

ADAPTATION KNOWLEDGE ACQUISITION: A CASE STUDY FOR CASE-BASED DECISION SUPPORT IN ONCOLOGY

MATHIEU D'AQUIN, JEAN LIEBER, AND AMEDEO NAPOLI

*Orpailleur research group, LORIA (UMR 7503 CNRS-INPL-INRIA-Nancy2-UHP),
Université Henri Poincaré for M. d'Aquin and J. Lieber, CNRS for A. Napoli
BP 239, 54506 Vandœuvre-lès-Nancy, France,
daquin@loria.fr, lieber@loria.fr, napoli@loria.fr*

KASIMIR is a case-based decision support system in the domain of breast cancer treatment. For this system, a problem is given by the description of a patient and a solution is a set of therapeutic decisions. Given a target problem, KASIMIR provides several suggestions of solutions, based on several justified adaptations of source cases. Such adaptation processes are based on adaptation knowledge. The acquisition of this kind of knowledge from experts is presented in this paper. It is shown how the decomposition of adaptation processes by introduction of intermediate problems can highlight simple and generalizable adaptation steps. Moreover, some adaptation knowledge units that are generalized from the ones acquired for KASIMIR are presented. This knowledge can be instantiated in other case-based decision support systems, in particular in medicine.

Key words: case-based decision support, adaptation, knowledge acquisition, breast cancer treatment, medical informatics.

1. INTRODUCTION

Case-based reasoning (CBR) consists of reusing the solutions of already solved problems in order to solve a new problem (Riesbeck & Schank, 1989). A CBR system exploits a case base, a case being a problem accompanied with one of its solutions. A case from the case base is called a *source case*. A CBR process usually relies on a retrieval step (selection of a source case) and an adaptation of the solution of the source case in order to solve the new problem. In many CBR systems, the adaptation is based on complex and domain-dependent adaptation knowledge that has to be acquired and modeled. This is the goal of adaptation knowledge acquisition (AKA).

This paper presents AKA from experts for the KASIMIR system whose application domain is breast cancer treatment: a problem is given by a description of a patient and a solution is a decision of a treatment for this patient. In this domain, the consequences of a wrong decision may be disastrous: it may severely endanger the health of the patient. This is why the KASIMIR system proposes several alternative adaptations, each of them being accompanied with explanations: as shown in (Doyle et al., 2005), explanations play a key role for the acceptance of a decision support system in medicine by the practitioners. Beyond this application, our ambition is to highlight, on the one hand, general adaptation knowledge units that can be instantiated for case-based decision support systems and, on the other hand, first elements of an AKA methodology.

Section 2 presents the KASIMIR project and the need for adaptation knowledge it involves. The principle of the AKA sessions is presented in section 3, together with a detailed example. Some general adaptation knowledge units for case-based decision support systems, in particular in medical domains, resulting from AKA sessions, are presented in section 4. Section 5 discusses this research work and the results, and section 6 concludes the paper.

2. CONTEXT: THE KASIMIR PROJECT

In the Lorraine region (East of France) the decisions about cancers for treatment, surveillance, etc., are based on decision protocols. For example, the protocol of the breast cancer treatment, simply called “the protocol” hereafter, associates to a patient description a treatment recommendation. It can be seen as a set of rules $R = (\text{Prem} \rightarrow \text{Cc1}^\circ)$ where Prem represents a class of patients and Cc1° represents a set of therapeutic decisions. For most medical cases (about 60 to 70%), the protocol is simply applied. The remaining cases, called the *out of protocol cases*, are examined by the *breast therapeutic decision committee* (BTDC). It has been shown that the BTDC generally does not solve the decision problems raised by the out of protocol cases from scratch, but *adapts* the protocol for them (Sauvagnac, 2000). More precisely, the BTDC selects a protocol rule $R = (\text{Prem} \rightarrow \text{Cc1}^\circ)$ such that Prem is close to the patient description and it adapts Cc1° to propose a treatment that is adapted to her/his specificities.

Thus, this adaptation of the protocol is a kind of CBR process: the problems to be solved are given by patient descriptions and their solutions are therapeutic decisions. Its peculiarity, as a CBR process, is that the source cases are general cases (according to the terminology of (Riesbeck & Schank, 1989), they are *ossified cases*): they *are* the rules $R = (\text{Prem} \rightarrow \text{Cc1}^\circ)$ and thus the case base *is* the protocol. Nevertheless, this peculiarity does not reduce the scope of the approach presented in this paper.

The goal of the KASIMIR project is decision knowledge management in oncology and its main subgoal is modeling of the physicians’ decision making. The KASIMIR system is a CBR system designed for this purpose and is also intended to be used as an intelligent assistant for physicians in their practice of decision making. Two versions of KASIMIR have been developed. The first one uses an object-based representation formalism (d’Aquin et al., 2004a), the second one is a semantic portal using standard knowledge languages dedicated to the semantic Web (d’Aquin et al., 2005a).

The adaptation step of the CBR system KASIMIR is essential since our aim is to model the process of protocol adaptation. The adaptation processes are based on complex and heterogeneous knowledge that has to be acquired.

3. AKA FROM EXPERTS: A CASE STUDY

This section aims at describing the activity of AKA for the KASIMIR system. Adaptation knowledge is a special form of knowledge in the sense that it is intended to be used during the adaptation step of the CBR cycle, in interrelation with the retrieval step. Adaptation knowledge units have to be elicited from real-world situations for becoming operational. Thus, a classical knowledge acquisition and modeling method, such as CommonKADS (Schreiber et al., 1999), cannot be directly carried out in the present research work. The adaptation principle on which this research is based is presented in section 3.1. The main steps of AKA are presented in section 3.2. An example is detailed in section 3.3.

3.1. Adaptation Principle

The principle of adaptation of the KASIMIR system has been developed during the conception and implementation of another system, RESYN/CBR, in the domain of synthesis planning in organic chemistry (Lieber & Napoli, 1996).

