The Problem of Induction and Artificial Intelligence

by Donald Gillies

(This is a paper which was given at a conference in La Coruña, Spain in 2003 on 'Karl R. Popper: Revision of his Legacy'. A Spanish version of it has appeared in the conference proceedings.)

1. Introduction

Although Popper tackled many philosophical problems, the one which perhaps engaged him most was the problem of induction. Chapter 1 of his 1972 book *Objective Knowledge* is devoted to issues about induction, and here on the very first page of the chapter and of the book, he actually claims to have solved the problem of induction. This is what he says (1972, p.1):

"I think that I have solved a major philosophical problem: the problem of induction. … This solution has been extremely fruitful, and it has enabled me to solve a good number of other philosophical problems."

Unfortunately hardly any other philosophers have accepted this claim of Popper's, as Popper himself admits for he goes on to say (1972, p. 1): "However, few philosophers would support the thesis that I have solved the problem of induction." Over thirty years have passed since the publication of this book of Popper's, but it still remains the case that very few philosophers indeed think that Popper solved the problem of induction. Indeed many regard Popper's claim to have done so as quite outrageous. Thus we have a striking assertion by an individual met with scepticism by the majority. Let us now look into the matter more closely. To do so it will be useful to consider first what exactly is the problem of induction, and this I will do in the next section.

2. The Problem of Induction

The phrase: '*the* problem of induction' is perhaps somewhat misleading, since there are several interconnected philosophical problems concerned with induction. Later, in section 5 of this lecture, I will distinguish two types of problem concerned with induction – one type is to do with *justification* by induction, while the other is to do with *discovery* by induction. Popper implicitly accepts the plurality of problems of induction because in chapter 1 of his 1972 book, he formulates no less than fourteen different problems of induction. Despite this plurality, there is perhaps one specific problem concerned with induction which many philosophers would indeed regard as *the* problem of induction. I will now try to formulate this problem using the standard example of black ravens. Suppose we observe several thousand black ravens. Let us call this observational evidence e. Then from e we could, quite reasonably, infer the prediction (d say) that the next raven we observe will be black, or the generalisation (h say) that all ravens are black. Although these seem to be reasonable inferences, it was Hume who showed that inference in this sense cannot be the say as logical deduction. If from premise A conclusion B follows by logical deduction, then if A is true, B must be true. However in the above example of ravens, e could be true, but d and hence h might still be false. It is quite possible that after observing thousands of black ravens, the next raven we observe turns out to be white. After all, thousands of white swans were observed in Europe, but when European explorers went to Australia they observed black swans for the first time. So the inference from e to d, or from e to h, though reasonable, could not be the same as logical deduction. It needed to be given some new name, and so was dubbed an inductive inference. In both everyday life and science, we are constantly inferring predictions and generalisations from past observations. It would thus seem that inductive inference plays a central role in both everyday life and science. However such inference appears somewhat problematic since there are many examples of inductive inferences leading to false conclusions. Thus, for example, the inductive inference based on numerous observations of European swans that all swans are white proved to be false. Granted then that inductive inference is rather questionable, the question arises of how we can justify our use of it. This is the problem of induction, which can therefore be formulated as the problem of justifying inductive inferences.

This formulation of the problem of induction is quite similar to what Popper calls (1972, p. 2): "*the traditional philosophical problem of induction*". Popper's second formulation of this is (1972, p. 2): "What is the justification for inductive inferences?"

3. Popper's Proposed Solution to the Problem of Induction

Having formulated the problem of induction, let us now examine Popper's proposed solution. Popper's basic strategy is to deny the existence of inductive inferences. Proponents of the concept of inductive inference claim that such inferences are essential both to everyday life and to science, but Popper, on the contrary, holds that such inferences are not necessary and are never made. According to Popper, the traditional problem of induction (1972, p. 2): "assumes that there are inductive inferences, and *rules* for drawing inductive inferences, and this … is an assumption which should not be made uncritically, and one which I … regard as mistaken." He puts the point in an even more striking fashion in an earlier work (1963, p. 53):

"Induction, i.e. inference based on many observations, is a myth. It is neither a psychological fact, nor a fact of ordinary life, nor one of scientific procedure."

Now if induction and inductive inferences do not exist, then there is clearly no need to try to justify them. So the problem of induction, that is to say the problem of justifying inductive inferences, disappears. This, in my opinion, is Popper's proposed solution to the problem of induction.

