Preference Learning using Statistical Methods for Label Ranking

Weiwei Cheng

amazon.com

Joint work with

March, 2014

Preferences are Ubiquitous

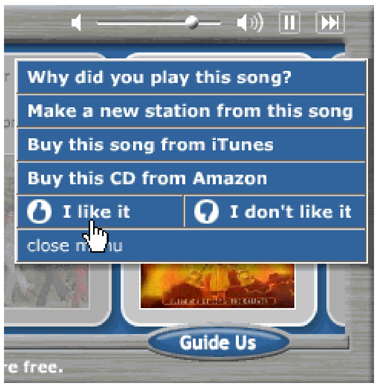

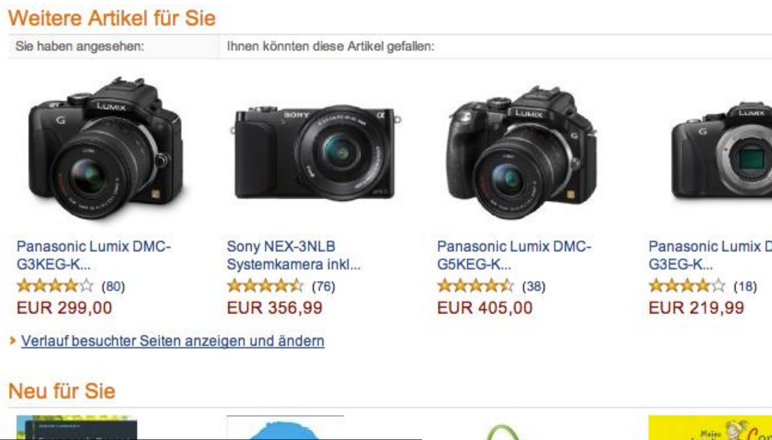

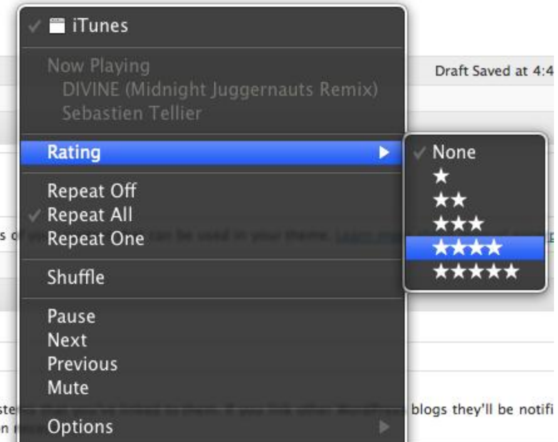

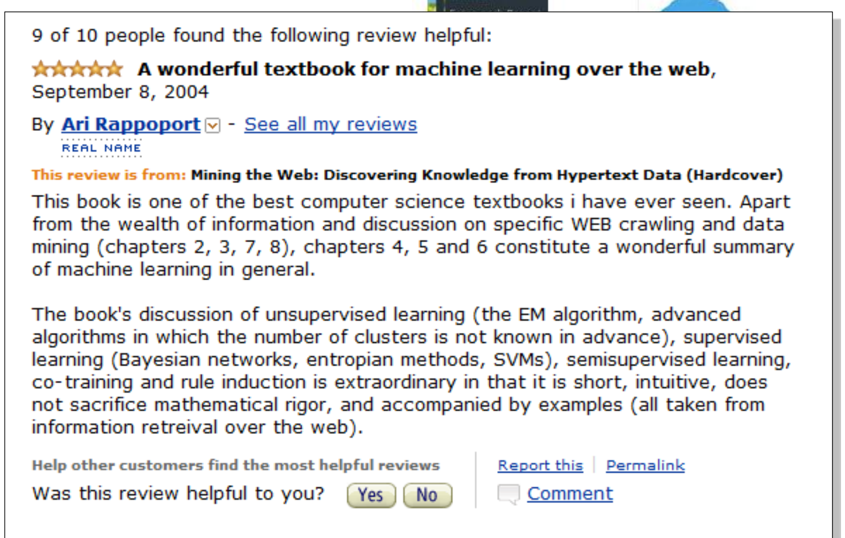

Preferences are Ubiquitous

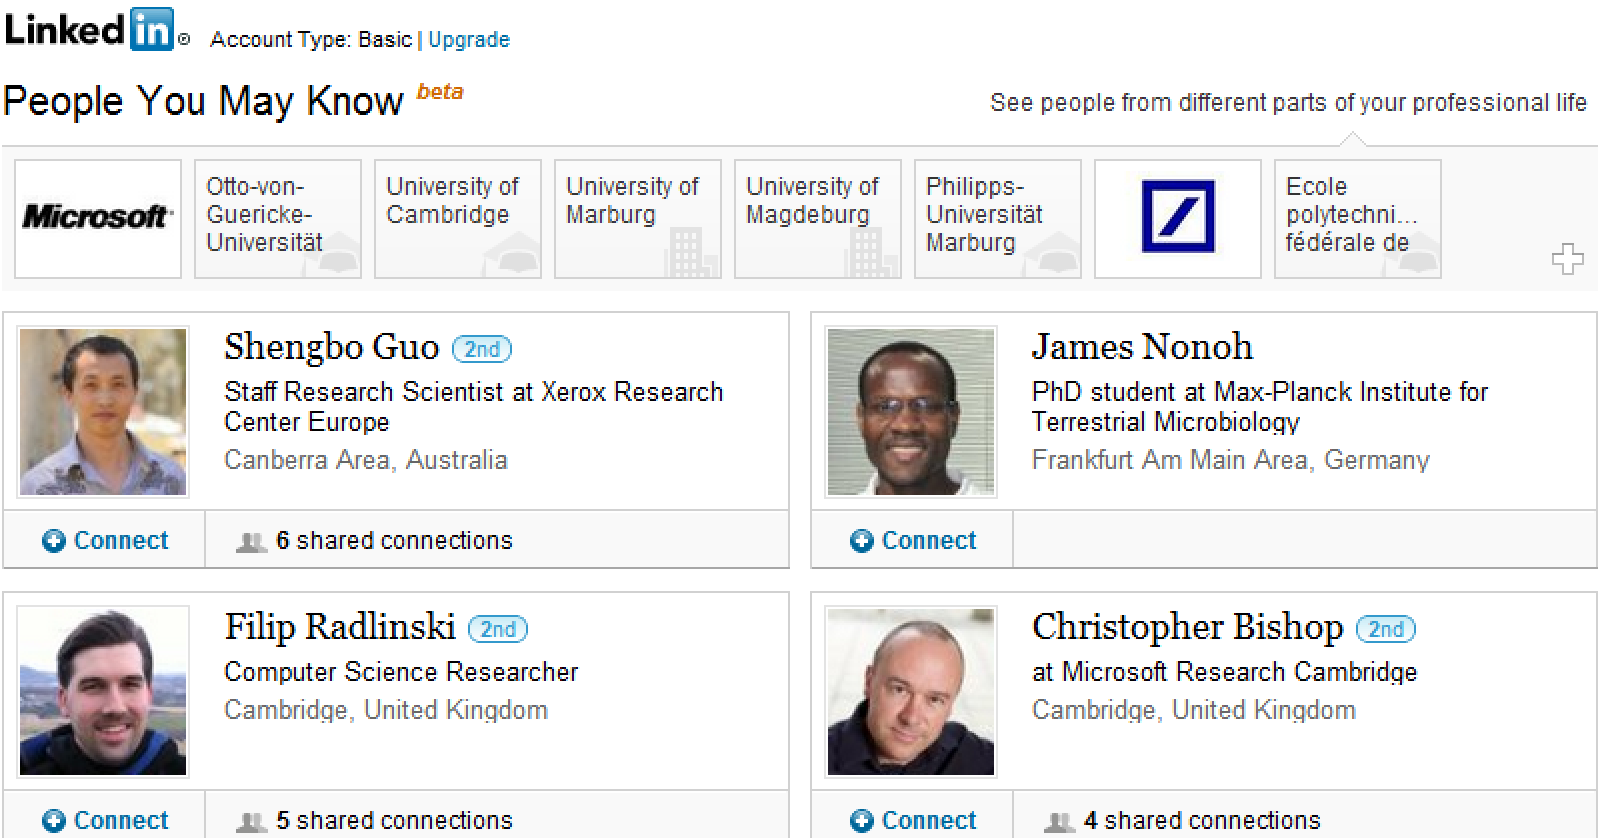

Preferences are Ubiquitous

Google

michael jordan

Web Images Maps Shopping News More Search tools

About 197,000,000 results (0.25 seconds)

Michael Jordan - Wikipedia, the free encyclopedia ← **NOT CLICKED**
en.wikipedia.org/wiki/**Michael_Jordan**
Michael Jeffrey Jordan (born February 17, 1963), also known by his initials, **MJ**, is an American former professional basketball player, entrepreneur, and majority ...
James R. Jordan, Sr. - List of career achievements - Emsley A. Laney High School

Michael Jordan | EECS at UC Berkeley ← **CLICKED**
www.eecs.berkeley.edu/Faculty/Homepages/**jordan**.html
Michael I. **Jordan** is the Pehong Chen Distinguished Professor in the Department of ...
F. R. Bach and M. **Jordan**, "**Learning** spectral clustering, with application to ...

