Preference Learning: A Tutorial Introduction

Johannes Fürnkranz
Knowledge Engineering
Dept. of Computer Science
Technical University Darmstadt, Germany

Eyke Hüllermeier
Computational Intelligence Group
Dept. of Mathematics and Computer Science
Marburg University, Germany

ECAI 2012, Montpellier, France, Aug 2012

Preferences are Ubiquitous

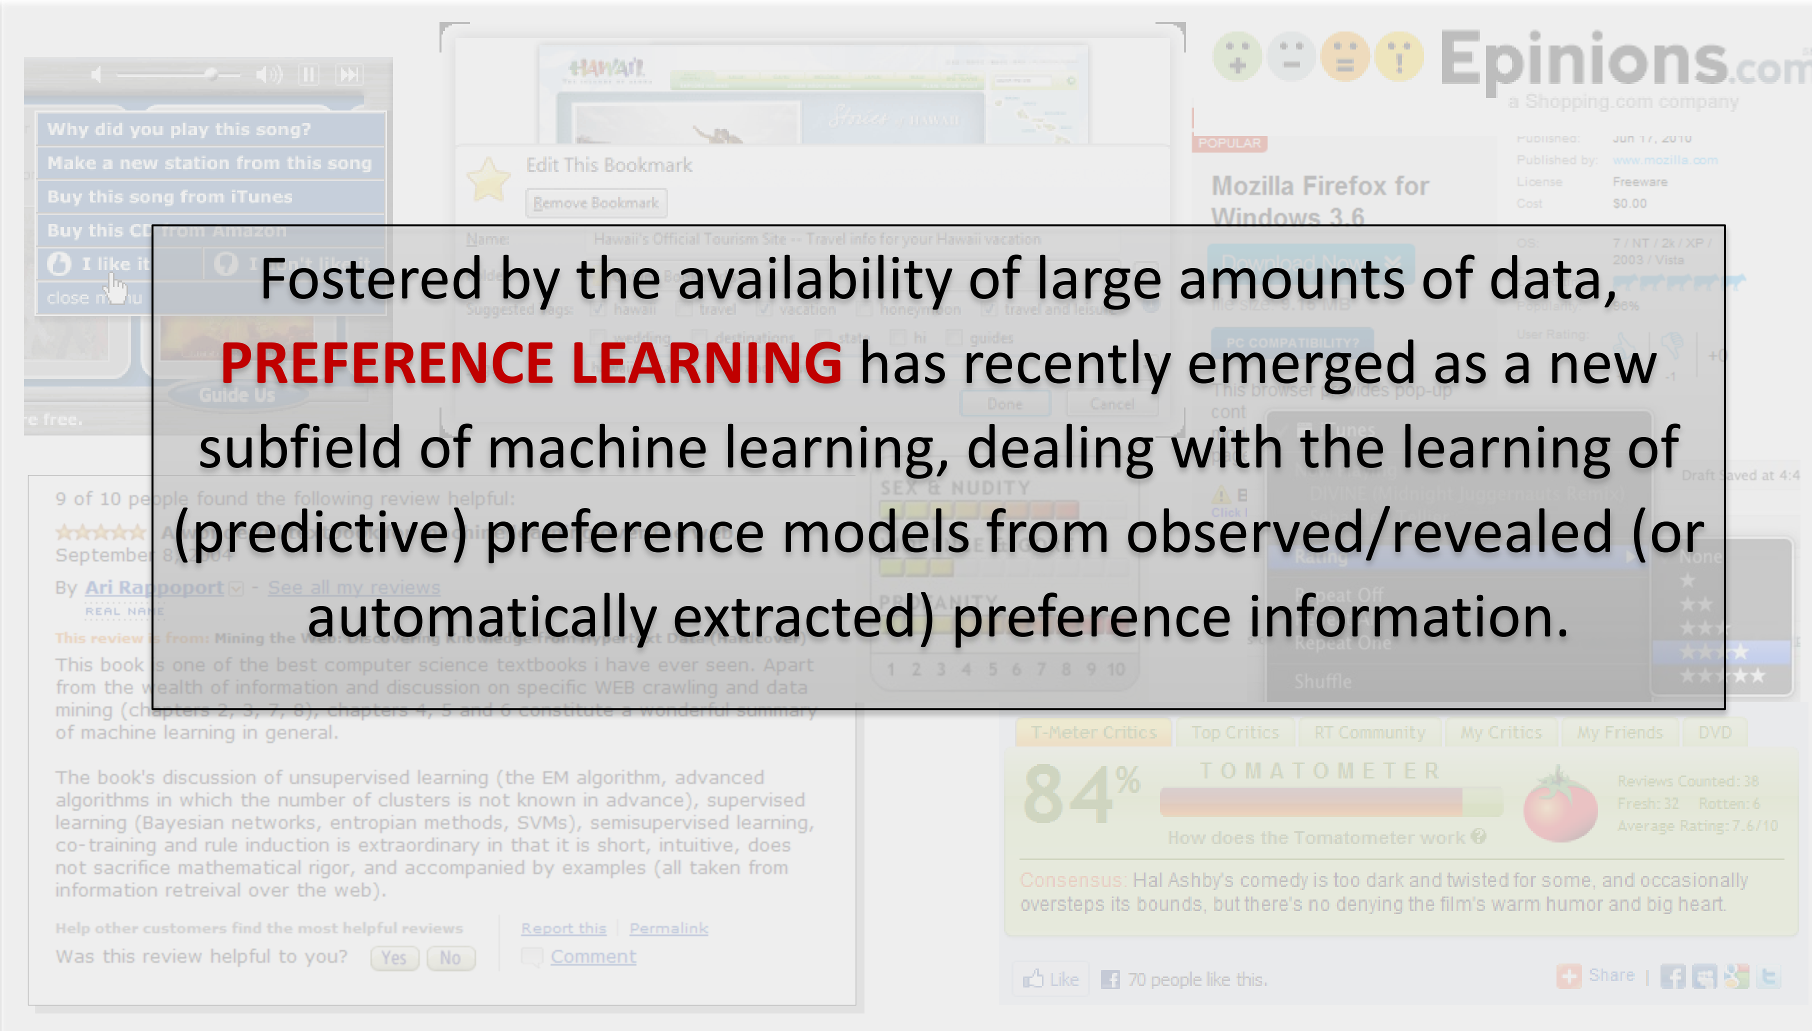

Fostered by the availability of large amounts of data, **PREFERENCE LEARNING** has recently emerged as a new subfield of machine learning, dealing with the learning of (predictive) preference models from observed/revealed (or automatically extracted) preference information.

What is Preference Learning?

