Preference-based CBR: General Ideas and Basic Principles

Eyke Hüllermeier and Weiwei Cheng
Department of Mathematics and Computer Science
University of Marburg, Germany
{eyke,cheng}@mathematik.uni-marburg.de

Abstract

Building on recent research on preference handling in artificial intelligence and related fields, our goal is to develop a coherent and generic methodological framework for case-based reasoning (CBR) on the basis of formal concepts and methods for knowledge representation and reasoning with preferences. A preference-based approach to CBR appears to be appealing for several reasons, notably because case-based experiences naturally lend themselves to representations in terms of preference or order relations. Moreover, the flexibility and expressiveness of a preference-based formalism well accommodate the uncertain and approximate nature of case-based problem solving. In this paper, we outline the basic ideas of preference-based CBR and sketch a formal framework for realizing these ideas.

1 Introduction

Case-based reasoning (CBR) is a problem solving paradigm built upon a rule of thumb suggesting that "similar problems tend to have similar solutions" [Kolodner, 1993; Aamodt and Plaza, 1994]. More specifically, the idea of CBR is to exploit the experience from similar problems in the past and to adapt then successful solutions to the current situation. Thus, the core of every case-based problem solver is the case base, which is a collection of memorized "chunks of experience", called cases. Besides, the concept of *similarity* plays a pivotal role in every CBR system.

Despite its great practical success, work on the theoretical foundations of CBR is still under way, and a coherent and universally applicable methodological framework is yet missing. The *CBR cycle* proposed by Aamodt and Plaza [1994] is a commonly accepted process model, which nicely illustrates the main aspects of the case-based problem solving paradigm. Likewise, the metaphor of *knowledge containers*, introduced by Richter [Richter, 1995], provides a general framework for the structuring of knowledge in CBR. However, both are high-levels models and still rather far from the conceptual realization and implementation of a case-based problem solver. On the other extreme, many CBR systems have been designed for tackling concrete problems. These, however, are mostly tailored for a specific purpose and not easily applicable to a wider range of problems.

In-between these two extremes, conceptual models and practical systems, there is arguably a need for developing CBR *methodologies* [Watson, 1998]. On the one hand, a CBR methodology should be sufficiently general and abstract, so as to allow for the development of generic algorithms, for analyzing formal properties, etc. On the other hand, it should also be sufficiently concrete, so as to support the development of specific applications. To make this idea more tangible, consider as an analogy the formalism of graphical models, by now an established methodology for the design of probabilistic expert systems [Koller and Friedman, 2009]. This class of models disposes of a formal theory and generic algorithms, but also tools for supporting the design of models for concrete applications.

In [Hüllermeier and Schlegel, 2011], we have made a first step toward a methodological framework for case-based reasoning on the basis of formal concepts and methods for knowledge representation and problem solving with *preferences*. The topic of preferences has recently attracted considerable attention in artificial intelligence (AI) research and plays an increasingly important role in several AI-related fields, including, e.g., agents, constraint satisfaction, decision theory, planning, machine learning, and argumentation [Doyle, 2004; Goldsmith and Junker, 2008; Domshlak *et al.*, 2011]. Preference-based methods are especially appealing from an AI perspective, notably as they allow one to specify desires in a declarative way, to combine qualitative and quantitative modes of reasoning and to deal with inconsistencies and exceptions in a quite flexible manner. Indeed, a preference can be considered as a relaxed constraint, which, if necessary, can be violated to some degree.

The important advantage of an increased flexibility of a preference-based problem solving paradigm is nicely explained by Brafman and Domshlak [2009]: "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."

Our claim is that the above insights do not only apply to AI in general but to CBR in particular. In fact, as will be argued in more detail below, case-based experience can be modeled in terms of preference information in a quite convenient way and, moreover, case-based inference can be realized quite elegantly in the form of preference processing. As pointed out in the above quotation, a key advantage in comparison to the classical (constraint-based) approach [Hüllermeier, 2007] is an increased flexibility and expressiveness, which appears to be especially advantageous for CBR.

