

Preference Learning

Johannes Fürnkranz • Eyke Hüllermeier
Editors

Preference Learning

Editors

Prof. Dr. Johannes Fürnkranz
Knowledge Engineering Group
Fachbereich Informatik
Technische Universität Darmstadt
Hochschulstr. 10
64289 Darmstadt
Germany
juffi@ke.tu-darmstadt.de

Prof. Dr. Eyke Hüllermeier
Knowledge Engineering & Bioinformatics
Fachbereich Mathematik und Informatik
Philipps-Universität Marburg
Hans-Meerwein-Str.
35032 Marburg
Germany
eyke@mathematik.uni-marburg.de

ISBN 978-3-642-14124-9 e-ISBN 978-3-642-14125-6
DOI 10.1007/978-3-642-14125-6
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2010937568

ACM Computing Classification (1998): I.2.6, H.2.8

© Springer-Verlag Berlin Heidelberg 2010

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Cover design: KuenkelLopka GmbH

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

The topic of preferences has attracted considerable attention in Artificial Intelligence (AI) research in previous years. Recent special issues of the *AI Magazine* (December 2008) and the *Artificial Intelligence Journal* (announced for 2010), both devoted to preferences, highlight the increasing importance of this area for AI. Representing and processing knowledge in terms of preferences is appealing as it allows one to specify desires in a declarative way, to combine qualitative and quantitative modes of reasoning, and to deal with inconsistencies and exceptions in a quite flexible manner.

Like in other subfields of AI, including autonomous agents, nonmonotonic reasoning, constraint satisfaction, planning and qualitative decision theory, researchers in machine learning have started to pay increasing attention to the topic of preferences. In fact, as witnessed by a number of dedicated events, notably several workshops on preferences, ranking, and related topics (held, e.g., at NIPS 2004 and 2005, ECML/PKDD 2008 and 2009, SIGIR 2008 and 2009), we currently observe the formation of “preference learning” as a new branch of machine learning and data mining. For the time being, there is still no stipulated demarcation of this emerging subfield, neither in terms of a list of relevant topics nor in terms of an intentional “definition”. Roughly, *preference learning* refers to the problem of learning from observations which reveal, either explicitly or implicitly, information about the preferences of an individual (e.g., a user of a computer system) or a class of individuals; the acquisition of this kind of information can be supported by methods for *preference mining*. Generalizing beyond the training data given, the models learnt are typically used for preference prediction, i.e., to predict the preferences of a new individual or the same individual in a new situation. The problem of “learning to rank” is a good example and an important special case; here, the goal is to predict preferences in the form of total orders of a set of alternatives (e.g., a personalized ranking of documents retrieved by a search engine).

This book, which is the first volume specifically dedicated to the topic of preference learning, distinguishes itself through the following features:

- It gives a comprehensive overview of the state-of-the-art in the field of preference learning.

- By including a number of survey chapters, it offers an introduction to the most important subfields of preference learning.
- By proposing a systematic categorization according to learning task and learning technique, along with a unified notation, it helps structuring the field; thereby, it is supposed to have a positive impact on future research.
- Through the selection of contributions, it emphasizes the interdisciplinary character of preference learning and establishes connections to related research fields, such as multicriteria decision-making and operations research.
- Last but not least, it highlights important applications of preference learning in different areas, such as information retrieval and recommender systems, thereby demonstrating its practical relevance.

Some chapters of the book are based on contributions selected from two successful workshops on preference learning that we organized as part of the ECML/PKDD conferences in 2008 and 2009. Besides, however, the material is complemented by a number of chapters that have been solicited explicitly for this book. Overall, we are quite confident that the book provides both a broad coverage and comprehensive survey of the field of preference learning as well as a useful guideline and good introduction to the most important directions in current research.

The origination of this book is largely due to our close collaboration in recent years, which in turn has greatly benefited from a joint research project funded by the German Science Foundation (DFG). This support is gratefully acknowledged. Moreover, we would like to thank Ronan Nugent and the Springer for providing excellent assistance and ready advice during the final stages of preparation.

Darmstadt Marburg
December 2009

*Johannes Fürnkranz
Eyke Hüllermeier*

Contents

Preference Learning: An Introduction	1
Johannes Fürnkranz and Eyke Hüllermeier	
A Preference Optimization Based Unifying Framework for Supervised Learning Problems	19
Fabio Aioli and Alessandro Sperduti	
Part I Label Ranking	
Label Ranking Algorithms: A Survey	45
Shankar Vembu and Thomas Gärtner	
Preference Learning and Ranking by Pairwise Comparison	65
Johannes Fürnkranz and Eyke Hüllermeier	
Decision Tree Modeling for Ranking Data	83
Philip L.H. Yu, Wai Ming Wan, and Paul H. Lee	
Co-Regularized Least-Squares for Label Ranking	107
Evgeni Tsivtsivadze, Tapio Pahikkala, Jorma Boberg, Tapio Salakoski, and Tom Heskes	
Part II Instance Ranking	
A Survey on ROC-Based Ordinal Regression	127
Willem Waegeman and Bernard De Baets	
Ranking Cases with Classification Rules	155
Jianping Zhang, Jerzy W. Bala, Ali Hadjarian, and Brent Han	

Part III Object Ranking

A Survey and Empirical Comparison of Object Ranking

Methods181

Toshihiro Kamishima, Hideto Kazawa, and Shotaro Akaho

Dimension Reduction for Object Ranking203

Toshihiro Kamishima and Shotaro Akaho

Learning of Rule Ensembles for Multiple Attribute

Ranking Problems217

Krzysztof Dembczyński, Wojciech Kotłowski, Roman Słowiński,
and Marcin Szelag

Part IV Preferences in Multi-Attribute Domains

Learning Lexicographic Preference Models251

Fusun Yaman, Thomas J. Walsh, Michael L. Littman, and Marie
desJardins

Learning Ordinal Preferences on Multiattribute Domains:

The Case of CP-nets273

Yann Chevaleyre, Frédéric Koriche, Jérôme Lang, Jérôme Mengin,
and Bruno Zanuttini

Choice-Based Conjoint Analysis: Classification vs. Discrete

Choice Models297

Joachim Giesen, Klaus Mueller, Bilyana Taneva, and Peter Zolliker

Learning Aggregation Operators for Preference Modeling317

Vicenç Torra

Part V Preferences in Information Retrieval

Evaluating Search Engine Relevance

with Click-Based Metrics337

Filip Radlinski, Madhu Kurup, and Thorsten Joachims

Learning SVM Ranking Functions from User Feedback

Using Document Metadata and Active Learning

in the Biomedical Domain363

Robert Arens

Part VI Preferences in Recommender Systems

Learning Preference Models in Recommender Systems	387
Marco de Gemmis, Leo Iaquina, Pasquale Lops, Cataldo Musto, Fedelucio Narducci, and Giovanni Semeraro	
Collaborative Preference Learning	409
Alexandros Karatzoglou and Markus Weimer	
Discerning Relevant Model Features in a Content-Based Collaborative Recommender System	429
Alejandro Bellogín, Iván Cantador, Pablo Castells, and Álvaro Ortigosa	
Subject Index	457
Author Index	465

Preference Learning: An Introduction

Johannes Fürnkranz and Eyke Hüllermeier

Abstract This introduction gives a brief overview of the field of preference learning and, along the way, tries to establish a unified terminology. Special emphasis will be put on learning to rank, which is by now one of the most extensively studied problem tasks in preference learning and also prominently represented in this book. We propose a categorization of ranking problems into object ranking, instance ranking, and label ranking. Moreover, we introduce these scenarios in a formal way, discuss different ways in which the learning of ranking functions can be approached, and explain how the contributions collected in this book relate to this categorization. Finally, we also highlight some important applications of preference learning methods.

