

A Lattice Machine Approach to Automated Casebase Design: Marrying Lazy and Eager Learning

Hui Wang, Werner Dubitzky, Ivo Diintsch, David Bell

School of Information and Software Engineering, University of Ulster

Newtownabbey, BT 37 OQB, N.Ireland

{H.Wang,W.Dubitzky,I.Duentsch,DA.Bell}@ulst.ac.uk

Abstract

Case-based reasoning (CBR) is concerned with solving new problems by adapting solutions that worked for similar problems in the past. Years of experience in building and fielding CBR systems have shown that the "case approach" is not free from problems. It has been realized that the knowledge engineering effort required for designing many real-world casebases can be prohibitively high. Based on the wide-spread use of databases and powerful machine learning methods, some CBR researchers have been investigating the possibility of designing casebases automatically. This paper proposes a flexible model for the automatic discovery of abstract cases from databases based on the Lattice Machine. It also proposes an efficient and effective algorithm for retrieving such cases. Besides the known benefits associated with abstract cases, the main advantages of this approach are that the discovery process is fully automated (no knowledge engineering costs).

Keywords: case-based reasoning, machine learning, knowledge acquisition, automated modeling

1 Introduction

Case-based reasoning (CBR) is concerned with solving new problems by adapting solutions that worked for similar problems in the past. CBR research is motivated by the desire to establish cognitive models to understand human thinking and behavior (psychology, cognitive science), and to build more effective, efficient, and robust computer systems that solve real-world problems (artificial intelligence). Since the early nineties, CBR has been successfully applied in a wide variety of areas. However, with almost ten years of both theoretical and applied experience in building and fielding case-based systems, it has been realized that the "case approach" is not free from problems [Leake, 1990]. Two of the more important issues that need to be addressed are outlined below:

- the implementation of *casebase maintenance*.

policies for revising the organization of casebases in order to facilitate effective and efficient future reasoning [Leake and Wilson, 1998]; and

- the *automation* of the *case, engineering* process, that is, the automatic generation of case-knowledge components from existing information (i.e., databases) [Patterson *et al.*, 1998]; the work presented in this paper will focus on this issue.

The term *case engineering* refers to the task of designing a casebase, that is, the processes of generating those components that represent the application-specific knowledge contained in a CBR system. In general, such knowledge structures describe the content-structure of cases, criteria for organizing cases within and retrieving cases from the case library (indexing scheme, similarity measures), rules for adapting case solutions, and case revision and retention schemes [Lenz *et al.*, 1998]. In addition to determining the fundamental knowledge elements, the generation process also entails the selection or extraction of those cases that will be used to populate the initial casebase.

The work presented in this paper proposes *Lattice Machine* a formal framework for learning from relations, and its application to automatically construct (case content description) and extract (case selection) abstract cases from databases. Furthermore, an efficient algorithm for the retrieval of such abstract cases is also proposed.

The remainder of the paper is organized as follows. Section 2 briefly discusses the role of abstract cases in CBR. Section 3 introduces the definitions and notational conventions, followed in Section 4 by a formal description of Lattice Machine, the pivotal part of the work in this paper. Sections 5 and 6 present efficient algorithms for extracting abstract cases and retrieving cases, based on the theoretical results in Lattice Machine. Experimental results are reported in Section 7. Section 8 concludes the paper.

2 Abstract Cases in CBR

Traditional CBR systems retrieve, reuse, and retain cases in a representation reflecting concrete problem-solving episodes. Recently, researchers have investigated

the role and use of abstract cases, e.g., [Bergmann and Wilke, 1996]. Abstract cases represent cases at a higher level of abstraction. Through abstract cases, the CBR process can be supported in several ways [Bergmann and Wilke, 1996]; those ways pertinent to the discussion in this paper are outlined below:

- Abstract cases can reduce the complexity of a casebase by substituting sets of concrete cases and thereby significantly reducing the size of the casebase. A drastically reduced casebase can improve retrieval efficiency, reduce maintenance costs, and eliminate or alleviate the notorious swamping problem [Smyth and Keane, 1995]
- Cases at higher levels of abstraction can serve as prototypes for indexing larger sets of more detailed cases. This can have profound effects on reducing retrieval times, maintenance costs, and it can promote a better user understanding of the easebase content, and facilitate explanations for the system's reasoning process.

