regain the information that was lost due to the introduction of the noise.

Table 2: Artificially created data sets with known ground truth and evaluation of the identified relevant features by the methods as compared to all relevant features. The data was generated and Gaussian noise (standard deviation $\sigma = 0.5$) was added to the predictors. The score is averaged over 30 independent runs. $M_e^{L_1+L_2}$ represents the surrogate model for the ElasticNet with RFECV.

<i>Metric</i>	<i>Dataset</i>	Clean			Noise		
		$M_e^{L_1+L_2}$	FRI_e	FRI_i	$M_e^{L_1+L_2}$	FRI_e	FRI_i
F1	Set 1	0.94	1.0	1.0	0.92	0.95	0.98
	Set 2	0.79	1.0	1.0	0.89	0.97	0.98
	Set 3	0.81	1.0	1.0	0.85	0.97	0.96
	Set 4	0.83	1.0	1.0	0.80	0.96	0.97
	Set 5	0.83	1.0	1.0	0.86	1.0	1.0
	Set 6	0.25	1.0	1.0	0.56	0.94	0.94
	Set 7	0.49	1.0	1.0	0.46	0.90	0.91
	Set 8	0.95	1.0	1.0	0.80	0.98	0.98
	Set 9	0.53	0.98	0.98	0.60	1.0	1.0
Precision	Set 1	0.90	1.0	1.0	0.87	1.0	1.0
	Set 2	0.86	1.0	1.0	0.86	1.0	1.0
	Set 3	0.95	1.0	1.0	0.90	1.0	1.0
	Set 4	0.95	1.0	1.0	0.91	1.0	1.0
	Set 5	0.89	1.0	1.0	0.81	1.0	1.0
	Set 6	1.0	1.0	1.0	1.0	1.0	1.0
	Set 7	0.97	1.0	1.0	0.84	1.0	1.0
	Set 8	0.91	1.0	1.0	0.95	1.0	1.0
	Set 9	1.0	0.97	0.97	1.0	1.0	1.0
Recall	Set 1	1.0	1.0	1.0	0.99	0.92	0.96
	Set 2	0.82	1.0	1.0	0.94	0.96	0.96
	Set 3	0.74	1.0	1.0	0.83	0.95	0.93
	Set 4	0.77	1.0	1.0	0.74	0.93	0.94
	Set 5	0.84	1.0	1.0	0.99	1.0	1.0
	Set 6	0.15	1.0	1.0	0.40	0.89	0.89
	Set 7	0.41	1.0	1.0	0.35	0.84	0.86
	Set 8	1.0	1.0	1.0	0.70	0.97	0.97
	Set 9	0.37	1.0	1.0	0.43	1.0	1.0

6.1.2. Benchmark Data

Here, we purely evaluate the model performance on benchmark data as described in [25, 39] without regarding feature selection. The imbalanced ordinal regression data sets used in the experiments are listed in Table 3. All samples are normalized to zero mean and unit variance.

Table 3: Real ordinal regression benchmark data sets with imbalanced classes taken from [39], where d is the number of features, and K is the number of classes.

Dataset	# Instances	d	K	Ordered Class Distribution
Automobile	205	71	6	(3,22,67,54,32,27)
Bondrate	57	37	5	(6,33,12,5,1)
Contact-lenses	24	6	3	(15,5,4)
Eucalyptus	736	91	5	(180,107,130,214,105)
Newthyroid	215	5	3	(30,150,35)
Pasture	36	25	3	(12,12,12)
Squash-stored	52	51	3	(23,21,8)
Squash-unstored	52	52	3	(24,24,4)
TAE	151	54	3	(49,50,52)
Winequality-red	1599	11	6	(10,53,681,638,199,18)

We replicate the experiments which have been presented in [5, 6] to evaluate the performance of our two possible underlying SVM models as stated in Section 2. Our models, which we will call $M_e^{L_1}$ and $M_i^{L_1}$ in the following, were tuned using 5-fold cross-validation and used all available features previous feature selection, i.e. the models do not use the procedure described in 5 and the scores are based on all features without retraining. The results are averaged over the same 30 folds as used in [6] and evaluation is based on the MMAE as defined in Equation 21. We compare our models with p-OGMLVQ and a-OGMLVQ, the best performing methods for the given data as stated in [5]. Results for the ElasticNet surrogate $M_e^{L_1+L_2}$ were omitted because they were nearly identical to $M_e^{L_1}$.

The outcomes are reported in Table 4. Overall the explicit variant $M_e^{L_1}$ outperforms the implicit variant $M_i^{L_1}$ in all cases except one when considering MMAE. Similarly, the runtime of $M_e^{L_1}$ is at least two times faster, in some cases even over 20 times faster. When comparing with the existing results of a-OGMLVQ, we can see $M_e^{L_1}$ outperforming it in 5 cases while being worse in 5 others, it can beat p-OGMLVQ in 6 cases and closely ties in one case (TAE).