The notions of *problem* and *solution* are domain-dependent. In a given application domain, let tgt , be a problem to be solved (a *target* problem). Let $(\text{srce}, \text{Sol}(\text{srce}))$ be a case retrieved from

the case base that must be adapted to solve tgt : srce is a problem and $\text{Sol}(\text{srce})$ is a solution of srce . Adapting $\text{Sol}(\text{srce})$ for solving tgt consists in building a solution $\text{Sol}(\text{tgt})$ of tgt derived from $\text{Sol}(\text{srce})$.

The first adaptation step consists in *matching* srce and tgt , i.e., in highlighting how these problems are similar and how they are dissimilar. In our approach, the matching result is a *similarity path*, i.e. a sequence

$$\text{pb}_0 \text{ r}_1 \text{ pb}_1 \text{ r}_2 \text{ pb}_2 \dots \text{pb}_{q-1} \text{ r}_q \text{ pb}_q$$

such that:

- The pb_i 's denote problems and the r_i 's denote binary relations between problems;
- $\text{pb}_0 = \text{srce}$ and $\text{pb}_q = \text{tgt}$;
- For each $i \in \{1, 2, \dots, q\}$, some *adaptation knowledge* is available to adapt the solution $\text{Sol}(\text{pb}_{i-1})$ of pb_{i-1} in a solution $\text{Sol}(\text{pb}_i)$ of pb_i .

The second adaptation step consists simply in “following” the similarity path in the solution space, involving the adaptation chain: 1°/ $\text{Sol}(\text{srce}) = \text{Sol}(\text{pb}_0)$ in $\text{Sol}(\text{pb}_1)$, 2°/ $\text{Sol}(\text{pb}_1)$ in $\text{Sol}(\text{pb}_2)$, ... q °/ $\text{Sol}(\text{pb}_{q-1})$ in $\text{Sol}(\text{pb}_q) = \text{Sol}(\text{tgt})$.

Implementing the adaptation function requires (a) design of a matching scheme which points out a similarity path, and (b) acquisition and modeling of adaptation knowledge. This adaptation knowledge, as seen above, aims at producing $\text{Sol}(\text{pb}_i)$ from $\text{Sol}(\text{pb}_{i-1})$, knowing on the one hand pb_{i-1} and pb_i , and on the other hand the relation r_i relating pb_{i-1} and pb_i . The relation r_i determines the adaptation function \mathcal{A}_{r_i} to be used:

$$\mathcal{A}_{\text{r}_i} : (\text{pb}_{i-1}, \text{Sol}(\text{pb}_{i-1}), \text{pb}_i) \mapsto \text{Sol}(\text{pb}_i)$$

Thus adaptation knowledge is composed of ordered pairs $(\text{r}_i, \mathcal{A}_{\text{r}_i})$ called *reformulations* in (Melis et al., 1998). A reformulation $(\text{r}, \mathcal{A}_{\text{r}})$ has to be understood as an “adaptation rule”:

if $\text{pb} \text{ r } \text{pb}'$ // pb is related to pb' by r
then $\mathcal{A}_{\text{r}}(\text{pb}, \text{Sol}(\text{pb}), \text{pb}') = \text{Sol}(\text{pb}')$ // $\text{Sol}(\text{pb})$ is adapted in $\text{Sol}(\text{pb}')$ by \mathcal{A}_{r}

The problems $\text{pb}_1, \text{pb}_2, \dots, \text{pb}_{q-1}$ are built during the matching process. For KASIMIR, these *intermediate problems* correspond to *virtual patients*: they are fictitious patients introduced during the reasoning.

Finally, it must be noticed that an adaptation has a *cost*, in general non-null. This cost characterizes the fact that the solution $\text{Sol}(\text{tgt})$ of tgt may be worse than the solution $\text{Sol}(\text{srce})$ of srce . The precise meaning of this cost depends on the CBR application. For KASIMIR, this cost is characteristic of the risk, taken during adaptation, of a wrong treatment choice. A reformulation can be accompanied by some information on its cost. In particular, a method for computing a numerical cost evaluating the adaptation is needed and it is used to select, during the retrieval step, the case that is less costly to adapt. Furthermore, some qualitative pieces of information about this cost can be useful for the presentation of the reasoning; they highlight the pros and cons of the application of a reformulation. For KASIMIR, these pieces of information are explanations about the reformulation.

It can be noticed that we adhere to the principle of adaptation-guided retrieval (Smyth & Keane, 1996), meaning that each retrieved source case is adaptable into a solution of the target problem. Moreover, the retrieval process points out a similarity path, and so, performs the first step of adaptation. Thus, the adaptation knowledge contributes to the retrieval knowledge and the AKA sessions contribute to the elicitation of this knowledge.

3.2. AKA Sessions

The adaptation of the protocol is performed during the meetings of the BTDC (cf. section 2). Minutes of these meetings have been written and analyzed from an ergonomic viewpoint (see (Sauvagnac, 2000)). The AKA sessions have consisted in the study of these minutes in presence of experts in cancerology, of a specialist in ergonomics and of computer specialists. Schematically, such a session has been composed of four phases:

phase 1: Presentation of the minutes by the specialist in ergonomics, with corrections and refinements from the experts.

phase 2: Discussion among the different participants so that each of them understands the reasoning process that has led to an adaptation.

phase 3: Re-description of this reasoning process by the computer specialists and discussions on the variations of this reasoning process.

phase 4: Analysis of the reasoning from the perspective of general adaptation knowledge propositions (this last phase usually takes place after the session).

It must be noticed that the specialist in ergonomics is also a physician, which facilitates her interactions with the experts and the communication between experts and computer specialists, giving her a status of interpreter. A previous work on a knowledge-based system in organic synthesis in chemistry has shown the usefulness of such an interpreter (Napoli et al., 1994). In these studies, it is important that the experts have some idea about the modeling and, thus, the domain knowledge representation. Indeed, conversely to the approach “cognitician-expert”, where the first person monopolizes the power related to the computer, it is essential that the expert has knowledge and consciousness of the available tools, of their advantages and limits, especially for knowledge representation and reasoning. So, during the transfer of expertise, the traditional problems of misunderstanding between computer specialists and experts are attenuated if not completely suppressed: the former cannot promise to the latter what the latter cannot hope to obtain. This approach is distributed and honest, in the sense that the modeling and design of the system are based on a real collaboration between the experts and the computer specialists.