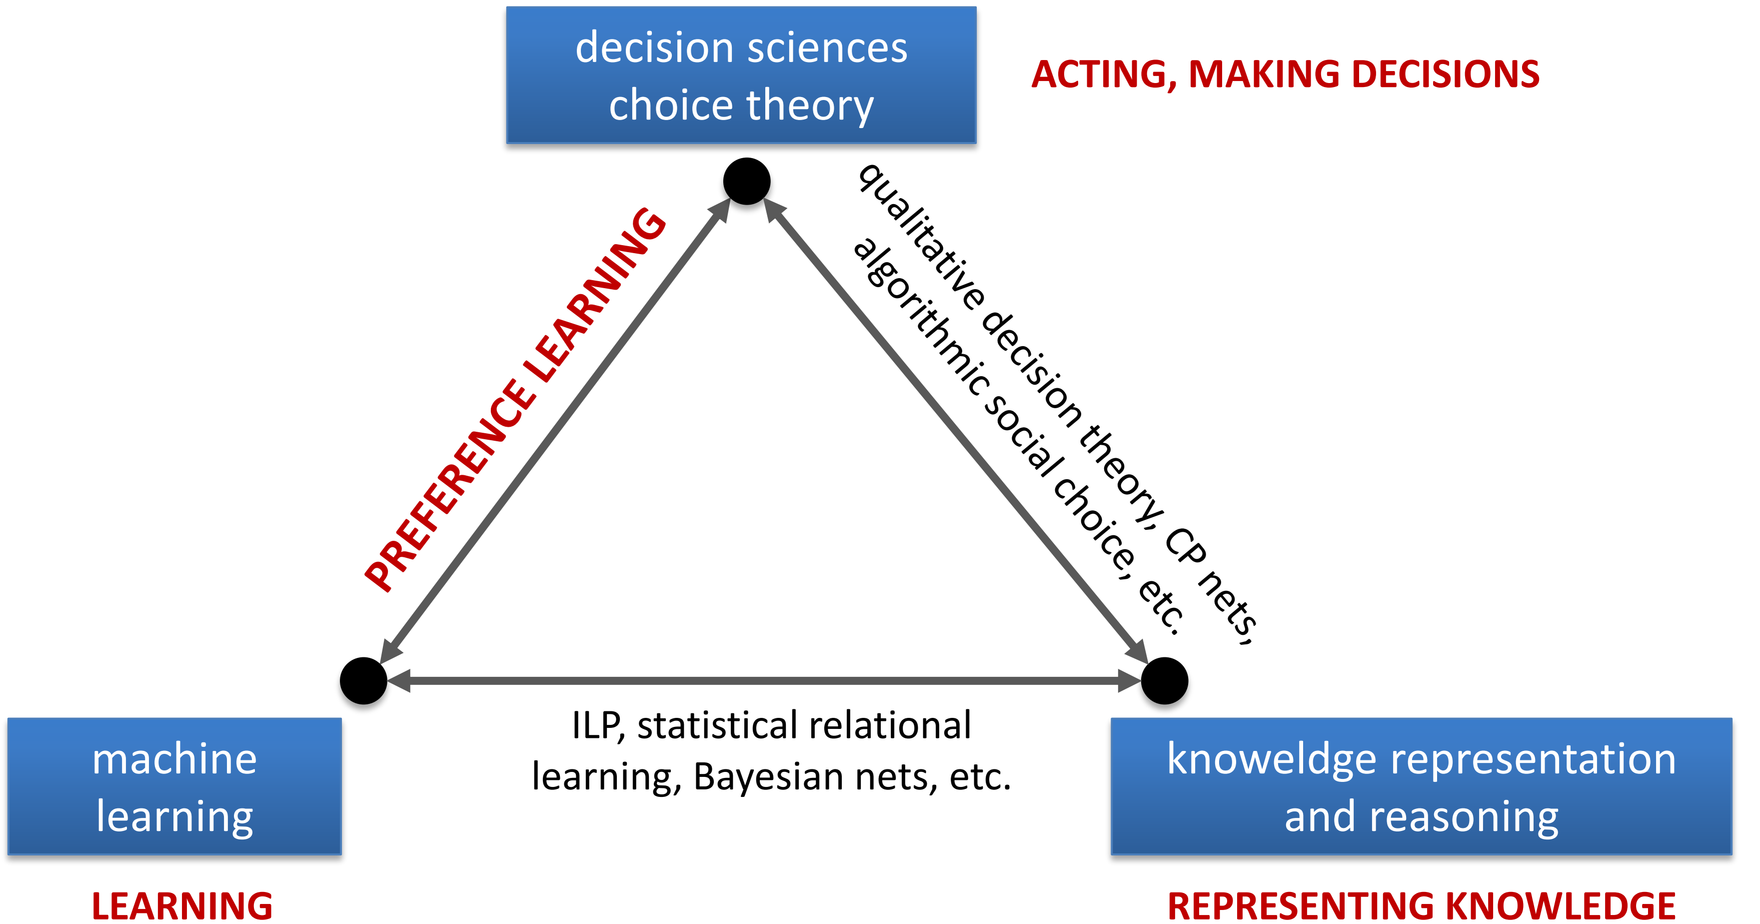

Preferences in AI

> “**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 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.**” [Brafman & Domshlak, 2009]

Preference learning: **From learning „the correct“ to learning „the preferred“**
(more flexible handling of training information and predictions)

Preferences in AI

User preferences play a key role in various fields of application:

- recommender systems,
- adaptive user interfaces,
- adaptive retrieval systems,
- autonomous agents (electronic commerce),
- games, …

Preferences in **AI research**:

- **preference representation** (CP nets, GAU networks, logical representations, fuzzy constraints, …)
- **reasoning** with preferences (decision theory, constraint satisfaction, non-monotonic reasoning, …)
- **preference acquisition** (preference elicitation, **preference learning**, ...)

Preference Learning vs. Preference Elicitation

- typically no user interaction
- holistic judgements
- fixed preferences but noisy data
- regularized models
- weak model assumptions, flexible (instead of axiomatically justified) model classes
- diverse types of training information
- computational aspects: massive data, scalable methods
- focus on predictive accuracy (expected loss)

MACHINE LEARNING	← **Preference Learning**	**Preference Elicitation** →	PREFERENCE MODELING and DECISION ANALYSIS
computer science artificial intelligence			operations research social sciences (voting and choice theory) economics and decision theory

Workshops and Related Events

- NIPS–01: New Methods for Preference Elicitation
- NIPS–02: Beyond Classification and Regression: Learning Rankings, Preferences, Equality Predicates, and Other Structures
- KI–03: Preference Learning: Models, Methods, Applications
- NIPS–04: Learning With Structured Outputs
- NIPS–05: Workshop on Learning to Rank
- IJCAI–05: Advances in Preference Handling
- SIGIR 07–10: Workshop on Learning to Rank for Information Retrieval
- **ECML/PDKK 08–10: Workshop on Preference Learning**
- NIPS–09: Workshop on Advances in Ranking
- American Institute of Mathematics Workshop in Summer 2010: The Mathematics of Ranking
- NIPS-11: Workshop on Choice Models and Preference Learning
- EURO-12: Special Track on Preference Learning

AGENDA

1. **Preference Learning Tasks**
2. Performance Assessment and Loss Functions
3. Preference Learning Techniques
4. Complexity of Preference Learning
5. Conclusions

