Relationships between Non-Monotonic Reasoning and Incremental Learning: Preliminary Outline of Invited Talk

Benjamin N. Grosof
IBM T. J. Watson Research Center
P.O. Box 704, Yorktown Heights, NY 10598
(914) 784-7100 ; Internet: grosof@watson.ibm.com

Abstract

We observe that non-monotonic reasoning arises in many kinds and aspects of incremental learning, including: induction, preference-type bias, and shift-able bias; learning by being told; knowledge assimilation / integration; knowledge-base refinement / theory revision; explanation-based learning; and sensory perception. We review what the field of formal non-monotonic logical knowledge representation has to offer the subject of incremental learning. This includes: algorithmic methods for inference and updating / revision; formal properties and guarantees; and concepts such as default, model-preference, and prioritization.

In our attached paper "Representing and Reasoning with Defaults For Learning Agents", we discuss in more detail a number of knowledge representation issues revolving around the use of defaults for learning agents.

1 Introduction; Learning Agents

By a **learning agent**, we mean a persisting agent that learns while using a knowledge base: it uses its current beliefs to support and constrain its current learning (i.e., it does *knowledge-based* learning), and it integrates the result of that learning, including perception and received communication, into its knowledge base (*knowledge integration / knowledge assimilation*).

Formally, we define a learning agent as an agent that, dynamically and on a continual basis, acquires new premise beliefs (e.g., from perception and received communication) and revises a *stored body of concluded beliefs* accordingly. Its conclusions are not necessarily a logically monotonic function of its premises; that is, its reasoning is, in general, logically *non-monotonic* when viewed declaratively.

In learning agents, almost all learning can be viewed as **incremental**; and almost all informational input (e.g., perception and received communication) can be viewed as **training** input.

From a declarative point of view: many, if not most, kinds of learning involve non-monotonic reasoning. This suggests applying available tools from the field of knowledge representation (KR) about non-monotonic reasoning (NMR) to both the knowledge representational and the computational aspects of incremental learning.

Why take a declarative view of the knowledge base (background and hypotheses) and reasoning activities of a learning agent?

1. Generally, to gain the advantages usual in AI of declarative approaches. (See sub-section 1.1 of our attached paper "Representing and Reasoning with Defaults For Learning Agents".)
2. More particularly, to exploit the concepts, formal properties and guarantees, and algorithmic methods that are available from NM KR.
3. In terms of incremental training: declarative methods are typically insensitive to the sequence of the premises, and thus are **independent of the ordering of training data**.

2 How Non-Monotonicity Arises in Several Kinds of Learning

Induction:
Inductive leaps, **shifts of bias**, and **preference-type bias** are typically logically non-monotonic. (By "inductive", we mean involving inductive leaps that go beyond what is known for-sure and posit a falsifiable hypothesis.) Shift-able / preference-type bias can be specified via defaults or other non-monotonic logical concepts. For more discussion of this point, see sections 1 through 3 of our attached paper "Representing and Reasoning with Defaults For Learning Agents" (i.e., [Grosof, 1992b]).

Learning by Being Told / Taking Advice:

(This is one type of **knowledge assimilation / integration**.) Advice that is potentially incorrect / errorful, or that may conflict with other advice or beliefs, is most appropriately treated as default in flavor. Non-monotonic logic offers principles for normative specification of how conflicts should be resolved. We formalize an approach to this issue in [Grosof, 1993b], which uses the DAP circumscription formalism [Grosof, 1993a] discussed in section 8 of our attached paper "Representing and Reasoning with Defaults For Learning Agents".

Knowledge-Base Refinement / Theory Revision:
Beliefs, e.g., rules, that are potentially inconsistent / incorrect / revisable are most appropriately treated as default in flavor. Revising the knowledge base / theory after an informational update (e.g., new test case) then often involves the need to retract previous rules, conclusions and explanations (justifications). In the field of logicist NMR, this is called the *belief revision* problem.

Explanation-Based Learning:
Explanation from beliefs that are potentially revisable / default-flavored is a kind of non-monotonic reasoning. The justification of a non-monotonic conclusion depends not only on what rules and facts appear positively in its proof, but also by the lack (**omission**) of any conflicting overriding rules and facts. This makes explanation-based generalization very much trickier than in the case of monotonic reasoning: there may be a huge or infinte space of such omission conditions.

Errorful Observations; Sensory Perception:
(This is one type of **knowledge assimilation / integration**.) Any kind of potentially errorful observational data may be viewed as only plausible in status, and thus appropriately represented with defaults or other non-monotonic logical concepts. In particular, reasoning of the kind: "if you see something that looks like a duck, then conclude it is a duck" involves logical non-monotonicity. More generally, justifying reasoning of the kind "if you see a bird and hear quacking, then conclude it is a duck" on the basis of the knowledge "ducks are birds that quack" is a kind of *abduction*; there may be other possible quacking birds or sources of quacks (e.g., hunters' whistles). Much of the initial stages of sensory perception involves non-monotonic reasoning from observed effect to conjectured cause. Musto and Konolige [1993] formalize an approach to this issue in perception.

Various Kinds of Knowledge:
In general for knowledge-based learning, NM KR appears appropriate for many different useful kinds of knowledge and reasoning, including: common-sense; causality; time, action, and plans; inheritance hierarchies with exceptions; logic programs with negation as failure; laws and regulations; design choices; goals; and much more. NM KR, especially the logical concept of a default, provides a formal declarative treatment of the revisable status of beliefs.

3 What Non-Monotonic KR Can Offer

Next, we describe some examples of what current NM KR can offer to the field of (incremental) learning. **Recall**: the list at the end of section 1.

