ICOT Technical Memorandum: TM-0689

TM-0689

Induction as Nonmonotonic Inference

by

Nicolas Helft

March, 1989

ICOT

Mita Kokusai Bldg. 21F
4-28 Mita 1-Chome
Minato-ku Tokyo 108 Japan

(03) 456-3191~5
Telex ICOT J32964

Institute for New Generation Computer Technology

Induction As Nonmonotonic Inference

Nicolas Helft
ICOT
1-4-28, Mita, Minato-ku, Tokyo 108, Japan
helft@icot.jp *

Abstract

This paper introduces a novel approach to similarity-based inductive reasoning. Induction is defined as inference in a nonmonotonic logic; this approach contrasts with the classical approach that consists of adding formulae to a theory in order to deduce other formulae. We point out problems arising in this setting and show how they are solved within our framework. Given a set of formulae Δ, we define the set Γ of inductive generalizations of Δ, and derive several of its properties.

1 Introduction

This paper introduces a novel approach to empirical (similarity-based) inductive reasoning. The model presented here contrasts with what we call the classical approach to induction: in this approach, a system is presented with information concerning a domain; its task is to infer hypotheses that allow it to "explain" what it observes. From a logical standpoint, what we informally call here "explain" is in fact "deduce". So the task of the system is to add formulae to a theory in order to be able to deduce other formulae. Deduction thus plays a key role in the definition of induction.

This situation can be formalized as follows:

> **Given** some background knowledge Δ and observations Θ, such that $\Delta \not\models \Theta$, **Find** Γ (called *generalizations of* Θ *with respect to* Δ) such that $\Delta \cup \Gamma \models \Theta$. (1)

(Although the problem is not always expressed in logical terms, it is always equivalent to this formulation. See for example the book by Genesereth and Nilsson; we omit additional details that are not relevant here.)

Now, this is certainly a satisfactory model of induction in the framework of scientific, rigourous thinking; but is does not seem to mirror accurately induction as the kind of ubiquitous reasoning of everyday life. For example, upon observing a number of birds and their ability to fly, people might generate the rule that *all birds fly* simply as a conclusion of the observations, grounded on their similarities, rather than as an explanation of the fact that, for example, Tweety flies knowing that it is a bird. No deductive step is involved here, so there is no reason for deduction playing such an important role in the definition of induction.

Contrasting with this, we argue that induction is a process of "jumping to conclusions" in the presence of partial information and thus a kind of inference under uncertainity. Predictably enough, it shares a basic property with certain kinds of default inference: induction assumes that the similarities between the observed data are representative of the rules governing them (we subsequently call it the ***similarity-assumption***). This assumption is like the one underlying default reasoning in that a priority is given to the information present in the database. In both cases, some form of "closing-off" the world is needed. However, there is a difference between these: loosely speaking, while in default reasoning the assumption is "what you are not told is false", in similarity-based induction, it is "what you are not told looks like what you are told".

This motivates the approach we introduce here in which, given a set of formulae, we infer other formulae called inductive generalizations of the former. Formally, the problem is

> **Given** a set of formulae Δ (we do not distinguish between background knowledge and observations), **Find** Γ (the *generalizations of* Δ) such that $\Delta \models_{IND} \Gamma$, where $\models_{IND}$ is a certain rule of inference that embodies the assumptions underlying induction.

Γ is supposed to represent all the regularities present in Δ, i.e. all the rules satisfied by its objects. In the machine learning terminology, this is often called "learning by observation and discovery", and is supposed to model a situation in which the learning system receives no assistance from a teacher. However, our aim here is not to model a particular learning situation, but rather to point out problems concerning the way inductive inference is currently formalized in

*This work is based on the author's doctoral research that was done at GRTC, Centre National de la Recherche Scientifique, Marseille, France

A Appendix: Proofs of Properties

Proposition 1: Γ contains only g-clauses.

