L&PS – Logic & Philosophy of Science
Vol. III, No. 1, 2005

Machine Learning from Examples: A Non-Inductivist Analysis

E. Datteri
Dip. di Filosofia, Università di Pisa
edoardo@mail-arts.sssup.it
H. Hosni
School of Mathematics, The University of Manchester
hykel@maths.man.ac.uk
G. Tamburrini
Dip. di Filosofia, Università di Pisa
gugt@fls.unipi.it

ABSTRACT. It has been suggested that AI investigations of mechanical learning undermine sweeping anti-inductivist views in the theory of knowledge and the philosophy of science. In particular, it is claimed that some mechanical learning systems perform epistemically justified inductive generalization and prediction. Contrary to this view, it is argued that no trace of such epistemic justification is to be found within a rather representative class of learning agents drawn from machine learning and robotics. Moreover, an alternative deductive account of these learning procedures is outlined. Finally, the opportunity of developing an induction-free logical analysis of non-monotonic reasoning in autonomous learning agents – capable of advancing and revising learning or background hypotheses – is emphasized by a broad reflection on some families of non-monotonic, albeit deductive, consequence relations.[1]

[1]We are grateful to Jeff Paris and David Makinson for useful comments and suggestions.

KEYWORDS. Induction, machine learning, behavior-based robotics, non-monotonic consequence relations, artificial intelligence.

1. Introduction

The idea that work in AI enables one to adjudicate epistemic issues about induction has attracted ever growing attention over the last two decades. Michalski made an initial suggestion to this effect in the early 1980s, at a time when machine learning was just beginning to be recognized as an independent field of investigation:

> [...] [T]here was even doubt whether it would ever be possible to formalize inductive inference and perform it on a machine [...]. The above pessimistic prospects are now being revised. With the development of modern computers and subsequent advances in artificial intelligence research, it is now possible to provide a machine with a significant amount of background information. Also the problem of automating inductive inference can be simplified by concentrating on the subject of hypothesis generation, while ascribing to humans the question of how to adequately validate them. ([30], pp. 87-88.)

More recently, Howson suggested that developments in machine learning provide significant material for philosophical and logical reflections on induction:

> [...] the problem of induction, considered as the problem of characterizing soundness for inductive inferences, has recently become hot (so to speak). People are now for the first time allocating [...] substantial intellectual and material resources to the design of intelligent machinery, and in particular machinery that will learn from data [...]. What is clear is that some logical ba-

We would also like to thank Alberto Mura and David Miller for fruitful discussions on this topic over the past few years.

> sis for learning will certainly have to be built into any successful system. ([18], p.3).[2]

And Gillies [16] claimed that recent advances in concept and rule learning impel a real turn in epistemological discussions of induction, for these results show that Popper's radical brand of anti-inductivism [36] is untenable.

We argue here that these various expectations and claims are not supported by current work on mechanical learning: no matter how significant for understanding or attaining mechanical intelligence, learning machines fail to bolster the inductivist case in epistemology and the philosophy of science.

To state more precisely this claim, and to set the stage for an alternative logical analysis of mechanical learning, to be sketched in the final part of this paper, let us preliminarily recall salient features of inductive inference and related epistemic issues.

> In its most basic sense, an induction is a non-deductive argument or inference from a sample to a conclusion which *projects* the sample in some way. One way in which a reasoner or arguer projects from a sample can be schematized as follows:
>
> 1. In sample S, the Fs are Gs,
> 2. So, Fs are Gs.
>
> A second way is this:
>
> 1. In the sample, Fs are Gs,
> 2. So, the next-observed F will be a G.

[2]Elsewhere, the idea that computational learning systems can aptly perform inductive inference is simply taken for granted:

> The Theory of Machine Inductive Inference (or "Computational Learning Theory", etc.) attempts to clarify the process by which a child or adult discovers systematic generalizations about her environment. ([19], p. 28.)

In the same work, it is assumed that human reasoners possess an inductive competence:

> The focus of our book is the inductive competence of scientists whose behavior can be simulated by computer. ([19], p. 61).

For a criticism of this standpoint on recursion-theoretic learning, see [41], section 3.

> The first example illustrates *generalization* by inductive inference, and the second illustrates *prediction* by the same means. ([42], pp. 106-107)

Let us note two central features of the above account of inductive inference. First, a rigid opposition between induction and deduction is evoked (it is clearly assumed that the distinction is indeed a partition of the domain of inference). Second, the term "projection" (of knowledge/information about the sample) is used (in a fairly metaphorical way) to distinguish induction from deduction, and projection is further qualified as a combination of "generalization" and "prediction". The traditional epistemic issue about these generalizations and projections is whether and what sorts of constraints make it reasonable to believe in their outcomes. This is indeed "the problem of characterizing soundness for inductive inference" mentioned by Howson in the passage quoted above.

We argue that current work in machine learning does not afford a positive solution to the epistemic problem of induction. Moreover, a variety of allegedly inductive procedures in learning machines, - i.e., learning procedures giving rise to the above mentioned "projective" behaviours from observed samples - can be accounted for in terms of default-based deductive reasoning, without appealing to as yet unjustified principles of mechanical induction.

This conclusion is reached by reference to representative classes of learning systems drawn from robotics and machine learning. Our first example, in section 2., is drawn from behaviour-based robotics, which is chiefly concerned with the design and implementation of autonomous systems that survive in realistic worlds. Some of these systems learn from experience, insofar as they acquire new sensorimotor capabilities and generalize environmental properties from observation. Analysis of a representative behaviour-based architecture reveals that even these rudimentary learning mechanisms embody crucial assumptions about their environment. Accordingly, the epistemic problem of induction is reformulated as the problem of whether these background assumptions make it reasonable to believe the associated projections about behavioural rules and environmental properties.

We address the epistemic problem of induction in section 3., by reference to the symbolically richer, ID3-style machine learning algorithms. These algorithms provide the main basis for Gillies's claim that radical anti-inductivism is untenable in the theory of knowledge and the philosophy of science (see Gillies [16]). However, a sweeping problem in learning from examples jeopardizes the idea that epistemically justified inductive processes are at work there. This is the overfitting of training data, which reminds one that a good approximation to the target concept or rule on training data is not, in itself, diagnostic of a good approximation over the whole instance space of that concept or rule. And the successful performances of these learning systems are of no avail either: a familiar regress in epistemological discussions of induction arises as soon as one appeals to past performances of these systems in order to conclude that good showings are to be expected in their future outings as well. Accordingly, we advance a different view of ID3-like projective behaviours, which essentially rely on deductive reasoning from a variety of heuristic hypotheses about concept spaces and current target concept.

