Abductive generalization and specialization

Chiaki Sakama
Department of Computer and Communication Sciences
Wakayama University, Sakaedani, Wakayama 640 8510, Japan
(sakama@sys.wakayama-u.ac.jp   http://www.sys.wakayama-u.ac.jp/~sakama)

Abstract. This chapter introduces new techniques called *abductive generalization* and *abductive specialization* in logic programs. Abductive generalization enables us to abduce not only specific facts but general rules accounting for positive observations. It is achieved by generalizing abductive explanations and incorporating them in a knowledge base. On the other hand, abductive specialization refines a knowledge base to recover consistency wrt negative observations. It is done by abductively finding the sources of inconsistency and specializing them with additional abductive hypotheses. A unique feature of our approach is that each procedure computes inductive generalization and specialization through abduction. The proposed techniques introduce learning ability to abduction and also realize efficient computation of induction.

1. Introduction

Abduction and induction both generate hypotheses to explain observed phenomena in an incomplete knowledge base, while they are distinguished in the following aspects. Abduction conjectures specific facts accounting for some particular observation. Those assumptions of facts are extracted using causal relations in the background knowledge base. As there are generally many possible facts which may imply the observation, candidates for hypotheses are usually pre-specified as *abducibles.* Then, the task is finding the *best* explanations from those candidates. By contrast, induction seeks regularities underlying the observed phenomena. The goal is not only explaining the current observations but discovering new knowledge for future usage. Hence induced hypotheses are general rules rather than specific facts. In constructing general rules, some constraints called *biases* are often used but candidates for hypotheses are not usually given in advance. The task is then forming *new* hypotheses using information in the background knowledge base.

Comparing two reasonings, abduction can compute explanations efficiently by specifying possible hypotheses in advance. Induction has a reasoning ability higher than abduction in the sense that it can produce new hypotheses. However, the computation of hypotheses will require a large search space and it is generally expensive. Thus abduction and induction have a trade-off between reasoning abilities and computational costs. Then, integrating two paradigms and taking advantages of each framework will provide a powerful methodology for hypothetical reasoning. Moreover, such transfers of techniques will benefit both abduction and induction. In abduction, introducing a mechanism of abducing not only facts but general rules will enhance the reasoning ability of abduction. In induction, on the other hand, it is provided a method of computing general rules abductively, which will make induction feasible.

In this chapter we propose new techniques called *abductive generalization* and *abductive specialization*. Abductive generalization provides a mechanism of abducing not only specific facts but general rules accounting for positive observations. It is achieved by computing abductive explanations and extending a knowledge base with generalized explanations. On the other hand, when a knowledge base is inconsistent with negative observations, abductive specialization refines a knowledge base to recover consistency. It is done by abductively finding the sources of inconsistency and specializing a knowledge base with additional abductive hypotheses. Abductive generalization and specialization provide methods for computing inductive hypotheses through abduction, thus contribute to a step of integrating abduction and induction in AI.

This chapter is organized as follows. Section 2 introduces an abductive framework used in this chapter. Section 3 presents a method of abductive generalization, and Section 4 provides a method of abductive specialization. Section 5 discusses related work and Section 6 concludes the chapter.

2. Preliminaries

2.1. Extended abduction

In this chapter we use an extended framework of abduction which is proposed by (Inoue and Sakama, 1995). [1]

A knowledge base K is a set of *definite clauses*

$$H \leftarrow B_1, \ldots, B_n$$

where H and B_i $(1 \leq i \leq n)$ are atoms. The atom H is the *head* and the conjunction $B_1, \ldots, B_n$ is the *body* of the clause. A clause with an empty body $H \leftarrow$ is called a *fact*. Each fact $H \leftarrow$ is identified with the atom H. A conjunction in the body is identified with the set of

[1] In (Inoue and Sakama, 1995) the framework is introduced for *nonmonotonic* theories. Here we use it for definite Horn theories with multiple observations.

atoms included in it. A clause (atom, literal) is *ground* if it contains no variable. Given a knowledge base K, a set of atoms $\mathcal{A}$ from the language of K is called *abducibles*. Abducibles specify a set of hypothetical facts. Any instance A of an element from $\mathcal{A}$ is also called an abducible and is written as $A \in \mathcal{A}$. Given a knowledge base K, its associated abducibles $\mathcal{A}$ are often omitted when their existence is clear from the context.

Let $\mathcal{O}$ be a set of ground literals. Each positive literal in $\mathcal{O}$ represents a *positive observation*, while each negative literal in $\mathcal{O}$ represents a *negative observation*. A positive observation presents an evidence that is known to be true, while a negative observation presents an evidence that is known to be false. An individual positive/negative observation is written by O^+/O^-, and the set of positive/negative observations from $\mathcal{O}$ is written by $\mathcal{O}^+/\mathcal{O}^-$, respectively.