Proof: Suppose not. Then Γ contains a clause $\phi = P \supset Q$, for which one of the following hold:

1. Q contains a variable that does not appear in P. In this case, let M be a minimal model of the initial set; the two following cases are possible:
   (a) There exists a ground instance P_i of P such that $M \models P_i$.
   Then, let x be a variable that appears in Q and not in P; if $M \models \phi$ then $M \models \forall x\, Q(x)$. Let $Q = Q1 \vee Q2$, where $Q1$ are the literals containing x and $Q2$ the rest of the literals of Q. Then $M \not\models \forall x\, Q1$, because M is finite. So $M \models Q2$. So $M \models P \supset Q2$, which subsumes ϕ. So each time a model satisfies such a clause ϕ, it satisfies a clause that subsumes it. So ϕ is not in Γ because of the last condition in the definition of generalizations.
   (b) Such an instance does not exist.
   So for every ground instance P_i of P, $M \not\models P_i$. So $Val(\phi, M) = 0$, and thus ϕ is not in Γ.
2. $\phi = P \supset Q$ contains a function symbol that is not a constant. Call l a literal in which such function symbol appears.
   In this case, if M is a minimal model of the original set, $M \not\models l$ because of [Bossu & Siegel, 1985] Property 3.2.1 which says that if a minimal model of a set of g-clauses Δ satisfies an atomic formula, this atomic formula contains only constants that appear in Δ. Now,
   (a) if l occurs in P, M satisfies no ground instance of P, so ϕ is not a generalization because $Val(\phi, M) = 0$.
   (b) if l occurs in Q, call $\phi' = \phi - \{l\}$. Then if $M \models \phi$, $M \models \phi'$. The conditions for ϕ and ϕ' to be generalizations being the same, ϕ cannot be one because it is subsumed by ϕ'. This concludes the proof.

Proposition 2: There are neither positive nor negative formulas in Γ.

Proof:

1. No positives: a positive formula is true in the minimal models of a set of formulae if and only if it is true in all models, i.e. if it can be deduced from such a set. Condition (b) in the definition of generalizations discards such formulas from Γ.
2. No negatives: if an interpretation M satisfies a negative clause $\neg P$ (i.e. $M \models \forall X \neg P$), it can never satisfy a ground instance of P.

Proposition 3: Every clause of Γ is linked.

Proof: Suppose not, let $\phi = p \wedge P \supset Q$ be such a clause, p being a non-linked literal, and call $\phi' = P \supset Q$.

If ϕ' is in Γ, ϕ is not since it is subsumed by a clause in Γ. Otherwise, one of the following conditions hold:

1. $Val(\phi', \Delta) = 0$.
   So for every minimal model M of Δ, $Val(\phi', M) = 0$. Again, one of the following must hold:
   (a) $M \not\models \phi'$. Then ϕ' has a ground instance $\phi'_i = P'_i \supset Q'_i$ not satisfied by M. Now consider the ground clause $\phi_i = p_i \wedge P'_i \supset Q'_i$, where p_i is some ground instance of literal p, not satisfied by M. (This is always possible, as M is finite). As p is not linked in ϕ, ϕ_i is necessarily a ground instance of ϕ. As M satisfies neither p_i nor ϕ'_i, it doesn't satisfy ϕ_i either. So M does not satisfy *phi* (as it does not satisfy one of its ground instances), and thus *phi* is not a generalization.
   (b) M satisfies no ground instance of P. Then M will not satisfy an instance of $p \wedge P$ either.
   (c) P is not injective. Then neither is $p \wedge P$ because p has no variable in common with P.
2. $\Delta \models \phi'$
   As $\Delta \models \phi'$ and $\phi' \models \phi$, $\Delta \models \phi$, so ϕ is not a generalization.
3. ϕ' is subsumed by a generalization.
   Then the same clause that subsumes ϕ' subsumes ϕ by transitivity, so ϕ is not a generalization.

Corollary 4:

1. No clause of Γ contains a negative ground literal.
2. No Horn clause of Γ contains a ground literal.

Proof:

1. Negative ground literals are not linked (as they contain no variables).
2. A Horn clause has only one positive literal. It cannot be a ground literal, because no negative literal would be linked. As there are no negative ground literals, there are no ground literals at all.

Proposition 5: Γ is finite.

Proof: Γ has a finite number of finite minimal models. We show that any of these can only satisfy a finite number of injective clauses.

Suppose not, and let $\{\phi_1, \phi_2, \ldots\}$ be an infinite set of such clauses. Then it is possible to construct a set $\{\psi_1, \psi_2, \ldots\}$ such that $\psi_i \subseteq \phi_i$ and all the $\{\psi_i\}$ have the same predicate symbol: as the number of predicate symbols is finite, in an infinite set there must be at least one that appears an infinite number of times.

derive some properties concerning the syntax of formulae in Γ.