The remainder of the paper is organized as follows. In the next section, the main ideas of our approach to preference-based CBR are outlined in an informal way. A formal framework of preference-based CBR as well as two core components of this framework, namely *inference* and *search*, is then introduced in Section 3. The paper ends with some concluding remarks and an outlook on future work in Section 4.

2 Preference-based CBR: General Ideas

Experience in CBR is most commonly (albeit not exclusively) represented in the form of problem/solution tuples $(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{X} \times \mathbb{Y}$, where $\boldsymbol{x}$ is an element of a problem space $\mathbb{X}$, and $\boldsymbol{y}$ an element of a solution space $\mathbb{Y}$. These two spaces can be as simple as Euclidean or categorical spaces (like in classification or regression problems), but may also be very complex; for example, in the problem of text summarization, $\mathbb{X}$ could be the "space" of scientific articles and $\mathbb{Y}$ the "space" of abstracts [Capus and Tourigny, 2003].

Despite its generality and expressiveness, the standard problem/solution representation exhibits some limitations, both from a knowledge acquisition and reuse point of view.

- *Existence of correct solutions*: It assumes the existence of a "correct" solution for each problem, and implicitly even its uniqueness. This assumption is often not tenable. In text summarization, for example, there is not a single "correct" abstract of an article. Instead, there will be many possible alternatives, maybe more or less preferred by the user.
- *Verification of optimality*: Even if the existence of a single correct solution for each problem could be assured, it will often be impossible to verify the optimality of the solution that has been produced by a CBR system. Storing a suboptimal solution $\boldsymbol{y}$ for a problem $\boldsymbol{x}$, however, and later on reusing this solution as if it were optimal, may mislead future problem solving. This issue is less critical, though does not dissolve completely, if only "acceptable" instead of optimal solutions are required.
- *Loss of information*: Storing only a single solution $\boldsymbol{y}$ for a problem $\boldsymbol{x}$, even if it can be guaranteed to be optimal, may come along with a potential loss of information. In fact, during a problem solving episode, one typically tries or at least compares several candidate solutions, and even if these solutions are suboptimal, preferences between them may provide useful information.
- *Limited guidance*: From a reuse point of view, a retrieved case $(\boldsymbol{x}, \boldsymbol{y})$ only suggests a single solution, namely $\boldsymbol{y}$, for a query problem $\boldsymbol{x}_0$. Thus, it does not imply a possible course of action in the case where the suggestion fails: If $\boldsymbol{y}$ is not a good point of departure, for example since it cannot be adapted to solve $\boldsymbol{x}_0$, there is no concrete recommendation on how to continue.

To avoid these problems, preference-based CBR replaces experiences of the form "solution $\boldsymbol{y}$ (optimally) solves problem $\boldsymbol{x}$" by weaker information of the form "$\boldsymbol{y}$ is better (more preferred) than $\boldsymbol{z}$ as a solution for $\boldsymbol{x}$", that is, by a preference between two solutions "contextualized" by a problem $\boldsymbol{x}$. More specifically, the basic "chunk of information" we consider is symbolized in the form $\boldsymbol{y} \succeq_{\boldsymbol{x}} \boldsymbol{z}$ and suggests that, for the problem $\boldsymbol{x}$, the solution $\boldsymbol{y}$ is supposedly at least as good as $\boldsymbol{z}$.

This type of knowledge representation overcomes many of the problems discussed above. As soon as two candidate solutions $\boldsymbol{y}$ and $\boldsymbol{z}$ have been tried as solutions for a problem $\boldsymbol{x}$, these two alternatives can be compared and, correspondingly, a strict preference in favor of one of them or an indifference can be expressed. To this end, it is by no means required that one of these solutions is optimal. It is worth mentioning, however, that knowledge about the optimality of a solution $\boldsymbol{y}^*$, if available, can be handled, too, as it simply means that $\boldsymbol{y}^* \succ \boldsymbol{y}$ for all $\boldsymbol{y} \neq \boldsymbol{y}^*$. In this sense, the conventional CBR setting can be considered as a special case of preference-based CBR.