1 Introduction

Reasoning with preferences has been recognized as a particularly promising research direction for artificial intelligence (AI) [15]. A preference can be considered as a relaxed constraint which, if necessary, can be violated to some degree. In fact, an important advantage of a preference-based problem solving paradigm is an increased flexibility, as nicely explained in [6]:

“Early work in AI focused on the notion of a goal – an explicit target that must be achieved – and this paradigm is still dominant in AI problem solving. But as application domains become more complex and realistic, it is apparent that the dichotomic notion of a goal, while adequate for certain puzzles, is too crude in general. The problem is that in many contemporary application domains . . . the user has little knowledge about the set of possible solutions or feasible items, and what she typically seeks is the best that’s out there. But

J. Fürnkranz (✉)
Technical University Darmstadt, Germany
e-mail: juffi@ke.tu-darmstadt.de

E. Hüllermeier
Philipps-Universität Marburg, Germany
e-mail: eyke@mathematik.uni-marburg.de

since the user does not know what is the best achievable plan or the best available document or product, she typically cannot characterize it or its properties specifically. As a result, she will end up either asking for an unachievable goal, getting no solution in response, or asking for too little, obtaining a solution that can be substantially improved.”

Drawing on past research on knowledge representation and reasoning, AI offers qualitative and symbolic methods for treating preferences that can reasonably complement traditional approaches that have been developed for quite a while in fields such as economic decision theory [37]. Needless to say, however, the acquisition of preferences is not always an easy task. Therefore, not only modeling languages and representation formalisms, but also methods for the automatic learning, discovery, and adaptation of preferences are needed. For example, computerized methods for discovering the preferences of individuals are useful in e-commerce and various other fields, where an increasing trend toward personalization of products and services can be recognized.

It is hence hardly surprising that methods for learning and predicting preferences in an automatic way are among the very recent research topics in disciplines, such as machine learning, knowledge discovery, and recommender systems. Approaches relevant to this area range from approximating the utility function of a single agent on the basis of an as effective as possible question-answer process (often referred to as *preference elicitation*) to collaborative filtering where a customer’s preferences are estimated from the preferences of other customers. In fact, problems of preference learning can be formalized within various settings, depending, e.g., on the underlying type of preference model or the type of information provided as an input to the learning system.

Roughly speaking, *preference learning* is about inducing predictive preference models from empirical data. In the literature on choice and decision theory, two main approaches to modeling preferences can be found, namely in terms of *utility functions* and in terms of *preference relations*. From a machine learning point of view, these two approaches give rise to two kinds of learning problems: learning utility functions and learning preference relations. The latter deviates more strongly than the former from conventional problems such as classification and regression, as it involves the prediction of complex structures, such as rankings or partial order relations, rather than single values. Moreover, training input in preference learning will not, as it is usually the case in supervised learning, be offered in the form of complete examples but may comprise more general types of information, such as relative preferences or different kinds of indirect feedback and implicit preference information.

This book tries to give a comprehensive overview of the state-of-the-art in the field of preference learning. Some of its chapters are based on selected contributions to two successful workshops on this topic [29,30], but the material is complemented with chapters that have been solicited explicitly for this book. Most notably, several survey chapters give a detailed account on ongoing research in various subfields of preference learning. Thus, we are confident that the book succeeds in giving a comprehensive survey of work on all aspects of this emerging research area.

In the remainder of this chapter, we shall briefly sketch some important branches of preference learning and, along the way, give pointers to the contributions in this volume. References to these contributions are indicated by capitalized author names, for example FÜRNKRANZ & HÜLLERMEIER.

2 Preference Learning Tasks

Among the problems in the realm of preference learning, the task of “learning to rank” has probably received the most attention in the machine learning literature in recent years. In fact, a number of different ranking problems have been introduced so far, though a commonly accepted terminology has not yet been established. In the following, we propose a unifying and hopefully clarifying terminology for the most important types of ranking problems, which will also serve as a guideline for organizing the chapters of the book. AIOLLI & SPERDUTI give an alternative unifying framework for learning to rank from preferences.

In general, a preference learning task consists of some set of items for which preferences are known, and the task is to learn a function that predicts preferences for a new set of items, or for the same set of items in a different context. Frequently, the predicted preference relation is required to form a total order, in which case we also speak of a ranking problem. In this book, we will frequently use the term “ranking” for categorizing different types of preference learning problems, but we note that the characterization mainly depends on the form of the training data and the required predictions, and not on the fact that a total order is predicted.¹

In the notation used in the remainder of this chapter (and throughout most of the book), our goal is to stick as much as possible to the terminology commonly used in supervised learning (classification), where a data object typically consists of an *instance* (the input, also called predictive or independent variable in statistics) and an associated *class label* (the output, also called target or dependent variable in statistics). The former is normally denoted by \mathbf{x} , and the corresponding instance space by \mathcal{X} , while the output space is denoted by \mathcal{Y} . Instances are often represented in the form of feature vectors, which means that \mathbf{x} is a vector

$$\mathbf{x} = (x_1, x_2, \dots, x_m) \in \mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2 \times \dots \times \mathcal{X}_m.$$

We distinguish three types of ranking problems, namely label ranking, instance ranking, and object ranking, which are described in more detail in the following.

¹ Besides, one should be aware of conflicts between terminology in different fields. In the field of operations research, for example, the term “ranking” is used for arranging a set of objects in a total order, while “sorting” refers to the assignment of objects to an ordered set of categories, a problem closely related to what is called “ordered classification” in machine learning.

2.1 Label Ranking

In label ranking, we assume to be given an instance space \mathcal{X} and a finite set of labels $\mathcal{Y} = \{y_1, y_2, \dots, y_k\}$. The goal is to learn a “label ranker” in the form of an $\mathcal{X} \rightarrow S_{\mathcal{Y}}$ mapping, where the output space $S_{\mathcal{Y}}$ is given by the set of all total orders (permutations) of the set of labels \mathcal{Y} (the notation is leaned on the common notation S_k for the symmetric group of order k). Thus, label ranking can be seen as a generalization of conventional classification, where a complete ranking

$$y_{\pi_x^{-1}(1)} \succ_x y_{\pi_x^{-1}(2)} \succ_x \dots \succ_x y_{\pi_x^{-1}(k)}$$

is associated with an instance x instead of only a single class label. Here, π_x is a permutation of $\{1, 2, \dots, k\}$ such that $\pi_x(i)$ is the position of label y_i in the ranking associated with x .

The training data \mathcal{T} of a label ranker typically consist of a set of pairwise preferences of the form $y_i \succ_x y_j$, suggesting that, for instance x , y_i is preferred to y_j . In other words, an “observation” consists of an instance x and an ordered pair of labels (y_i, y_j) . The label ranking problem is summarized in Fig. 1.