Michael Jordan Stats, Bio - ESPN
espn.go.com/nba/player/_/id/1035/**michael-jordan**
Get the latest news, career stats and more about guard **Michael Jordan** on ESPN. com.

Michael Jordan NBA & ABA Stats | Basketball-Reference.com
www.basketball-reference.com › Players › J
Michael Jeffrey **Jordan** (Air **Jordan**, M.J.). Position: Guard-Forward ▪ Height: 6-6 ▪ Weight: 195 lbs. Born: February 17, 1963 in Brooklyn, New York High School: ...

Preferences Learning Settings

- **binary vs. graded** (e.g., relevance judgments vs. ratings)
- **absolute vs. relative** (e.g., assessing single alternatives vs. comparing pairs)
- **explicit vs. implicit** (e.g., direct feedback vs. click-through data)
- **structured vs. unstructured** (e.g., ratings on a given scale vs. free text)
- **single user vs. multiple users** (e.g., document keywords vs. social tagging)
- **single vs. multi-dimensional**
- ...

A wide spectrum of learning problems!

Preference Learning Tasks

	representation		type of preference information		
task	**input**	**output**	**training**	**prediction**	**ground truth**
collaborative filtering	identifier	identifier	absolute ordinal	absolute ordinal	absolute ordinal
multi-label classification	feature	identifier	absolute binary	absolute binary	absolute binary
multilabel ranking	feature	identifier	absolute binary	ranking	absolute binary
graded multilabel classification	feature	identifier	absolute ordinal	absolute ordinal	absolute ordinal
label ranking	feature	identifier	relative binary	ranking	ranking
object ranking	feature	--	relative binary	ranking	ranking or subset
instance ranking	feature	identifier	absolute ordinal	ranking	absolute ordinal

generalized classification (collaborative filtering – label ranking)

ranking (label ranking – instance ranking)

Two main directions: (1) ranking and variants (2) generalizations of classification

Agenda

1. Introduction to Preference Learning
2. **Label Ranking**
3. Extensions and Applications
4. Conclusions

Label Ranking – An Example

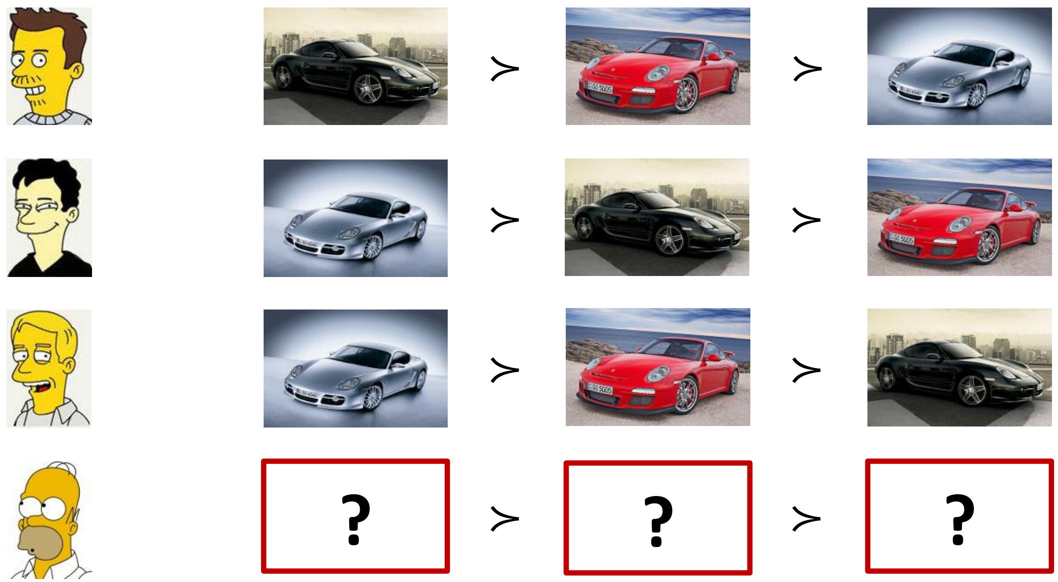

Instances are mapped to **total orders** over a fixed set of alternatives/labels.

Label Ranking: Training Data

TRAINING

X1	X2	X3	X4	Preferences
0.34	0	10	174	A ≻ B, C ≻ D
1.45	0	32	277	B ≻ C
1.22	1	46	421	B ≻ D, A ≻ D, C ≻ D, A ≻ C
0.74	1	25	165	C ≻ A, C ≻ D, A ≻ B
0.95	1	72	273	B ≻ D, A ≻ D
1.04	0	33	158	D ≻ A, A ≻ B, C ≻ B, A ≻ C

Instances are associated with pairwise preferences between labels.

... no demand for full rankings!

Label Ranking: Prediction

PREDICTION				A	B	C	D
0.92	1	81	382	?	?	?	?

new instance — ranking ?

Label Ranking: Prediction

PREDICTION				A	B	C	D
0.92	1	81	382	4	1	3	2

A ranking of all labels

new instance

$\pi(i)$ = position of *i*-th label

Label Ranking: Prediction

PREDICTION

0.92	1	81	382	4	1	3	2

A ranking of all labels

LOSS

GROUND TRUTH

0.92	1	81	382	2	1	3	4

SPEARMAN

$$L(\pi,\sigma) = \sqrt{\sum_{i=1}^{n} \left(\pi(i) - \sigma(j)\right)^2}$$

LOSS

$$\rho = 1 - \frac{6\, L^2(\pi,\sigma)}{n(n^2-1)}$$

RANK CORRELATION

Label Ranking: Prediction

PREDICTION

0.92	1	81	382	4	1	3	2

A ranking of all labels

LOSS

GROUND TRUTH

0.92	1	81	382	2	1	3	4

KENDALL

$$L(\pi, \sigma) = \sum_{1 \le i < j \le k} [\![\big(\pi(i) - \pi(j)\big) \cdot \big(\sigma(i) - \sigma(j)\big) < 0]\!]$$

LOSS

$$\tau = 1 - \frac{4\, L(\pi, \sigma)}{k(k-1)}$$

RANK CORRELATION

Learning Techniques

How to learn a label ranker $h : X \to S_n$?

The output space is complex ...

The Permutation Space

The output space is the class of permutations (symmetric group):