The above idea of a preference-based approach to knowledge representation in CBR also suggests a natural extension of the case retrieval and inference steps, that is, the recommendation of solutions for a new query problem: Instead of just proposing a single solution, it would be desirable to predict a *ranking* of several (or even all) candidate solutions, ordered by their (estimated) degree of preference:

$$\boldsymbol{y}_1 \succeq_{\boldsymbol{x}} \boldsymbol{y}_2 \succeq_{\boldsymbol{x}} \boldsymbol{y}_3 \succeq_{\boldsymbol{x}} \cdots \succeq_{\boldsymbol{x}} \boldsymbol{y}_n \quad (1)$$

Obviously, the last problem mentioned above, namely the lack of guidance in the case of a failure, can thus be overcome.

In order to realize an approach of that kind, a number of important questions need to be addressed, including the following: How to represent, organize and maintain case-based experiences, given in the form of preferences referring to a specific context, in an efficient way? How to select and access the experiences which are most relevant in a new problem solving situation? How to combine these experiences and exploit them to infer a solution or, more generally, a preference order on a set of candidate solutions, for the problem at hand?

3 A Formal Framework

In the following, we assume the problem space $\mathbb{X}$ to be equipped with a similarity measure $S_X : \mathbb{X} \times \mathbb{X} \to \mathbb{R}_+$ or, equivalently, with a (reciprocal) distance measure Δ_X :

$\mathbb{X} \times \mathbb{X} \to \mathbb{R}_+$. Thus, for any pair of problems $\boldsymbol{x}, \boldsymbol{x}' \in \mathbb{X}$, their similarity is denoted by $S_X(\boldsymbol{x}, \boldsymbol{x}')$ and their distance by $\Delta_X(\boldsymbol{x}, \boldsymbol{x}')$. Likewise, we assume that the solution space $\mathbb{Y}$ to be equipped with a similarity measure S_Y or, equivalently, with a (reciprocal) distance measure Δ_Y. In general, $\Delta_Y(\boldsymbol{y}, \boldsymbol{y}')$ can be thought of as a kind of adaptation cost, i.e., the (minimum) cost that needs to be invested to transform the solution $\boldsymbol{y}$ into $\boldsymbol{y}'$.

In preference-based CBR, problems $\boldsymbol{x} \in \mathbb{X}$ are not associated with single solutions but rather with preferences over solutions, that is, with elements from a class of preference structures $\mathfrak{P}(\mathbb{Y})$ over the solution space $\mathbb{Y}$. Here, we make the assumption that $\mathfrak{P}(\mathbb{Y})$ is given by the class of all weak order relations $\succeq$ on $\mathbb{Y}$, and we denote the relation associated with a problem $\boldsymbol{x}$ by $\succeq_{\boldsymbol{x}}$. More precisely, we assume that $\succeq_{\boldsymbol{x}}$ has a specific form, which is defined by an "ideal" solution[1] $\boldsymbol{y}^* \in \mathbb{Y}$ and the distance measure Δ_Y: The closer a solution $\boldsymbol{y}$ to $\boldsymbol{y}^* = \boldsymbol{y}^*(\boldsymbol{x})$, the more it is preferred; thus, $\boldsymbol{y} \succeq_{\boldsymbol{x}} \boldsymbol{z}$ iff $\Delta_Y(\boldsymbol{y}, \boldsymbol{y}^*) \leq \Delta_Y(\boldsymbol{z}, \boldsymbol{y}^*)$. In conjunction with the regularity assumption that is commonly made in CBR, namely that similar problems tend to have similar (ideal) solutions, this property legitimates a preference-based version of this assumption: *Similar problems are likely to induce similar preferences over solutions.*

3.1 Case-based Inference

The key idea of preference-based CBR is to exploit experience in the form of previously observed preferences, deemed relevant for the problem at hand, in order to support the current problem solving episode; like in standard CBR, the *relevance* of a preference will typically be decided on the basis of problem similarity, i.e., those preferences will be deemed relevant that pertain to similar problems. An important question that needs to be answered in this connection is the following: Given a set of observed preferences on solutions, considered representative for a problem $\boldsymbol{x}_0$, what is the underlying preference structure $\succeq_{\boldsymbol{x}}$ or, equivalently, what is the most likely ideal solution $\boldsymbol{y}^*$ for $\boldsymbol{x}_0$?

