Extensions of a qualitative approach to case-based decision making: Uncertainty and fuzzy quantification in act evaluation

Eyke Hüllermeier, Didier Dubois, Henri Prade
Institut de Recherche en Informatique de Toulouse
Université Paul Sabatier, France
{eyke,dubois,prade}@irit.fr

Abstract

We extend a recently proposed qualitative model of case-based decision making which can be seen as a "case-based" counterpart of (a possibilistic generalization of) the MAXIMIN decision rule. The idea underlying this approach is to give preference to acts which have always led to good results for problems which are similar to the current one. The extensions of the model can be motivated, among other things, by the idea of repeated decision making, which arises quite naturally within the context of case-based reasoning where new cases are encountered and, hence, experience is accumulated over time. Firstly, an agent will generally have made very few, if any, observations at the beginning of a decision sequence. This lack of experience does inevitably cause problems for a case-based approach to decision making which will always seem more or less arbitrary and, hence, will be open to criticism in such situations. We approach this problem by allowing for some kind of "hypothetical" reasoning in connection with a generalized evaluation of acts which allows for the representation of uncertainty. Secondly, the case-based valuation of acts according to the above-mentioned model appears to be rather drastic since it is based on a sort of worst case evaluation. We relax it by looking for acts which have yield good results at least in *most* cases in the past. Indeed, the original rule which requires good results in any (similar) case may lead to undesirable consequences if an agent has to act repeatedly. Even though qualitative decision rules are generally not intended for maximizing some kind of average performance, it is shown that the relaxation of the "always" requirement in the principle underlying the original decision criterion can be advantageous.

Keywords: decision theory, case-based reasoning, similarity, possibility theory, fuzzy quantifier, repeated decision making.

1 Introduction

Recently, Gilboa and Schmeidler [20] have proposed a case-based approach to decision making, called *case-based decision theory* (CBDT), which combines principles from decision theory and case-based reasoning (CBR). The main idea underlying their approach is the application of the CBR-principle to the problem of decision making. More precisely, they assume that an agent faced with a certain decision problem relies upon his experience from similar situations he encountered in the past. Loosely spoken, he chooses among (potential) acts based on the (cumulative or average) performance of these acts in connection with previous problems which are similar to the current one.

Putting it in a nutshell, the problem of case-based decision making (CBDM) 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 (finite) set of results or outcomes. We assume that choosing act $a \in \mathcal{A}$ for solving problem $p \in \mathcal{Q}$ leads to the outcome $r = r(p,a) \in \mathcal{R}$.[1] A utility function $u : \mathcal{R} \rightarrow U$ assigns utility values to such outcomes. Let $\sigma_{\mathcal{Q}} : \mathcal{Q} \times \mathcal{Q} \rightarrow [0,1]$ and $\sigma_{\mathcal{R}} : \mathcal{R} \times \mathcal{R} \rightarrow [0,1]$ be (reflexive and symmetric) similarity measures quantifying the similarity of problems and results, respectively. Suppose the decision making agent to have a (finite) memory $\mathcal{M} = \{(p_1,a_1,r_1),\ldots,(p_n,a_n,r_n)\}$ of cases at his 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 that he has to choose an act for a new problem $p_0 \in \mathcal{Q}$. If a certain act $a \in \mathcal{A}$ has not been applied to the problem p_0 so far (i.e., there is no case $(p_0,a,r) \in \mathcal{M}$) the agent will generally be uncertain about the result $r(p_0,a)$ and, hence, about the utility $u(r(p_0,a))$. According to the assumption underlying the paradigm of CBDM, he then evaluates an act based on

[1] The assumption that a problem/act tuple determines a unique outcome will be relaxed in Section 3.

its performance in similar problems in the past, as represented by (parts of) the memory $\mathcal{M}$.

In their original approach, Gilboa and Schmeidler assume the decision maker to choose an act maximizing the following valuation of an act $a \in \mathcal{A}$:

$$V(a) = V_{p_0,\mathcal{M}}(a) = \sum_{(p,a,r)\in\mathcal{M}} \sigma_{\mathcal{Q}}(p,p_0) \cdot u(r). \tag{1}$$

The summation over an empty set yields the "default value" 0 which plays the role of an "aspiration level." Alternatively, an "averaged similarity" version[2] of the linear functional (1) has been proposed, which results from replacing $\sigma_{\mathcal{Q}}$ in (1) by

$$\sigma(p,p_0) = \frac{\sigma_{\mathcal{Q}}(p,p_0)}{\sum_{(p',a,r')\in\mathcal{M}} \sigma_{\mathcal{Q}}(p',p_0)}. \tag{2}$$

Theoretical details of CBDT, including special assumptions concerning the memory $\mathcal{M}$ and an axiomatic characterization of decision principle (1), are presented in [20].

A main motivation behind CBDT is to provide a more faithful description of human decision making than expected utility theory (EUT) does. Indeed, in some situations this axiomatic theory seems rather restrictive. Particularly, it assumes the decision maker to have very detailed information at his disposal: a complete list of the *states of nature* with corresponding probabilities, a complete list of potential acts, and a numerical utility value for all act/state pairs. Moreover, some well-known paradoxes [1, 17] as well as psychological studies [28] show that EUT can be challenged as a *descriptive* theory of (human) decision making. A thorough discussion of the relation between CBDT and EUT can again be found in [20].