Preferences Learning Settings

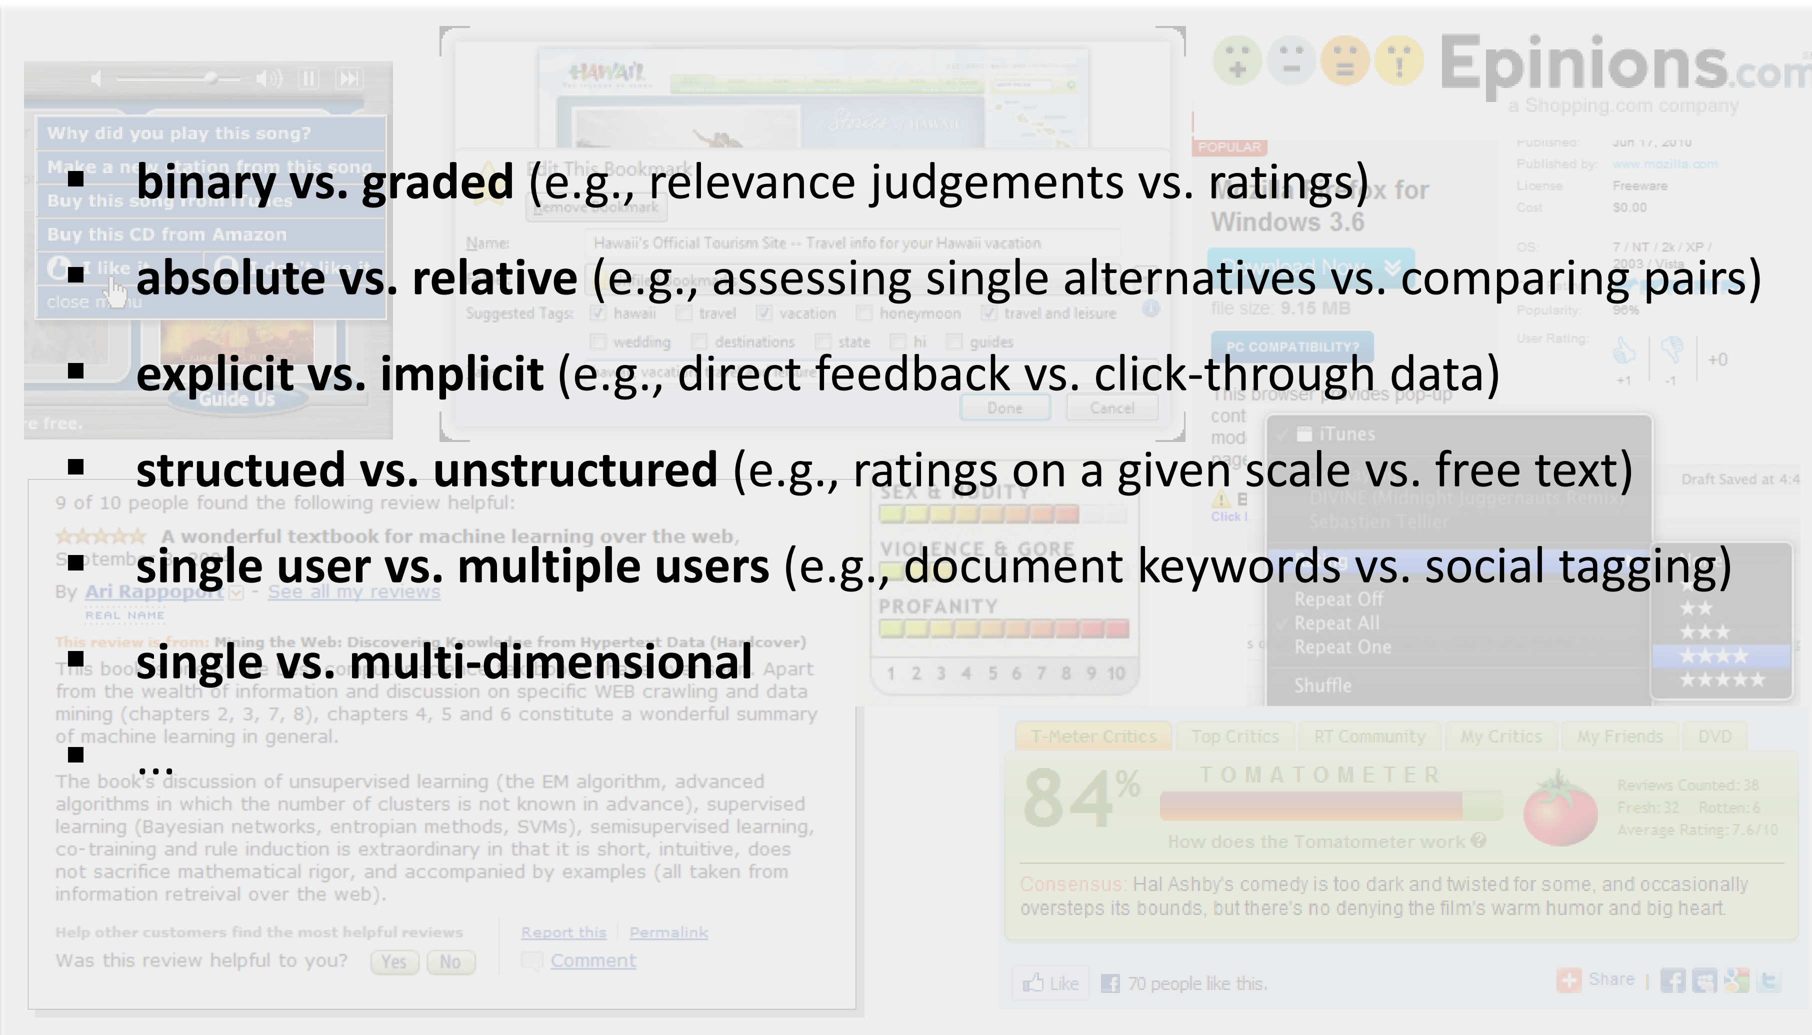

- **binary vs. graded** (e.g., relevance judgements vs. ratings)
- **absolute vs. relative** (e.g., assessing single alternatives vs. comparing pairs)
- **explicit vs. implicit** (e.g., direct feedback vs. click-through data)
- **structued vs. unstructured** (e.g., ratings on a given scale vs. free text)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

Preference Learning

Preference learning problems can be distinguished along several **problem dimensions**, including

- **representation of preferences, type of preference model**:
  - utility function (ordinal, numeric),
  - preference relation (partial order, ranking, …),
  - logical representation, …
- **description of individuals/users and alternatives/items**:
  - identifier, feature vector, structured object, …
- **type of training input**:
  - direct or indirect feedback,
  - complete or incomplete relations,
  - utilities, …
- …

