7. Case-Based Decision Making

Early work in AI has mainly focused on formal logic as a basis of knowledge representation and has largely rejected approaches from (statistical) decision theory as being intractable and inadequate for expressing the rich structure of (human) knowledge [193]. However, the recent development of more tractable and expressive decision-theoretic frameworks and inference strategies such as, e.g., graphical formalisms [292, 187], in combination with the analysis of restrictions of traditional AI reasoning techniques have stimulated renewed interest in decision theory. In fact, ideas from decision theory now play a predominant role in the modeling of *rationality*, one of the major topics of current research in AI [95]. Loosely speaking, the AI paradigm has undergone a shift from "acting logically" to "acting rationally" [322]. The related view of intelligent behavior deviates fundamentally from the classical "logicist" approach. While the latter emphasizes the ability to reach correct conclusions from correct premises, the decision-theoretic approach considers AI as the design of (limited) *rational agents* [324]. For this "agent-based" view of AI, intelligence is strongly related to the capacity of successful behavior in complex and uncertain environments and, hence, to *rational decision making*.[1]

Decision theory and AI can fertilize each other in various ways [298]. As already suggested above, classical decision theory provides AI with important ideas and concepts of rationality, thus contributing to a formal basis of intelligent agent design. Yet, it has been less concerned with computational and knowledge representational aspects. AI can particularly contribute in this direction. It has been realized very soon, for instance, that *perfect rationality*, in the sense of generating behavior which leads to maximal (expected) utility, cannot be achieved once computational aspects come into play [321]. In fact, an agent having to make a decision under limited computational (time, memory) resources not only has to reason about the decision itself but also about the computations it uses for deriving the decision: A more elaborated computation might yield a better decision but also requires more time (or other resources). Being perfectly rational in the aforementioned sense, it has to perform the reasoning about its computations in the same decision-theoretic way. This, however, leads to the problem of realizing some kind of *metalevel rationality* [23, 38, 324] and, hence, results in a conceptual regress. Problems of this kind have motivated the definition of alternative

[1] J. Doyle has suggested to define AI itself as the computational study of rational behavior [94].

concepts which are weakenings of perfect rationality. They serve as candidates for putting the agent-based understanding of intelligence and the related approach to the design of intelligent systems on a formal basis. Among the proposals, the concept of *bounded optimality* seems to be the one which is most relevant for practical as well as theoretical AI research [323].

As far as the aspect of knowledge representation is concerned, research in AI has shown various possibilities of extending the decision-theoretic frameworks usually considered in classical approaches. Recent developments include the modeling of decision problems within qualitative [52, 53, 123, 129] and constraint-based [143] settings and make use of formal logic in order to represent the knowledge of a decision maker in a more flexible way [39, 50, 98, 293, 326, 365, 366]. These approaches are intended to make decision-theoretic models more realistic, tractable and expressive.

In this chapter, we are mainly concerned with the idea of *case-based decision making* (CBDM) which is originally due to Gilboa and Schmeidler [167]. The notion CBDM stands for the application of the CBI principle in the context of decision making: An agent faced with a decision problem relies upon its experience from similar problems encountered in the past. Loosely speaking, it chooses an act based on the (cumulative or average) performance of (potential) acts in previous problems which are similar to the current one.

Even though the model in [167] has mainly been introduced with economic applications in mind, CBDM is particularly interesting from an AI perspective. Firstly, it combines principles from two important subfields of AI, namely decision theory and CBR. Secondly, it touches on interesting aspects of knowledge representation and reasoning. In fact, the mental notions of *preference* and *belief* constitute the main concepts of classical decision theories. Corresponding mathematical models are based on formalizations of these concepts, such as preference relations, utility functions, and probability distributions. The aforementioned approach of Gilboa and Schmeidler leads to a decision theory in which the cognitive concept of *similarity* plays a predominant role. Needless to say, incorporating this concept into formal approaches to decision making raises some interesting (semantical) questions. Particularly, it has to be clarified which role similarity plays and, hence, what the relation between this and other concepts such as preference and belief could be (cf. Section 7.6). Clearly, this question concerns basic assumptions underlying a decision-theoretic model. One should therefore not expect to find definite answers. Classical works by Ramsey [309], de Finetti [146], von Neumann and Morgenstern [278] and Savage [331] as well as recent developments in the field of decision theory, such as non-additive expected utility [334, 166] or qualitative decision making, show various ways of formalizing the notions of preference and belief (including measure-theoretic approaches, such as fuzzy measures [384] and different types of probability [145], as well as more logic-oriented symbolic methods [365]).[2] Moreover, a consensus concerning the actual meaning of the concept itself seems to exist even less in the case of similarity than in the case of preference or uncertainty. As will be seen, the approaches to case-based decision making discussed in this chapter not do only differ with respect to the mathematical formalization, they are also based on different principles and ideas for incorporating similarity and principles of CBI into decision making.

The remaining part of the chapter is organized as follows: In Section 7.1, we provide a brief review and discussion of case-based decision theory as introduced by Gilboa and Schmeidler. In Section 7.2, we consider the idea of case-based decision making in connection with the Nearest Neighbor principle which is commonly used in instance-based learning. A fuzzy set-based approach to CBDM which is due to Dubois and Prade [101] will be discussed in Section 7.3. A generalization of the latter is proposed in Section 7.4. Section 7.5 is devoted to an alternative framework of case-based decision making which is based on the methods of case-based inference proposed in previous chapters. A discussion of some selected aspects of CBDM models follows in Section 7.6. Finally, Section 7.7 introduces a framework of *experienced-based decision making* as a generalization of case-based decision making.

7.1 Case-based decision theory

This section gives a brief review of the model introduced by Gilboa and Schmeidler [167], referred to as *case-based decision theory* (CBDT) by the authors. Putting it in a nutshell, the setup they proceed from can be characterized as follows: Let $\mathcal{Q}$ and $\mathcal{A}$ be (finite) sets of problems and acts, respectively, and denote by $\mathcal{R}$ a set of outcomes (outputs) or results. Choosing act $a \in \mathcal{A}$ for solving problem $p \in \mathcal{Q}$ leads to the outcome $r = r(p, a) \in \mathcal{R}$. A utility function $u : \mathcal{R} \longrightarrow U$ resp. $u : \mathcal{Q} \times \mathcal{A} \longrightarrow U$ assigns utility values to such outcomes; the utility scale U is taken as the set of real numbers. Let

$$\sigma_{\mathcal{Q}} : \mathcal{Q} \times \mathcal{Q} \longrightarrow [0, 1], \quad \sigma_{\mathcal{R}} : \mathcal{R} \times \mathcal{R} \longrightarrow [0, 1]$$

be similarity measures quantifying the similarity of problems and outputs, respectively. Suppose the decision making agent to have a (finite) memory

$$\mathcal{M} = \{(p_1, a_1, r_1), \ldots, (p_n, a_n, r_n)\} \tag{7.1}$$

of cases at its disposal, where $(p_k, a_k) \in \mathcal{Q} \times \mathcal{A}$, $r_k = r(p_k, a_k)$ $(1 \leq k \leq n)$, and suppose furthermore that it has to choose an act for a new problem $p_0 \in \mathcal{Q}$. If a certain act $a_0 \in \mathcal{A}$ has not been applied to the problem p_0 so far (i.e.,

[2] Needless to say, a validation or comparison of decision-theoretic models is generally difficult, no matter whether from a descriptive or a normative point of view.

there is no case $(p_0, a_0, r) \in \mathcal{M}$) the agent will generally be uncertain about the result $r(p_0, a_0)$ and, hence, about the utility $u(r(p_0, a_0))$. According to the assumption underlying the paradigm of CBDT it then evaluates an act based on its performance in similar problems in the past, as represented by (parts of) the memory $\mathcal{M}$. More precisely, the decision maker is supposed to choose an act which maximizes a linear combination of the benefits experienced so far:

$$V(a_0) = V_{p_0,\mathcal{M}}(a_0) \stackrel{\mathrm{df}}{=} \sum_{(p,a_0,r) \in \mathcal{M}} \sigma_{\mathcal{Q}}(p, p_0) \cdot u(r). \tag{7.2}$$

The summation over an empty set yields the "default value" 0 which plays the role of an "aspiration level." Despite the formal resemblance between (7.2) and the well-known expected utility formula one should not ignore some substantial differences between CBDT and expected utility theory (EUT). This concerns not only the conceptual level but also mathematical aspects. Particularly, it should be noted that the similarity weights in (7.2) do not necessarily sum up to 1. Consequently, (7.2) must not be interpreted as an estimation of the utility $u(r(p_0, a_0))$.

As an alternative to the linear functional (7.2), an "averaged similarity" version has been proposed. It results from replacing $\sigma_{\mathcal{Q}}$ in (7.2) by the similarity measure

$$(p, p_0) \mapsto \sigma_{\mathcal{Q}}(p, p_0) \left(\sum_{(p',a_0,r') \in \mathcal{M}} \sigma_{\mathcal{Q}}(p', p_0) \right)^{-1} \tag{7.3}$$

whenever the latter is well-defined. (Note that this measure is defined separately for each act a_0.) Theoretical details of CBDT including an axiomatic characterization of decision principle (7.2) are presented in [167].

The basic model has been generalized with respect to several aspects. The problem of optimizing decision behavior by adjusting the aspiration level in the context of repeated problem solving is considered in [168] (see also Section 7.6). In [169], the similarity measure in (7.2) is extended to problem–act tuples: Given two similar problems, it is assumed that similar outcomes are obtained for *similar* acts (not only for the same act). Indeed, it is argued convincingly that a model of the form

$$V(a_0) = \sum_{(p,a,r) \in \mathcal{M}} \sigma_{\mathcal{Q} \times \mathcal{A}}((p, a), (p_0, a_0)) \cdot u(r), \tag{7.4}$$

where $\sigma_{\mathcal{Q} \times \mathcal{A}}$ is a (problem–act) similarity measure over $\mathcal{Q} \times \mathcal{A}$, is more realistic than (7.2). For example, an act a_0 which has not been applied as yet is generally not evaluated by the default utility 0 if experiences with a comparable act a have been made. In fact, an outcome $r(p, a)$ will then influence the rating of a_0 in connection with a problem p_0 which is similar to p. Besides, it should be noticed that (7.4) allows for realizing some kind of analogical reasoning. Suppose, for instance, that the effect expected from applying a_0 to p_0 is comparable to the effect of applying a to p. In that sense, (a_0, p_0) might appear to be quite similar to (a, p), although a and a_0 as well as p and p_0 as such are rather dissimilar.

With regard to alternative models of CBDM proposed in subsequent sections it is useful to picture again the following properties of the decision criteria outlined above:

- Accumulation/averaging: The criteria (7.2) and (7.4) realize a simple summation of (weighted) degrees of utility. Consequently, a decision maker might prefer an act a, which always brought about rather poor results, to an act a^* which has so far yielded very good results, simply because a has been tried more often than a^*. This effect is annulled by (7.3), where the use of a normalized similarity measure yields an average utility.
- Compensation: Both decision rules compensate between good results and bad results associated with an act a.

Gilboa and Schmeidler especially emphasize the cognitive plausibility of their model [171]. In fact, a main motivation behind CBDT is to provide a more faithful description of human decision making than EUT does. Indeed, in some situations this axiomatic theory seems rather restrictive. Particularly, it assumes the decision maker to have very detailed information at its disposal: a list of the *states of nature*, a probability distribution over these states, a list of potential acts, and a numerical utility value for all act–state pairs.[3] Since this information is generally not completely available, the decision maker is forced to engage in *hypothetical reasoning*.[4] Moreover, some well-known paradoxes [13, 140] as well as psychological studies [375] show that EUT can be challenged as a *descriptive* theory of (human) decision making. Still, it deserves mentioning that CBDT is not seen as a competing theory, but rather as an alternative (or complementary) "language" for modeling decision problems. It seems especially useful if a problem description is not naturally cast in the framework of decision making under risk or if a problem is very unfamiliar, in which case the modeling of states of nature and associated probabilities might be difficult. A thorough discussion of the relation between CBDT and EUT can again be found in [167].

Let us conclude with a remark on the concept of similarity as used in CBDT. One might argue that the measures $\sigma_{\mathcal{Q}}$ and $\sigma_{\mathcal{Q}\times\mathcal{A}}$ need not be interpreted as similarities at all: Basically, the valuation (7.2) can be seen as a weighted sum

$$V(a_0) = V_{p_0,\mathcal{M}}(a_0) = \sum_{(p,a,r)\in\mathcal{M}} \omega_{p_0,a_0,\mathcal{M}}(p,a) \cdot u(r) \tag{7.5}$$

of utility degrees encountered in the past,[5] where the weights reflect the *relevance* of a case. This relevance, however, might not only depend on similarity. Rather,

[3] Still, it has to be noticed that an unequivocal model does generally not exist. Rather, there is much freedom in the definition of, e.g., states and acts.

[4] What is the effect of choosing a certain act in a certain state of nature?

[5] The linearity of the representation (7.2) is mainly due to the separability axiom in [167].

it can capture other (or further) aspects as well and, hence, leaves much freedom for different types of cognitive interpretation.[6] In this connection, it is worth mentioning that the axiomatic frameworks in [167, 169] do not impose special restrictions (such as symmetry) on $\sigma_{\mathcal{Q}}$ and $\sigma_{\mathcal{Q}\times\mathcal{A}}$ which might appear natural when interpreting the latter as similarity measures.

The indexing of a weight $\omega_{p_0,a_0,\mathcal{M}}(p,a)$ in (7.5) suggests that the relevance of a case (p,a,r) is not necessarily a function of (p,a) and (p_0,a_0) alone but might also depend on other cases in the memory $\mathcal{M}$. An example of this type of "context-sensitive" relevance will be presented in the next section.

7.2 Nearest Neighbor decisions

Interestingly enough, the modification (7.3) of decision criterion (7.2) corresponds to a special version of a k-Nearest Neighbor approximation, namely Shephard's interpolation method which makes use of the complete set of observations [340]. It is used for making predictions in other CBI approaches as well (e.g., in the ELEM2-CBR system [61]). Indeed, case-based decision making can basically be seen as a special type of CBI or, more specifically, of case-based inference as discussed in previous chapters: Evaluating the act a_0 comes down to estimating the associated *outcome* $r(p_0,a_0)$ (resp. the utility thereof) when viewing a problem–act tuple (p_0,a_0) as an *input* in the sense of CBI. In this sense, a single decision problem gives rise to several CBI problems since a corresponding estimation has to be derived for all acts $a \in \mathcal{A}$. Of course, the estimation of an outcome can principally be realized by any method of instance-based prediction.[7] In particular, one might think of replacing (7.2) by the NN rule in its basic form, an idea that we shall discuss below.

7.2.1 Nearest Neighbor classification and decision making

Recall the problem–act similarity model (7.4) and let $\sigma_{\mathcal{S}} = \sigma_{\mathcal{Q}\times\mathcal{A}}$ denote a similarity measure over the set of inputs which now corresponds to the set $\mathcal{Q}\times\mathcal{A}$ of problem–act tuples. Moreover, let $\mathcal{M}^{\downarrow}$ be the projection of the memory $\mathcal{M}$ to $\mathcal{Q}\times\mathcal{A}$. The NN-based counterpart to the evaluation (7.4) of an act $a_0 \in \mathcal{A}$ is then given by

$$V(a_0) = u(r(\mathsf{NN}_{\mathcal{M}}(p_0,a_0))), \tag{7.6}$$

where $\mathsf{NN}_{\mathcal{M}}(p_0,a_0)$ is the nearest neighbor of the problem–act tuple (p_0,a_0):

$$\mathsf{NN}_{\mathcal{M}}(p_0,a_0) = \arg\max_{(p,a)\in\mathcal{M}^{\downarrow}} \sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)). \tag{7.7}$$

---

[6] Gilboa and Schmeidler fully agree in this point. See [71] for a related discussion and [36] for an application of CBDT where the notion of "relevance" might be preferred to that of "similarity."

[7] In fact, other machine learning methods could be used as well (cf. Section 7.7).

Of course, definition (7.7) should be refined in order to handle the non-uniqueness of the nearest neighbor. However, for the sake of simplicity we assume that each problem–act tuple (p_0, a_0) has a unique nearest neighbor in $\mathcal{M}^{\downarrow}$ (according to the similarity $\sigma_{\mathcal{Q}\times\mathcal{A}}$).

Observe that the CBDT criteria (7.2) and (7.4) use all cases in order to evaluate an act. As opposed to this, the decision maker concentrates completely on the most relevant experience when evaluating an act according to (7.6). More precisely, (7.6) corresponds to (7.5) with the relevance given by

$$\omega_{p_0,a_0,\mathcal{M}}(p,a) = \begin{cases} 1 & \text{if } (p,a) = \mathsf{NN}_{\mathcal{M}}(p_0,a_0) \\ 0 & \text{otherwise} \end{cases} .$$

On the one hand, some information is clearly lost by reducing the number of cases taken into account.[8] On the other hand, the nearest neighbor does generally provide the most relevant information, i.e., the loss of information is limited.[9] Moreover, (7.6) can be seen as an approximation of (7.4) which appears reasonable from a computational point of view. Indeed, since the retrieving of all previous cases might be very time consuming, a decision maker will generally not fall back on its entire experience when having to perform a prompt action. Besides, (7.6) might appear more natural in some situations since it avoids the accumulation and compensation effect produced by (7.2) and (7.4) (cf. Section 7.1). Particularly, the estimation (7.6) corresponds to the true utility if a_0 has already been applied to p_0 in the past (which means that $(p_0, a_0) \in \mathcal{M}^{\downarrow}$). The addition of further (weighted) utility degrees or any kind of averaging might then be counterproductive (cf. Section 7.6).