This learning scenario has a large number of practical applications. For example, it is relevant for the prediction of every sort of ordering of a fixed set of elements, such as the preferential order of a fixed set of products (e.g., different types of holiday apartments) based on demographic properties of a person, or the ordering of a set of genes according to their expression level (as measured by microarray analysis) based on features of their phylogenetic profile [1]. Another application scenario is meta-learning, where the task is to rank learning algorithms according to their suitability for a new dataset, based on the characteristics of this dataset [7]. Finally, every preference statement in the well-known CP-nets approach [3], a qualitative graphical representation that reflects conditional dependence and independence of

Given:

- a set of training instances $\{x_\ell \mid \ell = 1, 2, \dots, n\} \subseteq \mathcal{X}$ (each instance typically though not necessarily represented by a feature vector)
- a set of labels $\mathcal{Y} = \{y_i \mid i = 1, 2, \dots, k\}$
- for each training instance x_ℓ : a set of *pairwise preferences* of the form $y_i \succ_{x_\ell} y_j$

Find:

- a ranking function that maps any $x \in \mathcal{X}$ to a ranking \succ_x of \mathcal{Y} (permutation $\pi_x \in S_k$)

Performance measures:

- ranking error (e.g., based on rank correlation measures) comparing predicted ranking with target ranking
 - position error comparing predicted ranking with a target label
-

Fig. 1 Label ranking

preferences under a *ceteris paribus* interpretation, formally corresponds to a label ranking.

In addition, it has been observed by several authors [14, 18, 25] that many conventional learning problems, such as classification and multilabel classification, may be formulated in terms of label preferences:

- *Classification*: A single class label y_i is assigned to each example \mathbf{x}_ℓ . This implicitly defines the set of preferences $\{y_i \succ_{\mathbf{x}_\ell} y_j \mid 1 \leq j \neq i \leq k\}$.
- *Multilabel classification*: Each training example \mathbf{x}_ℓ is associated with a subset $P_\ell \subseteq \mathcal{Y}$ of possible labels. This implicitly defines the set of preferences $\{y_i \succ_{\mathbf{x}_\ell} y_j \mid y_i \in L_\ell, y_j \in \mathcal{Y} \setminus P_\ell\}$.

A general framework encompassing these and other learning problems can be found in the chapter by AIOLLI & SPERDUTI.

In each of the former scenarios, a ranking model $f : \mathcal{X} \rightarrow \mathcal{S}_k$ is learned from a subset of all possible pairwise preferences. A suitable projection may be applied to the ranking model (which outputs permutations) as a post-processing step, for example, a projection to the top-rank in classification learning where only this label is relevant.

To measure the predictive performance of a label ranker, a loss function on rankings is needed. In principle, any distance or correlation measure on rankings (permutations) can be used for that purpose, for example, the number of pairs of labels that are incorrectly ordered (i.e., the number of label pairs y_i and y_j such that y_i precedes y_j in the predicted ranking although y_j is actually preferred to y_i). Apart from this type of *ranking loss*, which compares a predicted ranking with a given target ranking, it is also possible to compare a predicted ranking with a single class label. For example, if this class label is the target one is looking for, then it makes sense to evaluate a predicted ranking by the position it assigns to the label; in [28], this type of error (measuring the distance of the assigned position from the top-rank) is called the *position error*.

A general survey of label ranking is given by VEMBU & GÄRNTNER. Another discussion of label ranking and related problems is given by FÜRNKRANZ & HÜLLERMEIER. This chapter is specifically focused on approaches that are based on the idea of *learning by pairwise comparison*, i.e., of decomposing the original problem into a set of smaller binary classification problems. YU, WAN & LEE show how decision-tree learning algorithms such as CART can be adapted to tackle label ranking learning problems by extending the concept of purity to label ranking data. TSIVTSIVADZE et al. show how an approach for minimizing an approximation of a ranking loss function can be extended with a semi-supervised learning technique that tries to improve predictions by minimizing the disagreement of several ranking functions, which have been learned from different views of the training data.

2.2 Instance Ranking

This setting proceeds from the setting of *ordinal classification*, where an instance $\mathbf{x} \in \mathcal{X}$ belongs to one among a finite set of classes $\mathcal{Y} = \{y_1, y_2, \dots, y_k\}$ and,

moreover, the classes have a natural order: $y_1 < y_2 < \dots < y_k$. Training data consists of a set \mathcal{T} of labeled instances. As an example, consider the assignment of submitted papers to categories *reject*, *weak reject*, *weak accept*, and *accept*.

In contrast to the classification setting, the goal is not to learn a classifier but a ranking function $f(\cdot)$. Given a subset $X \subset \mathcal{X}$ of instances as an input, the function produces a ranking of these instances as an output (typically by assigning a score to each instance and then sorting by scores).

For the case $k = 2$, this problem is well-known as the bipartite ranking problem. The case $k > 2$ has recently been termed *k-partite* [42] or *multipartite ranking* [19]. As an example, consider the task of a reviewer who has to rank the papers according to their quality, possibly though not necessarily with the goal of partitioning this ranking into the above four categories.

Thus, the goal of *instance ranking* – our proposal for a generic term of bipartite and multipartite ranking – is to produce a ranking in which instances from higher classes precede those from lower classes; see Fig. 2 for a formalization of this task. Different types of accuracy measures have been proposed for predictions of this kind. Typically, they count the number of ranking errors, that is, the number of pairs $(\mathbf{x}, \mathbf{x}') \in X \times X$ such that \mathbf{x} is ranked higher than \mathbf{x}' even though the former belongs to a lower class than the latter. In the two-class case, this amounts to the well-known AUC, the area under the ROC-curve [5], which is equivalent to the Wilcoxon–Mann–Whitney statistic [38, 47]. Its generalization to multiple (ordered) classes is known as the concordance index or C-index in statistics [22].

These measures and the multipartite ranking scenario are discussed in more detail by WAEGERMAN & DE BAETS. ZHANG et al. discuss different methods for employing rule learning algorithms for learning bipartite rankings. This scenario has been studied for decision-tree learning, but not yet for rule learning, where several additional problems have to be considered, such as how to combine estimates from overlapping rules into a single probability estimate.

Given:

- a set of training instances $\{\mathbf{x}_\ell \mid \ell = 1, 2, \dots, n\} \subseteq \mathcal{X}$ (each instance typically though not necessarily represented by a feature vector)
- a set of labels $\mathcal{Y} = \{y_1, y_2, \dots, y_k\}$ endowed with an order $y_1 < y_2 < \dots < y_k$
- for each training instance \mathbf{x}_ℓ an associated label y_ℓ

Find:

- a ranking function that allows one to order a new set of instances $\{\mathbf{x}_j\}_{j=1}^t$ according to their (unknown) preference degrees

Performance measures:

- the area under the ROC-curve (AUC) in the dichotomous case ($k = 2$)
 - generalizations such as the C-index in the polychotomous case ($k > 2$)
-

Fig. 2 Instance ranking

2.3 Object Ranking

In the setting of object ranking, there is no supervision in the sense that no output or class label is associated with an object. The goal in object ranking is to learn a ranking function $f(\cdot)$ which, given a subset Z of an underlying referential set \mathcal{Z} of objects as an input, produces a ranking of these objects as an output. Again, this is typically done by assigning a score to each instance and then sorting by scores.

Objects $z \in \mathcal{Z}$ are commonly though not necessarily described in terms of an attribute-value representation. As training information, an object ranker has access to exemplary rankings or pairwise preferences of the form $z \succ z'$ suggesting that z should be ranked higher than z' . This scenario, summarized in Fig. 3, is also known as “learning to order things” [12].

As an example consider the problem of learning to rank query results of a search engine [33, 41]. The training information is provided implicitly by the user who clicks on some of the links in the query result and not on others. This information can be turned into binary preferences by assuming that the selected pages are preferred over nearby pages that are not clicked on [34].

The performance of an object ranker can again be measured in terms of a distance function or correlation measure on rankings. In contrast to the setting of label ranking, however, the number of items to be ordered in the context of object ranking is typically much larger. Therefore, one often prefers measures that put more emphasis on the top of a ranking while paying less attention to the bottom [17]. In Web search, for example, people normally look at the top-10 results while ignoring the rest. Besides, the target is often not a “true” ranking but instead a single object or a subset of relevant objects, for example a set of documents relevant to a query. Evaluation measures especially tailored toward these types of requirements have been proposed in information retrieval. Typical examples include precision and recall as well as normalized discounted cumulative gain (NDCG) [32, 39].