But why does Popper think that induction is a myth? He has a number of arguments for this conclusion, and I will here consider two. First of all Popper holds that we cannot start with observations and infer theories, because observations themselves require a theoretical framework. This is how he puts the argument (1963, pp. 46-7):

"… the belief that we can start with pure observations alone, without anything in the nature of a theory, is absurd; as may be illustrated by the story of the man who dedicated his life to natural science, wrote down everything he could observe, and bequeathed his priceless collection of observations to the Royal Society to be used as inductive evidence. This story should show us that though beetles may profitably be collected, observations may not.

Twenty-five years ago I tried to bring home the same point to a group of physics students in Vienna by beginning a lecture with the following instructions: 'Take pencil and paper; carefully observe, and write down what you have observed! They asked, of course, *what* I wanted them to observe. Clearly the instruction, 'Observe!' is absurd. … Observation is always selective. It needs a chosen object, a definite task, an interest, a point of view, a problem. And its description presupposes a descriptive language, with property words; it presupposes similarity and classification, which in its turn presupposes interests, points of view, and problems. … objects can be classified, and can become similar or dissimilar, *only* in this way – by being related to needs and interests. … a point of view is provided … for the scientist by his theoretical interests, the special problem under investigation, his conjectures and anticipations, and the theories which he accepts as a kind of background: his frame of reference, his 'horizon of expectations'."

There is a great deal of truth in this passage from Popper, but it shows the need for an important reformulation of the character of inductive inferences rather than, as Popper hoped, the need for dispensing altogether with inductive inferences. Let e be a set of observation statements, and h a generalisation. Then we might initially take an inductive inference to be an inference from e to h. However Popper's argument shows that this is not correct, and that inductive inferences have to have the form: 'from K&e infer h' where K stands for 'background knowledge', 'frame of reference', etc. In particular K must contain a set of concepts to be used for the classification of objects. I will return to this 'need for K' later on.

Popper's second, and principal, argument against induction is that we can give an account of scientific and indeed common-sense reasoning which does not involve any of the supposed inductive inferences. This is Popper's famous model of conjectures and refutations. According to this model, science starts not with observations but with conjectures. These conjectures are then tested out against observation. If a conjecture is refuted, it has to be replaced by some new conjecture. If a conjecture has been confirmed by the tests so far, it may be tentatively accepted, but testing continues and no conjecture is ever regarded as definitely established. Now the process of testing and refutation does not require any inductive inferences but only the use of deductive logic. From our conjecture H we deduce logically that some observation statement O must hold. We then check empirically whether O is true or not. If observation shows that O is false, then our conjecture is refuted. This procedure requires inferring O from H which we do by deductive logic. It also requires inferring from H r O and not-O, that not-H. However this inference is again a law of deductive logic, known as *modus tollens*. Popper concludes therefore that only deductive logic is needed, or indeed used, so that inductive inferences are a myth and the problem of justifying such inferences, i.e. the problem of induction disappears.

4. Philosophical Criticisms of Popper's Proposed Solution

Popper's arguments concerning induction seem quite plausible. So why have they not been generally accepted by the philosophy community? In this section I will give some philosophical arguments against Popper's position. To begin with it might be argued that Popper's model of conjectures and refutations, and his falsifiability principle do not give an adequate account of science because of the Duhem-Quine problem. We considered this difficulty in Popper's position in the last lecture, and I do not want to pursue it further here. Let us assume then for the rest of this lecture that all the generalisations we are considering are what we called 'level 1 hypotheses', i.e. that they are falsifiable so that the method of conjectures and refutations applies. Many scientific laws are in fact of this character so that we are in effect confining our study of induction to an important but limited class of laws. If we can give a satisfactory account in this area, it could later be investigated whether it extends to the more abstract laws of level 2.

The question then becomes, whether, assuming we can apply the method of conjectures and refutations, this dispenses altogether with the need for inductive inferences. Let us suppose then that we conjecture a hypothesis (h say). h explains many existing observations and it passes a good number of severe tests which are designed and carried out. So h becomes well corroborated by the evidence. In these circumstances, it would seem sensible to accept h as a basis for practical action. Popper largely agrees with this for he writes (1972, pp. 21-2):

"Which theory should we prefer for practical action, from a rational point of view? … My answer … is … we should *prefer* as basis for action the best-tested theory."