Note that the NN-decision rule (7.6) partitions the set $\mathcal{A}$ into equivalence classes $[a]$, where

$$b \in [a] \;\Leftrightarrow\; a \sim b \;\Leftrightarrow\; \mathsf{NN}_{\mathcal{M}}(p_0, a) = \mathsf{NN}_{\mathcal{M}}(p_0, b).$$

In fact, two acts a and b are rated equally in the sense of (7.6) whenever $a \sim b$, i.e., as soon as both acts have the same nearest neighbor (in connection with a problem p_0). The criterion (7.6) hence ignores the actual degrees of similarity, a problem already mentioned in connection with the comparison of instance-based and kernel-based extrapolation of case-based information (cf. Section 5.3.5). This, however, does not appear reasonable from a decision making point of view. Consider, for instance, a case (p, a, r) with high utility $u(r)$. Moreover, let b and c be acts such that $\sigma_{\mathcal{Q}\times\mathcal{A}}((p, a), (p_0, b))$ is large and $\sigma_{\mathcal{Q}\times\mathcal{A}}((p, a), (p_0, c))$ is small. Still, assume that $\mathsf{NN}_{\mathcal{M}}(p_0, b) = \mathsf{NN}_{\mathcal{M}}(p_0, c) = (p, a)$. In this situation, a risk-averse decision maker will generally prefer b to c. The criterion (7.6), however, does not differentiate between these two acts. The NN principle (as any other estimation method) seems hence questionable in the context of decision making.

[8] Though formally only the relevance of some cases is set to 0.

[9] This claim can be proved in a formal way. The result in [74], for instance, can be interpreted as follows: Under certain technical assumptions, at least *half* of the information that a complete random sample contains about an outcome in already represented by the nearest neighbor of the query instance.

Indeed, at this point one should realize an important difference between decision making and *prediction*, the performance task which is commonly solved by NN algorithms: In a prediction problem, an estimation has to be derived for only *one* instance and this estimation is not considered as a valuation which supports any kind of comparison. Having to choose one among the potential candidates anyway, it might then be acceptable to base an estimation on the nearest neighbor even if it turns out to be quite dissimilar.

Let us mention that the averaged similarity criterion (7.3) suffers from a similar problem. In fact, it is readily seen that the valuation of an act according to (7.3) can be very large even though this act has only been applied in situations which are hardly similar to the current problem.

7.2.2 Nearest Neighbor decision rules

In order to overcome the aforementioned problem it seems natural to not only associate the utility v of the nearest neighbor $(p, a) \in \mathcal{M}^{\downarrow}$ with each act $a_0 \in \mathcal{A}$ (i.e., with the tuple (p_0, a_0)), but rather the tuple (v, σ), where σ denotes the similarity between (p_0, a_0) and (p, a). The preferences of an agent should then be expressed in terms of a preference relation over the class of such utility–similarity tuples. This is somewhat comparable to generalized decision rules which take not only the expected utility into account but also the variance (i.e. uncertainty) related to an act.

More specifically, one might think of the following generalization of (7.6):

$$V(a_0) = \sigma_{\mathcal{Q}\times\mathcal{A}}\left((p_0, a_0), \mathsf{NN}_{\mathcal{M}}(p_0, a_0)\right) \cdot u(r(\mathsf{NN}_{\mathcal{M}}(p_0, a_0)))\,. \tag{7.8}$$

This valuation, which represents a preference relation over the set of tuples (v, σ) by means of

$$(v, \sigma) \preceq (v', \sigma') \Leftrightarrow v \cdot \sigma \leq v' \cdot \sigma',$$

combines (7.4) and (7.6) to some extent. Again, it considers only one previous case (namely the nearest neighbor) rather than all cases when evaluating an act. The corresponding utility, however, is now weighted by the degree of similarity. In fact, (7.8) can be seen as a special version of (7.4) when interpreting $\sigma_{\mathcal{Q}\times\mathcal{A}}$ as a measure of relevance (cf. Section 7.1), which is then given by

$$\omega_{p_0,a_0,\mathcal{M}}(p, a) = \begin{cases} \sigma_{\mathcal{Q}\times\mathcal{A}}((p, a), (p_0, a_0)) & \text{if } (p, a) = \mathsf{NN}_{\mathcal{M}}(p_0, a_0) \\ 0 & \text{otherwise} \end{cases}\,. \tag{7.9}$$

According to (7.9), only the nearest neighbor is considered as a relevant observation. Of course, this idea might be generalized by taking the $k \geq 1$ nearest neighbors into account, or by introducing a threshold such that the relevance of an observation is set to 0 in case its similarity is too small.

The valuation (7.8) defines a reasonable tradeoff between the *goodness* (in terms of utility) and the *relevance* (in terms of similarity) of an experience. Still, it deserves mentioning that the degree of similarity is nothing else than a heuristic indication of the actual degree of uncertainty of an NN estimation. In fact, it is not true in general that a larger similarity comes along with a higher precision of an estimation.

It has already been mentioned that a reduction of observations as realized by (7.8) might be reasonable from a computational point of view. Particularly, this is true if the decision maker has a large memory of cases but a relatively small number of acts (and if it disposes of an efficient method of case retrieval). In the reverse case where the memory is small and the set of acts to be evaluated is large, a different strategy which passes through the set of cases, $\mathcal{M}$, rather than the set of acts, $\mathcal{A}$, might be preferred: Instead of considering the most relevant observation for each act one can proceed from an observation and attach the related experience to the most relevant act. This idea is realized by the following counterpart to (7.8):

$$V(a_0) = \sum_{(p,a)\in\mathcal{M}^{\downarrow}:\mathsf{NN}_{p_0,\mathcal{A}}(p,a)=a_0} \sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)) \cdot u(r(p,a)). \tag{7.10}$$

Here,

$$\mathsf{NN}_{p_0,\mathcal{A}}(p,a) = \arg\max_{a_0\in\mathcal{A}} \sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)) \tag{7.11}$$

denotes the problem–act tuple $(p_0,a_0) \in \{p_0\} \times \mathcal{A}$ which is maximally similar to the observation $(p,a) \in \mathcal{M}^{\downarrow}$. We assume (7.11) to be unique whenever some $a_0 \in \mathcal{A}$ exists such that $\sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)) > 0$; otherwise we let $\mathsf{NN}_{p_0,\mathcal{A}}(p,a) = \emptyset$ by definition.

7.2.3 An axiomatic characterization

In [169], an axiomatization of (7.4) is proposed which assumes a preference relation $\succeq_x \subset \mathcal{A} \times \mathcal{A}$ over the set of acts to be given. As suggested by the attached index, this preference relation depends on the experience of the decision maker: x defines a $\mathcal{M}^{\downarrow} \longrightarrow \mathfrak{R}^n$ function which assigns utility degrees to problem–act pairs. It can simply be thought of as the vector

$$x = (x_1, \ldots, x_n) = \big(u(r(p_1,a_1)), \ldots, u(r(p_n,a_n))\big)\,,$$

where $x_\imath = u(r(p_\imath, a_\imath)) \in \mathfrak{R}$ corresponds to the utility obtained in connection with the $\imath$-th problem–act tuple $(p_\imath, a_\imath)$. The vector x represents the history of the decision maker and determines the *context* of the new decision problem. The information available to a decision maker which has to evaluate an act $a \in \mathcal{A}$ might thus be illustrated in the form of a table as follows:

utility	x_1	x_2	$\ldots$	x_n
similarity	$\sigma_1(a)$	$\sigma_2(a)$	$\ldots$	$\sigma_n(a)$

(7.12)

The (case-based) rating of a will then be a function of the values in this table, namely the degrees of utility obtained so far and the similarities

$$\sigma_\imath(a) = \sigma_{\mathcal{Q}\times\mathcal{A}}((p_0, a), (p_\imath, a_\imath))$$

between the already encountered problem–act tuples and the new tuple (p_0, a). The criterion (7.4), for instance, is given by the weighted sum

$$V(a) = \sum_{\imath=1}^{n} \sigma_\imath(a) x_\imath .$$

Clearly, this criterion and table (7.12) remind one of expected utility theory. In fact, the context x plays formally the role of the probability distribution on the set of states of nature, and the degrees of similarity $\sigma_\imath$ correspond to degrees of utility in EUT.

For the NN-rules (7.8) and (7.10) we can show representation theorems similar to the one obtained in [169]. Consider the following axioms, which are basically formulated in terms of *contexts*[10] ($\succ_x$ and $\simeq_x$ denote the asymmetric and symmetric part of $\succeq_x$, respectively):

A1 Order: $\succeq_x$ is complete and transitive for all $x \in \mathfrak{R}^n$.

A2 Continuity: For all $(x^k)_{k\geq 1} \subset \mathfrak{R}^n$ and all $a, b \in \mathcal{A}$ it holds true that

$$\left(x^k \rightarrow x \ \wedge \ \forall\, k \geq 1 \ : \ a \succeq_{x^k} b\right) \quad \Rightarrow \quad a \succeq_x b.$$

A3 Additivity: For all $x, y \in \mathfrak{R}^n$ and $a, b \in \mathcal{A}$ it holds true that

$$a \succ_x b \wedge a \succeq_y b \ \Rightarrow \ a \succ_{x+y} b.$$

A4 Neutrality: For all $a, b \in \mathcal{A}$ it holds true that $a \simeq_{(0,\ldots,0)} b$.

A5 Diversity: For all distinct acts $a, b, c, d \in \mathcal{A}$ a vector $x \in \mathfrak{R}^n$ exists such that

$$a \succ_x b \succ_x c \succ_x d.$$

The following result has been shown in [169]: A1–A5 imply the existence of vectors $\omega(a) = (\omega_1(a), \ldots, \omega_n(a))$ for all $a \in \mathcal{A}$ such that

$$a \succeq_x b \ \Leftrightarrow \ \sum_{\imath=1}^{n} \omega_\imath(a) \cdot x_\imath \geq \sum_{\imath=1}^{n} \omega_\imath(b) \cdot x_\imath , \tag{7.13}$$

where the $x_\imath$ are the utility degrees in (7.12). Moreover, the vectors $\omega(a)$ are unique up to an affine transformation. Of course, the weights $\omega_\imath(a)$ can be interpreted as the similarity degrees $\sigma_\imath(a)$ in (7.12).

---

[10] This contrasts with classical decision-theoretic models which are formalized in terms of acts (in the SAVAGE setting) or probabilistic lotteries (in the VON NEUMANN-MORGENSTERN framework).

The valuations (7.8) and (7.10) are obviously special cases of the weighted sum in (7.13). In order to obtain a set of axioms which imply a nearest neighbor representation it is hence possible to extend A1–A5 in such a way that some of the weights $\omega_\imath$ become 0. Consider the following axiom (the k-th entry of the vector e_k is 1 and all other entries are 0):

A6 For all acts $a, b, c \in \mathcal{A}$, $x \in \mathfrak{R}^n$, $\gamma \geq 0$ and $1 \leq k \leq n$ it holds true that

$$c \succ_x a \ \wedge \ c \succ_x b \ \Rightarrow \ c \succ_{x+\gamma e_k} a \ \vee \ c \succ_{x+\gamma e_k} b. \tag{7.14}$$

In a certain sense, the meaning of A6 is opposite to that of axiom A5. The latter demands that a set of acts can be put in any order by defining the context appropriately. As opposed to this, A6 demands that a certain modification of the context, namely the increase of one utility degree x_k, can only have a limited influence: It can reverse but one of the preferences in the antecedent part of implication (7.14).

Lemma 7.1. Suppose A1-A6 to hold and let $1 \leq k \leq n$. The vector

$$\lambda = (\lambda_1, \ldots, \lambda_m) = (\omega_k(a_1), \ldots, \omega_k(a_m)),$$

where $m = \mathrm{card}(\mathcal{A}) \geq 4$, is of the form

$$\lambda = \alpha e_{\imath_0} + \beta \tag{7.15}$$

for some $1 \leq \imath_0 \leq m$, $\alpha \geq 0$ and $\beta \in \mathfrak{R}$. □

Proof. Consider a permutation π of $\{1, \ldots, m\}$ such that

$$\lambda_{\pi(1)} \geq \lambda_{\pi(2)} \geq \ldots \geq \lambda_{\pi(m)}. \tag{7.16}$$

We obviously have $\lambda_{\pi(2)} = \ldots = \lambda_{\pi(m)}$ if (7.15) holds. Suppose by way of negation that

$$\lambda_{\pi(2)} \geq \ldots \geq \lambda_{\pi(\jmath-1)} > \lambda_{\pi(\jmath)} \geq \ldots \geq \lambda_{\pi(m)}.$$

Axiom A5 guarantees the existence of $x \in \mathfrak{R}^n$ such that $a_{\pi(\jmath)} \succ_x a_{\pi(1)}$ and $a_{\pi(\jmath)} \succ_x a_{\pi(2)}$. Since $\omega_k(a_{\pi(\jmath)}) = \lambda_{\pi(\jmath)} < \lambda_{\pi(1)} = \omega_k(a_{\pi(1)})$ and $\omega_k(a_{\pi(\jmath)}) = \lambda_{\pi(\jmath)} < \lambda_{\pi(2)} = \omega_k(a_{\pi(2)})$, there is obviously some $\gamma > 0$ such that

$$\sum_{\imath=1}^{n} \omega_\imath(a_{\pi(\imath_0)}) \cdot (x_\imath + \gamma e_k) > \sum_{\imath=1}^{n} \omega_\imath(a_{\pi(\jmath)}) \cdot (x_\imath + \gamma e_k)$$

for $\imath_0 = 1$ and $\imath_0 = 2$. This means $a_{\pi(1)} \succ_{x+\gamma e_k} a_{\pi(\jmath)}$ and $a_{\pi(2)} \succ_{x+\gamma e_k} a_{\pi(\jmath)}$ according to (7.13) and, hence, contradicts A6. Consequently, the representation (7.15) must hold with $\imath_0 = \pi(1)$, $\alpha = \lambda_{\pi(1)} - \lambda_{\pi(2)}$ and $\beta = \lambda_{\pi(2)}$. □

Theorem 7.2. Consider a decision problem with $\text{card}(\mathcal{A}) \geq 4$. The preference relations $\succeq_x$ can be represented by (7.10) iff they satisfy A1–A6. □

Proof. It is readily verified that the preference relations $\succeq_x$ defined by (7.10) satisfy A1–A6. Concerning the converse direction, we make use of Lemma 7.1 and the fact that β in (7.15) can be set to 0 without loss of generality. In fact, the variation of β does not influence the relation on the right-hand side of (7.13). Thus, for each $1 \leq k \leq n$ there is at most one $1 \leq \imath_0 \leq m$ such that $\omega_k(a_{\imath_0}) \neq 0$. We hence obtain the representation (7.10) by letting

$$\mathsf{NN}_{p_0,\mathcal{A}}(p_\imath, a_\imath) = \{a \in \mathcal{A} \,|\, \omega_\imath(a) > 0\}$$

for all $(p_\imath, a_\imath) \in \mathcal{M}^\downarrow$. □

Note that a value $\omega_\imath(a) > 0$ is interpreted as the similarity between $(p_\imath, a_\imath)$ and $(p_0, a) = \mathsf{NN}_{p_0,\mathcal{A}}(p_\imath, a_\imath)$. It hence corresponds to the value $\sigma_\imath(a)$ in (7.12). It is clear, however, that the complete similarity relation $\sigma_{\mathcal{Q}\times\mathcal{A}}$ cannot be determined by the preferences $\succeq_x$. In fact, $\omega_\imath(b) = 0$ does not necessarily mean that $\sigma_{\mathcal{Q}\times\mathcal{A}}((p_\imath, a_\imath), (p_0, b)) = 0$ but only implies $\sigma_{\mathcal{Q}\times\mathcal{A}}((p_\imath, a_\imath), (p_0, b)) < \omega_\imath(a)$. This is caused by the behavior of a decision maker applying the NN principle. According to (7.9) it concentrates on the nearest neighbors of the observations but completely ignores other acts to which it assigns a relevance of 0. Thus, the preferences $\succeq_x$ can determine only the *relevance* of a case but not its *similarity* to (p_0, a_0).

Now, consider again the decision rule (7.8). Axiom A5 is obviously not satisfied in connection with this criterion. Indeed, we have

$$b \succeq_x c \succeq_x d \quad \text{or} \quad d \succeq_x c \succeq_x b$$

for all $x \in \Re^n$ if the acts $b, c, d \in \mathcal{A}$ have the same nearest neighbor (p, a) and if

$$\sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,b)) < \sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,c)) < \sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,d)).$$

Observe, however, that the act c will then be ignored by the decision maker in the sense that it is not chosen anyway (except perhaps if $V(b) = V(c) = V(d) = 0$). Besides, act b becomes interesting only if all acts have a negative (estimated) utility according to (7.8), a situation that can formally be avoided (see below). We can hence restrict the decision rule (7.8) to a set $\mathcal{A}_{p_0}$ of acts as follows: For the problem–act tuple $(p_\imath, a_\imath) \in \mathcal{M}^\downarrow$ define

$$A_{p_0}(p_\imath, a_\imath) = \arg \max_{a \in \mathcal{A} \,:\, \mathsf{NN}_{\mathcal{M}}(p_0, a) = (p_\imath, a_\imath)} \sigma_{\mathcal{Q}\times\mathcal{A}}((p_0, a), (p_\imath, a_\imath))$$

whenever the set on the right-hand side is not empty. For the sake of simplicity, we again assume $A_{p_0}(p_\imath, a_\imath)$ to be unique. The set $\mathcal{A}_{p_0}$ is then defined as

$$\mathcal{A}_{p_0} = \{A_{p_0}(p_\imath, a_\imath) \,|\, 1 \leq \imath \leq n, A_{p_0}(p_\imath, a_\imath) \text{ exists}\}. \tag{7.17}$$

It can be thought of as the set of *relevant* acts. As already suggested above, a decision based on (7.8) might appear somewhat peculiar if $V(a_0) < 0$ for all $a_0 \in \mathcal{A}$. Observe, however, that this problem can formally be avoided by means of a proper definition of acts. For instance, one might introduce a new act a^* which stands for "doing anything" or "trying something completely new." When adding a dummy case $(p^*, a^*, 0)$ to $\mathcal{M}$, $V(a^*) = 0$ is guaranteed by letting $\sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p^*,a^*)) = 1$ if $(p,a) = (p^*,a^*)$ and 0 otherwise. Thus, a^* is preferred to each act with a negative estimated utility. This is clearly in accordance with the idea of an aspiration level in [167].

