



Technische Universität Darmstadt
Knowledge Engineering Group
Hochschulstrasse 10, D-64289 Darmstadt, Germany



<http://www.ke.informatik.tu-darmstadt.de>

Technical Report TUD-KE-2007-04

Eyke Hüllermeier, Johannes Fürnkranz

**On Minimizing the Position Error
in Label Ranking**

A short version of this paper appeared as:
Eyke Hüllermeier, Johannes Fürnkranz. On Minimizing the Position Error in Label Ranking,
Proceedings of the 18th European Conference on Machine Learning (ECML-07),
Springer-Verlag, 2007.

On Minimizing the Position Error in Label Ranking

Eyke Hüllermeier

Department of Mathematics and Computer Science, Marburg University

EYKE@MATHEMATIK.UNI-MARBURG.DE

Johannes Fürnkranz

Department of Computer Science, TU Darmstadt

JUFFI@KE.INFORMATIK.TU-DARMSTADT.DE

Abstract

Conventional classification learning allows a classifier to make a one shot decision in order to identify the correct label. However, in many practical applications, the problem is not to give a single estimation, but to make repeated suggestions until the correct target label has been identified. Thus, the learner has to deliver a label ranking, that is, a ranking of all possible alternatives. In this paper, we discuss a loss function, called the position error, which is suitable for evaluating the performance of a label ranking algorithm in this setting. Moreover, we propose “ranking through iterated choice”, a general strategy for extending any multi-class classifier to this scenario. Its basic idea is to reduce label ranking to standard classification by successively predicting a most likely class label and retraining a model on the remaining classes. We demonstrate empirically that this procedure does indeed reduce the position error in comparison with a conventional approach that ranks the classes according to their estimated probabilities. Besides, we also address the issue of implementing ranking through iterated choice in a computationally efficient way.

1. Introduction

The main interest in the context of classification learning typically concerns the correctness of a prediction: A prediction is either correct or not and, correspondingly, is rewarded in the former and punished in the latter case. The arguably best-known loss function reflecting this problem conception is the misclassification or error rate of a classifier, that is, the probability of making an incorrect prediction. In many cases, however, error rate and variants thereof are not ideal performance measures. Just to give an example, in *ordinal* classification, where the class labels are taken from a totally ordered scale, the classification rate does not distinguish between a prediction A which is wrong and a prediction B which is even more wrong since, according to the ordering of the labels, A is between B and the true target label [10]. This is related to cost-sensitive learning, with the difference that one is not given numerical cost values but only knows that certain misclassifications are more expensive than others.

In this paper, we are interested in another scenario which motivates a generalization of the misclassification rate. As an illustration, consider a fault detection problem which consists of identifying the cause for the malfunctioning of a technical system. Suppose that a classifier has been trained to predict the true cause, e.g., on the basis of certain sensor measurements serving as input attributes (see, e.g., [1] for an application of that type). Now, if it turned out that a predicted cause is not correct, one cannot simply say that the

classification process terminated with a failure. Instead, since the cause must eventually be found, alternative candidates must be tried until the problem is fixed.

What is needed in applications of this type is not only a prediction in the form of a single class label but instead a *ranking* of all candidate labels. In fact, a ranking suggests a simple (trial and error) search process, which successively tests the candidates, one by one, until the correct cause is found. An obvious measure of the quality of a predicted ranking is a loss function that counts the number of futile trials made before the target label is identified. In Section 2 of this paper, we shall look at measures of that type in more detail.

Apart from a suitable loss function, one needs a learner that produces label rankings as outputs. In this regard, the most obvious idea is to use a scoring classifier which outputs a score for each label, which is then used for sorting the labels. In particular, one may use a probabilistic classifier that estimates, for every candidate label λ , the conditional probability of λ given the input \vec{x} . Intuitively, *probabilistic ranking* (PR), i.e., ordering the labels according to their respective probabilities of being the target label, appears to be a reasonable approach.