We will use a subset of clausal logic and consider *discriminant* interpretations. A property of these is to interpret different ground terms by different elements of the domain (this is equivalent to making the unique-names assumption). So we will identify an interpretation with a set, that of the ground atoms to wich it assigns the value true. Note however that we do not consider Herbrand interpretations, neither do we make the domain closure assumption, so the domain will be infinite, even in the case of a finite set of constants in the initial set and no function symbols.

We consider minimal models in which the extension of every relational symbol is minimalized, i.e., M is a minimal model of a set of formulae if for no other model M', $M' \subseteq M$.

But inductive reasoning occurs over a finite set of objects. We thus need to represent the initial set with a class of formulas for which minimal models always exist; moreover, we want these minimal models to be finite. Here is a class of formulae that has such properties:

Definition: A *groundable clause* (*g-clause* for short), is a clause that satisfies the following properties:

1. its function symbols are constants.
2. every variable that appears in a positive literal also appears in a negative one.

For example, $p(x,y) \supset q(y)$ is a g-clause, while $p(x,y) \supset q(z)$ is not.

The expected properties (proofs in [Bossu & Siegel, 1985]) are the following:

Proposition:

1. Every set of clauses has a minimal model.
2. A set of g-clauses has a finite set of finite minimal models.

The definition of generalization will be in two parts: we will first define the *value* of a clause w.r.t. an interpretation, and then define the generalizations.

We first need for the generalizations to verify a technical condition.

Definition: A clause $P \supset Q$ [2] is *injective* over a set of ground atomic formulae A, whenever there exists a substitution σ, mapping the literals of P onto elements of A, such that for every pair of variables x, y of P, $\sigma(x) \neq \sigma(y)$.

For example, if

$$A = \{hand(1, clubs, clubs), wins(1), hand(2, spade, spade), wins(2)\}$$

then the clause

$$hand(x,y,z) \supset wins(x)$$

is not injective over A because both

$$\sigma_1 = \{x/1, y/clubs, z/clubs\}$$

and

$$\sigma_2 = \{x/2, y/spade, z/spade\}$$

assign the same value to y and z. Of course

$$hand(x,y,y) \supset wins(x)$$

is injective over A.

The generalizations will have to satisfy an injectivity condition over the set of atomic formulas that can be deduced from the original set. There are two reasons for this: firstly, it avoids the introduction of "unnecessary" variables in the generalization; secondly, we will show it is a necessary condition to prove an important property: the set of generalizations is finite.

Definition:

1. Let M be an interpretation and $\phi = P \supset Q$ a clause. The *value of* ϕ *in* M, denoted $Val(\phi, M)$, is defined as follows:

   1 if $M \models \phi$, $M \models P_i$, a ground instance of P, and P is injective over M. [3]

   0 otherwise.

2. Let Δ be a set of formulae and ϕ a clause. The *value of* ϕ *in* Δ, denoted $Val(\phi, \Delta)$, is:

   1 if $Val(\phi, M) = 1$ for every minimal model M of Δ.

   0 if $Val(\phi, M) = 0$ for every minimal model M of Δ.

   $\frac{1}{2}$ otherwise.

Example: Let Δ be

$deputy(tom) \vee senator(tom)$
$deputy(x) \supset corrupt(x)$
$senator(x) \supset corrupt(x)$
$rich(tom)$ $rich(bill)$

Δ has two minimal models, M_1 and M_2 that assign true to the following formulae:

$$M1 = \{deputy(tom), corrupt(tom), rich(tom), rich(bill)\}$$

$$M2 = \{senator(tom), corrupt(tom), rich(tom), rich(bill)\}$$

Let
$\phi_1 = deputy(x) \supset rich(x)$
$\phi_2 = corrupt(x) \supset rich(x)$
$\phi_3 = rich(x) \supset corrupt(x)$

[2] Notation $P \supset Q$ means $P = p_1 \wedge \ldots \wedge p_n$, the *premise* and $Q = q_1 \vee \ldots \vee q_m$, the *conclusion*. Alternatively, we will write $\neg P \vee Q$.