Note that each act $a \in \mathcal{A}_{p_0}$ in (7.17) has a unique nearest neighbor in $\mathcal{M}^{\downarrow}$. Moreover, for each $(p_\imath, a_\imath) \in \mathcal{M}^{\downarrow}$ there is at most one act $a \in \mathcal{A}$ such that $(p_\imath, a_\imath)$ is the nearest neighbor of (p_0, a) (which implies $\text{card}(\mathcal{A}_{p_0}) \leq \text{card}(\mathcal{M})$). It is hence obvious that A5 is satisfied for $\mathcal{A}_{p_0}$. Besides, it is not difficult to show that the preference relations induced by (7.8) also satisfy the following axiom:

A7 For all $x, y \in \mathfrak{R}^n$ and $a, b \in \mathcal{A}$ it holds true that

$$a \succeq_x b \wedge a \succeq_y b \;\Rightarrow\; a \succeq_{\max\{x,y\}} b,$$

where the maximum of the vectors x and y is defined component-wise.

Theorem 7.3. Consider a decision problem with $\text{card}(\mathcal{A}) \geq 4$. The preference relations $\succeq_x$ can be represented by (7.8) iff they satisfy A1–A7. □

Proof. Again, A1–A7 are obviously satisfied when representing $\succeq_x$ by (7.8). In order to show the converse direction suppose A1–A7 to be satisfied. Given A1–A6, is has been shown in Theorem 7.2 that (7.13) holds in such a way that $\omega_\imath(a)\omega_\imath(b) = 0$ for all acts $a \neq b$. In order to establish a representation of $\succeq_x$ by (7.8), we further have to show that $\imath \neq \jmath \Rightarrow \omega_\imath(a)\omega_\jmath(a) = 0$ for all acts a. Thus, assume the existence of an act a such that $\omega_\imath(a) > 0$ and $\omega_\jmath(a) > 0$ for $1 \leq \imath \neq \jmath \leq n$. Moreover, let the contexts x and y be defined as follows:

$$x_k = \begin{cases} \omega_\jmath(a) & \text{if} \quad k = \imath \\ -\omega_\imath(a) & \text{if} \quad k = \jmath \\ 0 & \text{if} \quad \imath \neq k \neq \jmath \end{cases}, \quad y_k = \begin{cases} -\omega_\jmath(a) & \text{if} \quad k = \imath \\ \omega_\imath(a) & \text{if} \quad k = \jmath \\ 0 & \text{if} \quad \imath \neq k \neq \jmath \end{cases}.$$

It is readily verified that $V(a) = 0$ in both contexts. Moreover, $V(b) = 0$ does also hold true for all other acts since $b \neq a$ entails $\omega_\imath(b) = \omega_\jmath(b) = 0$. Thus, $b \succeq_x a$ and $b \succeq_y a$ for any act $b \neq a$. In the context $\max\{x,y\}$, however, we have $V(a) = 2\omega_\imath(a)\omega_\jmath(a) > 0$ and, hence, $a \succ_{\max\{x,y\}} b$. This contradicts A7. □

7.3 Fuzzy modeling of case-based decisions

Case-based decision making has been realized in [101] as a kind of case-based approximate reasoning. This approach is in line with methods of qualitative decision theory. In fact, the assumption that uncertainty and preference can be quantified by means of, respectively, a precise probability measure and a cardinal utility function (as it is assumed in classical decision theory) does often appear unrealistic. As opposed to (7.2), the approach discussed in this section only assumes an ordinal setting for modeling decision problems, i.e., ordinal scales for assessing preference and similarity. This interpretation should be kept in mind, especially since both scales will subsequently be taken as (subsets of) the unit interval.

7.3.1 Basic measures for act evaluation

Let $\rightsquigarrow$ be a multiple-valued implication connective. Given a memory $\mathcal{M}$ and a new problem p_0, the following (estimated) utility value is assigned to an act $a \in \mathcal{A}$:

$$V^{\downarrow}_{p_0,\mathcal{M}}(a) \stackrel{\text{df}}{=} \min_{(p,a,r)\in\mathcal{M}} \sigma_{\mathcal{Q}}(p,p_0) \rightsquigarrow u(r). \tag{7.18}$$

This valuation supports the idea of finding an act a which has *always* resulted in good outcomes for problems similar to the current problem p_0. Indeed, (7.18) can be considered as a generalized truth degree of the claim that "whenever a has been applied to a problem p similar to p_0, the corresponding outcome has yield a high utility." An essential idea behind (7.18) is that of avoiding the accumulation and compensation effect caused by the decision criterion (7.2) (cf. Section 7.1),[11] since these effects do not always seem appropriate (cf. Section 7.6).

As a special realization of (7.18) the valuation

$$V^{\downarrow}_{p_0,\mathcal{M}}(a) \stackrel{\text{df}}{=} \min_{(p,a,r)\in\mathcal{M}} \max\{n(h(\sigma_{\mathcal{Q}}(p,p_0))), u(r)\},$$

is proposed, where h is an order-preserving function which assures the linear scales of similarity and preference to be commensurable and n is the order-reversing function of the similarity scale. By taking n as $x \mapsto 1-x$ in $[0,1]$ and h as the identity, we obtain

$$V^{\downarrow}_{p_0,\mathcal{M}}(a) \stackrel{\text{df}}{=} \min_{(p,a,r)\in\mathcal{M}} \max\{1-\sigma_{\mathcal{Q}}(p,p_0), u(r)\}. \tag{7.19}$$

This criterion can obviously be seen as a qualitative counterpart to (7.2). Besides, the criterion

$$V^{\uparrow}_{p_0,\mathcal{M}}(a) \stackrel{\text{df}}{=} \max_{(p,a,r)\in\mathcal{M}} \min\{\sigma_{\mathcal{Q}}(p,p_0), u(r)\} \tag{7.20}$$

is introduced as an *optimistic* counterpart to (7.19). It can be seen as a formalization of the idea to find an act a for which there is at least one problem which is similar to p_0 and for which a has led to a good result. Again, let us mention that expressions (7.19) and (7.20) are closely related to decision criteria which

[11] Note that the accumulation effect is also the main motivation for the normalization (7.3).

have recently been derived in [123] in connection with an axiomatic approach to qualitative decision making under uncertainty.

In the more general context of problem–act similarity, the decision rules (7.19) and (7.20) become

$$V^{\downarrow}_{p_0,\mathcal{M}}(a_0) \stackrel{\text{df}}{=} \min_{(p,a,r)\in\mathcal{M}} \max\{1 - \sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)), u(r)\}, \tag{7.21}$$

$$V^{\uparrow}_{p_0,\mathcal{M}}(a_0) \stackrel{\text{df}}{=} \max_{(p,a,r)\in\mathcal{M}} \min\{\sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)), u(r)\}. \tag{7.22}$$

In order to make the basic principles underlying the above criteria especially obvious, suppose the qualitative utility scale to be given by $U = \{0, 1\}$. That is, only a crude distinction between "bad" and "good" outcomes is made. (7.21) and (7.22) can then be simplified as follows:

$$V^{\downarrow}_{p_0,\mathcal{M}}(a_0) \stackrel{\text{df}}{=} 1 - \max_{(p,a,r)\in\mathcal{M}\,:\,u(r)=0} \sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)), \tag{7.23}$$

$$V^{\uparrow}_{p_0,\mathcal{M}}(a_0) \stackrel{\text{df}}{=} \max_{(p,a,r)\in\mathcal{M}\,:\,u(r)=1} \sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)). \tag{7.24}$$

According to (7.23), the decision maker only takes cases (p, a, r) with bad outcomes into account. An act a_0 is discounted whenever (p_0, a_0) is similar to a corresponding problem–act tuple (p, a). Thus, the agent is cautious and looks for an act that it does not associate with a bad experience. According to (7.24), it only considers the cases with good outcomes. An act a_0 appears promising if (p_0, a_0) is similar to a tuple (p, a) which has yielded a good result. In other words, the decision maker is more adventurous and looks for an act that it associates with a good experience.

7.3.2 Modification of the basic measures

As noted in [125], (7.19) makes only sense if the memory contains at least one problem p such that $\sigma_{\mathcal{Q}}(p, p_0) = 1$ and a has been chosen for solving p. Otherwise, it may happen that (7.19) is very high even though none of the problems contained in the memory is similar to the current problem p_0.[12] Particularly,

$$\Big(\big\{p \in \mathcal{Q} \,|\, (p,a,r) \in \mathcal{M} \,\wedge\, \sigma_{\mathcal{Q}}(p,p_0) > 0\big\} = \emptyset\Big) \;\Rightarrow\; \Big(V^{\downarrow}_{p_0,\mathcal{M}}(a) = 1\Big),$$

which does not seem satisfactory.

Modifications of (7.19) and its optimistic counterpart have been proposed in order to cope with these difficulties. The modified version of (7.19) is based on some kind of *normalization* of the similarity function for each act a and a discounting

[12] Notice that the averaged similarity criterion (7.3) suffers from the same drawback.

which takes the absence of problems similar to p_0 into account. More precisely, the modified measure is given by

$$V^{\downarrow}_{p_0,\mathcal{M}}(a) = \min\left\{ h_{\mathcal{M}}(a,p_0), \min_{(p,a,r)\in\mathcal{M}} \max\{1-\sigma^a_{\mathcal{Q}}(p,p_0), u(r)\}\right\}, \qquad (7.25)$$

where

$$h_{\mathcal{M}}(a,p_0) = \max_{(p,a,r)\in\mathcal{M}} \sigma_{\mathcal{Q}}(p,p_0),$$

and $\sigma^a_{\mathcal{Q}}(\cdot,p_0)$ denotes a normalization[13] of $\sigma_{\mathcal{Q}}(\cdot,p_0)$, e.g.,

$$\sigma^a_{\mathcal{Q}}(p,p_0) = \begin{cases} 1 & \text{if} \quad \sigma_{\mathcal{Q}}(p,p_0) = h_{\mathcal{M}}(a,p_0) \\ \sigma_{\mathcal{Q}}(p,p_0) & \text{if} \quad \sigma_{\mathcal{Q}}(p,p_0) < h_{\mathcal{M}}(a,p_0) \end{cases}.$$

The idea behind (7.25) is that the willingness of a decision maker to choose act a is upper-bounded by the existence of problems which are completely similar to p_0 and to which a has been applied. Moreover, $\sigma_{\mathcal{Q}}(\cdot,p_0)$ is normalized in order to obtain a meaningful degree of inclusion. Thus, (7.25) corresponds to the compound condition that there are problems similar to p_0 to which act a has been applied and the problems which are most similar to p_0 are among the problems for which a has led to good results. Observe that (7.19) is retrieved from (7.25) as soon as $h_{\mathcal{M}}(a,p_0) = 1$. Moreover, note that a corresponding modification can also be defined for (7.20):

$$V^{\uparrow}_{p_0,\mathcal{M}}(a) = \max\left\{ 1-h_{\mathcal{M}}(a,p_0), \max_{(p,a,r)\in\mathcal{M}} \min\{\sigma^a_{\mathcal{Q}}(p,p_0), u(r)\}\right\}. \qquad (7.26)$$

The criteria (7.25) and (7.26) guarantee that $V^{\downarrow}_{p_0,\mathcal{M}}(a) \leq V^{\uparrow}_{p_0,\mathcal{M}}(a)$ which is not necessarily the case for (7.19) and (7.20).

7.3.3 Interpretation of the decision criteria

As opposed to (7.2), the criteria (7.19) and (7.20) do obviously not focus on some kind of average performance, which hardly makes sense within an ordinal setting. Rather, they should be considered from the same point of view as qualitative decision rules such as Maximin [123]. Indeed, the application of (7.18) seems reasonable, for instance, if an agent aims at minimizing the occurrence of worst case outcomes in competition with other agents or if only an ordinal preference relation on outcomes is assumed [52].

We shall now propose two interpretations of (7.19).[14] The first one is that of an approximation of a (generalized) Maximin evaluation: Observe that we can write (7.19) as

[13] Note that this normalization is again defined for each act individually.

[14] These interpretations can be transferred to (7.20) in a straightforward way.

$$V^{\downarrow}_{p_0,\mathcal{M}}(a) = \min_{0 \le k \le m} \max\{1 - \sigma_k, v_k\}, \tag{7.27}$$

where the values $0 = \sigma_0 < \sigma_1 < \ldots < \sigma_m = 1$ constitute the (finite) set $\{\sigma_{\mathcal{Q}}(p, p') \,|\, p, p' \in \mathcal{Q}\}$ of possible similarity degrees of problems and

$$v_k = \min V_k = \min\{u(r) \,|\, (p, a, r) \in \mathcal{M}, \sigma_{\mathcal{Q}}(p, p_0) = \sigma_k\}$$

is the lowest utility obtained in connection with act a for problems which are σ_k-similar to p_0. Moreover, $v_k = 1$ by definition if $V_k = \emptyset$ (which is just the reason for the problem that (7.19) becomes large if no similar observations have been made).

According to (7.27), the valuation (7.19) of an act is completely determined by the lower bounds v_k ($0 \le k \le m$) which are derived from the memory $\mathcal{M}$ (and discounted according to respective degrees of similarity). This reveals that (7.19) can indeed be seen as some kind of "experience-based" approximation of the MAXIMIN principle. The case in which all problems are completely similar makes this especially apparent. Then, (7.19) evaluates an act a simply according to the worst consequence observed so far. More generally, the value v_k can be seen as an estimation of the lower utility bound

$$w_k = \min\{u(r(p, a)) \,|\, p_0 \ne p \in \mathcal{Q}, \sigma_{\mathcal{Q}}(p, p_0) = \sigma_k\},$$

i.e., the smallest degree of utility which can be obtained in connection with act a for (not necessarily encountered) problems from $\mathcal{Q}$ which are σ_k-similar to p_0. Then, $V^{\downarrow}_{p_0,\mathcal{M}}(a)$ can be interpreted as an approximation of

$$W^{\downarrow}_{p_0}(a) = \min_{0 \le k \le m} \max\{1 - \sigma_k, w_k\},$$

which defines a case-based generalization of a MAXIMIN-evaluation. In fact, $W^{\downarrow}_{p_0}(a)$ is equal to $V^{\downarrow}_{p_0,\mathcal{M}}(a)$ if a has already been applied to all problems (up to p_0) from $\mathcal{Q}$, i.e., if $\{p \,|\, \exists r \in \mathcal{R} : (p, a, r) \in \mathcal{M}\} = \mathcal{Q} \setminus \{p_0\}$.

According to a second (more logic-oriented) interpretation, (7.18) might be seen as the (generalized) truth degree of a proposition characterizing the decision maker's preferences concerning acts. In our case, those acts are preferred which have always resulted in good outcomes for similar problems. Then, (7.18) defines the degree to which an act meets the requirements and, hence, induces a corresponding preference relation over acts. In a certain sense, this approach can be seen as a "compiled" decision model which skips the estimation of utility and relates similarity or, more generally, certain properties of an act to preference more directly. That is to say, the agent already knows which properties a preferred act should have. The idea of such a compiled model becomes even more obvious if we consider (crisp) rules of the form "if the problem has property x then choose an act with property y", such as "if it looks rainy then take an umbrella with you." Rules of this kind are often set up if a decision problem is solved frequently. They represent a sort of routine decision and reflect the agent's knowledge that a decision analysis, i.e., the estimation of utility degrees for all decisions, would result in choosing a certain act if the problem has a related property anyway.

Even though formally equivalent, the two interpretations are different from a semantical point of view. For instance, interpreting the value $V(a)$ assigned to an act a which has not yet been tried as a degree to which this act meets the agent's idea of an "ideal" decision seems less critical than viewing this value as an estimated utility. In fact, the latter is merely a "default utility." As opposed to this, the former interpretation does principally not assume observations at all. Rather, $V(a)$ can be seen as reflecting the agent's attitude toward uncertainty. Assigning a high default value to a then simply means that a not yet applied act seems attractive and, hence, amounts to model an uncertainty-prone decision maker who is willing to try new acts.

7.4 Fuzzy quantification in act evaluation

In some situations, the extremely pessimistic and optimistic nature of the criteria (7.19) and (7.20), respectively, might appear at least as questionable as the accumulation in (7.2). Here we shall propose a generalization of the decision rule (7.19) which is a weakening of the demand that an act has *always* produced good results for similar problems. In fact, one might already be satisfied if a turned out to be a good choice *for most* similar problems, thus allowing for a few exceptions [125]. In other words, the idea is to relax the universal "for all" quantifier. Observe that a similar generalization of (7.20), which replaces "there exists" by "there are at least several" and, hence, corresponds to a strengthening of this decision principle, seems reasonable as well. It can be obtained analogously.

Consider a finite set A of cardinality $m = |A|$. In connection with propositions of the form "most elements of A have property X" the fuzzy quantifier "most" can be formalized by means of a fuzzy set [132, 403],[15] the membership function $\mu : \{0, 1, \ldots, m\} \longrightarrow [0, 1]$ of which satisfies

$$\forall\, 1 \leq k \leq m-1 \,:\, \mu(k) \leq \mu(k+1) \qquad \text{and} \qquad \mu(m) = 1. \tag{7.28}$$

The special case "for all" then corresponds to $\mu(k) = 0$ for $0 \leq k \leq m-1$ and $\mu(m) = 1$. Given some μ satisfying (7.28), we define an associated membership function $\overline{\mu}$ by $\overline{\mu}(0) = 0$ and $\overline{\mu}(k) = 1 - \mu(k-1)$ for $1 \leq k \leq m$ (see e.g. [109]). A membership degree $\overline{\mu}(k)$ can then be interpreted as quantifying the importance that the property X is satisfied for k (out of the m) elements.