Case-based Inference as Probability Estimation

We approach this problem from a statistical perspective, considering the true preference model $\succeq_{\boldsymbol{x}_0} \in \mathfrak{P}(\mathbb{Y})$ associated with the query $\boldsymbol{x}_0$ as a random variable with distribution $\mathbf{P}(\cdot \,|\, \boldsymbol{x}_0)$, where $\mathbf{P}(\cdot \,|\, \boldsymbol{x}_0)$ is a distribution $\mathbf{P}_\theta(\cdot)$ parametrized by $\theta = \theta(\boldsymbol{x}_0) \in \Theta$. The problem is then to estimate this distribution or, equivalently, the parameter θ on the basis of the information available. This information consists of a set $\mathcal{D}$ of preferences of the form $\boldsymbol{y} \succ \boldsymbol{z}$ between solutions.

The basic assumption underlying nearest neighbor estimation is that the conditional probability distribution of the output given the input is (approximately) locally constant, that is, $\mathbf{P}(\cdot \,|\, \boldsymbol{x}_0) \approx \mathbf{P}(\cdot \,|\, \boldsymbol{x})$ for $\boldsymbol{x}$ close to $\boldsymbol{x}_0$. Thus, if the above preferences are coming from problems $\boldsymbol{x}$ similar to $\boldsymbol{x}_0$ (namely from the nearest neighbors of $\boldsymbol{x}_0$ in the case base), then this assumption justifies considering $\mathcal{D}$ as a representative sample of $\mathbf{P}_\theta(\cdot)$ and, hence, estimating θ via maximum likelihood (ML) inference by

$$\theta^{ML} = \arg\max_{\theta \in \Theta} \mathbf{P}_\theta(\mathcal{D}) \ . \qquad (2)$$

An important prerequisite for putting this approach into practice is a suitable data generating process, i.e., a process generating preferences in a stochastic way.

A Discrete Choice Model

Our data generating process is based on the idea of a discrete choice model as used in choice and decision theory [Peterson, 2009]. Recall that the (absolute) preference for a solution $\boldsymbol{y} \in \mathbb{Y}$ supposedly depends on its distance $\Delta_Y(\boldsymbol{y}, \boldsymbol{y}^*) \geq 0$ to an "ideal" solution $\boldsymbol{y}^*$, where $\Delta(\boldsymbol{y}, \boldsymbol{y}^*)$ can be seen as a "degree of suboptimality" of $\boldsymbol{y}$. As explained in [Hüllermeier and Schlegel, 2011], more specific assumptions on an underlying (latent) utility function on solutions justify the *logit* model of discrete choice:

$$\mathbf{P}(\boldsymbol{y} \succ \boldsymbol{z}) = \frac{\exp(-\beta(\Delta_Y(\boldsymbol{y}, \boldsymbol{y}^*))}{\exp(-\beta(\Delta_Y(\boldsymbol{z}, \boldsymbol{y}^*)) + \exp\big(-\beta\Delta_Y(\boldsymbol{y}, \boldsymbol{y}^*)\big)}$$

