

Representing Case Variations for Learning General and Specific Adaptation Rules

Fadi BADRA^a and Jean LIEBER^a

^a LORIA (UMR 7503 CNRS–INPL–INRIA–Nancy 2–UHP)
BP 239, 54 506 Vandœuvre-lès-Nancy, FRANCE
email: {badra,lieber}@loria.fr

Abstract. Adaptation is a task of case-based reasoning systems that is largely domain-dependant. This motivates the study of adaptation knowledge acquisition (AKA) that can be carried out thanks to learning processes on the variations between cases of the case base. This paper studies the representation of these variations and the impact of this representation on the AKA process, through experiments in an oncology domain.

Introduction

Case-based reasoning (CBR [7,9]) aims at solving a target problem thanks to a case base. A case is the description of a problem-solving episode which can be generally seen as a pair (problem,solution). A CBR system selects a case from the case base and then adapts the associated solution. Adaptation is a difficult task since it requires domain-dependant knowledge that needs to be acquired. The goal of adaptation knowledge acquisition is to detect and extract this knowledge. Some approaches have proved successful in learning such adaptation knowledge from the case base [2,3,6].

Our hypothesis is that this learning task can be improved by choosing an appropriate representation of the variations between cases in the case base. Introducing general knowledge on variations enables the extraction of different types of adaptation rules. In particular, a limited number of general adaptation rules may be presented to the analyst for validation. Afterwards, these rules are helpful to structure a set of specific adaptation rules.

The paper is organized as follows. Section 1 defines general notions on CBR that are used in the rest of the paper. The AKA approach of [6] is summarized in section 2. It involves some knowledge representation and learning issues — representing and learning variations and adaptation rules— that are addressed in section 3 independently of the case representation formalism. Then, this general framework is applied to an attribute-constraint formalism for representing cases (section 4). Section 5 presents some experiments in an oncology domain to

validate our hypothesis on the usefulness of choosing an appropriate representation of the variations. Section 6 discusses this work by comparing it to related work. Finally, section 7 concludes and points out some future work.

1. Definitions

Let \mathcal{L}_{pb} and \mathcal{L}_{sol} be two languages. A problem (resp. a solution) is by definition an element of \mathcal{L}_{pb} (resp. of \mathcal{L}_{sol}). The existence of a binary relation \rightsquigarrow on $\mathcal{L}_{pb} \times \mathcal{L}_{sol}$ is assumed, but not completely known in general. $Sol(pb)$ is a solution of pb if $pb \rightsquigarrow Sol(pb)$. A case is a pair $(pb, Sol(pb))$ such that $pb \rightsquigarrow Sol(pb)$. The case base is the finite set of available cases, called source cases and denoted by $(srce, Sol(srce))$. CBR aims at solving a *target problem* tgt thanks to a case base. It consists in general in retrieving a source case $(srce, Sol(srce))$ such that $srce$ is judged similar to tgt and in adapting this retrieved case in order to solve tgt : Adaptation: $(srce, Sol(srce), tgt) \mapsto Sol(tgt)$ where $(srce, Sol(srce), tgt)$ is an *adaptation problem* and $Sol(tgt)$ is a candidate solution for tgt (the relation $tgt \rightsquigarrow Sol(tgt)$ is not ensured since CBR is not a deductive reasoning). The adaptation step is based on adaptation knowledge that has to be acquired, which constitutes the adaptation knowledge acquisition issue.

2. AKA Principles

This section reformulates the main principle of the AKA approach proposed in [6]. The adaptation process is assumed to be composed of three steps:

1. **matching**: $(srce, tgt) \mapsto \Delta pb$. A representation Δpb of the variations from $srce$ to tgt is computed.
2. **AK**: $\Delta pb \mapsto \Delta sol$. The adaptation knowledge AK enables to build a representation Δsol of solution variations.
3. **modifying**: $(Sol(srce), \Delta sol) \mapsto Sol(tgt)$ such that $matching(Sol(srce), Sol(tgt)) = \Delta sol$. Δsol represents the variations from $Sol(srce)$ to the unknown solution $Sol(tgt)$ and thus enables to infer $Sol(tgt)$ from $Sol(srce)$.

Step 2 requires some adaptation knowledge. Conversely, if a set of pairs $(\Delta pb, \Delta sol)$ is available, some machine learning techniques may be used to learn AK. The idea is to exploit pairs of source cases to obtain such pairs of variations. If $(srce_k, Sol(srce_k))$ and $(srce_\ell, Sol(srce_\ell))$ are two source cases, let $\Delta pb_{k\ell} = matching(srce_k, srce_\ell)$ and $\Delta sol_{k\ell} = matching(Sol(srce_k), Sol(srce_\ell))$. Then, AK is learned with the training set $\{(\Delta pb_{k\ell}, \Delta sol_{k\ell})\}_{k\ell}$.

3. Representing and Learning Adaptation Rules

In order to make this AKA approach operational, it is necessary to be able to represent the variations Δpb and Δsol and to infer from ordered pairs of problems (resp., of solutions) Δpb (resp., Δsol), which constitutes the matching process.