In most situations, abstract cases are not readily available, they must be generated - manually or automatically - from concrete cases. To manually construct abstract cases will require a very high knowledge engineering effort for most applications. Whereas an automatic generation procedure requires general domain knowledge about ways of mapping concrete cases onto higher levels of abstraction.

The abstract case construct proposed in this paper is based on the concept of hypertuples. Hypertuples, i.e., abstract cases, are automatically generated through a so-called domain lattice, which is implied in a problem domain. The extraction and retrieval of abstract cases are achieved in domain lattice using the Lattice Machine.

3 Definitions and notation

To present our findings concisely and within the given page limit, we briefly introduce some notational conventions and definitions that are used throughout the paper.

3.1 Decision systems

An *information system* is a tuple $IT = \langle U, \Omega, V_x \rangle_{x \in \Omega}$, where $U = \{a_0, \dots, a_N\}$ is a nonempty finite set and $\Omega = \{x_0, \dots, x_T\}$ is a nonempty finite set of mappings $x_i : U \rightarrow V_{x_i}$.

We interpret U as a set of objects and Ω as a set of attributes or features, each of which assigns to an object a its value under the respective attribute. Let $\mathbf{V} \stackrel{\text{def}}{=} \prod_{x \in \Omega} V_x$. For $a \in U$, we let $\Omega(a) \stackrel{\text{def}}{=} \langle x(a) \rangle_{x \in \Omega} \in \mathbf{V}$. Each $\Omega(a)$ is called a *tuple*, and the collection of all tuples is denoted by D . Thus, for each $t \in D$, there is at least one $a \in U$ such that $\Omega(a) = t$.

¹Note V_{x_i} can be finite or infinite. For the latter case the domain lattice is infinite. However we are only interested in the *finite* sublattice generated from a finite casebase.

A *decision system* V is a pair (J, d) , where J is an information system as above, and $d : D \rightarrow V_d = \{d_0, \dots, d_K\}$ is an onto mapping, called a *labeling* of D ; the value $d(t)$ is called the *label* of t .

The mapping d induces a partition \mathcal{P}_d of D with the classes $\{D_0, \dots, D_K\}$, where $t \in D_i \iff d(t) = d_i$.

In this paper we consider a dataset represented as a decision system D , which can be regarded as an *initial easebase* consisting of concrete cases. Then D is the set of (descriptions of) concrete cases, d is the case solution, and V is the set of all possible concrete cases in a problem domain. Therefore each $t \in D$ is associated with a solution in the form of a class label $d(t)$.

3.2 Order and lattices

Let $\mathcal{P} = \langle P, \leq \rangle$ be a partially ordered set and $T \subseteq P$. We let $\downarrow T \stackrel{\text{def}}{=} \{y \in P : (\exists x \in T) y \leq x\}$. If no confusion can arise¹, we shall identify singleton sets with the element they contain.

Let L be a lattice, partially ordered by \leq . For $x, y \in L$, the least upper bound (or sum) is written by $x + y$ and the greatest lower bound (or product) by $x \times y$. For $A \subseteq L$, its least upper bound and greatest lower bound are denoted by $\text{lub}(A)$ and $\text{glb}(A)$ respectively. An element $a \in A$ is called *maximal* in A , if for all $x \in A$, $a \leq x$ implies $x = a$.

For $A, B \subseteq L$, we say that B *covers* A or A is *covered* by B , written as $A \preceq B$ if for each $s \in A$ there is some $t \in B$ such that $s \leq t$.

The *sublattice* of L generated from $M \subseteq L$, written by $[M]$, is $[M] = \{t \in L : \exists X \subseteq M \text{ such that } t = \text{lub}(X)\}$. The greatest element in $[M]$ is $\text{lub}(M)$. If M is finite, $[M]$ is also finite.

A com
found in

4 The Lattice Machine

This section introduces the Lattice Machine, a construct which facilitates the discovery of abstract cases from a given dataset. The discussion should also make apparent how the 'Lattice Machine' is linked to machine learning concepts.