Combining principles from case-based reasoning and decision making touches on aspects of knowledge representation and reasoning which are particularly interesting from the viewpoint of artificial intelligence (AI). A closely related line of research in AI which also aims at making decision theoretical representations more realistic, tractable, and expressive is that of *qualitative* decision theory [3, 4, 5, 6, 12, 15, 18, 25, 26, 27]. In fact, the assumption that uncertainty and preference can be quantified by means of a precise probability measure and a cardinal utility function, respectively, does often appear unrealistic. First steps toward a combination of a *qualitative* and a *case-based* decision methodology have recently been developed within the framework of fuzzy sets and possibility theory [13, 8, 9, 21, 24]. As opposed to (1), such approaches only assume an ordinal setting for modelling decision problems, i.e., ordinal scales for assessing preference and similarity.

Just as (1) may be seen as a (case-based) counterpart of classical expected utility, the qualitative model to CBDM which has been proposed in [13, 8], and which will be reviewed in Section 2 of this paper, can be seen as a counterpart of a generalization of the MAXIMIN decision rule which has recently been induced in qualitative decision making under uncertainty [10]. In Section 3 and Section 4 of this paper, we shall propose two extensions of this model, namely the handling of uncertainty and a weakening of the requirement that an act should *always* have led to good outcomes. As we shall see, these extensions can be motivated, among other things, by the idea of *repeated* decision making, which arises quite naturally within the context of case-based reasoning.

2 Fuzzy modelling of CBDM

Case-based decision making has been realized in [8] as some kind of similarity-based approximate reasoning. Let $\rightarrow$ be a multiple-valued implication connective. Given a memory $\mathcal{M}$ and a new problem p_0, an act $a \in \mathcal{A}$ is assigned the (estimated) utility value

$$V^{\downarrow}_{p_0,\mathcal{M}}(a) = \min_{(p,a,r)\in\mathcal{M}} \sigma_{\mathcal{Q}}(p,p_0) \rightarrow u(r). \tag{3}$$

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. An essential idea behind (3) is to avoid the following consequences of the decision criterion (1):

- A decision maker might prefer an act a which gave always rather poor results to an act a' which has yield very good results simply because a has been tried more often. This problem, which is caused by the accumulative nature of (1), is also the main motivation for the normalization (2).

[2] This version corresponds to a special case of a k-NEAREST NEIGHBOUR approximation which is also used in other CBR approaches such as, e.g., the ELEM2-CBR system [7].

- The valuation (1) compensates between good results and bad results associated with an act a, which does not always seem appropriate (even in the context of repeated problem solving).

As a special realization of (3) the valuation

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

is proposed, where h is an order-preserving function which assures the *commensuration* between the (finite) linear scales of similarity and preference, and n is the order-reversing function of the similarity scale. Taking n as $1-(\cdot)$ in $[0,1]$, and h as the identity, we obtain

$$V(a) = V^{\downarrow}_{p_0,\mathcal{M}}(a) = \min_{(p,a,r)\in\mathcal{M}} \max\{1-\sigma_{\mathcal{Q}}(p,p_0), u(r)\}. \tag{4}$$

Besides, the criterion

$$V(a) = V^{\uparrow}_{p_0,\mathcal{M}}(a) = \max_{(p,a,r)\in\mathcal{M}} \min\{\sigma_{\mathcal{Q}}(p,p_0), u(r)\} \tag{5}$$

is introduced as an *optimistic* counterpart of (4).[3] It can be seen as a formalization of the idea of finding an act a for which there is at least one problem similar to p_0 and for which a has led to a good result.

In some situations, the extremely pessimistic and optimistic nature of the criteria (4) and (5), respectively, might appear at least as questionable as the "averaging" of (1). In fact, the compensation between good and bad experiences might just as well be reasonable (if an agent has to act repeatedly over time.) However, as already pointed out in Section 1, (4) and (5) are not intended for maximizing some kind of average performance, which hardly makes sense within an ordinal setting. Rather, these decision critera should be seen from the same point of view as qualitative decision rules such as MAXIMIN [12]. Indeed, the application of (3) seems reasonable, for instance, if an agent aims at minimizing the occurence of worst case outcomes in competition with other agents, or if only an ordinal preference relation on outcomes can be assumed [5].

Alternative formalizations of CBDM have also been proposed in [10, 21, 23]. For estimating the utility of an act in connection with a new problem, these methods make use of observed cases by more indirect means than the approaches discussed so far. More precisely, a memory $\mathcal{M}$ of cases is used for deriving a quantification of how likely a certain act will yield a certain outcome. The (case-based reasoning) hypothesis underlying these approaches is the assumption that "the more similar two problems are, the more *likely* it is that an act leads to similar results." Within the framework of [21], where *likely* means *probable*, a probability distribution on the set $\mathcal{R}$ of outcomes is derived from a memory $\mathcal{M}$. Likewise, possibility distributions are obtained in connection with the possibilistic frameworks in [10] and [23], where *likely* means *possible* and *certain*, respectively.

3 Coping with uncertainty

There are several kinds of uncertainty which might become relevant in connection with CBDM. A first source of uncertainty concerns the observed cases. A problem which occurs frequently, e.g., in connection with experimental data, is that of imprecise observations, i.e., outcomes which cannot be observed exactly. This kind of uncertainty can be taken into account, e.g., by modelling outcomes as fuzzy sets $R \in \mathbb{F}(\mathcal{R})$, where $\mathbb{F}(\mathcal{R})$ denotes the set of all (normal) fuzzy subsets of the set $\mathcal{R}$ of results.