This interpretation of ID3 learning brings to the fore the central role of deductive trial and error-elimination processes in autonomous learning mechanisms, which interleave default-based introduction of projective hypotheses about observed samples, retraction of falsified hypotheses, and the selection of new default, background hypotheses for more effective learning. In section 4., the opportunity of developing an induction-free logical analysis of this multifaceted reasoning process is enhanced by reflecting on some families of non-monotonic, albeit deductive, consequence relations. These enable one to frame mechanical projections of learning hypotheses from observed samples into more comprehensive inference processes enabling agents to retract falsified learning hypotheses and to modify underlying knowledge bases in suitable ways.

2. Empirical assumptions of behaviour-based learning

Behaviour-based robotics is a relatively recent area of robotics research, which grew out of a widespread dissatisfaction with traditional robotic architectures for perception-action coordination in realistic environments. The question whether behaviour-based systems actually help overcome this dissatisfaction is still open. The point here is that some such systems, provided with suitable learning mechanisms, exhibit interesting aspects of the projective behaviour which is customarily associated to inductive agents. Of particular interest to us is their capability to "acquire new knowledge from experience".

Let us begin from a brief description of these systems, and then move on to analyze their allegedly inductive capabilities.

It is useful to describe the functional structure of behaviour-based agents as formed by a set of n layers, called *behaviours*, each layer computing a function from sensor data to motor actions [2]. The response r_i of behaviour b_i at time t is

$$r_i = \gamma \times b_i(s_t) \tag{1}$$

where s_t is the state perceived by the sensor associated with behaviour b_i at time t, γ is a real-valued gain and $0 \leq i \leq n$.

Additionally, to each behaviour is associated a set of preconditions, that is, of activation conditions. The architecture is parallel and asynchronous [6].

If m behaviours control the same actuator, a coordination function C, taking as input the outcomes of those m behaviours, outputs a single command for the actuator.

A typical coordination function (termed *cooperative*) blends together (typically sums) multiple responses, previously weighted by means of a set of m real valued gains.[3]

The function C, together with the various gains and mappings constituting the behaviours, determines how the system acts in response to perceived environmental conditions. Let this overall perception-action mapping be the agent's control function f [32]. Distinct sets of gains determine distinct reaction policies to environmental stimuli, i.e. distinct control functions.

Appropriate tuning of perception-action control algorithms is required to let these agents act properly and survive. When the appropriate perception-action control algorithm, i.e. the appropriate sets of gains (for given agents, environments, and tasks) is hard to choose *a priori*, then a viable option is to let the robot learn these values by itself:

[3] Another typical coordination function, termed *competitive*, selects one of the responses on the basis of some priority condition associated with each behaviour.

> One of the problems with behavior-based robots is that the component modules have to be laboriously programmed by a human designer. [...] If new behaviors could be learned, it would also free the designer from needing a deep understanding of the interactions between a particular robot and its application environment [25].

Learning, in behaviour-based agents, is ordinarily achieved by changing gain values over time, either those associated with each behaviour or those associated with coordination function C. Moreover, one can add new behaviours to the system. These changes result into a new system control function.

The most direct way to devise a learning algorithm for a behaviour-based system is to let the system learn the control function f from training examples, that is, from examples of how f acts on inputs $x_j, j = 1, \ldots, k$. In many behaviour-based systems, however, good training sets are not easily available. For example, a Mars Pathfinder should be capable of learning good perception-action policies on duty, since fully representative training sets containing pairs of "perception – (right) action" on Mars are not available.

Reinforcement learning is particularly suited for situations of this kind [20]. Started with default gains, the agent builds up an appropriate control function by stepwise parameter modifications, which depend on positive or negative action rewards, that is, on how the outcome of these actions are evaluated with respect to the agent's ultimate goal.

In a sense, reinforcement learning relieves system designers from the task of specifying the perception-action control function in every detail. At first sight, the system starts with no knowledge of the world, it is set free in its environment, and comes across positive and negative rewards. It learns "directly" from the environment, acquiring knowledge and skills in a completely autonomous and unbiased way, thereby exhibiting the main distinguishing features of a genuine "inductive learning mechanism"[4]. This *prima facie*

[4]See [20]):

> Reinforcement learning studies the problem of inducing by trial and error a policy from states to actions that maximizes a fixed performance measure (or reward).

plausible account, however, fails to underscore the crucial role played by *default assumptions* in this reinforcement learning mechanism. An understanding of this role, we submit, paves the way to a more satisfactory account of this learning mechanism in terms of a particular kind of deductive process.

To begin with, let us see how the learning mechanism is biased by *a priori* assumptions on how the result of the learning process should look like. As an example, consider behaviour-based system described in [7] (A from now on). It consists of three behaviours, each behaviour being associated with one or more gains. These gains play a role in behavioural assemblage which is achieved by summing the outcomes of each behaviour (*cooperative* coordination function). Here we list, for each behaviour, the parameters that are adjusted by the learning algorithm (omitting those additional parameters that never change):

- **move-to-goal:**
  - G_m (a gain)
- **avoid-obstacle:**
  - O_s (some sort of obstacle "influence sphere")
  - O_g (a gain)
- **noise:**
  - N_p (some sort of "noise persistence")
  - N_m (a gain)

The designers of *A* emphasized the need of endowing this system with learning capabilities:

> [...] the robot may be required to navigate in unfamiliar environments where the appropriate values for the behaviors cannot be known in advance. Our research has concentrated on extending reactive controllers to include the ability to learn these values[].

The learning strategy devised for this system is meant to allow A to learn the appropriate gain values (consequently, the function f) "assuming no knowledge of the world". It combines the principles of case-based reasoning [20] to a strategy named "learning momentum", that adjusts the gains at every perception-action step. Let us see in more details how this strategy works, by analyzing the learning algorithm embedded in its ADJUSTER module (L, for short).

Default behavioural parameters are set at the beginning of the navigation. Then, L calculates reinforcement values by evaluating the amount and speed of progress that A has made towards the achievement of its goal. In particular, it calculates the following values:

$\bar{m}$: the mean step size of the H_{steps} past steps ("step size" refers to the distance travelled in one perception-action cycle);

$\bar{p}$: the mean progress toward the goal, calculated as the ratio of distance travelled to the change in distance to the goal;

$\bar{o}$: the mean number of sensed obstacles.

At every step, L adjusts gains on the basis of the values of $\bar{m}, \bar{p}$ and $\bar{o}$. Four basic cases are considered, requiring distinct adjustment parameters:

1. **No-Movement**, in case step size is below threshold:
   $\bar{m} < T_m$

2. **Movement-Toward-Goal**, in case step size and progress to goal are above threshold:
   $\bar{m} > T_m$;
   $\bar{p} > T_p$

3. **No-Progress-With-Obstacles**, in case the robot is moving away from goal and there are several detectable obstacles:
   $\bar{m} > T_m$;
   $\bar{p} < T_p$

4. **No-Progress-No-Obstacles**, in case the robot is moving away from goal and there are no detectable obstacles:

$\bar{m} > T_m$;
$\bar{p} < T_p$;
$\bar{o} < T_o$

where T_m, T_p, T_o are fixed thresholds.