Given:

- a (potentially infinite) reference set of objects \mathcal{Z} (each object typically though not necessarily represented by a feature vector)
- a finite set of pairwise preferences $x_i \succ x_j, (x_i, x_j) \in \mathcal{Z} \times \mathcal{Z}$

Find:

- a ranking function $f(\cdot)$ that assumes as input a set of objects and returns a permutation (ranking) of this set

Performance measures:

- ranking error (e.g., based on rank correlation measures) comparing the predicted ranking with the target ranking
 - top-K measures comparing the top-positions of the rankings
 - retrieval measures such as precision, recall, NDCG
-

Fig. 3 Object ranking

An extensive survey of object ranking approaches is given by KAMISHIMA, KAZAWA & AKAHO. Subsequently, KAMISHIMA & AKAHO discuss dimensionality reduction methods for object ranking tasks, which retain the preference information as much as possible. They assume a scenario (which they call *supervised ordering*) in which total orders for multiple subsets of objects are given, and the goal is to predict an ordering of the full set of objects. DEMBCZYŃSKI et al. compare different approaches for rule-based learning of object ranking functions, namely one utility-based approach and one approach that directly learns the binary preference predicate (cf. also Sect. 3.3). An application to learning to rank documents in biomedical information retrieval is described by ARENS.

3 Preference Learning Techniques

All three of the basic learning tasks discussed in the previous section can be tackled by very similar basic techniques. In agreement with the distinction between using utility functions and binary relations for modeling preferences, two general approaches to preference learning have been proposed in the literature, the first of which is based on the idea of learning to evaluate individual alternatives by means of a utility function (Sect. 3.1), while the second one seeks to compare (pairs of) competing alternatives, that is, to learn one or more binary preference predicate (Sect. 3.2). Making sufficiently restrictive model assumptions about the structure of a preference relation, one can also try to use the data for identifying this structure (Sect. 3.3). Finally, local estimation techniques à la nearest neighbor can be used, which mostly leads to aggregating preferences in one way or the other (Sect. 3.4).

3.1 Learning Utility Functions

As mentioned previously, an established approach to modeling preferences resorts to the concept of a *utility function*. Such a function assigns an abstract degree of utility to each alternative under consideration. From a machine learning point of view, an obvious problem is to learn utility functions from given training data. Depending on the underlying utility scale, which is typically either numerical or ordinal, the problem becomes one of regression learning or ordered classification. Both problems are well-known in machine learning. However, utility functions often implicate special requirements and constraints that have to be taken into consideration such as, for example, monotonicity in certain attributes (DEMBZYŃSKI et al.).

Besides, as mentioned earlier, training data are not necessarily given in the form of input/output pairs, i.e., alternatives (instances) together with their utility degrees, but may also consist of qualitative feedback in the form of pairwise comparisons, stating that one alternative is preferred to another one and therefore has a higher utility degree. More generally, certain types of preference information can be

formalized in terms of constraints on one or more underlying utility functions. This idea forms the basis of the general framework presented by AIOLLI & SPERDUTI. Sometimes, of course, training data are less generic and more application-specific. In collaborative filtering, for example, it simply consists of an incomplete set of product ratings given by a set of users (see DE GEMMIS et al. in this volume).

In the instance and object preferences scenario, a utility function is a mapping $f : \mathcal{X} \rightarrow \mathbb{R}$ that assigns a utility degree $f(\mathbf{x})$ to each instance (object) \mathbf{x} and, hence, induces a complete order on \mathcal{X} . In the label preferences scenario, a utility function $f_i : \mathcal{X} \rightarrow \mathbb{R}$ is needed for each of the labels y_i ($i = 1, \dots, k$); alternatively, the functions can be summarized into a single function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ that maps instance/label tuples (\mathbf{x}, y) to real-valued scores (see AIOLLI & SPERDUTI).² Here, $f_i(\mathbf{x})$ is the utility assigned to alternative y_i by instance \mathbf{x} . To obtain a ranking for \mathbf{x} , the alternatives are sorted according to these utility scores, i.e., $\succ_{\mathbf{x}}$ is such that $y_i \succ_{\mathbf{x}} y_j \Rightarrow f_i(\mathbf{x}) \geq f_j(\mathbf{x})$.

In the setting of instance ranking, the training data consist of instances for which the sought utility scores are given. Thus, the learning problem can, in principle, be approached by means of classification or (ordinal) regression methods. As an important difference, however, note that the goal is not to maximize classification accuracy but ranking performance. Thus, conventional learning algorithms have to be adapted correspondingly. Approaches of this kind have, e.g., been proposed in [27, 33]. In this book, WAEGEMAN & DE BAETS discuss approaches that are based on the optimization of an extension of the binary AUC to a loss function for ordinal data.

In object and label ranking, training data typically originate from a kind of indirect supervision. Instead of the target scores of the utility function, the learner is given the constraints on the function, which are derived from comparative preference information of the form “This object (or label) should have a higher utility score than that object (or label)”. Thus, the challenge for the learner is to find a function which is as much as possible in agreement with these constraints.

For object ranking approaches, this idea has first been formalized by Tesauro under the name *comparison training* [44]. He proposed a symmetric neural network architecture that can be trained with representations of two states and a training signal that indicates which of the two states is preferable. The elegance of this approach comes from the property that one can replace the two symmetric components of the network with a single network, which can subsequently provide a real-valued evaluation of single states. Later works on learning utility function from object preference data include [24, 33, 46].

For the case of label ranking, a method for learning the functions $f_i(\cdot)$ ($i = 1, \dots, k$) has been proposed in the framework of *constraint classification* [25, 26]. Here, the authors proceed from linear utility functions and find a way to express a constraint of the form $f_i(\mathbf{x}) - f_j(\mathbf{x}) > 0$ (suggesting that $y_i \succ_{\mathbf{x}} y_j$) in the

² In a sense, this alternative is not just a formally equivalent rewriting. Instead, by considering an instance/label pair as an *object*, it suggests a natural way to unify the problems of object and label ranking.

form of a binary classification example in a newly constructed, higher-dimensional space. In other words, the original label ranking problem is transformed into a single binary classification problem. This problem is solved by fitting a separating hyperplane, and from this hyperplane, the linear utility functions (identified by corresponding weight vectors) can be reconstructed. An alternative approach, so-called *log-linear models for label ranking*, has been proposed in [14]. This approach is essentially equivalent to constraint classification, as it also amounts to learning linear utility functions for all labels. Algorithmically, however, the underlying optimization problem is approached in a different way, namely by means of a boosting-based algorithm that seeks to minimize a (generalized) ranking error in an iterative way. In this book, TSIVTSIVADZE et al. present an approach for learning a utility function for label ranking via minimization of a loss function that is based on a least-squares approximation of the ranking error.

3.2 Learning Preference Relations

The key idea of this approach is to learn a binary preference relation that compares pairs of alternatives (e.g., objects or labels). The training of a model thus becomes simpler, mainly because comparative training information (suggesting that one alternative is better than another one) can be used directly instead of translating it into constraints on a (latent) utility function. On the other hand, the prediction step may become more difficult, since a binary preference relation learned from data is not necessarily consistent in the sense of being transitive and, therefore, does normally not define a ranking in a unique way.