symmetric group S_3

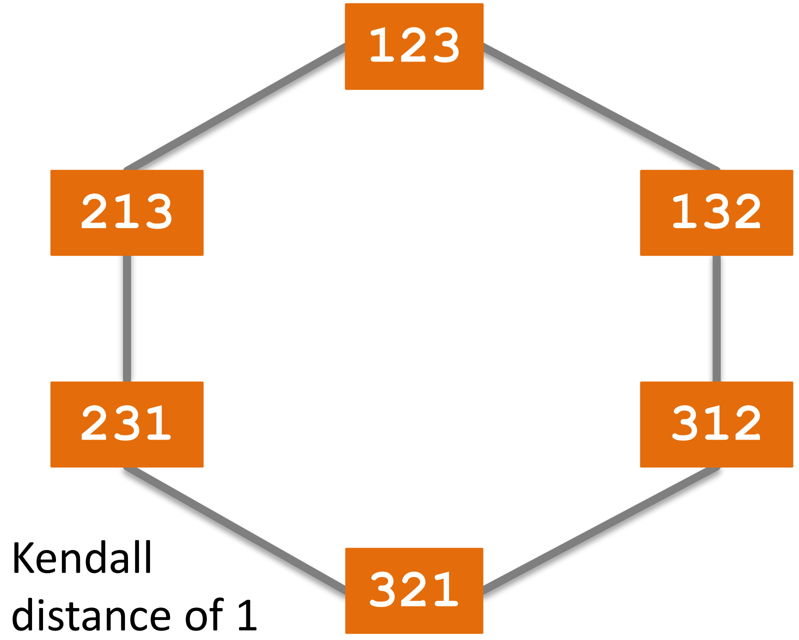

The Permutation Space

symmetric group S_4

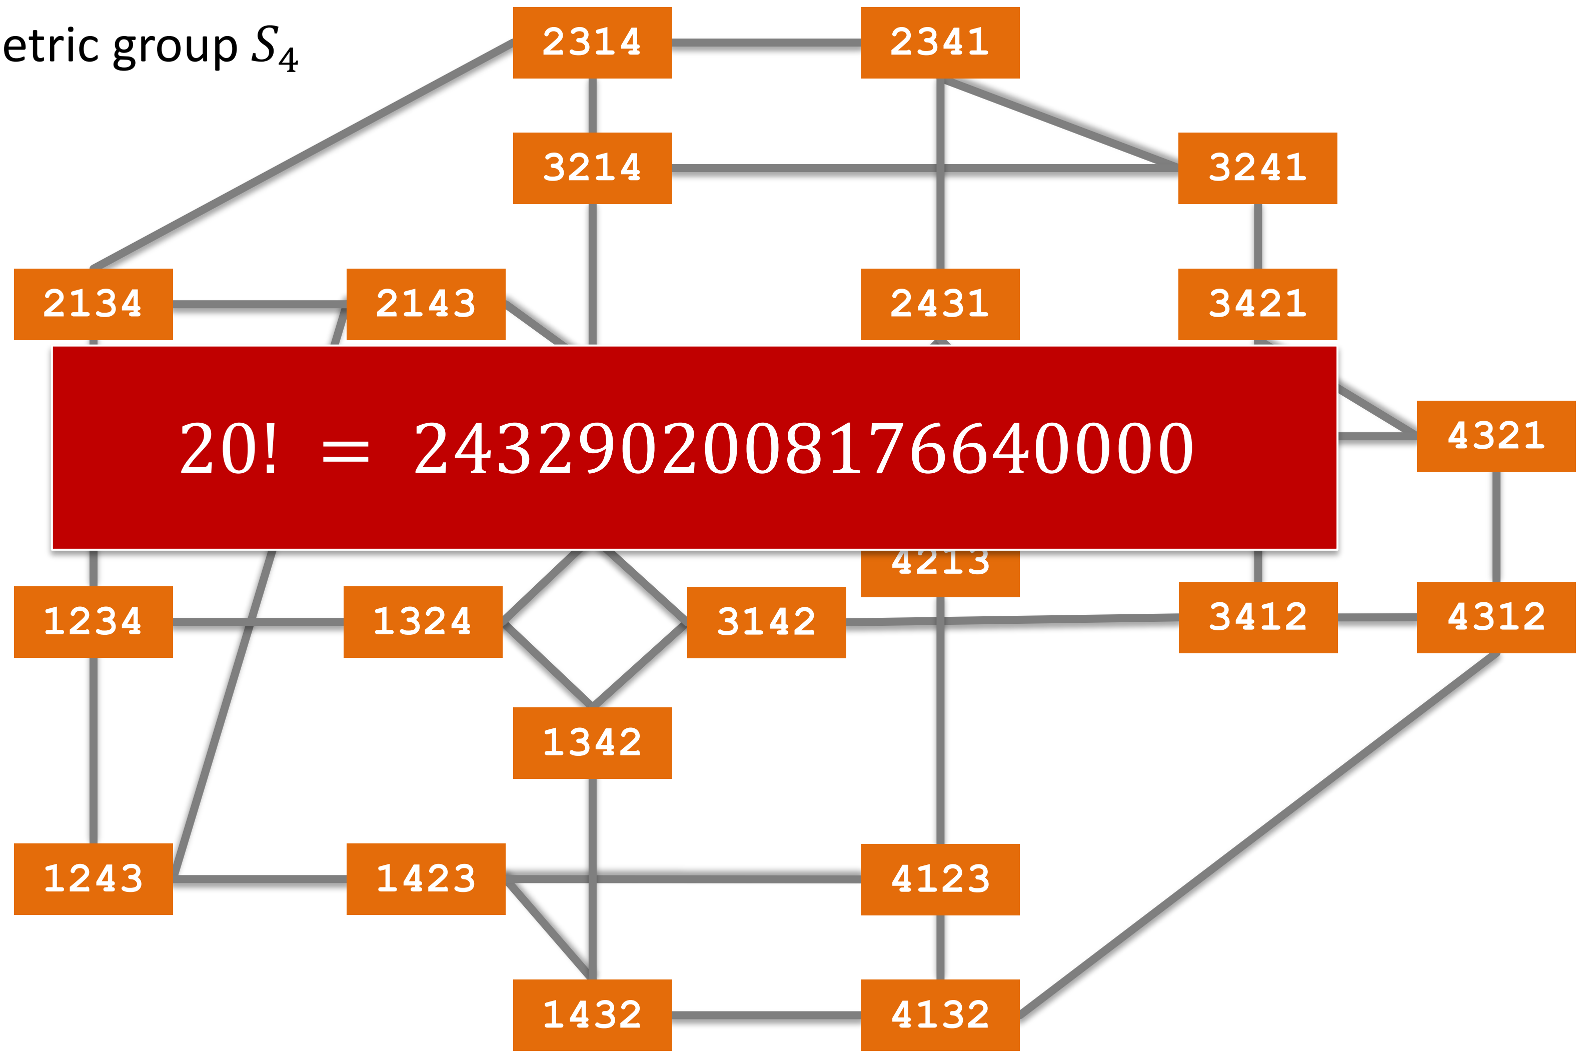

$$20! = 2432902008176640000$$

Learning Techniques

How to learn a label ranker $h : X \to S_n$?

Two approaches:

- Reduction to simpler problems (e.g., binary classification)
- Probabilistic modeling and statistical inference

Learning Techniques

reduction to binary classification	ranking by pairwise comparison [Hüllermeier et al., AI 08]	learning pairwise preferences
	constraint classification [Har-Peled et al., NIPS 02]	learning utility functions
boosting	log-linear models for label ranking [Dekel et al., NIPS 03]	
structured output prediction, margin maximization	structured output prediction [Vembu et al., UAI 09]	structured prediction
	local prediction (lazy learning) [Brinker et al. ECML 06 , Cheng et al., ICML 09]	
statistical inference	label ranking with probabilistic models [Cheng et al., ICML 09, Cheng et al., ICML 10]	

Learning Techniques

How to learn a label ranker $h : X \to S_n$?

Two approaches:

- Reduction to simpler problems (e.g., binary classification)
- Probabilistic modeling and statistical inference

Probabilistic Label Ranker

input

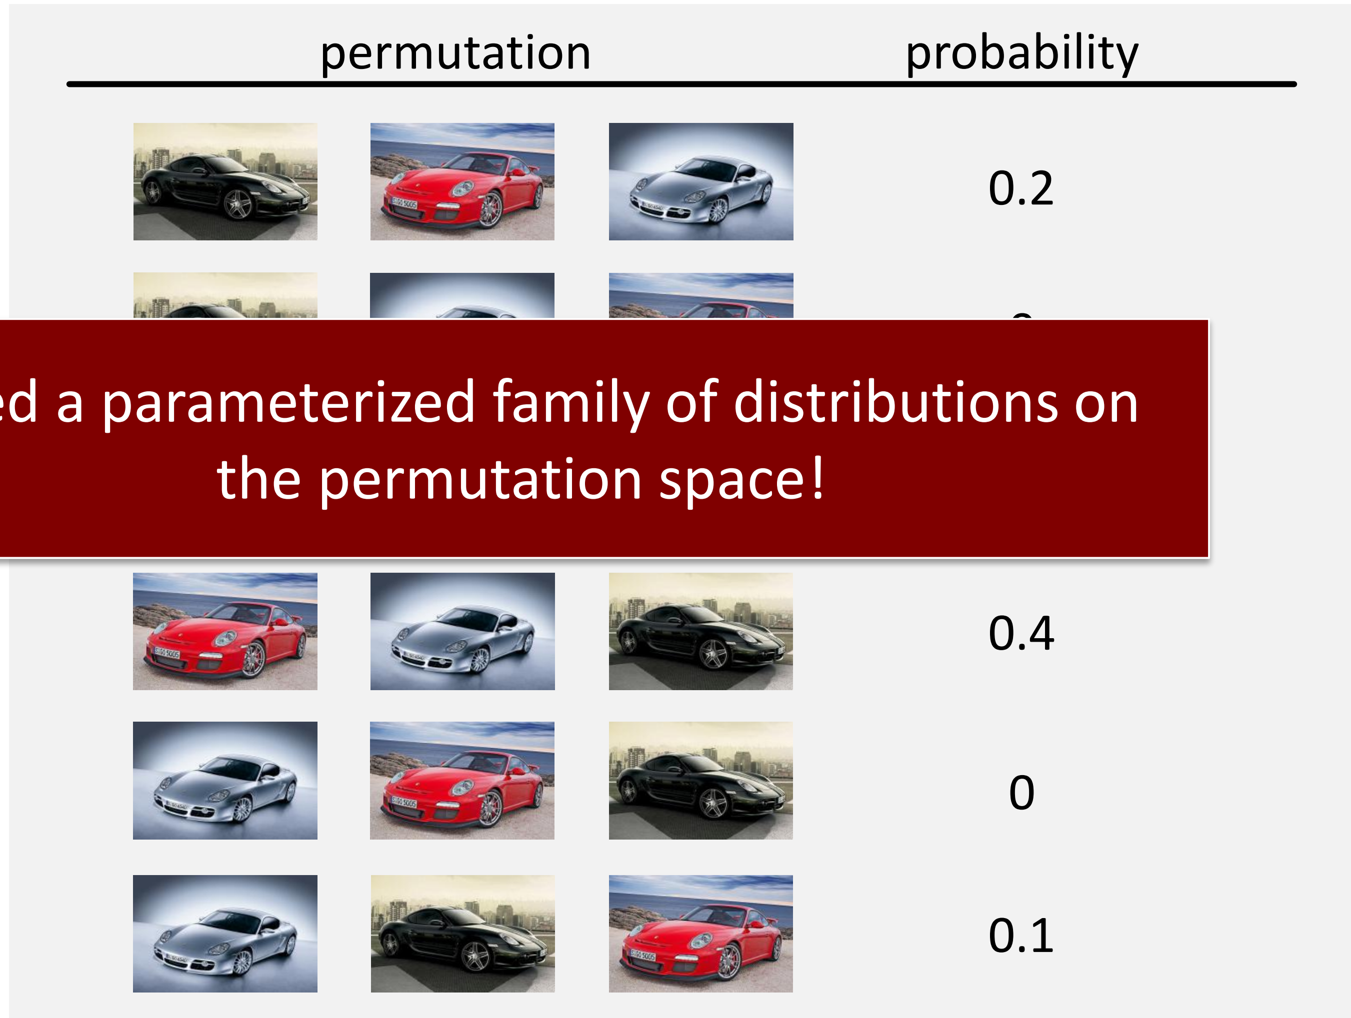

Need a parameterized family of distributions on the permutation space!