Given a knowledge base K with abducibles $\mathcal{A}$, and observations $\mathcal{O}$, a pair of sets of atoms (E, F) is an *explanation* of $\mathcal{O}$ in K if it satisfies the following conditions:

1. $(K \cup E) \setminus F \models O^+$ for every $O^+ \in \mathcal{O}^+$,
2. $(K \cup E) \setminus F \ \cup \mathcal{O}^-$ is consistent,
3. both E and F consist of ground instances of elements from $\mathcal{A}$.

That is, a knowledge base $(K \cup E) \setminus F$ derives every positive observation and is consistent with every negative observation.[2] It should be noted that in this extended framework hypotheses can not only be added to a knowledge base but also be discarded from it to explain observations. When $\mathcal{O}^+$ contains a single observation and $\mathcal{O}^-$ and F are empty, the above definition reduces to the traditional logical framework of abduction addressed by Flach and Kakas in the introduction of this volume.

An explanation (E, F) is *minimal* if for any explanation (E', F'), $E' \subseteq E$ and $F' \subseteq F$ imply $E' = E$ and $F' = F$. It holds that $E \cap F = \emptyset$ for any minimal explanation (E, F). In this chapter explanations mean minimal explanations unless stated otherwise.

EXAMPLE 2.1. Let K be the knowledge base

$$\begin{array}{l} driving(x) \leftarrow licensed(x),\ has\text{-}car(x), \\ licensed(John) \leftarrow, \quad licensed(Mary) \leftarrow, \\ has\text{-}car(John) \leftarrow \end{array}$$

[2] In (Inoue and Sakama, 1995), explanations for a negative observation are called *anti-explanations*.

with $\mathcal{A} = \{\, licensed(x),\ has\text{-}car(x) \,\}$. Suppose we observe that $Mary$ is driving but $John$ is not these days. The situation is represented as the set of observations $\mathcal{O} = \{\, driving(Mary),\ \neg\, driving(John) \,\}$. Then, $O^+ = driving(Mary)$ is explained by assuming $has\text{-}car(Mary)$, i.e., she got a car. On the other hand, $O^- = \neg\, driving(John)$ is explained by removing either $has\text{-}car(John)$ or $licensed(John)$ from K, i.e., he lost his car or license for some reason. As a result, $\mathcal{O}$ has two alternative explanations: $(E, F) = (\{\, has\text{-}car(Mary) \,\}, \{\, has\text{-}car(John) \,\})$ and $(\{\, has\text{-}car(Mary) \,\}, \{\, licensed(John) \,\})$.

2.2. Our goal

In extended abduction both positive and negative observations are explained by introducing/removing hypotheses to/from a knowledge base. However, explanations are still selected from the pre-specified abducible facts and no new rules are constructed like induction. Our goal in this chapter is to bridge the gap between abduction and induction, and to provide a method for abducing *new rules* which explain observations. The problem is formally stated as follows.

Given a knowledge base K (with abducibles $\mathcal{A}$) and positive/negative observations $\mathcal{O}$, abduce a new knowledge base K^* such that

1. $K^* \models O^+$ for every $O^+ \in \mathcal{O}^+$,
2. $K^* \cup \mathcal{O}^-$ is consistent.

To obtain K^* we use techniques for inductive generalization and specialization.

3. Generalizing knowledge bases through abduction

3.1. Abductive generalization

This section considers knowledge bases in which only positive observations are available. Since we consider monotonic definite theories, removing facts from a knowledge base does not increase proven facts. Hence, whenever a positive observation has an explanation (E, F), F is empty. Thus, an explanation $(E, \emptyset)$ is simply written as E in this section.

EXAMPLE 3.1. One can make a profit if he/she buys a stock and the stock price goes up. Now there are four persons a, b, c, d, and each one bought a stock e, f, g, h, respectively. The situation is represented as

$$K_1: \quad profit(x) \leftarrow stock(x,y),\ up(y),$$
$$stock(a,e) \leftarrow, \quad stock(b,f) \leftarrow,$$
$$stock(c,g) \leftarrow, \quad stock(d,h) \leftarrow .$$

Suppose that abducibles are specified as $\mathcal{A} = \{\, stock(x,y),\ up(y)\,\}$. Then, given the set of positive observations

$$\mathcal{O}^+ = \{\, profit(a),\ profit(b),\ profit(c)\,\},$$

abduction computes the explanation

$$E = \{\, up(e),\ up(f),\ up(g)\,\}.$$

Thus, abduction makes each observation derivable by introducing E to K_1. On the other hand, the observations present that every person except d has already made a profit. Then, one may consider that the market is rising and d will also make a profit. In this case, one can assume the optimistic rule

$$profit(x) \leftarrow stock(x,y),$$

rather than computing similar explanations for each observation. This inference is an *inductive generalization*, which is obtained from the original rule by *dropping conditions* (Michalski, 1983).

Our goal in this section is to compute such inductive generalization through abduction. That is, given a knowledge base and positive observations, we produce a generalized knowledge base which explains the observations.

