The Inductive Roots of Abduction: A Task Analysis

Chris Thornton
Cognitive and Computing Sciences
University of Sussex
Brighton
BN1 9QN
UK

Email: Chris.Thornton@cogs.susx.ac.uk
WWW: http://www.cogs.susx.ac.uk
Tel: (44)1273 678856

November 25, 1996

Abstract

The paper uses a task analysis of induction to justify the idea that induction and abduction are extreme points on a single dimension of learning processes.

Topic areas: frameworks of integration of abduction and induction, abduction in inductive Machine Learning

1 Introduction

The distinction between induction and abduction can sometimes seem unclear. In the call-for-papers for this workshop[1] the two processes are described thus.

> Abduction is used to generate a reason, an explanation, for the truth of the observation in terms of hypotheses which are typically specific to the situation and individual objects at hand ... On the other hand, induction is used when we want to synthesize the information conveyed by the observations into a hypothesis that can account for all the observations together in a common way.

[1]ECAI-96 workshop on the relationship between induction and abduction.

On this account, both processes ultimately involve the identification of hypotheses which account for (the truth of) observations. The uninitiated might therefore feel that the distinction being made is somewhat artificial. However, there are various ways in which we can make it more concrete.

In this paper I present a task analysis of induction and show how it leads to a model of the induction process which logically entails processes plausibly regarded as abductive. Although the approach I take may not be the only (or the best) way of formalising the connection between induction and abduction, it does have the advantage of relative simplicity.

2 Inductive Task Analysis

Imagine we have a body of data D, as shown in Table 1. Each datum in D (i.e., each row) is made up of the values of variables x_1, x_2, x_3, x_4 and x_5. One of the values of x_3 is missing (see the '?' in the x_3 column). Can we use the other data to predict this missing value? In other words, can we empirically *induce* the missing value from the data which are provided?

x_1	x_2	x_3	x_4	x_5
c	d	f	a	b
a	b	h	d	b
e	c	h	d	e
c	b	f	a	e
a	c	f	d	e
c	c	?	a	e
b	c	f	a	e
b	d	h	d	e
e	d	f	a	c
a	c	h	d	c
c	d	h	a	c

Table 1: Sample induction problem.

If we observe that every possible value of the relevant variable has the same probability then we clearly cannot make any prediction at all. If all values do *not* have the same probability then we will rationally predict the missing value to be the one which has the highest observed probability. However, there are several ways in which we can work out 'observed probabilities'. First, we can look at the unconditional probability of seeing a particular value v of x_i.

$$P(x_i = v)$$

In the present case this is not productive since both possible values of x_3 have the same unconditional probability. This is just the chance value of 0.5, i.e.,

$$P(x_i = v) = \frac{1}{|V|}$$

where V is the set of all possible values of x_i.

Second, we can look at the probability of seeing a particular value conditional on explicit instantiations of the other values, i.e.,

$$P(x_i = v_a | x_j = v_b ...)$$

where v_a and v_b are possible values and '...' denotes the optional inclusion of other instantiations. This is more rewarding since it turns out that

$$P(x_3 = f | x_4 = a) = 0.8$$

In other words, we see x_3 =f in 4 of the 5 cases where we see x_4 =a.[2]

Third and finally we can look at the probability of seeing a particular value conditional on there being an *implicit* property (i.e., a relationship) among the instantiations of other variables:

$$P(x_i = v | g(X) = v_g)$$

Here X is the entire datum and v_g is the value of an imaginary function g, which evaluates the relationship. This is more rewarding still since it turns out that

$$P(x_3 = h | duplicates(X)) = 1$$

where the duplicates function is a predicate which tests whether there are duplicated values in the datum. In other words, it turns out that we always see x_3 =h when there are duplicates among the other values.

These three formulae represent the *only* ways in which a particular guess might be empirically justified.[3] In fact, there are really only two formulae to consider since we can always regard an unconditional probability as a conditional probability with an empty condition. Thus the task analysis shows that there are really just two sources inductive justification: one based on *explicitly* observed probabilities and the other based on *implicitly* observed probabilities.[4]

[2] Of course this is not the only significant conditional probability.

[3] If this seems counter-intuitive note that the third formula acts as a kind of catch-all since it covers *any* computational, mathematical or functional justification for an inductive guess.

[4] The assumption is made that the concept of 'implicit property' is well defined. Without this assumption, any mapping over the data might be viewed as measuring an 'implicit property' and thus every guess would have a justification.