[3] As in [Shoham, 1987], if M is an interpretation, $M \models \phi$ means M satisfies ϕ

(a negative formula), or $deputy(x) \land friend(y,z) \supset rich(w)$ (which is not really meaningful), would not be intuitively acceptable as generalizations, as they can hardly represent interesting regularities present in the original set.

We now show that this will never occur, as the generalizations satisfy the following properties (proofs are given in the appendix):

Proposition 1: Γ contains only g-clauses.

Proposition 2: There are neither positive nor negative clauses in Γ. [5]

The following definition characterizes in a very natural way formulas that are "meaningful".

Definition:

1. The set of *linked literals* of a clause ϕ is the smallest set containing every literal l of ϕ having one of the following properties:
   (a) l is positive.
   (b) l shares a variable with a linked literal of ϕ.
2. A clause is linked if all of its literals are linked.

Loosely speaking, a link in a clause indicates a "path" between every negative literal and a positive one, through its variables. For example,

$$p(x) \land q(x,y) \supset r(y,z)$$

is linked, as $p(x)$ is linked to $q(x,y)$ which is linked to $r(y,z)$,but

$$p(x) \land q(y) \supset r(y)$$

is not, as $p(x)$ does not share a variable with a linked literal.

Note that, as this example shows, linked clauses are not necessarily g-clauses, and vice versa.

This definition is motivated by the following:

Proposition 3: Every clause of Γ is linked.

With these results it is easy to prove:

Corollary 4:

1. No clause of Γ contains a negative ground literal.
2. No Horn clause of Γ contains ground literals.

Altogether these results give a strong syntactic characterization of formulae in Γ.

Finally we prove:

Proposition 5: Γ is finite.

We motivate here the injectivity property. Consider the following set of clauses:

[5] A *positive (negative)* clause is a disjunction of positive (negative) literals.

$\phi_1 = p(x_0,x_1) \land p(x_1,x_0) \supset q(x_0)$
$\phi_2 = p(x_0,x_1) \land p(x_1,x_2) \land p(x_2,x_0) \supset q(x_0)$
$\ldots$
$\phi_k = p(x_0,x_1) \land \ldots \land p(x_k,x_0) \supset q(x_0)$

and the set

$\Delta = \{p(a,a), q(a)\}$.

The above clauses are are unordered w.r.t. model-inclusion, i.e. there is no pair of clauses ϕ_i and ϕ_j such that $\phi_i \models \phi_j$.[6] Moreover, they are all true in the only minimal model of Δ. Fortunately, none of these clauses satisfy the injectivity condition: the substitutions σ grounding its premises are such that $\sigma(x_i) = \sigma(x_j)$, for every pair of variables x_i, x_j of the clause, so none of these is a generalization of Δ. Of course $\phi_0 = p(x_0,x_0) \supset q(x_0)$ is such a generalization. Note that without the injectivity, ϕ_0 would not be in Γ, as it is subsumed by ϕ_1.

Plotkin [1970] also shows similar infinite sets of fuction-free clauses $\{\phi_1, \phi_2, \ldots\}$ ordered in a specific-to-general direction: $\phi_{i+1} \models \phi_i$. If, as in the example above, all these are true in the minimal model of some set of formulae, no generalization would existe as any clause would be subsumed by the next one.

6 Discussion and Comparison With Related Work

6.1 On the use of minimal models

In this section we explain the motivations for the definitions given so far; while doing so we examine alternatives to it and survey related work.

Given a set of formulae Δ, we want to produce all the "hidden" laws in it, i.e. all the rules verified by the objects in Δ. Within a first-order logic, these will then have the form $\forall X\ P \supset Q$. The problem is then to define the weakest conditions that constitute enough evidence to support such a rule.

These weakest conditions are:

- Δ contains an instance of P and Q.
- Δ does not contain an instance of P and $\neg Q$.

Such an approach is taken by [Delgrande, 1985]. In this situation, positive and negative information play a symmetric role, and this leads to a well-known problem in inductive logic, the Hempel paradox: a rule $P \supset Q$ being logically equivalent to its contrapositive $\neg P \supset \neg Q$, with the two conditions listed above one can generate rules with counter-intuitive support. The famous example is a white shoe being support for the rule *all crows are black*: it is supposed that we have information with which we can derive that if an

[6] Recall that without function symbols, model-inclusion is equivalent to subsumption and observe that for no pair of clauses there is a substitution that makes one clause a subset of the other.

