Predicting Partial Orders: Ranking with Abstention

Weiwei Cheng[1], Michaël Rademaker[2], Bernard De Baets[2], Eyke Hüllermeier[1]

[1] *Knowledge Engineering & Bioinformatics Lab, University of Marburg, Germany*

[2] *Research Unit Knowledge-Based Systems, Ghent University, Belgium*

RANKING WITH REJECT OPTION

Existing learning algorithms for ranking (*label ranking*, *object ranking*, *instance ranking*, etc.) produce a **total order** of alternatives. For example, consider the problem of learning a reviewer's preferences on papers:

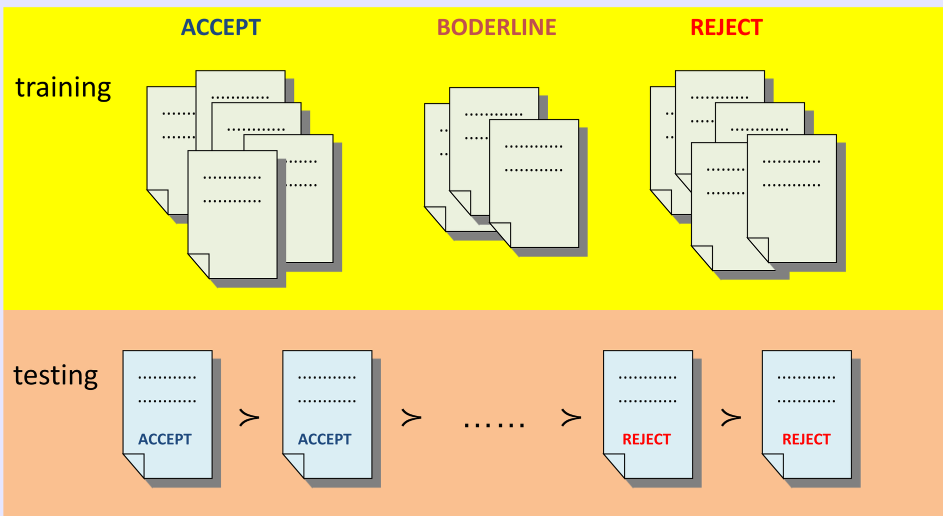

We propose learners that are able to partly **abstain**. For a pair of items a and b to be ranked, the learner can

- predict $a > b$ or $b > a$, or
- abstain from the prediction (reject option).

The learner should be consistent (**transitive** & **acyclic**).

The roadmap of our approach:

1. Predicting a binary preference relation P that specifies, for each pair of alternatives a and b, a degree of uncertainty regarding their relative comparison;
2. Deriving a (strict) partial order as much as possible in agreement with P.

EVALUATION MEASURES

Considering two parital order relations $\sqsupset_*$ and $\sqsupset$, we have

	$a \sqsupset_* b$	$b \sqsupset_* a$	$a \perp_* b$
$a \sqsupset b$	C	D	×
$b \sqsupset a$	D	C	×
$a \perp b$	×	×	×

C: concordant D: discordant

As ranker now has the ability to reject predictions, there is a trade-off between **correctness** $\frac{|C|-|D|}{|C|+|D|}$ and **completeness** $\frac{|C|+|D|}{|\sqsupset_*|}$.

PREDICTING A PREFERENCE RELATION

A preference relation $P : A \times A \longrightarrow [0,1]$ provides a measure of support for the pairwise preference $a > b$, with $P(a,b) = 1 - P(b,a)$ for all $a, b \in A$.

We use a generic approach that can turn every ranker into a partial ranker via **ensembling**.

1. With a ranker L, train k ranking models $M_1 \ldots M_k$ by resampling from the original data set, i.e., by k bootstrap samples. By querying these models, k rankings $\succ_1 \ldots \succ_k$ will be produced;
2. For each pair of alternatives a and b, we define the degree of preference $P(a,b) = \frac{1}{k}\left|\{i \,|\, a \succ_i b\}\right|$.