4.1 Domain lattice

We have found that, given a dataset expressed as a decision system, an elegant mathematical structure (lattice*) is implied. This structure makes it possible to investigate CBR, machine learning, as well the relationship between the two from an algebraic perspective. In the sequel, we shall use D as described above as a generic decision system representing a dataset.

Let $\mathcal{L} \stackrel{\text{def}}{=} \prod_{x \in \Omega} 2^{V_x}$. Then $t \in \mathcal{L}$ is a vector $\langle t(x) \rangle_{x \in \Omega}$, where $t(x) \subseteq V_x$ are sets of values². The elements of \mathcal{L} are called *hypertuples*; the elements t of \mathcal{L} with $|t(x)| = 1$ for all $x \in \Omega$ are called *simple tuples*. Any set of

²Note that if $t \in \mathcal{L}$ and $x \in \Omega$, then $t(x)$ is the projection of t to its x -th component.

hypertuples is called a *hyperrelation*³. Note that V is a set of *all* simple tuples for a given problem domain, and D is the set of simple tuples described in the dataset V .

C is a lattice under the ordering

$$(1) \quad t \leq s \iff t(x) \subseteq s(x)$$

with the sum and product operations, and the maximal element (i.e., 1) given by

$$(2) \quad t + s = \langle t(x) \cup s(x) \rangle_{x \in \Omega}.$$

$$(3) \quad t \times s = \langle t(x) \cap s(x) \rangle_{x \in \Omega}.$$

$$(4) \quad 1 = \langle V_x \rangle_{x \in \Omega}.$$

C is called *domain lattice* for V .

There is a natural embedding of D into \mathcal{E} by assigning

$$\Omega(a) \mapsto \langle \{x_0(a)\}, \{x_1(a)\}, \dots, \{x_T(a)\} \rangle$$

and we shall identify D with the image of this embedding. Then $\mathbf{D} \subseteq \mathbf{V} \subseteq \mathcal{L}$.

In the context of CBR, simple tuples are concrete cases whereas hypertuples are abstract cases since hypertuples cover multiple simple tuples hence they are "abstractions" of simple tuples.

Table 1(a) is a dataset (decision system) consisting of three simple tuples, where $V_{X_1} = \{a, b\}$ and $V_{X_2} = \{0, 1\}$, and d is the labeling. Table 1(b) and (c) are sets of hypertuples, which are the least and greatest E-sets respectively for the dataset, to be defined later.

$$\mathbf{D} = \langle \langle a, 0 \rangle, \langle a, 1 \rangle, \langle b, 0 \rangle \rangle;$$

$$\mathbf{V} = \langle \langle a, 0 \rangle, \langle a, 1 \rangle, \langle b, 0 \rangle, \langle b, 1 \rangle \rangle;$$

$$\mathcal{L} = \langle \langle \emptyset, \emptyset \rangle, \langle \emptyset, \{0\} \rangle, \langle \emptyset, \{1\} \rangle, \langle \emptyset, \{0, 1\} \rangle, \langle \{a\}, \emptyset \rangle, \langle \{a\}, \{0\} \rangle, \langle \{a\}, \{1\} \rangle, \langle \{a\}, \{0, 1\} \rangle, \langle \{b\}, \emptyset \rangle, \langle \{b\}, \{0\} \rangle, \langle \{b\}, \{1\} \rangle, \langle \{b\}, \{0, 1\} \rangle, \langle \{a, b\}, \emptyset \rangle, \langle \{a, b\}, \{0\} \rangle, \langle \{a, b\}, \{1\} \rangle, \langle \{a, b\}, \{0, 1\} \rangle \rangle.$$

U	X_1	X_2	d
u_0	a	0	α
u_1	a	1	α
u_2	b	0	β

(a)

U	2^{X_1}	2^{X_2}	d
u'_0	$\{a\}$	$\{0, 1\}$	α
u'_1	$\{b\}$	$\{0\}$	β

(b)

U	2^{X_1}	2^{X_2}	d
u'_0	$\{a\}$	V_{X_2}	α
u'_1	$\{b\}$	V_{X_2}	β


(c)

Table 1: (a) A set of simple tuples in a decision system, (b) A set of hypertuples as the least E-set. (c) A set of hypertuples as the greatest E-set.

³The concept of hyperrelation has been used before in e.g. [Orlowska, 1985; Wang et al., 1998].