But now if we adopt h as a basis for practical action, this can only be because we think that the guidance which h affords us as to the future is superior to the guidance offered by other theories. In other words we think that a well corroborated theory, i.e. one which has agreed very well with past evidence, will continue to give good results in the future. But this is surely an inductive assumption. We could call it the inductive interpretation of the corroboration function. Curiously enough there is one passage in which Popper comes close to accepting this approach to corroboration. It occurs in the new appendix *ix *Corroboration, the Weight of evidence, and Statistical tests*, where Popper is discussing his function *C(h, e)* which stands for the degree of corroboration of h given e. He writes (1959, p. 418):

"It might well be asked at the end of all this whether I have not inadvertently, changed my creed. For it may seem that there is nothing to prevent us from calling *C(h, e)* 'the inductive probability of *h*, given *e*' or – if this is felt to be misleading, in view of the fact that *C* does not obey the laws of the probability calculus – 'the degree of the rationality of our belief in *h*, given *e*'. A benevolent inductivist critic might even congratulate me on having solved, with my *C* function, the age-old problem of induction *in a positive sense* – on having finally established, with my *C* function, the validity of inductive reasoning."

However this passage is no doubt partly meant ironically, and, in any case, does not represent Popper's final opinion on the question since in chapter 1 of *Objective Knowledge*, he states very clearly that he favours a non-inductivist interpretation of corroboration (1972, p. 18):

"By the degree of corroboration of a theory I mean a concise report evaluating the state (at a certain time *t*) of the critical discussion of a theory, with respect to the way it solves its problems; its degree of testability; the severity of tests it has undergone; and the way it has stood up to these tests. Corroboration (or degree of corroboration) is thus an evaluating *report of past performance*. … Being a report of past performance only … *it says nothing whatever about future performance, or about the 'reliability' of a theory*."

It is nonetheless the case that we do rely on well corroborated theories as the basis of our practical actions. Thus we do assume that a high degree of corroboration is an indication that a theory will perform well in the future as a reliable guide to what will happen. For this reason, I do not regard Popper's 1972 account as a viable interpretation of corroboration, and I think that we have to adopt the inductivist interpretation of corroboration. This means, however, that Popper has not succeeded in eliminating inductive inference altogether and therefore that his solution of the problem of induction is not successful. As a 'benevolent inductivist critic', however, even though I cannot agree that Popper has 'solved … the age-old problem of induction *in a positive sense*', I think we should credit him with achieving an important and progressive problem-shift.

The traditional problem of induction was concerned with justifying inductive inferences similar to inferences about the colour of some or all ravens to be observed in future on the basis of the colour of those observed in the past. As a result of Popper's theorising, a rather different problem of induction has emerged. It is concerned with the inductive inference that an hypothesis h which has been well corroborated by the evidence so far will be a good guide to what will happen in the future and so constitute a reliable basis for practical action. This new problem of induction differs significantly from the old in a a number of respects. Among other things, it involves the concept of corroboration and raises the question of how we should measure degree of corroboration. If a particular measure of corroboration is chosen, the question arises of how, if at all, we can justify inductive inferences based on this measure. These are interesting and relatively novel problems. They are certainly very different from the problems discussed by Hume. However in the second part of this lecture, I do not want to pursue these questions further, but rather to consider how problems about induction are connected with new developments in artificial intelligence (AI). To provide a link with AI, we must make a distinction between two different, though connected, types of problem about induction. This will be done in the next section.

5. Justification by Induction versus Discovery by Induction

In this section I want to explain the distinction between justification by induction and discovery by induction. Let us start by considering justification by induction. Here we assume that we have already formulated and have in front of us an hypothesis h and a set of relevant evidence e. Our problem is whether this evidence e justifies us in accepting h as a reliable guide to the future and as the basis for some practical actions. If we can achieve such a justification, it would be a justification by induction. Now, however, suppose that we have a body of evidence e, but, as yet, no hypothesis h to explain e. Suppose we have some method by which we could obtain a suitable h from e. This would constitute discovery by induction. When Bacon advocated induction, he meant induction in the sense of discovery by induction. Indeed he hoped to formulate a quasi-mechanical procedure by which scientific hypotheses could be obtained from data. Those who like Bacon believe in discovery by induction will in general also believe in justification by induction. In fact suppose there is some method of discovering an hypothesis to explain e by inferring h inductively form e, then advocates of this method are likely to hold that h is justified by the evidence e from which it was obtained. It is, however, perfectly possible to believe in justification by induction without believing in discovery by induction. This was indeed Carnap's position, which he formulated explicitly as follows (1950, pp. 192-3):