Label Ranking with Probabilistic Models

statistical ranking models	**machine learning techniques**
Mallows model	instance-based learning
Plackett-Luce model	generalized linear model

The Mallows Model

... is a **distance-based** model from the exponential family:

$$\mathbf{P}(\sigma \mid \pi, \theta) = \frac{\exp(-\theta\, \Delta(\sigma, \pi))}{\phi(\theta)}$$

center ranking (π) — spread (θ) — normalization constant (ϕ(θ))

where $\Delta(\cdot,\cdot)$ is a metric (i.e., distance measure) on rankings.

The probability of a ranking is higher if it is close to the mode, i.e., the center ranking of the distribution.

Some Common Choices of Δ

Kendall's tau ★

$$\mathrm{T}(\pi,\sigma) = \sum_{i<j} [\![\big(\pi(i)-\pi(j)\big)\cdot\big(\sigma(i)-\sigma(j)\big) < 0]\!]$$

Spearman's rho

$$\mathrm{R}(\pi,\sigma) = \sqrt{\sum_i \big(\pi(i)-\sigma(j)\big)^2}$$

Spearman's footrule

$$\mathrm{F}(\pi,\sigma) = \sum_i \, |\, \pi(i)-\sigma(j)\, |$$

Hamming

$$\mathrm{H}(\pi,\sigma) = \sum_i \, [\![\pi(i) \neq \sigma(i)]\!]$$

For example:

$$\pi = (1\ 2\ 3\ 4), \sigma = (1\ 4\ 2\ 3)$$

$$\mathrm{T}(\pi,\sigma) = 2$$
$$\mathrm{R}(\pi,\sigma) = 2.45$$
$$\mathrm{F}(\pi,\sigma) = 4$$
$$\mathrm{H}(\pi,\sigma) = 3$$

Label Ranking with Probabilistic Models

statistical ranking models	**machine learning techniques**
Mallows model	instance-based learning
Plackett-Luce model	generalized linear model

The Plackett-Luce Model

… is a **multistage** model specified by a vector $\boldsymbol{v} = (v_1, \dots, v_n) \in \mathrm{R}_+^n$:

$$\mathbf{P}(\sigma \mid \boldsymbol{v}) = \prod_{i=1}^{n} \frac{v_{\sigma^{-1}(i)}}{v_{\sigma^{-1}(i)} + v_{\sigma^{-1}(i+1)} + \cdots + v_{\sigma^{-1}(n)}}$$

where $\sigma^{-1}(i)$ is the index of the label ranked at position i.

A ranking is produced by choosing labels one by one, with a probability proportional to their respective "skills".

The Plackett-Luce Model

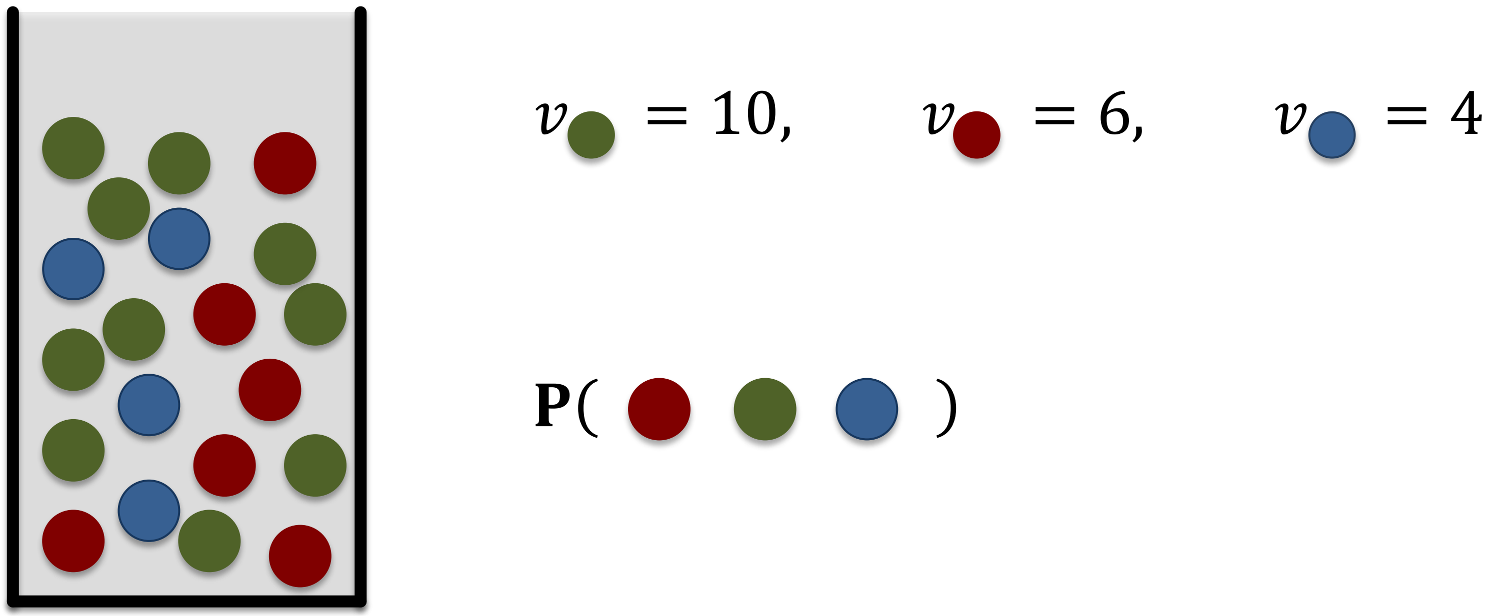

The Plackett-Luce Model

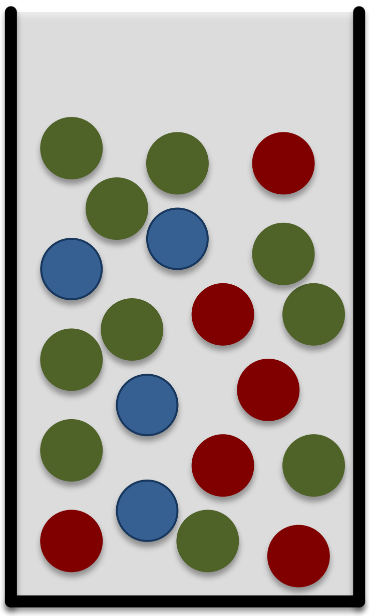

v $= 10,$ v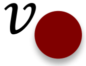 $= 6,$ v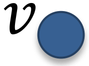 $= 4$

$\mathbf{P}($ 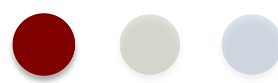 $) = \frac{6}{20}$

The Plackett-Luce Model

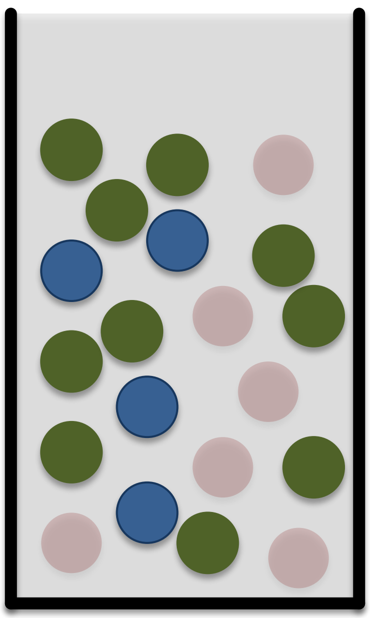

v  $= 10,$ v 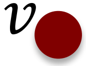 $= 6,$ v 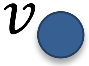 $= 4$

$\mathbf{P}($ 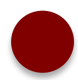 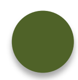 $) = \frac{6}{20} \times \frac{10}{14}$

The Plackett-Luce Model

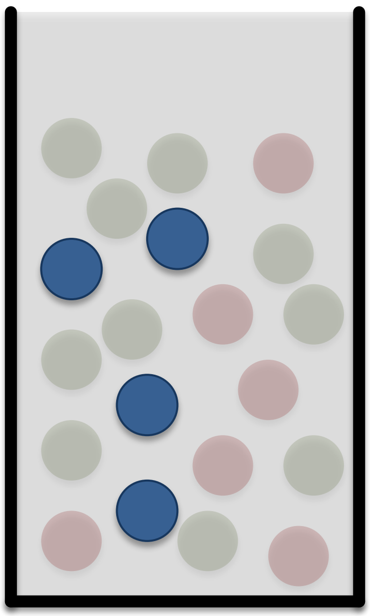

v  $= 10,$ v 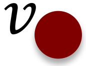 $= 6,$ v 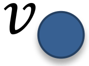 $= 4$

$\mathbf{P}($ 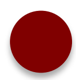 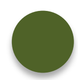 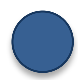 $) = \frac{6}{20} \times \frac{10}{14} \times \frac{4}{4}$