Preference Learning

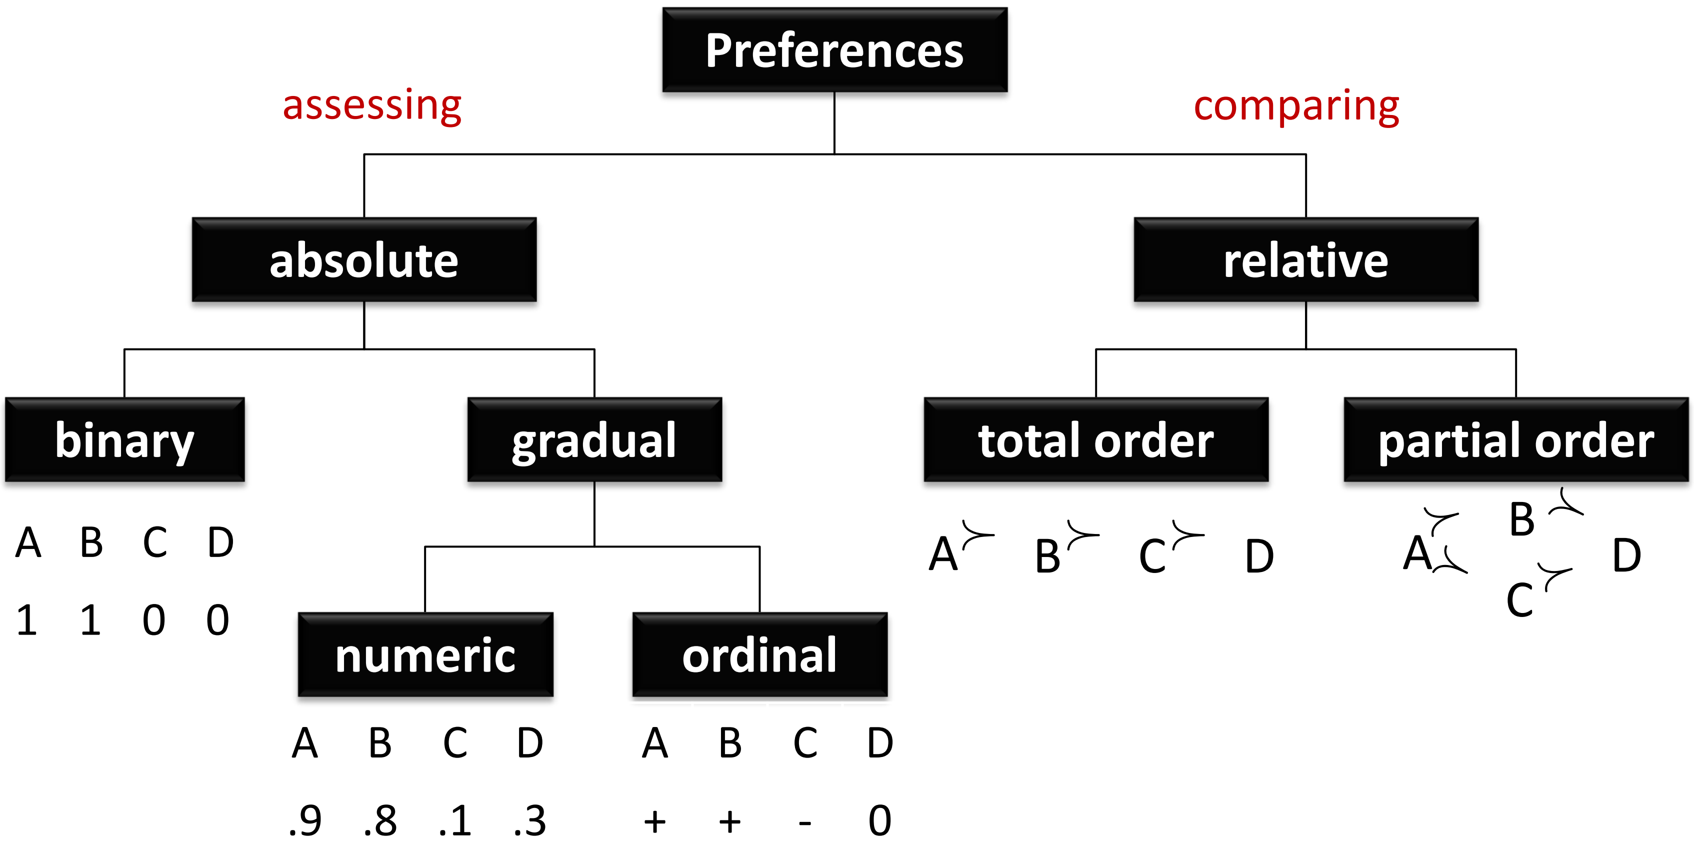

→ (ordinal) regression

→ classification/ranking

Structure of this Overview

(1) Preference learning as an extension of **conventional supervised learning**: Learn a mapping

$$\{\text{preference contexts}\} \rightarrow \{\text{preference models}\}$$

e.g., people, queries, etc.

e.g., rankings, partial orders, CP-nets, etc.

(→ connection to structured/complex output prediction)

(2) Other settings (object ranking, instance ranking, CF, …)

Structure of this Overview

(1) Preference learning as an extension of **conventional supervised learning**: Learn a mapping

$$\{\text{preference contexts}\} \rightarrow \{\text{preference models}\}$$

e.g., people, queries, etc.

e.g., rankings, partial orders, CP-nets, etc.