A set of values that the system uses to update gains N_p, N_m, G_m, O_g, O_s is associated to each of these four cases.

For example, in case of **No-Movement**, one can increase the gains associated with the **noise** behaviour (N_p and N_m) and decrease the other gains:

- N_p: +1.0
- N_m: +0.1
- G_m: -0.1
- O_g: -0.1
- O_s: -0.5

In this way, when A's behaviour is stagnant, the **noise** behaviour's gains are progressively increased. Accordingly, **noise** behaviour will eventually dominate every other system behaviour, possibly enabling A to escape from stalemate.

In case of **No-Progress-With-Obstacles**, two adjusting strategies suggest themselves. In so-called *ballooning strategy*, the gains of **noise** and of **avoid-obstacle** behaviours are increased:

- N_p: +1.0
- N_m: +0.05
- G_m: -0.05
- O_g: +0.01
- O_s: +0.5

The **avoid-obstacle** behaviour will drive the agent away from obstacles, inhibiting the effect of **move-to-goal**. Moreover, if the robot approaches a cluster of obstacles including no pathway through them, the high gain set for the **noise** behaviour will make the agent wander randomly, thus increasing its chances to get around the obstacles.

But consider now a very cluttered environment, in which obstacles are homogeneously distributed rather than grouped together. The distance between obstacles is small but sufficient for the agent to pass without colliding. Increasing O_s and O_g, as in the *ballooning* strategy, could have the effect of preventing the agent from getting close enough to find a path, as all free space should be inside the obstacles' sphere of influence. The *ballooning strategy* is clearly unsuitable for this environment. A better option is to decrease O_s and O_g in favour of the **move-to-goal** behaviour. As a result, A would no longer escape vigorously from obstacles, as determined by the *ballooning strategy*, and would rather look for a cleft among obstacles while still avoiding them. Let us call *squeezing* strategy this alternative modification of parameters.

The *ballooning strategy* and the *squeezing* strategy have enabled us to illustrate different learning biases, and to reveal the myth of *tabula rasa* learning in quite simple navigation tasks. Let us now assume that the two types of environments, the one full of sparse clusters of obstacles or canyons, the other one full of sparse, uniformly distributed obstacles, can be defined more precisely - say, taking into account the mean dimensions of obstacle clusters, the distance between their centers of mass, or other well-defined parameters. Let us call these types of environment E_1 and E_2 respectively. Furthermore, let L_1 and L_2 be the learning algorithms provided with the sets of adjustment values corresponding to the *ballooning* strategy, and the *squeezing* strategy, respectively. If e is the environment in which A is operating, conceptual analyses of L_1 and L_2 and related experimental results described in [39] suggest the following thesis

- due to its adjustment values, L_1 will output a control function adequate for A to survive in environments of type E_1: actions generated after learning in E_1 will be positively rewarded;
- due to its adjustment values, L_2 will output a control function adequate for A to survive in environments of type E_2: actions generated after

learning in E_2 will be positively rewarded.

These assumptions state a correlation between actions and positive rewards, as long as the right algorithm was chosen, based on the knowledge of e's properties:

> If e is of type E_1, then choose L_1, and if e is of type E_1, then choose L_2.

This discussion shows that no induction from *tabula rasa* is involved in the generation of a control function suitable for e. Indeed, the only way to make A survive in e is not just to let it act and learn. The fact that A will produce, at a certain point in time, a *positively rewarded action* (i.e., allowing its survival in e) crucially depends on the adequacy of the assumptions expressed above. In other words, the success of a certain action in a given environment e is conditional on the empirical adequacy of the default assumptions about the nature of the environment on which action execution ultimately depends.

Even though reinforcement learning algorithms relieve designers of the control function from the task of specifying detailed information about the environment, knowledge about the properties of the environment – and consequently about the properties of the appropriate target function(s) that the system should learn– is still necessary to devise an appropriate learning algorithm:

> The problem of programming however has not gone away. It has only shifted from specifying sets of reactive rules to specifying reinforcement policies that lead the learning system to discover sets of rules that accomplish the desired task. [8]

Addressing the epistemic problem of induction in these cases is tantamount to finding appropriate justification for the empirical adequacy (the plausibility, reliability, and so on) of newly learnt behavioural rules in the light of the default assumptions underlying chosen reinforcement policies. In the next section, we argue that this epistemic problem is not solved by representative inductive learning algorithms developed in machine learning.

3. Machine learning and the epistemic problem of induction

In AI agents that learn concepts from examples, hypothesis spaces are often construed as sets of Boolean-valued functions over concept instances. In order to learn target concept h, the algorithm examines a training set, that is a finite subset of the whole instance space X formed by positive or negative instances of h. [5]

The assumption that the projective behaviour of computational systems that learn concepts from examples is epistemically justified can be schematically stated as follows:

> (IC) Any hypothesis found to approximate the target function well over a sufficiently large set of training examples will also approximate the target function well over unobserved examples.[6]

A thorough examination of this assumption requires an extensive survey of learning systems that goes well beyond the scope of this paper. Hence, for present purposes, we will focus on versions of (IC) concerning the decision tree algorithm ID3. Decision tree learning is a widely used method in concept learning, and Quinlan's ID3 reflects crucial features of this method [37, 38]. Moreover, ID3 has been widely appealed to in order to undercut anti-inductivist claims in the theory of knowledge and the philosophy of science.

Thus, let us focus on the following:

> (IC-ID3) Any hypothesis constructed by ID3 which fits the target

[5]This theoretical framework is adequate to represent also some robotic behaviour-based learning systems [32]. In typical AI machine learning systems the "learning" phase is preliminary to the "testing" phase: the former runs through a set of pre-computed perception-action pairs whilst the latter consists in the application of the learnt function for action. In behaviour-based learning systems the learning and testing phases are usually mixed: every positively or negatively rewarded instance is used both for learning (because it causes a positive or negative reward) and for acting (because it leads to an action). Nevertheless, the two phases are conceptually distinct and easily identified in their algorithmic structure. The argument developed at the end of this section is applicable , with possible minor modifications, to many behaviour-based learning systems too.

[6]Cp. ([31], p. 23).

> function over a sufficiently large set of training examples will approximate the target function well over unobserved examples.