3.3. A Detailed Example

The example presented in this section is a real example with two modifications. First, the name of the patient has been changed to Jules. Second, the case has been modified to simplify the description of the corresponding adaptation. In reality, this case has been treated in its whole complexity. Furthermore, some pieces of information were omitted because they did not play any role in the reasoning.

Jules is a man with a cancer at the left breast. The first feature making him an out of protocol case is his sex. Indeed, the large majority of persons suffering from breast cancer are women, so the protocol—coming largely from statistical studies—has been elaborated for them. The idea is then to *do as if* Jules was a woman and to reason with this working hypothesis. Note that the use of expressions like “We do as if...” by the experts points out the possible presence of adaptation knowledge.

Another difficulty for Jules is that his tumor localization in his left breast is unknown. This raises a difficulty since it is important, for the radiotherapy, to know whether the tumor is external, central or internal. More precisely, the most pessimistic assumption—the one that takes the most precautions with radiotherapy—is that the tumor is internal or central. The experts make this assumption. So, if they are wrong, this would only mean that useless precautions have been taken.

To summarize, two difficulties that made Jules an out of protocol case have been successively (and temporarily) suppressed. This can be formulated by introducing two virtual patients: (1) a virtual patient Julie who is just like Jules but is a woman, (2) a virtual patient Juliette who is just like Julie except for the tumor localization (the localization of Julie’s tumor is unknown whereas it is known that the localization of Juliette’s is internal or central). Juliette is *in* the protocol, meaning that there is a rule of the protocol $R = (\text{Prem} \longrightarrow \text{Cc1}^\circ)$ such that Prem holds for Juliette—denoted by $\text{Prem} \Leftarrow \text{Juliette}$. Thus the following similarity path relates Jules to the protocol:

$$\text{Prem} \Leftarrow \text{Juliette} \text{ ps } \text{Julie} \text{ cs } \text{Jules}$$

where ps and cs are relations between problems that have to be modeled (see hereafter) and where the patients Jules, Julie and Juliette are described by:


Jules = (sex = male *and* tumor localization = unknown *and* ...)

Julie = (sex = female *and* tumor localization = unknown *and* ...)

Juliette = (sex = female *and* tumor localization = internal or central *and* ...)

(the dots correspond to features that play no role in this example). Prem is a generic patient (or a class of patients) for which the treatment $\text{Sol}(\text{Prem}) = \text{Cc1}^\circ$ is a radiotherapy taking into account the internal or central position of the tumor, and a hormonotherapy using tamoxifen (an anti-oestrogen drug).

When the similarity path is built—from Jules to Prem , reading from right to left—the reverse path in the solution space is followed, i.e., from the treatment $\text{Sol}(\text{Prem})$ of Prem to a treatment $\text{Sol}(\text{Jules})$ of Jules, reading from left to right:



3.3.1. Reformulation ($\Leftarrow, \mathcal{A}_{\Leftarrow}$)

The treatment $\text{Sol}(\text{Prem})$ can be applied to Juliette since $\text{Prem} \Leftarrow \text{Juliette}$. The knowledge unit reified by the reformulation ($\Leftarrow, \mathcal{A}_{\Leftarrow}$) can be written: “A treatment designed for a general case can be applied to a specific case of this general case.” (This reformulation is not a new knowledge unit: it is the basis of deductive reasoning in KASIMIR, that consists in applying the protocol, and not in adapting it.)

3.3.2. Reformulation ($\text{ps}, \mathcal{A}_{\text{ps}}$)

Juliette is a “pessimistic specialisation (ps)” of Julie: she is characterized by the fact that the tumor position of Julie has been precisely specified for Juliette and that this position is the one that makes the radiotherapy most complex (without modifying the other treatments). Therefore, the treatment $\text{Sol}(\text{Juliette})$ is copied for Julie. This reformulation is based on *Wald pessimistic criterion* (see section 4.3).

3.3.3. Reformulation (cs, \mathcal{A}_{cs})

Finally, questions are raised about the applicability of the treatment $\text{Sol}(\text{Julie})$ of Julie to Jules, her male equivalent. These questions deal with the consequences of the change of sex (cs) on the applicability of some treatment components. Following the principles developed in (Fuchs et al., 2000), we are interested on the *dependencies* between the descriptor “sex” of the problems and the descriptors “radiotherapy”, “hormonotherapy”, etc., of the solutions. In (Fuchs et al., 2000), the dependencies are defined by $\frac{\Delta y}{\Delta x}$ where Δx is the variation of a problem descriptor x and Δy is the variation of a solution descriptor y . For Julie and Jules, we are interested in $\frac{\Delta_{\text{radiotherapy}}}{\Delta_{\text{sex}}}$ and in $\frac{\Delta_{\text{hormonotherapy}}}{\Delta_{\text{sex}}}$. The knowledge acquisition from the experts indicates that these dependencies are null: the radiotherapy and the hormonotherapy recommended for Julie remain recommended for Jules.

The reformulation (cs, \mathcal{A}_{cs}) is based on the dependencies $\frac{\Delta \theta}{\Delta_{\text{sex}}}$, where θ is a particular treatment. The discussion on the variations (cf. phase 3 of section 3.2) makes it possible to precisely specify these dependencies. In this example, we try to establish the treatments “invariant by change of sex” and, for the other ones, how they can be adapted. For instance, the hormonotherapy consisting in an ablation of the ovaries is not invariant by change of sex. This treatment is replaced by a treatment that, for a man, brings similar expected benefits, e.g. a cure of tamoxifen.

4. RESULTS: ADAPTATION KNOWLEDGE FOR CASE-BASED DECISION SUPPORT

Several AKA sessions such as the one presented in the previous section have been carried out. They have led to a few domain-specific reformulations and also to *adaptation patterns*, i.e., general reformulations applicable to a variety of situations provided that they are correctly instantiated. These patterns can be reused for case-based decision support systems and are associated with explanations: in this way, justifications of the adaptation steps are provided to the physician who can either accept or reject them with a full knowledge of the facts, and this is particularly important in a domain, such as medicine, in which the impact of a wrong decision may be disastrous.