Table 4: Comparison of both proposed variants of ordinal regression models from Section 2. Benchmark on real ordinal datasets [39] by averaged MMAE and aggregated run time over 30 folds. Folds were identical to [6] and are comparable.

	MMAE				Run time	
	p-OGMLVQ	a-OGMLVQ	$M_e^{L_1}$	$M_i^{L_1}$	$M_e^{L_1}$	$M_i^{L_1}$
Automobile	0.482	0.446	0.532	0.516	151.6	876.8
Bondrate	0.768	0.737	0.939	0.949	49.7	133.6
Contact-lenses	0.243	0.221	0.190	0.265	23.7	53.9
Eucalyptus	0.450	0.477	0.390	0.390	768.7	3280.3
Newthyroid	0.124	0.097	0.043	0.045	37.5	92.3
Pasture	0.307	0.318	0.374	0.430	28.6	57.0
Squash-stored	0.415	0.411	0.371	0.371	36.0	68.9
Squash-unstored	0.488	0.228	0.280	0.300	35.9	69.4
TAE	0.553	0.537	0.552	0.664	43.3	83.4
Winequality-red	1.078	1.069	0.868	0.790	349.4	8359.4

With regard to feature relevance, no ground truth is available for the given data, rendering us unable to perform the same evaluation as for the artificial sets. We are only able to compare the amount of features provided by our method with feature selection (FRI) and the previously used model $M_e^{L_1+L_2}$ as a surrogate for EN with RFECV. Table 5 lists the average number of features identified as relevant for both techniques. For three data sets (Squash-stored, Squash-unstored, TAE), FRI identifies a smaller number of relevant features than the alternative, while yielding the same accuracy. For three further data sets (Automobile, Eucalyptus, Pasture), FRI identifies more (weakly relevant) features. In all cases, FRI potentially offers more information than EN by discriminating between weakly and strongly relevant features, and giving more candidate features to consider which can then be verified in practise.

Table 5: Mean feature set size of FRI model with explicit constraints and EN surrogate model ($M_e^{L_1+L_2}$) with RFECV on real datasets [25, 39]. FRI allows extra discrimination between strong (FRI^s) relevance and weak (FRI^w) relevance.

	Average Feature Set Size			
	FRI_e^s		FRI_e^w	$M_e^{L_1+L_2}$
Automobile	4.5	∪	12.6	4.0
Bondrate	0.0	∪	5.4	2.0
Contact-lenses	0.9	∪	1.1	2.0
Eucalyptus	2.1	∪	33.2	15.6
Newthyroid	0.0	∪	4.7	2.0
Pasture	0.0	∪	15.5	6.0
Squash-stored	2.4	∪	7.9	11.1
Squash-unstored	1.8	∪	3.3	8.0
TAE	1.9	∪	5.4	16.8
Winequality-red	0.0	∪	7.6	5.4

6.1.3. COMPAS Analysis

To showcase a possible application of our approach, we use FRI to examine the COMPAS dataset. This data was created by Propublica, a journalistic collective from New York, and consists of personal information regarding the criminal history of 11757 people from Broward County in Florida. Data like this has been used to predict an individual's risk of recidivism after a criminal offence. Hereby, previous analyses have shown [40] that racial bias is incorporated in at least one standard algorithmic prediction tool, meaning that African American individuals receive higher risk scores than Caucasian people. While it still remains an open research question if and how an algorithm should use socially sensitive attributes [41, 42] we are now interested which information is used by our linear ordinal regression model based on the FRI analysis on the given data. As such we try to find possible causes for direct or indirect discrimination [43] and facilitate careful model design, which seems to be necessary when aiming for long term impact of fair machine learning[44].

From the originally 28 features of the dataset, we scale down to ten by eliminating all identifying and time related information, which do not contribute information to the prediction task. These features are described in detail in Appendix C. We build a predictive model on the data, showing the relevancy of our features to that model. The result is shown in the upper plot in Figure 1. In this kind of plot, the relevance intervals are shown as vertical bars such that the maximum and minimum heights represent **maxrel** and **minrel**. For better comparison the values are normalized to the L_1 norm of the optimal model ($\|\tilde{\mathbf{w}}\|_1$). We also add the maximum element in $\Pi(\text{maxrel})$ as horizontal dashes, which represents the threshold which is used to classify between weakly relevant and irrelevant features.

The predictive accuracy is 66.73% which is directly inside the range of accuracies discussed in the Propublica analysis - note that the models used in practice deviate from the ones considered here, and the former are not available to us. Thus, we discuss properties of the linear models found by the proposed ORP only, not any other model. Two features are strongly relevant, namely, the count of prior charges and the age group 17-25 which show a big contribution in absolute terms. Many other features, such as the count of juvenile felonies and misdemeanors, or the degree of criminal changes are weakly relevant. More interestingly, socially sensitive features such as the sex and race are also considered weakly relevant. In the case of sex, both male and female exhibit the same maximal relevance which hints at the anti-correlation between the two features. In the case of race, being African-American, Caucasian or Native American is considered weakly relevant. When compared with the Propublica analysis, our relevance bounds are in line with their results.