Thus, the probability of observing the (revealed) preference $\boldsymbol{y} \succ \boldsymbol{z}$ depends on the degree of suboptimality of $\boldsymbol{y}$ and $\boldsymbol{z}$, namely their respective distances to the ideal solution, $\Delta_Y(\boldsymbol{y}, \boldsymbol{y}^*)$ and $\Delta_Y(\boldsymbol{z}, \boldsymbol{y}^*)$: The larger the difference $\Delta_Y(\boldsymbol{z}, \boldsymbol{y}^*) - \Delta_Y(\boldsymbol{y}, \boldsymbol{y}^*)$, i.e., the less optimal $\boldsymbol{z}$ in comparison to $\boldsymbol{y}$, the larger the probability to observe $\boldsymbol{y} \succ \boldsymbol{z}$; if $\Delta_Y(\boldsymbol{z}, \boldsymbol{y}^*) = \Delta_Y(\boldsymbol{y}, \boldsymbol{y}^*)$, then $\mathbf{P}(\boldsymbol{y} \succ \boldsymbol{z}) = 1/2$. The coefficient β can be seen as a measure of precision of the preference feedback. For large β, $\mathbf{P}(\boldsymbol{y} \succ \boldsymbol{z})$ converges to 0 if $\Delta_Y(\boldsymbol{z}, \boldsymbol{y}^*) < \Delta_Y(\boldsymbol{y}, \boldsymbol{y}^*)$ and to 1 if $\Delta_Y(\boldsymbol{z}, \boldsymbol{y}^*) > \Delta_Y(\boldsymbol{y}, \boldsymbol{y}^*)$; this corresponds to a deterministic (error-free) information source. The other extreme case, namely $\beta = 0$, models a completely unreliable source reporting preferences at random.

Maximum Likelihood Estimation

The probabilistic model outlined above is specified by two parameters: the ideal solution $\boldsymbol{y}^*$ and the (true) precision parameter $\beta^* \in \mathbb{R}_+$. Depending on the context in which these parameters are sought, the ideal solution might be unrestricted (i.e., any element of $\mathbb{Y}$ is an eligible candidate), or it might be restricted to a certain subset $\mathbb{Y}_0 \subseteq \mathbb{Y}$ of candidates.

Now, to estimate the parameter vector $\theta^* = (\boldsymbol{y}^*, \beta^*) \in \mathbb{Y}_0 \times \mathbb{R}^*$ from a given set $\mathcal{D} = \{\boldsymbol{y}^{(i)} \succ \boldsymbol{z}^{(i)}\}_{i=1}^N$ of observed preferences, we refer to the maximum likelihood estimation principle. Assuming independence of the preferences, the likelihood of $\theta = (\boldsymbol{y}, \beta)$ is given by

$$\ell(\theta) = \prod_{i=1}^N \mathbf{P}\left(\boldsymbol{y}^{(i)} \succ \boldsymbol{z}^{(i)} \,|\, \theta\right) \qquad (3)$$

The ML estimation $\theta_{ML} = (\boldsymbol{y}^{ML}, \beta^{ML})$ of θ^* is given by the maximizer of (3):

$$\theta_{ML} = \left(\boldsymbol{y}^{ML}, \beta^{ML}\right) = \arg\max_{\boldsymbol{y} \in \mathbb{Y}_0,\, \beta \in \mathbb{R}_+} \ell(\boldsymbol{y}, \beta) \qquad (4)$$

The problem of finding this estimation in an efficient way is discussed in [Hüllermeier and Schlegel, 2011].

[1]The solution $\boldsymbol{y}^*$ could be a purely imaginary solution, which may not exist in practice.

3.2 CBR as Preference-guided Search

Case-based inference as outlined above realizes a "one-shot prediction" of a promising solution for a query problem, given preferences in the context of similar problems encountered in the past. In a case-based problem solving process, this prediction may thus serve as an initial solution, which is then adapted step by step. Formally, an adaptation process of that kind can be formalized as a search process, namely a traversal of a suitable space of candidate solutions [Bergmann and Wilke, 1998].