It might also become necessary to give up the assumption that a problem $p \in \mathcal{Q}$ and an act $a \in \mathcal{A}$ determine a unique outcome $r(p,a)$. In [24], for instance, a case-based reasoning framework is considered in which results are treated as random variables. Again, there are different motivations for such a non-deterministic setting. For example, the process which determines the result associated with a problem p and an act a might indeed be subject to some random influences. A second motivation, which seems to be of considerable practical relevance, is related to the completeness, precision, and granularity of information. Even though the application of a certain act might *principally* determine the outcome, the characterization of the problem might be imprecise, incomplete, or not detailed enough. Thus, choosing an act a for repeatedly solving the (apparently) same

[3] As pointed out in [13, 8, 9] and discussed in Section 3, the above formulas are debatable if none of the already experienced problems is completely similar to the current problem.

problem p might result in different outcomes.[4] An uncertainty measure associated with a problem/act tuple (p, a) is then used for characterizing the true but unknown result. Consider the case where the description of the problem p contains missing attribute values as an example. The uncertainty concerning the outcome $r(p, a)$ is then directly related to the uncertainty concerning the values of these attributes. A further example is the problem of decision making in game playing. Namely, the outcome associated with a certain act will generally depend on the decision of the opponent as well. The latter, however, is not part of the problem description.

Here, we are particularly concerned with a second source of uncertainty which actually corresponds to a lack of information, and which is not related to the observed cases. Rather, it concerns the cases which have not been encountered so far. By this we mean the problem that a *case-based* decision procedure will inevitably get into trouble, or at least become dubious, if not enough cases have been observed. The assignment of the "default value" 0 in connection with (1), for instance, might appear somehow arbitrary. The alternative models of CBDM mentioned at the end of Section 2 seem advantageous with respect to this problem [22, 23]. The fact that no cases or, more generally, no similar cases have been observed so far can be modelled adequately by means of the possibility distribution $\pi \equiv 1$ on $\mathcal{R}$. Namely, the latter is an expression of complete ignorance, which cannot be depicted by less expressive scalar estimations such as (1) and (4).

Problems caused by a lack of information also occur in connection with (4). As pointed out in [8] this valuation only makes sense if the memory contains at least one problem p such that $\sigma(p, p_0) = 1$, and where a has been chosen for solving p. Otherwise, it may happen that (4) is very high even though none of the problems contained in the memory is similar to the current problem p_0. Particularly,

$$(\{p \in \mathcal{Q} \,|\, (p, a, r) \in \mathcal{M} \wedge \sigma_{\mathcal{Q}}(p, p_0) > 0\} = \emptyset) \rightarrow \left(V^{\downarrow}_{p_0,\mathcal{M}}(a) = 1\right),$$

which does not seem satisfactory.

Modifications of (4) and its optimistic counterpart have been proposed in order to cope with these difficulties. The modified measures are based on some kind of *normalization* of the similarity function for each act a, and a discounting of (4) and (5) which takes the absence of problems similar to p_0 into account. More precisely, the modified version of (4) is given by

$$V^{\downarrow}_{p_0,\mathcal{M}}(a) = \min\left\{h_{\mathcal{M}}(a, p_0), \min_{(p,a,r)\in\mathcal{M}} \max\{1 - \sigma'_{\mathcal{Q}}(p, p_0), u(r)\}\right\}, \tag{6}$$

where

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

and $\sigma'_{\mathcal{Q}}(\cdot, p_0)$ denotes a renormalization of $\sigma_{\mathcal{Q}}(\cdot, p_0)$ such as, e.g., $\sigma_{\mathcal{Q}}(\cdot, p_0)/h_{\mathcal{M}}(a, p_0)$ (assuming $h(a, p_0) > 0$.) The idea behind (6) 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 renormalized in order to obtain a meaningful degree of inclusion. Thus, (6) 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 (4) is retrieved from (6) as soon as $h_{\mathcal{M}}(a, p_0) = 1$.

We shall now propose a generalization of (4) which can handle the two above-mentioned sources of uncertainty in a unified way, and which is also able to express uncertainty in connection with the valuation of an act. To this end, it should first be noticed that we can write (4) as

$$V(a) = V^{\downarrow}_{p_0,\mathcal{M}}(a) = \min_{0 \le k \le m} \max\{1 - \sigma_k, v_k\}, \tag{7}$$

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

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

[4] It is, of course, possible to formally maintain the property $((p, a) = (p', a')) \rightarrow ((p, a, r) = (p', a', r'))$ for all $(p, a, r), (p', a', r') \in \mathcal{M}$. In fact, all problems might be considered as different by adding an arbitrary attribute such as, e.g., a consecutive number. Then, however, we should expect $\sigma_{\mathcal{Q}}(p, p') = 1$ if p coincides with p' up to this auxiliary attribute.

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 just leads to the problem that (4) becomes large if only few observations have been made.

According to (7), the valuation (4) of an act is completely determined by the lower bounds v_k $(0 \leq k \leq m)$, which are derived from the memory $\mathcal{M}$. This reveals that (4) can 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, (4) valuates 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

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

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

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

which defines a similarity-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 $\{p \,|\, \exists\, r \in \mathcal{R} : (p,a,r) \in \mathcal{M}\} = \mathcal{Q} \setminus \{p_0\}$.

The above considerations suggest a generalization which is obtained by replacing the scalar values v_k by fuzzy sets and applying the extension principle [30] to (7):

$$\mu_{V^{\downarrow}_{p_0,\mathcal{M}}(a)}(v) = \max\left\{ \min_{0 \leq k \leq m} \mu_{W_k}(v_k) \;\middle|\; v_1, \ldots, v_n \in \mathcal{R}, \min_{0 \leq j \leq m} \max\{1 - \sigma_j, v_j\} = v \right\}, \qquad (8)$$

where $W_k \in \mathbb{F}(U)$ $(0 \leq k \leq m)$, and $\mu_{V^{\downarrow}_{p_0,\mathcal{M}}(a)}$ denotes the membership function of $V^{\downarrow}_{p_0,\mathcal{M}}(a)$ which is now a fuzzy set. W_k represents the available information about w_k, and a value $\mu_{W_k}(v)$ is understood as the possibility that the lower bound w_k is given by v. The model (4) then corresponds to the special case where W_k is derived from $\mathcal{M}$ according to $\mu_{W_k} = \chi_{\{v_k\}}$.