3.1. Representing Variations

A variation from a problem to another is the representation of a binary relation r between problems. The language of such relations r is denoted by $\mathcal{L}_{\Delta pb}$. Thus, the semantics of any $r \in \mathcal{L}_{\Delta pb}$ is given by its extension: $\text{Ext}(r) \subseteq \mathcal{L}_{pb} \times \mathcal{L}_{pb}$. $\mathcal{L}_{\Delta pb}$ is assumed to contain the relation $\top_{\Delta pb}$ and the relation $\langle \text{srce}, \text{tgt} \rangle$ for each ordered pair $(\text{srce}, \text{tgt}) \in \mathcal{L}_{pb} \times \mathcal{L}_{pb}$. The semantics of these elements is: $\text{Ext}(\top_{\Delta pb}) = \mathcal{L}_{pb} \times \mathcal{L}_{pb}$ and $\text{Ext}(\langle \text{srce}, \text{tgt} \rangle) = \{(\text{srce}, \text{tgt})\}$.

Let \vDash be the entailment relation on $\mathcal{L}_{\Delta pb}$ (if $r, s \in \mathcal{L}_{\Delta pb}$, $r \vDash s$ means that $\text{Ext}(r) \subseteq \text{Ext}(s)$) and \equiv be the equivalence relation on $\mathcal{L}_{\Delta pb}$ ($r \equiv s$ iff $r \vDash s$ and $s \vDash r$). \vDash (and thus \equiv) is assumed to be computable. For $r \in \mathcal{L}_{\Delta pb}$, let $\text{DC}(r) = \{s \in \mathcal{L}_{\Delta pb} \mid r \vDash s\}$ be the *deductive closure* of r .

Matching two problems srce and tgt aims at identifying the relations $r \in \mathcal{L}_{\Delta pb}$ relating srce to tgt , that is, the set $\text{DC}(\langle \text{srce}, \text{tgt} \rangle)$. Since any relation $r \in \mathcal{L}_{\Delta pb}$ relating srce and tgt can be deduced from the relation $\langle \text{srce}, \text{tgt} \rangle$, matching is defined as:

$$\text{matching}(\text{srce}, \text{tgt}) = \langle \text{srce}, \text{tgt} \rangle$$

The representation of solution variations is similar: $\mathcal{L}_{\Delta \text{sol}}$ is the solution variation language containing $\top_{\Delta \text{sol}}$ and $\langle \text{Sol}(\text{srce}), \text{Sol}(\text{tgt}) \rangle$; \vDash , \equiv and DC are defined similarly on $\mathcal{L}_{\Delta \text{sol}}$; finally, $\text{matching}(\text{Sol}(\text{srce}), \text{Sol}(\text{tgt})) = \langle \text{Sol}(\text{srce}), \text{Sol}(\text{tgt}) \rangle$.

The Symmetry Hypothesis. The symmetry hypothesis is optional in our framework. It states that for each $r \in \mathcal{L}_{\Delta pb}$, there exists $r^{-1} \in \mathcal{L}_{\Delta pb}$ such that, for each $(\text{srce}, \text{tgt}) \in \mathcal{L}_{pb} \times \mathcal{L}_{pb}$, $(\text{srce}, \text{tgt}) \in \text{Ext}(r)$ iff $(\text{tgt}, \text{srce}) \in \text{Ext}(r^{-1})$. In particular, $\top_{\Delta pb}^{-1} \equiv \top_{\Delta pb}$ and $\langle \text{srce}, \text{tgt} \rangle^{-1} \equiv \langle \text{tgt}, \text{srce} \rangle$. A similar hypothesis can be made on $\mathcal{L}_{\Delta \text{sol}}$.

3.2. Representing Adaptation Rules

An *adaptation rule* is a piece of knowledge that can be used to solve adaptation problems $(\text{srce}, \text{Sol}(\text{srce}), \text{tgt}) \in \mathcal{L}_{pb} \times \mathcal{L}_{\text{sol}} \times \mathcal{L}_{pb}$: it aims at giving pieces of information about a solution $\text{Sol}(\text{tgt})$. The adaptation rules studied in this paper are ordered pairs $(r, R) \in \mathcal{L}_{\Delta pb} \times \mathcal{L}_{\Delta \text{sol}}$ that can be interpreted as follows:

$$\begin{array}{ll} \text{if} & \langle \text{srce}, \text{tgt} \rangle \vDash r \\ \text{then} & \text{Sol}(\text{tgt}) \text{ is such that } \langle \text{Sol}(\text{srce}), \text{Sol}(\text{tgt}) \rangle \vDash R \end{array} \quad (1)$$