$= \frac{3}{14}$

The Plackett-Luce Model

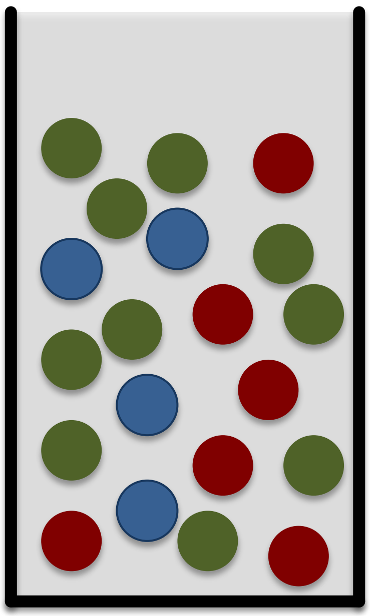

v $= 10,$ v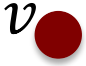 $= 6,$ v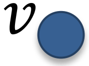 $= 4$

$\mathbf{P}($ 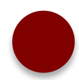 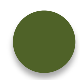 $) = \frac{6}{16} \times \frac{10}{10}$

$= \frac{3}{8}$

Label Ranking with Probabilistic Models

statistical ranking models	**machine learning techniques**
Mallows model	instance-based learning
Plackett-Luce model	generalized linear model

Results from: Cheng and Hüllermeier, ICML 09; Cheng et al., ICML 10

Instance-Based Approaches

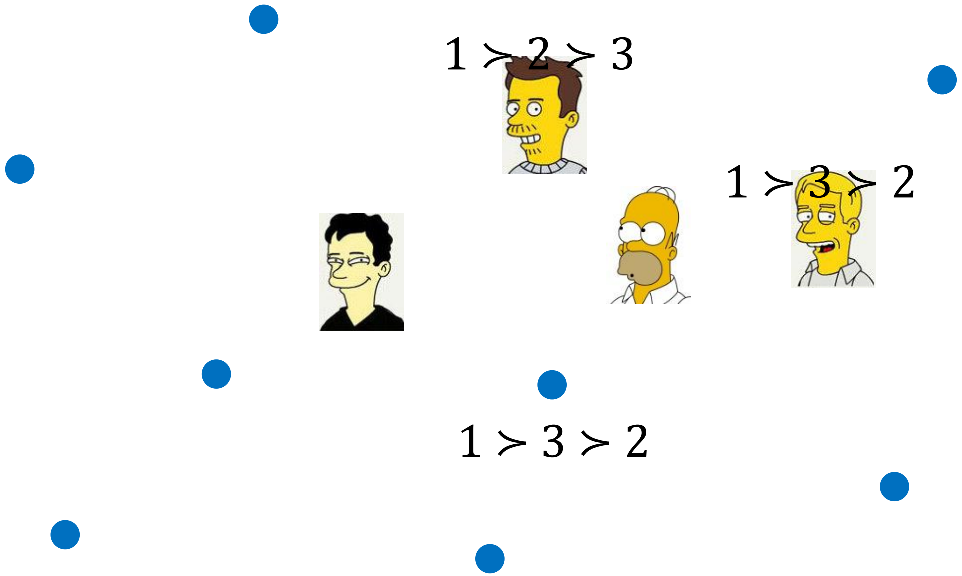

- Target function $X \to \Omega$ is estimated (on demand) in a local way.
- Distribution of rankings is (approx.) constant in a local region.
- Core part is to estimate the locally constant model.

Instance-Based Approaches

- Output (ranking) of an instance $\boldsymbol{x}$ is generated according to a distribution $\mathbf{P}(\cdot \mid \boldsymbol{x})$ on Ω.

- This distribution is (approximately) constant within the local region under consideration.

- Nearby preferences are considered as a sample generated by $\mathbf{P}$, which is estimated on the basis of this sample via maximum likelihood estimation. The likelihood function:

$$\mathbf{P}(\text{neighborhood data} \mid \text{parameters}) = \prod_{i=1}^{k} \mathbf{P}(\sigma_i \mid \boldsymbol{\omega})$$

Inference for Mallows (complete rankings)

Rankings $\boldsymbol{\sigma} = \{\sigma_1, \dots, \sigma_k\}$ observed locally

$$\mathbf{P}(\boldsymbol{\sigma} \mid \theta, \pi) = \prod_{i=1}^{k} \mathbf{P}(\sigma_i \mid \theta, \pi)$$

$$= \prod_{i=1}^{k} \frac{\exp(-\theta\ \mathrm{T}(\sigma_i, \pi))}{\phi(\theta)}$$

$$= \frac{\exp\left(-\theta\left(\mathrm{T}(\sigma_1, \pi) + \cdots + \mathrm{T}(\sigma_k, \pi)\right)\right)}{\phi^k(\theta)}$$

$$= \frac{\exp\left(-\theta\ \sum_{i=1}^{k} \mathrm{T}(\sigma_i, \pi)\right)}{\left(\prod_{j=1}^{n} \frac{1-\exp(-j\theta)}{1-\exp(-\theta)}\right)^k}$$

ML →

$$\hat{\pi} = \underset{\pi \in \Omega}{\operatorname{argmin}} \sum_{i=1}^{k} \mathrm{T}(\sigma_i, \pi)$$

↓

monotone in θ

$$\frac{1}{k} \sum_{i=1}^{k} \mathrm{T}(\sigma_i, \hat{\pi})$$

$$= \frac{n \exp(-\theta)}{1-\exp(-\theta)} - \sum_{j=1}^{n} \frac{j \exp(-j\theta)}{1-\exp(-j\theta)}$$

Probability of Incomplete Rankings

Given a probability $\mathbf{P}(\cdot)$ on S_n, the probability of an incomplete ranking σ is given by the probability of its linear extensions:

$$\mathbf{P}\left(\sigma\right) = \mathbf{P}\left(E(\sigma)\right) = \sum_{\pi \in E(\sigma)} P(\pi)$$

linear extensions

Probability of Incomplete Rankings

A	B	C	D	0.14
A	B	D	C	0.00
A	C	B	D	0.08
A	C	D	B	0.00
A	D	B	C	0.10
A	D	C	B	0.00
B	A	C	D	0.00
B	A	D	C	0.05
B	C	A	D	0.00
B	C	D	A	0.00
B	D	A	C	0.15
B	D	C	A	0.00
C	A	B	D	0.00
C	A	D	B	0.03
C	B	A	D	0.00
C	B	D	A	0.16
C	D	A	B	0.00
C	D	B	A	0.00
D	A	B	C	0.00
D	A	C	B	0.02
D	B	A	C	0.00
D	B	C	A	0.17
D	C	A	B	0.00
D	C	B	A	0.09

$$\mathbf{P}\,(\mathrm{A} \succ \mathrm{C}) =$$

Probability of Incomplete Rankings

A	B	C	D	0.14
A	B	D	C	0.00
A	C	B	D	0.08
A	C	D	B	0.00
A	D	B	C	0.10
A	D	C	B	0.00
B	A	C	D	0.00
B	A	D	C	0.05
B	C	A	D	0.00
B	C	D	A	0.00
B	D	A	C	0.15
B	D	C	A	0.00
C	A	B	D	0.00
C	A	D	B	0.03
C	B	A	D	0.00
C	B	D	A	0.16
C	D	A	B	0.00
C	D	B	A	0.00
D	A	B	C	0.00
D	A	C	B	0.02
D	B	A	C	0.00
D	B	C	A	0.17
D	C	A	B	0.00
D	C	B	A	0.09

$$\mathbf{P}\,(\mathrm{A} \succ \mathrm{C}) = 0.54$$

Inference for Mallows (incomplete rankings)

The corresponding likelihood:

$$\mathbf{P}(\boldsymbol{\sigma} \mid \theta, \pi) = \prod_{i=1}^{k} \mathbf{P}(E(\sigma_i) \mid \theta, \pi)$$

$$= \prod_{i=1}^{k} \sum_{\gamma \in E(\sigma_i)} \mathbf{P}(\gamma \mid \theta, \pi)$$

$$= \frac{\prod_{i=1}^{k} \sum_{\gamma \in E(\sigma_i)} \exp\big(-\theta \, \mathrm{T}(\gamma, \pi)\big)}{\left(\prod_{j=1}^{n} \frac{1 - \exp(-j\theta)}{1 - \exp(-\theta)}\right)^{k}}$$

Exact MLE $\left(\hat{\pi}, \hat{\theta}\right) = \operatorname*{argmax}_{\pi, \theta} \mathbf{P}(\boldsymbol{\sigma} \mid \theta, \pi)$ becomes infeasible when n is large. Approximation is needed.

Inference for Mallows (incompl

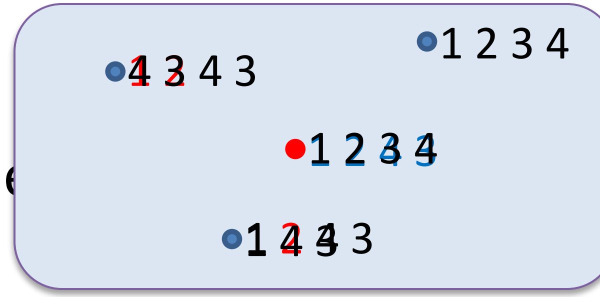

Approximation via a variant of EM, viewing the non-observed labels as hidden variables.

Key idea: replacing the E-step of EM algorithm with a maximization step (widely used in learning HMM, K-means clustering, etc.)