Consider a memory $\mathcal{M}$ of cases, a problem $p_0 \in \mathcal{Q}$, an act $a \in \mathcal{A}$, and let $\mathcal{M}_a = \{(p', a', r') \in \mathcal{M} \,|\, a = a'\}$. Moreover, let μ formalize the above-mentioned "for most" concept. A reasonable generalization of (7.19) is then given by

[15] Other possibilities of expressing a fuzzy quantifier exist as well, including the use of order-statistics [300] and an ordered weighted minimum or maximum [135].

$$V_{p_0,\mathcal{M}}(a) = \min_{0 \le k \le |\mathcal{M}_a|} \max\left\{1 - \overline{\mu}(k), \delta_a(k)\right\}, \tag{7.29}$$

where

$$\delta_a(k) = \max_{\mathcal{M}' \subset \mathcal{M}_a : |\mathcal{M}'| = k} \; \min_{(p,a,r) \in \mathcal{M}'} \max\{1 - \sigma_{\mathcal{Q}}(p, p_0), u(r)\}$$

defines the degree to which "the act a has induced good outcomes for similar problems k times." The extent to which a (small) degree $\delta_a(k)$ decreases the overall valuation of a is upper bounded by $1 - \overline{\mu}(k)$, i.e., by the respective level of (un-)importance. Observe that we do not have to consider all subsets $\mathcal{M}' \subset \mathcal{M}_a$ of size k for deriving $\delta_a(k)$. In fact, for computing $V_{p_0,\mathcal{M}}(a)$ it is reasonable to arrange the $m = |\mathcal{M}_a|$ values $v = \max\{1 - \sigma_{\mathcal{Q}}(p, p_0), u(r)\}$ in a non-increasing order $v_1 \ge v_2 \ge \ldots \ge v_m$. Then, (7.29) is equivalent to

$$V_{p_0,\mathcal{M}}(a) = \min_{0 \le k \le |\mathcal{M}_a|} \max\left\{1 - \overline{\mu}(k), v_k\right\},$$

where $v_0 = 1$.

The generalized criterion (7.29) can be useful, e.g., in connection with the idea of repeated decision making which arises quite naturally in connection with a case-based approach to decision making. We might think of different scenarios in which repeated problem solving becomes relevant. A simple model emerges from the assumption that problems are chosen repeatedly from $\mathcal{Q}$ according to some selection process which is not under the control of the agent, such as the repeated (and independent) selection of problems according to some probability measure. More generally, the problem faced next by the agent might depend on the current problem and the act which is chosen for solving it. A Markov Decision Process extended by a similarity measure over states (which correspond to problems) may serve as an example. Besides, we might consider case-based decision making as a reasonable strategy within a (repeated) game playing framework like the iterated prisoner's dilemma [19].

As a concrete example let us consider a very simple model of repeated decision making: Suppose that the agent faces the same problem p repeatedly and that the result associated with an act $a \in \mathcal{A} = \{a_1, a_2, a_3\}$ depends on a state of nature $\omega \in \Omega = \{\omega_1, \omega_2, \omega_3\}$. The state ω is assumed to be chosen randomly (every time) and is not part of the problem description. We assume the probability for $\omega = \omega_3$, which is also not known to the decision maker, to be positive but relatively small. Moreover, the results (= utilities) associated with act–state tuples shall be specified as follows:

	ω_1	ω_2	ω_3
a_1	1	1	0
a_2	1	0	0
a_3	0	0	0

Recall that 0 and 1 are interpreted as ordinal degrees of utility; they only indicate that one outcome is preferred to the other one, which might be encoded by -1 and 1 as well.[16]

Now, since a_1 dominates a_2 and a_3 (strictly), it is obviously the best choice. Observe, however, that the valuation of an act a according to (7.19) simply corresponds to the worst outcome observed in connection with this act, i.e.

$$V^{\downarrow}_{p_0,\mathcal{M}}(a) = \begin{cases} 0 & \text{if} \quad (p,a,0) \in \mathcal{M} \\ 1 & \text{if} \quad (p,a,0) \notin \mathcal{M} \end{cases} .$$

Thus, we have $V^{\downarrow}_{p_0,\mathcal{M}}(a_1) = 0$ as soon as a_1 has been selected for solving p and $\omega = \omega_3$. From this moment of time, a_1 and, sooner or later, a_2 and a_3 are rated equally and an act might be selected, e.g., by flipping a coin. In other words, the problem which occurs when basing decisions on (7.19) is the fact that this criterion does not, in the long run, discriminate between two acts even though the first one strictly dominates the second one. It is interesting to compare this with the MAXIMIN rule which also does not discriminate between a_1 and a_3.[17] This, however, seems to be acceptable more easily than the same property for (7.19): If used in connection with one-shot decisions, the MAXIMIN rule does not memorize experience from previous problem solving epochs. As opposed to a case-based decision rule, it does not have the opportunity of learning and experimenting in the course of a repeated problem solving process.[18]

The aforementioned drawback can be avoided by (7.29) in conjunction with a proper formalization of the "for most" concept. In fact, since (7.29) allows for a few exceptions (and ω_3 is assumed to occur but seldom) we will probably have $V_{p_0,\mathcal{M}}(a_2) = V_{p_0,\mathcal{M}}(a_3) = 0 < 1 = V_{p_0,\mathcal{M}}(a_1)$. Then, the relative frequency of selecting a_1 will converge toward 1 (instead of 1/3, as it would do in connection with a random choice between equally rated acts a_1, a_2, a_3). More precisely, suppose the "for all" quantifier to be defined such that it yields 1 if the property under consideration is satisfied in at least $100(1-\varepsilon)$ percent of the cases and 0 otherwise. In terms of our notation above, this means

$$\mu(k) = \begin{cases} 1 & \text{if} \quad k/m \geq 1-\varepsilon \\ 0 & \text{if} \quad k/m < 1-\varepsilon \end{cases} .$$

We will then have $V_{p_0,\mathcal{M}}(a_1) = 0$ if the proportion π_m of cases in which ω_3 has occurred in connection with a_1 exceeds ε, where m is the number of times a_1 has been chosen. Otherwise, we have $V_{p_0,\mathcal{M}}(a_1) = 1$. The probability that $\pi_m > \varepsilon$ and, hence, the probability that $V_{p_0,\mathcal{M}}(a_1) = 0$ will be small if ε is chosen sufficiently large in relation to the probability of the occurrence of ω_3. On the other hand, ε

[16] This clearly exemplifies that the application of (7.2) does hardly make sense.

[17] A discrimination can be achieved by extensions of MAXIMIN, such as the ordinal decision rules DISCRIMIN and LEXIMIN [152].

[18] This argument is no longer valid in a game playing context. Then, however, MAXIMIN can be justified by the assumption of an opponent acting optimally.

should not be made too large since otherwise $V_{p_0,\mathcal{M}}(a_2) = 1$ as well, which means that a_1 and a_2 are rated equally. An interesting idea arising in this context, which leads to a further extension of the model, is that of *learning* an optimal "for most" concept (from a parameterized class of membership functions). This can be seen as the counterpart to learning an optimal aspiration level in CBDT [168]. In our example, where the membership function μ depends only on ε, this parameter itself can be considered as an aspiration level.

Notice that the probability of $\pi_m > \varepsilon$ decreases with m if the probability that $\omega = \omega_3$ is smaller than ε. Thus, the probability of disqualifying a_1 is, if at all, relatively large at the beginning of a decision sequence, i.e., as long as a_1 has not been tried very often. This problem can be alleviated by means of a more flexible specification of the "for most" concept. Namely, the smaller the value of m, the less restrictive this concept should be specified in terms of the membership function μ. The definition above, for instance, could be generalized such that ε depends on m, i.e., $\mu(k) = 1$ if $k/m \geq \varepsilon_m$ and $\mu(k) = 0$ otherwise, with a non-increasing sequence $(\varepsilon_m)_{m\geq 0}$.

Let us now pass over from the (case-based) valuation of single acts (in the context of a certain problem) to the valuation of complete decision strategies. Of course, the question when to prefer a certain decision rule to an alternative criterion is by no means obvious in connection with the assumption of an ordinal setting for decision making. In fact, all kinds of "averaging" like, e.g., the derivation of the mean of the obtained utility values, are out of the question. Using the worst outcome, which might appear natural if (7.19) is seen as a kind of (case-based) analogue of the MAXIMIN decision rule, seems critical as well. In fact, within a *case-based* decision framework it is principally not possible to fully realize the idea underlying this (pessimistic) principle. Namely, an agent knows the possible consequences of a decision only *after* having applied the corresponding act. Then, however, the worst outcome has already occurred. In other words, it is impossible for a case-based decision maker to avoid the worst outcome in any case or to choose acts according to a (proper) MAXIMIN principle.

In connection with a model in which problems are chosen repeatedly according to some probability it seems reasonable to prefer a decision strategy S to a strategy S' if the former *dominates* the latter (stochastically) in the following sense: Let $U = \{u_1, u_2, \ldots, u_m\}$ such that $u_1 < u_2 < \ldots < u_m$ define the (linearly ordered) utility scale, and denote by $P_k^n(S)$ the probability of obtaining the utility u_k in the n-th step of a decision sequence if strategy S is used.[19] Then, S dominates S' (stochastically) if

$$\forall\, n \in \mathfrak{N}\; \forall\, 1 \leq k \leq m \,:\, \sum_{\imath=k}^{m} P_\imath^n(S') \leq \sum_{\imath=k}^{m} P_\imath^n(S). \tag{7.30}$$

[19] Observe that the sequences $(a(n))_{n\geq 1}$ of decisions and $(u(n))_{n\geq 1}$ of obtained outcomes resp. utility values are well-defined stochastic processes. In fact, for a (deterministic or stochastic) case-based decision procedure, the n-th decision is a function of the stochastic sequence of the first n problems $(p(1), \ldots, p(n))$.

For our example above, we have $U = \{0, 1\}$, i.e., $P_0^n(S)$ and $P_1^n(S)$ simply correspond to the probability of obtaining a "bad" and a "good" outcome, respectively, in connection with the n-th decision. Moreover, a decision criterion S is preferred to S' in the sense of (7.30) if $P_1^n(S) \geq P_1^n(S')$ for all $n \in \mathfrak{N}$.

Appendix F shows simulation results for different decision strategies S_ε which differ only with respect to the choice of ε, i.e., the definition of the "for most" quantifier. The states $\omega_1, \omega_2, \omega_3$ occur with probability 0.6, 0.3, and 0.1, respectively. Acts are evaluated according to (7.29), and ties between equally rated decisions are broken by coin flipping.

The results confirm the supposition that ε should satisfy $0.1 < \varepsilon < 0.4$. The critical values are $\varepsilon = 0.1$ and $\varepsilon = 0.4$. For $\varepsilon < 0.1$, the agent is overly ambitious, and all acts will sooner or later be judged equally and, hence, $P_1^n(S_\varepsilon) \rightarrow 1/2$ as $n \rightarrow \infty$. Letting $0.4 < \varepsilon$ is "too tolerant" in the sense that $V_{p_0,\mathcal{M}}(a_2) = 1$ in the long term, which means that (7.29) does not differentiate between a_1 and a_2 and, therefore, $P_1^n(S_\varepsilon) \rightarrow 3/4$ as $n \rightarrow \infty$. Note that the estimation of $P_1^n(S_\varepsilon)$ from the sequence $(u(n))_{n \geq 1}$ of obtained utility values is a good starting point for learning an optimal value for ε, i.e., for choosing an optimal "for most" concept from $\{\mu_\varepsilon \,|\, 0 \leq \varepsilon \leq 1\}$.

7.5 A CBI framework of CBDM

CBDT as introduced in [167] is largely motivated by practical problems arising in connection with EUT, notably the considerable need of precise information for modeling decision problems. Indeed, the specification of an EUT model might often be complicated and expensive, especially when having to solve relatively novel decision problems. In this section, we shall propose a framework of CBDM which also makes use of case-based reasoning in order to alleviate this problem, but which remains closer to classical decision theory. Loosely speaking, the idea is to apply the methods of case-based inference (CBI) discussed in previous chapters in order to support the modeling of decision problems.

7.5.1 Generalized decision-theoretic setups

The basic EUT setup (in the finite case) can be illustrated in the form of a table as follows:

$$\begin{array}{c|cccc} & \rho_1 & \rho_2 & \cdots & \rho_n \\ & \omega_1 & \omega_2 & \cdots & \omega_n \\ \hline a_1 & u_{11} & u_{12} & \cdots & u_{1n} \\ a_2 & u_{21} & u_{22} & \cdots & u_{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_m & u_{m1} & u_{m2} & \cdots & u_{mn} \end{array} \tag{7.31}$$

The ω_j constitute the set Ω of states of nature, and each ω_j is assumed to occur with probability ρ_j. Choosing act a_i yields utility u_{ij} if the state of nature is ω_j, which means that the expected utility of a_i is given by $\sum_{j=1}^{n} \rho_j u_{ij}$. The expected utility framework can be generalized in order to deal with infinite sets of acts and/or states of nature. Subsequently, however, we assume $\mathcal{A}$ and Ω to be finite.

When modeling a decision problem, some of the information in (7.31) might be incomplete or even missing. This concerns mainly the probability distribution on Ω and the utility function $u : \mathcal{A} \times \Omega \longrightarrow U$ which assigns a utility degree to each tuple consisting of an act and a state of nature. The basic idea which is discussed in this section and which characterizes CBDM is the use of case-based inference for deriving corresponding estimations. Of course, this approach presupposes the existence of *cases*. As will be seen, there are different possibilities for defining a case, each of which leads to a different extension of the basic EUT setup.

For instance, let $\mathcal{Q}$ be a set of problems and suppose an EUT setup (7.31) to be associated with each problem $p \in \mathcal{Q}$:

$$\begin{array}{c|cccc} & \rho_1^p & \rho_2^p & \cdots & \rho_n^p \\ & \omega_1 & \omega_2 & \cdots & \omega_n \\ \hline a_1 & u_{11}^p & u_{12}^p & \cdots & u_{1n}^p \\ a_2 & u_{21}^p & u_{22}^p & \cdots & u_{2n}^p \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_m & u_{m1}^p & u_{m2}^p & \cdots & u_{mn}^p \end{array} \tag{7.32}$$

A case is then defined as a triple (p, ρ^p, u^p), where ρ^p and u^p denote the probability distribution and utility function associated with the problem p, respectively. (The set of acts, $\mathcal{A}$, and the set of states of nature, Ω, are assumed to be fixed.) Within the framework of CBI, the problem p corresponds to an *input*. Moreover, ρ^p and u^p mark the *outcome* associated with a case, which can hence be written as a tuple $\langle p, (\rho^p, u^p)\rangle$. Note that such a case reduces to a tuple of the form $\langle p, \rho^p\rangle$ or $\langle p, u^p\rangle$ if either u^p or ρ^p is fixed in advance.

Suppose a decision maker to have a memory $\mathcal{M}$ of cases at its disposal. Given a new problem p_0, it can then make use of case-based inference in order to support the specification of a related EUT setup. This approach relies on the CBI assumptions

– that similar problems give rise to similar probability distributions on Ω, and/or

– that an act yields similar utilities for similar problems (under the same state of nature).

Example 7.4. Consider different types of urn experiments as an example: Let a state of nature, ω, correspond to the number of black balls in a random sample of size k. The sample is drawn from an urn which contains a large number K of balls, each of which is either black or white. An act a corresponds to an estimation of ω, and the utility $u(a, \omega)$ depends on the accuracy of this estimation, i.e., on the (absolute) difference $|a - \omega|$. Moreover, suppose that a problem is associated with the experimental conditions under which a sample is taken. The problems "simple selection with replacement" and "simple selection without replacement" can be considered as being similar if k/K is small. Indeed, the hypergeometric distribution which defines ρ in the latter case can then be approximated by the binomial distribution which is relevant if selected balls are replaced. From a CBI perspective, knowledge of the distribution ρ for the first problem can hence be seen as valuable information for defining the EUT setup for the (somewhat more complicated) second problem. Observe that the utility function is assumed to be known (and identical) for both problems. □

An alternative approach is to consider a setting in which the probability over Ω and/or the utility function depend not only on the problem but also on the act:

$$\begin{array}{c|cccc} & \rho_1^{(p,a)} & \rho_2^{(p,a)} & \dots & \rho_n^{(p,a)} \\ & \omega_1 & \omega_2 & \dots & \omega_n \\ \hline a & u_1^{(p,a)} & u_2^{(p,a)} & \dots & u_n^{(p,a)} \end{array} \tag{7.33}$$

A case can then be seen as a tuple $\langle (p, a), \mu_{p,a} \rangle$, where $\mu_{p,a}$ is a probability distribution on U. This definition is in accordance with the idea of a non-deterministic CBI setup as introduced in Section 2.4.2, where a random outcome is associated with each input. It can be considered as a generalization of CBDT which assumes the outcome associated with a problem–act tuple (p, a) to be deterministic. Thus, both frameworks (7.32) and (7.33) combine aspects from EUT and CBDT. The former, however, seems to be closer to EUT, whereas the latter is quite similar to CBDT.

Observe that a setup

$$\begin{array}{c|cccc} & \rho_1^{(p_0,a_0)} & \rho_2^{(p_0,a_0)} & \dots & \rho_n^{(p_0,a_0)} \\ \hline a_0 & u_1^{(p_0,a_0)} & u_2^{(p_0,a_0)} & \dots & u_n^{(p_0,a_0)} \end{array} \tag{7.34}$$

makes also sense within the original context of CBDT where a problem–act tuple has a unique outcome, i.e., if a case is a triple (p, a, r) resp. a tuple $\langle (p, a), r \rangle$. Then, however, an unknown outcome (or utility) is not considered as a random variable (in the proper sense), and an uncertainty measure η_{p_0,a_0} associated with a new problem p_0 and an act a_0 is interpreted as a quantification of a (subjective) belief concerning this outcome. Such a framework can be seen as defining an extended Bayesian approach in which CBI is used for assessing a (prior) measure of uncertainty over Ω. Symbolically, it can be illustrated as follows:

$$\left. \begin{array}{c} (p_1, a_1, r_1), \dots, (p_n, a_n, r_n) \\ p_0, a_0 \\ \sigma_{\mathcal{Q} \times \mathcal{A}}, \sigma_{\mathcal{R}} \end{array} \right\} \quad \stackrel{\text{CBI}}{\longrightarrow} \quad \eta_{p_0,a_0}. \tag{7.35}$$

7.5.2 Decision making using belief functions

The type of uncertainty measure derived in (7.35) depends on the way in which CBI is realized. Within the probabilistic framework of Section 4.5, for instance, the measure η_{p_0,a_0} takes the form of a belief function:

$$\eta_{p_0,a_0} = \mathsf{Bel}(H, \mathcal{M}, (p_0, a_0)) = \sum_{\imath=1}^{n} \alpha_\imath \cdot \mathsf{Bel}_\imath(H, (p_0, a_0)),$$

where H is a probabilistic similarity hypothesis and

$$\mathsf{Bel}_\imath(H, (p_0, a_0)) = \sigma_{\mathcal{R}}^{(-1)} \left(r_\imath, H(\sigma_{\mathcal{Q}\times\mathcal{A}}((p_0, a_0), (p_\imath, a_\imath))) \right)$$

denotes the belief function associated with the $\imath$-th case $(p_\imath, a_\imath, r_\imath) \in \mathcal{M}$. In this context, the probability distribution in (7.34) is replaced by a belief function. Consequently, the concept of an expected utility has to be generalized in order to evaluate an act. In other words, a framework of CBDM can be obtained by combining the CBI method of Section 4.5 and a generalization of expected utility based on belief functions. In recent years, several approaches to decision making on the basis of belief functions have been proposed in literature. Subsequently, we shall describe some of them very briefly.