In Section 3, we show that this approach is indeed optimal in a particular sense. Despite this theoretical optimality, however, an implementation of the approach turns out to be intricate in practice, mainly because estimating conditional probabilities is a difficult problem. In fact, it is well-known that most classification algorithms commonly used in the field of machine learning do not produce accurate probability estimates, even though they may have a strong hit rate. This motivates an alternative approach, to be introduced in Section 3.2, that we call *ranking through iterated choice* (RIC). The idea of this method is to employ a (multi-class) classifier as a choice function which, given a set of candidate labels and related training data, selects the most promising among these candidates. Roughly speaking, a label ranking is then obtained by repeated classification: In every iteration, the learning algorithm removes this label, and retrains a classifier for the remaining labels. Due to the retraining, RIC obviously comes along with an increased complexity. To overcome this problem, an efficient implementation of this approach, which is based on pairwise decomposition techniques, is proposed in Section 3.3. Experimental results, showing that RIC does indeed improve accuracy in comparison with PR, are presented in Section 4. Finally, we conclude the paper with a brief summary.

2. Label Ranking and Position Error

In this section, we will first formally define the label ranking problem (Section 2.1), discuss position error (Section 2.2) and some extensions (Section 2.3) as loss functions for this type of problem. We also briefly recapitulate related work in ranking (Section 2.4).

2.1 Label Ranking

We consider a learning problem which involves an input space \mathcal{X} and an output set $\mathcal{L} = \{\lambda_1 \dots \lambda_m\}$ consisting of a finite number of class labels. Assuming $\mathcal{X} \times \mathcal{L}$ to be endowed with a probability measure, one can associate a vector

$$p_{\vec{x}} = (\mathbb{P}(\lambda_1 | \vec{x}) \dots \mathbb{P}(\lambda_m | \vec{x})) \tag{1}$$

of conditional class probabilities with every input $\vec{x} \in \mathcal{X}$, where $\mathbb{P}(\lambda_i | \vec{x}) = \mathbb{P}(\lambda_i = \lambda_{\vec{x}})$ denotes the probability that \vec{x} belongs to class λ_i .

Given a set of training examples $\mathcal{D} = \{(\vec{x}_1, \lambda_{\vec{x}_1}) \dots (\vec{x}_n, \lambda_{\vec{x}_n})\} \subset (\mathcal{X} \times \mathcal{L})^n$, the learning problem is to induce a “label ranker”, which is a function that maps any input \vec{x} to a total order of the class labels, i.e., a complete, transitive, and asymmetric relation $\succ_{\vec{x}}$ on \mathcal{L} ; here, $\lambda_i \succ_{\vec{x}} \lambda_j$ means that λ_i precedes λ_j in the ranking associated with \vec{x} . Formally, a ranking $\succ_{\vec{x}}$ can be identified with a permutation $\tau_{\vec{x}}$ of $\{1 \dots m\}$, e.g., the permutation $\tau_{\vec{x}}$ satisfying

$$\lambda_{\tau_{\vec{x}}^{-1}(1)} \succ_{\vec{x}} \lambda_{\tau_{\vec{x}}^{-1}(2)} \succ_{\vec{x}} \dots \succ_{\vec{x}} \lambda_{\tau_{\vec{x}}^{-1}(m)}. \quad (2)$$

Here, $\tau_{\vec{x}}(i) = \tau_{\vec{x}}(\lambda_i)$ is the position of label λ_i in the ranking.

2.2 The Position Error

So far, we introduced the problem of predicting a label ranking in a formal way, but did not discuss the semantics of a predicted ranking. In fact, one should realize that such a ranking can serve different purposes. Needless to say, this point is of major importance for the evaluation of a predicted ranking.

In hitherto existing approaches to label ranking [7, 5], the quality of a prediction is measured in terms of a similarity or distance measure for rankings; for example, a commonly used measure for comparing a predicted ranking (permutation) $\tau_{\vec{x}}$ and a true ranking $\tau_{\vec{x}}^*$ is the Spearman rank correlation. Measures of that type take the position of *all* labels into account, which means, e.g., that swapping the positions of the two bottom labels is as bad as swapping the positions of the two top labels.

Measures such as Spearman rank correlation quantify, say, the *ranking error* of a prediction. In this paper, we are interested in an alternative type of measure, which is especially motivated by practical performance tasks where a prediction is used in order to support the search for a true target label. As outlined in the introduction, an obvious loss function in this context is the number of labels preceding that label in the predicted ranking. Subsequently, a deviation of the predicted target label’s position from the top-rank will be called a *position error*. Note that, while a ranking error relates to the comparison of two complete label rankings $\tau_{\vec{x}}$ and $\tau_{\vec{x}}^*$, the position error refers to the comparison of a label ranking $\tau_{\vec{x}}$ and a true class $\lambda_{\vec{x}}$.