To measure the contribution of the ethnic features in the model, we repeat the experiment with all those features removed. Hereby, the accuracy does not drop significantly, yielding 65.99%. The bottom plot of Figure 1 shows the relevance for all remaining features. Compared to the previous model, there are two notable changes. The count of juvenile offences and the information about

violent recidivism become relevant which are intuitively much more important to the problem at hand and do not reiterate a potential bias in society.




Figure 1: Relevance plots for the COMPASS dataset. Top: Relevance intervals (bars) for all features including ethnicity. Bottom: Relevance intervals for all features when ethnicity is eliminated from the data. Ethnicity is not a relevant factor for the model on top, so if those variables are eliminated, the relevancy of the other features do not change profoundly. The y-axis represents the computed feature relevance normalized to the L_1 norm of the optimal model.

6.2. Privileged Information

The following section evaluates our approach for the LUPI paradigm, i.e. our method handling privileged information, that we denote FRI^* . From here, we focus on the explicit variant, after showing its superiority over the implicit version in section 6.1.2 as regards computational complexity, leading to the notation FRI_e^* . Again, we show the quality of our feature selection by testing

on artificially created data with known ground truth. Due to a lack of specific LUPI benchmark datasets, we conclude our paper with a semantic analysis of a FRI_e^* model on one demonstrative example.

6.2.1. Artificial Data

We use the generation method presented in [28] to create artificial datasets containing regular as well as privileged information by sampling triplets (x_i, x_i^*, y_i) from:

$$\begin{aligned} x_i^* &\sim \mathcal{N}(0, I_d) \\ \varepsilon_i &\sim \mathcal{N}(0, I_d) \\ x_i &\leftarrow x_i^* + \varepsilon \\ y_i &\leftarrow f(\langle \omega, x_i^* \rangle), \end{aligned}$$

where f denotes a function that assigns the correct ordinal bin to the label y_i based on the value of the dot product between the weight vector and a privileged sample x_i^* .

Hereby, the privileged information X^* consists of clean versions of the noisy regular features X . Both, the regular and the privileged feature space, contain strong, weak and irrelevant features. These are created in the same way as described in section 6.1.1. The characteristics of the data used in our experiments are shown in Table 6. The last two sets differ from the generation method mentioned above. Their regular information is created similarly to the sets in Table 1, to which three irrelevant privileged features are added from random Gaussian noise. All features are normalized to zero mean and unit variance.

Table 6: Artificially created data with regular and privileged features under known ground truth. For the first six sets, the privileged features consist of clean versions of the regular information. The last two sets are regular ordinal regression sets with random noise as additional privileged information.

Dataset	#Instances	Regular Features			Privileged Features		
		#Str	#Weak	#Irr	#Str	#Weak	#Irr
Set 1	200	6	0	3	6	0	3
Set 2	200	0	12	3	0	12	3
Set 3	200	6	6	0	6	6	0
Set 4	200	3	6	0	3	6	0
Set 5	200	1	4	0	1	4	0
Set 6	200	1	40	10	1	40	10
Set 7	200	4	2	2	0	0	3
Set 8	200	0	4	2	0	0	3

Evaluation closely follows section 6.1.1. Again, we use the F-measure as a quantifying metric for the detection of the all relevant features set, and compare our method to the EN surrogate model $M_e^{L_1+L_2}$. While FRI_e^* differentiates

between the two feature spaces in the data, the EN receives both the regular and the privileged set as one. With that, we want to showcase the advantages of a LUPI model for feature selection over a purely regular model.

The results are given in Table 7. FRI_e^* achieves a perfect score on the regular feature set and only stumbles once, for set 6, on the privileged information. The EN on the other hand, performs considerably worse on the regular set but shows significant improvements on the privileged set, albeit it cannot match the performance of our method. The improvements on the privileged data are easy to explain since this information is the clear original information as opposed to the noisy features in the regular set.

Table 7: Artificially created datasets with known ground truth and evaluation of the identified relevant features by the methods as compared to all existing relevant features. The EN surrogate model ($M_e^{L_1+L_2}$) receives both feature sets as one but evaluation is done separately for the regular and privileged feature set. The score is averaged over 10 independent runs.