Binary preference relations can be turned into a ranking by finding a ranking that is maximally consistent with the corresponding pairwise preferences. The difficulty of this optimization problem depends on the concrete criterion, though many natural objectives (e.g., minimizing the number of object pairs whose ranks are in conflict with their pairwise preference) lead to NP-hard problems [12]. Fortunately, efficient techniques such as simple voting (known as the Borda count procedure in social choice theory) often deliver good approximations, sometimes even with provable guarantees [13]. Of course, one can also derive other, possibly more complex preference structures from a preference relation, for example weak instead of strict linear orders. In [45], a linear order with ties (indifference between two alternatives) is called a *bucket order* (a total order of “buckets”, where each bucket corresponds to an equivalence class), and a method is proposed to find an order of this type, which is maximally consistent with the data.

For object ranking problems, the relational approach has been pursued in [12]. The authors propose to solve object ranking problems by learning a binary preference predicate $Q(x, x')$, which predicts whether x is preferred to x' or vice versa. This predicate is trained on the basis of exemplary preferences of the form $x \succ x'$. A final ordering is found in a second phase by deriving (an approximation of) a ranking that is maximally consistent with these predictions. DEMBCZYŃSKI et al.

discuss this setting for rule learning and propose to combine the predictions using the Net Flow score proposed in [4]. They also compare this setting with an alternative approach that directly learns a utility function based on the preferences and monotonicity constraints.

For label ranking problems, the pairwise approach has been introduced by [18, 31], where it is referred to as *ranking by pairwise comparison*. The key idea is to learn, for each pair of labels (y_i, y_j) , a binary predicate $\mathcal{M}_{i,j}(\mathbf{x})$ that predicts whether $y_i \succ_{\mathbf{x}} y_j$ or $y_j \succ_{\mathbf{x}} y_i$ for an input \mathbf{x} . A label ranking is then derived from these pairwise preferences via weighted voting (generalized Borda counting).

Pairwise learning techniques for instance ranking have been proposed in [19]. More specifically, two approaches were developed and compared in that paper, one which trains binary models $\mathcal{M}_{i,j}$, one for each pair of labels y_i and y_j , and another one that trains models \mathcal{M}_i ($i = 1, \dots, k - 1$) to separate classes y_1, \dots, y_i from classes y_{i+1}, \dots, y_k . Given a new query instance \mathbf{x} , both approaches submit this instance to all models that have been learned and aggregate the corresponding predictions into an overall score. A set of instances is then ranked according to these scores.

An overview of work on learning binary preference relations for label and instance ranking is given by FÜRNKRANZ & HÜLLERMEIER.

3.3 Model-Based Preference Learning

Another approach to learning ranking functions is to proceed from specific model assumptions, that is, assumptions about the structure of the preference relations. This approach is less generic than the previous ones, as it strongly depends on the concrete assumptions made.

An example is the assumption that the target ranking of a set of objects described in terms of multiple attributes can be represented as a *lexicographic order*. YAMAN et al. address the learning of lexicographic orders in the context of object ranking. From a machine learning point of view, assumptions of the above type can be seen as an inductive bias restricting the hypothesis space. Provided the bias is correct, this is clearly an advantage, as it may simplify the learning problem. In the case of lexicographic orders, for example, a complete ranking of all objects is uniquely identified by a total order of the attributes plus a total order of each of the attribute domains. For example, suppose objects to be described in terms of (only) four binary attributes. Thus, there are $2^4 = 16$ objects and hence $16! \approx 2 \cdot 10^{13}$ rankings in total. However, only $(2^4) \cdot 4! = 384$ of these rankings can be expressed in terms of a lexicographic order.

Needless to say, the bias induced by the assumption of a lexicographic order is very strong and will be rarely justified in practical applications. In particular, preferences on individual attribute values will normally not be independent of each other. For example, red wine might be preferred as a beverage if the main dish is meat, while white wine might be preferred in the case of fish. As mentioned earlier,

CP-nets [3] offer a language for expressing preferences on the values of single attributes and provide a suitable formalism for modeling dependencies of this type. A compact representation of a complex preference (partial order) relation is achieved by making use of conditional independence relations between such preferences (which are interpreted in terms of a *ceteris paribus* semantics), in much the same way as Bayesian networks reduce complexity of probability models by exploiting conditional independence between random variables. The CP-net itself is a graphical representation of these (in)dependencies, and each node (belonging to a single variable) is associated with a function that assigns a preference relation on the values of that attribute to each combination of values of the parent attributes. In this volume, CHEVALEYRE et al. discuss the learnability of CP-networks, both in a passive and an active learning scenario.

If a ranking function is defined implicitly via an underlying utility (scoring) function, the latter is normally also restricted by certain model assumptions. For example, the approaches outlined in Sect. 3.1 make use of linear functions to represent scores, although mostly for algorithmic reasons. There are other approaches in which the choice of the underlying utility function is more intrinsically motivated and addressed in a more explicit way. For example, GIESEN et al. describe an approach for conjoint analysis, also called multiattribute compositional models, which originated in mathematical psychology and is nowadays widely used in the social sciences and operations research. Proceeding from the description of objects in terms of a set of attributes, they assume that an underlying utility function can be decomposed into a linear sum of individual utility functions, one for each attribute. These utility functions can then be learned efficiently from the data.

Like in the case of lexicographic orders, this model assumption is obviously quite restrictive. TORRA presents a complementary and more general approach. Starting with a discussion of general properties that aggregation operators for attribute-based utility functions should fulfill, he surveys different approaches for learning a complex aggregation operator in the form of a nonadditive integral.

3.4 Local Aggregation of Preferences

Yet another alternative is to resort to the idea of local estimation techniques as prominently represented, for example, by the nearest neighbor estimation principle: Considering the rankings observed in similar situations as representative, a ranking for the current situation is estimated on the basis of these “neighbored” rankings, typically using an averaging-like aggregation operator. This approach is in a sense orthogonal to the previous model-based one, as it is very flexible and typically comes with no specific model assumption (except the regularity assumption underlying the nearest neighbor inference principle).

For label ranking, the nearest neighbor (instance-based learning) approach was first used in [9, 10]. Roughly speaking, the idea is to identify the query’s k nearest neighbors in the instance space \mathcal{X} , and then to combine the corresponding rankings

into a prediction using suitable aggregation techniques. In [11], this approach was developed in a theoretically more sophisticated way, realizing the aggregation step in the form of maximum likelihood estimation based on a statistical model for rank data. Besides, this approach is also able to handle the more general case in which the rankings of the neighbored instances are only partially known.

In the same paper, the authors propose to use this estimation principle for decision tree learning, too, namely for aggregating the label rankings associated with the instances grouped in a leaf node of the tree. Indeed, decision tree learning can also be seen as a kind of local learning method, namely as a piecewise constant approximation of the target function.³ In this book, YU, WAN & LEE propose a similar method. They grow a decision tree and propose two splitting measures for label ranking data. The rankings that are predicted at the leaves of the trees are derived by aggregating the rankings of all training examples arriving at this leaf.

Aggregation techniques are also used for other types of preference learning problems, including object ranking. For example, assume the rankings of several subsets X_i of a reference set \mathcal{X} to be given. The learning task is to combine these rankings into a complete ranking of all objects in \mathcal{X} . A practical application of this setting occurs, e.g., in information retrieval, when different rankings of search results originating from different search engines should be combined into an overall ranking of all retrieved pages [16]. Among other things, the learning problem may involve the determination of suitable weights for the information sources (search engines), reflecting their performance or agreement with the preferences of the user [35]. Another example is the ranking of sports teams or players, where individual tournament results with varying numbers of participants have to be combined into an overall ranking [2], or where different rankings by different judges have to be aggregated into an overall ranking of the participants of a competition [23].