If there exists at most one $\text{Sol}(\text{tgt}) \in \mathcal{L}_{\text{sol}}$ such that (1) holds for any adaptation problem $(\text{srce}, \text{Sol}(\text{srce}), \text{tgt})$ then the adaptation rule (r, R) is said to be *specific*. Otherwise, it is a *general* adaptation rule that is not always sufficient to solve an adaptation problem. Given two adaptation rules $\text{AR}_1 = (r_1, R_1)$ and $\text{AR}_2 = (r_2, R_2)$, AR_1 is said to be less general than AR_2 —denoted by $\text{AR}_1 \vDash \text{AR}_2$ —if $r_1 \vDash r_2$ and $R_1 \vDash R_2$. This means that AR_1 can be applied on less adaptation problems but is more accurate in the sense that the constraint on $\text{Sol}(\text{tgt})$ is stronger. The most

general adaptation rule is $(\top_{\Delta pb}, \top_{\Delta sol})$. Given an adaptation rule AR , the deductive closure of AR is the set $DC(AR)$ of adaptation rules such that each $AR' \in DC(AR)$ is more general than AR (i.e., $AR \vDash AR'$).

Under the symmetry hypothesis (on $\mathcal{L}_{\Delta pb}$ and on $\mathcal{L}_{\Delta sol}$), an inverse rule $AR^{-1} = (r^{-1}, R^{-1})$ can be associated to each rule $AR = (r, R)$.

3.3. Learning Adaptation Rules

As stated in section 2, AKA consists in using as training set TS a set of pairs $(\Delta pb_{k\ell}, \Delta sol_{k\ell})$ with

$$\begin{aligned}\Delta pb_{k\ell} &= \langle srce_k, srce_\ell \rangle \in \mathcal{L}_{\Delta pb} \\ \Delta sol_{k\ell} &= \langle Sol(srce_k), Sol(srce_\ell) \rangle \in \mathcal{L}_{\Delta sol}\end{aligned}$$

Thus, an ordered pair $AR_{k\ell} = (\Delta pb_{k\ell}, \Delta sol_{k\ell}) \in TS$ is a specific adaptation rule that solves only the adaptation problem $(srce_k, Sol(srce_k), srce_\ell)$ in a solution $Sol(srce_\ell)$. AKA consists in highlighting some adaptation rules (r, R) that are more general than a “large” number of elements of TS . More formally, with AR an adaptation rule, let

$$\text{support}(AR) = \frac{\text{card} \{AR_{k\ell} \in TS \mid AR_{k\ell} \vDash AR\}}{\text{card } TS}$$

Given a uniform distribution of probability on TS , $\text{support}(AR)$ is the probability that a random variable on TS entails AR . Learning adaptation rules aims at finding the $AR = (r, R) \in \mathcal{L}_{\Delta pb} \times \mathcal{L}_{\Delta sol}$ such that $\text{support}(AR) \geq \sigma_s$, where $\sigma_s \in [0; 1]$ is a learning parameter called the support threshold.

It can be noticed that if $AR_1 \vDash AR_2$ then $\text{support}(AR_1) \leq \text{support}(AR_2)$. Therefore, a presentation of the learned adaptation rules by decreasing support, starting from $(\top_{\Delta pb}, \top_{\Delta sol})$ whose support is 1, presents any adaptation rule before all the adaptation rules that are more specific than it.

Under the symmetry hypothesis, $\text{support}(AR^{-1}) = \text{support}(AR)$ for each adaptation rule AR . It is suggested that the two rules AR and AR^{-1} are presented together to an expert for validation.

4. Application to an Attribute-Constraint Formalism

In this section, the general framework described above is applied to an attribute-constraint formalism that extends the attribute-value formalism frequently used in CBR [7].

4.1. Representing Cases

A problem instance is described by the values it takes for some attributes. Problems are defined by specifying some sets of values these attributes may range over. For example, the problem in the domain of breast cancer treatment,

$$\text{pb} = (\text{gender}, \{\text{female}\}) \wedge (\text{age}, [50; +\infty]) \wedge (\text{s}, [4; 7])$$

represents the class of women older than 50 and for which the tumor size s is such that $4 \leq s < 7$ cm.

More formally, a *problem instance* is an element of the Cartesian product $\text{Inst}_{\text{pb}} = V_1 \times \dots \times V_m$, where V_i is a set of values. In the examples, V_i is assumed to be either \mathbb{R} (the set of real numbers), \mathbb{Z} (the set of integers), $\mathbb{B} = \{\text{true}, \text{false}\}$ (the set of Boolean) or an enumerated set (e.g. $\{\text{red}, \text{green}, \text{blue}\}$). The *attribute* a_i ($i \in \{1, 2, \dots, m\}$) is the i^{th} projection of Inst_{pb} on V_i : $a_i(x_1, \dots, x_m) = x_i$.

