

Learning to Rank Based on Analogical Reasoning

Mohsen Ahmadi Fahandar, Eyke Hüllermeier

Department of Computer Science
Paderborn University
Pohlweg 49-51, 33098 Paderborn, Germany
ahmadim@mail.upb.de, eyke@upb.de

Abstract

Object ranking or “learning to rank” is an important problem in the realm of preference learning. On the basis of training data in the form of a set of rankings of objects represented as feature vectors, the goal is to learn a ranking function that predicts a linear order of any new set of objects. In this paper, we propose a new approach to object ranking based on principles of analogical reasoning. More specifically, our inference pattern is formalized in terms of so-called analogical proportions and can be summarized as follows: Given objects A, B, C, D , if object A is known to be preferred to B , and C relates to D as A relates to B , then C is (supposedly) preferred to D . Our method applies this pattern as a main building block and combines it with ideas and techniques from instance-based learning and rank aggregation. Based on first experimental results for data sets from various domains (sports, education, tourism, etc.), we conclude that our approach is highly competitive. It appears to be specifically interesting in situations in which the objects are coming from different subdomains, and which hence require a kind of knowledge transfer.

1 Introduction

Preference learning has received increasing attention in machine learning in recent years (Fürnkranz and Hüllermeier 2011). Roughly speaking, the goal in preference learning is to induce preference models from observational (or experimental) data that reveal information about the preferences of an individual or a group of individuals in a direct or indirect way; the latter typically serve the purpose of predictive modeling, i.e., they are then used to predict the preferences in a new situation.

In general, a preference learning system is provided with a set of items (e.g., products) for which preferences are known, and the task is to learn a function that predicts preferences for a new set of items (e.g., new products not seen so far), or for the same set of items in a different situation (e.g., the same products but for a different user). Frequently, the predicted preference relation is required to form a total order, in which case we also speak of a *ranking problem*. In fact, among the problems in the realm of preference learning, the task of “learning to rank” has probably received the most attention in the literature so far,

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

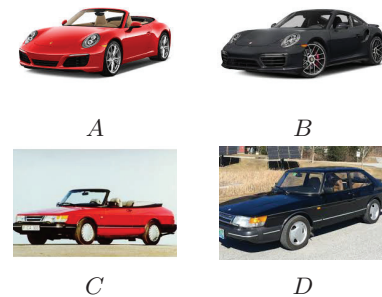


Figure 1: A is to B as C is to D (pictures from ImageNet).

and a number of different ranking problems have already been introduced. Based on the type of training data and the required predictions, Fürnkranz and Hüllermeier distinguish between the problems of object ranking (Cohen, Schapire, and Singer 1999; Kamishima, Kazawa, and Akaho 2010), label ranking (Har-Peled, Roth, and Zimak 2002; Cheng, Hühn, and Hüllermeier 2009; Vembu and Gärtner 2010), and instance ranking (Fürnkranz, Hüllermeier, and Vanderlooy 2009).

The focus of this paper is on object ranking. Given training data in the form of a set of exemplary rankings of subsets of objects, the goal in object ranking is to learn a ranking function that is able to predict the ranking of any new set of objects. Our main contribution is a novel approach that is based on the idea of *analogical reasoning*, and essentially builds on the following inference pattern: If object A relates to object B as C relates to D , and knowing that A is preferred to B , we (hypothetically) infer that C is preferred to D . Figure 1 provides an illustration with cars as objects. Cars A and B are more or less the same, except that A is a cabriolet and has color red instead of black, and the same is true for cars C and D . Thus, knowing that someone likes car A more than B , we may conjecture that the same person prefers car C to D . Our method applies this pattern as a main building block and combines it with ideas and techniques from instance-based learning and rank aggregation.

The rest of the paper is organized as follows. In the next section, we recall the setting of object ranking and formalize the corresponding learning problem. In the next section,

we present a formalization of analogical reasoning based on the concept of analogical proportion. Next, we introduce our approach to **analogy-based learning to rank** (able2rank). Finally, we present an experimental evaluation of this approach, prior to concluding the paper with a summary and an outline of future work.

2 Learning to Rank

Consider a reference set of objects, items, or choice alternatives \mathcal{X} , and assume each item $\mathbf{x} \in \mathcal{X}$ to be described in terms of a feature vector; thus, an item is a vector $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$ and $\mathcal{X} \subseteq \mathbb{R}^d$. The goal in object ranking is to learn a *ranking function* ρ that accepts any (query) subset

$$Q = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subseteq \mathcal{X}$$

of $n = |Q|$ items as input. As output, the function produces a ranking $\pi \in \mathbb{S}_n$ of these items, where \mathbb{S}_n denotes the set of all permutations of length n , i.e., all mappings $[n] \rightarrow [n]$ (symmetric group of order n); π represents the total order

$$\mathbf{x}_{\pi^{-1}(1)} \succ \mathbf{x}_{\pi^{-1}(2)} \succ \dots \succ \mathbf{x}_{\pi^{-1}(n)}, \quad (1)$$

i.e., $\pi^{-1}(k)$ is the index of the item on position k , while $\pi(k)$ is the position of the k th item \mathbf{x}_k (π is often called a *ranking* and π^{-1} an *ordering*). Formally, a ranking function is thus a mapping

$$\rho : \mathcal{Q} \rightarrow \mathcal{R}, \quad (2)$$

where $\mathcal{Q} = 2^{\mathcal{X}} \setminus \emptyset$ is the *query space* and $\mathcal{R} = \bigcup_{n \in \mathbb{N}} \mathbb{S}_n$ the *ranking space*. The order relation \succ is typically (though not necessarily) interpreted in terms of preferences, i.e., $\mathbf{x} \succ \mathbf{y}$ suggests that \mathbf{x} is preferred to \mathbf{y} .