More specifically, we define the position error of a prediction $\tau_{\vec{x}}$ as

$$\text{PE}(\tau_{\vec{x}}, \lambda_{\vec{x}}) \stackrel{\text{df}}{=} \tau_{\vec{x}}(\lambda_{\vec{x}}),$$

i.e., by the position of the target label $\lambda_{\vec{x}}$ in the ranking $\tau_{\vec{x}}$. To compare the quality of rankings of different problems, it is useful to normalize the position error for the number of labels. This *normalized position error* is defined as

$$\text{NPE}(\tau_{\vec{x}}, \lambda_{\vec{x}}) \stackrel{\text{df}}{=} \frac{\tau_{\vec{x}}(\lambda_{\vec{x}}) - 1}{m - 1} \in \{0, 1/(m - 1) \dots 1\}. \quad (3)$$

The position error of a label ranker is the *expected* position error of its predictions, where the expectation is taken with respect to the underlying probability measure on $\mathcal{X} \times \mathcal{L}$.

Compared with the conventional misclassification rate, the position error differentiates between “bad” predictions in a more subtle way: In the case of a correct classification, both

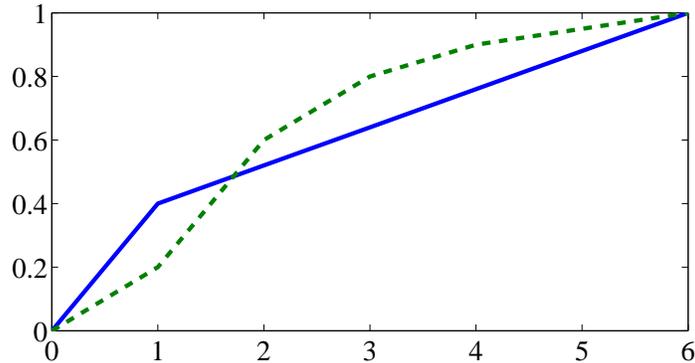


Figure 1: Exemplary C-distributions for two classifiers.

measures coincide. In the case of a wrong top label, however, the misclassification rate is 1, while the position error assumes values between 1 and m , depending on how “far away” the true target label actually is.

As most performance measures, the position error is a simple scalar index. To characterize a label ranking algorithm in a more elaborate way, an interesting alternative is to look at the mapping

$$C : \{1 \dots m\} \rightarrow \mathbb{R}, k \mapsto \mathbb{P}(\tau_{\vec{x}}(\lambda_{\vec{x}}) \leq k). \quad (4)$$

According to its definition, $C(k)$ is the probability that the target label is among the top k labels in the predicted ranking. In particular, $C(1)$ corresponds to the conventional classification rate. Moreover, $C(\cdot)$ is monotone increasing, and $C(m) = 1$. Formally, the mapping (4) is nothing else than the cumulative probability function of a random variable, namely the position of the target label, and the position error is the corresponding expected value. Of course, on the basis of the C -distribution (4), only a partial order can be defined on a class of learning algorithms: Two learners are incomparable in the case of intersecting C -distributions. Fig. 1 shows an example of that kind. The first learner (solid curve) is a good classifier, as it has a high classification rate. Compared with the second learner (dashed curve), however, it is not a good ranker, as its C -distribution has a rather flat slope.

2.3 Extensions of the Position Error

Depending on the concrete application scenario, various extensions and generalizations of the position error are conceivable. For example, imagine that, for what reason ever, it is essential to have the target label among the top k candidates in a predicted ranking. This goal might be captured most easily by means of a simple 0/1-loss function which yields 0 if $\text{PE}(\tau_{\vec{x}}, \lambda_{\vec{x}}) \leq k$ and 1 otherwise. More generally, one can use a (non-decreasing) weight function $w : \{1 \dots m\} \rightarrow \mathbb{R}$ and define the following weighted (transformed) version of the position error:

$$\text{PE}(\tau_{\vec{x}}, \lambda_{\vec{x}}) \stackrel{\text{df}}{=} w(\tau_{\vec{x}}(\lambda_{\vec{x}})).$$

Obviously, the standard position error is obtained for $w(i) \equiv i$, while the above 0/1-loss is recovered for the special case where $w(\cdot)$ is given by $w(i) = 0$ for $i \leq k$ and $w(i) = 1$ for $i > k$. Moreover, the normalized position error (3) can be modeled by the weighting scheme $w(i) = (i - 1)/m$.