Consider a belief function Bel on a set of outcomes, $\mathcal{R}$, and let m denote the mass distribution associated with Bel. Moreover, let $\mathcal{F}$ be the set of focal elements of m. A generalized expected utility can then be defined in terms of the Choquet integral

$$\int^{ch} u \, d\mathsf{Bel} = \int_0^\infty \mathsf{Bel}([u > t]) \, dt + \int_{-\infty}^0 (\mathsf{Bel}([u > t]) - 1) \, dt, \tag{7.36}$$

where $[u > t] \stackrel{\text{df}}{=} \{r \in \mathcal{R} \,|\, u(r) > t\}$. This approach is a pessimistic strategy in the sense that (7.36) is equal to the minimum (the infimum in the non-finite case [390, 389]) of a class of associated classical expected utilities:

$$\int^{ch} u \, d\mathsf{Bel} = \min_{\mu \in \mathcal{P}_{\mathsf{Bel}}} \int u \, d\mu, \tag{7.37}$$

where

$$\mathcal{P}_{\mathsf{Bel}} = \{\mu \in \mathcal{P}(\mathcal{R}) \,|\, \forall X \subset \mathcal{R} : \mathsf{Bel}(X) \le \mu(X)\} \tag{7.38}$$

is the set of probability measures over $\mathcal{R}$ compatible with Bel. As (7.38) reveals, this approach favors a lower probability interpretation of belief functions.

Choquet expected utility now plays an important role in research on axiomatic non-expected utility. This research direction is motivated by the paradoxes of ALLAIS [13] and ELLSBERG [140] which call the validity of the assumptions underlying EUT into question. A "behavioral foundation" of Choquet expected utility in the context of decision making under uncertainty has first been given by SCHMEIDLER [334], who uses the decision-theoretic setup of ANSCOMBE & AUMANN [15]. A corresponding extension of the approach of SAVAGE [331] has been proposed in [166]. These works have been refined by several authors. An appealing axiomatic characterization of non-additive expected utility somehow unifying [334] and [166] has been developed in [330]. In [190] it is shown that a common characterizing property of this line of research is a certain weakening of SAVAGE's axioms which essentially restricts the well-known *sure thing principle* to so-called comonotonic acts.

Related models for decision making with belief functions have also been proposed in [210]. The axiomatic theory developed in [212] gives a foundation to these decision models. Here, situations are considered in which information is ambiguous and not fully probabilizable. It is argued that entirely vague information should be processed according to the (objective) symmetry principles of *complete ignorance* [16, 69] (rather than to the principle of insufficient reason). Again, the most important aspect of the decision-theoretic framework developed in [212] is a natural weakening of SAVAGE's sure thing principle [331]. It is shown that, within the resulting axiomatic setting, decisions can be represented by belief functions on outcomes. More precisely, a representation of a preference relation on the set of acts is of the form

$$f \mapsto \sum_{F \in \mathcal{F}} \mathsf{m}_f(F)\, v(r_F, R_F) \ , \tag{7.39}$$

where m_f is the Möbius transform (mass distribution) associated with the belief function induced by the act $f : \mathcal{A} \longrightarrow \mathcal{R}$ on the set of outcomes. Moreover, r_F is the worst and R_F is the best outcome within $F \in \mathcal{F}$. As as special case of (7.39) the functional

$$\sum_{F \in \mathcal{F}} \mathsf{m}_f(F)\ \left(\alpha(r_F, R_F)\, u(r_F) + (1 - \alpha(r_F, R_F))\, u(R_F)\right) \tag{7.40}$$

is proposed, where u reflects the agent's attitude toward outcomes in decision under risk. The function α is interpreted as an index of the like or dislike of ambiguity.

A related generalization of the VON NEUMANN-MORGENSTERN framework has been proposed by JAFFRAY [211]. He combines the axioms of linear utility theory with axioms of rational decision making under *mixed uncertainty* [70] in order to justify a family of so-called Hurwicz α-criteria. According to these criteria, a belief function Bel over the set of outcomes $\mathcal{R}$ is evaluated by

$$\alpha \inf\{\mathbb{E}_\mu(u) \,|\, \mu \in \mathcal{P}_{\mathsf{Bel}}\} + (1 - \alpha) \sup\{\mathbb{E}_\mu(u) \,|\, \mu \in \mathcal{P}_{\mathsf{Bel}}\}, \tag{7.41}$$

where $\mathbb{E}_\mu(u)$ denotes the expected utility under the probability measure μ. The use of Hurwicz criteria is also advocated by STRAT [361].

YAGER [405] defines a generalized expected utility of the form

$$\sum_{F \in \mathcal{F}} \mathsf{m}(F) \cdot \phi(F) \tag{7.42}$$

which makes use of a set-function $\phi : 2^{\mathcal{R}} \longrightarrow \mathfrak{R}$. The problem of assigning a degree of utility, $\phi(F)$, to a focal set F is considered in the context of decision making under ignorance. It is proposed to solve this problem by applying an OWA (Ordered Weighted Average) operator[20] [404] to the collection $u(F)$ of utility degrees $u(r)$ $(r \in F)$.[21] That is, $\phi(F) = \mathrm{OWA}(u(F))$. Special cases of this operator include the well-known decision rules

$$\begin{aligned}\phi(F) &= \min u(F),\\ \phi(F) &= \alpha \min u(F) + (1-\alpha) \max u(F),\\ \phi(F) &= \sum_{r \in F} u(r) / \operatorname{card}(F).\end{aligned}$$

Note that the set-function ϕ in (7.42) allows one to model the agent's decision behavior under complete ignorance in a more general way than the extreme (pessimistic) valuation by means of the Choquet integral (where always the worst case is assumed) or the Hurwicz criteria (7.40) (where $\phi(F)$ depends only on the worst and the best element in F).

As (7.37) shows, the use of Choquet integration comes down to deriving a classical expected utility based on the selection of a probability measure compatible with the belief function. In [350] it has been proposed to apply a generalization of LAPLACE's insufficient reason principle in order to select a corresponding distribution:

$$\mu(\{r\}) = \sum_{F \in \mathcal{F}} \mathbb{I}_F(r)\, m(F) / \operatorname{card}(F). \tag{7.43}$$

The transformation (7.43), which corresponds to the betting function (4.22) introduced in Section 4.5.1, has been justified axiomatically in the context of the *transferable belief model* which favors a purely subjective (and non-probabilistic) interpretation of belief functions. Note that (7.43) is the distribution of maximum entropy among $\mathcal{P}_{\mathsf{Bel}}$, i.e., it can also be derived from the principle of maximum entropy.

7.5.3 Possibilistic decision making

In Chapter 6, we have proposed a possibilistic method of case-based inference which makes use of implication-based fuzzy rules. According to this approach, uncertainty concerning the outcome r_0 is characterized by means of a possibility distribution:

[20] Operators of this type are also known as linear order statistics in the field of robust statistics.

[21] Each outcome $r \in F$ contributes exactly one element to $u(F)$, i.e., the same utility degree might appear several times in $u(F)$.

$$\pi_{a_0,\mathcal{M}}(r') = \min_{(p,a,r)\in\mathcal{M}} \sigma_{\mathcal{Q}\times\mathcal{A}}((p_0,a_0),(p,a)) \rightsquigarrow \sigma_{\mathcal{R}}(r,r') \tag{7.44}$$

for all $r' \in \mathcal{R}$, where $\rightsquigarrow$ is a generalized implication operator. An alternative approach using conjunction-based (example-based) fuzzy rules has been outlined in Chapter 5. It leads to the possibility distribution[22]

$$\pi_{a_0,\mathcal{M}}(r') = \max_{(p,a,r)\in\mathcal{M}} \min\{\sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)), \sigma_{\mathcal{R}}(r,r')\}. \tag{7.45}$$

Suppose that outcomes are directly given in terms of utilities, i.e. $U = \mathcal{R}$. (Otherwise, a possibility distribution on the set of utility degrees can be obtained via $v \mapsto \max_{r\,:\,u(r)=v} \pi_{a_0,\mathcal{M}}(r)$.) The problem of choosing an act then turns out as one of choosing among the possibility distributions

$$\{\pi_{a_0,\mathcal{M}} \,|\, a_0 \in \mathcal{A}\}. \tag{7.46}$$

This situation is quite similar to decision under risk where the agent has to choose among probability distributions (lotteries).

There are different ways of realizing a corresponding selection. We can, for instance, adopt a quantitative point of view and interpret possibility degrees as upper probabilities. A possibility distribution then corresponds to a special type of plausibility measure, which means that the methods discussed in Section 7.5.2 can be applied.

We can, however, also interpret the possibilistic approach in a purely qualitative way. Dubois and Prade [123] have recently proposed a qualitative decision theory in which uncertainty and utility are represented by possibility measures and qualitative utility functions, respectively. The corresponding decision criteria are derived from an axiomatic framework which can be seen as a qualitative counterpart to the axioms of Von Neumann and Morgenstern's expected utility theory.

Let $\sqsubseteq$ be a preference relation on the class Π of normalized possibility measures on a finite set $\mathcal{R} = \{r_1, \ldots, r_n\}$ of outcomes. As usual, denote by $\sim$ and $\sqsubset$ the symmetric and anti-symmetric part of $\sqsubseteq$, respectively. Moreover, let V be a finite linear scale of uncertainty such that $\min V = 0$ and $\max V = 1$. Likewise, let U be a finite linear scale of preference such that $\min U = 0$ and $\max U = 1$. The commensurability between the ordinal scales U and V is achieved via an order-preserving mapping h from the plausibility scale to the preference scale which satisfies $h(0) = 0$ and $h(1) = 1$. For $\lambda, \mu \in V$ with $\max\{\lambda,\mu\} = 1$ the possibilistic mixture $(\lambda/\pi, \mu/\pi')$ of two possibility distributions π and π' again defines a possibility distribution:

$$\forall\, r \in \mathcal{R} \,:\, (\lambda/\pi, \mu/\pi')(r) \stackrel{\text{df}}{=} \max\{\min\{\lambda, \pi(r)\}, \min\{\mu, \pi'(r)\}\} \,.$$

In [103], the following axiomatic system P has been proposed:

[22] Note that in Chapter 5 this distribution has been denoted by δ instead of π. Here, this distinction is not needed.

P1 $\sqsubseteq$ is a total preorder.

P2 $\pi \leq \pi' \Rightarrow \pi' \sqsubseteq \pi$ (uncertainty aversion).

P3 $\pi_1 \sim \pi_2 \Rightarrow (\lambda/\pi_1, \mu/\pi) \sim (\lambda/\pi_2, \mu/\pi)$ (independence).

P4 $\forall \pi \in \Pi \; \exists \lambda \in V \; : \; \pi \sim (1/r^*, \lambda/r_*)$, where r^* and r_* denote a maximal and a minimal element of $\mathcal{R}$, respectively.[23]

Based on this set of axioms, the existence of a utility function $u : \mathcal{R} \longrightarrow U$ and the following *pessimistic* decision criterion, which represents the preference relation $\sqsubseteq$, are derived:

$$\mathrm{QU}^-(\pi) \stackrel{\mathrm{df}}{=} \min_{r \in \mathcal{R}} \max \{n(h(\pi(r))), u(r)\} \; . \tag{7.47}$$

That is $\pi \sqsubseteq \pi' \;\Leftrightarrow\; \mathrm{QU}^-(\pi) \leq \mathrm{QU}^-(\pi')$. Here, n is the order-reversing function on U.

As an alternative model, an axiomatic system O has been proposed in which the uncertainty aversion axiom P2 is replaced by an uncertainty-prone postulate. Moreover, P4 is slightly modified:

O1 $\sqsubseteq$ is a total preorder.

O2 $\pi \leq \pi' \Rightarrow \pi \sqsubseteq \pi'$.

O3 $\pi_1 \sim \pi_2 \Rightarrow (\lambda/\pi_1, \mu/\pi) \sim (\lambda/\pi_2, \mu/\pi)$.

O4 $\forall \pi \in \Pi \; \exists \lambda \in V \; : \; \pi \sim (\lambda/r^*, 1/r_*)$.

Based on these axioms one obtains the *optimistic* decision criterion

$$\mathrm{QU}^+(\pi) \stackrel{\mathrm{df}}{=} \max_{r \in \mathcal{R}} \min \{h(\pi(r)), u(r)\} \; . \tag{7.48}$$

If we interpret the approach to CBI outlined in Chapters 5 and 6 as purely qualitative ones (and also assure the commensurability of the plausibility scale and the preference scale), the decision theory of [123] can be applied to the distributions (7.44) or (7.45). That is, the decision criteria derived from the above axioms can be used in order to choose the most preferred distribution from the set (7.46), and, hence, the most preferred act. Applying (7.47) resp. (7.48) leads to the following valuations of an act $a \in \mathcal{A}$:

$$\mathrm{V}^{\downarrow}(a) = \mathrm{QU}^-(\pi_{a,\mathcal{M}}) = \min_{r \in \mathcal{R}} \max \{n(h(\pi_{a,\mathcal{M}}(r))), u(r)\}, \tag{7.49}$$

$$\mathrm{V}^{\uparrow}(a) = \mathrm{QU}^+(\pi_{a,\mathcal{M}}) = \max_{r \in \mathcal{R}} \min \{h(\pi_{a,\mathcal{M}}(r)), u(r)\} \; . \tag{7.50}$$

Let us finally mention that the uncertainty averse and uncertainty prone postulates P2 and O2 can be replaced by (intuitively plausible and somewhat more

[23] We extend $\sqsubseteq$ to $\mathcal{R}$ in the usual way: $r \sqsubseteq r'$ iff $\pi_r \sqsubseteq \pi_{r'}$, where $\pi_r = \mathbb{I}_{\{r\}}$ and $\pi_{r'} = \mathbb{I}_{\{r'\}}$.

appealing) *possibilistic dominance* criteria which are possibilistic counterparts to the well-known concept of probabilistic dominance. This result is proved in Appendix A.

7.6 CBDM models: A discussion of selected issues

In this section, we shall discuss some selected issues in case-based decision making. Our emphasis is on pointing out some principal differences between the models outlined in previous sections. In order to demarcate the different approaches, we shall reserve the acronym CBDM mainly for the framework presented in Section 7.5. Since the methods from Sections 7.1–7.4 are closer to case-based decision theory originally introduced by GILBOA and SCHMEIDLER, they will be referred to as CBDT.

7.6.1 The relation between similarity, preference, and belief

A main difference between the models outlined in Section 7.5 (CBDM) and the approaches of previous sections concerns the way in which the concepts of belief, preference, and similarity are related. The approaches of Sections 7.1–7.4 make use of a decision-theoretic setup which is based on the concepts of similarity and utility alone. As opposed to this, the framework of CBDM in Section 7.5 makes also explicit the concept of belief and can thus be seen as an extension of classical (statistical) decision-theoretic models. In fact, this approach realizes a two-stage process, in which the actual decision problem is only solved in the second stage by means of (more or less) common techniques from decision theory. Case-based reasoning is not used for selecting an act directly. Rather, it has influence on the formation of the *belief* of the decision maker. This belief is represented in the form of a belief function or possibility distribution on the set of outcomes, $\mathcal{R}$. The cases contained in a memory $\mathcal{M}$ are treated as observations. For instance, observing that an act a has led to a good result for a similar problem will increase the agent's belief that a is also a good choice for the problem at hand.

The derivation of (7.2) in [167] shows that an agent with a utility function u, who obeys the respective axioms, behaves *as if* it had a similarity measure over $\mathcal{Q}$ and evaluates acts according to (7.2). This way, similarity is directly related to utility and indirectly to preference. The formal resemblance of (7.2) and the EUT formula, i.e., the expected utility of an act, suggests that the meaning of similarity in CBDT is to some extent comparable to the role that probability plays in EUT.

Most approaches to decision making evaluate acts by combining preference and belief in some way, where preference is quantified in the form of a utility function. In fact, for estimating the utility one obtains when choosing a certain act it seems natural to consider the set V of possible utility degrees,[24] to modify each degree v in accordance with an associated degree of belief, and to aggregate these modified utilities.[25] In expected utility theory, for instance, degrees of belief associated with $v \in V$ (and an act a) correspond to probabilities $p_a(v)$, and modification and aggregation are realized by multiplication and addition, respectively:

$$V(a) = \sum_{v \in V} p_a(v) \cdot v. \tag{7.51}$$

Within the qualitative approach proposed in [103, 123], belief is represented by possibility degrees $\pi_a(v)$, modification corresponds to bounding the impact of less possible utility degrees upon the valuation of an act, and the min-operator is used as an aggregation function:

$$V(a) = \min_{v \in V} \max\{1 - \pi_a(v), v\}. \tag{7.52}$$

Observe that the averaged similarity version of (7.2) corresponds to the expected utility model (7.51) if the probability $p_a(v)$ is estimated according to

$$p_a(v) = \frac{\sum_{(p,a,r) \in \mathcal{M}, u(r)=v} \sigma_{\mathcal{Q}}(p, p_0)}{\sum_{(p,a,r) \in \mathcal{M}} \sigma_{\mathcal{Q}}(p, p_0)}. \tag{7.53}$$

Likewise, (7.19) is equivalent to (7.52) with

$$\pi_a(v) = \max_{(p,a,r) \in \mathcal{M}, u(r)=v} \sigma_{\mathcal{Q}}(p, p_0). \tag{7.54}$$

As can be seen, based on the idea that similarity is used for assessing a degree of belief, namely (7.53) resp. (7.54), it is possible to interpret the approaches (7.2) and (7.18) within an extended decision-theoretic framework which combines similarity, preference, and belief, even though the latter only appears implicitly.