A ranking function ρ is learned on a set of training data that consists of a set of rankings

$$\mathcal{D} = \{(Q_1, \pi_1), \dots, (Q_M, \pi_M)\}, \quad (3)$$

where each ranking π_j defines a total order of the query set Q_j . Once a ranking function has been learned, it can be used for making predictions for new query sets Q (see Figure 2). Such predictions are evaluated in terms of a suitable loss function or performance metric. A common choice is the (normalized) *ranking loss*, which counts the number of inversions between two rankings π and π' :

$$d_{RL}(\pi, \pi') = \frac{\sum_{1 \leq i, j \leq n} [\pi(i) < \pi(j)] [\pi'(i) > \pi'(j)]}{n(n-1)/2},$$

where $[\cdot]$ is the indicator function.

2.1 Methods

The ranking function (2) sought in object ranking is a complex mapping from the query to the ranking space. An important question, therefore, is how to represent a “ranking-valued” function of that kind, and how it can be learned efficiently. As will be seen later on, our approach avoids an explicit representation of this function, and instead implements a query-specific (transductive) inference, in very much the same way as instance-based learning represents a predictor (through the sample data) in an indirect way.

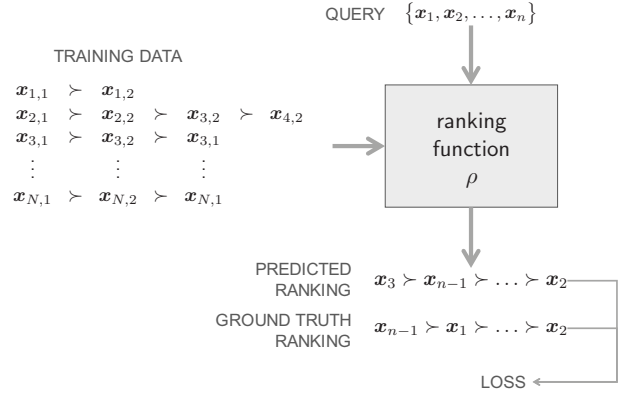


Figure 2: The setting of object ranking.

Most commonly, a ranking function is represented by means of an underlying scoring function

$$U : \mathcal{X} \rightarrow \mathbb{R},$$

so that $\mathbf{x} \succ \mathbf{x}'$ if $U(\mathbf{x}) > U(\mathbf{x}')$. In other words, a ranking-valued function is represented through a real-valued function. Obviously, U can be considered as a kind of utility function, and $U(\mathbf{x})$ as a latent utility degree assigned to an item \mathbf{x} . Seen from this point of view, the goal in object ranking is to learn a latent utility function on a reference set \mathcal{X} .

The representation of a ranking function in terms of a real-valued (utility) function also suggests natural approaches to learning. In particular, two such approaches are prevailing in the literature. The first one reduces the original ranking problem to *regression*; as it seeks a model that assigns appropriate scores to individual items \mathbf{x} , it is referred to as the *pointwise* approach in the literature. The second idea is to reduce the problem to *binary classification*; here, the focus is on pairs of items, which is why the approach is called the *pairwise* approach.

As a representative of the first category, we will include *expected rank regression* (ERR) in our experimental study later on (Kamishima and Akaho 2006; Kamishima, Kazawa, and Akaho 2010). ERR reduces object ranking to standard linear regression. To this end, every training example (Q, π) is replaced by a set of data points $(\mathbf{x}_i, y_i) \in \mathcal{X} \times \mathbb{R}$. Here, the target y_i assigned to object $\mathbf{x}_i \in Q$ is given by

$$y_i = \frac{\pi(i)}{|Q| + 1}.$$

This is justified by taking an expectation over all (complete) rankings of \mathcal{X} and assuming a uniform distribution. In spite of this apparently oversimplified assumption, and the questionable transformation of ordinal ranks into numerical scores, ERR has shown quite competitive performance in empirical studies, especially when all rankings in the training data (3) are of approximately the same length (Melnikov et al. 2016).

Given a ranking (1) as training information, the pairwise approach extracts all pairwise preferences $\mathbf{x}_{\pi^{-1}(i)} \succ$

$\mathbf{x}_{\pi^{-1}(j)}$, $1 \leq i < j \leq n$, and considers these preferences as examples for a binary classification task. This approach is especially simple if U is a linear function of the form $U(\mathbf{x}) = \mathbf{w}^\top \mathbf{x}$. In this case, $U(\mathbf{x}) > U(\mathbf{x}')$ if $\mathbf{w}^\top \mathbf{x} > \mathbf{w}^\top \mathbf{x}'$, which is equivalent to $\mathbf{w}^\top \mathbf{z} > 0$ for $\mathbf{z} = \mathbf{x} - \mathbf{x}' \in \mathbb{R}^d$. Thus, from the point of view of binary classification (with a linear threshold model), \mathbf{z} can be considered as a positive and $-\mathbf{z}$ as a negative example. In principle, any binary classification algorithm can be applied to learn the weight vector \mathbf{w} from the set of examples produced in this way. As a representative of this class of methods, we will use support vector machines in our experiments; more specifically, we include Ranking SVM (Joachims 2002) as a state-of-the-art baseline to compare with.