4 Applications of Preference Learning

Preference learning problems in general, and ranking problems in particular, arise quite naturally in many application areas. For example, a search engine should rank Web pages according to a user's preferences. Likewise, cars can be ranked according to a customer's preferences on characteristics that discriminate different models. Another example is the ranking of possible keywords according to their relevance for an article. A few more examples have already been given in this article, and many more could be found.

In particular in the field of information retrieval, ranking applications occur quite naturally. Two particularly interesting problems are learning to rank the results of a query to a search engine, and learning to rank possible recommendations for new products. We will briefly discuss research in these areas below.

³ More general approximations can be realized by labeling a leaf node with a nonconstant function, for example a linear function in regression learning.

4.1 Learning to Rank Search Results

A widely studied preference learning problem is the ranking of retrieval results of a search engine. Roughly, the problem is the following: given a query \mathbf{q} and a set of documents \mathcal{D} , find a ranking of the documents in \mathcal{D} that corresponds to their relevance with respect to \mathbf{q} . This ranking is based on an unknown preference relation $\succ_{\mathbf{q}}$, which ought to be learned from user feedback on past rankings of retrieval results for different queries. An elaborate survey of current research in this area can be found in [36].

Current research focuses particularly on suitable ways of characterizing the queries that allow one to transfer the ranking from one query to another [20, 40]. In some sense, a query may be considered as a context for a ranking, and the task is to learn a function that allows one to transfer the ranking from one context to the other. Much of the research is conducted on the LETOR (LEARNING TO Rank for information retrieval) collection, a package of datasets containing a wide variety of queries with user feedback in several domains.⁴

Users can provide explicit feedback by labeling the retrieved pages with their degree of relevance. However, users are only willing to do this for a limited number of pages. It would be better if feedback could be collected in a way that is transparent to the user. RADLINSKI & JOACHIMS discuss a variety of techniques that allow one to collect the user feedback implicitly via their clicking behavior. Alternatively, ARENS proposes the use of active learning techniques which help minimizing the burden on the user by a careful automatic selection of suitable training examples for the ranking algorithm. He illustrates his technique in an application which learns a ranking function for the PubMed search engine for the MEDLINE database of biomedical literature.

4.2 Recommender Systems

Nowadays, *recommender systems* [43] are frequently used by online stores to recommend products to their customers. Such systems typically store a data table with products over users, which records the degree of preference of a user for this product. A customer can provide this preference degree explicitly by giving some sort of feedback (e.g., by assigning a rating to a movie) or implicitly (e.g., by buying the DVD of the movie). The elegant idea of collaborative filtering systems [21] is that recommendations can be based on user similarity, and that user similarity can in turn be defined by the similarity of their recommendations. Alternatively, recommender systems can also be based on item similarities, which are defined via the recommendations of the users that recommended the items in question. Yet other approaches try to learn models that capture the preference information contained in the matrix. A very good comparison of these approaches can be found in [8].

⁴ <http://research.microsoft.com/en-us/um/beijing/projects/letor/>

In this book, DE GEMMIS et al. given an extensive survey of recommender systems. Subsequently, KARATZOGLOU & WEIMER describe the use of matrix factorization methods for compressing the information contained in the user/product matrix into a product of two lower-dimensional matrices. The dimensions over which the product is computed may be viewed as hidden concepts, which can be used to categorize the interests of a user. Interestingly, only very few (in the order of 10) such concepts are enough for a sufficiently accurate representation of large numbers of users and products. Finally, BELLOGIN et al. describe an approach that uses decision tree learning for identifying features of recommendation models that influence the quality of the predicted preference ranking.

5 Conclusions

In this introductory chapter, we have tried to give an overview of different approaches to preference learning, categorized by the learning task (label, instance, or object ranking) and the learning technique (learning utility functions, learning binary preference relations, learning preference models having a specific structure, or using local estimation and preference aggregating methods). In principle, all task/technique combinations are conceivable and can be found in the literature. We also highlighted important application areas, in particular in ranking search results and product recommendations.

Throughout the chapter, pointers to the remaining articles in this book were given. They could be characterized along a variety of dimensions, including all of those mentioned above. We have adopted a grouping that more or less corresponds to the sections in this survey. As the categorization is multi-dimensional, most articles fit into several of these categories, and, in some cases, the choice was not entirely clear. This is why we do not have separate parts on learning of a utility function or a binary preference relation, for example. In fact, almost all approaches presented in this book follow either of the two approaches.

Acknowledgements This research has been supported by the German Science Foundation (DFG).

References

1. R. Balasubramanian, E. Hüllermeier, N. Weskamp, J. Kämper, Clustering of gene expression data using a local shape-based similarity measure. *Bioinformatics* **21**(7), 1069–1077 (2005)
2. A. Birlutiu, T. Heskes, Expectation propagation for rating players in sports competitions, in *Proceedings of the 11th European Symposium on Principles of Knowledge Discovery in Databases (PKDD-07)*, ed. by J.N. Kok, J. Koronacki, R. López de Mántaras, S. Matwin, D. Mladenić, A. Skowron (Springer, Warsaw, Poland, 2007), pp. 374–381
3. C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, D. Poole, CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res.* **21**, 135–191 (2004)
4. D. Bouyssou, Ranking methods based on valued preference relations: A characterization of the net flow method. *Eur. J. Oper. Res.* **60**(1), 61–67 (1992)