The problem of uncertainty due to the absence of solved problems which are (completely) similar to p_0 can now be handled in a more flexible way. Consider, for instance, the case where $V_k = \emptyset$, i.e., no problem has been encountered so far which is σ_k-similar to p_0 and to which act a has been applied. As already mentioned above, the original approach (4) does then (implicitely) estimate the lower bound w_k by $v_k = 1$, whereas (7) is able, e.g., to express complete ignorance via $\mu_{W_k} \equiv 1$. Particularly, letting $W_m \equiv 1$ in the case where $V_m = \emptyset$ implies that $\mu_{V^{\downarrow}_{p_0,\mathcal{M}}(a)}(0) = 1$, i.e., the fact that act a should be assigned the valuation 0 seems completely possible.

The above modelling of ignorance concerning the lower bound w_k might be generalized to the case where $V_k \neq \emptyset$ by means of $\mu_{W_k}(v) = 1$ if $v \leq v_k$, and $\mu_{W_k}(v) = 0$ otherwise.[5] It seems reasonable, however, to think of more general definitions of W_k. Seen from the viewpoint of CBR, for instance, the memory $\mathcal{M}$ may provide evidence concerning w_k even if $V_k = \emptyset$. In this connection it seems particularly interesting to combine (3) with the possibilistic methods mentioned at the end of Section 2, which leads to a more "hypothetical" specification of the W_k. Suppose, for example, the CBR principle that acts lead to similar outcomes for similar problems to be strongly supported by the observations which have been made so far. Moreover, suppose that a certain act a has often led to good results for problems which are *very* (but not perfectly) similar to the problem p_0 under consideration. It seems, then, likely that a also leads to good outcomes if applied to problems which are *completely* similar to p_0. Thus, $\mu_{W_m}(v)$ should be small for small utility values v, even though a case (p,a,r) such that $\sigma_{\mathcal{Q}}(p,p_0) = \sigma_m = 1$ has not yet been encountered.

Observe that (8) also allows for the utilization of background knowledge which is not derived from the memory $\mathcal{M}$. It might be known from a further information source, for instance, that w_k does definitely not fall below a certain bound v'_k, or at least that $w_k < v'_k$ is unlikely, which leads to $\mu_{W_k} = 0$ resp. $\mu_{W_k}(v) \ll 1$ for $v < v'_k$.

Based on (8) in conjunction with a generalization of the CBDM framework outlined in Section 1 we can also approach the first type of uncertainty, which has been mentioned at the beginning of this section. To this

[5] A distinction between two kinds uncertainty associated with an act a should be emphasized. The uncertainty concerning the prediction of a lower utility bound which is caused by a lack of information will be modelled, e.g., by $W_k \equiv 1$. Uncertainty concerning the prediction of a utility degree which is caused by the fact that a has led to very different outcomes for σ_k-similar problems will be reflected by a membership function μ_{W_k} having large values for small utilities and small values for large utilities. This example shows that we might be relatively certain about the value of a lower utility bound while being completely uncertain about the result that an act a will yield for a certain problem.

end, we extend the set of outcomes to $\mathbb{F}(\mathcal{R})$, i.e., the set of all (normal) fuzzy subsets of the set $\mathcal{R}$ of results. A representation W_k of knowledge about the lower bound w_k is then derived from "fuzzy" cases of the form $(p, a, R) \in \mathcal{Q} \times \mathcal{A} \times \mathbb{F}(\mathcal{R})$. A value $\mu_R(r)$ is interpreted as the possibility $\pi(X = r)$ that the (unknown, not precisely observed) outcome X is given by $r \in \mathcal{R}$. The derivation of W_k from "fuzzy" cases can be realized by applying the extension principle to a derivation of W_k from "crisp" cases.

It has already been hinted at in the introduction that Gilboa and Schmeidler's approach to case-based decision making is partly motivated by the idea of avoiding any kind of "hypothetical" reasoning. As poinetd out in [20], such reasoning might become necessary in connection with EUT since the decision maker has to know, e.g., all outcomes associated with act/state pairs. It should, therefore, be mentioned that the hypothetical reasoning in connection with (8) is by far less demanding. Particularly, it does not require any knowledge which is not available. On the contrary, nothing prevents us from using W_k in order to express complete ignorance. If available, however, general background knowledge (including hypothetical knowledge "derived" from the CBR assumption) should be utilized, and (8) presents the opportunity for doing this.

Comparing acts in the context of the generalized model (8) turns out as the problem of comparing fuzzy sets resp. possibility distributions. Needless to say, such a comparison is less straightforward than the comparison of scalar values. In fact, there are different possibilities to approach this problem [14] which are, however, not further discussed here.

4 Repeated decision making

The generalization of (4) which we will propose in this section is a weakening of the demand that an act has *always* given good results for similar problems. In fact, one might be satisfied if a turned out to be a good choice *for most* similar problems, thus allowing for a few exceptions [13]. In other words, the idea is to relax the universal quantifier "for all." Observe that a similar generalization of (5), which replaces "there exists" by "there are at least several" and, hence, corresponds to a strengthening of this decision principle, seems reasonable as well.