(a negative formula), or $deputy(x) \wedge friend(y,z) \supset rich(w)$ (which is not really meaningful), would not be intuitively acceptable as generalizations, as they can hardly represent interesting regularities present in the original set.

We now show that this will never occur, as the generalizations satisfy the following properties (proofs are given in the appendix):

Proposition 1: Γ contains only g-clauses.

Proposition 2: There are neither positive nor negative clauses in Γ. [5]

The following definition characterizes in a very natural way formulas that are "meaningful".

Definition:

1. The set of *linked literals* of a clause ϕ is the smallest set containing every literal l of ϕ having one of the following properties:
   (a) l is positive.
   (b) l shares a variable with a linked literal of ϕ.
2. A clause is linked if all of its literals are linked.

Loosely speaking, a link in a clause indicates a "path" between every negative literal and a positive one, through its variables. For example,

$$p(x) \wedge q(x,y) \supset r(y,z)$$

is linked, as $p(x)$ is linked to $q(x,y)$ which is linked to $r(y,z)$,but

$$p(x) \wedge q(y) \supset r(y)$$

is not, as $p(x)$ does not share a variable with a linked literal.

Note that, as this example shows, linked clauses are not necessarily g-clauses, and vice versa.

This definition is motivated by the following:

Proposition 3: Every clause of Γ is linked.

With these results it is easy to prove:

Corollary 4:

1. No clause of Γ contains a negative ground literal.
2. No Horn clause of Γ contains ground literals.

Altogether these results give a strong syntactic characterization of formulae in Γ.

Finally we prove:

Proposition 5: Γ is finite.

We motivate here the injectivity property. Consider the following set of clauses:

$$\phi_1 = p(x_0,x_1) \wedge p(x_1,x_0) \supset q(x_0)$$
$$\phi_2 = p(x_0,x_1) \wedge p(x_1,x_2) \wedge p(x_2,x_0) \supset q(x_0)$$
$$\ldots$$
$$\phi_k = p(x_0,x_1) \wedge \ldots \wedge p(x_k,x_0) \supset q(x_0)$$

and the set

$$\Delta = \{p(a,a), q(a)\}.$$

The above clauses are are unordered w.r.t. model-inclusion, i.e. there is no pair of clauses ϕ_i and ϕ_j such that $\phi_i \models \phi_j$.[6] Moreover, they are all true in the only minimal model of Δ. Fortunately, none of these clauses satisfy the injectivity condition: the substitutions σ grounding its premises are such that $\sigma(x_i) = \sigma(x_j)$, for every pair of variables x_i, x_j of the clause, so none of these is a generalization of Δ. Of course $\phi_0 = p(x_0,x_0) \supset q(x_0)$ is such a generalization. Note that without the injectivity, ϕ_0 would not be in Γ, as it is subsumed by ϕ_1.

Plotkin [1970] also shows similar infinite sets of fuction-free clauses $\{\phi_1, \phi_2, \ldots\}$ ordered in a specific-to-general direction: $\phi_{i+1} \models \phi_i$. If, as in the example above, all these are true in the minimal model of some set of formulae, no generalization would existe as any clause would be subsumed by the next one.

6 Discussion and Comparison With Related Work

6.1 On the use of minimal models

In this section we explain the motivations for the definitions given so far; while doing so we examine alternatives to it and survey related work.

Given a set of formulae Δ, we want to produce all the "hidden" laws in it, i.e. all the rules verified by the objects in Δ. Within a first-order logic, these will then have the form $\forall X\ P \supset Q$. The problem is then to define the weakest conditions that constitute enough evidence to support such a rule.

These weakest conditions are:

- Δ contains an instance of P and Q.
- Δ does not contain an instance of P and $\neg Q$.

Such an approach is taken by [Delgrande, 1985]. In this situation, positive and negative information play a symmetric role, and this leads to a well-known problem in inductive logic, the Hempel paradox: a rule $P \supset Q$ being logically equivalent to its contrapositive $\neg P \supset \neg Q$, with the two conditions listed above one can generate rules with counter-intuitive support. The famous example is a white shoe being support for the rule *all crows are black*: it is supposed that we have information with which we can derive that if an

[5] A *positive (negative)* clause is a disjunction of positive (negative) literals.