<i>Metric</i>	<i>Dataset</i>	Regular Features		Privileged Features	
		$M_e^{L_1+L_2}$	FRI_e^*	$M_e^{L_1+L_2}$	FRI_e^*
F1	Set 1	0.44	1.0	0.89	1.0
	Set 2	0.48	1.0	0.85	1.0
	Set 3	0.65	1.0	0.91	1.0
	Set 4	0.58	1.0	0.88	1.0
	Set 5	0.67	1.0	0.92	1.0
	Set 6	0.40	1.0	0.69	0.99
	Set 7	0.93	1.0	1.0	1.0
	Set 8	0.70	1.0	1.0	1.0
Precision	Set 1	0.72	1.0	0.91	1.0
	Set 2	0.75	1.0	0.98	1.0
	Set 3	1.0	1.0	1.0	1.0
	Set 4	0.90	1.0	1.0	1.0
	Set 5	0.80	1.0	1.0	1.0
	Set 6	0.98	1.0	0.97	1.0
	Set 7	0.94	1.0	1.0	1.0
	Set 8	1.0	1.0	1.0	1.0
Recall	Set 1	0.37	1.0	0.88	1.0
	Set 2	0.38	1.0	0.78	1.0
	Set 3	0.52	1.0	0.84	1.0
	Set 4	0.48	1.0	0.80	1.0
	Set 5	0.62	1.0	0.88	1.0
	Set 6	0.26	0.99	0.54	0.98
	Set 7	0.93	1.0	1.0	1.0
	Set 8	0.55	1.0	1.0	1.0

6.2.2. Semantic Analysis

Performing evaluations similar to sections 6.1.2 and 6.1.3 is not possible because of the lack of public LUPI benchmark. Therefore, we consider one illus-

trative example to demonstrate the semantic implications of the FRI framework for LUPI. We generate a set with 400 samples and six features. Initially, there are three strongly relevant features and three irrelevant ones drawn from random Gaussian noise. We divide the samples into four groups, each with 100 members. The first group has Gaussian noise with a standard deviation of 0.1 added to the first strongly relevant feature. The second group has a noise level of 0.5 added to the second feature. Similarly, the third one, has Gaussian noise on the last strong feature with a standard deviation of 2. The data in the last group is noise free. The idea is to provide the insight which samples of the dataset are hard to classify as privileged information to the model. Therefore, the privileged set consists of three features, incorporating the noise that was added to the groups, with the first privileged feature corresponding to the first group and so on.

The plots in Figure 2 show the relevancy for the regular features (a) as well as for the privileged features (b). Our method correctly dismissed the three irrelevant features and also classifies all strongly relevant features. More importantly, all privileged features were also correctly classified, and their relevance correlates with the noise level. With that, we show that FRI_e^* can discriminate between the usefulness of multiple privileged features and utilize those that are necessary in this setting.




Figure 2: Relevance plots for the semantic analysis. (a) Relevance of the regular features for the LUPI model. (b): Relevance of the privileged features for the LUPI model.

7. Conclusions

In this paper we presented the adaption of the feature relevance bounds approach to ordinal regression data using the *explicit* order variant. The optimization problem was phrased by approximating the generalization ability of the model with a bound on the L_1 -margin. The resulting problem can be transferred to a linear problem. For its solution, we used another approximation by splitting the objective into the margin and slack variables separately, for larger

robustness. Further, we proposed a resampling-based procedure to determine which values correspond to no information of the features, to automatically set situation-dependent thresholds. Based on the experiments we showed that the *explicit* variant is comparable to the *implicit* variant for this use case on the given data as regards the accuracy and more efficient. Our method can provide a near perfect all-relevant feature set approximation while being significantly faster than the other variant. Although not many feature selection approaches exist for that specific context we could also showcase the feature selection performance in comparison with another popular approach on toy and real data. The feature sets produced by our approach represents additional information useful in analytic use cases for model and experiment design, subject for further evaluation, and it constitutes a possible starting point to investigate, e.g. the information which restricted or protected features can provide for the class of linear ORP models.

Furthermore, we also provided a definition for feature relevance bounds when additional information is present in the context of learning using privileged information. Here we defined a features relevant in relation to the training phase itself. Similar to the classical context, our method achieved very good feature selection sensitivity in both the regular and privileged feature set, this way enabling a strategy to choose suitable features or teacher information to facilitate training.

References

- [1] F. M. Harper, J. A. Konstan, The movielens datasets: History and context, *ACM transactions on interactive intelligent systems* 5 (4) (2016) 19.
- [2] E. Frank, M. Hall, A Simple Approach to Ordinal Classification, in: L. De Raedt, P. Flach (Eds.), *Machine Learning: ECML 2001*, 2001, pp. 145–156.
- [3] A. Shashua, A. Levin, Ranking with Large Margin Principle: Two Approaches, in: *Proceedings of the 15th International Conference on Neural Information Processing Systems*, NIPS’02, MIT Press, 2002, pp. 961–968.
URL <http://dl.acm.org/citation.cfm?id=2968618.2968738>
- [4] W. Chu, S. S. Keerthi, Support Vector Ordinal Regression, *Neural Comput.* 19 (3) (2007-03) 792–815. doi:10.1162/neco.2007.19.3.792.
- [5] S. Fouad, P. Tiño, Adaptive Metric Learning Vector Quantization for Ordinal Classification, *Neural computation* 24 11 (2012) 2825–51.
- [6] F. Tang, P. Tiño, Ordinal regression based on learning vector quantization, *Neural Networks* 93 (2017) 76–88. doi:10.1016/j.neunet.2017.05.006.
- [7] C. Li, M. de Rijke, Incremental sparse Bayesian ordinal regression 106 294–302. doi:10.1016/j.neunet.2018.07.015.
URL <https://linkinghub.elsevier.com/retrieve/pii/S0893608018302144>
- [8] W. Cao, V. Mirjalili, S. Raschka, Consistent rank logits for ordinal regression with convolutional neural networks, *CoRR abs/1901.07884* (2019). arXiv:1901.07884.
URL <http://arxiv.org/abs/1901.07884>
- [9] F. Pedregosa, F. R. Bach, A. Gramfort, On the consistency of ordinal regression methods, *CoRR abs/1408.2327* (2014). arXiv:1408.2327.
URL <http://arxiv.org/abs/1408.2327>