Still, there are several motivations for modeling the (causal) relation between similarity and belief in a more explicit way, as we have done in Section 7.5. Firstly, viewing the cases of a memory as an (additional) information source which has an effect on the agent's belief and, hence, utilizing case-based reasoning for decision making only indirectly leads to a more expressive approach which also avoids some technical difficulties. This becomes obvious, for instance, when considering the extreme example of a memory that does not contain any case similar to the current problem, which means that the memory is effectively empty. If, however, no cases exist, it seems somewhat peculiar that a *case-based* (similarity-based) reasoning procedure should be used for estimating the utility of choosing some act for solving the problem. Instead of assigning a "default utility," it appears more natural

---

[24] For the sake of simplicity suppose this set to be finite.

[25] Note that the consideration of single utility degrees may not be enough if belief is formalized by means of non-additive measures of uncertainty [166, 330, 334].

to expect the result of case-based (similarity-based) reasoning to be *complete ignorance* about utilities, which is adequately reflected, e.g., by the possibility distribution $\pi \equiv 1$ on the set of outcomes. Needless to say, an uncertainty measure like a probability distribution, a belief function or a possibility distribution, is able to reproduce certain characteristics of a memory $\mathcal{M}$ better than a "point estimation." The averaged similarity version of (7.2), for instance, can be seen as a kind of weighted mean. It is unable, however, to represent the *variance* of utility degrees associated with a certain act.

Secondly, making uncertainty related to decision problems explicit allows for taking the agent's attitude toward uncertainty into account. Otherwise, this attitude has to be encoded in the similarity measure or the utility function. Suppose, for example, that a decision maker (repeatedly) faced with a problem p can choose between two acts a and b. Act a yields utility 0 with certainty. The more risky act b yields either an extremely high utility M or an extremely low utility $-M$, where the high utility occurs with a fixed but unknown probability every time b is chosen. The willingness of an (uncertainty averse) agent to choose b will then depend on the number of times the cases (p, b, M) and $(p, b, -M)$ have been observed.[26] The memory $\mathcal{M} = \{(p, b, M)\}$ containing only one case, for instance, might not be convincing enough, even though $V(b) = M > 0 = V(a)$ according to (7.2).

Thirdly, the distinction between two "mental" levels, one for representing knowledge and one for making decisions, seems to have advantages with respect to the design of intelligent systems [130], i.e., if a decision-theoretic model is understood as a language for modeling the problem solving behavior of an agent. The integration of different information sources at the decision level, for instance, would require a related extension of the underlying decision theory and seems to be more difficult than doing the same at the knowledge representation level. Consider again the approach of Section 7.5.3 as an example. There, case-based reasoning takes place at the knowledge representation level and yields a possibility distribution on the set of outcomes. It is hence not difficult to combine this case-based knowledge with general background knowledge represented, e.g., in the form of fuzzy rules. In fact, the possibility distributions associated with such rules can simply be combined (via intersection) with the distribution(s) originating from CBI.

Let us finally mention that the (causal) relation SIMILARITY $\rightarrow$ BELIEF is also supported by psychological evidence. In fact, the finding that people rely on similarity as a heuristic principle for assessing the *probability* of an uncertain event or the value of an uncertain quantity was made by TVERSKY and KAHNEMAN in various psychological studies [374]. The authors call this heuristic approach the *representativeness principle*. For example, the probability that a person has a certain job seems to be assessed by the degree to which this person is similar to the stereotype of a person having this job.

[26] In other words, the agent estimates the unknown probability that M occurs by the corresponding (relative) frequency of occurrence.

7.6.2 The effect of observed cases

The impact that case-based information has upon the evaluation of acts is rather different for the decision models discussed in this chapter. A major difference concerns the question whether experienced cases can be compensated by other cases, e.g., whether a good experience can compensate for a bad one, or whether several moderately similar cases can outweigh one completely similar case.

Let us consider the last point first. Due to the accumulation of utility degrees in (7.2), a good experience with a very similar problem can be compensated by several good experiences with less similar problems. This contrasts, e.g., with the NEAREST NEIGHBOR decision rule (7.8) which takes only one (namely the most similar) observed case into account, i.e., which fully relies on the most relevant experience. Needless to say, the adequacy of the two principles will strongly depend on the application or, more precisely, the extent to which experience with a certain act can be transferred from one problem to a similar one. Consider, for instance, a medical agent having to choose between treatments T_1 and T_2. The successful application of therapy T_1 to several diseases with somewhat similar symptoms will generally not compensate for T_2's curing exactly the same symptoms, even if T_2 has not been applied to any other disease.

Now, consider the second point, i.e., the question whether good experiences can compensate for bad ones and vice versa. The CBDT decision rule (7.2) as well as the averaged similarity version (7.28) do obviously allow for such a compensation, and the same is true for the NEAREST NEIGHBOR decisions in Section 7.2. As opposed to this, the criteria (7.19) and (7.20) compensate in only one direction: According to (7.19), an observed case can only *decrease* the evaluation of an act, which reflects the pessimistic or cautious character of this decision rule. Consequently, a positive experience cannot compensate for a negative one. Contrariwise, each observation can only positively influence the evaluation according to (7.20), i.e., a good experience is never annulled by a bad one.

Again, different evaluation principles will be adequate for different applications. In this connection, it should be noted that (7.19) and (7.20) assume an ordinal setting, whereas the addition and multiplication operators used by the CBDT criteria (7.2) and (7.3) make sense only for cardinal utility and similarity functions. Indeed, (7.19) and (7.20) might be preferred whenever it is difficult to define such functions. Consider again an example from the medical domain: Treatments T_1 and T_2 usually have the same effect. On the one hand, T_2 is less expensive than T_1. On the other hand, it is also more risky in the sense that is has already caused the death of some patients, whereas T_1 cures the disease with certainty. In this situation, it will of course be difficult to come up with a reasonable utility function, or to fix a minimal success rate of T_2 as a decision criterion.[27] Rather, one will generally be cautious and decide in favor of T_1, a decision behavior which is perfectly in line with (7.19).

[27] Extreme examples of this kind are often raised against expected utility theory.

It was pointed out above that (7.19) and (7.20) reflect very opposite attitudes of a decision maker. Let us finally mention that a similar remark applies to the indirect evaluations which deduce the *possibility* of outcomes. When using the criterion (7.49) in connection with the possibility distribution (7.44), the decision maker considers all outcomes as being fully possible as long as it has not made any observations. Each new case serves as a constraint and decreases the possibility of certain results. By applying the same decision rule to (7.45), the decision maker starts with the possibility distribution $\pi \equiv 0$.[28] Each new case serves as evidence for certain results and increases the possibility correspondingly. Loosely speaking, the agent learns what *can* happen, whereas it learns what *can* or *should not* happen if it relies on (7.44). The difference between (7.44) and (7.45) becomes also clear if we realize that (7.44) is based on the idea of an implication-based fuzzy rule, whereas (7.45) is related to the concept of a possibility rule, i.e., an example-based (conjunction-based) fuzzy rule.

7.6.3 Dynamic aspects of decision making

Since CBI is closely related to the idea of repeated problem solving and aspects of learning it seems natural to consider a CBDM agent acting over time in a certain environment. The question, then, is how successful a CBDM strategy proves to be. Since the acquisition of experience in the form of cases is an inherent part of CBDM, investigating a CBDM strategy in the context of repeated problem solving seems to be the only reasonable way of judging its efficiency.[29] Such an analysis, which will have much in common with the analysis of heuristic problem solving methods [291], is principally possible. For instance, given (among other things) the precise specification of a stochastic environment in which the agent acts as well as the specification of utilities of histories (which correspond to paths in this environment), the *expected* performance of a CBDM strategy is well-defined (cf. Section 7.4).

Let us mention, however, an interesting aspect of CBDM which makes the analysis of a given strategy as well as the selection of an optimal strategy difficult. Namely, a single decision at a certain point of time does not only affect the expected utility and future states of an agent directly. Rather, it has also an impact on the agent's *experience* and, hence, changes its future decision behavior. For the analysis of a given decision strategy this means that it has to take the (expected) evolution of the agent's memory into account. For the development of an optimal strategy it means that a single decision should not only be judged on the basis of some estimated (immediate) utility. Since a more experienced agent

[28] Again, note that this is actually a distribution of *guaranteed possibility*, denoted by the symbol δ in Section 5.

[29] Considering something such as the performance of a decision strategy makes sense if we concede to CBDM a *normative* character in connection with the idea of heuristic problem solving. Other criteria become relevant if a case-based decision theory serves as a *descriptive* theory of (human) decision making [171].

will probably make better decisions in the future it should also take the aspect of broadening experience into account. Informally speaking, the agent has to find a tradeoff between the *exploitation* of its past experience and the *exploration* of new decisions. This, of course, requires some kind of metalevel reasoning quite comparable to the concept of metalevel rationality in connection with expected utility theory (cf. the remarks on page 253). The aforementioned exploration–exploitation tradeoff is also well-known in fields like optimization and machine learning.

The above idea can be illustrated nicely for the model (7.2). As an interesting consequence of this decision principle it has been pointed out in [167] that it can be seen as a theory of "bounded rationality" formalizing SIMON's idea of "satisficing" [260, 344]. Suppose, for example, that the selection of a certain act a_* for a problem p has led to a positive utility $u(r(p, a_*))$. When faced with the same problem again, the decision maker will prefer this act to all other acts to which the default utility 0 is assigned (since they have not been tried yet). More generally, the agent may try several acts until one results in a positive utility, but it will not attempt to maximize utility. Now, an intuitively reasonable modification of the decision behavior prescribed by (7.2) is to try a new act a from time to time. This way, an act a^* such that $u(r(p, a^*)) > u(r(p, a_*))$ might eventually be found. Since the agent will then go on choosing a^*, this would have a positive impact on its (future) "welfare," a prospect that justifies to put up with the risk of sometimes obtaining a smaller utility. See [168] for a related strategy of realizing an "experimenting" agent in CBDT.[30] The idea is to adapt the aspiration level α in the generalization

$$V_{p_0,\mathcal{M}}(a) \overset{\text{df}}{=} \sum_{(p,a,r)\in\mathcal{M}} \sigma_{\mathcal{Q}}(p, p_0) \cdot (u(r) - \alpha)$$

of (7.2) by choosing an act at random from time to time. This way, the agent can avoid to get stuck in a suboptimal strategy.[31]

7.7 Experience-based decision making

Case-based decision making, as presented in different versions in previous sections, can basically be seen as a two-step procedure:

I. Estimation/evaluation: Given a set of experiences in the form of triples (p, a, u) and a new problem p_0, one estimates the utility $u(p_0, a_0)$ for each act $a_0 \in \mathcal{A}$.

[30] As the "conservative" decision strategy of always choosing the act "go to a restaurant which has not been tried yet" shows, a careful distinction between the agent's decisions and its actual behavior has to be made. Particularly, a satisficing decision strategy does not necessarily entail conservative behavior.

[31] The same idea is also present in several approaches to learning in game theory.

II.Decision: An (apparently) optimal act is then chosen on the basis of these estimations.

It deserves mentioning that one actually has to distinguish between the *estimation* of a utility degree and the *evaluation* of an act. As opposed to (7.3), for instance, the value $V(a_0)$ in (7.2) is obviously not an estimation of the utility $u(p_0, a_0)$ but still an evaluation of the act a_0. As can be seen, many possibilities of act evaluation exist, although the estimation of an induced utility might be regarded as the most natural one. Subsequently, we make the reasonable assumption that the agent bases its decision on estimations of the utilities of acts a_0. That is, we assume that the agent is an *estimated* utility maximizer, just like a decision maker applying EUT is an *expected* utility maximizer.

Experience-based decision making [196] generalizes case-based decision making in the sense that the estimation of utility degrees as part of the above two-step procedure is realized by any learning method, not necessarily a case-based one. Note that EBDM thus defined is an *indirect* approach in which an approximation Δ to the *optimal decision function*

$$\Delta^* : \mathcal{P} \longrightarrow \mathcal{A}, \tag{7.55}$$

which maps problems to (optimal) decisions, is derived from an estimation $\hat{u}$ of the utility function $u : \mathcal{P} \times \mathcal{A} \longrightarrow \mathfrak{R}$:

$$\Delta : d \mapsto \arg\max_{a \in \mathcal{A}} \hat{u}(p, a).$$

An obvious alternative is to realize EBDM in a more *direct* way. In this case, the agent tries to learn the decision function (7.55) directly, without estimating the utility function as an intermediate step.[32] This kind of EBDM, which appears especially appealing from an efficiency point of view, will be discussed in detail in Section 7.7.1.

Case-based decision making, as case-based reasoning in general, is closely related to learning from experience in the form of examples or facts. Investigating the link between factual knowledge and beliefs derived from that knowledge, this relation is also emphasized in [173], where the axiomatic foundation of CBDT is developed in a more general context, not restricted to decision making.

In the field of machine learning, several standard types of learning problems are distinguished. In this connection, it is interesting to note that the indirect approach to EBDM (as realized by CBDT) is closely related to *reinforcement learning*, at least from a formal point of view. There are different settings for reinforcement learning, most of which fall back on concepts from Markov Decision Processes: The decision making agent acts in some unknown environment defined by a set of states S. At each point of time, the agent finds itself in a state $s \in S$, where an action has to be performed. The consequences of performing action a in

[32] Such agents are often called *reflex agents* in artificial intelligence.

state s are determined by a *reward function*, r, and a *transition function*, δ: The agent receives an immediate reward,[33] $r(s, a)$, and moves from state s to state $\delta(s, a) \in S$. This process is repeated until eventually a *final state* is reached.

As can be seen, the notions of state and reward in reinforcement learning play the roles, respectively, of the concepts of problem and utility in CBDT. Moreover, the optimal decision function Δ^* in EBDM basically corresponds to what is called an *optimal policy* in reinforcement learning. The basic difference between CBDT and reinforcement learning (sequential decision making) concerns the generation of problems resp. states. In sequential decision making, the next state is a (perhaps non-deterministic) function of the current state s and the action a. Consequently, an action does not only determine the immediate reward, but also the next decision problem and, thereby, the prospect of future rewards. A "myopic" decision maximizing only the immediate reward $r(s, a)$ is hence not necessarily optimal. Rather, an optimal action should be one that maximizes the sum of the immediate reward and the (expected) future rewards.[34] A function taking this into account is the so-called Q-function, that can be defined as follows:

$$Q(s, a) = r(s, a) + \gamma \sum_{s' \in S} p(s, a, s') \cdot \max_{a'} Q(s', a'), \tag{7.56}$$

where $0 \leq \gamma \leq 1$ is a discounting factor. Moreover, $p(s, a, s')$ is the probability that act a yields s' as the successor state of s (here the transition function δ is non-deterministic). This leads to BELLMANN's optimality equations

$$V^*(s) = \mathbb{E}\left[r(s, \Delta^*(s)) + \gamma V^*(\delta(s, \Delta^*(s)))\right], \tag{7.57}$$

which determine the optimal decision function Δ^*. Thus, the value of being in state s, $V^*(s)$, is the expected sum of the immediate reward and the discounted future rewards under optimal behavior, as prescribed by Δ^*.

In CBDT, the action chosen for a problem p does not affect the occurrence of future problems, which are not under the control of the decision maker. Thus, an optimal decision is simply one that maximizes $u(p, a)$. Note that the same strategy, namely maximizing $r(s, a)$, is also optimal in sequential decision problems if either future rewards (utilities) are discounted by means of a discounting factor of $\gamma = 0$ or if the transition function δ (as realized, e.g., by the probability function p in (7.56)) does not depend on a. In other words, there are two possibilities of viewing CBDT, at least formally, as a special type of reinforcement learning: Either the case-based decision maker follows a myopic strategy, or future states (problems) do not depend on actions.

The main objective in reinforcement learning is to estimate the Q-function on the basis of rewards obtained so far.[35] If $\widehat{Q}$ is such an estimation, an apparently

[33] In a more general version, feedback can also be delayed (e.g., the win or loss of a game).

[34] The addition of rewards might be replaced by an alternative aggregation operator, of course.

[35] The reward function r and probability function p in (7.56) are assumed to be unknown. Otherwise, classical approaches (e.g., dynamic programming techniques in the case of finite horizon decision problems) can be used to find an optimal policy on the basis of (7.57).

optimal strategy is given by the decision function

$$\Delta : s \mapsto \max_a \widehat{Q}(s, a).$$

Even though $\Delta(s)$ does indeed maximize the estimated rewards, the strategy of persistently choosing actions $\Delta(s)$ so as to maximize the current estimation of Q is again somehow shortsighted and not necessarily optimal. Namely, the agent has to bear in mind that it keeps learning over time: The estimation of the Q-function is permanently revised in the light of new observations, and the improvement of $\widehat{Q}$ might be larger for an alternative action $a' \neq \Delta(s)$. Consequently, choosing a' might lead to better decisions in the future, even though the immediate reward might be better for a. In other words, the agent has to find an *exploration–exploitation tradeoff* (cf. page 287): It has to trade off (estimated) rewards against the potential for learning useful new information. One strategy, for instance, is to make random choices, where the probability of an action is proportional to its estimated value. This way, preference is given to higher valued acts, but apparently suboptimal acts are not completely ignored. Another possibility is to assign relatively large default values to yet unknown states (or state–action pairs) so as to offer an incentive for exploring such states.

As already mentioned in Section 7.6.3, the exploration–exploitation problem is solved in a very similar way in CBDT, namely by assigning default utilities to problem–act tuples for which relevant experience is not available as yet. The induced satisficing behavior of a decision maker can be seen as a special exploration strategy: The agent tries new actions until a satisfying one has been found.