[6] Recall that without function symbols, model-inclusion is equivalent to subsumption and observe that for no pair of clauses there is a substitution that makes one clause a subset of the other.

derive some properties concerning the syntax of formulae in Γ.

We will use a subset of clausal logic and consider *discriminant* interpretations. A property of these is to interpret different ground terms by different elements of the domain (this is equivalent to making the unique-names assumption). So we will identify an interpretation with a set, that of the ground atoms to wich it assigns the value true. Note however that we do not consider Herbrand interpretations, neither do we make the domain closure assumption, so the domain will be infinite, even in the case of a finite set of constants in the initial set and no function symbols.

We consider minimal models in which the extension of every relational symbol is minimalized, i.e., M is a minimal model of a set of formulae if for no other model M', $M' \subseteq M$.

But inductive reasoning occurs over a finite set of objects. We thus need to represent the initial set with a class of formulas for which minimal models always exist; moreover, we want these minimal models to be finite. Here is a class of formulae that has such properties:

Definition: A *groundable clause* (*g-clause* for short), is a clause that satisfies the following properties:

1. its function symbols are constants.
2. every variable that appears in a positive literal also appears in a negative one.

For example, $p(x,y) \supset q(y)$ is a g-clause, while $p(x,y) \supset q(z)$ is not.

The expected properties (proofs in [Bossu & Siegel, 1985]) are the following:

Proposition:

1. Every set of clauses has a minimal model.
2. A set of g-clauses has a finite set of finite minimal models.

The definition of generalization will be in two parts: we will first define the *value* of a clause w.r.t. an interpretation, and then define the generalizations.

We first need for the generalizations to verify a technical condition.

Definition: A clause $P \supset Q$ [2] is *injective* over a set of ground atomic formulae A, whenever there exists a substitution σ, mapping the literals of P onto elements of A, such that for every pair of variables x, y of P, $\sigma(x) \neq \sigma(y)$.

For example, if

$$A = \{hand(1, clubs, clubs), wins(1), hand(2, spade, spade), wins(2)\}$$

then the clause

$$hand(x,y,z) \supset wins(x)$$

is not injective over A because both

$$\sigma_1 = \{x/1, y/clubs, z/clubs\}$$

and

$$\sigma_2 = \{x/2, y/spade, z/spade\}$$

assign the same value to y and z. Of course

$$hand(x,y,y) \supset wins(x)$$

is injective over A.

The generalizations will have to satisfy an injectivity condition over the set of atomic formulas that can be deduced from the original set. There are two reasons for this: firstly, it avoids the introduction of "unnecessary" variables in the generalization; secondly, we will show it is a necessary condition to prove an important property: the set of generalizations is finite.

Definition:

1. Let M be an interpretation and $\phi = P \supset Q$ a clause. The *value of* ϕ *in* M, denoted $Val(\phi, M)$, is defined as follows:
   - 1 if $M \models \phi$, $M \models P_i$, a ground instance of P, and P is injective over M. [3]
   - 0 otherwise.
2. Let Δ be a set of formulae and ϕ a clause. The *value of* ϕ *in* Δ, denoted $Val(\phi, \Delta)$, is:
   - 1 if $Val(\phi, M) = 1$ for every minimal model M of Δ.
   - 0 if $Val(\phi, M) = 0$ for every minimal model M of Δ.
   - $\frac{1}{2}$ otherwise.

Example: Let Δ be

$deputy(tom) \vee senator(tom)$
$deputy(x) \supset corrupt(x)$
$senator(x) \supset corrupt(x)$
$rich(tom)$ $rich(bill)$

Δ has two minimal models, M_1 and M_2 that assign true to the following formulae:

$$M1 = \{deputy(tom), corrupt(tom), rich(tom), rich(bill)\}$$

$$M2 = \{senator(tom), corrupt(tom), rich(tom), rich(bill)\}$$

Let
$\phi_1 = deputy(x) \supset rich(x)$
$\phi_2 = corrupt(x) \supset rich(x)$
$\phi_3 = rich(x) \supset corrupt(x)$

[2] Notation $P \supset Q$ means $P = p_1 \wedge \ldots \wedge p_n$, the *premise* and $Q = q_1 \vee \ldots \vee q_m$, the *conclusion*. Alternatively, we will write $\neg P \vee Q$.