Consider a finite set A and let $m = |A|$. 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 [16, 29], the membership function $\mu : \{0, 1, \ldots, m\} \rightarrow [0, 1]$ of which satisfies

$$(\forall\, k \in \{1, \ldots, m-1\} : \mu(k) \leq \mu(k+1) \qquad \text{and} \qquad \mu(m) = 1. \tag{9}$$

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 (9), we define an associated membership function $\overline{\mu}$ by $\overline{\mu}(k) = 1 - \mu(k-1)$ for $1 \leq k \leq m$ and $\overline{\mu}(0) = 0$ (see, e.g., [11]). 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.

Now, 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 (4) is then given by

$$V_{p_0,\mathcal{M}}(a) = \min_{0 \leq k \leq |\mathcal{M}_a|} \max \left\{ 1 - \overline{\mu}(k), \max_{\mathcal{M}' \subset \mathcal{M}_a \,:\, |\mathcal{M}'| = k} \; \min_{(p,a,r) \in \mathcal{M}'} \max\{1 - \sigma_{\mathcal{Q}}(p, p_0), u(r)\} \right\}. \tag{10}$$

The second expression in braces on the right-hand side of (10) defines the degree $\delta(k)$ to which "the act a has induced good outcomes for similar problems k times." The extent to which a (small) degree $\delta(k)$ decreases the overall valuation of a is 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 $\delta(k)$. In fact, for computing $V_{p_0,\mathcal{M}}(a)$ one will arrange the $m = |\mathcal{M}_a|$ values $v = \max\{1 - \sigma_{\mathcal{Q}}(p, p_0), u(r)\}$ in a non-increasing order $v_1 \geq v_2 \geq \ldots \geq v_m$. Then, (10) is equivalent to

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

where $v_0 = 1$.

The generalized criterion (10) can be useful, e.g., in connection with the idea of repeated decision making. As already mentioned above, this idea 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, e.g., 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 such as, e.g., the iterated prisoner's dilemma [2].

As a concrete example let us consider a very simple model of repeated decision making: Suppose the agent to face the same problem p over and over again, 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 (cf. Section 3). 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 the utility values 0 and 1 only represent the fact that one outcome is preferred to the other, and that we might encode this by -1 and 1 as well.[6]

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 (4) simply corresponds to the worst outcome observed in connection with this act, i.e., $V^{\downarrow}_{p_0,\mathcal{M}}(a) = 0$ if $(p, a, 0) \in \mathcal{M}$, and $V^{\downarrow}_{p_0,\mathcal{M}}(a) = 1$ otherwise. 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 a_3 (and, sooner or later, a_2) 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 (4) is the fact that this criterion does not, in the long term, discriminate between two acts even though the first one strictly dominates the second one. Observe that the MAXIMIN rule does also not discriminate between a_1 and a_3.[7] This, however, seems to be acceptable more easily than the same property for (4): If used in connection with one-shot decisions, the MAXIMIN rule does not memorize experience from previous problem solving. That is, it does not have the opportunity of learning and experimenting in the course of a repeated problem solving process.[8]

It is just the above-mentioned drawback which can be avoided by (10) in conjunction with a proper formalization of the "for most" concept. In fact, since (10) allows for a few exceptions (and ω_3 is assumed to occur seldomly) 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 $\mu(k) = 1$ if $k/m \geq 1 - \varepsilon$, and $\mu(k) = 0$ otherwise. Then, we will have $V_{p_0,\mathcal{M}}(a_1) = 0$ if the fraction π_m of cases in which ω_3 has occured 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 occurence of ω_3. On the other hand, ε should not be made too large, since otherwise $V_{p_0,\mathcal{M}}(a_3) = 1$ as well, which means that a_1 and a_3 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 parametrized class of membership functions.)

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., if a_1 has not yet been tried very often. This problem, which is also related to the general problem of case-based decision making with only few observations (cf. Section 3), 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

[6] This clearly exemplifies that the application of (1) does hardly make sense.

[7] This can be achieved by extensions of MAXIMIN such as the ordinal decision rules DISCRIMIN and LEXIMIN [19].

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

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 valuation of acts (in the context of a certain problem) to the valuation of complete decision strategies. Of course, the question of when to prefer a certain decision rule to an alternative criterion, i.e., the question concerning the valuation of a (case-based) decision strategy, is by no means obvious in connection with the assumption of an ordinal setting for decision making. In fact, all kinds of "averaging" such as, e.g., the derivation of the mean of obtained utility values, are out of the question. Using the worst outcome, which might appear natural if (4) is seen as some kind of (case-based) analogue of the MAXIMIN decision rule, seems critical as well. Namely, 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 occured. In other words, it is anyway impossible for a case-based decision strategy to avoid the worst outcome completely, 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, for instance, 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 $u_1 < u_2 < \ldots < u_m$, and denote by $P_k^n(S)$ the probability of obtaining the utility u_k in the nth step of a decision sequence if strategy S is used.[9] Then, S dominates S' (stochastically) if

$$(\forall\, n \in \mathbb{N}, 1 \leq k \leq m) \,:\, \sum_{i=k}^{m} P_k^n(S') \leq \sum_{i=k}^{m} P_k^n(S). \tag{11}$$

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 receiving a "bad" and a "good" outcome, respectively, in connection with the nth decision. Moreover, a decision criterion S is preferred to S' in the sense of (11) if $P_1^n(S) \geq P_1^n(S')$ for all $n \in \mathbb{N}$. Appendix A shows simulation results for different decision strategies S_ε which only differ 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 valuated according to (10), 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 $\varepsilon < 0.1$, 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 (10) does not differentiate between a_1 and a_2 and, therefore, $P_1^n(S_\varepsilon) \rightarrow 3/4$ as $n \rightarrow \infty$. Observe 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 learning an optimal "for most" concept.