3 Analogical Reasoning

Analogical reasoning has a long tradition in artificial intelligence research, and various attempts at formalizing analogy-based inference can be found in the literature. We are building on a recent approach that is based on the concept of *analogical proportion* (Miclet and Prade 2009; Prade and Richard 2017), which has already been used successfully in different problem domains, including classification (Bounhas, Prade, and Richard 2014a), recommendation (Hug et al. 2016), preference completion (Pirlot, Prade, and Richard 2016), decision making (Billingsley et al. 2017) and solving IQ tests (Beltran, Prade, and Richard 2016).

Consider four values a, b, c, d from a domain \mathbb{X} . The quadruple (a, b, c, d) is said to be in analogical proportion, denoted $a : b :: c : d$, if “ a relates to b as c relates to d ”. A bit more formally, this idea can be expressed as

$$\mathcal{R}(a, b) \sim \mathcal{R}(c, d), \quad (4)$$

where the relation \sim denotes the “as” part of the informal description. \mathcal{R} can be instantiated in different ways, depending on the underlying domain \mathbb{X} . Specifically relevant for us are the cases of Boolean variables, where $\mathbb{X} = \mathbb{B} = \{0, 1\}$, and numerical variables, where $\mathbb{X} = \mathbb{R}$. In the latter case, we will distinguish between *arithmetic proportions*, where $\mathcal{R}(a, b) = a - b$, and *geometric proportions*, where $\mathcal{R}(a, b) = a/b$.

Independently of the concrete choice of \mathbb{X} and \mathcal{R} , the following axioms can reasonably be required (assuming an expression $a : b :: c : d$ to be either true or false):

- $a : b :: a : b$ (reflexivity)
- $a : b :: c : d \Rightarrow c : d :: a : b$ (symmetry)
- $a : b :: c : d \Rightarrow a : c :: b : d$ (central permutation)

3.1 Boolean Variables

Consider the case of Boolean variables with values in $\mathbb{B} = \{0, 1\}$. There are $2^4 = 16$ instantiations of the pattern $a : b :: c : d$. The smallest (and hence most informative) subset of *logical proportions* satisfying the above axioms consists of the following 6 instantiations:

a	b	c	d
0	0	0	0
0	0	1	1
0	1	0	1
1	0	1	0
1	1	0	0
1	1	1	1

This formalization captures the idea that a differs from b (in the sense of being “equally true”, “more true”, or “less true”) exactly as c differs from d , and vice versa. This can also be expressed by the following formula:

$$((a \Rightarrow b) \equiv (c \Rightarrow d)) \wedge ((b \Rightarrow a) \equiv (d \Rightarrow c)) \quad (5)$$

We remark that the above definition further satisfies the properties of independence with regard to positive or negative encoding of features ($a : b :: c : d \Rightarrow \neg a : \neg b :: \neg c : \neg d$) as well as transitivity ($a : b :: c : d$ and $c : d :: e : f \Rightarrow a : b :: e : f$). For an in-depth discussion of logical proportions, we refer to (Prade and Richard 2013).

3.2 Real-Valued Variables

To extend analogical proportions to real-valued variables X , we assume these variables take values in $[0, 1]$. Practically, this may require a normalization of the original variable in a preprocessing step (a discussion of this will follow below).

Values in the unit interval can be interpreted as generalized (graded) truth degrees. For instance, if x is the speed of a car, the normalized value $x = 0.8$ can be interpreted as the degree to which the property “high speed” applies, i.e., the truth degree of the proposition “the car has high speed”. The idea, then, is to generalize the formalization of analogical proportions in the Boolean case using generalized logical operators (Dubois, Prade, and Richard 2016). Naturally, the analogical proportion itself will then become a matter of degree, i.e., a quadruple (a, b, c, d) can be in analogical proportion *to some degree*; in the following, we will denote this degree by $v(a, b, c, d)$.

For example, taking (5) as a point of departure, and generalizing the logical conjunction $x \wedge y$ by the minimum operator $\min(x, y)$, the implication $x \Rightarrow y$ by Lukasiewicz implication $\max(1 - x + y, 0)$, and the equivalence $x \equiv y$ by $1 - |x - y|$, one arrives at the following expression:

$$v_A(a, b, c, d) = 1 - |(a - b) - (c - d)|$$

if $\text{sign}(a - b) = \text{sign}(c - d)$ and

$$v_A(a, b, c, d) = 1 - \max(|a - b|, |c - d|)$$

otherwise. We also consider a variant $v_{A'}$ of this expression, which is given by $v_{A'}(a, b, c, d) = v_A(a, b, c, d)$ if $\text{sign}(a - b) = \text{sign}(c - d)$ and $v_{A'}(a, b, c, d) = 0$ otherwise.

Based on alternative interpretations of \mathcal{R} and \sim in (4), other proposals for measuring the degree of analogical proportion can be found in the literature:

- Geometric proportion (Beltran, Jaudoin, and Pivert 2014):

$$v_G(a, b, c, d) = \frac{\min(ad, bc)}{\max(ad, bc)},$$

if $\text{sign}(a - b) = \text{sign}(c - d)$ and $\max(ad, bc) > 0$, and 0 otherwise.