- [10] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, D. Pedreschi, A survey of methods for explaining black box models, *ACM Comput. Surv.* 51 (5) (2018) 93:1–93:42. doi:10.1145/3236009. URL <http://doi.acm.org/10.1145/3236009>
- [11] X. Geng, T.-Y. Liu, T. Qin, H. Li, Feature Selection for Ranking, in: *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '07*, ACM, 2007, pp. 407–414. doi:10.1145/1277741.1277811.
- [12] S. Baccianella, A. Esuli, F. Sebastiani, Feature Selection for Ordinal Regression, in: *Proceedings of the 2010 ACM Symposium on Applied Computing, SAC '10*, ACM, 2010, pp. 1748–1754. doi:10.1145/1774088.1774461.
- [13] L. Yu, H. Liu, Efficient feature selection via analysis of relevance and redundancy, *J. Mach. Learn. Res.* 5 (2004) 1205–1224. URL <http://dl.acm.org/citation.cfm?id=1005332.1044700>
- [14] M. B. Kursa, W. R. Rudnicki, Feature Selection with the **Boruta** Package, *Journal of Statistical Software* 36 (11) (2010). doi:10.18637/jss.v036.i11.
- [15] C. Göpfert, L. Pfannschmidt, J. P. Göpfert, B. Hammer, Interpretation of Linear Classifiers by Means of Feature Relevance Bounds, *Neurocomputing* (accepted).
- [16] M. Kearns, Fair algorithms for machine learning, in: *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17*, ACM, New York, NY, USA, 2017, pp. 1–1. doi:10.1145/3033274.3084096. URL <http://doi.acm.org/10.1145/3033274.3084096>
- [17] V. Vapnik, A. Vashist, A new learning paradigm: Learning using privileged information, *Neural Networks* 22 (5) (2009-07-01) 544–557. doi:10.1016/j.neunet.2009.06.042.
- [18] V. Vapnik, R. Izmailov, Learning using privileged information: Similarity control and knowledge transfer., *Journal of machine learning research* 16 (2023-2049) (2015) 2.
- [19] F. Tang, P. Tino, P. Gutierrez, H. Chen, The Benefits of Modelling Slack Variables in SVMs, *Neural computation* 27 (4) (2015) 954–981.
- [20] D. Pechyony, V. Vapnik, On the theory of learning with privileged information, in: *Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010. Proceedings of a meeting held 6-9 December 2010, Vancouver, British Columbia, Canada., 2010*, pp. 1894–1902.
- [21] R. Tibshirani, Regression Shrinkage and Selection via the Lasso, *Journal of the Royal Statistical Society. Series B (Methodological)* 58 (1) (1996) 267–288. arXiv:2346178.
- [22] W. Chu, S. S. Keerthi, New Approaches to Support Vector Ordinal Regression, in: *Proceedings of the 22nd International Conference on Machine Learning, ACM, 2005*, pp. 145–152.
- [23] S. Agarwal, Generalization Bounds for Some Ordinal Regression Algorithms, in: *ALT, 2008*.
- [24] T. Zhang, Covering number bounds of certain regularized linear function classes, *Journal of Machine Learning Research* 2 (2002) 527–550.
- [25] H. Zou, T. Hastie, Regularization and variable selection via the elastic net, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67 (2) (2005) 301–320. doi:10.1111/j.1467-9868.2005.00503.x.