Another interesting extension is related to the idea of *cost-sensitive* classification. In this respect, one usually associates a cost-value with each pair of class labels (λ_i, λ_j) , reflecting the cost incurred when erroneously predicting λ_i instead of λ_j . In the context of our scenario, it makes sense to associate a cost value $c(i)$ with each individual label λ_i , reflecting the cost for *verifying* whether or not λ_i is the correct label. Then, the cost induced by a predicted label ranking $\tau_{\vec{x}}$ is given by the cost for testing the target label plus the labels preceding the target in the ranking:

$$\sum_{i=1}^{\tau_{\vec{x}}(\lambda_{\vec{x}})} c(\tau_{\vec{x}}^{-1}(i)).$$

Finally, we note that the idea of the position error can of course be generalized to multi-label (classification) problems which assume several instead of a single target label for each instance. In this connection, it makes a great difference whether one is interested in having at least one of the targets on a top rank (e.g., since finding one solution is enough), or whether all of them should have high positions (resp. none of them should be ranked low). In the latter case, it makes sense to count the number of non-target labels ranked above the lowest target label (an application of that type has recently been studied in [3]), while in the former case, one will look at the number of non-target labels placed before the highest-ranked target label.

2.4 Related Work on Ranking

Before proceeding, we briefly comment on related work on ranking in the field of machine learning, especially since the term “ranking” is used in different ways in the literature. As mentioned previously, the problem of *label ranking* has already been addressed in other works, albeit within a different problem setting (minimizing a ranking error versus minimizing a position error). Even though we use the same name for both settings, they should be well distinguished [8].

Alternatively, the term “ranking” has been used in the sense of *object ranking* or “learning to order things” (see, e.g., [2]). Here, the problem is to learn a function that, given a subset of objects \mathcal{O} from a reference set \mathcal{X} as input, outputs a ranking of these objects. Training data typically consists of exemplary pairwise preferences of the form $\vec{x} \succ \vec{x}'$, with $\vec{x}, \vec{x}' \in \mathcal{X}$. As a key difference between object and label ranking, note that the latter associates a ranking of a fixed number of labels with every instance \vec{x} , whereas the former is interested in ranking the instances themselves.

Yet another use of the term “ranking”, which is less related to this work, refers to ordering of classified examples (in the dichotomous case) according to the scores they received [9]: The higher (lower) the score, the more safely an example can be classified as positive (negative). A good ranker, in this sense, is a classifier that well separates between positive and negative examples, i.e., that systematically ranks positive examples ahead of negative ones. This type of ranking performance has been studied in the context of ROC analysis.

3. Minimizing the Position Error

What kind of ranking procedure should be used in order to minimize the risk of a predicted ranking with respect to the position error as a loss function? As mentioned before, an intuitively plausible idea is to order the candidate labels λ according to their probability $\mathbb{P}(\lambda = \lambda_{\vec{x}})$ of being the target label. In fact, this idea is not only plausible but also provably correct. Even though the result is quite obvious, we state its formally as a theorem.

Theorem 1 *Given a query instance $\vec{x} \in \mathcal{X}$, ranking the labels $\lambda \in \mathcal{L}$ according to their (conditional) probabilities of being the target class $\lambda_{\vec{x}}$ yields a risk minimizing prediction with respect to the position error (3) as a loss function. That is, the expected loss*

$$\mathbb{E}(\tau_{\vec{x}}) = \frac{1}{m-1} \sum_{i=1}^m (i-1) \cdot \mathbb{P}(\tau_{\vec{x}}(\lambda_{\vec{x}}) = i)$$

becomes minimal for any ranking $\succ_{\vec{x}}$ such that $\mathbb{P}(\lambda_i = \lambda_{\vec{x}}) > \mathbb{P}(\lambda_j = \lambda_{\vec{x}})$ implies $\lambda_i \succ_{\vec{x}} \lambda_j$.