" … in one point the present opinions of most philosophers and scientists seem to agree, namely, that the inductive procedure is not, so to speak, a mechanical procedure prescribed by fixed rules. If, for instance, a report of observational results is given, and we want to find a hypothesis which is well confirmed and furnishes a good explanation for the events observed, then there is no set of fixed rules which would lead us automatically to the best hypothesis or even a good one. It is a matter of ingenuity and luck for the scientist to hit upon a suitable hypothesis; … . This point, the impossibility of an automatic inductive procedure, has been especially emphasized, among others, by Karl Popper … who also quotes a statement by Einstein … . The same point has sometimes been formulated by saying that it is not possible to construct an inductive machine. The latter is presumably meant as a mechanical contrivance which, when fed an observational report, would furnish a suitable hypothesis, just as a computing machine when supplied with two factors furnishes their product. I am completely in agreement that an inductive machine of *this* kind is not possible."

Carnap and Popper were noted opponents on questions concerned with induction. Yet the above quotation shows that they agreed on at least one point, namely that the discovery of hypotheses by a mechanical process of induction was not possible. Advances in artificial intelligence have, however, shown that they were both wrong on this point. In fact programs have been written which enable computers, when fed with data, to generate suitable hypotheses for explaining that data. Moreover this new kind of computer induction has resulted in the discovery of important and previously unknown scientific laws. In the next two sections I will give a brief description of some of these remarkable successes of artificial intelligence, and will also examine the implications of these new results for Popper's philosophical ideas about induction.

6. Successes of Machine Learning: the Example of GOLEM

Since the late 1970s a new branch of AI, known as machine learning, has come into existence and has been developed with considerable success. The aim of machine learning is to do precisely what both Carnap and Popper believed in the 1950s to be impossible – namely to induce hypotheses automatically from data. In my 1996 (particularly chapters 2 and 3) I trace the history of machine learning, giving a variety of examples of successful machine learning programs, and examining their implications for the views of Bacon and other philosophers. In the present lecture I will limit the task by describing only one successful machine learning system, and examining its implications for the views of just one philosopher: Popper.

The machine learning system, known as GOLEM, was developed by Stephen Muggleton and his colleagues in the early 1990s. The key paper describing it is Muggleton and Feng 1992. Subsequently Muggleton has developed an improved machine learning system known as PROGOL, but his earlier one suffices for the philosophical points I want to make.

GOLEM was set to work on an investigation of protein structure. A good introduction to this field is Branden and Tooze, 1991, while GOLEM's contribution is described in Muggleton, King, and Sternberg, 1992. 20 different amino acids constitute the building blocks of all proteins. These are joined end to end, so that a protein consists of a sequence of amino acid residues. Now it is relatively easy to discover the sequence of residues in a protein. This is known as the protein's *primary structure*. Unfortunately knowledge of the primary structure of a protein is not sufficient to understand the biological properties of the protein, for these depend crucially on the 3-dimensional shape of the protein. Now proteins as they form fold up into complicated 3-dimensional structures, which are known as the protein's *secondary structure*. The secondary structure of a protein can be determined by X-ray crystallography or NMR techniques; but it is a long and costly business. So far the secondary structures of about 500 proteins have been determined. Progress in biochemistry would become much easier and quicker if it were possible to predict the secondary structure of a protein (which is difficult to determine but biologically crucial), from the primary structure which is easy to determine, but not so significant biologically. As Branden and Tooze put it (1991, p. 3):

> "To understand the biological function of proteins we should therefore like to be able to deduce or predict the three-dimensional structure from the amino acid sequence. This we cannot do. In spite of considerable efforts over the last 25 years, this folding problem is still unsolved and remains one of the most basic intellectual challenges in molecular biology."

Machine learning techniques have advanced to the point at which they can make a contribution to one of leading problems of modern natural science.