To each a_i is associated a language of constraints \mathcal{L}_{a_i} : $C \in \mathcal{L}_{a_i}$ is interpreted as a subset $\text{Ext}(C)$ of V_i . As an example, the constraint $[50, +\infty[$ stated on $V_i = \mathbb{Z}$ represents the set of integers that are greater or equal to 50. \mathcal{L}_{a_i} is assumed (1) to contain the constraints \top_i and \perp_i such that $\text{Ext}(\top_i) = V_i$ and $\text{Ext}(\perp_i) = \emptyset$, (2) to be closed under conjunction (if $C, D \in \mathcal{L}_{a_i}$ there exists $E \in \mathcal{L}_{a_i}$ such that $\text{Ext}(C) \cap \text{Ext}(D) = \text{Ext}(E)$). Besides, it is assumed that two distinct constraints cannot denote the same set of objects: if $\text{Ext}(C) = \text{Ext}(D)$ then $C = D$.

A *problem descriptor* is a pair $d = (a_i, C)$ where $i \in \{1, \dots, m\}$ and $C \in \mathcal{L}_{a_i}$. It is interpreted as the set of problem instances which value for the i^{th} attribute a_i is in the extension of C : $\text{Ext}((a_i, C)) = a_i^{-1}(\text{Ext}(C)) = \{x \in \text{Inst}_{\text{pb}} \mid a_i(x) \in \text{Ext}(C)\}$.

Finally, a problem is an expression of the form $\text{pb} = d_1 \wedge \dots \wedge d_p$ where the d_i are problem descriptors. A problem is interpreted by an intersection: $\text{Ext}(d_1 \wedge \dots \wedge d_p) = \text{Ext}(d_1) \cap \dots \cap \text{Ext}(d_p)$.

A problem pb is satisfiable if $\text{Ext}(\text{pb}) \neq \emptyset$. A problem pb is under normal form if $\text{pb} = d_1 \wedge \dots \wedge d_m$ where, for each $i \in \{1, \dots, m\}$, $d_i = (a_i, C)$ for some $C \in \mathcal{L}_{a_i}$. It can be shown that a problem in normal form pb is satisfiable iff it contains no occurrence of \perp_i . Moreover, every satisfiable problem pb is equivalent to a unique problem pb' in normal form: $\text{Ext}(\text{pb}) = \text{Ext}(\text{pb}')$. If $\text{pb} = (a_1, C_1) \wedge \dots \wedge (a_m, C_m)$ is satisfiable and in normal form, then $a_i(\text{Ext}(\text{pb})) = \text{Ext}(C_i)$, which justifies the notation $a_i(\text{pb}) = C_i$. For the sake of simplicity, a conjunct of a problem of the form (a_i, \top_i) may be omitted in the notation: $\text{pb} = (a_1, C_1) \wedge (a_2, \top_2) = (a_1, C_1)$.

Solutions are defined analogously with a set of attributes A_j with $j \in \{1, \dots, n\}$.

4.2. Representing Variations

In this section, problems and solutions are assumed to be satisfiable and in normal form. The problem variation language $\mathcal{L}_{\Delta\text{pb}}$ is defined by four constructors. The constructors $\top_{\Delta\text{pb}}$ and $\langle \text{srce}, \text{tgt} \rangle$ have already been defined in the general framework (cf. section 3.1). The third constructor is the conjunction: if $r, s \in \mathcal{L}_{\Delta\text{pb}}$ then $r \wedge s \in \mathcal{L}_{\Delta\text{pb}}$ and $(\text{pb}_1, \text{pb}_2) \in \text{Ext}(r \wedge s)$ iff $(\text{pb}_1, \text{pb}_2) \in \text{Ext}(r)$ and $(\text{pb}_1, \text{pb}_2) \in \text{Ext}(s)$. The last constructor is a_i^δ , where a_i is an attribute and δ represents a binary relation on \mathcal{L}_{a_i} ; δ is chosen in a language $\mathcal{L}_{\Delta a_i}$. The semantics of a_i^δ is as follows: two problems srce and tgt are related by a_i^δ if $a_i(\text{srce})$ and $a_i(\text{tgt})$ are related by δ . More formally:

$$\text{Ext}(a_i^\delta) = \{(\text{srce}, \text{tgt}) \in \mathcal{L}_{\text{pb}} \times \mathcal{L}_{\text{pb}} \mid (a_i(\text{srce}), a_i(\text{tgt})) \in \text{Ext}(\delta)\}$$

For example, let us consider the formalism \mathcal{L}_{pb} based on the attributes $a_1 = \text{gender}$, $a_2 = \text{age}$, and $a_3 = \text{s}$ ($V_1 = \{\text{female}, \text{male}\}$, $V_2 = \mathbb{Z}$, and $V_3 = \mathbb{R}$). It is

assumed that it is an attribute-*value* formalism, meaning that \mathcal{L}_{a_i} contains only \top_i , \perp_i , and the singletons $\{x\}$, for $x \in V_i$. All the $\mathcal{L}_{\Delta a_i}$ share the relations = and \neq . $\mathcal{L}_{\Delta \text{age}}$ and $\mathcal{L}_{\Delta s}$ share the relations $\delta \in \{<, \leq, \geq, >\}$ on singletons defined by $\{x\} \delta \{y\}$ if $x \delta y$ (e.g., $\{3\} < \{4\}$ since $3 < 4$). Thus, if

$$\text{srce} = (\text{gender}, \{\text{female}\}) \wedge (\text{age}, \{65\}) \wedge (s, \{3\})$$

$$\text{tgt} = (\text{gender}, \{\text{female}\}) \wedge (\text{age}, \{70\}) \wedge (s, \{2\})$$

then

$$\langle \text{srce}, \text{tgt} \rangle \models \text{gender}^= \wedge \text{age}^< \wedge \text{age}^{\leq} \wedge \text{age}^{\neq} \wedge \text{age}^> \wedge \text{s}^{\geq} \wedge \text{s}^{\neq}$$