Proof: This result follows almost by definition. In fact, note that we have

$$\mathbb{E}(\tau_{\vec{x}}) \propto \sum_{i=1}^m \mathbb{P}(\tau_{\vec{x}}(\lambda_{\vec{x}}) > i)$$

and that, for each position i , the probability to exceed this position when searching for the target $\lambda_{\vec{x}}$ is obviously minimized when ordering the labels according to their (conditional) probabilities. □ □

3.1 Conventional Conditioning

According to the above result, the top rank (first position) should be given to the label λ_{\top} for which the estimated probability is maximal. Regarding the second rank, recall the fault detection metaphor, where the second hypothesis for the cause of the fault is only tested in case the first one turned out to be wrong. Thus, when having to make the next choice, one principally has additional information, namely that λ_{\top} *was not the correct label*. Taking this information into account, the second rank should not simply be given to the label with the second highest probability according to the original probability measure, say, $\mathbb{P}_1(\cdot) = \mathbb{P}(\cdot)$, but instead to the label that maximizes the *conditional* probability $\mathbb{P}_2(\cdot) = \mathbb{P}(\cdot | \lambda_{\vec{x}} \neq \lambda_{\top})$ of being the target label given that the first proposal was incorrect.

At first sight, passing from $\mathbb{P}_1(\cdot)$ to $\mathbb{P}_2(\cdot)$ may appear meaningless from a ranking point of view, since standard probabilistic conditioning yields

$$\mathbb{P}_2(\lambda) = \frac{\mathbb{P}_1(\lambda)}{1 - \mathbb{P}_1(\lambda_{\top})} \propto \mathbb{P}_1(\lambda) \tag{5}$$

for $\lambda \neq \lambda_{\top}$, and therefore does not change the order of the remaining labels. And indeed, in case the original $\mathbb{P}(\cdot)$ is a proper probability measure and conditioning is performed according to (5), the predicted ranking will not be changed at all.

3.2 Empirical Conditioning

One should realize, however, that standard conditioning is not an incontestable updating procedure in our context, simply because $\mathbb{P}_1(\cdot)$ is not a “true” probability measure over the class labels. Rather, it is only an estimated measure coming from a learning algorithm, perhaps one which is not a good probability estimator. In fact, it is well-known that most machine learning algorithms for classification perform rather poorly in probability estimation, even though they may produce good classifiers. Thus, it seems sensible to perform “conditioning” not on the measure itself, but rather on the learner that produced the measure. What we mean by this is retraining the learner on the original data without the λ_\top -examples, an idea that could be paraphrased as “empirical conditioning”.

This type of conditioning depends on the data \mathcal{D} and the model assumptions, that is, the hypothesis space \mathcal{H} from which the classifier is taken. To emphasize this dependence and, moreover, to indicate that it concerns an *estimated* (“hat”) probability, the conditional measure $\mathbb{P}_2(\cdot)$ could be written more explicitly as

$$\mathbb{P}_2(\cdot) = \widehat{\mathbb{P}}(\cdot | \lambda_{\vec{x}} \neq \lambda_\top, \mathcal{D}, \mathcal{H}).$$

To motivate the idea of empirical conditioning, suppose that the estimated probabilities come from a classification tree. Of course, the original tree trained with the complete data will be highly influenced by λ_\top -examples, and the probabilities assigned by that tree to the alternatives $\lambda \neq \lambda_\top$ might be inaccurate. Retraining a classification tree on a reduced set of data might then lead to more accurate probabilities for the remaining labels, especially since the multi-class problem to be solved has now become simpler (as it involves fewer classes).

Fig. 2 shows a simpler example, where the hypothesis space \mathcal{H} is given by the class of decision stumps (univariate decision trees with only one inner node, i.e., axis-parallel splits in the case of numerical attributes). Given the examples from three classes (represented, respectively, by the symbols \diamond , \star , and \bullet), the best model corresponds to the split shown in the left picture. By estimating probabilities through relative frequencies in the leaf nodes of the decision stump, one derives the following estimates for the query instance, which is marked by a \oplus symbol: $\widehat{\mathbb{P}}(\diamond | \oplus) = 12/15$, $\widehat{\mathbb{P}}(\star | \oplus) = 2/15$, $\widehat{\mathbb{P}}(\bullet | \oplus) = 1/15$; thus, the induced ranking is given by $\diamond \succ \star \succ \bullet$. Now, suppose that the top label \diamond turned out to be an incorrect prediction. According to the above ranking (and probabilistic conditioning), the next label to be tested would be \star . However, when fitting a new model to the training data without the \diamond -examples, the preference between \star and \bullet is reversed, because the query instance is now located “on the \bullet -side” of the decision boundary. Roughly speaking, conditioning by “taking a different look” at the data, namely a look that suppresses the \diamond examples, gives a quite different picture (shown on the right-hand side of Fig. 2) of the situation. In fact, one should realize that, in the first model, the preference between \star and \bullet is strongly biased by the \diamond -examples: The first decision boundary is optimal only because it classifies all \diamond -examples correctly, a property that loses importance once it turned out that \diamond is not the true label of the query.

