Preference Learning

Johannes Fürnkranz, Eyke Hüllermeier

By combining practical relevance with novel types of prediction problems, the learning from/of preferences has recently received a lot of attention in the machine learning literature. Just as other types of complex learning tasks, preference learning deviates strongly from the standard problems of classification and regression. It is particularly challenging because it involves the prediction of complex structures, such as weak or partial order relations, rather than single values. This article aims at conveying a first idea of typical preference learning problems. To this end, two particular learning scenarios will be sketched, namely learning from label preferences and learning from object preferences. Both scenarios can be handled in two fundamentally different ways: by evaluating individual candidates (using a utility function) or by comparing competing candidates (using a binary "is preferred to" predicate).

1 Introduction

Recently, the topic of *preferences* has attracted considerable attention in Artificial Intelligence (AI) research, notably in fields such as agents, non-monotonic reasoning, constraint satisfaction, planning, and qualitative decision theory [2]. Preferences provide a means for specifying desires in a declarative way, which is a point of critical importance for AI. In fact, consider AI's paradigm of a rationally acting (decision-theoretic) agent: The behavior of such an agent has to be driven by an underlying preference model, and an agent recommending decisions or acting on behalf of a user should clearly reflect that user's preferences. Therefore, the formal modeling of preferences can be considered an essential aspect of autonomous agent design.

Drawing on past research on knowledge representation and reasoning, AI offers qualitative and symbolic methods for treating preferences that can reasonably complement standard approaches from economic decision theory, namely numerical *utility functions* and binary *preference relations*. Needless to say, however, the acquisition of preferences is not always an easy task. Therefore, not only are modeling languages and representation formalisms needed, but also methods for the automatic learning, discovery and adaptation of preferences. For example, computerized methods for discovering the preferences of individuals are useful in e-commerce and various other fields where an increasing trend toward personalization of products and services can be recognized.

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

Needless to say, this short article can neither provide a systematic exposition of different types of preference learning problems nor a comprehensive survey of recent literature in the field. In the remainder, we rather restrict ourselves to the discussion of two particular learning problems that appear to be especially interesting from a machine learning point of view.

2 Learning Label Preferences

The first learning scenario studies the problem of predicting, for any instance x (e.g. a person) from an instance space $\mathcal{X}$, a preference relation $\mathcal{P}_x \subseteq \mathcal{L} \times \mathcal{L}$ among a finite set $\mathcal{L}$ of labels or alternatives λ (e.g. politicians in an election). Here, $(\lambda, \lambda') \in \mathcal{P}_x$ means that the instance x prefers the label λ to the label λ', also written as $\lambda \succ_x \lambda'$.[1] The training information consists of a set of instances for which (partial) knowledge about the associated preference relation is available. More formally:

Given:

- a set of training instances $\{x_k \,|\, k = 1 \dots n\} \subseteq \mathcal{X}$ (encoded in an attribute-value representation)
- a set of labels $\mathcal{L} = \{\lambda_i \,|\, i = 1 \dots c\}$
- for each training instance x_k: a set of *pairwise preferences* of the form $\lambda_i \succ_{x_k} \lambda_j$

Find: a function that predicts $\mathcal{P}_x \subseteq \mathcal{L} \times \mathcal{L}$ for any query $x \in \mathcal{X}$

This learning setting can be seen as a generalization of several standard settings. In particular, the following problems are special cases of preference learning in the above sense:

[1] Extensions that distinguish between weak preferences ($\succeq$) and strict preferences ($\succ$), or allow the specification of indifference, are possible. Usually, $\lambda \succ_x \lambda'$ denotes strict preference and is short for $(\lambda \succeq_x \lambda') \wedge (\lambda' \not\succeq_x \lambda)$; moreover, $\lambda \sim_x \lambda'$ denotes indifference and is short for $(\lambda \succeq_x \lambda') \wedge (\lambda' \succeq_x \lambda))$.

- *Classification:* A single class label λ_i is assigned to each example x_k. This implicitly defines the set of preferences $\{\lambda_i \succ_{x_k} \lambda_j \,|\, 1 \leq j \neq i \leq c\}$.
- *Multi-label classification:* Each training example x_k is associated with a subset $S_k \subseteq \mathcal{L}$ of possible labels. This implicitly defines the set of preferences $\{\lambda_i \succ_{x_k} \lambda_j \,|\, \lambda_i \in S_k, \lambda_j \in \mathcal{L} \setminus S_k\}$.
- *Ranking:* Each training example x_k is associated with a total order of the labels, i.e., $\succ_{x_k}$ is a transitive relation such that $\lambda_i \succ_{x_k} \lambda_j$ or $\lambda_j \succ_{x_k} \lambda_i$ holds for each pair of labels (λ_i, λ_j), $i \neq j$.