Some terminologies are introduced from (Plotkin, 1970). Two atoms are *compatible* if they have the same predicate and the same number of arguments. Let S be a set of compatible atoms. For A_1, $A_2 \in S$, A_1 is *more general* than A_2 (written $A_1 \leq A_2$) if $A_1\theta = A_2$ with some substitution θ. An atom A is a *least generalization*[3] of S if (i) $A \leq A_i$ for every $A_i \in S$, and (ii) if $A_j \leq A_i$ holds for every $A_i \in S$, then $A_j \leq A$. If A and A' are two least generalizations of S, A and A' are alphabetic variants. Given a set of atoms S, consider a decomposition $S = S_1 \cup \cdots \cup S_k$ where each S_i is a set of compatible atoms and no two atoms $A \in S_i$ and $B \in S_j$ $(i \neq j)$ are compatible. When an atom A is a least generalization of S_i, it is written as $lg(S_i) = \{\, A\,\}$. Then, the least generalization $lg(S)$ of S is defined as $lg(S) = lg(S_1) \cup \cdots \cup lg(S_k)$.

[3] In ILP literature, it is also called a least general generalization. But we use the term from (Plotkin, 1970) in this chapter.

DEFINITION 3.1. Let K be a knowledge base and $\mathcal{O}^+$ a set of positive observations. Then the following procedure computes an *abductive generalization* K^+ of K wrt. $\mathcal{O}^+$. First, put $K^+ = K$.

1. Compute an explanation E of $\mathcal{O}^+$ and its least generalization $lg(E)$.
2. For any clause C from K^+ whose body has atoms unifiable with atoms in $lg(E)$, produce a new clause C^+ by resolving C with $lg(E)$ on every such atom.[4]
3. If $C^+\theta \subseteq C$ holds for some substitution θ, replace C by C^+ in K^+. Otherwise, add C^+ to K^+.

The procedure consists of two generalization processes. The first one is the generalization of abduced explanations, and the second one is the generalization of a knowledge base. Abductive generalization weakens the conditions of existing clauses by the least generalization of the abduced explanations. The knowledge base K^+ is also an inductive generalization of K, which explains the observations $\mathcal{O}^+$.

EXAMPLE 3.2. In Example 3.1, the least generalization of E is $lg(E) = \{\, up(y)\, \}$. As the clause $C_1 : profit(x) \leftarrow stock(x,y),\ up(y)$ contains the atom $up(y)$, resolving C_1 with $up(y)$ produces the clause

$$C_1^+ :\ profit(x) \leftarrow stock(x,y).$$

Since the original clause C_1 is subsumed by the produced clause C_1^+, K_1^+ is obtained from K_1 by replacing C_1 with C_1^+:

$$\begin{array}{ll} K_1^+ : & profit(x) \leftarrow stock(x,y), \\ & stock(a,e) \leftarrow,\quad stock(b,f) \leftarrow, \\ & stock(c,g) \leftarrow,\quad stock(d,h) \leftarrow . \end{array}$$

A generalized knowledge base K^+ is also considered as a theory which is obtained from K by *partial evaluation* with respect to abduced explanations. That is, instead of explicitly introducing abductive hypotheses to a knowledge base, corresponding hypotheses are implicitly incorporated in their general forms. As a result, each observation is derived from K^+ without introducing the abduced explanation E.

THEOREM 3.1. *Let K be a knowledge base, $\mathcal{O}^+$ a set of positive observations, and E an explanation of $\mathcal{O}^+$. Then, for any ground atom A such that $A \notin E$, $K \cup E \models A$ implies $K^+ \models A$.*

---

[4] Resolving C with $lg(E)$ means resolution between C and an atom in $lg(E)$.

Proof. Let us identify knowledge bases $K \cup E$ and K^+ with their ground instances. Then, any ground clause $H \leftarrow B$ from K such that $B \cap E \neq \emptyset$ is transformed to a ground clause $H \leftarrow B'$ in K^+ where $B' = B \setminus E$. Hence, $K \cup E \models H$ implies $K^+ \models H$. Therefore, any ground atom A s.t. $A \notin E$, which is derived from $K \cup E$, is also derived from K^+.

When $A \in E$, the relation $K \cup E \models A$ does not necessarily imply $K^+ \models A$. This is because K^+ may have no clause defining A.

COROLLARY 3.2. *For any $O^+ \in \mathcal{O}^+$, $K \cup E \models O^+$ and $O^+ \notin E$ imply $K^+ \models O^+$.*

Proof. The result follows from Theorem 3.1.

By Theorem 3.1, any fact which is not in an explanation and is derived from the prior knowledge base together with the explanation, is also derived from the generalized knowledge base. This is especially the case for observations (Corollary 3.2). Note that since we consider minimal explanations, $O^+ \in E$ implies $E = O^+$. In this case, O^+ is explained by itself and K^+ does not necessarily entail O^+ in such a trivial case.