If we want to make an inductive guess regarding the missing value of x_3, we therefore must exploit some combination of these two sources. In practice, we must either exploit the **statistical effects** underlying explicitly observed probabilities, or the **relational effects** underlying implicitly observed probabilities. In the present example we will probably guess that x_3 =h since the highest probability we have unearthed (so far) is based on the observation that this value occurs in every case where there are duplicates among the remaining values (i.e., a relational effect).

Methods which attempt to discover and exploit such probabilities for inductive purposes — without using any other source of information — are **empirical learning** algorithms. The development of these methods is the concern of several research communities including Machine Learning and Connectionism (see (1, 2, 3)).

3 Complexity

A method that attempts to exploit explicit probabilities confronts a relatively easy task. Only cases that are explicitly observed in the data need to be taken into account. If the dataset if finite, there are a finite number of these. The task thus involves deriving frequency statistics (probabilities) over a *finite* dataset.

A method that attempts to exploit implicit probabilities, on the other hand, confronts a harder task since it has to first identify the appropriate evaluation function for the implicit property (i.e., it has to guess what the relational effect is). If functions with an infinite range are to be considered, then the task is *infinitely* hard, since there are clearly an infinite number of such functions. Even if we restrict attention to functions with a finite range, the task is still *hard* since the number of functions to be considered is exponentially related to the number of observed cases. Consider the simplest case. We have n variables each of which takes m values and we consider only functions with a binary range and minimum arity n. The number of possible functions is then 2^{m^n}.

The general implication is that statistical effects are more easily exploited than relational ones. This reinforces the long-standing belief among Machine Learning researchers that 'learning relationships is hard' (cf. (4)). Statistical learning methods (i.e., methods which exploit statistical effects only) thus execute an easier task than relational methods, i.e., methods which exploit relational effects. (In practice, learning methods often attempt to exploit both types of effect.)

4 Discussion

Since the space of relationships is, in general, infinitely large, relational learners always and necessarily incorporate some form of *bias* [5], i.e., they have a predisposition to focus attention on certain types of relationship. We can regard this as a form of knowledge or a set of hypotheses about the data. The process of relational learning can then be viewed as the task of assimilating observations to a specific set of hypotheses, i.e., as a form of abduction. The process of statistical learning, on the other hand, involves exploitation of observed statistical associations and is thus more easily viewed as a form of induction. The task analysis, then, suggests that, to a first approximation, abduction can be equated with relational learning while induction can be equated with statistical learning.

However, there is an alternative way of utilising the analysis which leads to a more interesting viewpoint. Relational learners are always potentially 'recursive'. The identification of any set of relational effects involves the application of evaluations (functions) to the original data. This effectively creates new values and thus new data. These new data can themselves be processed for statistical and relational effects in a recursive manner. A full-blown relational learner thus operates recursively and necessarily generates *structured hypotheses* about the original data. Statistical learning carried out at higher levels within such structures has the effect of assimilating raw data within (or 'to') the relevant structured hypothesis. This is a process which appears to be strongly abductive in character.

The general implication of the analysis, then, is that abduction is what happens when we move beyond exploitation of statistical associations into the realms of knowledge-based discovery. It suggests that the distinction between induction and abduction is not black-and-white but rather a question of degree. Inductive processes can be more or less abductive and vice versa. The more relational the learning is, and the deeper the generated recursion, the more 'abductive' is the underlying learning (or reasoning) process. If there is little use of relational bias and no recursion then the process is 'minimally abductive'. In extreme cases, we may want to say that learning is exclusively inductive or exclusively abductive. But in general we will need to think in terms of a combination of the two processes.

References

[1] Shavlik, J. and Dietterich, T. (Eds.) (1990). *Readings in Machine Learning.* San Mateo, California: Morgan Kaufmann.

[2] Michalski, R., Carbonell, J. and Mitchell, T. (Eds.) (1983). *Machine Learning: An Artificial Intelligence Approach.* Palo Alto: Tioga.

[3] Michalski, R., Carbonell, J. and Mitchell, T. (Eds.) (1986). *Machine Learning: An Artificial Intelligence Approach: Vol II.* Los Altos: Morgan Kaufmann.

[4] Dietterich, T., London, B., Clarkson, K. and Dromey, G. (1982). Learning and inductive inference. In P. Cohen and E. Feigenbaum (Eds.), *The Handbook of Artificial Intelligence: Vol III.* Los Altos: Kaufmann.

[5] Utgoff, P. (1986). *Machine Learning of Inductive Bias.* Kluwer International Series in Engineering and Computer Science, Vol. 15, Kluwer Academic.