Sub-structures of a protein structure are usually of one of two basic types: α*-helixes* and β*–strands*. α–helixes were first described in 1951 by Linus Pauling, who predicted that such a structure would be stable and favourable in proteins. This remarkable prediction almost immediately received strong experimental support from diffraction patterns obtained by Max Perrutz in Cambridge. Proteins can accordingly be classified into three domains according to their secondary structure. These are (i) alpha type domains, in which the proteins have only α-helixes (or at least have very few β-strands), (ii) beta type domains, and (iii) alpha/beta type domains, where α-helixes alternate with β-strands. The simplest prediction problem is obtained by restricting the proteins considered to those of alpha type domain, and then attempting to predict from the primary structure whether a particular residue belongs to an α-helix or not.

GOLEM was applied to this problem in the following fashion. 12 non-homologous proteins of known structure and alpha type domain, involving 1612 residues were selected as the training set. From this training set and background knowledge, GOLEM learned a small set of rules for predicting which residues are part of α-helices. The rules were then tested on 4 independent non-homologous proteins of known structure and alpha type domain, involving 416 residues. The accuracy of the rules was 81% (± 2%). These then were the overall procedure and results.

Let us now look at one of the rules produced by GOLEM. I have selected Rule 12 which is stated in PROLOG format by Muggleton, King, and Sternberg (1992, p. 655). If translated into something closer to normal English, it runs as follows:

GOLEM'S Rule 12 regarding Protein Secondary Structure

There is an α-helix residue in protein A at position B if
(i) the residue at B-2 is not proline,
(ii) the residue at B-1 is neither aromatic nor proline,
(iii) the residue at B is large, not aromatic, and not lysine,
(iv) the residue at B+1 is hydrophobic and not lysine,
(v) the residue at B+2 is neither aromatic nor proline,
(vi) the residue at B+3 is neither aromatic nor proline, and either small or polar, and
(vii) the residue at B+4 is hydrophobic and not lysine.

Some readers may feel rather disappointed with this rule, which is rather long, cumbersome, and specific. It was, however, 95% accurate on the training set, and 81% accurate on the test set. It was not known before being produced by GOLEM, and it makes a contribution to an important current problem in the natural sciences. It seems to me fair, therefore, to credit GOLEM with the discovery of a law of nature.

In the next (and final) section of the lecture I will describe in outline how GOLEM works and this will enable us to assess its implications for Popper's philosophy. Here, however, I would like to make a preliminary point. In order to get GOLEM to work, a good deal of background knowledge concerning proteins and their structure had to be coded into it. In this respect GOLEM was no different from a human research scientist. Research scientists have to learn the existing knowledge in a field before they can advance to the discovery of new results. In particular, as we can see from Rule 12, GOLEM had to be provided with the standard classificatory predicates of the subject such as aromatic, hydrophobic, proline, etc. In addition GOLEM had to be provided with the generally accepted principle that the character of a particular residue (whether it belonged to an α-helix or not) was likely to be determined by the basic properties of the residue itself, and of the four residues on each side of it. Technically this was done by providing GOLEM with a predicate octf describing 9 sequential positions, so that e.g. octf(19, 20, 21, 22, 23, 24, 25, 26, 27). The existence of this predicate enabled GOLEM to search for laws having the general form of Rule 12. It was from background knowledge of this kind together with data about the training set that GOLEM automatically induced results like Rule 12. I will now describe in more detail how this was done.

7. How GOLEM works, and implications for Popper's Philosophy

To give a general idea of how GOLEM works, let us consider how a particular human (Ms A say) might infer inductively that generalisation beloved of philosophers, viz. all European swans are white, and all Australian swans are black. Ms A obtains this result by making cautious inductive inferences from her observations. Thus, on observing one white swan on the river Thames near London, she infers that all swans on the Thames near London are white. On examining swans on other parts of the Thames she infers that all swans on the Thames are white. After looking at swans on other rivers and lakes in England, she infers that all English swans are white. Investigations of swans in France, Italy, Germany, etc. yields the conclusion that all European swans are white. Then she sails to Australia, and ends up with the generalisation that all European swans are white, and all Australian swans are black.

Let us now see how something like this procedure might be carried out by a computer. Alan Robinson in his 1965 introduced a form of logic suitable for theorem proving on the computer. This is known as the *clausal form of logic*. Plotkin at Edinburgh (see his 1970, and 1971a & b) applied this form of logic to machine learning. He made use of the concept of *subsumption*, which had also been introduced by Robinson (1965, 7.2, pp. 38-9), to obtain a partial ordering of the clauses in terms of generality. This ordering enabled Plotkin to define what he called the *relative least general generalisation* (r.l.g.g.) of a set of data points.