4.2 Equilabelledness and generalization

A dataset imposes a labeling d of D on the domain lattice C . Thus all elements in D are *labeled*, and those in $\mathbf{V} \setminus \mathbf{D}$ are *unlabeled*. This labeling can be generalized to elements in the lattice which cover D . This generalization must be consistent with d in the sense that the generalized labeling must be the same as d for $t \in D$. This renders only those generalizations acceptable which generalize d to *equilabeled* elements. Intuitively, an equilabeled element is $t \in C$ which covers at least one labeled element and all labeled elements covered by t have the same label.



Legend:

○ hypertuple ○ simple tuple ○ unseen simple tuple
+ - class labels

Figure 1: A running example.

For an illustrative example, consider the diagram in Figure 1, which depicts a small part of a domain lattice. The bottom elements (thin-lined circles) in the diagram represent simple tuples. All other elements (bold-lined circles) represent hypertuples. The original labeling d is defined only for the elements $\{H, I, J, A, A'\}$. Here, B is equilabeled as H and I are all labeled positive by d while G is unlabelled. In fact, all elements in Figure 1 are equilabeled except A, G and L : A covers elements with different labels whereas G and L cover no labeled element.

Formally, we call an element $r \in \mathcal{L}$ *equilabeled* with respect to \mathbf{D}_q , if $\emptyset \neq r \cap \mathbf{D}, \downarrow r \cap \mathbf{D} \subseteq \mathbf{D}_q$. In other words, r is equilabeled if $\downarrow r$ intersects D , and every element in this intersection is labeled d_q for some $q \leq \mathbf{K}$. Recall that \mathbf{K} is the number of classes. In this case, we say that r *G-belongs* to \mathbf{D}_q . We denote the set of all equilabeled elements G -belonging to \mathbf{D}_q by \mathcal{E}_q , and let \mathcal{E} be the set of all equilabeled elements. Note that $\mathbf{D} \subseteq \mathcal{E}$, and that $q, r \leq \mathbf{K}, q \neq r$ implies $\mathcal{E}_q \cap \mathcal{E}_r = \emptyset$.

We will now extend d over all of \mathcal{E} by setting

$$d(r) = \begin{cases} d_q, & \text{if } r \in \mathcal{E}_q, \\ \text{unknown}, & \text{otherwise.} \end{cases}$$

Now E , along with the extended labeling, can be regarded as a casebase of abstract cases (hypertuples). This is clearly too large. Since the elements in \mathcal{E} are partially ordered - some are covered by some others

we need only look at those which are not covered by any; they are *maximal*. Our wish to find maximal elements in some context leads to the following notions.

Def. 4.1. A (*generalization*) *context* (for learning) is a set P such that $\mathbf{D} \subseteq P \subseteq \mathbf{V}$. We let

$$\mathbf{M}(P) \stackrel{\text{def}}{=} \{h \in \mathcal{E} : \exists X \subseteq P, h = \text{lub}(X)\}$$

$$\mathbf{E}(P) \stackrel{\text{def}}{=} \{t : t \text{ is maximal in } \mathbf{M}(P)\}$$

$\mathbf{E}(P)$ is called the *E set* for P , and $t \in \mathcal{C}$ is said to be *in context* P if $t \in \mathbf{M}(P)$.

We observe the following lemma⁴:

Lemma 4.2. $A \subseteq B \subseteq \mathcal{L}$ implies $\mathbf{M}(A) \preceq \mathbf{M}(B)$ and $\mathbf{E}(A) \preceq \mathbf{E}(B)$.

This lemma implies $\mathbf{E}(D) \preceq \mathbf{E}(P) \preceq \mathbf{E}(V)$ for $D \subseteq P \subseteq V$. We then call $\mathbf{E}(D)$ *least E-set* and $\mathbf{E}(V)$ *greatest E-set*, which are denoted by \mathbf{E} and \mathbf{E} respectively. It is not hard to see that \mathbf{E} is the set of all maximal elements in \mathcal{E} .

Consider Figure 1 again. $\mathbf{E}(\{P, \dots, K, M\}) = \{C, M\}$, and $\mathbf{E}(\{G, \dots, M\}) = \{B, C, D, M\}$.