1. Start with an initial center ranking (via *generalized Borda count*)
2. Replace an incomplete observation with its most probable extension (*first M-step*, can be done efficiently)
3. Obtain MLE as in the complete ranking case (*second M-step*)
4. Replace the initial center ranking with current estimation
5. Repeat until convergence

Inference for Plackett-Luce

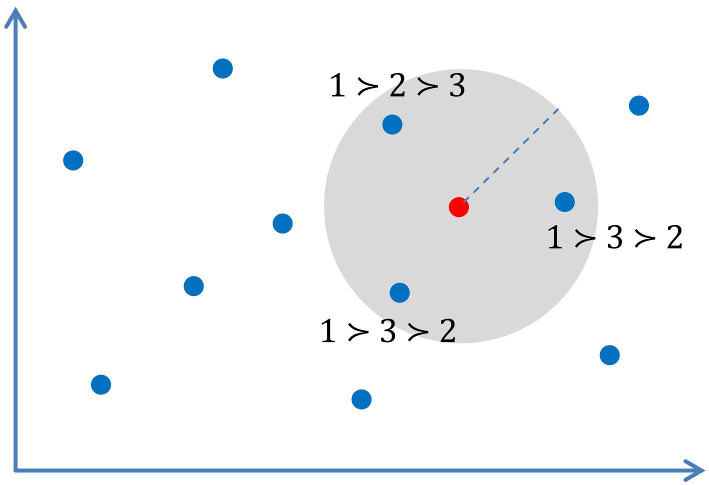

The probability to observe the rankings $\boldsymbol{\sigma} = \{\sigma_1, \ldots, \sigma_k\}$ in the neighborhood:

$$\mathbf{P}(\boldsymbol{\sigma} \mid \boldsymbol{v}) = \prod_{i=1}^{k} \prod_{j=1}^{n_i} \frac{v_{\sigma_i^{-1}(j)}}{v_{\sigma_i^{-1}(1)} + \cdots + v_{\sigma_i^{-1}(n_i)}}$$

Corresponding MLE can be efficiently done through, e.g., MM (minorization and maximization) algorithm.

Sushi Data Set

Rankings of 10 types of sushi by 5000 customers.
Each customer is characterized by 11 features.

Collected by Kamishima et al. Preprocessed by Grbovic.

Experimental Results

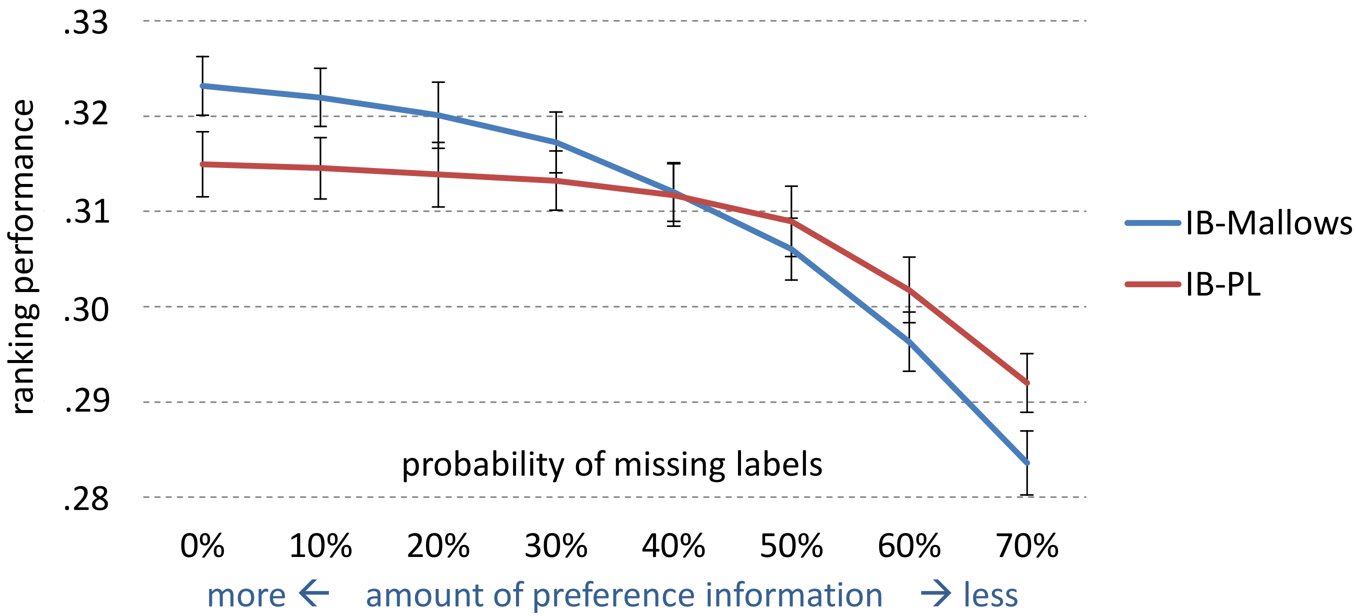

Main observation

– Mallows vs PL model: the former is better for complete rankings and the latter is better for incomplete ones.

Experimental Results

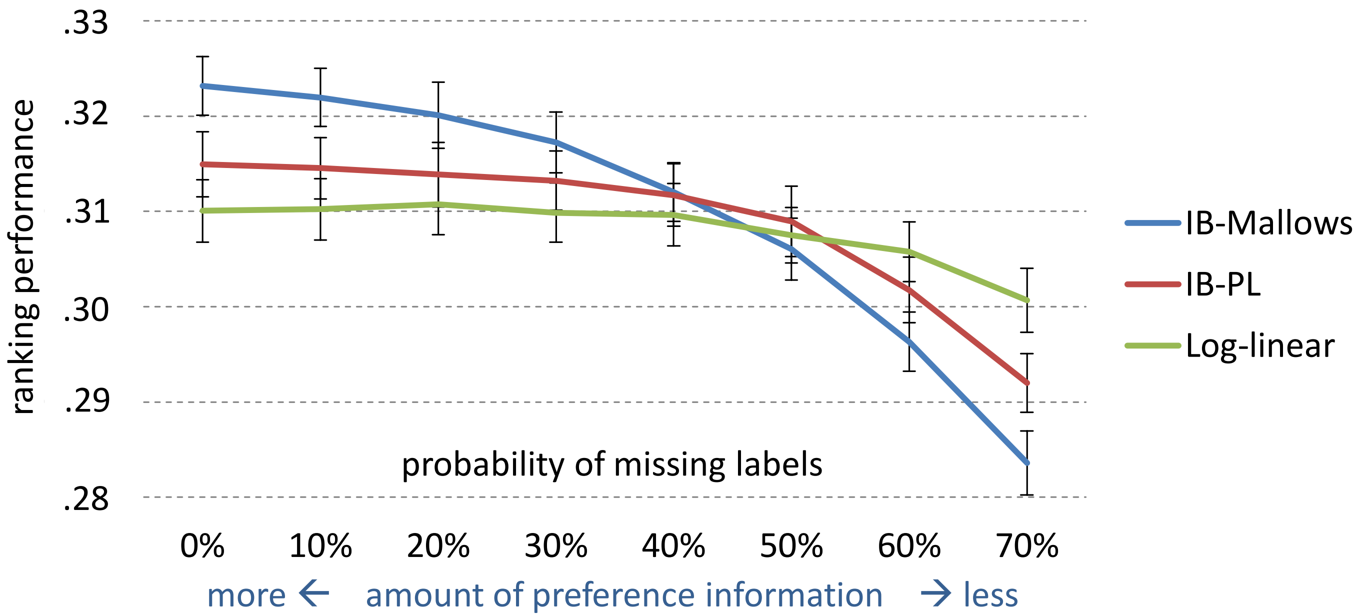

Main observation

- Mallows vs PL model: the former is better for complete rankings and the latter is better for incomplete ones.
- Instance-based methods are more *flexible* and have *higher variance* and *lower bias* compared to the log-linear approach.

Label Ranking with Probabilistic Models

statistical ranking models	**machine learning techniques**
Mallows model	instance-based learning
Plackett-Luce model	generalized linear model

Results from: Cheng et al., ICML 10

A Generalized Linear Model based on PL

Recall the PL model:

$$\mathbf{P}(\pi \mid \boldsymbol{v}) = \prod_{i=1}^{n} \frac{v_{\sigma^{-1}(i)}}{v_{\sigma^{-1}(i)} + v_{\sigma^{-1}(i+1)} + \cdots + v_{\sigma^{-1}(n)}}$$

We model the parameter v_i as a linear function of the features describing the instance:

$$v_i = \exp\left(\sum_{j=1}^{d} \alpha_j^{(i)} \cdot x_j\right), 1 \leq i \leq n, 1 \leq j \leq d$$

Maximum Likelihood Estimation

Given training data $D = \left\{\left(\boldsymbol{x}^{(k)}, \sigma^{(k)}\right)\right\}_{k=1}^{m}$ with $\boldsymbol{x}^{(k)} = \left(x_1^{(k)}, \ldots, x_d^{(k)}\right)$, the log-likelihood function is

$$\mathbf{P}(D \mid \boldsymbol{\alpha}) = \sum_{k=1}^{m} \sum_{i=1}^{n_k} \left[\log v\left(\tilde{\sigma}^{(k)}(i), k\right) - \log \sum_{j=i}^{n_k} v\left(\tilde{\sigma}^{(k)}(j), k\right) \right]$$

where $\tilde{\sigma}(i) = \sigma^{-1}(i)$ is the index of the label ranked at position i, n_k is the number of labels in the ranking $\sigma^{(k)}$, and

$$v(i, k) = \exp\left(\sum_{j=1}^{d} \alpha_j^{(i)} \cdot x_j^{(k)} \right).$$

It is convex!