- [26] P. Zhao, B. Yu, On Model Selection Consistency of Lasso, *J. Mach. Learn. Res.* 7 (2006-12) 2541–2563.
URL <http://dl.acm.org/citation.cfm?id=1248547.1248637>
- [27] S. Shalev-Shwartz, S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*, Cambridge University Press, 2014.
- [28] D. Lopez-Paz, L. Bottou, B. Schölkopf, V. Vapnik, Unifying distillation and privileged information, in: *ICLR*, 2016.
URL <https://arxiv.org/abs/1511.03643v3>
- [29] S. Geisser, *Predictive Inference*, CRC Press, 1993-06-01.
- [30] D. François, F. Rossi, V. Wertz, M. Verleysen, Resampling methods for parameter-free and robust feature selection with mutual information 70 (7) 1276–1288. doi:10.1016/j.neucom.2006.11.019.
URL <http://www.sciencedirect.com/science/article/pii/S0925231206004796>
- [31] J. C. Platt, *Sequential Minimal Optimization: A Fast Algorithm for Training Support Vector Machines*.
- [32] N. Karmarkar, A new polynomial-time algorithm for linear programming 4 (4) 373–395. doi:10.1007/BF02579150.
URL <https://doi.org/10.1007/BF02579150>
- [33] P. M. Vaidya, Speeding-up linear programming using fast matrix multiplication, in: *30th Annual Symposium on Foundations of Computer Science*, pp. 332–337. doi:10.1109/SFCS.1989.63499.
- [34] M. B. Cohen, Y. T. Lee, Z. Song, Solving Linear Programs in the Current Matrix Multiplication Time [arXiv:1810.07896](https://arxiv.org/abs/1810.07896).
URL <http://arxiv.org/abs/1810.07896>
- [35] T. Joachims, Training linear SVMs in linear time, in: *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '06*, ACM Press, p. 217. doi:10.1145/1150402.1150429.
URL <http://portal.acm.org/citation.cfm?doid=1150402.1150429>
- [36] C.-J. Hsieh, K.-W. Chang, C.-J. Lin, S. S. Keerthi, S. Sundararajan, A dual coordinate descent method for large-scale linear SVM, in: *Proceedings of the 25th International Conference on Machine Learning - ICML '08*, ACM Press, pp. 408–415. doi:10.1145/1390156.1390208.
URL <http://portal.acm.org/citation.cfm?doid=1390156.1390208>
- [37] L. Pfannschmidt, C. Göpfert, U. Neumann, D. Heider, B. Hammer, FRI - Feature Relevance Intervals for Interpretable and Interactive Data Exploration.
URL <https://pub.uni-bielefeld.de/record/2935456>
- [38] C. Göpfert, L. Pfannschmidt, B. Hammer, Feature Relevance Bounds for Linear Classification, in: M. Verleysen (Ed.), *Proceedings of the ESANN, 24th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, Ciaco - i6doc.com, 2017, pp. 187–192.
- [39] J. Sánchez-Monedero, Gutiérrez, P. Antonio, P. Tino, C. Hervás-Martínez, Exploitation of Pairwise Class Distances for Ordinal Classification, *Neural Computation* 25 (9) (2013).
- [40] J. Angwin, J. Larson, S. Mattu, L. Kirchner, *How we analyzed the compas recidivism algorithm* (2016).
URL <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>

- [41] S. Hajian, J. Domingo-Ferrer, A Methodology for Direct and Indirect Discrimination Prevention in Data Mining, IEEE Transactions on Knowledge and Data Engineering 25 (7) (2013-07) 1445–1459. doi:10.1109/TKDE.2012.72.
- [42] M. Hardt, E. Price, N. Srebro, Equality of Opportunity in Supervised Learning (2016-10-07). arXiv:1610.02413.
URL <http://arxiv.org/abs/1610.02413>
- [43] D. Pedreschi, S. Ruggieri, F. Turini, Discrimination-aware Data Mining 9.
- [44] L. T. Liu, S. Dean, E. Rolf, M. Simchowitz, M. Hardt, Delayed Impact of Fair Machine Learning (2018-03-12). arXiv:1803.04383.
URL <http://arxiv.org/abs/1803.04383>
- [45] A. Domahidi, E. Chu, S. Boyd, ECOS: An SOCP solver for embedded systems, in: European Control Conference (ECC), pp. 3071–3076.

Appendix A. Proof of Theorem 2

We proof Theorem 2. We rely on Theorem 4 in [15], which states the following: Assume two optimization problems

$$\text{Problem A : } \min_x h_1(x) \text{ s.t. } x \in A_1$$

$$\text{Problem B : } \min_y h_2(y) \text{ s.t. } y \in A_2$$

Assume mappings $f : A_1 \rightarrow A_2$ and $g : A_2 \rightarrow A_1$ exist such that for all $x \in A_1, y \in A_2$

$$\begin{aligned} h_2(y) < h_2(f(x)) &\Rightarrow h_1(g(y)) < h_1(x) \\ h_1(x) < h_1(g(y)) &\Rightarrow h_2(f(x)) < h_2(y) \end{aligned}$$

Then the two problems A and B are equivalent in the sense that the mappings f and g establish direct correspondences of their global optima.

Equivalence of $\text{minrel}(\mathcal{I})$ and $\text{minrel}^(\mathcal{I})$*

Solutions of $\text{minrel}(\mathcal{I})$ have the form

$$\mathbf{w} = (w_1, \dots, w_n), b = (b_1, \dots, b_{l-1}), \chi = (\chi_1^1, \dots, \chi_{m_l-1}^{l-1}), \xi = (\xi_1^2, \dots, \xi_{m_l}^l)$$

$\text{minrel}^*(\mathcal{I})$ combines this form with an additional vector $\hat{\mathbf{w}} = (\hat{w}_1, \dots, \hat{w}_n)$. Define the mapping

$$f : (\mathbf{w}, b, \chi, \xi) \mapsto (\mathbf{w}, \hat{\mathbf{w}} = |\mathbf{w}| := (|w_1|, \dots, |w_n|), b, \chi, \xi)$$

Define the mapping

$$g : (\mathbf{w}, \hat{\mathbf{w}}, b, \chi, \xi) \mapsto (\mathbf{w}, b, \chi, \xi)$$

f is obviously a mapping in between feasible sets. The same holds for g , since the constraints (11) ensure $\sum_k \hat{w}_k \geq \|\mathbf{w}\|_1$.