PREDICTING A STRICT PARTIAL ORDER

Based on P, we seek to induce a (partial) order relation, $\mathcal{R} : \mathbf{A} \times \mathbf{A} \to \{0,1\}$ with a threshold α:

$$\mathcal{R}_\alpha = \{(a,b) \,|\, P(a,b) \geq \alpha\}$$

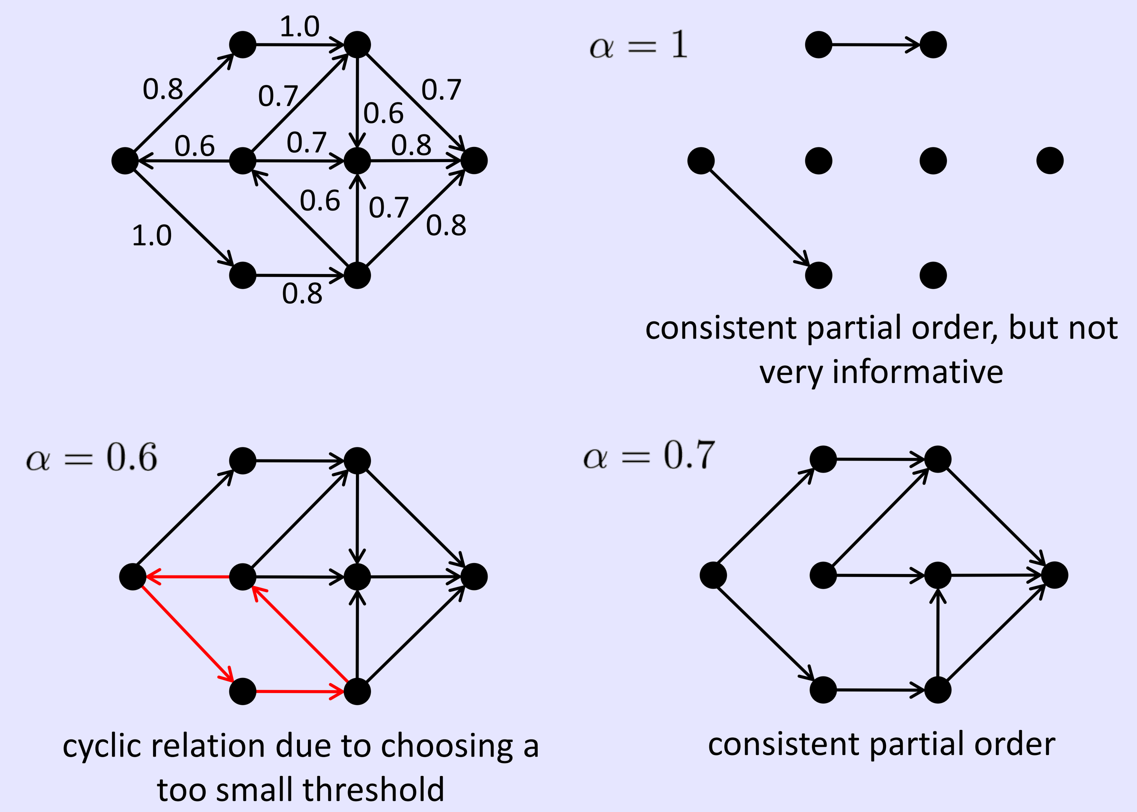

Finding a minimal α (denoted as α^*) such that the transitive closure of $\mathcal{R}_\alpha$ (denoted as $\overline{\mathcal{R}}_\alpha$) is a strict partial order relation:

- The domain of α can be restricted to $\{0, 1/k, 2/k, \ldots, 1\}$;
- If $\mathcal{R}_\alpha$ is cyclic, $\mathcal{R}_\beta$ is cyclic as well, unless $\beta > \alpha$.