To begin with, let us recall some distinctive features of (the ID3) decision tree learning. Decision trees provide classifications of concept instances in a training set, formed by conjunctions of attribute/value pairs. Each non-terminal node in the tree stands for a test on some attribute, and each branch descending from that node stands for one of the possible values assumed by that attribute. Each path in the tree represents a classified instance. The terminal node of each path in the tree is labeled with the yes/no classification. The learnt concept description can be read off from the paths which terminate into a "yes" leaf. Such description can be expressed as a disjunction of conjunctions of attribute/value pairs.[7] An instance in the training set is classified by starting at the root of the tree, testing the attribute associated to this node, selecting the descending branch associated to the value assumed by this attribute in the instance under examination, repeating the test on the successor node along this branch, and so on until a leaf is reached. Each concept instance in the training set is associated to a path in a tree, which is labeled "yes" or "no" at the terminal node. ID3 places closer to the tree root attributes which better classify positive and negative examples in the training set. This is done by associating to each attribute P mentioned in the training set a measure of how well P alone separates the training examples according to their being positive or negative instances of the target concept. Let us call this preference in tree construction the ID3 "*informational* bias".

There is another bias characterizing the ID3 construction strategy. ID3 stops expanding a decision tree as soon as an hypothesis accounting for training data is found. In other words, simpler hypotheses (shorter decision trees) are singled out from the set of hypotheses that are consistent with training data, and more complicated ones (longer decision trees) are discarded. On account of this *simplicity* bias, longer decision trees that are compatible with the training set are not even generated, and thus no conflict resolution strategy is needed to choose between competing hypotheses.

[7]Concept descriptions that make essential use of relational predicates (such as "ancestor") cannot be learnt within this framework. Hence ID3 decision trees amount to nothing but propositional binary decision diagrams.

We are now in the position to state more precisely inductive claim (IC-ID3), by reference to the main background hypotheses used by ID-3 to reduce its hypothesis space:

> (IC-ID3: second version): Any hypothesis constructed by ID3 on the basis of its informational and simplicity biases which fits the target function over a sufficiently large set of training examples will also approximate the target function well over unobserved examples.

Scepticism about this claim is fostered by the *overfitting problem*. An hypothesis $h \in H$ is said to overfit the training set if another hypothesis $h' \in H$ performs better than h on the instance space X, even though h' does not fit the training set better than h. Overfitting in ID3 trees commonly occurs when the training set contains an attribute P unrelated to the target concept, which happens to separate well the training instances. In view of this "informational gain" P is placed close to the tree root.

> Overfitting is a significant practical difficulty for decision tree learning and many other learning methods. For example, in one experimental study of ID3 involving five different learning tasks with noisy, nondeterministic data,[...] overfitting was found to decrease the accuracy of learnt decision trees by 10-25% on most problems. ([31], p. 68)

Unprincipled expansions of the original training set may not prevent the generation of overfitting trees, for a larger training set may bring about additional noise and coincidental regularities. Accordingly, claim (IC-ID3) is to be further qualified: the "sufficiently large set of training examples" mentioned there must be "sufficiently representative of the target concept" as well. This means that conjectures about the representativeness of concept instance collections play a central role in successful ID3 learning.

Consider, in this connection, the post-pruning of overfitting decision trees ([31], pp. 67-72). In post-pruning, one constructs a "validation set", which differs from both training and test sets. The validation set can be used to remove a subtree of the learnt decision tree: this is actually done if the pruned tree performs at least as well as the original tree on the validation set. Expectations of a good performance of the pruned tree on as yet unobserved instances rely on the assumption that the validation set is more representative of the target concept than the training set. Thus, the sceptical challenge directed at (IC-ID3) can be iterated after post-pruning, just by noting the conjectural character of this assumption.

In order to counter this sceptical challenge to (IC-ID3), one should look more closely at the criteria used for judging the representativeness of training and validation examples. But additional problems arise here. These criteria may vary over concepts, and are not easily stated in explicit form. In expert systems, for example, the introspective limitations of human experts is a major bottleneck in system development. The process of extracting rules from human experts turns out to be an extremely time consuming and often unrewarding task. These subjects can usually pick out significant examples of rules or concepts, but are often unable to state precisely the criteria underlying these judgments. Accordingly, automatic learning from examples is more likely to be adopted when criteria for selecting significant concept or rule instances are not easily supplied by human experts; and yet an examination of these criteria is just what is needed to support inductive claim (IC-ID3) by appeal to the representativeness of training examples.

Confronted with these various difficulties, that the sceptic consistently interprets as symptoms that inductive claim (IC-ID3) cannot be convincingly argued for, let us try and see the extent to which ID3 fits into a deductive framework.

We have already formed a vague picture of ID3 as a component of a trial and error-elimination cycle: on the basis of assumptions guiding both training set construction and selection of some concept c, ID3 makes predictions about the classification of concept instances that are not included in the training set.

Let us now provide a deductive account of ID3 predictive behaviour, drawing on the above distinction between the preferences or biases embedded in ID3 proper (which determine both the language for expressing concepts and the construction of decision trees) on the one hand, and the presuppositions that are used to select training sets on the other hand. If the presuppositions of the first kind (ID3 biases) are suitably stated in declarative form, a concept learning algorithm such as ID3 can be redescribed as a theorem prover. This is brought out by the following definition of the inductive bias of a concept learning algorithm (see [31], p. 43).

> Definition: Consider a concept learning algorithm L for the set of instances X. Let c be an arbitrary concept defined over X, and let $D_c = \{\langle x, c(x)\rangle\}$ be an arbitrary set of training examples of c. Let $L(x_i, D_c)$ denote the classification assigned to the instance x_i by L after training on the data D_c. The inductive bias of L is any minimal set of assertions B such that, for any target concept c and corresponding training examples Dc,
>
> $(\forall x_i \in X)[L(x_i, D_c)$ is logically derivable from $(B \wedge D_c \wedge x_i)]$.

One is provisionally entitled to preserve B and D_c as long as the classifications coming in through L or its equivalent deductive system are satisfactory. Suppose, however, that for given i's $L(x_i, D_c)$ is an incorrect prediction. Then, this consequence of B and D_c is to be retracted, and either B or D_c are to be appropriately *revised* in order to obtain a correct classification in those cases. This trial and error-correction behaviour is essentially non-monotonic. ID3-like systems perform only the "trial" part of this process, but a fully autonomous learning machine should be capable of carrying out the "error-correction" part as well. As a consequence, these machines should be capable of performing non-monotonic reasoning, insofar as autonomous learning involves, in addition to advancing learning hypotheses, the possibility of retracting empirically inadequate hypotheses and revising one's knowledge base accordingly.

One may wonder whether this requirement for non-monotonic reasoning takes autonomous learning machines outside the realm of deductive reasoners. If this were the case, the behaviour of fully autonomous learning machines, unlike ID3-like systems, could not be accounted for in terms of deductive procedures only. However, as we shall presently see, research efforts converging in the theories of *non-monotonic reasoning* and *belief revision*, which provide the more comprehensive and logically well-understood models of inference processes under conjectural knowledge, pave the way to a deductive account of autonomous learning machines too.