(→ connection to structured/complex output prediction)

The output space consists of preference models over a fixed set of alternatives (classes, labels, …) represented in terms of an identifier
→ *extensions of multi-class classification*

Multilabel Classification [Tsoumakas & Katakis 2007]

Training

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

Binary preferences on a fixed set of items: liked or disliked

Prediction

0.92	1	81	382	**0**	**1**	**0**	**1**

LOSS

Ground truth

0.92	1	81	382	**1**	**1**	**0**	**1**

Multilabel Ranking

Training

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

Binary preferences on a fixed set of items: liked or disliked

Prediction

$B \succ D \succ C \succ A$

0.92	**1**	**81**	**382**	**4**	**1**	**3**	**2**

A ranking of all items

LOSS

Ground truth

0.92	**1**	**81**	**382**	**1**	**1**	**0**	**1**

Graded Multilabel Classification [Cheng et al. 2010]

Training

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

Ordinal preferences on a fixed set of items: liked, disliked, or something in-between

Prediction

0.92	**1**	**81**	**382**	**--**	**+**	**0**	**++**

LOSS

Ground truth

0.92	**1**	**81**	**382**	**0**	**++**	**--**	**+**

Graded Multilabel Ranking

Training

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

Ordinal preferences on a fixed set of items: liked, disliked, or something in-between

Prediction

$B \succ D \succ C \succ A$

0.92	**1**	**81**	**382**	**4**	**1**	**3**	**2**

A ranking of all items

LOSS

Ground truth

0.92	**1**	**81**	**382**	**0**	**++**	**--**	**+**

Label Ranking [Hüllermeier et al. 2008]

Training

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

Instances are associated with pairwise preferences between labels.

Prediction

$B \succ D \succ C \succ A$

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**

Calibrated Label Ranking [Fürnkranz et al. 2008]

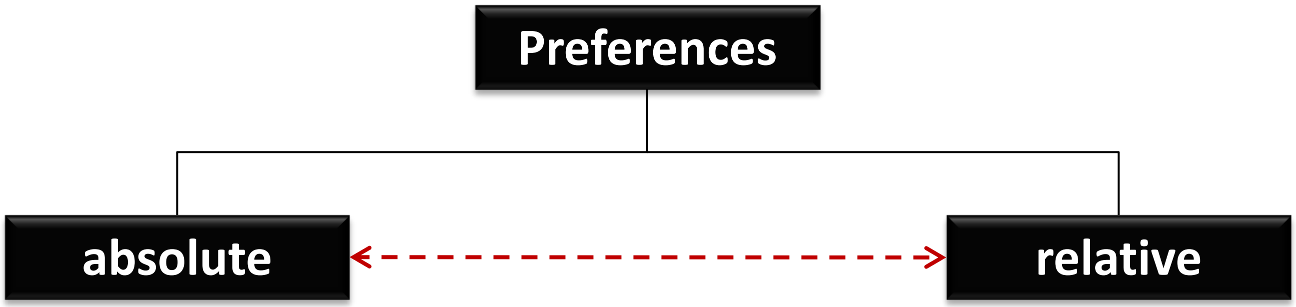

Combining absolute and relative evaluation:

$$\underbrace{a \succ b \succ c}_{\substack{\text{relevant}\\ \text{positive}\\ \text{liked}}} \;\big|\; \underbrace{d \succ e \succ f \succ g}_{\substack{\text{irrelevant}\\ \text{negative}\\ \text{disliked}}}$$

Classes of Methods to Tackle these Problems

Reduction to binary classification	Ranking by pairwise comparison [Hüllermeier et al. 08]]	Learning pairwise preferences
	Constraint classification [Har-Peled et al. 03]	Learning utility functions
Boosting	Log-linear models for label ranking [Dekel et al. 04]	
Structured output prediction, margin maximization	Structured output prediction [Vembu et al. 09] Local prediction (lazy learning) [Brinker & EH , Cheng et al. 09]	Structured prediction
Statistical inference	Statistical models for label ranking [Cheng et al. 09, Cheng et al. 10]	

Structure of this Overview

(1) Preference learning as an extension of **conventional supervised learning**: Learn a mapping

$$\{\text{preference contexts}\} \rightarrow \{\text{preference models}\}$$

e.g., people, queries, etc.

e.g., rankings, partial orders, CP-nets, etc.

(→ connection to structured/complex output prediction)

(2) Other settings:

object ranking, instance ranking,
collaborative filtering, dyadic prediction

Object Ranking [Cohen et al. 99]

Training

$$(0.74, 1, 25, 165) \succ (0.45, 0, 35, 155)$$
$$(0.47, 1, 46, 183) \succ (0.57, 1, 61, 177)$$
$$(0.25, 0, 26, 199) \succ (0.73, 0, 46, 185)$$

 $\succ$ 

Pairwise preferences between objects (instances)

Prediction (ranking a new set of objects)

$$\mathcal{Q} = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \boldsymbol{x}_3, \boldsymbol{x}_4, \boldsymbol{x}_5, \boldsymbol{x}_6, \boldsymbol{x}_7, \boldsymbol{x}_8, \boldsymbol{x}_9, \boldsymbol{x}_{10}, \boldsymbol{x}_{11}, \boldsymbol{x}_{12}, \boldsymbol{x}_{13}\}$$

$$\boldsymbol{x}_{10} \succ \boldsymbol{x}_4 \succ \boldsymbol{x}_7 \succ \boldsymbol{x}_1 \succ \boldsymbol{x}_{11} \succ \boldsymbol{x}_2 \succ \boldsymbol{x}_8 \succ \boldsymbol{x}_{13} \succ \boldsymbol{x}_9 \succ \boldsymbol{x}_3 \succ \boldsymbol{x}_{12} \succ \boldsymbol{x}_5 \succ \boldsymbol{x}_6$$