- Min-Max proportion (Bounhas, Prade, and Richard 2014b):

$$v_{MM}(a, b, c, d) = 1 - \max(|\min(a, d) - \min(b, c)|, |\max(a, d) - \max(b, c)|).$$

- Approximate equality proportion (Beltran, Jaudoin, and Pivert 2014):

$$v_{AE}(a, b, c, d) = \max(\llbracket a \approx b \rrbracket \llbracket c \approx d \rrbracket, \llbracket a \approx c \rrbracket \llbracket b \approx d \rrbracket),$$

where $x \approx y$ is true if $|x - y| \leq \epsilon$ for a fixed threshold $\epsilon \in [0, 1]$. We will also use a graded variant of $v_{AE}(\cdot)$, denoted as $v_{AE'}$, as defined above. To this end, we generalize the indicator function and evaluate the approximate equality of values a, b as $\llbracket a \approx b \rrbracket = \max(1 - |a - b|/\epsilon, 0)$.

3.3 Extension to Feature Vectors

The degree of analogical proportion defined above can be generalized from individual values to objects in the form of vectors $\mathbf{x} = (x_1, \dots, x_d)$, where $x_i \in \mathbb{X}_i \in \{\mathbb{B}, \mathbb{R}\}$ —in the context of object ranking, \mathbf{x} is a feature vector characterizing a choice alternative. To this end, the degrees of analogical proportion for the individual entries are aggregated into a single degree:

$$v(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}) = \text{AGG} \{v(a_i, b_i, c_i, d_i) \mid i = 1, \dots, d\} \quad (6)$$

The aggregation operator AGG can be specified in different ways. For example, a conjunctive combination is obtained with $\text{AGG} = \min$. This aggregation is very strict, however, and does not allow for compensating a low value of analogical proportion on a single feature by large values on other features. In our current implementation of analogy-based object ranking, we therefore use the arithmetic average as an aggregation.

4 Analogy-based Object Ranking

Recall that, in the setting of learning to rank, we suppose to be given a set of training data in the form

$$\mathcal{D} = \{(Q_1, \pi_1), \dots, (Q_M, \pi_M)\},$$

where each π_m defines a ranking of the set of objects Q_m . If $z_i, z_j \in Q_m$ and $\pi_m(i) < \pi_m(j)$, then $z_i \succ z_j$ has been observed as a preference. In the following, we will denote by

$$\mathcal{D}_{pair} = \bigcup_{m=1}^M \bigcup_{1 \leq i < j \leq |Q_m|} (z_{\pi^{-1}(i)}, z_{\pi^{-1}(j)})$$

the set of all pairwise preferences that can be extracted from \mathcal{D} . The goal is to generalize beyond the data \mathcal{D} so as to be able to predict a ranking of any query set

$$Q = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}.$$

Our analogy-based approach to object ranking (able2rank) consists of two main steps:

- First, for each pair of objects $\mathbf{x}_i, \mathbf{x}_j \in Q$, a degree of preference $p_{i,j} \in [0, 1]$ is extracted from \mathcal{D} . If these degrees are normalized such that $p_{i,j} + p_{j,i} = 1$, they define a reciprocal preference relation

$$P = (p_{i,j})_{1 \leq i, j \leq n}.$$

Normalization is not a strict requirement, however; in our concrete implementation, pairwise preference are expressed in terms of absolute frequencies (see below).

- Second, the preference relation P is turned into a ranking π using a suitable ranking procedure.

Both steps will be explained in more detail further below. Before, however, we briefly return to the issue of data preprocessing.

4.1 Preprocessing

Recall that we assume objects to be represented in terms of feature vectors $\mathbf{x} = (x_1, \dots, x_d)$. As already mentioned, analogical reasoning based on analogical proportions assumes real-valued features to be normalized, with values in the unit interval $[0, 1]$. Therefore, if the k th feature is real-valued, we apply the linear transformation

$$x_k \leftarrow \frac{x_k - \min_k}{\max_k - \min_k}$$

where \min_k and \max_k denote, respectively, the smallest and largest value of that feature in the data. This transformation is applied separately to the training and the test data.

Prior to normalization, we apply a logarithmic transformation to some of the features. For each feature, we compute a numeric measure of the skewness of its distribution and the distribution of its log-transform (Joanes and Gill 1998). If the skewness of the latter is smaller than the skewness of the former, the log-transform is adopted, i.e., each value x_k is replaced by $\log(x_k)$; otherwise, the feature is left unchanged.

4.2 Analogical Prediction of Pairwise Preferences

The first step of able2rank, prediction of pairwise preferences, is based on analogical reasoning. The basic idea is as follows: Consider any pair of query objects $\mathbf{x}_i, \mathbf{x}_j \in Q$. Moreover, suppose we find a pair of objects $(z, z') \in \mathcal{D}_{pair}$ in the training data, i.e., a preference $z \succ z'$, such that $(z, z', \mathbf{x}_i, \mathbf{x}_j)$ are in analogical proportion. Then, this is taken as an indication in favor of the preference $\mathbf{x}_i \succ \mathbf{x}_j$. We refer to this principle as *analogical transfer* of preferences, because the observed preference $z \succ z'$ between objects z, z' is (hypothetically) transferred to $\mathbf{x}_i, \mathbf{x}_j$.