4. From learning projections to constrained monotonicity

We have just seen that the allegedly inductive behaviours of machine learning algorithms *à la* ID3, involving generalizations and predictions from sample data, are sensibly construed as default-based, deductive inferences from theories and past observations. More specifically, these projections depend on background assumptions capturing both the inductive bias of ID3, which is invariant over the class of ID3 learnable concepts, and more "local" assumptions about the target concept c under examination. These local assumptions are embodied into the system's empirical experience, that is, in the training set selected for c. Thus, the retaining of any ID3 identification of target concept is non-monotonically conditional on the empirical adequacy of both sorts of assumptions. Let us notice, moreover, that in the behaviour-based systems examined above, similar issues about the empirical adequacy of background assumptions arise: one can choose between different learning strategies about obstacle avoidance and navigation on the basis of local information/hypotheses about spatial distribution of obstacles in the currently explored environment.

The aim of this section is to suggest that some families of non-monotonic consequence relations provide an appropriate *deductive* logical framework for investigating reasoning patterns of this sort which, in addition to ID3-like generalization and prediction, allow one to retract learning hypotheses and to revise background theories.

The present approach conflicts with the rather common view that non-monotonic inference falls outside the realm of deductive reasoning. Indeed, non-monotonic inference is occasionally dubbed as"quasi-deductive" [9], and more customarily as an outright non-deductive mode of inference, – the latter classification being based on the traditional partition of the domain of inference into deductive and non-deductive.[8] In particular, the present account

[8]This partition is assumed, for example, in a recent collection of cognitive science studies on human rationality:

> Other models of reasoning, which are not deductively valid, are "non-monotonic" –adding premises can lead to conclusions being withdrawn. An important example is induction, in which general laws or regularities are inferred from particular observations. [. . .] Another example is abduction, which

challenges the view that non-monotonic inference, unlike deduction, fails to capture "sound inference". In fact, the non-monotonic rules of inference discussed below, when interpreted on the basis of appropriate semantic notions, are provably sound, just like the inference rules of classical logic: in each of these rules, in fact, the truth of the premises (classically) implies the truth of the conclusion.[9]

A bit of formal setting is needed to proceed further on in our discussion. Let $\mathcal{SL}$ be the set of sentences built up from a classical propositional language $\mathcal{L}$ in the usual way. We use small Greek letters for elements of $\mathcal{SL}$ and block capital Greek letters for subsets of $\mathcal{SL}$. We denote by $2^{\mathcal{L}}$ the set of classical valuations on $\mathcal{L}$, that is the set of all maps from $\mathcal{L}$ to $2 = \{0, 1\}$. Valuations extend uniquely to $\mathcal{SL}$ in the usual, recursive way. We write, as usual, Cn for the *classical* consequence operation (sometimes called Tarskian), namely, an operation which is reflexive, transitive, and monotonic, the corresponding notation for the classical consequence relation being $\vdash$. As we shall be concerned with *finitary* consequence only (i.e. with premises being finite sets of $\mathcal{SL}$) we adopt the usual convention of freely swapping the notions of consequence operation and consequence relations.[10] Hence Cn and $\vdash$ represent classical logical inference which we assume to be deductive. Consequence relations on $2^{\mathcal{SL}} \times \mathcal{SL}$ other than the classical one will be denoted by $\mid\!\sim$, possibly with decorations. Note that all consequence relations we are going to consider below rest on the same underlying language $\mathcal{L}$.

Let us now begin our logical analysis from an examination of what is bad *and what is good* in monotonic reasoning, as far as the reasoning of AI agents is concerned.

If one assumes that subsets of $\mathcal{SL}$ represent agents' information, then

> typically involves inferring causes from their effects. [...] Thus, abduction is also nonmonotonic and hence not deductive. ([33] p.176)

[9]It is easily shown, though space constraints prevent us from delving into the details here, that *semantically*, non-monotonic consequence relations arise naturally from considerations of truth-preservance that *generalize* those introduced by A. Tarski in the 1930's, the back-bone of the formal characterization of deduction. We will cursorily say something more about the semantics of non-monotonic reasoning below.

[10]Indeed, given $\Gamma \subseteq \mathcal{SL}$, $\theta \in \mathcal{SL}$ and the operation Cn, one can define a relation $\vdash \subseteq 2^{\mathcal{SL}} \times \mathcal{SL}$ as the set of ordered pairs $\langle \Gamma, \theta \rangle$ such that $\theta \in Cn(\Gamma)$, and conversely $Cn(\Gamma) = \{\theta \mid \Gamma \vdash \theta\}$.

monotonicity just says that the introduction of additional information to the premises of some derivation does not force one to reject any of the conclusions previously drawn from the initial set of premises, *no matter what this new information turns out to be*. More compactly:

$$\text{if} \quad \theta \in Cn(\Gamma) \quad \text{then} \quad \theta \in Cn(\Gamma \cup \Delta). \qquad \text{(MON)}$$

The emphasized lack of qualification is an immediate consequence of the fact that no formal constraint is imposed on the additional set Δ. We call this the *unconstrained* form of monotonicity.

Since monotonicity does not impose any constraint on the enlargement of the set of premises, Δ can indeed be *any* subset of $\mathcal{SL}$. Hence, any possible addition to Γ will be *a priori* irrelevant. To put it into more graphic terms, once a monotonic agent draws a conclusion from a given set of premises, nothing the agent might possibly come to learn will cause it changing its mind. In realistic (i.e. "world-like") environments, however, agents do change their mind.

Still, as far as the characterization of intelligent reasoning is concerned, monotonic patterns of reasoning have a number of desirable properties, so that throwing away monotonicity as a whole will not do. A purely anti-monotonic agent, indeed, would waste an enormous amount of resources questioning over and over again each of the previously drawn conclusions. Thus, precepts of informational economy suggest a sensible balance between unconstrained monotonicity and anti-monotonic behaviours: *agents should not change their mind unless they have good reasons to do so*.[11]

The approach to nonmonotonic reasoning based on consequence relations focuses precisely on the characterization of suitable constraints which capture intuitively appealing precepts of informational economy. The seminal papers by Gabbay [12], Makinson [26] and Kraus, Lehmann and Magidor [23] have provided an abstract (proof-theoretic and model-theoretic) framework, suited for the logical investigation of nonmonotonic inference in general, and nonmonotonic consequence relations in particular.