The converse of Theorem 3.1 or Corollary 3.2 does not hold in general. Indeed, K^+ possibly derives facts that are not derived from $K \cup E$. For instance, in Example 3.2, $profit(d)$ is derived from K_1^+ but not from $K_1 \cup E$. Such an increase of proven facts other than observations is called an *inductive leap*, which is a feature of inductive generalization.

3.2. Some remarks on abductive generalization

Abductive generalization introduces an inductive mechanism to abduction by constructing general rules which explain observations. From induction viewpoint, generalization K^+ is computed by modifying existing clauses in the background knowledge base K. Restricting dropping atoms to abducibles is a kind of *bias*, which reduces the number of possible generalization. Dropping abducibles is also semantically justified, since any rule containing hypotheses is considered incomplete and is subject to change.

The reliability of abductive generalization increases in proportion to the number of (compatible) positive observations. When $\mathcal{O}^+$ does not have more than one compatible observation, the procedure generalizes a knowledge base to the smallest extent. For example, if the single observation $\mathcal{O}^+ = \{ profit(a) \}$ is given to K_1 of Example 3.1, its explanation is $E = \{ up(e) \}$. In this case $lg(E) = E$, and resolving $C_1 : profit(x) \leftarrow stock(x, y),\ up(y)$ with $lg(E)$ produces the clause

$$C^+_{1'} :\ profit(x) \leftarrow stock(x, e),$$

which presents that one can make a profit if he/she buys the stock e. Since $C^+_{1'}$ does not subsume the original clause C_1, it is just added to K_1 and abductive generalization produces $K^+_{1'} = K_1 \cup \{ C^+_{1'} \}$. This is a technique of *introduction of clauses*, which is also used in inductive generalization.

Abductive generalization reduces nondeterminism in induction. There may be many possible inductive generalizations that comply with observations, then abduction leads us to hypotheses on which a knowledge base should be repaired. However, when positive observations $\mathcal{O}^+$ have multiple explanations $E_1, \ldots, E_n$ in K, a generalization K^+ exists with respect to each $lg(E_i)$.

EXAMPLE 3.3. Let K be the knowledge base

$$\begin{aligned} & p(x) \leftarrow q(x),\ s(x), \\ & q(x) \leftarrow r(x),\ t(x), \\ & s(a) \leftarrow,\ \ s(b) \leftarrow,\ \ s(c) \leftarrow, \\ & t(a) \leftarrow,\ \ t(b) \leftarrow \end{aligned}$$

with $\mathcal{A} = \{ q(x),\ r(x) \}$. Given $\mathcal{O}^+ = \{ p(a),\ p(b) \}$, there are two explanations $E_1 = \{ q(a),\ q(b) \}$ and $E_2 = \{ r(a),\ r(b) \}$. Using E_1, abductive generalization becomes

$$\begin{aligned} K^+_{E_1} : \quad & p(x) \leftarrow s(x), \\ & q(x) \leftarrow r(x),\ t(x), \\ & s(a) \leftarrow,\ \ s(b) \leftarrow,\ \ s(c) \leftarrow, \\ & t(a) \leftarrow,\ \ t(b) \leftarrow . \end{aligned}$$

On the other hand, using E_2 it becomes

$$\begin{aligned} K^+_{E_2} : \quad & p(x) \leftarrow q(x),\ s(x), \\ & q(x) \leftarrow t(x), \\ & s(a) \leftarrow,\ \ s(b) \leftarrow,\ \ s(c) \leftarrow, \\ & t(a) \leftarrow,\ \ t(b) \leftarrow . \end{aligned}$$

Here, $p(c)$ is derived from $K^+_{E_1}$ but not from $K^+_{E_2}$.

Thus, there are generalizations according to different explanations, and each generalization produces different leaps in general. This kind of nondeterminism could be reduced if further observations on the leaps are available. For instance, if $p(c)$ is known to be false, $K^{+}_{E_1}$ does not reflect the situation and $K^{+}_{E_2}$ is chosen.

Some additional condition is considered for performing abductive generalization. Suppose that the positive observations $\mathcal{O}^{+} = \{\, profit(k), profit(h) \,\}$ are given to the knowledge base K_1 of Example 3.1. Since there is no fact on k's and h's stock, abduction computes the explanation $E = \{\, stock(k, t_1),\ stock(h, t_2),\ up(t_1),\ up(t_2) \,\}$ for some instances t_1 and t_2. In this case, by the least generalization $lg(E) = \{\, stock(x, y),\ up(y) \,\}$, both $stock(x, y)$ and $up(y)$ are dropped from the body of the clause $profit(x) \leftarrow stock(x, y),\ up(y)$. The generalized clause then becomes

$$profit(x) \leftarrow,$$

saying that everyone makes a profit. To avoid such over-generalization, it is effective to restrict dropping conditions only when generalized clauses are range-restricted.[5]

4. Specializing knowledge bases through abduction

4.1. Abductive specialization

This section considers a situation where negative observations are given to a knowledge base. In a definite theory whenever a negative observation has an explanation (E, F), E is empty. This is because introducing facts to a definite theory does not help to recover consistency with respect to negative observations. Thus, an explanation $(\emptyset, F)$ is simply written as F in this section.