There are two natural ways for representing preferences with regard to a set of alternatives, namely to *evaluate* individual candidates and to *compare* competing candidates. Correspondingly, there are two natural approaches to preference learning in the above setting. A first idea is to estimate a kind of utility degree (score) $f_i(x)$ for each alternative λ_i. To obtain a ranking, for example, the alternatives can then be ordered according to these utility degrees. A corresponding method for learning the functions $f_i(\cdot)$, $i = 1 \ldots c$, from training data has been proposed in [4].

Following the second approach, [3] propose to learn, for each pair of labels (λ_i, λ_j), a binary predicate $Q_{ij}(x)$ that predicts whether $\lambda_i \succ_x \lambda_j$ or $\lambda_j \succ_x \lambda_i$ for an input x. In order to rank the labels for a new object, predictions for all pairwise label preferences are obtained and a ranking that is maximally consistent with these preferences is derived. This approach is a natural extension of pairwise classification, i.e., the idea to tackle a multi-class classification problem by learning separate theories for each pair of classes.

3 Learning Object Preferences

The second scenario studies the problem of "learning to order things" such as, e.g., web pages retrieved by a search engine [6]. More precisely, the problem is to learn a function that is able to rank any subset O of an underlying class of objects C. The objects themselves are typically characterized by a finite set of features as in conventional attribute-value learning. Again, the training data consists of a set of pairwise preferences. More formally:

Given:
- a set of training objects $D = \{o_k \,|\, k = 1 \ldots n\} \subseteq C$ (encoded in an attribute-value representation)
- a set of pairwise preferences $P \subseteq D \times D$, where $(o_i, o_j) \in P$ indicates that o_i is preferred to o_j (written as $o_i \succ o_j$)

Find: a function that rank orders any set of objects $O \subseteq C$

Again, the two basic strategies discussed above can be applied in this setting: Typically, the problem is solved by learning a utility function $C \rightarrow \mathbb{R}$ that assigns a utility score to each object and to rank order the objects in O according to their scores [7, 5]. The second strategy has been applied in [1]. Here, a binary preference predicate $Q(o, o')$ is learned, which predicts whether o is preferred to o' or vice versa. A final ordering is found by deriving a ranking that is maximally consistent with these predictions.

4 Concluding Remarks

The goal of this article is to motivate preference learning as a theoretically interesting and practically relevant subfield of machine learning. To convey a first idea of problem types arising in this field, we have briefly outlined two closely related learning tasks. We also sketched two approaches that have been applied to both scenarios in the recent literature.

Of course, the field as such is still in its beginning. In fact, open lines of research in preference learning range from establishing a solid theoretical foundation for particular types of learning problems to putting things into practice by adapting idealized settings to more realistic learning scenarios. The increasing activity in this area is also witnessed by several workshops that have been devoted to preference learning and related topics, such as those at the NIPS-02, KI-03, SIGIR-03, NIPS-04, and GfKl-05 conferences.

References

[1] W.W. Cohen, R.E. Schapire, and Y. Singer. Learning to order things. *Journal of Artificial Intelligence Research*, 10:243–270, 1999.

[2] J. Doyle. Prospects for Preferences. *Computational Intelligence*, 20(2):111–136, 2004.

[3] J. Fürnkranz and E. Hüllermeier. Pairwise preference learning and ranking. In Proc. ECML–03, pages 145–156, Cavtat, Croatia, 2003. Springer-Verlag.

[4] S. Har-Peled, D. Roth, and D. Zimak. Constraint classification: A new approach to multiclass classification. In Proc. ALT–02, pages 365–379, Lübeck, Germany, 2002. Springer.

[5] R. Herbrich, T. Graepel, P. Bollmann-Sdorra, and K. Obermayer. Supervised learning of preference relations. In FGML–98, *Fachgruppentreffen Maschinelles Lernen*, pages 43–47, 1998.

[6] T. Joachims. Optimizing search engines using clickthrough data. Proc. KDD–02, 2002.

[7] Gerald Tesauro. Connectionist learning of expert preferences by comparison training. In D. Touretzky, editor, *Advances in Neural Information Processing Systems 1 (NIPS-88)*, pages 99–106. Morgan Kaufmann, 1989.

Contact

Johannes Fürnkranz
TU Darmstadt, FB Informatik, FG Knowledge Engineering
Hochschulstr. 10, 64289 Darmstadt
Email: `fuernkranz@informatik.tu-darmstadt.de`
`http://www.ke.informatik.tu-darmstadt.de`

Eyke Hüllermeier
Philipps-Universität Marburg, FB Mathematik/Informatik
Hans-Meerwein Straße, Lahnberge, 35032 Marburg
Email: `eyke@mathematik.uni-marburg.de`