According to the above idea, a classifier is used as a *choice function*: Given a set of potential labels with corresponding training data (and a new query instance \vec{x}), it selects the most likely candidate among these labels. We refer to the process of successively selecting alternatives by estimating top-labels from (conditional) probability measures $\mathbb{P}_1(\cdot), \mathbb{P}_2(\cdot) \dots \mathbb{P}_m(\cdot)$

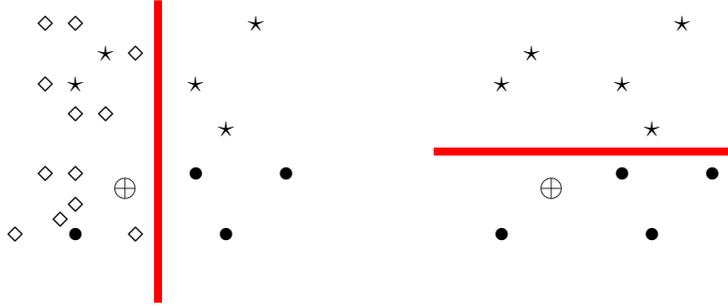


Figure 2: Example of empirical conditioning: The optimal model (decision stump) for the complete training data (left) and the data omitting the examples of the top label (\diamond).

as *ranking through iterated choice* (RIC). As an important advantage, note that this approach can be used to turn any multi-class classifier into a label ranker. In principle, it is not required that a corresponding classifier outputs a score, or even a real probability, for every label. In fact, since only a simple decision in favor of a single label has to be made in each iteration, any classifier is good enough. In this regard, let us note that, for the ease of exposition, the term “probability” will subsequently be used in a rather informal manner.

Regarding its effect on label ranking accuracy, one may expect the idea of RIC to produce two opposite effects:

- *Information loss*: In each iteration, the size of the data set to learn from becomes smaller.
- *Simplification*: Due to the reduced number of classes, the learning problems become simpler in each iteration.

The first effect will clearly have a negative influence on generalization performance, as a reduction of data comes along with a loss of information. In contrast to this, the second effect will have a positive influence: The classifiers will become increasingly simple, because it can be expected that the decision boundary for separating m classes is more complex than the decision boundary for separating $m' < m$ classes of the same problem. The hope is that, in practice, the second (positive) effect will dominate the first one.

3.3 Efficient Implementation

An obvious disadvantage of RIC concerns its computational complexity. In fact, since empirical conditioning essentially means classifying on a subset of \mathcal{L} , the number of models needed in (potentially) of the order $2^{|\mathcal{L}|}$. To overcome this problem, we propose the use of pairwise decomposition techniques.

The idea of pairwise learning is well-known in the context of classification [4], where it allows one to transform a polychotomous classification problem, i.e., a problem involving $m > 2$ classes $\mathcal{L} = \{\lambda_1 \dots \lambda_m\}$, into a number of *binary* problems. To this end, a separate

model (base learner) \mathcal{M}_{ij} is trained for each *pair* of labels $(\lambda_i, \lambda_j) \in \mathcal{L}$, $1 \leq i < j \leq m$; thus, a total number of $m(m-1)/2$ models is needed. \mathcal{M}_{ij} is intended to separate the objects with label λ_i from those having label λ_j . Depending on the classifier used, an output $\mathcal{M}_{ij}(\vec{x})$ can be interpreted, e.g., as the conditional probability

$$p_{ij} = \mathbb{P}(\lambda_{\vec{x}} = \lambda_i \mid \lambda_{\vec{x}} \in \{\lambda_i, \lambda_j\}, \vec{x})$$

In a second step, an estimation of the probability vector (1), i.e., of the individual probabilities $p_i = \mathbb{P}(\lambda_{\vec{x}} = \lambda_i \mid \vec{x})$, has to be derived from these pairwise probabilities. To this end, different techniques have been developed (e.g., [6]). Here, we resorted to the approach proposed in [11], which derives the p_i as a solution of a system of linear equations, S , that includes one equation for every label.