Since the training data is not necessarily coherent, both preferences $\mathbf{x}_i \succ \mathbf{x}_j$ and $\mathbf{x}_j \succ \mathbf{x}_i$ might be supported (via analogical transfer). We define

$$p_{i,j} = \frac{c_{i,j}}{c_{i,j} + c_{j,i}},$$

where $c_{i,j}$ is the number of preferences that support $\mathbf{x}_i \succ \mathbf{x}_j$, and $c_{j,i}$ the number of preferences supporting $\mathbf{x}_j \succ \mathbf{x}_i$. Here, instead of counting all preferences, we only consider

the most relevant ones, i.e., those with the largest degrees of analogical proportion. More specifically, let $V = V_{i,j} \cup V_{j,i}$ with

$$V_{i,j} = \bigcup_{(\mathbf{z}, \mathbf{z}') \in \mathcal{D}_{pair}} v(\mathbf{z}, \mathbf{z}', \mathbf{x}_i, \mathbf{x}_j) \quad (7)$$

$$V_{j,i} = \bigcup_{(\mathbf{z}, \mathbf{z}') \in \mathcal{D}_{pair}} v(\mathbf{z}, \mathbf{z}', \mathbf{x}_j, \mathbf{x}_i) \quad (8)$$

be the set of analogy scores for \mathbf{x}_i and \mathbf{x}_j , and $V^* \subset V$ the k largest scores in V . Then

$$c_{i,j} = |V_{i,j} \cap V^*|, \quad c_{j,i} = |V_{j,i} \cap V^*|. \quad (9)$$

The degrees of analogical proportion in (7–8) are given by (6) with $v \in \{v_A, v_{A'}, v_G, v_{MM}, v_{AE}, v_{AE'}\}$. Thus, v plays the role of a parameter of our method, just like the value k of relevant proportions taken into account; in our current implementation, we choose $k \in \{10, 15, 20\}$. Algorithm (1) provides a summary of this part of the method.

Algorithm 1 Analogy-based Pairwise Preferences (APP)

Require: \mathcal{D}_{pair}, Q

```

1:  $lstX_iX_j \leftarrow []$  // initialize the list
2: for all  $\mathbf{x}_i, \mathbf{x}_j \in Q$  do
3:   // binary vector: supporting  $\mathbf{x}_i \succ \mathbf{x}_j$  as 1 and 0 otherwise
4:    $X_iX_j \leftarrow []$ 
5:    $scores \leftarrow []$  // numeric vector: the degree of support
6:   for all  $\mathbf{z}, \mathbf{z}' \in \mathcal{D}_{pair}$  do
7:      $s_{ij} \leftarrow v(\mathbf{z}, \mathbf{z}', \mathbf{x}_i, \mathbf{x}_j)$ 
8:      $s_{ji} \leftarrow v(\mathbf{z}, \mathbf{z}', \mathbf{x}_j, \mathbf{x}_i)$ 
9:     if  $s_{ij} > s_{ji}$  then
10:        $scores.append(s_{ij})$ 
11:     else
12:        $scores.append(s_{ji})$ 
13:     end if
14:      $X_iX_j.append(\text{not}(\text{xor}([\mathbf{z} \succ \mathbf{z}'], [s_{ij} > s_{ji}])))$ 
15:   end for
16:   // re-arrange  $X_iX_j$  based on positions of sorted scores
17:    $X_iX_j \leftarrow X_iX_j(\text{arg sort}(scores))$ 
18:    $lstX_iX_j.append(X_iX_j)$ 
19: end for
20: return  $lstX_iX_j$ 

```

4.3 Rank Aggregation

To turn pairwise preferences into a total order, we make use of a rank aggregation method. More specifically, we apply the Bradley-Terry-Luce (BTL) model, which is well-known in the literature on discrete choice (Bradley and Terry 1952). It starts from the parametric model

$$\mathbf{P}(\mathbf{x}_i \succ \mathbf{x}_j) = \frac{\theta_i}{\theta_i + \theta_j}, \quad (10)$$

where $\theta_i, \theta_j \in \mathbb{R}_+$ are parameters representing the (latent) utility $U(\mathbf{x}_i)$ and $U(\mathbf{x}_j)$ of \mathbf{x}_i an \mathbf{x}_j , respectively. Thus, according to the BTL model, the probability to observe a preference in favor of a choice alternative \mathbf{x}_i , when being compared to any other alternative, is proportional to θ_i .

Given the data (9), i.e., the numbers $c_{i,j}$ informing about how often every preference $\mathbf{x}_i \succ \mathbf{x}_j$ has been observed, the parameter $\theta = (\theta_1, \dots, \theta_n)$ can be estimated by likelihood maximization:

$$\hat{\theta} \in \arg \max_{\theta \in \mathbb{R}^n} \prod_{1 \leq i \neq j \leq n} \left(\frac{\theta_i}{\theta_i + \theta_j} \right)^{c_{i,j}}$$

Finally, the predicted ranking π is obtained by sorting the items \mathbf{x}_i in descending order of their estimated (latent) utilities $\hat{\theta}_i$ (see Algorithm 2).

We note that many other rank aggregation techniques have been proposed in the literature and could principally be used as well; see e.g. (Ahmadi Fahandar, Hüllermeier, and Couso 2017). However, since BTL seems to perform very well, we did not consider any other method.

Algorithm 2 Rank Aggregation (RA)

Require: $lstX_iX_j, k$

```

1: Initialize comparison matrix  $C_{n \times n}$ 
2: for all  $X_iX_j$  in  $lstX_iX_j$  do
3:    $C_{i,j} \leftarrow \sum_{r=1}^k \mathbb{1}[X_iX_j(r) = 1]$ 
4:    $C_{j,i} \leftarrow \sum_{r=1}^k \mathbb{1}[X_iX_j(r) = 0]$ 
5: end for
6:  $\hat{\theta} \leftarrow \text{BTL}(C)$ 
7:  $\hat{\pi} \leftarrow \text{arg sort}(\hat{\theta})$ 
8: return  $\hat{\pi}$ 

```

5 Experiments

In order to study the practical performance of our proposed method, we conducted experiments on several real-world data sets, using ERR and SVM-Rank (cf. Section 2.1) as baselines for comparison.