[3] As in [Shoham, 1987], if M is an interpretation, $M \models \phi$ means M satisfies ϕ

A Appendix: Proofs of Properties

Proposition 1: Γ contains only g-clauses.

Proof: Suppose not. Then Γ contains a clause $\phi = P \supset Q$, for which one of the following hold:

1. Q contains a variable that does not appear in P. In this case, let M be a minimal model of the initial set; the two following cases are possible:
   (a) There exists a ground instance P_i of P such that $M \models P_i$.
   Then, let x be a variable that appears in Q and not in P; if $M \models \phi$ then $M \models \forall x\, Q(x)$. Let $Q = Q1 \vee Q2$, where $Q1$ are the literals containing x and $Q2$ the rest of the literals of Q. Then $M \not\models \forall x\, Q1$, because M is finite. So $M \models Q2$. So $M \models P \supset Q2$, which subsumes ϕ. So each time a model satisfies such a clause ϕ, it satisfies a clause that subsumes it. So ϕ is not in Γ because of the last condition in the definition of generalizations.
   (b) Such an instance does not exist.
   So for every ground instance P_i of P, $M \not\models P_i$. So $Val(\phi, M) = 0$, and thus ϕ is not in Γ.
2. $\phi = P \supset Q$ contains a function symbol that is not a constant. Call l a literal in which such function symbol appears.
   In this case, if M is a minimal model of the original set, $M \not\models l$ because of [Bossu & Siegel, 1985] Property 3.2.1 which says that if a minimal model of a set of g-clauses Δ satisfies an atomic formula, this atomic formula contains only constants that appear in Δ. Now,
   (a) if l occurs in P, M satisfies no ground instance of P, so ϕ is not a generalization because $Val(\phi, M) = 0$.
   (b) if l occurs in Q, call $\phi' = \phi - \{l\}$. Then if $M \models \phi$, $M \models \phi'$. The conditions for ϕ and ϕ' to be generalizations being the same, ϕ cannot be one because it is subsumed by ϕ'. This concludes the proof.

Proposition 2: There are neither positive nor negative formulas in Γ.

Proof:

1. No positives: a positive formula is true in the minimal models of a set of formulae if and only if it is true in all models, i.e. if it can be deduced from such a set. Condition (b) in the definition of generalizations discards such formulas from Γ.
2. No negatives: if an interpretation M satisfies a negative clause $\neg P$ (i.e. $M \models \forall X \neg P$), it can never satisfy a ground instance of P.

Proposition 3: Every clause of Γ is linked.

Proof: Suppose not, let $\phi = p \wedge P \supset Q$ be such a clause, p being a non-linked literal, and call $\phi' = P \supset Q$.

If ϕ' is in Γ, ϕ is not since it is subsumed by a clause in Γ. Otherwise, one of the following conditions hold:

1. $Val(\phi', \Delta) = 0$.
   So for every minimal model M of Δ, $Val(\phi', M) = 0$. Again, one of the following must hold:
   (a) $M \not\models \phi'$. Then ϕ' has a ground instance $\phi'_i = P'_i \supset Q'_i$ not satisfied by M. Now consider the ground clause $\phi_i = p_i \wedge P'_i \supset Q'_i$, where p_i is some ground instance of literal p, not satisfied by M. (This is always possible, as M is finite). As p is not linked in ϕ, ϕ_i is necessarily a ground instance of ϕ. As M satisfies neither p_i nor ϕ'_i, it doesn't satisfy ϕ_i either. So M does not satisfy phi (as it does not satisfy one of its ground instances), and thus phi is not a generalization.
   (b) M satisfies no ground instance of P. Then M will not satisfy an instance of $p \wedge P$ either.
   (c) P is not injective. Then neither is $p \wedge P$ because p has no variable in common with P.
2. $\Delta \models \phi'$
   As $\Delta \models \phi'$ and $\phi' \models \phi$, $\Delta \models \phi$, so ϕ is not a generalization.
3. ϕ' is subsumed by a generalization.
   Then the same clause that subsumes ϕ' subsumes ϕ by transitivity, so ϕ is not a generalization.