A crucial difference between monotonic and nonmonotonic logics, as far as consequence relations are concerned, is that the latter admit no "smallest"

[11] In the literature on Belief Revision, informational economy (see e.g. [14]) requires that the revision of a belief set should preserve old beliefs "as much as possible".

consequence relation: various notions of nonmonotonic consequence arise as soon as certain conditions are fixed.[12] However, this distinguishing feature of nonmonotonic logics, and the corresponding variety of formal developments, are not symptomatic of a lack of shared underlying intuitions about core features of nonmonotonic reasoning. On the contrary, there is wide consensus that formal systems like the one based on Preferential consequence relations (so-called system P [23] or, after the initials of its authors' names, system KLM) do capture key features of nonmonotonic reasoning.[13] In the first place, the intuitive justification of KLM-like formal constraints (to be briefly presented below) in terms of "rationality", "minimality" and "informational economy" connects the investigations on constrained monotonicity to commonsense approaches to uncertain reasoning (see, e.g. [34, 35]), where rational reasoning under imperfect information is mathematically characterized as obedience to "common-sense" principles, some of which turn out to be remarkably close the KLM rules/conditions.[14] Moreover, the KLM system can be proved to be sound and complete with respect to various semantic interpretations of reasoning, respectively based on usualness, typicality, normality, very high probability, and so on.[15] Finally, there is a tight connection, indeed a genuine translation, between non-monotonic rules of inference and the classic postulates for *Belief Revision* introduced in [1].[16]

Let us now briefly recall the formal rules/conditions characterizing preferential consequence relations.[17] The inference system is presented here, as

[12](with the remarkable property, however, that the intersection of those various relations brings one back to classical monotonic logic [27, 28], the logic used in metatheoretic investigations of nonmonotonic logics.)

[13]For detailed treatments, the interested reader is referred to [4, 40, 21].

[14]The first formal aspects of this connection have been put forward by relating nonmonotonic logics to maximum entropy reasoning. See e.g. [21, 17].

[15]Friedman and Halpern recently suggested a formal explanation of this deep connection between the proof-theoretic characterization and a wide class of semantics for which nonmonotonic consequence is complete, based on the general framework of *plausibility measures* [11]. Particularly relevant to us is the semantical interpretation according to which the relation $\theta \mid\sim_P \phi$ holds if ϕ is (classically) satisfied in all the *most preferred* worlds in which θ is (classically) satisfied.

[16]The main idea underlying this translation is the so-called Makinson-Gärdenfors Identity, according to which inferring nonmonotonically a conclusion from a sentence θ amounts to revising a certain set of background assumptions by θ. For details, see [29, 15, 40, 21, 4].

[17]Note that the following rules of inference are generally understood also as constraints that

usual, as a Gentzen-style system with individual sentences as arguments of the consequence relations.[18] Relations satisfying the following conditions are called *rational consequence relations*, whereas relations satisfying all but condition RMO are called *preferential consequence relations*.

Reflexivity is the only axiom scheme:

$$\theta \mid\sim \theta \quad (\mathbf{REF}).$$

Three kinds of rules (or conditions) are imposed on relation $\mid\sim$. Rules of the first kind are "pure conditions": Left Logical Equivalence and Right Weakening:[19]

$$\frac{\vdash \theta \leftrightarrow \phi, \quad \theta \mid\sim \psi}{\phi \mid\sim \psi} \quad (\mathbf{LLE})$$

$$\frac{\theta \mid\sim \phi, \quad \phi \vdash \psi}{\theta \mid\sim \psi} \quad (\mathbf{RWE})$$

Rules of the second kind are the usual Conjunction in the conclusions and Disjunction in the premises:

$$\frac{\theta \mid\sim \phi, \quad \theta \mid\sim \psi}{\theta \mid\sim \phi \wedge \psi} \quad (\mathbf{AND})$$

$$\frac{\theta \mid\sim \psi, \quad \phi \mid\sim \psi}{\theta \vee \phi \mid\sim \psi} \quad (\mathbf{OR})$$

Rules of third kind express the formal constraints on monotonicity. These are Cautious Monotonicity and Rational Monotonicity:

---

consequence relations should satisfy in order to account for sensible patterns of reasoning. The latter being their intended interpretation as far as their justification is concerned.

[18]The first generalization to the infinitary case was studied in [10].

[19]Note that any consequence relation $\mid\sim^*$ such that $\vdash \subseteq \mid\sim^*$ is called *supraclassical*, and Reflexivity and Right Weakening entail that. Intuitively then, a supraclassical consequence relation extends the deductive power of classical reasoning. Very roughly, the desirability of this property is a consequence of the intuition that reasoning under perfect information is just a limiting case of reasoning under uncertainty. This extension, however, is not priceless. Supraclassical relations, in fact, generally have a "less regular behaviour", in Makinson's phraseology. See [28, 40] for more on this.

$$\frac{\theta \mathrel{|\!\sim} \phi, \quad \theta \mathrel{|\!\sim} \psi}{\theta \wedge \phi \mathrel{|\!\sim} \psi} \quad (\mathbf{CMO})$$

$$\frac{\theta \mathrel{|\!\sim} \phi, \quad \theta \not\mathrel{|\!\sim} \neg\psi}{\theta \wedge \psi \mathrel{|\!\sim} \phi} \quad (\mathbf{RMO})$$

Cautious monotonicity puts a very natural condition on the application of monotonicity: the addition of a sentence that was previously derived from some given premise should not induce one to abandon any consequence of that premise only. This is an important conservative feature of monotonicity that need not be rejected in characterizations of sensible reasoning.[20] Rational Monotonicity puts much a stronger constraint on the applicability of monotonicity. CMO is meant to filter out "irrelevant information", where an irrelevant sentence is understood, roughly speaking, as a sentence added to given premises which does not invalidate any of the conclusions drawn from those premises. A different interpretation of what it means for a sentence (ψ) to be irrelevant is provided by the requirement that one should rule out the possibility for the added hypothesis (ψ) to clash with conclusion (ϕ). This is accomplished, as far as RMO is concerned, by requiring that $\neg\psi$ *cannot* be (nonmonotonically) proved from θ.[21]

For the purpose of illustrating how the reasoning of autonomous learning agents can be interpreted in the light of these nonmonotonic consequence

[20] More extended discussion of Cautious Monotonicity, or Rational Monotonicity (see below), falls outside the scope of this paper. The interested reader is referred to [23, 24] for underlying motivations and justifications, and to [40, 21, 4] for general surveys of principles of nonmonotonic inference. [13] emphasizes (see especially chapter 3) the benefits of conservative reasoning for realistic agents.

[21] There is a widespread orientation, however, that weaker conditions might capture more adequately this idea of relevance. Rational Monotonicity is often criticized (see, e.g. [27, 40]) on the grounds that it fails to be a Horn-condition, thus preventing the resulting consequence relation from being closed under intersection. There are also concerns on the semantics of rational consequence relations which require a total ordering (or equivalent conditions) among preferences in order to yield soundness and completeness results for the axiom system. (See the seminal [23, 24] for a fully detailed account on (the semantics of) rational consequence relations). As far as "positive" consequence of "positive" knowledge is concerned, however, the significance of the latter concern is undermined by the characterization results by Lehmann and Magidor [24] according to which preferential and rational consequence relations are in fact equivalent. See [5] for more on this.

relations, we consider now Makinson's approach to bridging classical and (various) nonmonotonic consequence relations arising from the imposition of suitable constraints on monotonicity.[22] In particular, we focus on Makinson's method of generating nonmonotonic consequence relations from the monotonic, *pivotal-assumption* consequence relation.