Object Ranking [Cohen et al. 99]

prediction — TOTAL ORDER

1	x_{11}
2	x_7
3	x_4
4	x_2
5	x_{10}
6	x_1
7	x_8
8	x_{12}
9	x_9
10	x_6
11	x_3
12	x_5

ground truth — TOTAL ORDER

1	x_7
2	x_6
3	x_3
4	x_9
5	x_1
6	x_8
7	x_2
8	x_{10}
9	x_{11}
10	x_4
11	x_5
12	x_{12}

ground truth — TOP-K RANKING

1	x_7
2	x_{10}
3	x_1
4	x_{11}
5	x_9

ground truth — RELEVANCE RATING

x_{11}	$\oplus$
x_7	$\oplus$
x_4	$\ominus$
x_2	$\ominus$
x_{10}	$\oplus$
x_1	$\ominus$
x_8	$\oplus$
x_{12}	$\ominus$
x_9	$\ominus$
x_6	$\ominus$
x_3	$\oplus$
x_5	$\ominus$

Instance Ranking [Fürnkranz et al. 2009]

Training

	X1	X2	X3	X4	class
$\boldsymbol{x}_1$	0.34	0	10	174	--
$\boldsymbol{x}_2$	1.45	0	32	277	0
$\boldsymbol{x}_3$	0.74	1	25	165	++
	…	…	…	…	…
$\boldsymbol{x}_n$	0.95	1	72	273	+

Absolute preferences on an ordinal scale.

Prediction (ranking a new set of objects)

$$\mathcal{Q} = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \boldsymbol{x}_3, \boldsymbol{x}_4, \boldsymbol{x}_5, \boldsymbol{x}_6, \boldsymbol{x}_7, \boldsymbol{x}_8, \boldsymbol{x}_9, \boldsymbol{x}_{10}, \boldsymbol{x}_{11}, \boldsymbol{x}_{12}, \boldsymbol{x}_{13}\}$$

$$\boldsymbol{x}_{10} \succ \boldsymbol{x}_4 \succ \boldsymbol{x}_7 \succ \boldsymbol{x}_1 \succ \boldsymbol{x}_{11} \succ \boldsymbol{x}_2 \succ \boldsymbol{x}_8 \succ \boldsymbol{x}_{13} \succ \boldsymbol{x}_9 \succ \boldsymbol{x}_3 \succ \boldsymbol{x}_{12} \succ \boldsymbol{x}_5 \succ \boldsymbol{x}_6$$

Ground truth (ordinal classes)

$\boldsymbol{x}_{10}$	$\boldsymbol{x}_4$	$\boldsymbol{x}_7$	$\boldsymbol{x}_1$	$\boldsymbol{x}_{11}$	$\boldsymbol{x}_2$	$\boldsymbol{x}_8$	$\boldsymbol{x}_{13}$	$\boldsymbol{x}_9$	$\boldsymbol{x}_3$	$\boldsymbol{x}_{12}$	$\boldsymbol{x}_5$	$\boldsymbol{x}_6$
+	0	++	++	--	+	0	+	--	0	0	--	--

Instance Ranking [Fürnkranz et al. 2009]

Extension of AUC maximization to the polytomous case, in which instances are rated on an ordinal scale such as { **bad, medium, good**}