RIC can then be realized as follows: First, the aforementioned system of linear equations is solved, and the label λ_i with maximal probability p_i is chosen as the top-label λ_{\top} . This label is then removed, i.e., the corresponding variable p_i and its associated equation are deleted from S . To find the second best label, the same procedure is then applied to the reduced system S' thus obtained, i.e., by solving a system of $m-1$ linear equations and $m-1$ variables. This process is iterated until a full ranking has been constructed.

This approach reduces the training effort from an exponential to a quadratic number of models. Roughly speaking, a classifier on a subset $\mathcal{L}' \subseteq \mathcal{L}$ of classes is efficiently assembled “on the fly” from the corresponding subset of pairwise models $\{\mathcal{M}_{ij} \mid \lambda_i, \lambda_j \in \mathcal{L}'\}$. Or, stated differently, the *training* of classifiers is replaced by the *combination* of associated binary classifiers.

The hope that empirical conditioning improves accuracy in comparison with conventional probabilistic conditioning is essentially justified by the aforementioned *simplification effect* of RIC. Note that this simplification effect is also inherently present in pairwise learning. Here, the simplification due to a reduction of class labels is already achieved at the very beginning and, by decomposing the original problem into *binary* problems, carried to the extreme. Thus, if the simplification effect is indeed beneficial in the original version of RIC, it should also have a positive influence in the pairwise implementation (RIC-P). These are exactly the two conjectures to be investigated empirically in the next section: (i) Empirical conditioning (RIC) pays off with respect to accuracy, and (ii) the increased efficiency of the pairwise implementation, RIC-P, is achieved without sacrificing this gain in accuracy.

4. Empirical Results

In order to investigate the practical usefulness of empirical conditioning and the related RIC procedure, we compare the corresponding strategy to the most obvious alternative, namely ordering the class labels right away according to the respective probabilities produced by a multi-class classifier (probabilistic ranking, PR). So, given any multi-class classifier, capable of producing such probabilities, as a base learner, we consider the following three learning strategies:

- **PR:** A ranking is produced by applying the base learner to the complete data set only once and ordering the class labels according to their probabilities.

- **RIC:** This version refers to the *ranking through iterated choice* procedure outlined in Section 3.2, using the multi-class classifier as a base learner.
- **RIC-P:** This is the pairwise implementation of RIC as introduced in Section 3.3 (again using as base learners the same classifiers as RIC and PR).

In connection with selecting the top-label or ordering the labels according to their probability, ties are always broken through coin flipping.

Table 1 shows the results that we obtained for a number of benchmark data sets from the UCI repository and the StatLib archive¹, using two widely known machine learning algorithms as base learners: C4.5 and Ripper. For comparison purpose, we also derived results for the naive Bayes (NB) classifier, as this is one of the most commonly used “true” probabilistic classifiers. Note that, since conditional probabilities in NB are estimated individually for each class, empirical conditioning is essentially the same as conventional conditioning, i.e., RIC is equivalent to PR; this is why the results for RIC and RIC-P are omitted.

1. <http://www.ics.uci.edu/~mlearn>, <http://stat.cmu.edu/>

Table 1: Position error for conventional probabilistic ranking (PR), ranking through iterated choice (RIC), and its pairwise implementation (RIC-P), using C4.5, Ripper, and naive Bayes as base learners.

data	m	C4.5			Ripper			NB
		PR	RIC	RIC-P	PR	RIC	RIC-P	PR
abalone	28	4,650	4,004	3,552	4,667	4,358	3,500	4,346
anneal	6	1,023	1,028	1,024	1,031	1,028	1,017	1,150
audiology	24	2,310	2,186	3,190	2,394	3,274	3,270	3,102
autos	7	1,273	1,293	1,502	1,449	1,376	1,449	1,771
balance-scale	3	1,397	1,326	1,294	1,406	1,325	1,256	1,170
glass	7	1,547	1,486	1,449	1,612	1,486	1,463	1,855
heart-c	5	1,231	1,231	1,224	1,218	1,218	1,218	1,165
heart-h	5	1,197	1,197	1,197	1,187	1,187	1,187	1,16
hypothyroid	4	1,005	1,007	1,008	1,012	1,011	1,007	1,054
iris	3	1,073	1,053	1,053	1,067	1,073	1,073	1,047
lymph	4	1,270	1,250	1,236	1,284	1,277	1,297	1,189
primary-tumor	22	4,254	3,764	3,531	4,478	4,316	3,472	3,248
segment	7	1,135	1,042	1,042	1,131	1,075	1,060	1,258
soybean	19	1,205	1,113	1,085	1,220	1,123	1,073	1,136
vehicle	4	1,411	1,309	1,313	1,489	1,449	1,343	1,831
vowel	11	2,314	1,274	1,309	2,501	1,516	1,423	1,555
zoo	7	1,238	1,099	1,149	1,307	1,327	1,188	1,069
letter	26	2,407	1,279	1,202	2,168	1,375	1,188	2,515

Table 2: Win/loss statistics for each pair of methods, using C4.5 (left) and Ripper (right) as base learners.