EXAMPLE 4.1. Consider the knowledge base $K_2 = K_1^{+}$ of Example 3.2. When the negative observation $\mathcal{O}^{-} = \{\, \neg profit(d) \,\}$ is provided, $K_2 \cup \mathcal{O}^{-}$ is inconsistent. To recover consistency of K_2 wrt. $\mathcal{O}^{-}$, abduction computes the explanation $F = \{\, stock(d, h) \,\}$.

Thus, abduction recovers consistency by removing hypothetical facts from a knowledge base. Our goal in this section is to achieve the same effect not by removing hypotheses but by specializing clauses. That is, given a knowledge base and negative observations, we produce a specialized knowledge base which is consistent with the observations.

DEFINITION 4.1. Let K be a knowledge base and $\mathcal{O}^{-}$ a set of negative observations. Then, the following procedure computes an *abductive specialization* K^{-} of K wrt. $\mathcal{O}^{-}$. First, put $K^{-} = K$.

[5] A clause is range-restricted if any variable in the clause occurs in the body.

1. Compute an explanation F of $\mathcal{O}^-$.

2. For every $A \in F$, replace the corresponding fact $C : A \leftarrow$ in K^- with the clause
$$C^- : A \leftarrow A'$$
where A' is a newly introduced abducible uniquely associated with A.

Abductive specialization abductively finds facts which are the sources of inconsistency. Then those facts are specialized by introducing newly *invented* abducibles to their conditions. The specialized knowledge base K^- is consistent with $\mathcal{O}^-$.

THEOREM 4.1. *Let K be a knowledge base and $\mathcal{O}^-$ negative observations. If $\mathcal{O}^-$ has an explanation F, $K^- \cup \mathcal{O}^-$ is consistent.*

Proof. For any $O^- = \neg G$ from $\mathcal{O}^-$, $K \setminus F \not\models G$ holds by definition. When $K \cup \{O^-\}$ is inconsistent, G is derived from K using each atom A in F. Then, rewriting every corresponding fact $A \leftarrow$ in K with $A \leftarrow A'$ in K^-, G is not derived from K^-.

EXAMPLE 4.2. Consider the knowledge base K_2 and $\mathcal{O}^- = \{\, \neg profit(d)\, \}$ of Example 4.1. By the explanation $F = \{stock(d,h)\}$, the corresponding fact $C_2 : stock(d,h) \leftarrow$ in K_2 is specialized to
$$C_2^- : \; stock(d,h) \leftarrow stock'(d,h).$$
As a result, K_2^- becomes
$$\begin{aligned} K_2^- : \quad & profit(x) \leftarrow stock(x,y), \\ & stock(a,e) \leftarrow, \quad stock(b,f) \leftarrow, \quad stock(c,g) \leftarrow, \\ & stock(d,h) \leftarrow stock'(d,h), \end{aligned}$$
where $K_2^- \cup \{\, \neg profit(d)\, \}$ is consistent. In the specialized knowledge base K_2^-, an additional hypothesis $stock'(d,h)$ is requested to conclude that d bought a (good) stock h.

Note that abduction removes explanatory facts from a knowledge base, while abductive specialization keeps information on them. This is useful for recovering the previous state of a knowledge base. For instance, if the stock h later rises and $profit(d)$ turns positive, K_2 is reproduced from K_2^- using abductive generalization, i.e., dropping the condition $stock'(d,h)$ in C_2^-.

Abductive specialization recovers consistency by modifying facts while retaining general knowledge. This is also the case of updates in *deductive databases* where every fact in a database is considered an abducible which is subject to change (Kakas and Mancarella, 1990). On the other hand, when one wants to specialize not only facts but rules in a knowledge base, abductive specialization is applied in the following manner.

Given a knowledge base K with abducibles $\mathcal{A}$, we first select *hypothetical clauses* from K which are subject to change. For any hypothetical clause

$$C_i : \; H \leftarrow B$$

in K, we consider the clause

$$C_i' : \; H \leftarrow B,\, A_i$$

where A_i is a new abducible uniquely associated with each C_i.[6] Then we consider the knowledge base

$$K' = (K \setminus \bigcup_i \{C_i\}) \cup \bigcup_i \{C_i'\} \cup \bigcup_i \{\, A_i\theta_j \leftarrow \},$$

where $A_i\theta_j$ is any ground instantiation of A_i. Abducibles associated with this new theory K' are defined as

$$\mathcal{A}' = \mathcal{A} \cup \bigcup_i \{\, A_i \,\}.$$

Then, we apply abductive specialization to K' with the following policy. If we want to specialize C_i and negative observations $\mathcal{O}^-$ have an explanation F containing $A_i\theta_j$, then we take the explanation F and specialize the corresponding fact $A_i\theta_j \leftarrow$ in K'. The resulting knowledge base K'^- has the same effect as specializing C_i in K.