5.1 Data

The data sets¹ are collected from various domains (e.g., sports, education, tourism) and comprise different types of feature (e.g., numeric, binary, ordinal). Table 1 provides a summary of the characteristics of the data sets. Here is a detailed description:

- **Decathlon:** This data contains rankings of the top 100 men’s decathletes worldwide in the years 2005 and 2006, with 10 numeric features associated with each athlete. Each feature is the performance achieved by the athlete in the corresponding discipline (e.g., the time in 100 meter race). The ranking of the decathletes is based on a scoring scheme, in which each performance is first scored in terms of a certain number of points (the mapping from performance to scores is non-linear), and the total score is obtained as the sum of the points over all 10 disciplines. In addition, the results of Olympic games Rio de Janeiro 2016 (24 instances) as well as under-20 world championships 2016 (22 instances) are considered. The data are extracted from the Decathlon2000 web site².

¹available at <https://cs.uni-paderborn.de/is/>

²www.decathlon2000.com

Table 1: Properties of data sets.

data set	domain	# instances	# features	numeric	binary	ordinal	name
Decathlon	Year 2005	100	10	x	–	–	D1
	Year 2006	100	10	x	–	–	D2
	Olympic games Rio de Janeiro 2016	24	10	x	–	–	D3
	Under-20 world championships 2016	22	10	x	–	–	D4
Bundesliga	Season 15/16	18	13	x	–	–	B1
	Season 16/17	18	13	x	–	–	B2
	Mid-Season 16/17	18	7	x	–	–	B3
	Season 16/17 Away	18	6	x	–	–	B4
FIFA	Year 2016	100	40	36	1	3	F1
	Year 2017	100	40	36	1	3	F2
	Year 2016 (Position:Streaker)	50	40	36	1	3	F3
	Year 2017 (Position:Streaker)	50	40	36	1	3	F4
Hotels	Düsseldorf	110	28	x	x	x	H1
	Frankfurt	149	28	x	x	x	H2
Uni. Rankings	Year 2015	100	9	x	–	–	U1
	Year 2014	100	9	x	–	–	U2
Volleyball WL	Group 3	12	15	x	–	–	V1
	Group 1	12	15	x	–	–	V2
Netflix	Germany	9	7	x	x	–	N1
	USA	13	7	x	x	–	N2

- **FIFA:** The FIFA Index website³ ranks the best football players in the world based on different metrics each year. These metrics belong to different categories, such as ball skills (ball control, dribbling, etc.), physical performance (acceleration, balance, etc.), defence (marking, slide tackling, etc.), and so on. We considered the list of top 100 footballers in the years 2016 and 2017, where each player is described by 40 attributes. Since the metrics of different types of players are not comparable, the overall evaluation of a player depends on his position (goal keeper, defender, streaker, etc.). Obviously, predicting the overall ranking is therefore a difficult task. In addition to the full data sets, we therefore also considered two position-specific data sets, namely the rankings for players with position streaker in the years 2016 and 2017.
- **Hotels:** This data set contains rankings of hotels in two major German cities, namely Düsseldorf (110 hotels) and Frankfurt (149 hotels). These rankings have been collected from TripAdvisor⁴ in September 2014. Each hotel is described in terms of a feature vector of length 28 (e.g., distance to city center, number of stars, number of rooms, number of user rating, etc). The way in which a ranking is determined by TripAdvisor on the basis of these features is not known (and one cannot exclude that additional features may play a role).
- **University Rankings:** This data includes the list of top 100 universities worldwide for the years 2014 and 2015. It is published by the Center for World University Rankings (CWUR). Each university is represented by 9 features such as national rank, quality of education, alumni employment, etc. Detailed information about how the ranking is determined based on these features can be found on

the CWUR website⁵.

- **Bundesliga:** This data set comprises table standings of 18 football teams in the German Bundesliga (German football league) for the seasons 2015/16 and 2016/17.⁶ Each team is described in terms of 13 features, such as matches, win, loss, draw, goals-for, goals-against, etc. To study the ability of knowledge transfer, we also included the table standing for the season 2016/17, in which only the statistics for away matches are considered (with 6 features). Another case is the table standing in the mid-season 2016/17 (i.e., only the first half of the season) with 7 features.
- **Volleyball WL:** This data contains the table standing for Group1 (statistics of 12 teams divided into subgroups, each with 9 matches) and Group3 (statistics of 12 teams, each with 6 matches) of volleyball world league 2017 extracted from the FIVB website⁷. There are 12 features in total, such as win, loss, number of (3-0, 3-1, 3-2, etc.) wins, sets win, sets loss, etc.
- **Netflix:** This data set includes the Netflix ISP speed index (extracted from Netflix website⁸ in August, 2017) for Germany and USA with 9 and 13 Internet providers, respectively. The Netflix ISP Speed Index is a measure of prime time Netflix performance on particular internet service providers (ISPs) around the globe. The rankings are represented by 7 (binary and numeric) features like speed, speed of previous month, fiber, cable, etc.