7.7.1 Compiled decision models

The modification (7.3) of evaluation (7.2) can be seen as a special version of estimated utility maximization as discussed in Section 7.7. In fact, (7.3) is nothing else than the application of SHEPHARD's interpolation method [340], a special type of NEAREST NEIGHBOR (NN) estimation [76].[36] This method is well-known in machine learning, and it is used for making predictions in other CBR approaches as well (e.g., in the ELEM2-CBR system [61]).[37]

Of course, SHEPHARD's interpolation method is not the only way of realizing the estimation step in EBDM. Principally, it could be replaced by any machine learning method. In this connection, it is especially interesting to distinguish between *instance-based* and *model-based* approaches to (supervised) machine learning [79]. In particular, our discussion in Section 2.1 has shown that instance-based learning, as a lazy approach, is easy and quite appealing from a knowledge revision and adaptation point of view, but not very efficient in the prediction phase. In

[36] This is the weighted k-NEAREST NEIGHBOR approximation with $k = n$, i.e., it makes use of the complete set of observations.

[37] See the paper [173] of Gilboa and Schmeidler for a comparison between the nearest neighbor method and their case-based approach.

the context of EBDM, this means that a case-based approach appears reasonable if the decision maker disposes of a *limited* number of experiences. If a large set of observed cases is available, however, a *compressed* representation of the agent's knowledge in the form of a *model* might be more efficient. The aspect of efficiency becomes especially relevant if deliberation time is costly or strictly limited as in real-time decision making [64]. Then, model-based learning might be preferred to instance-based learning, since it is the decision process itself rather than the learning process that is time critical.

In this connection, it is convenient to classify decision problems with respect to their novelty. According to a crude distinction, we can differentiate between problems which are solved frequently and hence become almost automated, problems for which deliberation is required but which are still familiar, and problems which are completely unfamiliar [167]. These problem types might be tackled most efficiently by means of different approaches to learning and knowledge representation:

– instance-based learning of the utility function for unfamiliar problems,

– model-based learning of the utility function for familiar problems,

– and a "compiled decision model" approach for routine decisions.

As already mentioned before, the idea of a compiled decision model is to learn the optimal decision function (7.55) directly, rather than making a detour by learning the utility function. In [325], *compilation* is understood as a method for omitting intermediate computations in some input–output relation. Thus, a compiled model is an execution architecture computing the original mapping, but doing so in a more efficient way. This approach will now be discussed in more detail.

When the decision maker tries to learn the utility function $u : \mathcal{P} \times \mathcal{A} \longrightarrow \mathfrak{R}$, a case (p, a, u) can be considered as an example of the form (x, u), where the input $x = (p, a)$ is a problem–act pair and the outcome $y = u(p, a)$ is a utility degree. Thus, learning the utility function fits the framework of *supervised* machine learning. Still, a case (p, a, u) can also be interpreted in a different way, namely as a *valued example*. That is, $(x, y) = (p, a)$ is an example and u an evaluation thereof. The target function is now the (optimal) decision function $\Delta^* : \mathcal{P} \longrightarrow \mathcal{A}$. Roughly speaking, the utility $u = u(p, a)$ indicates the quality of an associated example (p, a).

The compiled model approach thus necessitates an extension of standard (supervised) learning methods which takes the valuation u of an example (p, a) into account. To this end, we shall fall back on the idea of "satisficing" as discussed above in connection with the model of CBDT: A "satisficing" decision maker discriminates between only two types of decisions, namely acceptable and non-acceptable ones. As will be seen below, the problem of inducing a (satisficing) decision model thus comes close to the standard setting of supervised learning.

7.7.2 Satisficing decision trees

In this section, we are going to propose a concrete approach to learning compiled decision models, namely a modified version of decision tree induction. Thus, the idea is to implement the decision function $\Delta : \mathcal{P} \longrightarrow \mathcal{A}$ as a decision tree[38] resp. a set of *condition–action rules* that yields as a classification the action to be chosen, given a new problem (condition). This approach is adequate if the set $\mathcal{A}$ of available acts is relatively small. Moreover, it assumes that problems are represented by attribute–value pairs with discrete-valued attributes.[39] Before presenting our approach, we give a brief introduction to decision tree learning.

The basic principle underlying most decision tree learners, well-known examples of which include the ID3 algorithm [304] and its successors C4.5 and C5.0 [306] as well as the CART system [55], is that of partitioning the set of given examples, S, in a recursive manner. Each inner node η of the decision tree defines a partition of a subset $S_\eta \subseteq S$ of examples assigned to that node. This is done by classifying elements $s \in S_\eta$ according to the value of a specific attribute T. The attribute is selected according to a measure of effectiveness in classifying the examples, thereby supporting the overall objective of constructing a small tree.

A widely applied "goodness of split" measure is the *information gain*, $G(S,T)$, which is defined as the expected reduction in "impurity" (with regard to the class distribution) which results from partitioning S according to T:

$$G(S,T) = I(S) - \sum_t \frac{|S_t|}{|S|}\, I(S_t), \tag{7.58}$$

where S_t denotes the set of elements $s \in S$ whose value for attribute T is t. Moreover, $I(\cdot)$ is a measure of impurity, such as the GINI function [55]

$$I(S) = \sum_{c \neq c' \in \mathcal{C}} q_c q_{c'} = 1 - \sum_{c \in \mathcal{C}} (q_c)^2 \tag{7.59}$$

with q_c the proportion of elements $s \in S$ having class c. Besides, a number of alternative (im)purity measures, such as entropy, have been devised. See [268] for an empirical comparison of splitting measures.

Suppose a set $\mathcal{X}$ of instances to be given, where each instance is characterized by several attribute values. Moreover, each $x \in \mathcal{X}$ belongs to one class $c = \text{class}(x) \in \mathcal{C}$. Given a set of training samples $S = \{(x_1, c_1), \ldots, (x_n, c_n)\} \subseteq \mathcal{X} \times \mathcal{C}$, the basic ID3 algorithm derives a decision tree as follows:

– The complete set of training samples, S, is assigned to the root of the decision tree.

[38] Decision tree learning is actually a classification method. Even though classifications can be considered as decisions, it is not specifically used for decision making in the proper sense. Therefore, one might prefer the alternative terms discimination or classification tree.

[39] Continuous-valued attributes can be discretized before or during the learning of a decision tree [93].

– A node η becomes a leaf (answer node) of the tree if all associated samples S_η belong to the same class or if all attributes have already been used along the path from the root of the tree to η.

– Otherwise, node η becomes a decision node: It is split by partitioning the associated set S_η of examples. This is done by selecting an attribute as described above and by classifying the samples $s \in S_\eta$ according to the value for that attribute. Each element of the resulting partition defines one successor node.

Once the decision tree has been constructed, each path can be considered as a rule. The precedent part of a rule is a conjunction of conditions of the form $T = t$, where T is an attribute and t a specific value thereof. The consequent part determines a value for the class variable. New examples are then classified on the basis of these rules, i.e., by looking at the class label of the leaf node whose attribute values match the description of the example. Notice that a unique class label is associated with each answer node if the data is not noisy and, hence, the original sample does not contain any *clashes* (cases with identical attribute vectors but different classes). Otherwise, the distribution of class labels at the leaf can be used for deriving a probabilistic estimate. Quite often, the induced tree undergoes further (post-)processing [267]. Here, the objective is to *prune* large trees in order to guarantee transparency. Moreover, pruning counteracts the problem of overfitting.

An incremental decision tree algorithm has been proposed in [377]. Given the same training data, this algorithm induces the same tree as ID3. Now, however, instances are processed in a serial way, that is, the current decision tree is updated each time a new example arrives. Since algorithmic aspects are not our main concern, we refrain from describing the algorithm here. It should be noted, however, that an incremental approach to learning has considerable advantages, especially in the context of decision making. In fact, the idea of learning and improving performance over time is one of the major aspects of case-based or experience-based decision making.

In this connection let us also mention a method that combines decision tree learning and lazy learning [155]: Given a set of training data, new instances are classified by means of a decision tree. However, a new tree is built for each individual instance (as in lazy learning, the complete data thus needs to be stored). Loosely speaking, this algorithm induces decision trees which are optimal for the individual instances, whereas a usual decision tree is good *on average*. The algorithm is efficient due to the fact that actually only one path of a tree is constructed, namely the one needed to classify the new instance. Besides, computational efficiency is improved by means of a caching scheme.

Let $\mathcal{P} = T_1 \times T_2 \times \ldots \times T_m$ be a set of potential problems, where $T_\imath$ denotes the (finite) domain of the $\imath$-th attribute. Thus, each problem $p \in \mathcal{P}$ is represented as a vector $p = (t_1, \ldots, t_m)$ of attribute values. Moreover, let $\mathcal{A} = \{\alpha_1, \ldots, \alpha_k\}$ be a set of available actions. Finally, utility degrees are again measured on the real number line.

Assume a memory of cases

$$\mathcal{M} = \big\{(p_1, a_1, u_1), \ldots, (p_n, a_n, u_n)\big\} \in \mathcal{P} \times \mathcal{A} \times \mathfrak{R}$$

to be given, where $u_\imath = u(p_\imath, a_\imath)$ is the utility that has resulted from applying act $a_\imath$ to problem $p_\imath$. Let $\mathcal{M}_\mathcal{P}$, $\mathcal{M}_\mathcal{A}$, and $\mathcal{M}_{\mathcal{P}\times\mathcal{A}}$ denote the projection of $\mathcal{M}$ to $\mathcal{P}$, $\mathcal{A}$, and $\mathcal{P} \times \mathcal{A}$, respectively. The idea pursued here is to compile this case base into a decision tree which is then used for solving future decision problems.

Let $u^* \in \mathfrak{R}$ be a utility threshold defined by the decision maker. This threshold corresponds to the aspiration level in the CBDT model of Gilboa and Schmeidler: An action a is *acceptable* for a problem p if $u(p, a) \geq u^*$, and it is not acceptable if $u(p, a) < u^*$.

In a first step, each case $(p_\imath, a_\imath, u_\imath)$ is changed into an *example* $(p_\imath, a_\imath)$. The latter is called a *positive example* if $u_\imath \geq u^*$ and a *negative example* if $u_\imath < u^*$. In a second step, a *modified memory*, S^*, is derived from $\mathcal{M}$. For each problem $p \in \mathcal{M}_\mathcal{P}$ it contains a *generalized example* (p, A_p), where A_p denotes the set of *feasible acts* for problem p. This set is defined as follows:

$$a \in A_p \Leftrightarrow \begin{cases} u(p,a) = u_{max}(p, \mathcal{M}) & \text{if } \; u_{max}(p, \mathcal{M}) \geq u^* \\ (p,a) \notin \mathcal{M}_{\mathcal{P}\times\mathcal{A}} & \text{if } \; u_{max}(p, \mathcal{M}) < u^* \end{cases},$$

for all $a \in \mathcal{A}$, where

$$u_{max}(p, \mathcal{M}) \stackrel{\text{df}}{=} \max_{(p,b) \in \mathcal{M}_{\mathcal{P}\times\mathcal{A}}} u(p, b).$$

In plain words, an action a is feasible for p if it belongs to the best among the actions known to be acceptable for p, or if no acceptable action is known and a has not been tried as yet. Notice that $A_p = \emptyset$ if all available acts have been applied to p but none of them was acceptable, that is, if an acceptable act for p does actually not exist. In this case, the decision maker should reduce the utility threshold u^*.[40] Subsequently, we assume $A_p \neq \emptyset$ for all $p \in \mathcal{P}$.

Before proceeding, let us point to a meaningful weakening of the above feasibility condition $u(p, a) = u_{max}(p, \mathcal{M})$. In fact, this condition could be replaced by $u(p, a) \geq u_{max}(p, \mathcal{M}) - \varepsilon$, where $\varepsilon \geq 0$ is a tolerance threshold. Here, the idea is that an action is acceptable even if its utility is slightly below the utility of the best (known) action. Of course, a decision maker being less ambitious in this sense will usually be able to induce simpler decision functions, i.e., to gain efficiency at the cost of decision quality.

We are now ready to formulate a generalized version of the decision tree learning problem whose objective is to induce a decision tree that prescribes, for any

[40] It would also be possible to maintain individual thresholds for the problems.

problem, an acceptable action:[41] Given a set of generalized examples

$$S^* = \{(p_1, A_{p_1}), (p_2, A_{p_2}), \ldots, (p_n, A_{p_n})\} \subseteq \mathcal{P} \times 2^{\mathcal{A}}, \tag{7.60}$$

induce a decision tree which implements a decision function $\Delta : \mathcal{P} \longrightarrow \mathcal{A}$ such that

$$\forall\, p \in \mathcal{M}_{\mathcal{P}} \,:\, \Delta(p) \in A_p.$$

In this problem, splitting a set of examples (7.60) is no longer necessary if

$$A(S^*) = \bigcap_{\imath=1}^{n} A_{p_\imath} \neq \emptyset, \tag{7.61}$$

hence, (7.61) defines a natural stopping condition. The corresponding node η in the decision tree then becomes a leaf, and any action $a \in A(S^*)$ can be chosen as the prescribed decision a_η associated with that node.

The main modification concerns the "goodness of split" measure. Let $G(\cdot)$ denote the information gain (7.58) as used for classical decision tree learning. That is, $G(S,T)$ quantifies the quality of the split of a (standard) sample S induced by the attribute T. Now let the class of *selections*, $\mathcal{F}(S^*)$, of the generalized set of examples (7.60) be given by the class of samples

$$\{(p_1, a_1), (p_2, a_2), \ldots, (p_n, a_n)\} \subseteq \mathcal{P} \times \mathcal{A}$$

such that $a_\imath \in A_{p_\imath}$ for all $1 \leq \imath \leq n$. We extend the measure $G(\cdot)$ to generalized samples S^* as follows:

$$G(S^*, T) \stackrel{\text{df}}{=} \max_{S \in \mathcal{F}(S^*)} G(S, T). \tag{7.62}$$

As can be seen, the extended measure (7.63) is the ordinary measure obtained for the most favorable instantiation of the generalized examples (p, A_p) and hence defines a "potential" goodness of split. It corresponds to the "true" measure that would have been derived for the attribute T if this instantiation is compatible with the ultimate decision tree. Taking this optimistic attitude is clearly justified since the tree is indeed constructed in an (apparently) optimal manner (hence averaging would hardly make sense).

Computing (7.62) comes down to solving a combinatorial optimization problem, namely to finding

$$I(S^*) \stackrel{\text{df}}{=} \min_{S \in \mathcal{F}(S^*)} I(S) \tag{7.63}$$

for (sub-)sets S^* of extended examples, where $I(\cdot)$ is a measure of impurity. It might hence be regarded as critical from a time complexity point of view,

[41] An alternative approach would be to learn, for any action, the class of decision problems to which it can be applied. This type of problem fits into the framework of multi-label classification in machine learning. However, it does not provide efficient condition-action rules.

especially in the context of real-time decision making. One should keep in mind, however, that not the construction (or revision) of the decision tree is time critical but rather its application. In fact, real-time decision making must not be confused with real-time learning; as in other decision support systems, learning will rather be realized as an "off-line" procedure [64].

In Appendix G, we present a heuristic search method for computing (7.63) based on a branch & bound technique. The efficiency of the method depends critically on the number of actions (which is the maximal branching factor of the search tree) and the average size of the sets A_p (which determines the average branching factor). Even without a detailed complexity analysis, experience has shown that no problems occur if the number of actions is small. For example, for six actions and sample sizes up to $n = 500$ the generalized splitting measure can be computed within the bounds of seconds. Still, if the number of actions is too large, the exact computation of (7.63) becomes intractable.

As an alternative we therefore suggest the following heuristic approximation of (7.63):

$$I(S^*) = I(S_*), \tag{7.64}$$

where the selection $S_* \in \mathcal{F}(S^*)$ is defined as follows: Let $q_\imath$ be the frequency of the action $\alpha_\imath$ in the set of examples S^*, i.e., the number of examples $(p_\jmath, A_{p_\jmath})$ such that $\alpha_\imath \in A_{p_\jmath}$. The $\alpha_\imath$ are first "preferentially ordered" according to their frequency $q_\imath$ (ties are broken by coin flipping), starting with the most frequent one. Then, the most preferred action $\alpha_\imath \in A_{p_\imath}$ is chosen for each example $p_\imath$. Clearly, the idea underlying this selection is to make the distribution of labels as skewed (non-uniform) as possible, since distributions of this type are favored by the impurity measure. In [204], we found that the measure (7.64) yields very good results in practice and compares favorably with alternative extensions of splitting measures.

Let us finally mention that the adequacy of a decision tree as a representation of the decision function Δ does of course depend on the structure of the optimal decision function Δ^*. In fact, since a decision tree, at least in its standard version, partitions the problem (attribute) space $\mathcal{P}$ by means of axis-parallel decision boundaries, good results (in terms of both complexity and accuracy) are to be expected only if Δ^* is at least approximately compatible with this type of inductive bias.

7.7.3 Experimental evaluation

In order to get an idea of how the satisficing decision tree approach performs we have employed a procedure that generates decision problems in a systematic way. A *decision environment* is specified by the following parameters:

– The number m of attributes describing a decision problem.

– The number k of possible actions.

– The number of values for each attribute. For the sake of simplicity, we assume that each attribute has the same number v of values. Without loss of generality these values are represented by natural numbers, that is $T_\jmath = \{1, 2, \ldots, v\}$ for all $1 \leq \jmath \leq m$.

– The utility function u that assigns a utility degree $u(p, a)$ to each problem $p \in \mathcal{P}$ and action $a \in \mathcal{A}$. The generation of u is realized by a random procedure which is under the control of a complexity parameter γ, as described below.