Moreover, we can show that

$$\alpha_u = 1 \geq \alpha^* \geq \alpha_l = \tfrac{1}{k} + \max_{a,b} \min(P(a,b), P(b,a)).$$

An algorithm of complexity $\mathcal{O}(|\mathbf{A}|^3)$:

Repeat until $\alpha_u - \alpha_l < 1/k$

1. Set α to the middle point between α_u and α_l
2. Compute $\mathcal{R}_\alpha$
3. Compute $\overline{\mathcal{R}}_\alpha$ (e.g., using the Floyd-Warshall algorithm)
4. If $\overline{\mathcal{R}}_\alpha$ is a partial order, set α_u to α
5. Else set α_l to α

EXPERIMENTAL RESULTS

Results on bi-partite ranking problems

data set	#attributes	#instances	with abstention	w/o abstention	completeness
breast	9	286	0.330±0.150	0.318±0.141	0.578±0.074
breast-w	9	699	0.988±0.014	0.987±0.015	0.982±0.015
horse colic	22	368	0.734±0.135	0.697±0.142	0.790±0.044
credit rating	15	690	0.858±0.062	0.827±0.065	0.888±0.038
credit german	20	1000	0.610±0.088	0.568±0.084	0.741±0.060
pima diabetes	8	768	0.684±0.084	0.666±0.086	0.819±0.047
heart statlog	13	270	0.811±0.102	0.797±0.101	0.890±0.060
hepatitis	19	155	0.709±0.292	0.697±0.271	0.797±0.084
ionosphere	34	351	0.771±0.174	0.722±0.190	0.814±0.098
kr-vs-kp	36	3196	0.992±0.006	0.980±0.007	0.991±0.006
labor	16	57	0.990±0.049	0.985±0.060	0.989±0.052
mushroom	22	8124	1.000±0.000	1.000±0.000	0.808±0.017
thyroid disease	29	3772	0.890±0.071	0.883±0.070	0.928±0.040
sonar	60	206	0.684±0.224	0.575±0.271	0.575±0.056
tic-tac-toe	9	958	0.253±0.127	0.221±0.120	0.908±0.013
vote	16	435	0.981±0.032	0.976±0.036	0.913±0.035

Trade-off between correctness and completeness by varying α

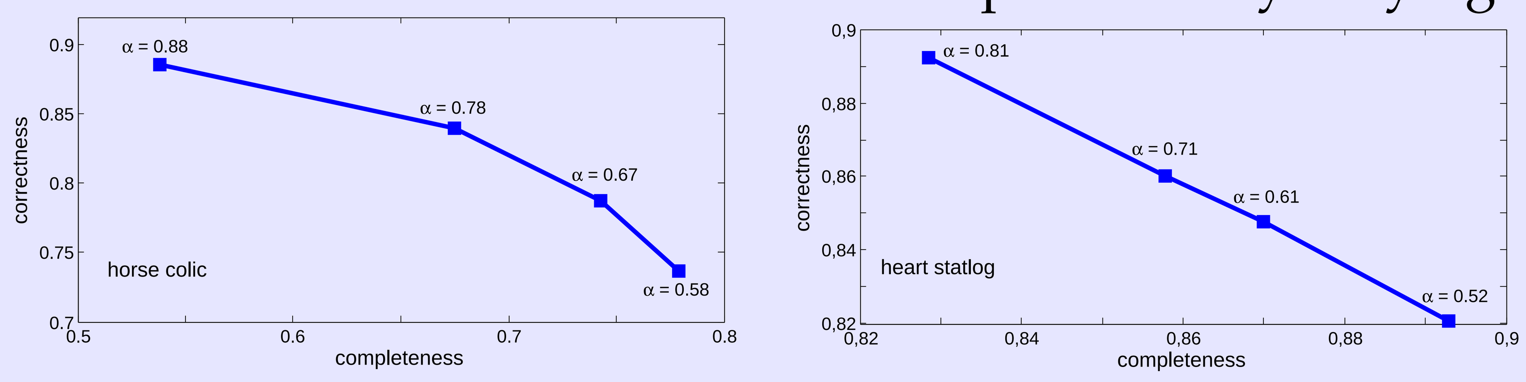

ACKNOWLEDGEMENT: Weiwei Cheng was financially supported by the UNESCO ECML PKDD 2010 Conference Grant

ECML PKDD 2010, *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, Barcelona, September 2010