$\mathcal{L}_{\Delta \text{sol}}$ is defined thanks to four similar constructors.

The adaptation rule language is $\mathcal{L}_{\Delta \text{pb}} \times \mathcal{L}_{\Delta \text{sol}}$. If $\text{AR} = (r, R)$ and $\text{AR}' = (r', R')$ are two adaptation rules then $\text{AR} \wedge \text{AR}'$ denotes the adaptation rule $(r \wedge r', R \wedge R')$, which is consistent with the semantics of adaptation rules given by (1).

Examples of δ . The definition of $\mathcal{L}_{\Delta \text{pb}}$ has been reduced above to the definition of a $\mathcal{L}_{\Delta a_i}$, for each attribute a_i . Although the definition of $\mathcal{L}_{\Delta a_i}$ is a knowledge acquisition issue, some examples of relations $\delta \in \mathcal{L}_{\Delta a_i}$ that may be useful are presented here.

When a set V_i is associated to an algebraic structure, the latter may be reused on singletons. For example, since \leq is a relation on \mathbb{Z} , it can be used as a relation between singletons of integers, as already mentioned above ($\{3\} < \{4\}$). Another example is related to the law $+$ on \mathbb{Z} , that is used to define the binary relation $\delta = \text{add}(\alpha)$ on \mathbb{Z} (for each $\alpha \in \mathbb{Z}$): $x \text{ add}(\alpha) y$ if $x + \alpha = y$.

Since δ relates two (representations of) sets, the classical binary relations between sets ($\subseteq, \subset, =, \supseteq, \supset$) can be suggested for elements of $\mathcal{L}_{\Delta a_i}$. Additionally, let $C \in \mathcal{L}_{a_i}$ and let $(\subseteq C \supseteq)$ be defined, for $C_1, C_2 \in \mathcal{L}_{a_i}$ by $C_1 (\subseteq C \supseteq) C_2$ if $\text{Ext}(C_1) \subseteq \text{Ext}(C) \supseteq \text{Ext}(C_2)$: C_1 and C_2 share the constraint C . Another relation is \ominus (resp., \oplus) defined by $C \ominus D$ if $C \neq \top_i$ and $D = \top_i$ (resp., $C = \top_i$ and $D \neq \top_i$). Note that $a_i^\ominus \models a_i^{\subseteq}$ and $a_i^\oplus \models a_i^{\supseteq}$. These relations can be applied in particular when $\mathcal{L}_{a_i} = 2^{V_i}$, where V_i is an enumerated set. They can also be applied when $\mathcal{L}_{\Delta a_i}$ is a finite set of atomic constraints organized in a hierarchy of root \top_i . Finally, they can be applied on intervals on, e.g., \mathbb{Z} or \mathbb{R} .

Other relations between intervals may be defined thanks to the reuse of Allen relations on temporal intervals [1]. For example, if $C_1 = [x_1, y_1]$ and $C_2 = [x_2, y_2]$ are two closed intervals on \mathbb{R} , $C_1 b C_2$ if $y_1 < x_2$, $a = b^{-1}$, $C_1 m C_2$ if $y_1 = x_2$, etc. (b , a , and m stand for *before*, *after*, and *meets*). These *qualitative* relations may be completed with *quantitative* relations such as $\text{addToBound}(\alpha, \beta)$ defined by $C_1 \text{ addToBound}(\alpha, \beta) C_2$ if $x_1 + \alpha = x_2$ and $y_1 + \beta = y_2$.

5. Experiments

Some experiments have been carried out in the oncology domain in order to evaluate the benefit of choosing an appropriate representation of the variations between cases.

5.1. Learning algorithm: CHARM

CHARM [11] is a data-mining algorithm that efficiently performs the extraction of *frequent closed itemsets* (FCIs). Given a finite set \mathcal{P} (the set of *properties* or *items*), an itemset I is a subset of \mathcal{P} : $I \in 2^{\mathcal{P}}$. The input of CHARM is a set of itemsets called the *transactions*. The support of an itemset I is the proportion $\text{support}(I)$ of the transactions T that contain I ($T \supseteq I$). An itemset I is *frequent* with respect to the threshold $\sigma_s \in [0; 1]$ if $\text{support}(I) \geq \sigma_s$. I is closed if adding to it any property alters its support: $\text{support}(I) > \text{support}(I \cup \{x\})$ for any $x \in \mathcal{P}$ such that $x \notin I$.