It has already been mentioned that the adequacy of a decision tree representation depends on the structure of the decision function Δ^*. In fact, Δ^* can be represented by small trees if its structure is in agreement with a decision tree-like partitioning of the feature space; otherwise, the decision tree model might become rather complex. In order to control the complexity of the decision environment we have generated a utility function u as follows: In a first step, an optimal decision tree is generated at random. This is done in a recursive manner by starting with the root of the tree and deciding for each node whether it is an inner node or a leaf of the tree. The probability of a node to become an inner node is specified by a parameter $0 < \gamma < 1$. Each inner node at level $\imath$ has v successors, each of which corresponds to a value of the $\imath$-th attribute. If the tree has been generated, each leaf η covers a subset $\mathcal{P}_\eta$ of the set of problems $\mathcal{P}$, namely those problems which match the attribute values associated with η. The leaf node η is then assigned an optimal decision a_η at random. From the resulting optimal decision tree, the utility function is finally derived by letting $u(p, a) = 1$ if a is the optimal solution to p, i.e., if $a = a_\eta$, where η is the leaf node that covers p. For all other (suboptimal) actions b the utility $u(p, b)$ is defined as a decreasing function of the distance between a and b (where the distance between act $a_\imath$ and act $a_\jmath$ is $|\imath - \jmath|$). Note that the complexity is completely determined by the parameter γ: The larger γ, the larger the expected size of the optimal decision tree, i.e., the more complex the decision environment (at least for an agent that employs a decision tree representation of its decision model).

After having specified a concrete decision environment by generating the utility function u at random, a simulation experiment is performed as follows: At the beginning, the memory of the decision maker is empty, and its (satisficing) decision tree corresponds to a single node. In the $\imath$-th decision epoch, a decision problem $p_\imath$ is chosen at random from $\mathcal{P}$, according to a uniform distribution. For this problem, the decision maker selects an action $a_\imath$ according to the current decision tree model (ties are broken by coin flipping). The new case, consisting of the problem $p_\imath$, the act $a_\imath$, and the experienced utility $r_\imath = u(p, a)$, is added to the memory. Moreover, the satisficing decision tree is updated whenever necessary. The simulation stops after the L-th decision epoch.

Illustrating example. To illustrate, we present a simple example step by step. Let $m = 6$, $k = 3$, $v = 3$, $L = 10$, $u^* = 0.7$ and consider the sequence of decision problems in Table 7.1 (and disregard, for the time being, the actions and utility degrees shown in the same table).

problem						act	utility
t_1	t_2	t_3	t_4	t_5	t_6		
2	1	2	1	1	2	α_1	1
2	1	2	3	3	1	α_1	1
1	1	2	3	2	3	α_1	0.2
1	3	3	3	3	3	α_2	0.8
1	1	3	2	1	3	α_2	1
2	2	1	2	1	3	α_1	0.9
3	2	2	2	2	3	α_1	0
3	1	2	2	3	2	α_2	0.5
3	1	1	3	3	1	α_3	0
1	3	2	1	1	2	α_1	0.8

Table 7.1. Sequence of decision problems specified by the values of six attributes (columns 1–6), the action performed by the decision maker, and the resulting utility degree.

For the first decision problem, the agent chooses an action at random. As shown in Table 7.1, this is action α_1, for which it receives a utility of 1. Thus, the agent generates a decision tree which corresponds to the following rules:

t_1	t_2	t_3	t_4	t_5	t_6	action
?	?	?	?	?	?	α_1

This tree prescribes action α_1 regardless of the attribute values. Thus, for the second problem the agent chooses again α_1, and again obtains a utility of 1. Consequently, it does not modify the decision tree. For the third problem, however, α_1 yields a utility of 0.2 which falls below the utility threshold u^*. Therefore, the agent changes the decision tree according to the procedure outlined above:

t_1	t_2	t_3	t_4	t_5	t_6	action
1	?	?	?	?	?	α_2
2	?	?	?	?	?	α_1

This tree prescribes to choose α_1 if the value of the first attribute is 2, but α_2 if this value is 1. Thus, the agent's hypothesis is that α_1 yields bad outcomes if $t_1 = 2$ (note that t_5 and t_6 might have been chosen as splitting attributes as well). The next update occurs after the 7-th problem. Since the decision tree does not prescribe an action for $t_1 = 3$, the agent chooses α_1 at random. This leads to a utility of 0. The new decision tree contains the following rules:

t_1	t_2	t_3	t_4	t_5	t_6	action
1	?	?	?	?	?	α_2
2	?	?	?	?	?	α_1
3	?	?	?	?	?	α_2

This tree is changed into

t_1	t_2	t_3	t_4	t_5	t_6	action
1	?	?	?	?	?	α_2
2	?	?	?	?	?	α_1
3	?	?	?	?	?	α_3

after the 8-th problem, since α_2 yields $u = 1/2 < u^*$ for a problem with $t_1 = 3$. Notice that, so far, the decision is completely determined by the first attribute. After the 9-th problem, however, the agent realizes that this is not enough. The new decision tree also involves attribute t_6:

t_1	t_2	t_3	t_4	t_5	t_6	action
1	?	?	?	?	3	α_2
2	?	?	?	?	3	α_1
3	?	?	?	?	3	α_2
?	?	?	?	?	1	α_1
?	?	?	?	?	2	α_1

General findings. More generally, we were interested in effects of the complexity of the decision environment and of the aspiration level of the decision maker. Regarding the first factor, more complex decision environments are expected to entail larger decision trees and smaller average utilities over time. As concerns the aspiration level, it is to be expected that higher levels will probably guarantee higher utilities on average but, at the same time, lead to more complex decision models. To illustrate, consider the problem of choosing the optimal dose of a drug for different patients. The simple decision tree shown in Fig. 7.1 might lead to satisfying results (the utility of a decision depends on the patient's state of health after the treatment). Still, even better results might be obtained by differentiating more precisely between patients, taking further attributes such as weight into account. This would of course mean using a more complex decision tree.

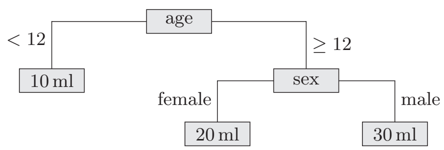

Fig. 7.1. Decision tree implementing a simple strategy for choosing the dose of a drug.

One might furthermore suspect that the effect of increasing the threshold u^* is not independent of the complexity of the decision environment. Consider the following example: Suppose that exactly one optimal action with utility 1 exists for each problem. An act applied to a problem for which it is not optimal yields a utility of $0 < \alpha < 1$. Now, with a threshold $u^* \leq \alpha$ the decision maker is always satisfied, regardless of what action it applies to a problem. In fact, its decision model consists of only one rule which prescribes to choose act a, where a is the act that has been applied to the first problem, perhaps randomly. The average utility is exactly $\alpha + (1-\alpha)/k$. If the utility threshold u^* exceeds α, the decision maker is satisfied only with the optimal acts, and it will spend enormous effort in finding these acts. The difficulty of this venture in turn depends on the generalization capability of the induced decision trees. If a decision tree is indeed a good model for the application at hand, the agent might succeed very quickly. Otherwise, it might try several actions for each individual problem before finally finding the optimal one it seeks for. Anyway, the decision model will become much more complex in this case. Of course, this model will finally come up with an optimal act for each problem. It should be noted, however, that it might take a long time and many unsuccessful attempts before this model is constructed. Therefore, the gain in utility might be poor over a limited time horizon and might hence not compensate for the increased complexity.

For different combinations of utility thresholds u^* and complexity measures γ, we have performed 1,000 simulation runs with $m = 6$, $k = 4$, $v = 4$, $L = 100$, respectively. For each simulation, the average utility $(r_1 + \ldots + r_{100})/100$ and the average size of the decision tree have been computed. (The size was measured in terms of the number l of leaf nodes.) The corresponding results, documented in Table 7.2 and Appendix H, permit the following conclusions which confirm our above suppositions: Increasing u^* always leads to more complex decision models. However, it yields an improvement in average utility only if the decision environment is not too complex. Roughly speaking, the decision maker should be demanding for simple environments, where decision trees provide an adequate model and, hence, looking for better decision models is likely to be successful. If the environment is complex, however, it is urged to be modest: Searching for better models will generally increase the size of decision trees but hardly the quality of decisions.

	$\gamma = 0.5$	$\gamma = 1$
$u^* = 0.2$	0.59	0.59
$u^* = 0.9$	0.79	0.56

	$\gamma = 0.5$	$\gamma = 1$
$u^* = 0.2$	8.84	15.76
$u^* = 0.9$	19.82	48.79

Fig. 7.2. Average values of the (average) utility degrees (left) and (average) number of leaf nodes (right), taken over the 1,000 simulation runs.

7.8 Summary and remarks

Summary

– We have briefly reviewed the original idea of case-based decision making due to Gilboa and Schmeidler (Section 7.1) as well as an alternative (fuzzy set-based) model proposed by Dubois and Prade (Section 7.3). Rather than concentrating on the accumulated or average performance of acts, the latter gives preference to acts which have always led to good results for problems which are similar to the current one.

– Methods of CBDM on the basis of the Nearest Neighbor principle have been investigated and characterized axiomatically in Section 7.2. NN decision rules can be seen as approximations of the decision criteria in CBDT. They can be motivated, among other things, for reasons of computational efficiency.

– The fuzzy set-based approach to CBDM has been generalized in Section 7.4. The extreme (worst case) valuation in the original model has been relaxed by looking out for acts which have yielded good results at least in *most* (rather than all) cases in the past. It has been shown that the relaxation of the "always" requirement in the principle underlying the original decision criterion can be advantageous in the context of repeated decision making.

– Section 7.5 has outlined an alternative CBDM framework. Corresponding methods combine results of previous chapters and generalized decision theories which have recently been proposed in literature in order to realize case-based decision making. These methods are *case-based* in the sense that an agent makes use of case-based reasoning (in the form of case-based inference) in order to support the modeling of a new decision problem, notably the specification of an uncertainty measure over possible outcomes. Since the latter is not necessarily a probability measure, the concept of an expected utility has to be generalized in order to compare acts. Two concrete methods have been discussed: The CBI approaches of Section 4.5 and Chapters 5 and 6 give rise to decision making with "case-based" belief functions and "case-based" possibility distributions, respectively.

– In Section 7.7, we introduced a framework of *experienced-based decision making* as an extension of case-based decision making. In EBDM, an agent faced with a new decision problem acts on the basis of experience gathered from previous problems in the past, either through predicting the utility of potential actions or through establishing a direct relationship between decision problems and appropriate actions. A realization of the latter approach has been proposed in the form of "satisficing decision trees".

Remarks

- A representation of cases which is similar to the one proposed by Gilboa and Schmeidler was already suggested by Kolodner [234]. Apart from a problem and a solution she introduced a third component of a case: the *outcome* is thought of as the state of the world under the condition that the corresponding solution is applied and usually comprises some kind of feedback (see also [30]). According to this point of view, a triple (p, a, r) is seen as an extended decription of a case, i.e., a usual case (p, a) supplemented by some valuation r. By using the notation $\langle (p, a), r \rangle$ we have suggested a second interpretation in this chapter: p and a are taken together and constitute the first component of an *ordinary* case. This component is now partly under the control of the agent which can choose a. The second component is the outcome associated with the problem–act tuple (p, a). Even though formally equivalent to the first notation, considering a case as a tuple $\langle (p, a), r \rangle$ seems more natural in the context of Section 7.5 where case-based reasoning (case-based inference) is used in its basic form, namely for predicting the outcomes associated with inputs (= problem–act tuples).
- In [172], Gilboa and Schmeidler provide an interesting comparison between case-based and rule-based knowledge representation, with special emphasis on the problem of induction. This article also contains further examples showing that the linearity of the CBDT functionals will often be too restrictive in practice. Particularly, this seems to be true if the decision maker is allowed to *learn* a similarity function resp. the importance of cases.[42] For instance, if experience is better represented by *subsets* of cases, the weight of an individual case depends on other observations as well. This effect, however, cannot be captured by the (additively) separable CBDT functionals but rather calls for the use of non-additive set-functions.
- As the summation of (weighted) degrees of *utility* in (7.2) reveals, CBDT actually assumes that the application of an act to similar problems yields similar utilities rather than similar outcomes. Of course, the two principles are only equivalent if outcomes are directly given in terms of utilities. Otherwise, the use of utility degrees in (7.2) has to be justified by the additional assumption that similar outcomes have similar utilities.
- The memory (7.1) of cases represents the experience of the decision maker. This does not mean, however, that all cases have been collected by the agent itself, or that the agent has made all related decisions. In fact, cases can also be experienced in a passive way or might even be the product of some kind of hypothetical reasoning.
- The simple accumulation of utility degrees in (7.2) does not always appear plausible, of course. Let us mention, however, that it might well be reasonable in connection with certain applications, such as the modeling of consumer behavior

[42] The adaptation of the similarity function is interpreted as some kind of *second-order induction* in [172].

in economics [170]. Interestingly enough, some undesirable effects of the accumulative nature of (7.2) can also be avoided by using the more general approach (7.5): The relevance of an observation might be reduced if the same case has already been encountered before. This idea seems quite plausible from a cognitive point of view. In fact, it again reveals the advantage of a "relevance-based" decision theory which is more general than a "similarity-based" approach.

- It has been mentioned that CBDT should not be seen as a competing theory, but as an alternative model which complements expected utility theory in a reasonable way. The claim that neither of them is superior in general and that the adequacy of a model strongly depends on the kind of problem under consideration is supported by a theoretical result of Matsui [262]. He shows that EUT and (a slight modification of) Gilboa and Schmeidler's CBDT are equivalent in the sense that each EUT model can be represented in the framework of CBDT and vice versa.[43] The embedded model, however, might be much more complex than the original model.
- Notwithstanding the cognitive appeal of CBDT, one might feel some uneasiness concerning the manifold possibilities for defining a case-based decision model. CBDT basically suggests that the current decision is a function of the agent's *experience*, considered against the background of a similarity relation between inputs. The experience, as represented by the history of cases, is an element of a quite complex and high-dimensional space on which various decision functions can be defined. Moreover, similarity is a rather vague concept, and it is by no means obvious what a reasonable similarity function should look like. In this respect, expected utility theory appears more restrictive. In fact, a decision is derived from a utility function[44] and a probability function which can be seen as an (information-compressed) *statistic* of the agent's experience (at least if probabilities are obtained from relative frequencies). Besides, the linear combination of probability and utility by means of the expected utility formula seems more straightforward than a similarity-based evaluation. Loosely speaking, EUT determines the information to be extracted from the agent's experience and the way in which this information is to be used more strictly.
- We have assumed the nearest neighbor (7.11) to be unique. The case of non-uniqueness could be handled by means of a set-valued generalization in the Dempster-Shafer style. Then, (7.11) defines the set of nearest neighbors, thus playing a role somewhat comparable to a focal element of a belief structure over $\mathcal{A}$. (Observe, however, that the weights in (7.10) are utility degrees which do not necessarily sum up to 1.) Moreover, (7.10) becomes

[43] Consequently, the two theories are *observationally* equivalent.

[44] Note that a utility function is principally required in CBDT as well. There, however, the function needs to be known only partially, namely for the observed outcomes.

$$V(a_0) = \sum_{(p,a) \in \mathcal{M}^{\downarrow} : \mathsf{NN}_{p_0,\mathcal{A}}(p,a) \ni a_0} \sigma_{\mathcal{Q}\times\mathcal{A}}((p,a),(p_0,a_0)) \cdot u(r(p,a)),$$

i.e., $V(a_0)$ defines the counterpart to the plausibility of a value $a_0 \in \mathcal{A}$.

- In the context of CBDM, the decision maker treats an uncertainty measure derived via CBI as some kind of "intermediate result" of the complete decision procedure. In Section 7.5.3, for instance, the possibility distributions (7.46) are taken as primitives in the second step of this procedure, namely the ranking of acts according to (7.49) or (7.50). In order to apply these qualitative decision criteria, the agent has to consider the distributions as being objectively given. In fact, the axiomatic framework in [123] is set up in the style of Von Neumann and Morgenstern: A utility function is derived from preferences, but the concept of belief in the form of a possibility distribution on outcomes is assumed to be given.[45] The two-stage procedure realized by CBDM might appear vulnerable from this point of view, particularly since the meaning of objectivity seems less obvious in the case of a possibility distribution than in the case of a probability [194].
- The idea of relating similarity and uncertainty (cf. Section 7.6.1) is also realized in the theory of counterfactuals proposed by Lewis [250], where the plausibility of an imaginary input is determined by its similarity to the current input.
- A combination of the concepts of similarity, preference (utility), and belief (probability) has also been outlined in [319]. However, this approach is quite different from the ideas discussed in this chapter. Particularly, it is not related to case-based reasoning.
- In connection with NN decision rules (Section 7.2) it has been mentioned that a decision maker will generally not utilize its complete memory when having to perform a prompt action. This consideration reveals the importance of efficient memory organization and case retrieval strategies. Needless to say, a computationally efficient (and cognitively plausible) case-based decision theory has to take these aspects into account.
- The methods proposed in Section 7.5 are based on *generalizations* of expected utility theory. Let us mention that one could also think of other ways of combining case-based inference and EUT. The constraint-based approach to CBI discussed in Chapter 3, for instance, can be used in order to suggest a subset of acts or states of nature which should be taken into account. EUT can then be applied to the reduced setup. Not only is an approach of this kind computationally efficient, it also appears cognitively plausible. In fact, human decision makers will generally concentrate on a small number of acts and disregard states of nature which are considered as being impossible anyway.
- The property of bounded optimality mentioned at the beginning of this chapter can be paraphrased as "the optimization of computational utility given a set of

[45] See [128] for an axiomatization of qualitative decision making in the style of Savage.

assumptions about expected problems and constraints in reasoning resources" [192]. According to [323], a program exhibits bounded optimality if it "is a solution to the constrained optimization problem presented by its architecture." A relaxation of this concept is *asymptotic* bounded rationality [323] which to some extent parallels the idea of asymptotic complexity. It aims at supporting a constructive theory of bounded rationality which makes the design of bounded optimal agents largely independent of the architecture of the computational environment.

The fact that computational aspects of rational decision making have only recently become a focus of research should not give rise to the impression that related problems have been ignored before. Indeed, classical decision theory has well been aware of computational problems [255]. See, for instance, [182] for a generalization of the axioms of subjective probability taking related aspects into account.