EXAMPLE 4.3. Let K be the knowledge base

$$\begin{array}{l} \mathit{flies}(x) \leftarrow \mathit{bird}(x), \\ \mathit{bird}(\mathit{tweety}) \leftarrow, \quad \mathit{bird}(\mathit{polly}) \leftarrow \end{array}$$

with $\mathcal{A} = \{\, \mathit{bird}(x) \,\}$. Suppose that the first clause is a hypothetical clause which we want to revise. First, K is transformed to K':

$$\begin{array}{ll} K' : & \mathit{flies}(x) \leftarrow \mathit{bird}(x),\, p(x), \\ & p(\mathit{tweety}) \leftarrow, \quad p(\mathit{polly}) \leftarrow, \\ & \mathit{bird}(\mathit{tweety}) \leftarrow, \quad \mathit{bird}(\mathit{polly}) \leftarrow \end{array}$$

[6] This technique is called *naming* in (Poole, 1988). When C_i contains n distinct free variables $\mathbf{x} = x_1, \ldots, x_n$, an abducible $A_i = p_i(\mathbf{x})$ is associated with C_i where p_i is an n-ary predicate appearing nowhere in K.

with $\mathcal{A}' = \{\, bird(x),\, p(x) \,\}$. Given $\mathcal{O}^- = \{\, \neg\, flies(tweety) \,\}$, it has two explanations $F_1 = \{\, bird(tweety) \,\}$ and $F_2 = \{\, p(tweety) \,\}$. According to the policy, F_2 is chosen then K'^- becomes[7]

$$\begin{aligned} K'^- : \quad & flies(x) \leftarrow bird(x),\, p(x), \\ & p(tweety) \leftarrow p'(tweety), \\ & p(polly) \leftarrow, \\ & bird(tweety) \leftarrow, \quad bird(polly) \leftarrow . \end{aligned}$$

Note that K'^- has the effect of specializing the first clause in K wrt. $\mathcal{O}^-$. The revised knowledge base means that a bird flies if it satisfies an additional property p (normality or something). But tweety fails to satisfy the property by the presence of the unproved condition p'.

4.2. Combining abductive generalization and specialization

Finally, we consider combining abductive generalization and specialization in the presence of both positive and negative observations. Abductive generalization often produces an overly general theory which is inconsistent with some negative observations. Let us consider the knowledge base K_1 of Example 3.1 in which the observations $\mathcal{O} = \{\, profit(a),\, profit(b),\, profit(c),\, \neg\, profit(d) \,\}$ are given. By the positive observations $\mathcal{O}^+$ from $\mathcal{O}$, abductive generalization produces $K_2 = K_1^+$ of Example 3.2 which explains $\mathcal{O}^+$. As K_2 is inconsistent with the negative observation $\mathcal{O}^-$ from $\mathcal{O}$, abductive specialization produces K_2^- of Example 4.2 which is consistent with $\mathcal{O}^-$. Thus, in the presence of both positive and negative observations, it is considered that first generalizing a theory to derive positive observations, then specializing the theory to satisfy negative observations. Note that in this example each positive observation in $\mathcal{O}^+$ is still derived from K_2^-. However, the specialization may affect the derivation of positive observations in general.

Given a knowledge base K and positive/negative observations $\mathcal{O}$, let $K^{\pm}$ be a knowledge base obtained by combining the procedures of Definitions 3.1 and 4.1. When a generalization K^+ is obtained by $\mathcal{O}^+$, a *necessary set* of $\mathcal{O}^+$ in K^+ is defined as a minimal set F of facts such that $K^+ \setminus F \not\models G$ for some $G \in \mathcal{O}^+$. Such a necessary set is computed using abduction. That is, putting $\neg\, G$ as a negative observation in K^+, a necessary set of $\{G\}$ is obtained as an explanation F of $\neg\, G$.

[7] When there are (infinitely) many ground instantiations of $p(x)$, the set of facts $p(t) \leftarrow$ other than $p(tweety) \leftarrow$ is shortly written as $p(x) \leftarrow x \neq tweety$.