4.3 Interpretation of a domain lattice

Given a hypertuple t in a context P (i.e., $t \in \mathbf{M}(P)$), we need a calculus to get the remaining hypertuples in the same context (see Theorem 4.1 below). We therefore need to interpret the domain lattice in a suitable way.

Formally, let $t = \langle e_0, \dots, e_T \rangle$ be a hypertuple, with $e_i = \{e_{i_0}, \dots, e_{i_{(i)}}\}$. We think of t as an (undeterministic) description of some object a in the following way:

$$(x_0(a) = e_{0_0} \vee \dots \vee x_0(a) = e_{0_{i(0)}}) \wedge \dots \\ \dots \wedge (x_T(a) = e_{T_0} \vee \dots \vee x_T(a) = e_{T_{i(T)}}).$$

In short we have $t = e_0 \wedge \dots \wedge e_T$

Each e_i stands as a hypertuple on its own, which is interpreted as $\langle V_0, \dots, V_{i-1}, e_i, V_{i+1}, \dots, V_T \rangle$; and $e_j \wedge e_k$ is interpreted as $\text{glb}(e_j, e_k)$.

Consider a hypertuple $t = \langle e_0, \dots, e_T \rangle$. The *relative complement* or, simply, *R-complement* of t with respect to the greatest element in \mathcal{L} (i.e., 1) is $\bar{t} = \bar{e}_0 \vee \dots \vee \bar{e}_T$ and \bar{e}_j is interpreted as $V_j \setminus e_j$.

Given a hyperrelation $R = \{t_0, \dots, t_n\}$, where $t_i = \langle e_{i_0}, \dots, e_{i_T} \rangle$, we wish to calculate its R-complement \bar{R} in such a way that $R \cup \bar{R} \preceq \mathbf{V}$ and $\downarrow R \cap \downarrow \bar{R} = \emptyset$. For this purpose we interpret R as $t_0 \vee \dots \vee t_n$. Using the propositional calculus we can calculate the R-complement of R as follows. Let $X = \{\bigwedge_{i=0}^n \bar{e}_{i_j} : \text{for } j_i \in \{0, \dots, T\}\}$ $\bar{R} = \bigwedge_{i=0}^n \bar{t}_i = \bigvee_{x \in X} x$. It follows $|X| = (T+1)^{(n+1)}$. Clearly the calculation of R-complement is exponential, which makes this not practical for large R .

The above R-complement calculus can be generalized as follows. Let $t = \langle e_{t_0}, \dots, e_{t_T} \rangle$ and $h = \langle e_{h_0}, \dots, e_{h_T} \rangle$. If $t \leq h$ then the R-complement of t with respect to h , denoted by $t|h$, is calculated similarly except that $\bar{e}_{t_j} = e_{h_j} \setminus e_{t_j}$.

⁴Due to lack of space, we omit all proofs throughout the paper.

Lemma 4.3. For $R \subseteq \mathcal{C}$,

$$1. R \cup \bar{R} \preceq \mathbf{V}.$$

$$2. \downarrow R \cap \downarrow \bar{R} = \emptyset.$$

3. Let $W \stackrel{\text{def}}{=} \{w \in \mathcal{L} : \downarrow w \cap \downarrow R = \emptyset\}$. Then \bar{R} is the set of all maximal elements in W .

The first two properties guarantee that R is exclusively complementary to \bar{R} with respect to the maximal element of the domain lattice. For an example, consider Table 1(a). The table is interpreted as $u_0 \vee u_1 \vee u_2$. u_0 is interpreted as $(X_1 = a) \wedge (X_2 = 0)$. $\bar{u}_0 = \overline{\langle a, 0 \rangle} = \bar{a} \vee \bar{0} = \langle \bar{a}, V_2 \rangle \vee \langle V_1, \bar{0} \rangle$. Let $V_1 = \{a, b\}$ and $V_2 = \{0, 1\}$. Then $\bar{u}_0 = \langle b, \{0, 1\} \rangle \vee \langle \{a, b\}, 1 \rangle$, which clearly covers u_1 and u_2 , as well as an unseen tuple $(b, 1)$.

The third property says that R is the set of maximal elements not covered by \bar{R} . Consider Figure 1. Let $R = \{E\}$. Then $\bar{R} = \{G, D, M\}$.