5. A.P. Bradley, The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognit.* **30**(7), 1145–1159 (1997)
6. R.I. Brafman, Preferences, planning and control, in *Proceedings of the 11th Conference on Principles of Knowledge Representation and Reasoning (KR-08)*, ed. by G. Brewka, J. Lang (AAAI, Sydney, Australia, 2008), pp. 2–5
7. P.B. Brazdil, C. Soares, J.P. da Costa, Ranking learning algorithms: Using IBL and meta-learning on accuracy and time results. *Mach. Learn.* **50**(3), 251–277 (2003)
8. J.S. Breese, D. Heckerman, C. Kadie, Empirical analysis of predictive algorithms for collaborative filtering, in *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI-98)*, ed. by G.F. Cooper, S. Moral (Morgan Kaufmann, Madison, WI, 1998), pp. 43–52
9. K. Brinker, E. Hüllermeier, Case-based label ranking, in *Proceedings of the 17th European Conference on Machine Learning (ECML-06)*, ed. by J. Fürnkranz, T. Scheffer, M. Spiliopoulou (Springer, Berlin, Germany, 2006), pp. 566–573
10. K. Brinker, E. Hüllermeier, Case-based multilabel ranking, in *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, ed. by M.M. Veloso (Hyderabad, India, 2007), pp. 702–707
11. W. Cheng, J. Hühn, E. Hüllermeier, Decision tree and instance-based learning for label ranking, in *Proceedings of the 26th International Conference on Machine Learning (ICML-09)*, ed. by A. Danyluk, L. Bottou, M.L. Littman (Montreal, Canada, 2009)
12. W.W. Cohen, R.E. Schapire, Y. Singer, Learning to order things. *J. Artif. Intell. Res.* **10**, 243–270 (1999)
13. D. Coppersmith, L. Fleischer, A. Rudra, Ordering by weighted number of wins gives a good ranking for weighted tournaments, in *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA-06)* (2006), pp. 776–782
14. O. Dekel, C.D. Manning, Y. Singer, Log-linear models for label ranking, in *Advances in Neural Information Processing Systems (NIPS-03)*, ed. by S. Thrun, L.K. Saul, B. Schölkopf (MIT, Cambridge, MA, 2004), pp. 497–504
15. J. Doyle, Prospects for preferences. *Comput. Intell.* **20**(2), 111–136 (2004)
16. C. Dwork, R. Kumara, M. Naor, D. Sivakumar, Rank aggregation methods for the Web, in *Proceedings of the 10th International World Wide Web Conference (WWW-01)* (ACM, Hong Kong, China, 2001), pp. 613–622
17. R. Fagin, R. Kumar, D. Sivakumar, Comparing top k lists. *SIAM J. Discrete Math.* **17**(1), 134–160 (2003)
18. J. Fürnkranz, E. Hüllermeier, Pairwise preference learning and ranking, in *Proceedings of the 14th European Conference on Machine Learning (ECML-03)*, vol. 2837, *Lecture Notes in Artificial Intelligence*, ed. by N. Lavrač, D. Gamberger, H. Blockeel, L. Todorovski (Springer, Cavtat, Croatia, 2003), pp. 145–156
19. J. Fürnkranz, E. Hüllermeier, S. Vanderlooy, Binary decomposition methods for multipartite ranking, in *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD-09)*, vol. Part I, ed. by Wray L. Buntine, M. Grobelnik, D. Mladenic, J. Shawe-Taylor (Springer, Bled, Slovenia, 2009), pp. 359–374
20. X. Geng, T.-Y. Liu, T. Qin, A. Arnold, H. Li, H.-Y. Shum, Query dependent ranking using k -nearest neighbor, In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-08)*, ed. by S.-H. Myaeng, D.W. Oard, F. Sebastiani, T.-S. Chua, M.-K. Leong (ACM, Singapore, 2008), pp. 115–122
21. D. Goldberg, D. Nichols, B.M. Oki, D. Terry, Using collaborative filtering to weave and information tapestry. *Commun. ACM* **35**(12), 61–70 (1992)
22. M. Gönen, G. Heller, Concordance probability and discriminatory power in proportional hazards regression. *Biometrika* **92**(4), 965–970 (2005)
23. S. Gordon, M. Truchon, Social choice, optimal inference and figure skating. *Soc. Choice Welfare* **30**(2), 265–284 (2008)
24. P. Haddawy, V. Ha, A. Restificar, B. Geisler, J. Miyamoto, Preference elicitation via theory refinement. *J. Mach. Learn. Res.* **4**, 317–337 (2003)

25. S. Har-Peled, D. Roth, D. Zimak, Constraint classification: A new approach to multiclass classification, in *Proceedings of the 13th International Conference on Algorithmic Learning Theory (ALT-02)*, ed. by N. Cesa-Bianchi, M. Numao, R. Reischuk (Springer, Lübeck, Germany, 2002), pp. 365–379
26. S. Har-Peled, D. Roth, D. Zimak, Constraint classification for multiclass classification and ranking, in *Advances in Neural Information Processing Systems 15 (NIPS-02)*, ed. by S. Becker, S. Thrun, K. Obermayer (2003), pp. 785–792
27. R. Herbrich, T. Graepel, K. Obermayer, Large margin rank boundaries for ordinal regression, in *Advances in Large Margin Classifiers*, ed. by P.J. Bartlett, B. Schölkopf, D. Schuurmans, A.J. Smola (MIT, 2000), pp. 115–132
28. E. Hüllermeier, J. Fürnkranz, Learning label preferences: Ranking error versus position error, in *Advances in Intelligent Data Analysis: Proceedings of the 6th International Symposium (IDA-05)* (Springer, Madrid, Spain, 2005), pp. 180–191
29. E. Hüllermeier, J. Fürnkranz (eds.), *Proceedings of the ECML/PKDD-08 Workshop on Preference Learning* (Antwerp, Belgium, 2008)
30. E. Hüllermeier, J. Fürnkranz (eds.), *Proceedings of the ECML/PKDD-09 Workshop on Preference Learning* (Bled, Slovenia, 2009)
31. E. Hüllermeier, J. Fürnkranz, W. Cheng, K. Brinker, Label ranking by learning pairwise preferences. *Artif. Intell.* **172**, 1897–1916 (2008)
32. K. Järvelin, J. Kekäläinen, Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst.* **20**(4), 422–446 (2002)
33. T. Joachims, Optimizing search engines using clickthrough data, in *Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-02)* (ACM, 2002), pp. 133–142
34. T. Joachims, L. Granka, B. Pan, H. Hembrooke, G. Gay, Accurately interpreting clickthrough data as implicit feedback, in *Proceedings of the 28th Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR-05)* (2005), pp. 154–161
35. G. Lebanon, J. Lafferty, Cranking: Combining rankings using conditional probability models on permutations, in *Proceedings of the 19th International Conference on Machine Learning (ICML-02)*, ed. by C. Sammut, A. Hoffmann (Sydney, Australia, 2002), pp. 363–370
36. T.-Y. Lu, Learning to rank for information retrieval. *Found. Trends Inf. Retrieval* **3**(3), 225–331 (2009)
37. R. Duncan Luce, H. Raiffa, *Games and Decisions: Introduction and Critical Survey* (Wiley, New York, NY, 1957)
38. H.B. Mann, D.R. Whitney, On a test of whether one of two random variables is stochastically larger than the other. *Ann. Math. Stat.* **18**(50), 50–60 (1947)
39. C.D. Manning, P. Raghavan, H. Schütze, *Introduction to Information Retrieval* (Cambridge University Press, 2008)
40. W. Ni, Y. Huang, M. Xie, A query dependent approach to learning to rank for information retrieval, in *Proceedings of the 9th International Conference on Web-Age Information Management (WAIM-08)* (IEEE, Zhangjiajie, China, 2008), pp. 262–269
41. F. Radlinski, T. Joachims, Learning to rank from implicit feedback, in *Proceedings of the ACM Conference on Knowledge Discovery and Data Mining (KDD-05)* (2005), pp. 239–248
42. S. Rajaram, S. Agarwal, Generalization bounds for k -partite ranking, in *Proceedings of the NIPS 2005 Workshop on Learning to Rank*, ed. by S. Agarwal, C. Cortes, R. Herbrich, (Whistler, BC, Canada, 2005), pp. 28–23
43. P. Resnick, H.R. Varian, Special issue on recommender systems. *Commun. ACM* **40**(3), (1997)
44. G. Tesauro, Connectionist learning of expert preferences by comparison training. in *Advances in Neural Information Processing Systems 1 (NIPS-88)*, ed. by D. Touretzky (Morgan Kaufmann, 1989), pp. 99–106
45. A. Ukkonen, K. Puolamäki, A. Gionis, H. Mannila, A randomized approximation algorithm for computing bucket orders. *Inf. Process. Lett.* **109**, 356–359 (2009)
46. J. Wang, Artificial neural networks versus natural neural networks: A connectionist paradigm for preference assessment. *Decision Support Syst.* **11**, 415–429 (1994)
47. F. Wilcoxon, Individual comparisons by ranking methods. *Biometrics Bull.* **1**, 80–83 (1945)

A Preference Optimization Based Unifying Framework for Supervised Learning Problems

Fabio Aioli and Alessandro Sperduti

Abstract Supervised learning is characterized by a broad spectrum of learning problems, often involving structured types of prediction, including classification, ranking-based predictions (label and instance ranking), and (ordinal) regression in its various forms. All these different learning problems are typically addressed by specific algorithmic solutions.

In this chapter, we propose a general preference learning model (GPLM), which gives an easy way to translate any supervised learning problem and the associated cost functions into sets of preferences to learn from. A large margin principled approach to solve this problem is also proposed.