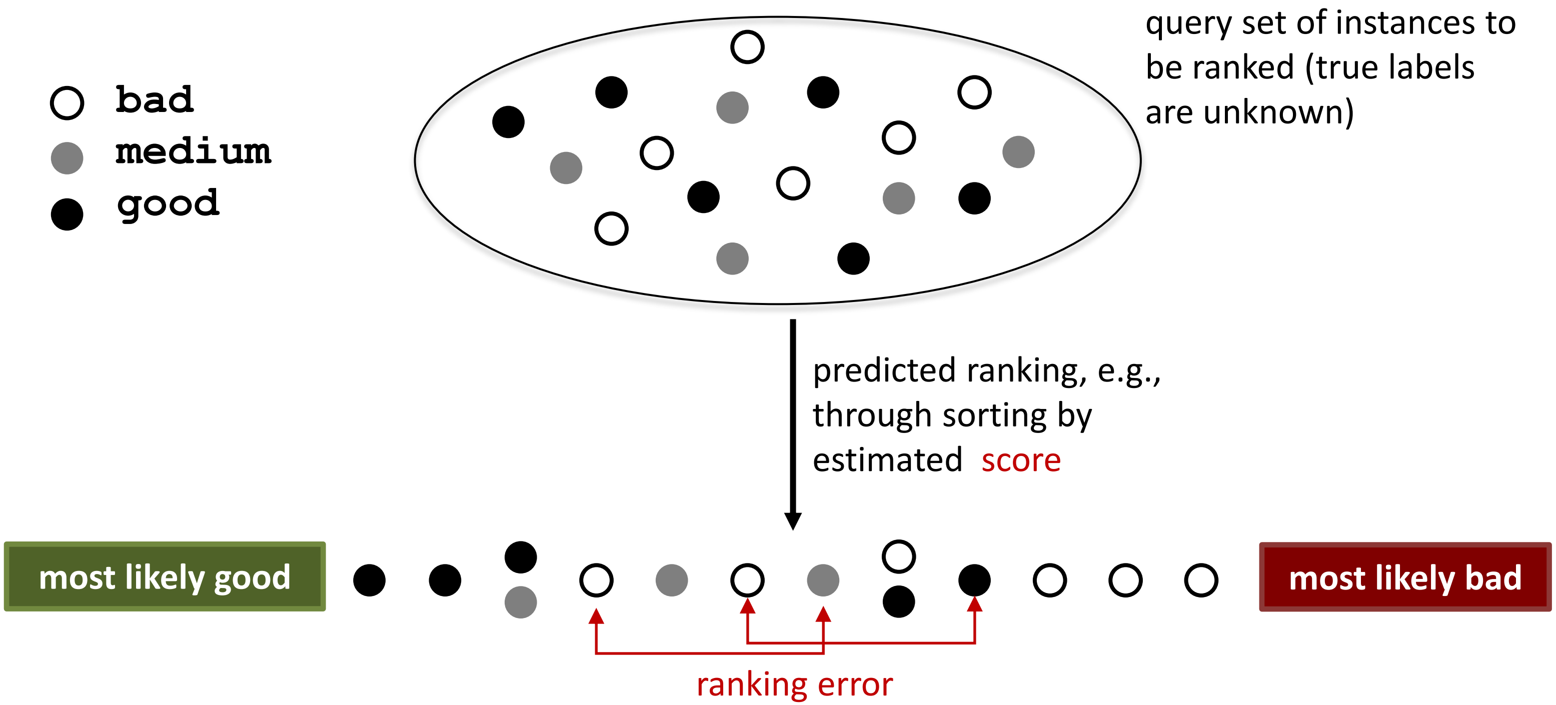

Collaborative Filtering [Goldberg et al. 1992]

PRODUCTS

USERS

	P1	P2	P3	…	P38	…	P88	P89	P90
U1	1		4	…		…		3	
U2		2	2	…		…	1		
…				…		…			
U46	**?**	**2**	**?**	**…**	**?**	**…**	**?**	**?**	**4**
…				…		…			
U98	5			…		…	4		
U99			1	…		…		2	

1: very bad, 2: bad, 3: fair, 4: good, 5: excellent

Inputs and outputs as identifiers, absolute preferences in terms of ordinal degrees.

Dyadic Prediction [Menon & Elkan 2010]

Additional **side-information**: observed features + latent features of users and items

					?	?	?	?	?	?	?	?	?
					10	14	45	32	52	61	16	33	53
					P1	**P2**	**P3**	**…**	**P38**	**…**	**P88**	**P89**	**P90**
?	?	1	5	**U1**	1		4	…		…		3	
?	?	0	4	**U2**		2	2	…		…	1		
?	?	0	6	**…**				…		…			
?	?	1	5	**U46**	**?**	**2**	**?**	**…**	**?**	**…**	**?**	**?**	**4**
?	?	1	7	**…**				…		…			
?	?	0	6	**U98**	5			…		…	4		
?	?	1	6	**U99**			1	…		…		2	

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
multilabel 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

ranking

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

Loss Functions

Specification of a machine learning problem

- What kind of **training data** is offered to the learning algorithm?
- What **type of model** (prediction) is the learner supposed to produce?
- What is the nature of the **ground truth**,
- and how is a prediction assessed (**loss function**)?
  → *part 2*

Loss Functions

Things to be compared:

absolute utility degree	⟷	absolute utility degree	standard comparison of scalar predictions
subset of preferred items	⟷	subset of preferred items	non-standard comparisons
fuzzy subset of preferred items	⟷	fuzzy subset of preferred items	