THEOREM 4.2. *Let K be a knowledge base in which positive and negative observations $\mathcal{O}$ have an explanation, and $\mathcal{O}^+ \cap \mathcal{A} = \emptyset$. For any necessary set F_1 of $\mathcal{O}^+$ in K^+ and an explanation F_2 of $\mathcal{O}^-$ in K^+, suppose $F_1 \cap F_2 = \emptyset$. Then, there is a knowledge base $K^\pm$ such that*

1. $K^\pm \models O^+$ *for every* $O^+ \in \mathcal{O}^+$,
2. $K^\pm \cup \mathcal{O}^-$ *is consistent.*

Proof. First, by the transformation from K to K^+, $K^+ \models O^+$ for every $O^+ \in \mathcal{O}^+$ (Corollary 3.2). Second, by the transformation from K^+ to $K^\pm$ using F_2, $K^\pm \cup \mathcal{O}^-$ is consistent (Theorem 4.1). In this transformation, each fact $A \in F_2$ is transformed to $A \leftarrow A'$ in $K^\pm$. As $A \notin F_1$, this rewriting does not affect the derivation of any O^+ from K^+. Therefore, $K^\pm \models O^+$ for every $O^+ \in \mathcal{O}^+$.

When a necessary set of $\mathcal{O}^+$ and every explanation of $\mathcal{O}^-$ in K^+ have an atom in common, the result of Theorem 4.2 does not hold in general.

EXAMPLE 4.4. Let K be the knowledge base

$$\begin{aligned} &p(x) \leftarrow q(x),\\ &q(x) \leftarrow r(x), s(x),\\ &s(a) \leftarrow, \quad s(b) \leftarrow \end{aligned}$$

with $\mathcal{A} = \{\, r(x),\, s(x) \,\}$. Given $\mathcal{O} = \{\, p(a),\, p(b),\, \neg\, q(a) \,\}$, the positive observations $\mathcal{O}^+ = \{\, p(a),\, p(b) \,\}$ have the explanation $E = \{\, r(a),\, r(b) \,\}$ and K is generalized to K^+:

$$\begin{aligned} &p(x) \leftarrow q(x),\\ &q(x) \leftarrow s(x),\\ &s(a) \leftarrow, \quad s(b) \leftarrow . \end{aligned}$$

In this K^+, $O^+ = p(a)$ has the necessary set $F_1 = \{\, s(a) \,\}$. On the other hand, the negative observation $\mathcal{O}^- = \{\, \neg\, q(a) \,\}$ has the explanation $F_2 = \{\, s(a) \,\}$ which is equivalent to F_1. Then, the specialized knowledge base $K^\pm$:

$$\begin{aligned} &p(x) \leftarrow q(x),\\ &q(x) \leftarrow s(x),\\ &s(a) \leftarrow s'(a),\\ &s(b) \leftarrow \end{aligned}$$

does not derive O^+.

Note that $\mathcal{O}$ has no explanation in K in the above example. Indeed, when some facts are necessary to explain a positive observation, those facts cannot be removed to satisfy negative observations. This is a necessary condition which a knowledge base should satisfy to have an explanation for both positive and negative observations. This condition is expressed as $F_1 \cap F_2 = \emptyset$ in Theorem 4.2.

5. Related work

Several systems use abduction in the process of induction. Ourston and Mooney (1990) introduce a theory refinement system called EITHER. To generalize a theory, it abductively searches the cause of failed positive observations and generalization is done by inserting a new clause or dropping most specific explanations from the conditions of a clause. Specialization is done by removing a clause or adding new antecedents to a clause. In this volume, Mooney also presents a theory refinement algorithm which generalizes a theory by deleting abduced literals or inducing new clauses. In AUDREY (Wogulis, 1991), abduction is used for identifying assumptions on which the domain theory should be repaired. In AUDREY II (Wogulis, 1991), abduced assumptions are deleted from a theory unless it covers negative examples. In EITHER the domain theory is a propositional Horn theory, while AUDREY (II) considers predicate Horn theories. Our use of abduction is similar to these systems, but is different from them on the following points. First, we generalize a theory with least generalized explanations. Second, we use extended abduction for not only generalizing theories but also specializing them. In CLINT (De Raedt and Bruynooghe, 1992) and RUTH (Adé, *et al.*, 1994) abduction supplements induction by hypothesizing factual knowledge and is also used for diagnosing a cause of integrity violation in a knowledge base. SIERES (Wirth and O'Rorke, 1992) uses abduction to infer training sets of assumptions for inductively invented predicates.

Integration of abduction and induction is also investigated by some researchers. Adé and Denecker (1995) introduce a procedure called SLDNFAI. It inductively constructs hypothetical clauses to cover positive observations and uncover negative observations. Inductive hypotheses constructed in this manner, called *inducibles*, are used for explaining observations in the abductive procedure. Dimopoulos and Kakas (1996) construct hypothetical rules to explain both positive and negative observations. Abduction is used for restricting a background theory to a relevant part on which induction is based. Abducibles are introduced to the body of a clause to refine hypotheses. Lamma *et al.*

and Inoue and Haneda in this volume introduce systems for *learning abductive logic programs*, which produce a new abductive program from observations and a background abductive theory. The above presented systems consider a general induction problem, namely, learning concepts possibly having no definitions in the background knowledge base. They use typical induction algorithms like MIS (Shapiro, 1981) or FOIL (Quinlan, 1990), in which hypotheses are constructed from scratch. However, such general induction algorithms require exhaustive search and would produce many useless hypotheses. By contrast, we restricted the problem setting and assumed a prior knowledge base having imperfect rules defining the learning concept. Then we used abduction to revise such incomplete knowledge rather than inducing arbitrary new knowledge.