Given an element of the feasible set of the two problems, denoted by $x := (\mathbf{w}^A, b^A, \chi^A, \xi^A)$ and $y := (\mathbf{w}^B, \hat{\mathbf{w}}^B, b^B, \chi^B, \xi^B)$, respectively. Assume $h_2(y) < h_2(f(x))$, i.e. $\hat{w}_T^B < |w_T^A|$. Then constraints (11) ensure $|w_T^B| \leq \hat{w}_T^B$, hence $|w_T^B| < |w_T^A|$, i.e. $h_1(g(y)) < h_1(x)$.

Conversely, $h_1(x) < h_1(g(y))$ implies $|w_T^A| < |w_T^B|$ hence constraints (11) ensure $|w_T^A| < \hat{w}_T^B$, i.e. $h_2(f(x)) < h_2(y)$.

Equivalence of $\max\text{rel}(\mathcal{I})$ and the optimum of $\max\text{rel}_{\text{pos}}^(\mathcal{I})$ and $\max\text{rel}_{\text{neg}}^*(\mathcal{I})$*

We consider two problems which are associated to $\max\text{rel}(\mathcal{I})$, $\mathbf{maxrel}_{\text{pos}}(\mathcal{I})$ equals $\max\text{rel}(\mathcal{I})$ with the additional constraint $w_{\mathcal{I}} \geq 0$, and $\mathbf{maxrel}_{\text{neg}}(\mathcal{I})$ equals $\max\text{rel}(\mathcal{I})$ with the additional constraint $w_{\mathcal{I}} \leq 0$. Since these two auxiliary problems decompose the feasible set of the original one into two halves, we can solve these two auxiliary problem and take whichever solution is best instead of solving $\max\text{rel}(\mathcal{I})$. Thus, we can show equivalence of these two sub problems to the versions as introduced in Theorem (2). Instead of maximization, we can focus on the minimization of the respective negative of the original objectives, to phrase the setting within the notation of Theorem 4 in [15].

We show equivalence of $\max\text{rel}_{\text{pos}}(\mathcal{I})$ and $\max\text{rel}_{\text{pos}}^*(\mathcal{I})$. Define the mapping f as identity for $(\mathbf{w}, b, \chi, \xi)$ and $\hat{\mathbf{w}} = |\mathbf{w}| := (|w_1|, \dots, |w_n|)$. Define the mapping g as projection of $(\mathbf{w}, \hat{\mathbf{w}}, b, \chi, \xi)$ onto all elements but $\hat{\mathbf{w}}$. f and g are obviously mappings in between the feasible sets. Note that constraints $w_{\mathcal{I}} \geq 0$ and $\hat{w}_{\mathcal{I}} \leq w_{\mathcal{I}}$ are required at this step.

Given elements of the feasible sets of the problems $x := (\mathbf{w}^A, b^A, \chi^A, \xi^A)$ and $y := (\mathbf{w}^B, \hat{\mathbf{w}}^B, b^B, \chi^B, \xi^B)$. Assume $h_2(y) < h_2(f(x))$, i.e. $-\hat{w}_{\mathcal{I}}^B < -|w_{\mathcal{I}}^A|$. Then the constraints (11) and (12) ensure $\hat{w}_{\mathcal{I}}^B = w_{\mathcal{I}}^B$ and $\hat{w}_{\mathcal{I}}^B \geq 0$, hence $-|w_{\mathcal{I}}^B| < -|w_{\mathcal{I}}^A|$, i.e. $h_1(g(y)) < h_1(x)$.

Conversely, $h_1(x) < h_1(g(y))$ implies $-|w_{\mathcal{I}}^A| < -|w_{\mathcal{I}}^B|$ hence $-|w_{\mathcal{I}}^A| < -\hat{w}_{\mathcal{I}}^B$, i.e. $h_2(f(x)) < h_2(y)$ due to constraints (11) and (12).

Similarly, equivalence of $\max\text{rel}_{\text{neg}}(\mathcal{I})$ and $\max\text{rel}_{\text{neg}}^*(\mathcal{I})$ can be shown. f and g are as above. These are mappings in between feasible sets.

f and g are obviously mappings in between the feasible sets. Note that constraints $w_{\mathcal{I}} \geq 0$ and $\hat{w}_{\mathcal{I}} \leq w_{\mathcal{I}}$ are required at this step.