	PR	RIC	RIC-P	PR	RIC	RIC-P
PR	—	3/13	4/13	—	3/13	3/12
RI	13/3	—	7/8	13/3	—	2/13
RIC-P	13/4	8/7	—	12/3	13/2	—

For each data set and each method we estimated the mean (absolute) position error using leave-one-out cross validation, except for the data set `letter`, for which we used the predefined separation into training and test data. The results are summarized in Table 1.

From the win-loss statistics for NB in comparison with PR using, respectively, C4.5 (10/8) and Ripper (10/8), there is no visible difference between these multi-class classifiers in terms of label ranking accuracy. Important are the win-loss statistics summarized in Table 2. These results perfectly support the two conjectures raised above. First, RIC significantly outperforms PR: According to a simple sign test for the win-loss statistic, the results are significant at a level of 2%. Second, RIP-P is fully competitive to RIC (and actually shows a better performance in the case of Ripper as a base learner).

5. Concluding Remarks

In the context of the label ranking problem, an interesting and practically relevant extension of conventional classification, we have discussed the position error as an alternative loss function. To minimize this loss function, we proposed *ranking through iterated choice* (RPC), a strategy that essentially reduces label ranking to repeated classification. In each iteration, RPC performs *empirical conditioning*, which in turn requires the retraining of classifiers. To avoid the need for training a potentially large number of models, we used a pairwise implementation in which retraining is done implicitly, namely by combining the outputs of certain pairwise models.

In an experimental study, RPC was compared to standard probabilistic ranking, where the class labels are ranked according to the originally estimated probabilities. Our results suggest that retraining (empirical conditioning) does indeed reduce the expected loss when using standard multi-class classifiers as base learners, and that this gain in accuracy is preserved by the pairwise implementation.

This work can be extended in various directions, both theoretically and practically. For example, one important aspect of future work is to generalize our framework to variants of the position error as outlined in Section 2.4.

Acknowledgements

This research is supported by the *German Science Foundation (DFG)*.

References

- [1] C. Alonso, J.J. Rodríguez, and B. Pulido. Enhancing consistency based diagnosis with machine learning techniques. In *Current Topics in AI*, vol. 3040 of LNAI, 312–321. Springer, 2004.
- [2] W.W. Cohen, R.E. Schapire, and Y. Singer. Learning to order things. *Journal of Artificial Intelligence Research*, 10:243–270, 1999.
- [3] K. Crammer and Y. Singer. A family of additive online algorithms for category ranking. *Journal of Machine Learning Research*, 3:1025–1058, 2003.
- [4] J. Fürnkranz. Round robin classification. *J. of Mach. Learn. Res.*, 2:721–747, 2002.
- [5] J. Fürnkranz and E. Hüllermeier. Pairwise preference learning and ranking. In *Proc. ECML-03*, Cavtat-Dubrovnik, Croatia, 2003.
- [6] T. Hastie and R. Tibshirani. Classification by pairwise coupling. In *Proc. NIPS-97*, pages 507–513, 1998.
- [7] S. Har-Peled, D. Roth, and D. Zimak. Constraint classification: a new approach to multiclass classification. In *Proc. ALT-02*, pp. 365–379, Lübeck, 2002.
- [8] E. Hüllermeier and J. Fürnkranz. Learning label preferences: Ranking error versus position error. In *Proc. IDA-2005*, Madrid, 2005.
- [9] F. Provost and P. Domingos. Tree induction for probabilistic ranking. *Machine Learning*, 52(3):199–215, 2003.
- [10] J.D.M. Rennie and N. Srebro. Loss Functions for Preference Levels: Regression with Discrete Ordered Labels. *Proc. IJCAI Multidisciplinary Workshop on Advances in Preference Handling*. 2005
- [11] T.F. Wu, C.J. Lin, and R.C. Weng. Probability estimates for multi-class classification by pairwise coupling. *J. Machine Learning Res.*, 5:975–1005, 2004.