More precisely, clause C is said to be the least general generalisation of e_1 and e_2 relative to K whenever C is the least general clause in the partial ordering of the clauses for which K&C entails e_1&e_2, where C is used only once in the derivation of both e_1 and e_2. This corresponds to the idea of our human (Ms A) making a cautious generalisation from two pieces of data. Note however that the r.l.g.g. is relative to background knowledge K. GOLEM, and indeed all other existing machine learning programs, do not make their inductive inferences just from data, but from a combination of data and background knowledge.

Having introduced the notion of relative least general generalisation (r.l.g.g.), Plotkin encountered a problem. His r.l.g.g.'s could in the worst case become infinitely long, and, in general tended to grow exponentially with the number of examples involved. For these reasons the approach was abandoned until it was taken up again by Muggleton and Feng. They managed to introduce restrictions which caused the resulting r.l.g.g.'s to be not just finite, but of reasonable length. Using such r.l.g.g.'s it was possible to construct a machine learning program (GOLEM) operating according to the following iterative procedure.

GOLEM is provided with a set of positive examples, and a set of negative examples. It begins by taking a random sample of pairs of positive examples. It constructs the r.l.g.g. of each such pair. GOLEM takes each such r.l.g.g. and computes the number of examples which it could be used to predict. Clearly a given r.l.g.g. might predict some examples which are false. GOLEM therefore chooses the r.l.g.g. which predicts the most true examples while predicting less than a predefined threshold of false examples. Having found the pair with the best r.l.g.g. (S say), GOLEM then takes a further random sample of the as yet unpredicted positive examples, and forms the r.l.g.g of S and each of the members of this new random sample. These new r.l.g.g's are evaluated as before, and the process continues until no improvement in prediction is produced.

There is no doubt that the successes of GOLEM and other machine learning programs undermine Popper's project for eliminating induction from science. I have already quoted Popper's emphatic statement (1963, p. 53):

"Induction, i.e. inference based on many observations, is a myth. It is neither a psychological fact, nor a fact of ordinary life, nor one of scientific procedure."

This view can no longer be maintained in the light of programs such as GOLEM which do make inductive inferences based on many observations and have become a part of scientific procedure.

Despite this blow to some of Popper's more extreme claims, it by no means follows that the new results of AI are all and unequivocally against Popper. On the contrary some of Popper's ideas do find support from the advances in machine learning. More specifically there seem to me to be three points which can be made in favour of Popper, and I will conclude this section and the lecture by considering these three points in turn.

The first point concerns Popper's emphasis on the need for background knowledge, and here GOLEM and other machine learning programs, far from contradicting Popper, bear him out completely. Popper argues that all observation presupposes a theoretical background, and says, more specifically, that (1963, p. 46): "Observation … presupposes a descriptive language, with property words; it presupposes … classification …" As we have seen, in order for GOLEM to operate in the domain of protein folding, it was necessary to code in the standard predicates of the field such as α-helix, aromatic, lysine, proline, etc. It was moreover a standard piece of background knowledge in the field that the character of a particular residue (whether it belonged to an α-helix or not) would generally be determined by the basic properties of the residue itself, and of the four residues on each side of it. This item of background knowledge was also, as we saw, coded into GOLEM. Not just GOLEM, but all the successful machine learning programs make use of background knowledge K, which is coded into the program and plays an essential role in guiding the computer's search for hypotheses. Indeed computer inductive inferences really take the form: from K&e infer h, rather than the form : from e infer h.

A second point in favour of Popper is that GOLEM and many other machine learning programs make use of Popper's model of conjectures and refutations. As we saw, GOLEM forms an hypothesis from part of the data, and this hypothesis is then tested out against the remainder of the data. Any hypothesis which predicts more than a predefined threshold of false examples is rejected, so that, in effect, a principle of falsifiability is applied. This procedure for eliminating hypotheses which have been falsified by the data agrees very well with Popper's philosophy. There is however the difference that, in Popper's original model, the conjectures or hypotheses were invented using human ingenuity and imagination. In the machine learning case, however, the hypotheses are obtained mechanically from data and background knowledge using an inductive rule of inference. In GOLEM this inductive rule of inference was that of generating the relative least general generalisation of two pieces of evidence.