Algorithmic Methods

A large variety of algorithms are available for various kinds of NMR, including:

- **Logic programs** with negation-as-failure.
- **Truth Maintenance Systems (TMS)**, including both Doyle-style [1979] and Assumption-based [de Kleer, 1986]. Many TMSes can be viewed as a special case of logic programs. TMSes can be used as a core sub-procedure to support many different kinds of NMR tasks, including the finding of contradictions.
- Default **inheritance**, i.e., hierarchical inheritance with exceptions.
- Query-answering in many different **highly expressive** NM logical formalisms, including circumscription, model-preference logics, Default Logic, Autoepistemic Logic, conditional logics, modal logics, possibilistic logics, and others.
- Updating in Gärdenfors-style so-called "belief revision" systems [1988] that are characterized by a particular principle of "**minimal change**" in the retraction phase of updating.
- **Locality methods** in inference and updating, including for: prioritized circumscription [Grosof, 1992c] [Grosof, 1992a]; and a "context"-based formalism [Giunchiglia, 1992].

Formal Properties and Guarantees

NM KR can provide various useful formal properties and guarantees, both semantically and computationally. Most basically, these include soundness and completeness results for reasoning algorithms. Next, we list a few others:

- Constrained expressive classes (e.g., Horn, etc.) with better worst-case **computational cost**.
- **Monotonicities of updating** (e.g., [Grosof, 1992c]): guarantee that some categories of previous beliefs will not need to be retracted for particular categories of updates. For example, this lowers the cost of testing revisions (e.g., updated hypotheses) against the overall KB (e.g., previous observations / test cases).
- Reformulation and transformation theorems.

Concepts

NM KR also provides a number of tools for conceptualizing incrementally learning agents' knowledge and reasoning. One such tool is the concept of a **default**. More generally, NM KR offers various paradigms for **specifying preferences** among alternative candidate revisions / belief sets. One such paradigm is to specify a **transitive preference relation between individual models** in classical or modal logic.

Another useful concept is that of *causal* versus *evidential* directionality in reasoning. E.g., ducks cause quacking sounds is a causal statement. But non-monotonically inferring that quacking sounds indicate the presence of a duck is reasoning in the *evidential* direction, i.e., from effect to cause.

Prioritization-type precedence between defaults is still another useful concept. For one default to have precedence over a second means that in cases of conflict between them, the first "wins": its conclusion goes through. E.g., if the default that Republicans are not Pacifists has precedence over the default that Quakers are indeed Pacifists, then the conclusion that Nixon (who is both a Quaker and a Republican) is *not* a Pacifist is sanctioned.

There are a number of NM KR concepts that need further development in order to deal properly with the issues posed by incrementally learning agents. We discuss several such NM KR issues in our attached paper "Representing and Reasoning with Defaults For Learning Agents", including: meta-reasoning about prioritization and adoption of defaults; and integration of probabilistic / statistical reasoning with defaults.

References

J. de Kleer. An assumption-based truth maintenance system. *Artificial Intelligence*, 28:127–162, 1986.

J. Doyle. A truth maintenance system. *Artificial Intelligence*, 12:231–272, 1979.

Peter Gärdenfors. *Knowledge In Flux: Modeling the Dynamics of Epistemic States.* MIT Press, Bradford Books, Cambridge, Mass., 1988.

Fausto Giunchiglia. Contextual reasoning. *To appear in: Epistemologia, speical issue on I Linguaggi e Le Macchine*, 1992. Also available as IRST-Technical Report 9211-20, IRST, Trento, Italy.

Benjamin N. Grosof. Reformulating non-monotonic theories for inference and updating. In Michael Lowry, editor, *Proceedings of the 1992 Workshop on Change of Representation and Problem Reformulation*, Moffett Field, California, April 1992. NASA Ames Research Center Technical Report FIA-92-06. Workshop held Asilomar, California, April 28 – May 1, 1992. Also available as IBM Research Report RC17955.

Benjamin N. Grosof. Representing and reasoning with defaults for learning agents. In Diana Gordon, editor, *Proceedings of the ML92 Workshop on Biases in Inductive Learning.* Available as an IBM Research Report. Copies of Proceedings are available through the editor, at Naval Research Lab, Code 5510, 4555 Overlook Ave. S.W., Washington, D.C. 20375-5000., 1992.

Benjamin N. Grosof. *Updating and Structure in Non-Monotonic Theories.* PhD thesis, Computer Science Dept., Stanford University, Stanford, California 94305, October 1992. Published by University Microfilms, Inc.. Also available as Research Report from Stanford University Computer Science Dept. and/or IBM T.J. Watson Research Center.

Benjamin N. Grosof. Defeasible and pointwise prioritization: Preliminary report. Technical report, IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598, January 1993. Revised from Working Paper of April 1992.

Benjamin N. Grosof. Prioritizing multiple, contradictory sources in common-sense learning by being told; or, advice-taker meets bureaucracy. In Leora Morgenstern, editor, *Proceedings of the Second Symposium on Logical Formalizations of Common-Sense Reasoning (Common-Sense '93)*, IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598, 1993. Available as an IBM Research Report. Copies of the Proceedings are available via the editor (wider publication being arranged). Held Guest Quarters Hotel, Austin, Texas, Jan. 11-13, 1993.

Daniela Musto and Kurt Konolige. Acquisition of beliefs through perception. In Leora Morgenstern, editor, *Proceedings of the Second Symposium on Logical Formalizations of Common-Sense Reasoning (Common-Sense '93)*, IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598, 1993. Copies of Proceedings are available via the editor (wider publication being arranged). Held Guest Quarters Hotel, Austin, Texas, Jan. 11-13, 1993.