Suppose that a distinguished subset of $\mathcal{SL}$ –K– sums up an agent *background assumptions*. We could intuitively think of K as to the corpus of sentences which is taken for granted by the agent when deriving a certain conclusion from an explicit set of premises. The *pivotal-assumption* consequence relation is constructed by reserving to the distinguished set K a special role. Indeed, it is just a classical consequence relation *modulo K*. Formally:

$$\Delta \vdash_K \phi \quad \text{iff for no} \quad v \in 2^{\mathcal{L}}, \quad v(\Delta \cup K) = 1 \quad \text{and} \quad v(\phi) = 0.$$

In other words, ϕ is a pivotal-assumption consequence of the set Δ relative to the pivotal set K, just if ϕ is a *classical* consequence of $\Delta \cup K$. Thus, a consequence relation is called a pivotal-assumption consequence if it is of the form $\vdash_K$ for some $K \subseteq \mathcal{SL}$. [23]

[22] In [28], D. Makinson examines three (monotonic) consequence relations which naturally provide such connection. More precisely, Makinson examines

> three different ways of getting out of a set of premises more than is authorized by straightforward application of classical consequence, without amplifying the language in which these premises are stated, which remains that of classical logic. ([28], p.74).

Another interesting example is discussed by Beirle and Kern-Isberner [3] who investigate the relation between classical and nonmonotonic reasoning (where the latter is mainly discussed through conditionals and probabilistic semantics) within the category-theoretic context of *institutions*.

[23] An important observation about pivotal-assumption consequence relations is that they behave in a decidedly classical way to the effect that the logico-mathematical properties of $\vdash_K$ cast no doubts about its deductive nature. Indeed, as determined by a characterization theorem due to Rott ([40], p.117), any closure operation which is also supraclassical, compact and satisfies the classical "disjunction in the premises" *is* a pivotal-assumption consequence. Note however that pivotal-assumption consequence differ from the classical one in that the former need not satisfy uniform substitution of arbitrary formulae for the elementary letters in a formula. Such consequences relations are called *paraclassical* by Makinson. See [28] for fuller details. Hence this relation provides a deductive bridge between monotonic and non-monotonic reasoning.

According to Makinson, the relation based on $\vdash_K$ allows one "to take some of the mystery out of what is known as nonmonotonic logic", and to reach an unprecedented understanding of nonmonotonic consequence relations.

A crucial motivation underlying the idea of marking K as a distinguished subset of $\mathcal{SL}$ is that of reserving to K the special role of a relatively entrenched set of background assumptions (the general heuristics or inductive bias for constructing any ID3-learning trees being a case in point in view of their independence of local information about particular target concepts). The set on the left of $\vdash_K$ is the set of "local" assumptions (or hypotheses), that is, information specific to a certain situation, choice, decision, etc.. (in ID3-learning, this role is played by specific training sets). A characteristic feature of $\vdash_K$, then, is that the set of background assumptions is assumed *not* to vary in relation to the local assumptions. This follows from the fact that each K determines a distinct relation $\vdash_K$. Up to suitable abstraction, $\vdash_K$ characterizes any agent whose behaviour is stricly determined by heuristics that are invariant over sets of local assumptions. In case of ID3, for example, a change of training set for concept c determined by a revision of local assumptions about the more representative training data for c does not affect the inductive bias of ID3 proper.

So far we have focused only on monotonic, assumption-based reasoning, which is characterized by assigning a special status to some unchanging set of background assumptions K. Far more interesting, however, are learning systems that adapt the set of background assumptions according to their actual "experience". These systems would qualify as more flexible and adaptable learners. And yet, under certain formal conditions, adaptable learners of this kind do not trespass the boundaries of deductive inference. This is indeed the scenario in which constrained-monotonic inference is at its best. The nonmonotonic consequence relation $\mid\!\sim_K$, in fact, can be derived from $\vdash_K$ just by allowing the set of background assumptions to vary according to the local assumptions. This "small" variation is all that there is between monotonic and non-monotonic consequence relations.

The main rationale for allowing the set of background assumptions to vary according to "experience" is that a certain set of local hypotheses Δ might be logically inconsistent with K. Consider the case of a behaviour-based system which is capable of detecting that current perceptual data clash with some background hypothesis driving its learning procedure, say, with the hypothesis that the environment is cluttered with obstacles. If such inconsistencies arise, $\mid\!\sim_K$ behaves non-monotonically by adjusting K so as to remove the inconsistency with local premises. The idea of *default assumption consequences*, thus, is that K should be replaced by its maximal subsets consistent with Δ.[24] Formally:

$$\Delta \mid\!\sim_K \phi \quad \text{iff} \quad K' \cup \Delta \vdash \phi, \quad \forall K' \subseteq K \quad \text{which is maxiconsistent with} \quad \Delta.$$

Among the various properties of $\mid\!\sim_K$, of particular interest to us is the fact that it satisfies Cautious Monotonicity, i.e. what is usually regarded as the key ingredient of "core" nonmonotonic reasoning.[25] Subsequent expansions of K' (say, for effective learning by the above behaviour-based system in an uncluttered environment) should be taken care of by means of appropriate heuristics.

In conclusion, the general notion of an autonomous learning system which, in addition to advancing learning hypotheses, is capable of retracting empirically inadequate hypotheses and revising its background knowledge base accordingly, does not require one to introduce notions of consequence that exceed or otherwise cannot be captured within the framework of deductive reasoning. And indeed, it is precisely the *extra-logical* consideration about the certainty, or the lack thereof, of information that underlies what is traditionally seen as a logical difference between classical deduction and various forms of uncertain reasoning. In the case of the classical consequence relation, what warrants the validity of the inference is its truth-preservation, or more precisely the preservation of the postulated truth of the initial statements. The same can –and ought to– be required from logical systems which are intended to capture patterns of reasoning in which the premises are not

[24] A subset $Y \subseteq X$ is maximal if and only if for all $Z \subseteq X$, $Y \subseteq Z$ implies $Y = Z$.

[25] Preferential reasoning is accounted for also semantically by means of paraclassical, monotonic bridges, and in particular the *pivotal-valuation* consequence discussed in section 3 of [28].

necessarily taken as stable truths, but only as elements of a revisable body of knowledge. This fact changes –and usually considerably complicates– the sort of constraints that are to be formalized, yet it does *not* change the logical form of the problem.[26]

References

[1] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50:510–530, 1985.

[2] R.C. Arkin. *Behavior-based robotics*. The MIT Press, May 1998.

[3] C. Beierle and G. Kern-Isberner. Using institutions for the study of qualitative and quantitative conditional logics. In S. Flesca, S. Greco, and N. Leone, editors, *JELIA 02*, volume 2371 of *Lecture Notes in Artificial Intelligence*, pages 161–172. Springer, 2002.