5 Summary

We have considered a recently proposed decision model which appears particularly interesting from the viewpoint of knowledge representation and reasoning since it combines ideas from *qualitative* and *case-based* decision making. Such an approach, however, does inevitably raise some difficulties. Firstly, a *case-based* approach to decision making seems problematic if an agent suffers from a lack of experience, in the sense that he has not yet encountered enough cases. In order to alleviate this problem we have replaced the scalar evaluation of an act by means of a "fuzzy" evaluation, which allows for the expression of uncertainty concerning the usefulness of an act. Besides, this approach allows for the processing of observations with imprecise outcomes, and for making use of background knowledge which is not necessarily based on observed cases.

Secondly, a *qualitative* setting for decision making rules out the idea of "averaging" in connection with the valuation of an act and, hence, does generally favour rather extreme decision rules. Just as, e.g., MAXIMIN, the case-based valuation of acts according to the above-mentioned model is extremely pessimistic. We have, therefore, proposed a formalization of a relaxed version of this ordinal (heuristic) decision principle. Instead of requiring that an act has *always* led to good outcomes for previous problems which are similar to the problem under consideration, the proposed version is less demanding and allows for some exceptions. This tolerance toward exceptions seems to be advantageous, e.g., if an agent has to act repeatedly over time. Moreover, it

[9] 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 nth decision is a function of the stochastic sequence of the first n problems $(p(1), \ldots, p(n))$.

makes the model more flexible and provides the opportunity of adapting it to a certain class of problems. This aspect is closely related to the idea of (case-based) *learning*, which is an interesting topic of future work.

Acknowledgements. The first author greatfully acknowledges financial support in form of a TMR research grant funded by the European Commission.

References

[1] M. Allais. Le comportement de l'homme rationnel devant le risque: Critique des postulates et axioms de l'ecole americaine. *Econometrica*, 21:503–546, 1953.

[2] R. Axelrod. *The Evolution of Cooperation*. Basic Books, Inc., New York, 1984.

[3] B. Bonet and H. Geffner. Arguing for decisions: a qualitative model for decision making. In *Proceedings* UAI-96, pages 98–105, 1996.

[4] C. Boutilier. Toward a logic for qualitative decision theory. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Proceedings* KR-94, *4.th International Conference on Principles of Knowledge Representation and Reasoning*, pages 75–86, Bonn, Germany, 1994.

[5] R. Brafmann and M. Tennenholtz. On the foundations of qualitative decision theory. In *Proceedings* AAAI-96*, 13th National Conference on Artificial Intelligence*, pages 1291–1296. AAAI-Press, 1996.

[6] R. Brafmann and M. Tennenholtz. On the axiomatization of qualitative decision criteria. In *Proceedings* AAAI-97*, 14th National Conference on Artificial Intelligence*, pages 76–81. AAAI-Press, 1997.

[7] N. Cercone, A. An, and C. Chan. Rule-induction and case-based reasoning: hybrid architectures appear advantageous. *IEEE Transactions on Knowledge and Data Engineering*, 11(1):166–174, 1999.

[8] D. Dubois, F. Esteva, P. Garcia, L. Godo, R.L. de Mantaras, and H. Prade. Fuzzy modelling of case-based reasoning and decision. In D.B. Leake and E. Plaza, editors, *Case-based Reasoning Research and Development, Proceedings* ICCBR-*97*, pages 599–610. Springer-Verlag, 1997.

[9] D. Dubois, F. Esteva, P. Garcia, L. Godo, R.L. de Mantaras, and H. Prade. Case-based reasoning: a fuzzy approach. In A. Ralescu, editor, *Proceedings* IJCAI-97 *Workshop on Fuzzy Logic in Artificial Intelligence*, number 1566 in Lecture Notes in Artificial Intelligence. Springer-Verlag, 1999.

[10] D. Dubois, L. Godo, H. Prade, and A. Zapico. Making decision in a qualitative setting: from decision under uncertainty to case-based decision. In A.G. Cohn, L. Schubert, and S.C. Shapiro, editors, *Proceedings of the 6th International Conference on Principles of Knowledge Represenation and Reasoning (*KR-98*)*, pages 594–605, Trento, Italy, 1998.

[11] D. Dubois, M. Nakata, and H. Prade. Extended divisions for flexible queries in relational databases. Technical Report 97-43 R, IRIT - Institut de Recherche en Informatique de Toulouse, Université Paul Sabatier, September 1997.

[12] D. Dubois and H. Prade. Possibility theory as a basis for qualitative decision theory. In *Proceedings 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 1924–1930, Montreal, 1995.

[13] D. Dubois and H. Prade. A fuzzy set approach to case-based decision. In R. Felix, editor, EFDAN-97*, 2nd European Workshop on Fuzzy Decision Analysis and Neural Networks for Management, Planning and Optimization*, pages 1–9, Dortmund, Germany, 1997.

[14] D. Dubois and H. Prade. A unified view of ranking techniques for fuzzy numbers. In *Proceedings* FUZZ-IEEE-99, Seoul, 1999. To appear.

[15] D. Dubois, H. Prade, and R. Sabbadin. Qualitative decision theory with sugeno integrals. In *Proceedings* UAI-94*, 14th Conference on Uncertainty in Artificial Intelligence*, pages 121–128. Morgan Kaufmann, 1998.