4.4 Hypothesis space and casebase

Given a dataset we wish to have a *hypothesis* to replace the entire dataset. The hypothesis should not only cover the dataset but also generalize it. As discussed earlier, a dataset V imposes a labeling on the underlying domain lattice. The labeling can then be generalized to elements in $\mathbf{M}(P)$ for a given context P . However we do not need to use the whole $\mathbf{M}(P)$ as the hypothesis; a proper subset of it will suffice. Then a hypothesis is just a set of hypertuples, each of which is more informative than the simple tuples in the original dataset. Therefore, it is possible to consider a *casebase* as a hypothesis for a dataset. In this section we will introduce and justify some concepts through which we can precisely describe what kind of hypothesis we are aiming for.

Note that, by Lemma 4.2, $\mathbf{M}(D) \preceq \mathbf{M}(P) \preceq \mathbf{M}(V)$ for $D \subseteq P \subseteq V$. Therefore $H \subseteq \mathbf{M}(P) \Rightarrow H \subseteq \mathbf{M}(V)$. Then we have

Def. 4.4. A *hypothesis* for D is a $H \subseteq \mathbf{M}(V)$ such that $D \preceq H$. We use $\text{GEN}(D)$ to denote the set of all hypotheses for D .

Similarly, we define a *hypothesis for \mathbf{D}_q* . Note that for a hypothesis H for D , $H \cap \mathcal{E}_q$ is a hypothesis for \mathbf{D}_q . Conversely, if H_q is a hypothesis for \mathbf{D}_q for each $q \leq \mathbf{K}$, then $H \stackrel{\text{def}}{=} \bigcup_{q \leq \mathbf{K}} H_q$ is a hypothesis for D .

Since $\mathbf{M}(P)$ contains only equilabeled elements, H is *consistent* with the dataset. Since $D \preceq H$, H covers all simple tuples in the dataset.

Def. 4.5. Let H_j and H_k be two hypotheses for D . Then H_j is more general than H_k if and only if $H_k \preceq H_j$. H_j is (strictly) more general than H_k , written $(H_k \prec H_j)$, if and only if $(H_k \preceq H_j) \wedge (H_j \not\preceq H_k)$.

A hypothesis H for D is *maximally general* if and only if $H \in \text{GEN}(D)$ and there is no $H' \in \text{GEN}(D)$ such that $H \prec H'$. We denote by G the set of all maximally general hypotheses for D . In [Mitchell, 1997] G is called the *general boundary* for D .

The following lemma establishes the equivalence between G and the greatest E-set (\mathbf{E}).

Lemma 4.6. $G = \{E\}$.

This lemma says that, although there are many possible hypotheses for a given dataset, there is only one maximally general hypothesis — the greatest E-set. This hypothesis is consistent with the dataset but has the maximal coverage of unseen simple tuples (because it has the maximal context V). General boundary is a well established concept in the field of concept learning, and it has been used as an inductive bias in some concept learning algorithms [Mitchell, 1997]. The equivalence of the greatest E-set and the general boundary enables us to use the same inductive bias in automatic casebase design. Therefore our objective is to find the greatest E-set for a dataset. Our casebase design and retrieval are both associated with the greatest E-set.

However, as shown in [Haussler, 1988], the size of the general boundary can grow exponentially in the number of training examples. In the context of domain lattice, calculating the greatest E-set needs V , which is not explicitly available; instead it has to be calculated from D using the R-complement calculus discussed above. This involves calculating D , which has been shown exponential in $|D|$. It is then not practical to directly use the greatest E-set as the casebase. The following theorem guarantees that we can use a much smaller hypothesis as the casebase, but we can still use the general boundary as inductive bias. In effect it establishes the relationship between the casebase and the greatest E-set.

Theorem 4.1. *Let H be hypothesis for V , as defined in Def. 4.4. Let $V_E = \downarrow E \cap V$, and $V_H = \{t \in V : \exists h \in H, \exists g \in (h|h+t) \cup h, \text{ such that, } t \leq g, \text{ and } g \text{ is equilabeled, or unknown}\}$. Then $V_E \subseteq V_H$*

This theorem says that if a case (simple tuple) t is covered by E , then there must be h such that $t \leq h$ or $t \leq g$ for $g \in h|h+t$ with g being either equilabeled or unknown. Note that $h|h+t$ is the R-complement of h with respect to $h+t$, which is the set of all elements covered by $h+t$ but not covered by h . The CASERETRIEVE algorithm in Section 6 exploits this theorem to retrieve cases to classify new cases.