In the spirit of preference-based CBR, we implement case-based problem solving as a search process that is guided by preference information collected in previous problem solving episodes. The type of application we have in mind is characterized by two important properties:

- *The evaluation of candidate solutions is expensive.* Therefore, only relatively few candidates can be considered in a problem solving episode before a selection is made. Typical examples include cases where an evaluation requires time-consuming simulation studies or human intervention (for example, the reading of a text summary).
- *The quality of candidate solutions is difficult to quantify.* Therefore, instead of asking for numerical utility degrees, we make a much weaker assumption: Feedback is only provided in the form of pairwise comparisons, informing about which of two candidate solutions is preferred (for example, which of two text summaries is better). Formally, we assume the existence of an "oracle" (for example, a user or a computer program) which, given a problem $\boldsymbol{x}_0$ and two solutions $\boldsymbol{y}$ and $\boldsymbol{z}$ as input, returns a preference $\boldsymbol{y} \succ \boldsymbol{z}$ or $\boldsymbol{z} \succ \boldsymbol{y}$ as output.

We assume the solution space $\mathbb{Y}$ to be equipped with a topology that is defined through a *neighborhood structure*: For each $\boldsymbol{y} \in \mathbb{Y}$, we denote by $\mathcal{N}(\boldsymbol{y}) \subseteq \mathbb{Y}$ the neighborhood of this candidate solution. The neighborhood is thought of as those solutions that can be produced through a single modification of $\boldsymbol{y}$, i.e., by applying one of the available adaptation operators to $\boldsymbol{y}$ (for example, adding, removing or modifying a single sentence in an abstract).

Our case base $\mathbf{CB}$ stores problems $\boldsymbol{x}_i$ together with a set of preferences $\mathcal{P}(\boldsymbol{x}_i)$ that have been observed for these problems. Thus, each $\mathcal{P}(\boldsymbol{x}_i)$ is a set of preferences of the form $\boldsymbol{y} \succ_{\boldsymbol{x}_i} \boldsymbol{z}$, which are collected while searching for a good solution to $\boldsymbol{x}_i$.

We conceive preference-based CBR as an iterative process in which problems are solved one by one. In each problem solving episode, a good solution for a new query problem is sought, and new experiences in the form of preferences are collected. In what follows, we give a high-level description of a single problem solving episode:

(i) Given a new query problem $\boldsymbol{x}_0$, the K nearest neighbors $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_K$ of this problem (i.e., those with smallest distance in the sense of Δ_X) are retrieved from the case base $\mathbf{CB}$, together with their preference information $\mathcal{P}(\boldsymbol{x}_1), \ldots, \mathcal{P}(\boldsymbol{x}_K)$.

(ii) This information is collected in a single set of preferences $\mathcal{P}$, which is considered representative for the problem $\boldsymbol{x}_0$ and used to guide the search process.

(iii) The search for a solution starts with an initial candidate $\boldsymbol{y}^* \in \mathbb{Y}$, namely the "one-shot prediction" (4) based on $\mathcal{P}$, and iterates L times. Restricting the number of iterations by an upper bound L reflects our assumption that an evaluation of a candidate solution is costly.

(iv) In each iteration, a new candidate $\boldsymbol{y}^{query}$ is determined, again based on (4), and given as a query to the oracle, i.e., the oracle is asked to compare $\boldsymbol{y}^{query}$ with the current best solution $\boldsymbol{y}^*$. The preference reported by the oracle is memorized by adding it to the preference set $\mathcal{P}_0 = \mathcal{P}(\boldsymbol{x}_0)$ associated with $\boldsymbol{x}_0$, as well as to the set $\mathcal{P}$ of preferences used for guiding the search process. Moreover, the better solution is retained as the current best candidate.

(v) When the search stops, the current best solution $\boldsymbol{y}^*$ is returned, and the case $(\boldsymbol{x}_0, \mathcal{P}_0)$ is added to the case base.

The preference-based guidance of the search process is realized in (iii) and (iv). Here, our case-based inference method is used to find the most promising candidate among the neighborhood of the current solution $\boldsymbol{y}^*$, based on the preferences collected in the problem solving episode so far. By providing information about which of these candidates will most likely constitute a good solution for $\boldsymbol{x}_0$, it (hopefully) points the search into the most promising direction.