Experimental Results

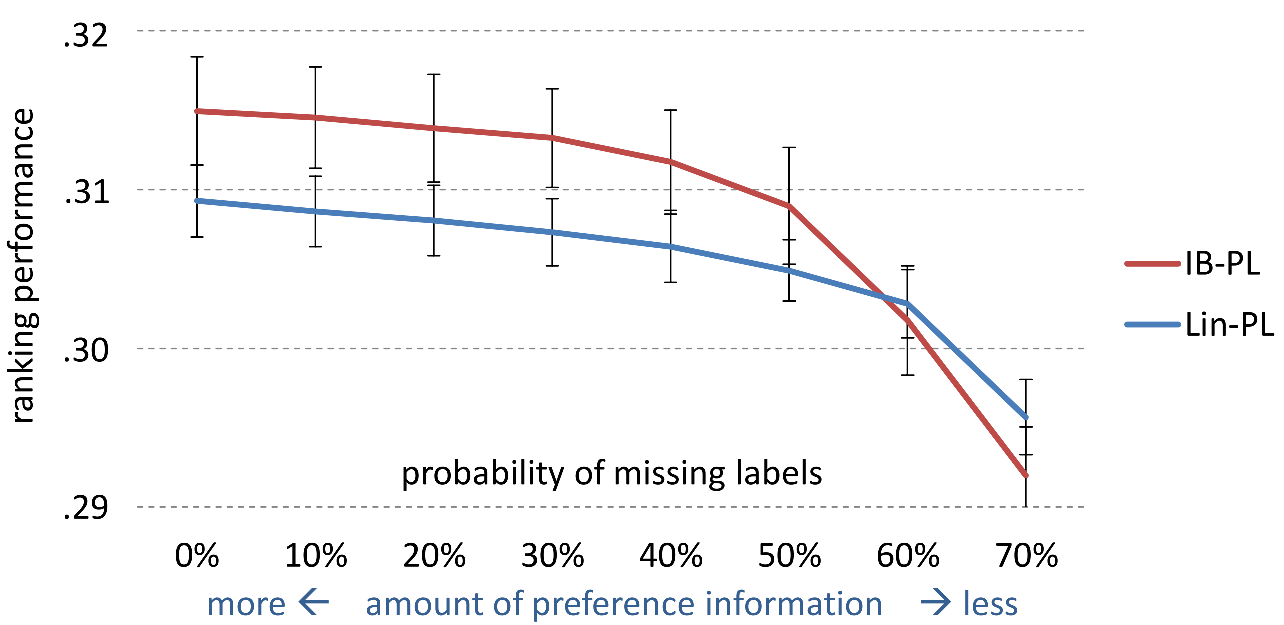

Main observation

Instance-based methods are more *flexible*; generalized linear models are more *robust*.

Agenda

1. Introduction to Preference Learning
2. Label Ranking
3. **Extensions and Applications**
4. Conclusions

Learning with Reject Option

To train a learner that is able to say "I don't know".

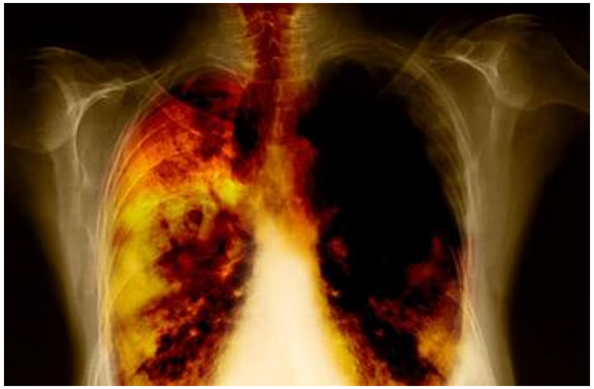

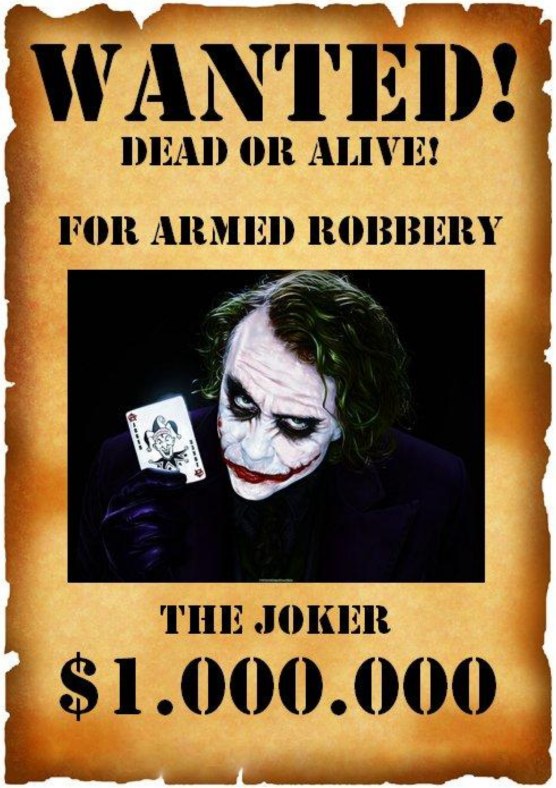

Label Ranking with Rejects

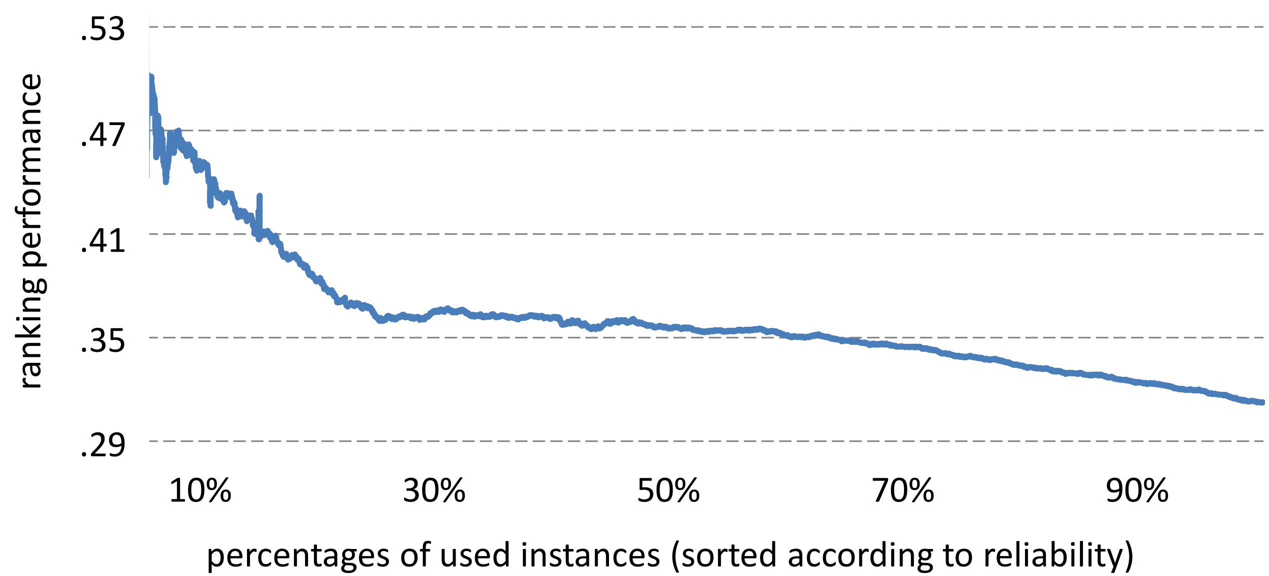

The above accuracy-rejection curve confirms the outputs of the probabilistic models can be used as a *reliability* measure.

From Total to Partial Order Relations

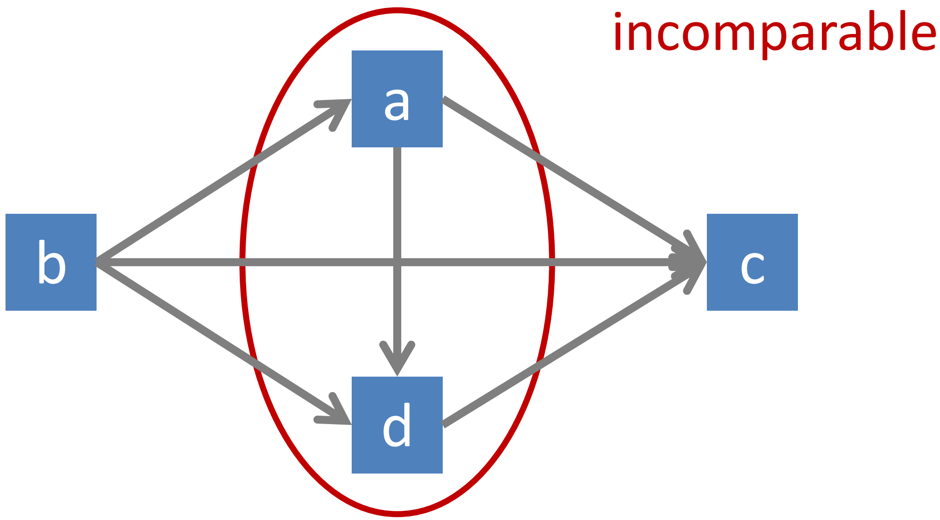

Partial abstention:
The target is a total order, and a predicted partial order expresses incomplete knowledge about the target .