[16] D. Dubois, H. Prade, and C. Testemale. Weighted fuzzy pattern matching. *Fuzzy Sets and Systems*, 28:313–331, 1988.

[17] D. Ellsberg. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.

[18] H. Fargier, J. Lang, and T. Schiex. Mixed constraint satisfaction: a framework for decision problems under incomplete knowledge. In *Proceedings* AAAI-96*, 13th National Conference on Artificial Intelligence*, pages 175–180, Portland, Oregon, 1996.

[19] P. Fortemps and M. Pirlot. Axiomatic characterization of some ordinal choice procedures. In B. De Baets, J. Fodor, and L.T. Kóczy, editors, *Proceedings* EUROFUSE-SIC-99, pages 122–125, Budapest, Hungary, 1999.

[20] I. Gilboa and D. Schmeidler. Case-based decision theory. *Quarterly Journal of Economics*, 110(4):605–639, 1995.

[21] E. Hüllermeier. A Bayesian approach to case-based probabilistic reasoning. In *Proceedings IPMU-98, Seventh International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, pages 1296–1303, Paris, La Sorbonne, July 1998. Editions E.D.K.

[22] E. Hüllermeier. A possibilistic formalization of case-based reasoning and decision making. Technical Report 54, University of Paderborn, Department of Economics, 1998.

[23] E. Hüllermeier. A possibilistic formalization of case-based reasoning and decision making. In *Proceedings of the 6th International Conference on Computational Intelligence*, Dortmund, Germany, May 1999. To appear.

[24] E. Hüllermeier. A probabilistic approach to case-based inference. Technical Report 99-02 R, IRIT - Institut de Recherche en Informatique de Toulouse, Université Paul Sabatier, January 1999.

[25] J. Pearl. From qualitative utility to conditional "ought to". In D. Heckerman and H. Mamdani, editors, *Proceedings 9th International Conference on Uncertainty in Artificial Intelligence*, pages 12–20, San Mateo, CA, 1993. Morgan Kaufmann.

[26] S.W. Tan and J. Pearl. Qualitative decision theory. In *Proceedings* AAAI-94*, 11th National Conference on Artificial Intelligence*, pages 928–933, Seattle, WA, 1994.

[27] S.W. Tan and J. Pearl. Specification and evaluation of preferences under uncertainty. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Proceedings* KR-94*, 4.th International Conference on Principles of Knowledge Representation and Reasoning*, pages 530–539. Bonn, Germany, 1994.

[28] A. Tversky and D. Kahnemann. Rational choice and the framing of decisions. *Journal of Business*, 59:251–278, 1986.

[29] R.R. Yager. Aggregating evidence using quantified statements. *Information Sciences*, 36:179–206, 1985.

[30] L.A. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1965.

A Simulation results

The pictures on the left plot the performance $e(n)$ of a decision strategy for several simulation runs, where $e(n) = \sum_{k=1}^{n} u(k)/k$, and $u(k) \in U = \{0, 1\}$ denotes the outcome for the kth decision. The pictures on the right plot the respective frequencies $f(n) = \sum_{k=1}^{n} c(k)/n$ of "correct" decisions, i.e., $c(k) = 1$ if the kth decision was a_1, and $c(k) = 0$ otherwise.