Theorem 4.1 says that any hypothesis can be used as a casebase that serves as an intermediary between a new case and the greatest E-set. Classification of a new case can then be made based on its relationship to the greatest E-set, which employs the general boundary inductive bias.

5 Case extraction

As indicated in Theorem 4.1, an expected casebase can be any hypothesis as defined in Def. 4.4. The simplest one is the dataset itself. However, checking whether the conditions are satisfied requires computing R-complements of the tuples. It is usually the case that datasets are large hence the computation cost is high. Therefore we need an algorithm to efficiently find a hypothesis, other than the dataset itself, satisfying the con-

ditions. The least E-set seems ideal since it is the E-set in the minimal context (D). However calculating the least E-set is computationally expensive. The following algorithm, CASEEXTRACT, finds, given the minimal context D , the set of elements in $M(D)$ which have disjoint coverage of D .

Given D_q and E_q as defined above.

- Initialization: let $X = D_q, H = \emptyset$.
- Repeat until X is empty:
 1. Let $h \in X$ and $X = X \setminus \{h\}$.
 2. For $g \in X$, let $X = X \setminus g$. If $h+g$ is equilabeled then $h = h+g$.
 3. Let $H = H \cup \{h\}$.

This algorithm bi-partitions X into a set of elements the sum of which is an equilabeled element, and a new X consisting of the rest of the elements. The new X is similarly bi-partitioned until X becomes empty. This process leads to a binary tree, the depth of which is a measure of the time complexity of the algorithm. In the worst case the time complexity for building the casebase for class D_q is in the order $O(|D_q|)$. Therefore the worst case complexity for building the whole casebase is $O(K \times |D_q|)$, where K is the number of classes.

Consider Figure 1. CASEEXTRACT gives the hypothesis $\{E, F, M\}$ which has disjoint coverage of the labeled elements.

6 Case retrieval

The retrieval of relevant cases from the casebase is arguably the most important process in CBR. In this section we discuss how to retrieve cases from a casebase to classify new instances. Having a casebase H as discovered by the CASEEXTRACT algorithm, we can associate a new instance $t \in V$ with a case $h \in H$ by checking whether t is covered by E through h . Then t is regarded as being in the same class as h . The CASERETRIEVE algorithm is as follows.

- Sort the elements in H in decreasing order of $|A|$ for $t \in H$ and $h = \text{lub}(X)$, which results in $h = \{h_0, \dots, h_n\}$ with h_0 having the largest coverage of D elements.
- t is classified by the first, h , in the sorted H such that the conditions in Theorem 4.1 are satisfied.
- If there is no such h_i , label t by $d(h_0)$.

The time complexity of this algorithm is dominated by calculating the R-complement of $h \in H$, as needed in Theorem 4.1. This is in the order of $O(T)$, where T is the number of attributes. In the worst case we need to do so for all h_i , in H , $i = 0, \dots, n$. Therefore the overall time complexity of the algorithm is $O(nT)$ in the worst case.

To illustrate the CASERETRIEVE algorithm, consider Figure 1. The casebase is now $\{E, F, M\}$, as discovered by CASEEXTRACT. Consider a new case G . The sum of G and E is B , which is equilabeled. Then E is retrieved

and G is labeled as positive. Clearly $G \leq B$, hence $G \not\leq E$.

7 Experiment

CASEEXTRACT and CASERETRIEVE are implemented in our CBR system, called LM. We compared LM with C4.5 using public datasets. The datasets are described in Table 2. Datasets in the upper half are from UC Irvine Machine Learning Repository; and those in the lower half are collections of documents which are used as benchmark for text mining study [Cohen and Hirsh, 1998]. The results are shown in Table 3.