A benefit of using CHARM lies in its efficiency with large sets of transactions. A difficulty is that, since it operates on data (and not on pieces of knowledge), a translation of the training set TS into a set of transactions is required. Since the training set is constituted by specific adaptation rules $\text{AR}_{k\ell}$ that have to be apprehended modulo the deduction relation \vDash , the idea is to translate $\text{AR}_{k\ell}$ into the transaction $\text{DC}(\text{AR}_{k\ell})$. To be consistent with this definition, the set \mathcal{P} is set to $\bigcup_{\text{AR}_{k\ell} \in \text{TS}} \text{DC}(\text{AR}_{k\ell})$. Now, let AR_1 , AR_2 , and AR_3 be three adaptation rules and $I = \{\text{AR}_1, \text{AR}_2, \text{AR}_3\}$. If there exist exactly n $\text{AR}_{k\ell} \in \text{TS}$ such that $\text{AR}_{k\ell} \vDash \text{AR}_1 \wedge \text{AR}_2 \wedge \text{AR}_3$ then the itemset I is frequent iff $n \geq \sigma_s \times \text{card}(\text{TS})$. If, for the *same* n $\text{AR}_{k\ell} \in \text{TS}$, $\text{AR}_{k\ell} \vDash \text{AR}$ with $\text{AR} \notin I$, then I is not closed. In other words, if I is not closed, this means that I is an over-generalization that can be specialized in the adaptation rule language without loss of the coverage of the rule in the training set.

A practical problem is raised when some $\text{DC}(\text{AR}_{k\ell})$ are not finite: this leads to an infinite \mathcal{P} that CHARM cannot manage. The idea is then to restrict \mathcal{P} to a finite set and each deductive closure to $\text{DC}(\text{AR}_{k\ell}) \cap \mathcal{P}$.

5.2. Experimental Setup

The application domain of the experiments is breast cancer treatment: a problem describes a class of patients ill with breast cancer and a solution is a treatment. A problem is represented in an attribute-constraint formalism, with 22 attributes with various constraint languages \mathcal{L}_{a_i} : $2^{\mathbb{B}}$, singletons, numerical intervals, and atomic constraint hierarchies. A solution is also represented in an attribute-constraint formalism with 65 attributes. The case base contains 44 cases.

Variation languages are defined in the formalism of section 4.2, with the relations δ given as examples. The training set TS is translated into $44 \times (44 - 1) = 1892$ transactions with $\text{card } \mathcal{P} = 300$.

5.3. Results

From this set of transactions, CHARM extracted 342,994 itemsets in about 1 minute on a current PC. About 84% of this result set corresponds to itemsets with a support lower than 3% (i.e., generalizing less than 56 transactions).

Examples of learned rules. The extracted adaptation rules are organized in a hierarchy for \vDash . The expert that has to validate these rules navigates in this hierarchy. For example, the expert has found the following adaptation rule:

$$\begin{aligned} \text{AR} &= (\text{age}^{\neq} \wedge \text{ctxt}, \\ &\quad \text{nb-of-FEC-cycles}^{\neq} \wedge \text{dose-of-FEC}^{\neq} \wedge \text{Ctxt}) \end{aligned}$$

where ctxt denotes some common context the two problems srce and tgt have to share. More precisely, ctxt is a conjunction of a_i^- (and, thus, $\text{ctxt}^{-1} \equiv \text{ctxt}$), which involves that $a_i(\text{tgt}) = a_i(\text{srce})$ is a condition of the adaptation rule. Similarly, Ctxt denotes some common context for the two solutions $\text{Sol}(\text{srce})$ and $\text{Sol}(\text{tgt})$ and is a conjunction of A_j^- , which involves that $A_j(\text{Sol}(\text{tgt})) = A_j(\text{Sol}(\text{srce}))$. FEC is the name of a drug for chemotherapy that is given in several cycles (attribute nb-of-FEC-cycles with values in \mathbb{Z}) with a fixed dose in each cycle (attribute dose-of-FEC with values in \mathbb{R}).

This rule expresses that the choice of the FEC treatment depends on the age of the patient, but does not make this dependency explicit. A navigation down the hierarchy gives the following pair of rules:

$$\begin{aligned} \text{AR}_1 &= (\text{age}^b \wedge \text{ctxt}', \\ &\quad \text{nb-of-FEC-cycles}^> \wedge \text{dose-of-FEC}^> \wedge \text{Ctxt}') \\ \text{AR}_2 &= (\text{age}^a \wedge \text{ctxt}', \\ &\quad \text{nb-of-FEC-cycles}^< \wedge \text{dose-of-FEC}^< \wedge \text{Ctxt}') \end{aligned}$$

where $\text{AR}_2 \equiv \text{AR}_1^{-1}$, $\text{ctxt}' \vDash \text{ctxt}$ and $\text{Ctxt}' \vDash \text{Ctxt}$. Each of these rules states that the dependency pointed out thanks to AR is decreasing: when the age increases ($\text{age}(\text{srce})$ b $\text{age}(\text{tgt})$) the number of cycles and the dose per cycle decrease.