Making a decision is choosing or designing an action that modifies the state of the world. A therapeutic decision, for example, leads to a treatment that modifies the state of the patient. An action (and, by extension, a decision) may be applicable or not, and may have positive and/or negative consequences. Section 4.1 describes an adaptation pattern of an inapplicable decision. Section 4.2 describes two adaptation patterns based on the consequences of a decision. A by-product of the AKA sessions is the discovery of knowledge units for the *retrieval* step of CBR. Section 4.3 briefly presents some examples of knowledge units for retrieval.

In this section, a problem pb is the description of a situation, and a solution $\text{Sol}(pb)$ of pb is modeled by a set of decisions dec that are applied to the situation pb . For KASIMIR, pb represents a patient and $\text{Sol}(pb)$ represents a set of therapeutic decisions (e.g., surgery, chemotherapy, etc.). In some situations, adaptation requires a more precise modeling of solutions—e.g., taking into account temporal constraints between decisions dec —but such situations are not presented below.

4.1. Adaptation of an Inapplicable Decision

Let us consider a situation for which $\text{Sol}(srce)$ cannot be applied to solve tgt because one of the decisions $dec \in \text{Sol}(srce)$ is not applicable in the framework of tgt (or, is judged to be

too difficult to be applied). In this case, one way to perform adaptation consists in identifying the decision $\underline{dec} \in \text{Sol}(\text{srce})$ that is not applicable to tgt , and to replace \underline{dec} by a decision \underline{dec}' that is applicable and that has similar positive consequences as \underline{dec} (for KASIMIR, \underline{dec} and \underline{dec}' must have similar expected therapeutic benefits). This replacement can be seen as a removal of \underline{dec} followed by an addition of \underline{dec}' , this addition aiming at *compensating* for the removal of the positive consequences of \underline{dec} . The adaptation pattern of figure 1 describes this kind of adaptation.

<p>Parameters $\underline{dec}, \underline{dec}'$: two decisions having similar positive consequences when they are applicable and \underline{R}: a reason why a decision is not applicable</p> <p>If $\underline{dec} \in \text{Sol}(\text{srce})$ and \underline{dec} is not applicable (or difficult to apply) to tgt because of \underline{R} and \underline{dec}' is applicable on tgt</p> <p>then replace \underline{dec} by \underline{dec}' in $\text{Sol}(\text{srce})$ for obtaining $\text{Sol}(\text{tgt})$:</p> $\text{Sol}(\text{tgt}) = (\text{Sol}(\text{srce}) \setminus \{\underline{dec}\}) \cup \{\underline{dec}'\}$ <p>Explanation The decision \underline{dec} being non applicable to tgt because of \underline{R}, it can be replaced by \underline{dec}' that is applicable, and leads to positive consequences on tgt similar to the one of \underline{dec} on srce.</p>

Figure 1: An adaptation pattern for an inapplicable decision.

The AKA sessions for KASIMIR have highlighted adaptations that instantiate this pattern. These adaptations are performed at different levels of granularity of decisions. One example at a fine-grained level considered during an AKA session is related to the case of an aged patient, with low mobility and living far from a radiotherapy center. The decision reached by applying the protocol is to give a daily session of radiotherapy, during 3 weeks, which involves too much travel for her. The adaptation has consisted of raising the dose d by session and reducing the number n of sessions, keeping the overall dose $n \times d$ constant: the decreasing of n was compensated for by the increase of d (it was also necessary to take into account the maximal daily dose d_{\max}). This adaptation instantiates the adaptation pattern of figure 1 by:

$$\begin{aligned} \underline{dec} &= n \text{ sessions of radiotherapy, with dose } d \\ \underline{dec}' &= n' \text{ sessions of radiotherapy, with dose } d' \text{ with } n' < n, n' \times d' = n \times d \text{ and } d' \leq d_{\max} \\ \underline{R} &= \text{low mobility of the patient} \end{aligned}$$

At a coarser level, there is the replacement of a treatment of a type by a treatment of another type. For example, a patient had, for psychological reasons, her arms crossed in front of her chest. It was possible to perform a surgery on the breast under general anesthesia, but, it was practically impossible to give her radiotherapy treatments. The goal of such radiotherapy treatments is to kill the cancerous cells near the surgical zone (after the surgery). The adaptation consisted of not using any radiotherapy and compensating for this with a surgery larger than the one recommended by the protocol, including the zone that would have been the goal of radiotherapy. This adaptation instantiates the pattern of figure 1 by:

$$\begin{aligned} \underline{dec} &= \text{radiotherapy of the surgery zone} \\ \underline{dec}' &= \text{enlarged surgery (compared to the surgery of } \text{Sol}(\text{srce})) \\ \underline{R} &= \text{practical impossibility to perform radiotherapy sessions on the patient} \end{aligned}$$

In section 3.3.3, it was noted that the treatment by ablation of the ovaries can be adapted to men by the domain-dependent reformulation (cs, \mathcal{A}_{cs}). This adaptation can also be managed by the following instantiation of the figure 1 adaptation pattern:

$$\begin{aligned} \underline{dec} &= \text{ablation of the ovaries} \\ \underline{dec}' &= \text{tamoxifen} \\ \underline{R} &= \text{men have no ovaries} \end{aligned}$$

4.2. Adaptation based on the Consequences of a Decision

The *therapeutic index* is a conceptual tool used by physicians (in particular, oncologists involved in the KASIMIR project). Given a patient and a therapeutic decision, the therapeutic index is the rate $benefits/undesirable_effects$ where *benefits* is the measure of the expected therapeutic benefits of the treatment and *undesirable_effects* is the measure of its undesirable effects. The idea is that the higher this index is, the better the treatment is for the patient. This index is sometimes used in a quantitative way but can also be used qualitatively:

- (A) The undesirable effects being constant, the index increases when the expected benefits are improved.
- (B) The expected benefits being constant, the index increases when the undesirable effects are reduced.
- (C) When the expected benefits are improved and the undesirable effects are reduced, the index increases.

The notion of therapeutic index can be reused in other decision support applications: instead of using the oncology-related notions of expected therapeutic benefits and of undesirable effects, the domain-independent notions of positive and negative consequences of a decision can be used.

The adaptation aims at finding the decision that gives a high index value, given the target problem. The above assertions (A), (B) and (C) can be used for this purpose.