Corollary 4:

1. No clause of Γ contains a negative ground literal.
2. No Horn clause of Γ contains a ground literal.

Proof:

1. Negative ground literals are not linked (as they contain no variables).
2. A Horn clause has only one positive literal. It cannot be a ground literal, because no negative literal would be linked. As there are no negative ground literals, there are no ground literals at all.

Proposition 5: Γ is finite.

Proof: Γ has a finite number of finite minimal models. We show that any of these can only satisfy a finite number of injective clauses.

Suppose not, and let $\{\phi_1, \phi_2, \ldots\}$ be an infinite set of such clauses. Then it is possible to construct a set $\{\psi_1, \psi_2, \ldots\}$ such that $\psi_i \subseteq \phi_i$ and all the $\{\psi_i\}$ have the same predicate symbol: as the number of predicate symbols is finite, in an infinite set there must be at least one that appears an infinite number of times.

Induction As Nonmonotonic Inference

Nicolas Helft
ICOT
1-4-28, Mita, Minato-ku, Tokyo 108, Japan
helft@icot.jp *

Abstract

This paper introduces a novel approach to similarity-based inductive reasoning. Induction is defined as inference in a nonmonotonic logic; this approach contrasts with the classical approach that consists of adding formulae to a theory in order to deduce other formulae. We point out problems arising in this setting and show how they are solved within our framework. Given a set of formulae Δ, we define the set Γ of inductive generalizations of Δ, and derive several of its properties.

1 Introduction

This paper introduces a novel approach to empirical (similarity-based) inductive reasoning. The model presented here contrasts with what we call the classical approach to induction: in this approach, a system is presented with information concerning a domain; its task is to infer hypotheses that allow it to "explain" what it observes. From a logical standpoint, what we informally call here "explain" is in fact "deduce". So the task of the system is to add formulae to a theory in order to be able to deduce other formulae. Deduction thus plays a key role in the definition of induction.

This situation can be formalized as follows:

> **Given** some background knowledge Δ and observations Θ, such that $\Delta \not\models \Theta$, **Find** Γ (called *generalizations of* Θ *with respect to* Δ) such that $\Delta \cup \Gamma \models \Theta$. (1)

(Although the problem is not always expressed in logical terms, it is always equivalent to this formulation. See for example the book by Genesereth and Nilsson; we omit additional details that are not relevant here.)

Now, this is certainly a satisfactory model of induction in the framework of scientific, rigourous thinking; but is does not seem to mirror accurately induction as the kind of ubiquitous reasoning of everyday life. For example, upon observing a number of birds and their ability to fly, people might generate the rule that *all birds fly* simply as a conclusion of the observations, grounded on their similarities, rather than as an explanation of the fact that, for example, Tweety flies knowing that it is a bird. No deductive step is involved here, so there is no reason for deduction playing such an important role in the definition of induction.

Contrasting with this, we argue that induction is a process of "jumping to conclusions" in the presence of partial information and thus a kind of inference under uncertainity. Predictably enough, it shares a basic property with certain kinds of default inference: induction assumes that the similarities between the observed data are representative of the rules governing them (we subsequently call it the *similarity-assumption*). This assumption is like the one underlying default reasoning in that a priority is given to the information present in the database. In both cases, some form of "closing-off" the world is needed. However, there is a difference between these: loosely speaking, while in default reasoning the assumption is "what you are not told is false", in similarity-based induction, it is "what you are not told looks like what you are told".

This motivates the approach we introduce here in which, given a set of formulae, we infer other formulae called inductive generalizations of the former. Formally, the problem is

> **Given** a set of formulae Δ (we do not distinguish between background knowledge and observations), **Find** Γ (the *generalizations of* Δ) such that $\Delta \models_{IND} \Gamma$, where $\models_{IND}$ is a certain rule of inference that embodies the assumptions underlying induction.

Γ is supposed to represent all the regularities present in Δ, i.e. all the rules satisfied by its objects. In the machine learning terminology, this is often called "learning by observation and discovery", and is supposed to model a situation in which the learning system receives no assistance from a teacher. However, our aim here is not to model a particular learning situation, but rather to point out problems concerning the way inductive inference is currently formalized in

*This work is based on the author's doctoral research that was done at GRTC, Centre National de la Recherche Scientifique, Marseille, France