My third and final point in favour of Popper is an historical one. I have not been able to find a convincing example of mechanical induction in science prior to the appearance of the first successful machine learning programs in the late 1970s and early 1980s. I do not want to maintain this position dogmatically since it is difficult to survey the whole history of science over many hundreds of years, and some convincing example of mechanical induction may come to light. However, even if one or two such examples do exist, they are undoubtedly rare and exceptional. So Popper, when he wrote in 1963 that induction was a myth, was largely correct in maintaining that mechanical induction had not been used as a method of discovery in science up to that time. It would in principle have been possible at an earlier date to have devised methods for generating hypotheses mechanically and for systematically testing out these hypotheses against data. But before the development of the electronic computer such procedures would hardly have been practical, and would certainly have been less effective than the more intuitive ways of thinking which are undoubtedly more natural to human beings. So the practical realisation of mechanical induction is an innovation of the computer era.

Now that computers and artificial intelligence are beginning to play a part in scientific procedure, however, it is predictable that they will transform the way in which science is done, and lead to a whole host of remarkable and surprising discoveries. An analogy will help to illustrate what may lie in store. Before the year 1609 observations in science had been made exclusively with unaided human sense organs. In that year, Galileo for the first time used an instrument to make scientific observations, namely the telescope to survey the heavens. His discoveries were truly remarkable. He was able to observe mountains on the moon, and could see at least ten times as many stars as had previously been known. He found that the Milky Way consisted of innumerable stars, and discovered that the planet Jupiter had moons circling round it.

All this showed the enormous advantage of improving naked-eye observation by the use of an instrument. Soon other instruments, such as the microscope were devised, and nowadays hardly any science is carried out which does not use instruments. The development of instruments to assist the human sense organs has changed the way in which science is done, and brought about a vast extension of human knowledge. Is it not reasonable to suppose, therefore, that the development of instruments to help the human mind – i.e. computers furnished with AI programs – will have a similar effect on science?

It is important to realise that scientific method is not something fixed and ahistorical. Rather as Bacon says (1620, Book I, CXXX, p. 301): " ... the art of discovery may advance as discoveries advance." Popper was not wrong, when, writing before the late 1970's, he denied that mechanical induction formed part of scientific procedure. What he said was true of science as it had been practised up to that time. However, just as earlier the use of instruments to assist observation altered the way in which science was done, so the current development of computers and AI is also destined to change science, and in such a way that mechanical induction becomes a standard part of scientific procedure.

References

Bacon, F. (1620) *Novum Organum*, English translation in R. L. Ellis and J. Spedding (eds.), *The Philosophical Works of Francis Bacon*, routledge, 1905, pp. 212-387.

Branden, C., and Tooze, J. (1991) *Introduction to Protein Structure*, Garland Publishing.

Carnap, R. (1950) *Logical Foundations of Probability*, 2nd Edition, University of Chicago Press, 1963.

Gillies, D.A. (1996) *Artificial Intelligence and Scientific Method*. Oxford University Press.

Muggleton, S., and Feng. C. (1992) 'Efficient Induction of Logic Programs' in S. Muggleton (ed.), *Inductive Logic Programming*, Academic Press, Ch. 13, pp. 281-98.

Muggleton, S., King, R.D., and Sternberg, M.J.E. (1992) 'Protein Secondary Structure Prediction using Logic-Based Machine Learning', *Protein Engineering*, **5**/**7**, pp. 647-57.

Plotkin, G.D. (1970) 'A Note on Inductive Generalisation', in B. Meltzer and D. Michie (eds.). *Machine Intelligence*, **5**, Elsevier North-Holland, pp. 153-63.

Plotkin, G.D. (1971a) 'Automatic Methods of Inductive Inference', Ph.D. thesis, Edinburgh University.

Plotkin, G.D. (1971b) 'A Further Note on Inductive Generalisation' in B. Meltzer and D. Michie (eds.), *Machine Intelligence*, **6**, Elsevier North-Holland, pp. 101-24.

Popper, K.R. (1959) *New Appendices to the Logic of Scientific Discovery*. 6th revised impression of the 1959 English translation, Hutchinson, 1972, pp. 309-464.

Popper, K.R. (1963) *Conjectures and Refutations*. *The Growth of Scientific Knowledge*. Routledge & Kegan Paul.

Popper, K.R. (1972) *Objective Knowledge. An Evolutionary Approach*. Oxford.

Robinson, J.A. (1965) 'A Machine-Oriented Logic based on the Resolution Principle', *Journal of the Association for Computing Machinery*, **12**, pp. 23-41.