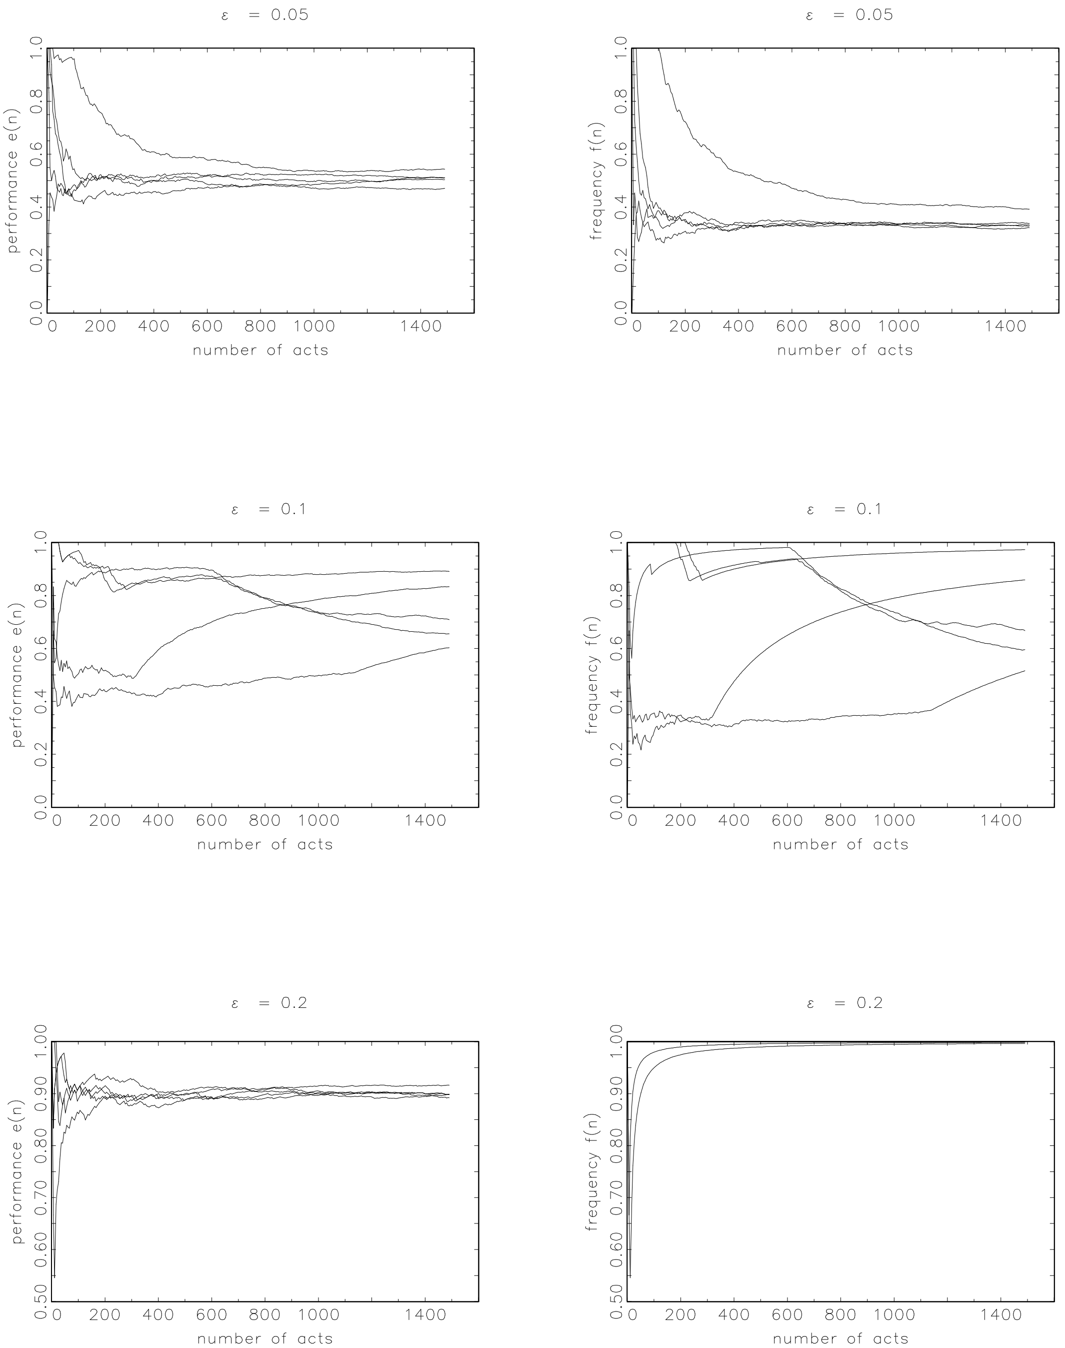

ε = 0.05
performance e(n)
number of acts
ε = 0.05
frequency f(n)
number of acts
ε = 0.1
performance e(n)
number of acts
ε = 0.1
frequency f(n)
number of acts
ε = 0.2
performance e(n)
number of acts
ε = 0.2
frequency f(n)
number of acts

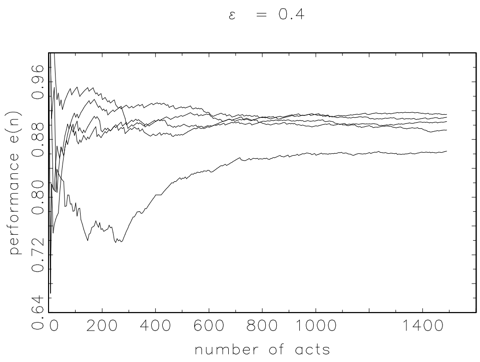
ε = 0.4
performance e(n)
number of acts
0.64
0.72
0.80
0.88
0.96
0
200
400
600
800
1000
1400

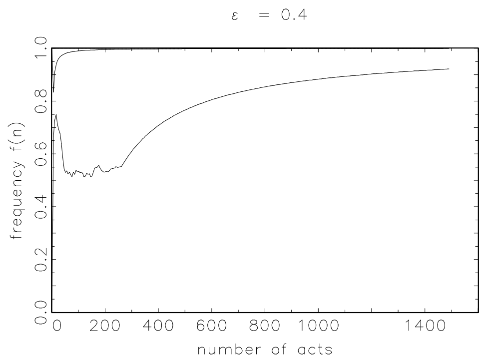
ε = 0.4
frequency f(n)
number of acts
0.0
0.2
0.4
0.6
0.8
1.0
0
200
400
600
800
1000
1400

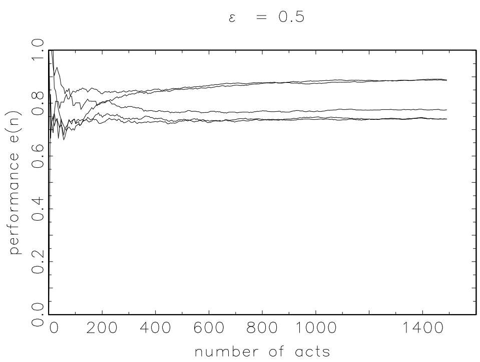
ε = 0.5
performance e(n)
number of acts
0.0
0.2
0.4
0.6
0.8
1.0
0
200
400
600
800
1000
1400

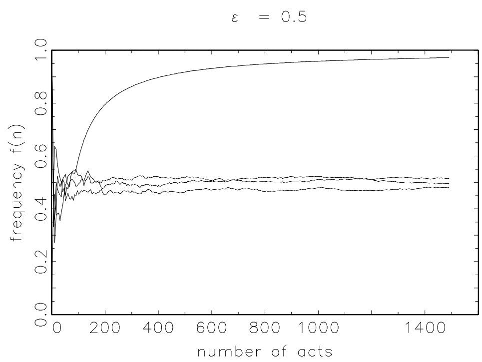
ε = 0.5
frequency f(n)
number of acts
0.0
0.2
0.4
0.6
0.8
1.0
0
200
400
600
800
1000
1400