6. Concluding remarks

This chapter introduced new techniques of abductive generalization and abductive specialization. They provide methods for revising a knowledge base in the face of positive and negative observations, and compute inductive generalization and specialization through abduction. Although the proposed techniques are still restrictive compared with general induction systems, they enhance the reasoning ability of abduction and also realize efficient induction. Our system is realized using the procedure of extended abduction (Inoue and Sakama, 1998).

According to Peirce (1932), "*if we are ever to learn anything or to understand phenomena at all, it must be by abduction that this is to be brought about*." In this respect, abduction is considered as a step to induction. Abductive generalization and specialization are captured as techniques based on this view, and there are possibilities of exploiting further techniques in this direction. On the application side, it is known that abduction is useful for database update and theory revision where extensional facts are subject to change (Kakas and Mancarella, 1990; Inoue and Sakama, 1995). By contrast, abductive generalization/specialization constructs intensional rules for new information, so it has potential applications to rule updates in knowledge bases. Future research also includes extending the techniques to nonmonotonic knowledge bases.

ACKNOWLEDGMENTS

The author thanks Katsumi Inoue for comments on an earlier draft of this paper.

References

Adé, H., Malfait, B. and De Raedt, L.: RUTH: an ILP theory revision system. In: *Proc. 8th Int'l Symp. on Methodologies for Intelligent Systems, Lecture Notes in Artificial Intelligence* 869, pp. 336–345, Springer-Verlag, Berlin, 1994.

Adé, H. and Denecker, M.: AILP: abductive inductive logic programming. In: *Proc. 14th Int'l Joint Conf. on Artificial Intelligence*, pp. 1201–1207, Morgan Kaufmann, California, 1995.

De Raedt, L. and Bruynooghe, M.: An overview of the interactive concept-learner and theory revisor CLINT. In: S. Muggleton (ed.), *Inductive Logic Programming*, pp. 163–191, Academic Press, London, 1992.

Dimopoulos, Y. and Kakas, A.: Abduction and inductive learning. In: L. De Raedt (ed.), *Advances in Inductive Logic Programming*, pp. 144–171, IOS Press, Amsterdam, 1996.

Inoue, K. and Sakama, C.: Abductive framework for nonmonotonic theory change. In: *Proc. 14th Int'l Joint Conf. on Artificial Intelligence*, pp. 204–210, Morgan Kaufmann, California, 1995.

Inoue, K. and Sakama, C.: Specifying transactions for extended abduction. In: *Proc. 6th Int'l Conf. on Principles of Knowledge Representation and Reasoning*, pp. 394-405, Morgan Kaufmann, California, 1998.

Kakas, A. and Mancarella, P.: Database updates through abduction. In: *Proc. 16th Int'l Conf. on Very Large Databases*, pp. 650–661, Morgan Kaufmann, California, 1990.

Michalski, R. S.: A theory and methodology of inductive learning. In: R. S. Michalski, *et al.* (eds.), *Machine Learning: An Artificial Intelligence Approach*, pp. 83–134, Morgan Kaufmann, California, 1983.

Ourston, D. and Mooney, R. J.: Changing the rules: a comprehensive approach to theory refinement. In: *Proc. 8th National Conf. on Artificial Intelligence*, pp. 815–820, MIT Press, Cambridge, 1990.

Peirce, C. S.: *Collected Papers of Charles Sanders Peirce*, vol.2, Hartshorne, C. and Weiss, P. (eds.), Harvard University Press, Cambridge, 1932.

Plotkin, G. D.: A note on inductive generalization. In: B. Meltzer and D. Michie (eds.), *Machine Intelligence* 5, pp. 153–163, Edinburgh University Press, 1970.

Poole, D.: A logical framework for default reasoning. *Artificial Intelligence*, vol.36, No.1, pp. 27–47, 1988.

Quinlan, J. R.: Learning logical definitions from relations. *Machine Learning*, vol.5, pp. 239–266, 1990.

Shapiro, E. Y.: Inductive inference of theories from facts. Research Report 192, Yale University, 1981.

Wirth, R. and O'Rorke, P.: Constraints for predicate invention. In: S. Muggleton (ed.), *Inductive Logic Programming*, pp. 299–318, Academic Press, London, 1992.

Wogulis, J.: Revising relational domain theories. In: *Proc. 8th Int'l Workshop on Machine Learning*, pp. 462–466, Morgan Kaufmann, California, 1991.

Wogulis, J. and Pazzani, M. J.: A methodology for evaluating theory revision systems: results with Audrey II. In: *Proc. 13th Int'l Joint Conference on Artificial Intelligence*, pp. 1128–1134, Morgan Kaufmann, California, 1993.