The expert explains this rule by (1) the fact that the growth rate of the tumour is higher for younger patients and thus requires higher doses of chemotherapy and (2) the necessity to make a compromise taking into account the life expectation (that decreases with the age) and the life quality (that decreases with the dose of FEC).

These rules could be learned only because an expressive language was chosen to represent the variations between cases.

Navigating in the results. Using an expressive language to represent the variations between cases also allowed to structure the result set and to provide the expert with efficient means of navigation in it.

Among extracted rules, only a few like AR were extracted with a high support value. These rules constitute a good starting point because their interpretation is quite straightforward, which makes the expert's work much easier during the validation phase. However, they often appear to be too general and their contexts of validity need to be refined. Most of the valid rules that were found had a fairly low support. Some rules are even valid only locally, that is in a very specific context.

To discover these rules, the generality relation \vDash on adaptation rules was used to structure the result set and provide the expert with efficient means of navigation in it. For example, a filter on the result set has been implemented that enables the expert to have access to all rules that are more specific (resp., more general) than a given rule. Using this filter, the expert has for instance been able

to visualize the set of all rules that are more specific than AR , among which are the rules AR_1 and AR_2 (since $AR_1 \vDash AR$ and $AR_2 \vDash AR$).

Towards a methodology for result exploration. A methodology for the exploration of candidate adaptation rules has emerged from these experiments. The exploration starts with a phase of elaboration in which a search context $ctxt$ is set up together with the expert. This search context is used to restrict the set of rules to search into by considering only the rules that apply to a particular medical situation. $ctxt$ takes the form of a conjunction of a_i^- with which the result set is filtered to retain only the rules AR such that $AR \vDash ctxt$. The expert is then provided with means of navigation in the remaining rules and may choose to restrict his/her search to the set of rules that are more specific than a particular rule.

6. Discussion and Related Work

Discussion. The learning process presented above is sensitive to the choice of the variation languages $\mathcal{L}_{\Delta pb}$ and $\mathcal{L}_{\Delta sol}$. This choice should be made according to a bias/variance compromise [10]. If the languages are too rich, it may prevent from learning. For example, if the language of adaptation rules $\mathcal{L}_{\Delta pb} \times \mathcal{L}_{\Delta sol}$ is closed under disjunction, then the process learns “by heart” the adaptation rule $\bigvee_{AR_{k\ell} \in TS} AR_{k\ell}$. If this language is too poor for expressing relevant adaptation rules, these latter cannot be learned.

In the context of a given application domain, what are the relevant languages $\mathcal{L}_{\Delta pb}$ and $\mathcal{L}_{\Delta sol}$? Consider, in the breast cancer treatment domain, two problems $srce$ and tgt such that $age(srce) = \{50\}$ and $age(tgt) = \{70\}$. Both relations $r_1 = age^{\delta_1}$ and $r_2 = age^{\delta_2}$, with $\delta_1 = add(20)$ and $\delta_2 = multiplyBy(1.4)$, relate $srce$ to tgt , but only r_1 is relevant for the domain experts. Indeed, the comparison of patient ages is sometimes based on differences, never (or rarely) on ratios. Then, choosing $r_1 \in \mathcal{L}_{\Delta pb}$ (and, more generally, $add(\alpha) \in \mathcal{L}_{\Delta age}$ for $\alpha \in \mathbb{Z}$) and $r_2 \notin \mathcal{L}_{\Delta pb}$, can be justified by the assumption that the adaptation rules must be expressed with a language compatible with the way the expert expresses comparisons between cases. Therefore, the choice of a language bias is a knowledge acquisition process that is similar to the process of acquiring the vocabulary for representing cases: the latter consists in pointing out the entities for representing cases whereas the former consists in pointing out the entities for representing case variations.

Related Work. The approach of adaptation learning from the case base presented in this paper is inspired from the seminal work of Kathleen Hanney and Mark T. Keane [5,6], that presents some general principles and tests them successfully in two domains. The cases of these domains are expressed in an attribute-value formalism. More precisely, with the notations defined above, for each i , $V_i = \mathbb{Z}$ or V_i is a finite (and rather small) interval of \mathbb{Z} , and \mathcal{L}_{a_i} contains only singletons. The variations are expressed by differences, i.e., with our notations, they are conjunctions of $a_i^{add(\alpha)}$ ($\alpha \in \mathbb{Z}$). Therefore, the current paper may be seen as a formalized generalization of [6]. A difference in methodology is the following: in [6], some adaptation rules are generated that are rather specific, and then, they

are generalized using some of the R. S. Michalski's generalization rules [8]. In our approach, specific and general adaptation rules are generated at the same time, which enables to organize all these rules in a hierarchy for ε , that makes navigation among them easier.