Parameters $\underline{dec}, \underline{dec}'$: two decisions having similar positive consequences when they are effective and \underline{f} : a problem feature

If $\underline{dec} \in \text{Sol}(srce)$ **and** the only difference between $srce$ and tgt is the difference of their feature \underline{f} ($\underline{f}(srce) \neq \underline{f}(tgt)$) that makes \underline{dec} ineffective while it does not prevent \underline{dec}' from being effective.

then replace \underline{dec} by \underline{dec}' in $\text{Sol}(srce)$ for obtaining $\text{Sol}(tgt)$:

$$\text{Sol}(tgt) = (\text{Sol}(srce) \setminus \{\underline{dec}\}) \cup \{\underline{dec}'\}$$

Explanation Since the decision \underline{dec} is not effective on tgt because of the problem feature \underline{f} , it can be replaced by \underline{dec}' that has similar positive consequences on tgt as \underline{dec} has on $srce$.

Figure 2: An adaptation pattern of a decision having insufficient positive consequences.

(A) is useful in particular whenever the positive consequences of $\text{Sol}(srce)$ on tgt are lower than the positive consequences of $\text{Sol}(srce)$ on $srce$. The extreme case occurs when a decision has no positive consequences on tgt . An adaptation pattern corresponding to this extreme case is the one of figure 2 (an instantiation of this pattern is presented in section 5.3).




Figure 3: An adaptation pattern of a decision having too many negative consequences.

(B) is useful for KASIMIR in particular to manage the contraindications of a treatment. The pattern of figure 3 describes this kind of adaptation in the general framework of case-based decision support. For example, let us consider the example of a patient having a liver disease making the hormonotherapeutic treatment by tamoxifen contraindicated. Since the protocol does not take into account this contraindication, it must be adapted. A possible adaptation consists in replacing tamoxifen by anti-aromatases, that are drugs with similar expected benefits as tamoxifen. This adaptation instantiates the pattern of figure 3 by:

$$\begin{aligned} \underline{dec} &= \text{hormonotherapy with tamoxifen} \\ \underline{dec}' &= \text{hormonotherapy with anti-aromatases} \\ \underline{R} &= \text{allergy to tamoxifen} \end{aligned}$$

Some adaptation processes based on (B) consist in adding to the treatment a decision making the therapy tolerable. For example, let us consider a patient having a tumor at the left breast and for whom the protocol proposes a radiotherapy of this breast, after the surgery. Moreover, let us assume that this patient wears an artificial pacemaker. The radiations of the radiotherapy may disturb the mechanism of the pacemaker. One adaptation consists in moving the pacemaker before the radiotherapy cure, in order to avoid this disturbance. This adaptation instantiates the pattern of figure 3 by:

$$\begin{aligned} \underline{dec} &= \text{perform a radiotherapy of the left breast} \\ \underline{dec}' &= \text{move the pacemaker and then perform a radiotherapy of the left breast} \\ \underline{R} &= \text{the radiations may disturb the pacemaker} \end{aligned}$$

(C) can be seen as the composition of (A) and (B): an adaptation based on (C) can be seen as an adaptation along a similarity path of length 2 such that one of its step is based on (A) and the other one is based on (B).

4.3. Retrieval Knowledge Acquired

During the AKA session, knowledge units for case retrieval have also been acquired. This is not surprising since, as mentioned at the end of section 3.1, we adhere to the adaptation-guided retrieval principle. Below, two kinds of retrieval knowledge units are presented.

Taking into Account Missing Data. It may occur that pieces of information that are useful for decision problem solving are missing. It occurs in the framework of KASIMIR, in particular when an examination cannot be performed on a patient without endangering her/him. In this context, the Wald pessimistic criterion (Wald, 1950; Dubois et al., 2001) can be applied. According to this criterion, decisions must be evaluated on the basis of their worst possible consequences (w.p.c.): dec_1 is preferred to dec_2 if the w.p.c. of applying dec_1 to tgt is preferred to the w.p.c. of applying dec_2 to tgt . In (d’Aquin et al., 2005b), the use of this criterion for the retrieval step of case-based decision support systems is studied. This article presents an approach using a rather small amount of knowledge and a way to acquire it during the use of a CBR system, according to what Kristian J. Hammond calls *learning by remembering*: when the system needs a preference between decision consequences and does not have knowledge to infer it, it asks the expert and will be able to reuse her/his answer later (Hammond, 1990).

Hereafter, an example is presented. Let us consider the case of a patient having a 2 centimeters tumor. Given other features, the protocol recommends a partial mastectomy. Now, the radiography shows some white dots on the image that are rather far away from the tumor and that may be either (a) cancerous cells, or (b) something harmless. Under assumption (a), a radical mastectomy is recommended—decision $dec_{(a)}$. Under assumption (b), a partial mastectomy is recommended—decision $dec_{(b)}$. If no examination before surgery can indicate which of the hypotheses (a) and (b) is correct, the question that is raised is to know whether it is better to do

- ($dec_{(a)}/b$) A radical mastectomy $dec_{(a)}$ under assumption (b)—and thus, a larger surgery than necessary—or
- ($dec_{(b)}/a$) A partial mastectomy $dec_{(b)}$ under assumption (a)—which would leave cancerous cells in the body of the patient.

Whenever additional knowledge is available telling that ($dec_{(a)}/b$) has to be preferred to ($dec_{(b)}/a$), the relevant source case for case retrieval is the one with $dec_{(a)}$ in its solution.