⁵www.cwur.org

⁶www.bundesliga.com

⁷worldleague.2017.fivb.com

⁸ispspeedindex.netflix.com

³www.fifaindex.com

⁴www.tripadvisor.com

5.2 Experimental Setting

Given the data sets for training and testing, able2rank first preprocesses the individual attributes as explained above (log-transformation and normalization). We also apply a standard normalization for the baseline methods (ERR and SVM), transforming each real-valued feature by standardization:

$$x \leftarrow \frac{x - \mu}{\sigma},$$

where μ and σ denote the empirical mean and standard deviation, respectively.

Recall that able2rank has two parameters to be tuned: The type of analogical proportion $v \in \mathcal{S}_v$, where $\mathcal{S}_v = \{v_A, v_{A'}, v_G, v_{MM}, v_{AE}, v_{AE'}\}$, and the number $k \in \mathcal{S}_k$ of relevant proportions considered for estimating pairwise preferences, where $\mathcal{S}_k = \{10, 15, 20\}$. We fixed these parameters in an (internal) 2-fold cross validation (repeated 5 times) on the training data, using simple grid search on $\mathcal{S}_v \times \mathcal{S}_k$ (i.e., trying all combinations). The combination (v^*, k^*) with the lowest cross-validated error d_{RL} is eventually adopted and used to make predictions on the test data (using the entire training data). The complexity parameter C of SVM is fixed in a similar way using an internal cross-validation.

5.3 Results

In our experiments, predictions were produced for certain parts \mathcal{D}_{test} of the data we collected, using other parts \mathcal{D}_{train} as training data; an experiment of that kind is denoted $\mathcal{D}_{train} \rightarrow \mathcal{D}_{test}$. The results of the conducted experiments are summarized in Table 2, with the best performance on each problem displayed in bold-face. Additionally, we report the parameters used by able2rank.

From the results, we conclude that able2rank is quite competitive in terms of predictive accuracy. In 8 of 11 cases, it achieves the best performance in terms of the average ranking loss d_{RL} (which translates to a p-value of around 10% when comparing able2rank and SVM with a pairwise sign test, and of around 0.5% for the comparison with ERR).

6 Conclusion and Future Work

This paper advocates the use of analogical reasoning in the context of preference learning. Building on the notion of analogical proportion, we formalize the heuristic principle suggesting that, if an alternative A is preferred to B , and C relates to D as A relates to B , then C is likely to be preferred to D . Based on this formalization, we develop a concrete method, able2rank, for the problem of object ranking. First experimental results on real-world data from different domains are quite promising and suggest that able2rank is competitive to state-of-the-art methods for object ranking.

In future work, we plan to elaborate more closely on the setting of *transfer learning*, because we believe our analogical approach to preference learning, like analogical reasoning in general, to be specifically useful for this purpose. In fact, preference transfer as defined in this paper only requires the relation \mathcal{R} to be evaluated separately for *source* objects A and B on the one side and *target* objects C and

Table 2: Results in terms of loss d_{RL} on the test data.

train \rightarrow test	(v^*, k^*)	able2rank	ERR	SVM
D1 \rightarrow D2	$v_{A'}, 10$	0.066	0.064	0.025
D3 \rightarrow D4	$v_{MM}, 10$	0.139	0.147	0.152
B1 \rightarrow B2	$v_A, 15$	0.046	0.144	0.059
B4 \rightarrow B2	$v_A, 20$	0.078	0.078	0.026
B3 \rightarrow B2	$v_A, 10$	0.000	0.033	0.013
F1 \rightarrow F2	$v_G, 20$	0.225	0.259	0.310
F3 \rightarrow F4	$v_G, 10$	0.158	0.364	0.169
H1 \rightarrow H2	$v_A, 20$	0.060	0.084	0.072
U1 \rightarrow U2	$v_{A'}, 15$	0.073	0.163	0.154
V1 \rightarrow V2	$v_{MM}, 10$	0.030	0.485	0.015
N1 \rightarrow N2	$v_A, 10$	0.013	0.090	0.077

D on the other side, but never between sources and targets; in principle, different specifications of \mathcal{R} could even be used for the source and the target. In the experiments conducted in this paper, some kind of knowledge transfer was already required, but source and target were still closely related and essentially from the same domain. As already said, we plan to extend these experiments toward problems where source and target are from different domains, and to generalize able2rank correspondingly.

Besides, the basic version of able2rank as presented in this paper ought to be improved and extended in different directions. For example, for the time being, the analogical proportion of feature vectors in able2rank is simply the average of the analogical proportions on the individual attributes. However, since different features may have a different influence on preferences and analogies, the selection or weighting of features appears to be quite important. Therefore, the development of methods for feature selection and weighting will be addressed in future work.

Furthermore, the current version of able2rank ignores possible interactions between features. For example, while a person may prefer color red to black for T-shirts, this preference might be reversed in the case of pullovers. Learning such interactions and capturing them in analogical inference is a highly non-trivial problem.

Last but not least, issues related to computational efficiency ought to be addressed. Since able2rank follows the lazy learning paradigm (Aha 1997), just like instance-based learning and nearest neighbor prediction, it can learn very easily (simply by storing observed preferences) but is relatively costly at prediction time. In comparison to nearest neighbor classification, for example, the cost is even higher due to the need to iterate over *pairs* of objects in the training data to find the highest analogical proportions. Thus, implementing “highest analogy” search in a naive way, the complexity scales quadratically with the size of the training data. Consequently, there is a need to reduce this complexity by means of suitable data structures and algorithms for efficient retrieval of analogy pairs. Besides, other means for complexity reduction could be considered, such as instance selection or editing strategies like those used in case-based reasoning (McKenna and Smyth 2000; Delany and Cunningham 2004).