Examples of how the proposed framework has been effectively used by us to address non-standard real-world applications are reported showing the flexibility and effectiveness of the approach.

1 Introduction

Supervised learning is probably the most commonly used learning paradigm and a large spectrum of learning algorithms have been devised for different learning tasks in the last decades. The need for such a large spectrum of learning algorithms is, in part, due to the many real-world learning problems, that are characterized by heterogeneous tasks and problem-specific learning algorithms for their solution. These include classification and regression problems (including multilabel and multiclass classification, and multivariate regression), as well as ranking-based (either label or instance ranking) and ordinal regression problems. Typically, the approach followed to deal with a nonstandard problem is to map it into a series of simpler, well-known problems and then to combine the resulting predictions. Often, however, this type

F. Aioli (✉) and A. Sperduti

Department of Pure and Applied Mathematics - Padova - Italy, Via Trieste 63, 35131 Padova, Italy
e-mail: aioli@math.unipd.it, sperduti@math.unipd.it

of methodology lacks a principled theory supporting it and/or requires too much computational resources to be practical for real-world applications.

In this chapter, we give a survey of a quite general framework, which is able to generalize different types of supervised learning settings into a common preference optimization task. In particular, this is done by considering supervision as a set of order preferences over the predictions of the learner. More generally, we show that supervised learning problems can be characterized by considering two main dimensions, the type of prediction and the type of supervision involved in the problem to be solved. Then, based on this characterization, we are able to map any of these learning problems into a simple preference learning task. From a practical point of view, we show how all these supervised tasks can also be addressed in a simple linear setting, where any problem formulation can be transformed into a binary problem defined on an augmented space, thus allowing the exploitation of very simple optimization procedures available for the binary case. We also stress the flexibility of the preference model, which allows a user to optimize the parameters on the basis of a proper evaluation function. In fact, while in general the goal of a problem in terms of its evaluation function is clear, a crucial issue in the design of a learning algorithm is how to get a theoretical guarantee that the defined learning procedure actually minimizes the target cost function. One advantage of the framework reviewed in this chapter is that it defines a very natural and uniform way to devise and code a cost function into a learning algorithm.

Examples of real-world applications are then discussed. In particular, two recent applications are discussed in more detail. The first application concerns the problem to select the best candidate for a job role. This is an instance ranking problem, where, however, only binary supervision from the past history is available. The second application concerns a patent classification task, where patent applications have to be associated with primary categories as well as secondary categories. This is an example of a label ranking task, which cannot be properly addressed by an ordinal regression approach.

In Sect. 2, we review the general preference learning model (GPLM). Specifically, we show how the preference model generalizes the supervised learning setting by considering supervision as a partial order of (soft) constraints over the learner predictions. In addition, we show (Sect. 2.2) how the suggested generalization can be instantiated to well-known supervised learning problems. In the same section, we also discuss how cost functions for learning problems can be cast by using preferences (Sect. 2.3) and a simple linear model for the learner (Sect. 2.4). Quite general optimization procedures for training models within the proposed framework are also presented (Sect. 2.5). In Sect. 3, different application scenarios are described and discussed. In particular, it is discussed how the GPLM applies to a job candidate selection task and to a patent classification task. In Sect. 4, related works are sketched and a discussion about the proposed approach is given. Finally, in Sect. 5, some future extensions to the preference framework are suggested and final conclusions are drawn.

2 GPLM: A General Model for Supervised Learning

2.1 The Learning Domain

Let us consider a very general domain with a space of instances \mathcal{X} and a space of class labels \mathcal{Y} . For example, this could be the domain of a recommender system where instances might correspond to customers while labels to products, or the domain of an information retrieval system where instances could correspond to documents while labels to queries.

The basic idea underpinning our general preference learning model is that we want to learn the set of parameters of a real valued relevance (or scoring) function defined on instance label pairs

$$f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R},$$

which should approximate the actual *target* function. In a recommender system task, for example, this target function would represent the actual rating (a real value) a customer would give to a given product. Similarly, in the information retrieval example, the target function could represent the log-ratio of the probability of relevance of a document given a query.

We can easily note that, once such a scoring function is computed, a predictor will be able to order instances in \mathcal{X} based on their relevance once any label $y \in \mathcal{Y}$ is selected, and similarly, to order class labels in \mathcal{Y} based on their relevance once any instance $x \in \mathcal{X}$ is selected.

2.2 Prediction and Supervision

In supervised learning, supervision is assumed to be provided according to an unknown probability distribution \mathcal{D} over pairs, where the first member is a description of a domain object (instance) and the second member is the corresponding expected prediction (target label). We generalize this setting by considering supervision as (soft) constraints over the learner predictions, that is constraints whose violation entails a cost, or penalty, for the solution. Specifically, we assume a learner makes its predictions on the basis of a set of parameters Θ , characterizing its *hypothesis space*. Each supervision constraint S , that cannot be satisfied makes the learner suffer a cost $c(S|\Theta)$. It is easy to notice that this generalizes the above-mentioned case of supervision as instance-label pairs. In fact, this is obtained back when a unitary cost is given to hypotheses generating incorrect labeling.

Now, we are able to show that, by using the setting presented above, it is possible to cast the main types of supervised learning tasks into a taxonomy on the basis of their expected prediction and supervision feedback. To this end, let us first recall the definition of order relations.

2.2.1 Definition

A *partial order* is a pair (\mathcal{P}, \succeq) in which \mathcal{P} is a set and \succeq is a reflexive, antisymmetric, and transitive binary relation. A *partial ranking* of length r is a partial order in which the set \mathcal{P} can be partitioned in r sets $\mathcal{P}_1, \dots, \mathcal{P}_r$ such that $z \in \mathcal{P}_i, z' \in \mathcal{P}_j, i < j$, implies $z \succeq z'$ and no further information is conveyed about the ordering within subsets \mathcal{P}_k . A *full order* on \mathcal{P} is defined as a partial ranking of length $|\mathcal{P}|$. We denote by $PO(\mathcal{P})$, $PR(\mathcal{P})$, and $FO(\mathcal{P})$ the set of partial orders, partial rankings, and full orders over the set \mathcal{P} , respectively.

2.2.2 Label Rankings as Qualitative Preferences

A first important family of supervised learning tasks is related to the ordering of the classes on the basis of their relevance for an instance, and thus they are characterized by the fact that predictions should be based on a full order over the labels. This family of problems is referred to as *label rankings*. Supervision is in the form of partial orders over the classes. In our notation, we have supervision $S \in PO(\mathcal{Y})$ and predictions in $FO(\mathcal{Y})$. Different settings can be obtained corresponding to different types of supervision. A few well-known instances are listed in the following:

Category Ranking (CR)

In this setting, the goal is to order categories on the basis of their relevance for an instance. As an example, in a collaborative filtering setting, users could correspond to our instances and the different movies to our classes. Then, one could be interested in the ordering (by relevance) of the set of movies based on user preferences. This is trivially a particular case of label ranking where supervision is given as full orders over \mathcal{Y} .

Bipartite Category Ranking (BCR)

In this task, supervision is given as two groups of classes and it is required to predict full orders in which the first group of classes is ranked over the second. As a leading example, in information retrieval, given a document, one might have to rank the available topics with the aim to return the most relevant topics on the top of the list. This is again a specific case of label ranking where supervision is given as partial rankings of length two. This task has been also referred to as category ranking in literature [10]. Here a different terminology is adopted to avoid confusion between these two different tasks.¹

¹ Note that this task and the two that follow are conceptually different from the task to decide about the membership of an instance. Here, supervision only gives *qualitative* information about the fact that some classes are more relevant than others.