Taking into account the Threshold Effect. When a numerical patient feature (e.g., the age) is close to a decision threshold of the protocol, the simple application of the protocol raises a problem. For example, let $srce_1$, $srce_2$ and tgt be the following problems:

$$\begin{aligned} srce_1 &= (\text{sex} = \text{female} \text{ and } \text{tumor-size} \in [0; 4[) \\ srce_2 &= (\text{sex} = \text{female} \text{ and } \text{tumor-size} \in [4; 7[) \\ tgt &= (\text{sex} = \text{female} \text{ and } \text{age} = 56 \text{ and } \text{tumor-size} = 3.8) \end{aligned}$$

$srce_1$ (resp., $srce_2$) is assumed to be a problem of the protocol and $Sol(srce_1)$ (resp., $Sol(srce_2)$) is assumed to be the solution of $srce_1$ (resp., of $srce_2$) in the protocol. Moreover, it is assumed that $Sol(srce_1) \neq Sol(srce_2)$. tgt is a target problem. The answer to the question “What solution should be associated with tgt ?” when applying the protocol is $Sol(srce_1)$ and not $Sol(srce_2)$, because $tgt \leftarrow srce_1$ and $tgt \not\leftarrow srce_2$. But, since the size of the tumor of the patient associated with tgt , 3.8 cm, is close to the threshold $\tau = 4$ cm, this decision is not certain, for two reasons. First, the value of the decision threshold τ is uncertain, second, the measure 3.8 cm may be imprecise. A better idea is to propose to the user of the KASIMIR system *both* solutions $Sol(srce_1)$ and $Sol(srce_2)$, with a preference for the first one.

More generally, AKA sessions have shown that when a patient numerical feature f is close to a decision threshold τ (f may be the tumor size, the age of the patient, etc.) then the experts

often raise the problem of choosing between the two decisions: the one for $f < \tau$ and the one for $f \geq \tau$. This has been modeled in KASIMIR thanks to the substitution of *fuzzy thresholds* for absolute thresholds, and this has required the development of a new reasoner for the object-based version of KASIMIR (Lieber, 2002; d’Aquin et al., 2004a). The development of a fuzzy reasoner for the semantic Web portal of KASIMIR is planned (d’Aquin et al., 2006).

5. DISCUSSION

5.1. AKA is an Inductive Process

Contrasting with more classical knowledge acquisition and modeling techniques, such as CommonKADS (Schreiber et al., 1999), AKA for KASIMIR relies on the one hand on documents related to decision support in oncology (the protocols) and, on the other hand, on the description of real adaptations. Thus, the generic patterns that have been introduced before are extracted from and are abstraction of specific adaptations, in accordance with the domain knowledge given by the protocols. Therefore, the AKA process can be seen as an inductive learning process: extracting generic knowledge units from specific knowledge units. One further step would be to create a template knowledge model for AKA in the same way as it is done for general tasks such as, e.g., classification, assessment, diagnosis (Schreiber et al., 1999). This would make concrete the transformation between the present symbolic level, i.e., taking into account specific out of protocol cases, towards a generic knowledge level, making explicit a reusable methodology for AKA. A proposal for first elements of such a future methodology are introduced in next section.

5.2. Elements of a Methodology for AKA from Experts

This section gathers first elements of a methodology for AKA from experts.

The first issue—maybe the most important—is the decomposition of adaptation based on the notions of similarity path and of intermediate problems between the source and target problems, pointing out simple adaptation steps that can be generalized in reformulations. The AKA we have described is based on informal descriptions of adaptation processes performed by experts. For each of these adaptation processes, the knowledge acquisition steps are as follows:

- Re-description of the adaptation process in several steps by introducing intermediate problems $pb_1, pb_2, \dots, pb_{q-1}$ and their respective solutions $Sol(pb_1), Sol(pb_2), \dots, Sol(pb_{q-1})$. Recall that $pb_0 = srce$ is the source problem and that $pb_q = tgt$ is the target problem.

The elicitation of the intermediate problems is often made from the right to the left, i.e., from pb_i to pb_{i-1} . For example, when the expert associates to pb_i a working hypothesis (“We do as if some conditions on pb_i were changed”), it can be expressed by introducing the problem pb_{i-1} .

- For each $i \in \{1, 2, \dots, q\}$, analysis of the adaptation step

$$(pb_{i-1}, Sol(pb_{i-1}), pb_i) \mapsto Sol(pb_i)$$

This analysis aims at giving a reformulation (r_i, \mathcal{A}_{r_i}) which is either a reformulation already in the adaptation knowledge base, or a new one.

The second issue is linked to problem and solution representations. Indeed, it is useful not only to represent what a solution is but also how it does or does not contribute to solving the

problem. For KASIMIR, the knowledge linked with the therapeutic index (i.e., expected benefits and undesirable effects) is a part of this knowledge.

The third issue concerns the dependencies between problem descriptors x and solution descriptors y , as seen above in section 3.3, about the reformulation (cs, \mathcal{A}_{cs}) . These dependencies can be symbolized by $\frac{\Delta y}{\Delta x}$ and involve questions such as “How does y vary when x varies?” that are useful to ask the expert.

5.3. Different Types of AKA Approaches

We distinguish three kinds of AKA: AKA from experts, (semi-)automatic AKA, and mixed AKA. The approach presented in this paper is of the first kind. Béatrice Fuchs and Alain Mille describe different knowledge types useful for CBR and, in particular, for the adaptation phase (Fuchs & Mille, 1999). A specification of the adaptation task is proposed together with a decomposition of adaptation in several subtasks (add, suppress, substitute, reorganize, etc.). This work presents adaptation at a general level useful as a guide and the way it is instantiated in several applications. Diane E. Oliver et al. represent knowledge about changes in a medical context (Oliver et al., 1999). This work is very different from ours since changes of knowledge are at the level of domain vocabulary (addition replacement and suppression of terms, changes in the term hierarchy, etc.), whereas our approach concerns therapeutic adaptations, and therefore, changes in treatment rules. In (Leake et al., 1996), AKA from experts is performed by retaining adaptations performed thanks to an interaction with the user, in order to reapply them later.

David B. Leake et al. also present an automatic AKA that consists in retaining adaptations performed automatically by the system (Leake et al., 1996). Jacek Jarmulak and Susan Craw present another approach of automatic AKA based on retaining adaptation cases (Jarmulak et al., 2001). By contrast, automatic AKA of (Hanney & Keane, 1996) learns adaptation rules. The source of this AKA is the case base: the idea is to induce adaptation rules that reflect variations between source cases. This work has inspired ongoing research in the KASIMIR project, applying principles and techniques of knowledge discovery from databases to semi-automatic AKA (d’Aquin et al., 2004b). It has appeared from our experience that these two AKAs are complementary: AKA from experts provide adaptation knowledge that is human-understandable but that needs to be instantiated while semi-automatic AKA provides operational adaptation knowledge that may lack explanation.