Datasets	#Features #Terms	#Train	#Test	#Class
Annealing	38	798	5cv	6
Auto	25	205	5cv	6
Diabetes	8	768	5cv	2
Glass	9	214	5cv	6
Iris	4	150	5cv	3
Sonar	60	208	5cv	2
Vote	18	232	5cv	2
Memos	1014	334	10cv	11
CDroms	1133	798	10cv	6
Birdcom	674	914	10cv	22
Birdsci	1738	914	10cv	22

Table 2: Description of the datasets.

Dataset	Prediction accuracy	
	C4.5	LM
Annealing	91.8	93.6
Auto	72.2	76.1
Diabetes	72.9	71.7
Glass	81.3	82.7
Iris	91.0	96.0
Sonar	69.4	69.7
Vote	95.1	97.0
memos	57.5	59.8
cdroms	39.2	40.0
birdcom	79.6	90.4
birdsci	83.3	92.3
Average	76.0	79.0

Table 3: Prediction accuracy of C4.5 and LM

8 Summary and conclusion

The paper proposed a promising model for automating the design of CBR systems. Revolving around the notion of hypertuples (abstract cases), the proposed model presents a successful attempt at combining powerful eager techniques from machine learning with the flexible "defer-processing" philosophy characteristic for lazy methods [Aha, 1997]. On the basis of concise formal argument and empirical evaluation, it has been demonstrated that, the Lattice Machine approach constitutes an effective; and efficient mechanism to discover abstract cases in a given dataset. Abstract cases have been shown to be an effective alternative to representing the knowledge held in CBR systems [Bergmann and Wilke, 1996].

They can provide answers to issues such as casebase complexity, maintenance costs, retrieval efficiency, and user acceptance. In addition to the discovery of abstract cases, an algorithm was presented, which employs the general boundary inductive bias and ensures that the retrieval of relevant abstract cases is within the limits of reasonable time constraints. The main contribution of this work lies in the Lattice Machine's ability to discover abstract cases within a given dataset *without* requiring difficult-to-obtain domain knowledge.

References

- [Aha, 1997] D. Aha, editor. *Lazy Learning*. Kluwer Academic Pub., 1997.
- [Bergmann and Wilke, 1996] R. Bergmann and W. Wilke. On the role of abstraction in case-based reasoning. In *Proc. Advances in Case-Based Reasoning, 3rd EWCBR-96*, pages 28-41, 1996.
- [Cohen and Hirsh, 1998] William W. Cohen and Haym Hirsh. joins that generalize: Text classification using whirl. In *Proc. KDD-98, New York*, 1998.
- [Gratzer, 1978] George Gratzer. *General Lattice Theory*. Birkhauser, Basel, 1978.
- [Haussler, 1988] I). Haussler. Quantifying inductive bias: Ai learning algorithms and valiant's learning framework. *Artificial Intelligence*, 36:177-221, 1988.
- [Leake and Wilson, 1998] D.B. Leake and D.C. Wilson. Categorizing case base maintenance: Dimensions and directions. In *Proc. Advances in Case-Based Reasoning, 4th EWCBR-98*, pages 196-207, 1998.
- [Leake, 1996] D.B. Leake, editor. *Case-Based Reasoning: Experiences, Lessons & Future Directions*. MIT Press, MA, 1996.
- [Lenz et al, 1998] M. Lenz, B. Bartsch-Spoil, H-D. Burkhard, and S. Wess, editors. *Case-Based Reasoning Technology: From Foundations to Applications*. Springer-Verlag, 1998.
- [Mitchell, 1997] T. M. Mitchell. *Machine Learning*. The McGraw-Hill Companies, Inc, 1997.
- [Orlowska, 1985] Ewa Orlowska. Logic of nondeterministic information. *Studia Logica*, 44:93-102, 1985.
- [Patterson et al, 1998] D. Patterson, W. Dubitzky, S.S. Anand, and J.G. Hughes. On the automation of case base development from large databases. In *Proc. AAAI Workshop: Case-Based Reasoning Integrations*, pages 126-130, 1998.
- [Smyth and Keane, 1995] B. Smyth and M.T. Keane. Remembering to forget: A competence-preserving case deletion policy for case-based reasoning systems. In *Proc. 14th IJCAI-95*, pages 337-382, 1995.
- [Wang et al, 1998] Hui Wang, Ivo Diintsch, and David Bell. Data reduction based on hyper relations. In *Proceedings of KDD98, New York*, pages 349-353, 1998.