Appendix B. Feature Relevance Bounds for Ordinal Regression with Implicit Order

In the following we are defining the relevance bounds for the implicit variant from Section 2. The definition is very similar to Section 3, and the following will be very concise.

Assume a training set X . Denote an optimum solution of problem (3) as $(\tilde{\mathbf{w}}, \tilde{b}_j, \tilde{\xi}_i^j, \tilde{\chi}_i^j)$. We define

$$L := \sum_{j=1}^{l-1} \left(\sum_{k=1}^j \sum_{i=1}^{n^k} \chi_{ki}^j + \sum_{k=j+1}^l \sum_{i=1}^{n^k} \xi_{ki}^j \right)$$

as the sum of all slack variables. The optimum solution induces the value

$$\mu_X := \frac{1}{2} \|\tilde{\mathbf{w}}\|_1 + C \cdot L$$

which is uniquely determined by X .

The class of equivalent good hypotheses is proxied by

$$F_{\delta}(X) := \left\{ \mathbf{w} \in \mathbb{R}^n \mid \exists \xi, \chi, \mathbf{b} \text{ such that constraints in (3) hold, } \right. \\ \left. \frac{1}{2} \|\mathbf{w}\|_1 + C \cdot L \leq (1 + \delta) \cdot \mu_X \right\}$$

Problem minrel(\mathcal{I}):

$$\begin{aligned} & \min_{\mathbf{w}, \mathbf{b}, \chi, \xi} && |w_{\mathcal{I}}| && \text{(B.1)} \\ \text{s.t. for all } i, j &&& \text{conditions in (3) hold} && \\ &&& \frac{1}{2} \|\mathbf{w}\|_1 + C \cdot L \leq (1 + \delta) \cdot \mu_X && \text{(B.2)} \end{aligned}$$

Problem maxrel(\mathcal{I}):

$$\begin{aligned} & \max_{\mathbf{w}, \mathbf{b}, \chi, \xi} && |w_{\mathcal{I}}| && \text{(B.3)} \\ \text{s.t. for all } i, j &&& \text{conditions in (3) hold} && \\ &&& \frac{1}{2} \|\mathbf{w}\|_1 + C \cdot L \leq (1 + \delta) \cdot \mu_X && \text{(B.4)} \end{aligned}$$




Figure B.3: Plot of runtime with respect to number of instances (a) and number of features (b). Additionally, both show comparison between single thread (1 CPU) and multi threaded run (8 CPU).

As before, this problem can be equivalently phrased as a linear program.

Appendix B.1. Scaling

Here we evaluate the scaling of our implementation in the setting without privileged information. We already discussed the theoretical time complexity bounds in Section 5.1 where we concluded that the overall method with feature selection is in $\mathcal{O}(n^3 + (dz + c)n^{2.5})$. We now run two separate experiments where we generate artificial sets as described earlier and scale up their size with respect to the number of instances n and number of features d . In the first experiment we set $d = 20$ and scale n between 10 and 10000 and in the second we set $n = 500$ and scale d between 10 and 500. Our implementation is using the high level library *cvxpy*³ and the ECOS solver [45] and presents no specific adaption for the problems at hand. The implementation runtime is measured on a modern Intel Xeon processor. Additionally, because relevance bounds can be computed in parallel, we run both experiments with one single thread and 8 threads in parallel.

In Figure B.3 results for both experiments are given. One can see that the complexity can limit the application of the method to small to medium-sized problems. This is in line with other *all* relevant feature selection methods [37] which exhibit much higher runtimes than simple sparse methods. While slightly bigger sets with, e.g. $n > 10^4$ or $d > 500$ are feasible, multiprocessing is recommended. For bigger data sets, further optimization or filtering of the feature space is necessary.

Appendix C. Features of the COMPAS dataset

This section gives a description of all the features of the COMPAS dataset that we use in our analysis in 6.1.3. The features are listed in Table C.8. All categorical variables are One-Hot-Coded for the analysis. The ethnicities are one of {African-American, Caucasian, Hispanic, Asian, Native American, Other}, the sexes are male or female, the age is grouped into {less than 25, 25-45, greater than 45} and the charge can be one of {felony, misdemeanour, offence}. The total number of features fed into the first model is 20. After eliminating all ethnic information, the count reduces to 14.

³<https://www.cvxpy.org>

Table C.8: Description of features of the COMPAS dataset used for the analysis in 6.1.3.

Feature Name	Type	Description	One-Hot
Juv_fel_count	Numerical	# Felonies as a juvenile	No
Juv_misd_count	Numerical	# Misdemeanour as a juvenile	No
Juv_other_count	Numerical	# Offences as a juvenile	No
Priors_count	Numerical	# Prior convictions	No
Is_recid	Binary	If recidivism happened	No
Is_violent_recid	Binary	If violent recidivism happened	No
Ethnicity	Categorical	One of 6 ethnicities	Yes
Sex	Categorical	One of 2 sexes	Yes
Age	Categorical	One of 3 age groups	Yes
C_charge	Categorical	One of 3 charge groups	Yes