Partial Orders from Pairwise Comparisons

only rely on most confident comparisons → thresholding the relation

	a	b	c	d
a		0.3	0.8	0.4
b	0.7		0.9	0.7
c	0.2	0.1		0.7
d	0.6	0.3	0.3	

$\mathrm{P}(\mathrm{a}, \mathrm{d}) = \mathbf{P}(\mathrm{a} \succ \mathrm{d})$

thresholding at 0.5 →

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

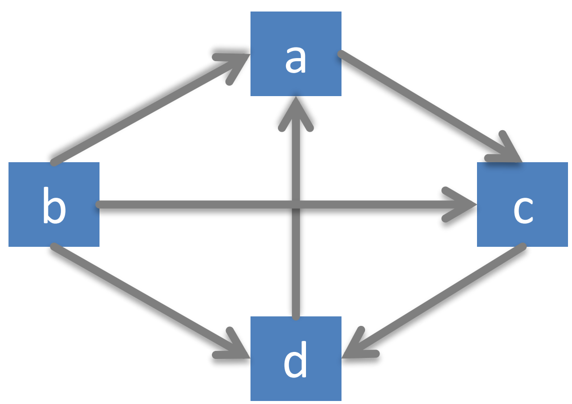

Inconsistent!

Partial Orders from Pairwise Comparisons

only rely on most confident comparisons → thresholding the relation

	a	b	c	d
a		0.3	0.8	0.4
b	0.7		0.9	0.7
c	0.2	0.1		0.7
d	0.6	0.3	0.3	

thresholding at 1 →

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

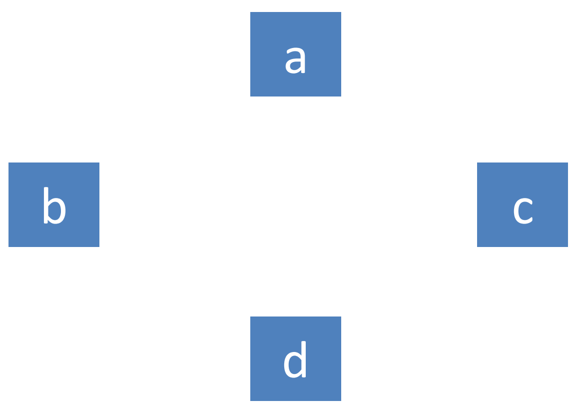

complete abstention

Partial Orders from Pairwise Comparisons

only rely on most confident comparisons → thresholding the relation

	a	b	c	d
a		0.3	0.8	0.4
b	0.7		0.9	0.7
c	0.2	0.1		0.7
d	0.6	0.3	0.3	

thresholding at 0.6 →

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

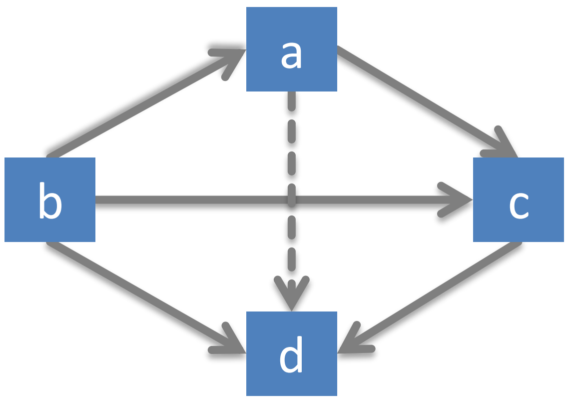

Consistent, but not a partial order!

Our Ideas & Results

Can we restrict $\mathrm{P}(\cdot,\cdot)$ to exclude the possibility of cycles and violations of transitivity from the very beginning?

- We make use of label ranking methods that produce probability distributions $\mathbf{P}$ over the ranking space Ω.

- We show that thresholding pairwise preferences induced by certain distributions yields partial order relations.

Results from: Cheng et al., NIPS 2012.

Theoretical Results

Let the preference relation P be given by a probability distribution $\mathbf{P}$ on Ω, that is $\mathrm{P}(y_i, y_j) = \mathbf{P}(y_i \succ y_j) = \sum_{\sigma \in \mathrm{E}(y_i, y_j)} \mathbf{P}(\sigma)$.

Theorem Let $\mathbf{P}$ be

(1) the Plackett-Luce model or

(2) the Mallows model with a distance Δ having the transposition property.

Moreover, let Q be the thresholded relation

$\mathrm{Q}(y_i, y_j) = 1$ if $\mathrm{P}(y_i, y_j) > q$ and

$\mathrm{Q}(y_i, y_j) = 0$ otherwise.

Then Q defines a proper partial order relation for all $q \in [1/2, 1)$.

Theoretical Results

Theorem Let $\mathcal{R}_{\mathrm{M}}$ denote the set of different partial orders (up to isomorphism) that can be represented as a thresholded relation Q, where P is derived according to the Mallows model with Kendal distance. For any given threshold $q \in [1/2, 1)$, the cardinality of this set $|\mathcal{R}_{\mathrm{M}}| = n$.

Theorem Let $\mathcal{R}_{\mathrm{PL}}$ denote the set of different partial orders (up to isomorphism) that can be represented as a thresholded relation Q, where P is derived according to the Plackett-Luce model. For any given threshold $q \in [1/2, 1)$, the cardinality of this set is given by the n^{th} Catalan number:

$$|\mathcal{R}_{\mathrm{PL}}| = \frac{1}{n+1}\binom{2n}{n}$$

Multi-Label Classification

– An instance can belong to multiple classes.
– Complex structured information may exist (e.g., label dependency).

X1	X2	X3	Y1	Y2	Y3	Y4
0.34	0	10	1	0	1	0

ranking with ties

It can be solved by (1) **label ranking** and then (2) **grouping**, e.g., [Fürnkranz et al., ML 08]

Reinforcement Learning

- Learning to behave optimally in **uncertain dynamic environments**.
- A **policy** maps **states** to **actions**.
- **Feedback** is often of a **qualitative** nature!

[Cheng et al., ECMLPKDD 11]

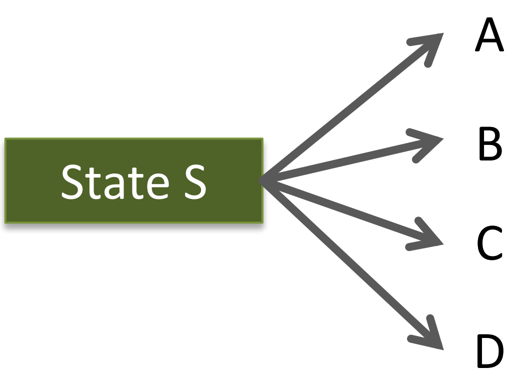

$$S \mapsto C \succ A \succ B \succ D$$

LABEL RANKING

[Fürnkranz et al., ML 11]

Reinforcement Learning with Qualitative Feedback (DFG).

Summary

- Preference learning is
  - **methodologically** interesting,
  - **theoretically** challenging,
  - and **practically** useful, with many potential **applications**;
  - more **general** than could be shown in this talk („preferences“ in the broad sense, standard ML problems as special cases, ...); in fact, a flexible machine learning framework for learning from **weak supervision**;
  - **interdisciplinary** (connections to operations research, decision sciences, economics, social choice, recommender systems, information retrieval, ...).
- We discuss **label ranking**, which, albeit being a specific type of preference learning problem, shares commonalities with other problems in this field.
- Label ranking with probabilistic models; predicting partial orders via thresholding; applications …

Some Related Publications

W. Cheng, E. Hüllermeier, W. Waegeman, V. Welker. **Label ranking with partial abstention based on thresholded probabilistic models.** NIPS 2012. Lake Tahoe, USA. December 2012.

J. Fürnkranz, E. Hüllermeier, W. Cheng, S.-H. Park. **Preference-based reinforcement learning: a formal framework and a policy iteration algorithm.** Machine Learning 89.

W. Cheng, J. Fürnkranz, E. Hüllermeier, S.-H. Park. **Preference-based policy iteration: leveraging preference learning for reinforcement learning.** ECMLPKDD 2011. Athens, Greece. September 2011.

W. Cheng, M. Rademaker, B. De Baets, E. Hüllermeier. **Predicting partial orders: ranking with abstention.** ECMLPKDD 2010. Barcelona, Spain. September 2010.

W. Cheng, K. Dembczyński, E. Hüllermeier. **Label ranking methods based on the Plackett-Luce model.** ICML 2010. Haifa, Israel. June 2010.

W. Cheng, J. Hühn, E. Hüllermeier. **Decision tree and instance-based learning for label ranking.** ICML 2009. Montreal, Canada. June 2009.

E. Hüllermeier, J. Fürnkranz, W. Cheng, K. Brinker. **Label ranking by learning pairwise preferences.** Artificial Intelligence 172.

More related publications can be found at my Google Scholar page.