A future work in the KASIMIR project is mixed AKA, i.e., a combination of AKA from experts and semi-automatic AKA, with the aim of having human-understandable *and* operational adaptation knowledge. For example, let us consider the following reformulation automatically extracted from the KASIMIR case base (for the sake of clarity, this reformulation has been simplified):

if the only difference between *srce* and *tgt* is that $HR(srce) = -$ and $HR(tgt) = +$
and *tamoxifen* \in $Sol(srce)$
then $Sol(tgt) = (Sol(srce) \setminus \{tamoxifen\}) \cup \{FEC\}$

where HR stands for the hormone receptor feature of a patient and FEC is a chemotherapy drug. This reformulation instantiates the adaptation pattern of figure 2 by:

$$\underline{f} = HR \quad \underline{dec} = tamoxifen \quad \underline{dec}' = FEC$$

Based on this instantiation, the adaptation of the explanation becomes:

Since the decision *tamoxifen* is not effective on *tgt* because of problem feature HR, it can be replaced by FEC that has similar positive consequences on *tgt* as *tamoxifen* has on *srce*.

This instantiated explanation is useful first to the expert in order to validate (or correct) the reformulation and, then, it can be used by the CBR system to give an explanation when this reformulation is used.

5.4. Related work in medical CBR

This work shares some features with other studies in medical CBR. For instance, the system CARE-PARTNER applies CBR in the domain of stem-cell transplant (Bichindaritz et al., 1998). This system relies on rules, generalized cases (called *pathways*) and specific cases. Rules and pathways are applied and customized to the target problem, whereas specific cases are adapted. By contrast, in KASIMIR, general cases are adapted (currently, this system does not manage specific cases). The need of explanations in medicine is crucial: without them, the solutions provided by a CBR process are hardly accepted by physicians (Doyle et al., 2005). That is why explanations have been associated to the adaptation knowledge that is acquired for Kasimir. In these medical CBR systems, adaptation has often a key role to play, as argued in (Schmidt & Vorobieva, 2005), in which several kinds of adaptations are presented. One of them is constituted by “adaptation operators or rules”, that are confronted to the issue of knowledge acquisition bottleneck. The work presented here addresses this issue, considering reformulations as adaptation rules.

5.5. Implementation and Evaluation

Two versions of the KASIMIR system have been developed for the application of the protocols. The first one relies on an object-based representation and on an ad-hoc inference engine implementing hierarchical classification (d’Aquin et al., 2004a). In this version, a user interface is automatically generated for the elaboration of the problem—the description of the patient—and the visualization of the solution—the recommended treatment (see figure 4). Some studies, carried out by the physicians of the KASIMIR project, have already shown the strength of this system in the field of breast cancer surveillance (Rios et al., 2003). More precisely, they have pointed out a statistically significant improvement of observance of medical standards by physicians. The KASIMIR system has also recently been deployed on the Web for the purpose of an evaluation on a larger scale.

The second version of KASIMIR has been developed as a semantic portal, following the technologies and principles of the semantic Web (d’Aquin et al., 2005). It is based on a formal representation of the protocols in the ontology representation language OWL, and on standard description logic reasoning. Thus, a strong point of this version is the integration of CBR within semantic Web technologies (d’Aquin et al., 2005a). Indeed, an adaptation knowledge model has been formalized in OWL on the basis of reformulations and then used for the representation of the adaptation knowledge acquired during the AKA sessions. A prototype of a CBR service using adaptation knowledge in OWL has been implemented and is planned to be validated in a near future.

6. CONCLUSION

This paper presents AKA from experts for the system KASIMIR. This system adapts a breast cancer treatment protocol for medical cases not covered by an application of the protocol. The notions of similarity path, of intermediate problem and of reformulation play an important role for adaptation knowledge acquisition and modeling: similarity paths and intermediate problems




Figure 4: The user interface of the object-based version of KASIMIR (the characteristics are entered on the left panels and the therapeutic recommendations are displayed on the right panel).

(corresponding to virtual patients) enable decomposition of adaptations performed in simple steps, that will be modeled by reformulations containing general adaptation knowledge.

Then, this article presents elements of the adaptation knowledge that has been acquired. We have chosen to present general knowledge that can be reused for other case-based decision support systems. The adaptation patterns that have been detailed are about adaptation of inapplicable decisions and adaptation of decisions whose consequences raise problems (lack of positive consequences, too many negative consequences or a combination of both).

As presented in the discussion, semi-automatic and mixed AKAs are ongoing research directions of the KASIMIR project. Other future work concerns the use of “adaptation cases” (Leake et al., 1996) i.e., description of adaptations as they have been performed. In the framework of this project, adaptation cases could be representation of summaries of BTDC meetings. The goal of this future work is to see what these adaptation cases bring in our framework, compared to adaptation knowledge that is “compiled” in reformulations.

ACKNOWLEDGMENTS

The authors wish to thank the experts involved in the knowledge acquisition sessions presented in this paper, and in particular doctor Maria Rios and doctor Catherine Sauvagnac. They also wish to thank the three referees for their helpful comments.