subset of preferred items	⟷	ranking of items	
ranking of items	⟷	ordered partition of items	
ranking of items	⟷	ranking of items	

prediction **ground truth**

References

- W. Cheng, K. Dembczynski and E. Hüllermeier. *Graded Multilabel Classification: The Ordinal Case*. ICML-2010, Haifa, Israel, 2010.
- W. Cheng and E. Hüllermeier. *Predicting partial orders: Ranking with abstention*. ECML/PKDD-2010, Barcelona, 2010.
- Y. Chevaleyre, F. Koriche, J. Lang, J. Mengin, B. Zanuttini. *Learning ordinal preferences on multiattribute domains: The case of CP-nets*. In: J. Fürnkranz and E. Hüllermeier (eds.) Preference Learning, Springer-Verlag, 2010.
- W.W. Cohen, R.E. Schapire and Y. Singer. *Learning to order things*. Journal of Artificial Intelligence Research, 10:243–270, 1999.
- J. Fürnkranz, E. Hüllermeier, E. Mencia, and K. Brinker. *Multilabel Classification via Calibrated Label Ranking*. Machine Learning 73(2):133-153, 2008.
- J. Fürnkranz, E. Hüllermeier and S. Vanderlooy. *Binary decomposition methods for multipartite ranking*. Proc. ECML-2009, Bled, Slovenia, 2009.
- D. Goldberg, D. Nichols, B.M. Oki and D. Terry. *Using collaborative filtering to weave and information tapestry*. Communications of the ACM, 35(12):61–70, 1992.
- E. Hüllermeier, J. Fürnkranz, W. Cheng and K. Brinker. *Label ranking by learning pairwise preferences*. Artificial Intelligence, 172:1897–1916, 2008.
- G. Tsoumakas and I. Katakis. *Multi-label classification: An overview*. Int. J. Data Warehouse and Mining, 3:1–13, 2007.
- A.K. Menon and C. Elkan. *Predicting labels for dyadic data*. Data Mining and Knowledge Discovery, 21(2), 2010