For a detailed description of the above search procedure, as well as first experimental studies, we refer to [Abdel-Aziz *et al.*, 2013].

4 Conclusion

In this paper, we have presented the idea of a preference-based approach to CBR: Experience is represented in the form of preferences, which are "contextualized" by a previously solved problem, and these preferences are used to direct the search for a good solution to a new problem. This is an interesting alternative to conventional CBR whenever solution quality is a matter of degree and feedback is only provided in an indirect or qualitative way. We have introduced a general framework of preference-based CBR, in which *prediction* is realized through statistical inference and *problem solving* is formalized as (heuristic) search.

For future work, we plan to generalize our framework in various directions, and extend it by adding components that are essential for a complete and effective CBR system. For example, since the number of preferences collected in the course of time may become rather large, effective methods for case base maintenance ought to be developed [Smyth and McKenna, 2001]. Besides, we are of course interested in testing and evaluating our approach in different application domains.

Acknowledgments

This work has been supported by the German Research Foundation (DFG).

References

[Aamodt and Plaza, 1994] A. Aamodt and E. Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, 7(1):39–59, 1994.

[Abdel-Aziz *et al.*, 2013] A. Abdel-Aziz, W. Cheng, M. Strickert, and E. Hüllermeier. Preference-based CBR: A search-based problem solving framework. 2013. Submitted.

[Bergmann and Wilke, 1998] R. Bergmann and W. Wilke. Towards a new formal model of transformational adaptation in case-based reasoning. In H. Prade, editor, ECAI-98*, 13th European Conference on Artificial Intelligence*, pages 53–57, 1998.

[Brafman and Domshlak, 2009] R.I. Brafman and C. Domshlak. Preference handling–an introductory tutorial. *AI Magazine*, 30(1), 2009.

[Capus and Tourigny, 2003] L. Capus and N. Tourigny. A case-based reasoning approach to support story summarization. *International Journal of General Systems*, 18(8):877–891, 2003.

[Domshlak *et al.*, 2011] C. Domshlak, E. Hüllermeier, S. Kaci, and H. Prade. Preferences in AI: An overview. *Artificial Intelligence*, 2011.

[Doyle, 2004] J. Doyle. Prospects for preferences. *Comput. Intell.*, 20(2):111–136, 2004.

[Goldsmith and Junker, 2008] J. Goldsmith and U. Junker. Special issue on preference handling for Artificial Intelligence. *Computational Intelligence*, 29(4), 2008.

[Hüllermeier and Schlegel, 2011] E. Hüllermeier and P. Schlegel. Preference-based CBR: First steps toward a methodological framework. In A. Ram and N. Wiratunga, editors, *Proceedings ICCBR–2011, 19th International Conference on Case-Based Reasoning*, pages 77–91. Springer-Verlag, 2011.

[Hüllermeier, 2007] E. Hüllermeier. *Case-Based Approximate Reasoning*, volume 44 of *Theory and Decision Library, Series B: Mathematical and Statistical Methods*. Springer-Verlag, Heidelberg, Berlin, 2007. 370 pages.

[Koller and Friedman, 2009] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, Massachusetts, 2009.

[Kolodner, 1993] J.L. Kolodner. *Case-based Reasoning*. Morgan Kaufmann, San Mateo, 1993.

[Peterson, 2009] M. Peterson. *An Introduction to Decision Theory*. Cambridge Univ. Press, 2009.

[Richter, 1995] M.M. Richter. The knowledge contained in similarity measrues. Invited talk at ICCBR-95. `wwwagr.informatik.uni-kl.de/~lsa/cbr/richtericcbr95remarks.html`, 1995.

[Smyth and McKenna, 2001] B. Smyth and E. McKenna. Competence models and the maintenance problem. *Computational Intelligence*, 17(2):235–249, 2001.

[Watson, 1998] I. Watson. Case-based reasoning is a methodology not a technology. In R. Mile, M. Moulton, and M. Bramer, editors, *Research and Development in Expert Systems XV*, pages 213–223. London, 1998.