References

- Aha, D., ed. 1997. *Lazy Learning*. Kluwer Academic Publ.
- Ahmadi Fahandar, M.; Hüllermeier, E.; and Couso, I. 2017. Statistical inference for incomplete ranking data: The case of rank-dependent coarsening. In Precup, D., and Teh, Y., eds., *Proceedings ICML–2017, 34th International Conference on Machine Learning*.
- Beltran, W. C.; Jaudoin, H.; and Pivert, O. 2014. Analogical prediction of null values: The numerical attribute case. In *Proc. ADBIS, 18th East European Conference on Advances in Databases and Information Systems*, 323–336. Ohrid, Macedonia: Springer.
- Beltran, W. C.; Prade, H.; and Richard, G. 2016. Constructive solving of raven’s IQ tests with analogical proportions. *Int. J. Intell. Syst.* 31(11):1072–1103.
- Billingsley, R.; Prade, H.; Richard, G.; and Williams, M. 2017. Towards analogy-based decision – A proposal. In *Proc. FQAS, 12th International Conference on Flexible Query Answering Systems*, 28–35.
- Bounhas, M.; Prade, H.; and Richard, G. 2014a. Analogical classification: A new way to deal with examples. In *Proc. ECAI, 21st European Conference on Artificial Intelligence*, 135–140.
- Bounhas, M.; Prade, H.; and Richard, G. 2014b. Analogical classification: Handling numerical data. In *Proc. SUM, 8th International Conference on Scalable Uncertainty Management*, 66–79.
- Bradley, R., and Terry, M. 1952. The rank analysis of incomplete block designs I. The method of paired comparisons. *Biometrika* 39:324–345.
- Cheng, W.; Hühn, J.; and Hüllermeier, E. 2009. Decision tree and instance-based learning for label ranking. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML ’09*, 161–168. New York, NY, USA: ACM.
- Cohen, W. W.; Schapire, R. E.; and Singer, Y. 1999. Learning to order things. *J. Artif. Int. Res.* 10(1):243–270.
- Delany, S., and Cunningham, P. 2004. An analysis of case-base editing in an Spam filtering system. In *Proc. ECCBR–2004, European Conference on Case-Based Reasoning*, 128–141. Heidelberg: Springer-Verlag.
- Dubois, D.; Prade, H.; and Richard, G. 2016. Multiple-valued extensions of analogical proportions. *Fuzzy Sets and Systems* 292:193–202.
- Fürnkranz, J., and Hüllermeier, E. 2011. *Preference Learning*. Springer-Verlag.
- Fürnkranz, J.; Hüllermeier, E.; and Vanderlooy, S. 2009. *Binary Decomposition Methods for Multipartite Ranking*. Berlin, Heidelberg: Springer Berlin Heidelberg. 359–374.
- Har-Peled, S.; Roth, D.; and Zimak, D. 2002. *Constraint Classification: A New Approach to Multiclass Classification*. Berlin, Heidelberg: Springer Berlin Heidelberg. 365–379.
- Hug, N.; Prade, H.; Richard, G.; and Serrurier, M. 2016. Analogy in recommendation. Numerical vs. ordinal: A discussion. In *IEEE International Conference on Fuzzy Systems*, 2220–2226.
- Joachims, T. 2002. Optimizing search engines using click-through data. In *Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-02)*, 133–142. ACM Press.
- Joanes, D. N., and Gill, C. A. 1998. Comparing measures of sample skewness and kurtosis. *Journal of the Royal Statistical Society: Series D (The Statistician)* 47:183–189.
- Kamishima, T., and Akaho, S. 2006. Supervised ordering by regression combined with thurstone’s model. *Artificial Intelligence Review* 25(3):231–246.
- Kamishima, T.; Kazawa, H.; and Akaho, S. 2010. A survey and empirical comparison of object ranking methods. In Fürnkranz, J., and Hüllermeier, E., eds., *Preference Learning*. Berlin, Heidelberg: Springer-Verlag. 181–202.
- McKenna, E., and Smyth, B. 2000. Competence-guided editing methods for lazy learning. In *Proceedings ECAI–2000, 14th European Conference on Artificial Intelligence*, 60–64.
- Melnikov, V.; Gupta, P.; Frick, B.; Kaimann, D.; and Hüllermeier, E. 2016. Pairwise versus pointwise ranking: A case study. *Schedae Informaticae* 25:73–83.
- Miclet, L., and Prade, H. 2009. *Handling Analogical Proportions in Classical Logic and Fuzzy Logics Settings*. Berlin, Heidelberg: Springer Berlin Heidelberg. 638–650.
- Pirlot, M.; Prade, H.; and Richard, G. 2016. Completing preferences by means of analogical proportions. In *Proc. MDAI, 13th International Conference on Modeling Decisions for Artificial Intelligence*, 135–147.
- Prade, H., and Richard, G. 2013. From analogical proportion to logical proportions. *Logica Universalis* 7(4):441–505.
- Prade, H., and Richard, G. 2017. Analogical proportions and analogical reasoning – An introduction. In *Proceedings ICCBR, Int. Conf. on Case-Based Reasoning*, 16–32.
- Vembu, S., and Gärtner, T. 2010. Label ranking: a survey. In Fürnkranz, J., and Hüllermeier, E., eds., *Preference Learning*. Springer-Verlag.