Another method for learning adaptation knowledge from the case base is presented in [2], in which the use of the extracted knowledge significantly improved the CBR process. In this work, supervised learning methods such as decision trees are applied to learn predictive models from a set of adaptation cases. These models are later used to reduce the set of adaptation cases to be used in a case-based adaptation process. There are two main differences between [2] and the work presented in this paper. First, the goal of [2] is to find adaptation knowledge in the form of adaptation cases, whereas we search for adaptation rules. Second, in [2] —and also in [6]— only pairs of *similar* source cases are considered for the learning process, whereas we favour the selection of every distinct ordered pairs of cases. This is motivated by the will of not adding another bias to the learning process.

Finally, in [3], the authors present a knowledge discovery approach for adaptation knowledge acquisition. They use a simple description logic for representing cases, that can be likened to the attribute-constraint formalism presented above. The main difference with the work presented in this paper is the richness of the variation language —and, consequently, of the adaptation rule language— we use. Indeed, with our notations, the comparisons of two cases in [3] are only based on a_i^- , a_i^o , and a_i^+ (with a_i : a problem or solution attribute).

7. Conclusion and Future Work

This paper presents a formalization of the task of adaptation rule learning from variations in the case base and shows, through experiments in the oncology domain, the benefit of an appropriate representation of variations for the purpose of this learning process. Such a representation consists in a set of binary relations between pairs of problems and pairs of solutions. When an attribute-constraint formalism (e.g., an attribute-value formalism) is used to represent cases, these relations can be based on binary relations δ between constraints C_1 and C_2 associated to the same attribute a_i . Examples of such relations δ are presented.

An ongoing work aims at reducing the number of candidate adaptation rules to be examined by the experts, while keeping the same adaptation knowledge. It is based on the notion of adaptation rule composition: the composition of (r_1, R_1) and (r_2, R_2) is the rule $(r_2 \circ r_1, R_2 \circ R_1)$, provided that $\mathcal{L}_{\Delta pb}$ and $\mathcal{L}_{\Delta sol}$ are closed under binary relation composition \circ . If S_{CAR} is the set of candidate adaptation rules that have been learned, a generative family of S_{CAR} is a set $G \subseteq S_{CAR}$ such that its closure under adaptation rule composition contains S_{CAR} . The design of algorithms giving G with high rates $\frac{\text{card } S_{CAR}}{\text{card } G}$ has begun. The first algorithms implemented give a rate close to 2 (i.e., the expert's work is divided by 2).

Assessing the quality of the learned adaptation rules may be achieved by defining an objective measure of the quality of a rule, as it is usually done for association rules [4]. A future work aims at defining such measures, and studying

whether they are useful to identify the most relevant adaptation rules in the result set. For this purpose, some of the measures of association rules, such as confidence or interest, may be adapted to the context of adaptation rules.

References

- [1] J. F. Allen, Maintaining knowledge about temporal intervals, *Communications of the ACM* **26**(11) (1983), 832–843.
- [2] S. Craw, N. Wiratung, and R. Rowe, Learning adaptation knowledge to improve case-based reasoning, *Artificial Intelligence*, **170**(16-17) (2006), 1175–1192.
- [3] M. d’Aquin, F. Badra, S. Lafrogne, J. Lieber, A. Napoli, and L. Szathmary, Case base mining for adaptation knowledge acquisition, in *Proceedings of the International Conference on Artificial Intelligence, IJCAI’07*, 750–756, 2007.
- [4] F. Guillet and H. J. Hamilton, *Quality Measures in Data Mining (Studies in Computational Intelligence)*, Springer-Verlag New York, Inc., 2007.
- [5] K. Hanney, *Learning Adaptation Rules from Cases*, Master’s thesis, Trinity College, Dublin, 1997.
- [6] K. Hanney and M. T. Keane, Learning Adaptation Rules from Cases, in *Proceedings of the 3rd European Workshop on Case-Based Reasoning, EWCBR-96*, eds., I. Smith and B. Falting, volume 1168 of LNAI, Springer, 1996.
- [7] J. Kolodner, *Case-Based Reasoning*, Morgan Kaufmann Publishers, Inc., 1993.
- [8] R. S. Michalski, A Theory and Methodology of Inductive Learning, in *Machine Learning*, 83–134, Springer-Verlag, 1983.
- [9] C. K. Riesbeck and R. C. Schank, *Inside Case-Based Reasoning*, Lawrence Erlbaum Associates, Inc., Hillsdale, New Jersey, 1989.
- [10] D. H. Wolpert, On Bias Plus Variance, *Neural Computation*, **9**(6) (1997), 1211–1243.
- [11] M. J. Zaki and C.-J. Hsiao, CHARM: An efficient algorithm for closed itemset mining, in *Proceedings of the Second SIAM International Conference on Data Mining, Arlington, VA, USA, April 11-13, 2002*, eds., Robert L. Grossman, Jiawei Han, Vipin Kumar, Heikki Mannila, and Rajeev Motwani, SIAM, 2002.