[4] A. Bochman. *A Logical Theory of Nonmonotonic Inference and Belief Change*. Springer, 2001.

[5] R. Booth and J. B. Paris. A note on the rational closure of knowledge bases with both positive and negative knowledge. *Journal of Logic, Language, and Information*, 7(2):165–190, 1998.

[26] This had already been effectively pointed out long ago by Kleene, when discussing in his *Mathematical Logic* applications of propositional logic to ordinary reasoning:

> Whether $\theta_1, \theta_2, \ldots, \theta_n$ are true or not may be a matter of empirical fact, or of belief, or may rest on earlier assumptions under which the argument is being pursued and which make $\theta_1, \theta_2, \ldots, \theta_n$ available for the purpose of the argument. Soundness is thus relative to whatever criteria or standards are being presupposed in the claim of the argument that $\theta_1, \theta_2, \ldots, \theta_n$ are available, and a full statement on the matter of soundness would include such reference. Also it seems convenient to recognize graduations by calling an argument simply *plausible* when it is valid but we can only say that $\theta_1, \theta_2, \ldots, \theta_n$ are plausible. ([22], 67–68).

[6] R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, pages 14–23, 1986.

[7] R.J. Clark, R.C. Arkin, and A. Ram. Learning momentum: online performance enhancement for reactive systems. In *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, volume 1, pages 111–116, Nice, France, May 1992.

[8] A.H. Fagg, D. Lotspeich, and G.A. Bekey. A reinforcement-learning approach to reactive control policy design for autonomous robots. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, volume 1, 1994.

[9] P.A. Flach. Modern logic and its role in the study of knowledge. In D. Jacquette, editor, *A Companion to Philosophical Logic*, pages 680–693. Blackwell Publishers, 2002.

[10] M. Freund and D. Lehmann. Nonmonotonic inference operations. *Bullettin of the Interest Group in Pure and Applied Logics*, 1(1):23–68, 1993.

[11] N. Friedman and Y. Halpern. Plausibility measures and default reasoning. *Journal of the ACM*, 48(4):648–685, 2001.

[12] D.M. Gabbay. *Theoretical foundations for non-monotonic reasoning in expert systems*, pages 439–457. Proceedings NATO Advanced Institute on Logics and Models of Cuncurrent Systems. Springer, Berlin, 1985.

[13] D.M. Gabbay and J. Woods. *Agenda Relevance: A Study in Formal Pragmatics*. North-Holland, 2003.

[14] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Bradford Books, 1988.

[15] P. Gärdenfors and D. Makinson. Nonmonotonic inferences based on expectations. *Artificial Intelligence*, 55:1–60, 1994.

[16] D. Gillies. *Artificial Intelligence and Scientific Method*. Oxford University Press, 1996.

[17] L.C. Hill and J.B. Paris. When maximising entropy gives the rational closure. *Journal of Logic and Computation*, 12(1):1–19, 2002.

[18] C. Howson. *Hume's Problem: Induction and the Justification of Belief*. Oxford University Press, 2000.

[19] S. Jain, D. Osherson, J.S. Royer, and A. Sharma. *Systems that learn*. MIT Press, Cambridge, 1999.

[20] L.P. Kaelbling, M.L. Littman, and A.P. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.

[21] G. Kern-Isberner. *Conditionals in Nonmonotonic Reasoning and Belief Revision - Considering Conditionals as Agents*, volume 2087 of *Lecture Notes in Computer Science*. Springer, 2001.

[22] S.C. Kleene. *Mathematical Logic*. Wiley and Sons, 1967.

[23] S. Kraus, D Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative inference. *Artificial Intelligence*, 44(1):167–207, 1990.

[24] D. Lehmann and M Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992.

[25] S. Mahadevan and J. Connell. Automatic programming of behavior-based robots using reinforcement learning. In *Proceedings of AAAI-91*, pages 8–14, Pittsburgh, PA, 1991.

[26] D. Makinson. *General theory of cumulative inference*, volume n.346 of *Lecture Notes in Artificial Intelligence*, pages 1–18. Springer, 1989.

[27] D. Makinson. General patterns in nonmonotonic reasoning. In D.M. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, pages 35–110, 1994.

[28] D. Makinson. Bridges between classical and nonmonotonic logic. *Logic Journal of the IGPL*, 11(1):69–96, 2003.

[29] D. Makinson and P. Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In A. Fuhrmann and M. Morreau, editors, *Proceedings of the Workshop on The Logic of Theory Change*, volume 465 of *Lecture Notes in Computer Science*, pages 185–205, Konstanz, FRG, 1991. Springer-Verlag.

[30] R. S. Michalski. A theory of methodology of inductive learning. In R.S. Michalski, J. Carbonell, and T. M. Mitchell, editors, *Machine Learning: An Artificial Intelligence Approach*, pages 83–134. Springer Verlag, Berlin, 1984.

[31] T.M. Mitchell. *Machine Learning*. McGraw Hill, 1997.

[32] U. Nehmzow and T. Mitchell. The prospective student's introduction to the robot learning problem. *Technical Report Series, UMCS95 -12-6, the Department of Computer Science, Manchester University*, 1995.

[33] M. Oaksford and N. Chater. Commonsense reasoning, logic and human rationality. In R. Elio, editor, *Common sense, Reasoning, and Rationality*, pages 174–214. Oxford University Press, New York, 2002.

[34] J.B. Paris. *The Uncertain Reasoner's Companion: A Mathematical Perspective*. Cambridge University Press, Cambridge, England, 1994.

[35] J.B Paris. Comon sense and maximum entropy. *Synthese*, 117(1):73–93, 1999.

[36] K.R. Popper. *Objective Knowledge: An Evolutionary Approach*. Clarendon Press, Oxford, 1972.

[37] J.R. Quinlan. Induction of decision trees. *Machine Learning*, 1(1):81–106, 1983.

[38] J.R. Quinlan. *C.4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, 1993.

[39] A. Ram, R.C. Arkin, K. Moorman, and R.J. Clark. Case-based reactive navigation: a method for on-line selection and adaptation of reactive robotic control parameters. *IEEE Transactions on Systems, Man and Cybernetics, Part B*, 27(3):376–394, June 1997.

[40] H. Rott. *Change, Choice and Inference : A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford University Press, 2001.

[41] G. Tamburrini. AI and Popper's solution to the problem of induction. In I. Jarvie, K. Milford, and D. Miller, editors, *Karl Popper, A Centenary Assessment*. Routledge, forthcoming, 2004.

[42] J. Woods. Standard logics as theories of argument and inference: Induction. In D.M. Gabbay, R.H. Johnson, H.J. Ohlbach, and J. Woods, editors, *Handbook of the Logic of Argument and Inference*, pages 105–171. North-Holland, 2002.