REFERENCES

- Bichindaritz, I., Kansu, E., & Sullivan, K. M. (1998). Case-Based Reasoning in CARE-PARTNER: Gathering Experience for Evidence-Based Medical Practice. In Smyth, B., & Cunningham, P. (Eds.), *Fourth European Workshop on Case-Based Reasoning, EWCBR-98*, Lecture Notes in Artificial Intelligence 1488, pp. 334–345. Springer.
- d’Aquin, M., Bouthier, C., Brachais, S., Lieber, J., & Napoli, A. (2005). Knowledge Edition and Maintenance Tools for a Semantic Portal in Oncology. *International Journal of Human Computer Studies*, 62(5), 619–638.
- d’Aquin, M., Brachais, S., Lieber, J., & Napoli, A. (2004a). Decision Support and Knowledge Management in Oncology using Hierarchical Classification. In Kaiser, K., Miksch, S., & Tu, S. W. (Eds.), *Proceedings of the Symposium on Computerized Guidelines and Protocols (CGP-2004)*, Vol. 101 of *Studies in Health Technology and Informatics*, pp. 16–30. IOS Press.
- d’Aquin, M., Brachais, S., Lieber, J., & Napoli, A. (2004b). Vers une acquisition automatique de connaissances d’adaptation par examen de la base de cas — une approche fondée sur des techniques d’extraction de connaissances dans des bases de données. In Kanawati, R., Salotti, S., & Zehraoui, F. (Eds.), *Actes du douzième atelier raisonnement à partir de cas, RàPC’04*, pp. 41–52.
- d’Aquin, M., Lieber, J., & Napoli, A. (2005a). Decentralized Case-Based Reasoning for the Semantic Web. In Gil, Y., & Motta, E. (Eds.), *Proceedings of the 4th International Semantic Web Conference (ISWC 2005)*, LNCS 3729. Springer.
- d’Aquin, M., Lieber, J., & Napoli, A. (2005b). Vers une utilisation du critère pessimiste de Wald pour l’aide à la décision à partir de cas. In *actes du treizième atelier raisonnement à partir de cas, RàPC’05, plate-forme AFIA*.
- d’Aquin, M., Lieber, J., & Napoli, A. (2006). Towards a Semantic Portal for Oncology using a Description Logic with Fuzzy Concrete Domains. In Sanchez, E. (Ed.), *Fuzzy Logic in the Semantic Web* (to appear). Elsevier.
- Doyle, D., Cunningham, P., & Walsh, P. (2005). An Evaluation of the Usefulness of Explanation in a CBR System for Decision Support in Bronchiolitis Treatment. In Isabelle Bichindaritz and Cindy Marling (Ed.), *Proceedings of the Workshop on CBR in the Health Sciences of the 6th International Conference on Case-Based Reasoning (ICCBR-05)*.
- Dubois, D., Prade, H., & Sabbadin, R. (2001). Decision-theoretic foundations of qualitative possibility theory. *European Journal of Operational Research*, 128, 459–478.

- Fuchs, B., Lieber, J., Mille, A., & Napoli, A. (2000). An Algorithm for Adaptation in Case-Based Reasoning. In *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI-2000)*, Berlin, Germany, pp. 45–49.
- Fuchs, B., & Mille, A. (1999). A Knowledge-Level Task Model of Adaptation in Case-Based Reasoning. In Althoff, K.-D., Bergmann, R., & Branting, L. K. (Eds.), *Case-Based Reasoning Research and Development — Third International Conference on Case-Based Reasoning (ICCB-99)*, Lecture Notes in Artificial Intelligence 1650, pp. 118–131. Springer, Berlin.
- Hammond, K. J. (1990). Explaining and Repairing Plans That Fail. *AI Magazine*, 45(1–2), 173–228.
- Hanney, K., & Keane, M. T. (1996). Learning Adaptation Rules From a Case-Base. In Smith, I., & Faltings, B. (Eds.), *Advances in Case-Based Reasoning – Third European Workshop, EWCBR'96*, LNAI 1168, pp. 179–192. Springer Verlag, Berlin.
- Jarmulak, J., Craw, S., & Rowe, R. (2001). Using Case-Base Data to Learn Adaptation Knowledge for Design. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI'01)*, pp. 1011–1016. Morgan Kaufmann, Inc.
- Leake, D. B., Kinley, A., & Wilson, D. C. (1996). Acquiring Case Adaptation Knowledge: A Hybrid Approach. In *AAAI/IAAI*, Vol. 1, pp. 684–689.
- Lieber, J. (2002). Strong, Fuzzy and Smooth Hierarchical Classification for Case-Based Problem Solving. In van Harmelen, F. (Ed.), *Proceedings of the 15th European Conference on Artificial Intelligence (ECAI-02)*, Lyon, France, pp. 81–85. IOS Press, Amsterdam.
- Lieber, J., & Napoli, A. (1996). Using Classification in Case-Based Planning. In Wahlster, W. (Ed.), *Proceedings of the 12th European Conference on Artificial Intelligence (ECAI'96)*, Budapest, Hungary, pp. 132–136. John Wiley & Sons, Ltd.
- Melis, E., Lieber, J., & Napoli, A. (1998). Reformulation in Case-Based Reasoning. In Smyth, B., & Cunningham, P. (Eds.), *Fourth European Workshop on Case-Based Reasoning, EWCBR-98*, Lecture Notes in Artificial Intelligence 1488, pp. 172–183. Springer.
- Napoli, A., Laurenço, C., & Ducournau, R. (1994). An object-based representation system for organic synthesis planning. *International Journal of Human-Computer Studies*, 41(1/2), 5–32.
- Oliver, D. E., Shahar, Y., Musen, M. A., & Shortliffe, E. H. (1999). Representation of Change in Controlled Medical Terminologies. *Artificial Intelligence in Medicine*, 15(1), 53–76.
- Riesbeck, C. K., & Schank, R. C. (1989). *Inside Case-Based Reasoning*. Lawrence Erlbaum Associates, Inc., Hillsdale, New Jersey.
- Rios, M., Desandes, E., Bresson, B., Klein, I., Lesur, A., Boisson, F., Demange, V., & Bey, P. (2003). Référentiels de bonnes pratiques cancérologiques : étude comparative de trois supports d'aide à la proposition thérapeutique pour les cancers du sein et de la prostate en Lorraine. *Bulletin du cancer*, 90(4).

- Sauvagnac, C. (2000). *La construction de connaissances par l'utilisation et la conception de procédures. Contribution au cadre théorique des activités métafonctionnelles*. Thèse d'Université, Conservatoire National des Arts et Métiers.
- Schmidt, R., & Vorobieva, O. (2005). Adaptation and Medical Case-Based Reasoning, Focusing on Endocrine Therapy Support. In *Artificial Intelligence in Medicine (AIME'05)*. Springer, Berlin.
- Schreiber, G., Akkermans, H., Anjewierden, A., de Hoog, R., Shadbolt, N., de Velde, W. V., & Wielinga, B. (1999). *Knowledge Engineering and Management – The CommonKADS Methodology*. MIT Press.
- Smyth, B., & Keane, M. T. (1996). Using adaptation knowledge to retrieve and adapt design cases. *Knowledge-Based Systems*, 9(2), 127–135.
- Wald, A. (1950). *Statistical Decision Functions*. Wiley, New York.