

An Exploration into Incremental Learning :

the INFLUENCE system

Antoine CORNUEJOLS

Laboratoire de Recherche en Informatique, UA 410 du CNRS
Université de Paris-sud, Orsay
Bâtiment 490, 91405 ORSAY (France)

Abstract

Although some recent works are concerned with the incremental acquisition of knowledge, these have often been conceived as modifications of existing algorithms or models added afterwards, not as investigations in their own right in a central problem of learning : how is a system to learn incrementally through experiences and time, and how is it to make do with limited computational and memory resources ? In other words, how is the temporal and dynamic dimension of learning to be taken into account ?

This paper proposes a clearcut and operational definition of "incremental learning" as opposed to "direct learning" and "sequential learning". In particular this definition leads to the consideration of the "best education" of a learning system, that is how to best organize available data so as to guide the system towards a desired functionality.

A system called INFLUENCE is then described that adapts incrementally its model of the world. An important characteristic of this system that is both significant and novel, is that it is naturally functioning out of equilibrium, allowing it to respond in real time to the environment, contrary to the all-pervasive quasi-static systems, but giving rise to new problems on the best learning rate.

In that way, INFLUENCE allows experiments encompassing the whole spectrum of problems (ordering of data and rate of presentation) specific to a learning process that takes place over time. Results obtained so far are reported in the last section of this paper, and the conclusion underlines lessons and perspectives for this research.

1. Motivation

Consider your learning of the language Lisp, or the way you came to distinguish between wines or music styles, look at a child developing and growing in its grasp of the world, remember the way your computer science department had to go to master, maintain and adapt to different hardware and software configurations over time. It is apparent that most realistic learning tasks are characterized by their unfolding over a certain duration, which, often, extend over the entire life of the learning system. Yet, except for the pioneering works of the earlier cyberneticians and connectionists and of some precursors in symbolic learning, the vast majority of the research efforts has been concerned "only" with "direct learning". Direct, because as in direct access memory, the whole set of relevant data is supposed to be available at once, with the additional assumption that the system is able to grasp its totality in a single learning step. This state of affair is quite understandable in view of the difficulties faced even in this simplified situation. However, we feel that on the one hand too a great ignorance of the specific problems imposed by incremental learning could lead to a prejudicial forgetting for reasons that will be soon developed, and on the other hand that we can start interesting investigations in this domain now.

Most current research on learning (empirical or explanation-based) still relies implicitly on three assumptions:

- unlimited or at least sufficient computing power to treat all data at once
- unlimited memory resources that allow to keep data in their original form for treatment
- all pertinent data available for the learning task

All three assumptions are doomed to be dropped in face of realistic learning tasks where computing power is limited compared to the wealth of data and the time requirements, memory space is limited conducting to pre-processing stages and some form of data compression hence forgettings, and decisions and actions have to be taken before all pertinent data are available or can be analysed. Therefore a capacity for incremental learning is necessary. But what really is *incremental learning*, and is it fundamentally different from *direct learning*, or can we find some transformation that make them equivalent ?

The naive approach is to call incremental any learning process where the training data are provided in a sequential manner to the learning system. Obviously this definition cannot found a distinction between learning types since one can always allow the system to store all data before starting to process them. A more sophisticated approach is to consider systems where data are processed and their raw form forgotten as they become available as incremental learners. I have shown elsewhere (Cornuéjols, 1989) that this distinction is still insufficient since one can exhibit simple learning systems where data are considered one at a time and not kept in memory in their explicit form, and still producing a behaviour that is quite undistinguishable from a system that would consider all data at once¹. I call this type of learning *sequential*. A stronger restriction is thus required, a restriction that takes into account differences in behaviours. I therefore propose the following definition:

A learning process is incremental when the "history" of the training period is influential on the function reached by the system.

That means, for instance, that different orders of presentation of the same training data can produce different functionalities of the training system.

Using this vocabulary, we get the following picture of the possible learning situations :




Figure 1: There are fundamentally two types of learning process, the *direct* one, and the *incremental* one where teaching history has an impact on performance. This last one covers most of the learning situations. *Sequential* learning

¹Besides, who would be responsible for judging if the form under which information is kept in the system is fundamentally different from the form of the raw data ? Systems are to become increasingly opaque to the observation of their internal workings as shown for instance with connectionist systems.

can be of either one type or the other depending on the commutativity of the messages received from the environment or the teacher.

Most learning tasks certainly implies sequential learning as we have seen. Now, *one immediate question is whether it is possible to make any incremental learning process a direct one* (where order of data presentation does not affect performance), *and, if yes, under which conditions*. A preliminary theoretical study of this question can be found in (Cornuéjols, 1989). It shows that this transformation, when possible, imposes very stringent conditions on the types of messages coming from the environment, and on the amount of information that has to be kept by the system. *A second question, of practical interest, is how to determine the best, or at least a good, teaching strategy in case of sequential learning*. This question covers the choice of the data to be presented to the system as well as the order of this presentation. This is akin to the controllability problem in mathematical system theory, but it is even more difficult to formalize and to solve in Artificial Intelligence. We can try however to get a grasp on this problem by experimenting with systems that are sensitive to their "education trajectory". This is, in part, what we have done with the system INFLUENCE described in the following.

2. Previous related works

Very few works exist on the specific problem of optimizing an Artificial Intelligence system education. One can cite the paper of (Mac Gregor, 1988) which deals more particularly with monotonic learning tasks where no revision of interpretation has to be done. Strangely, people that study connectionist systems or Genetic Algorithms, which are intrinsically alike dynamical systems, have not, to the best of my knowledge, delved (yet) into inquiries on the temporal aspect of learning. Likewise, but for more understandable reasons, people working on non-monotonic logics and reason maintenance systems have ignored this question. A slight departure from these logicist studies is (Elgot-Drapkin et. al., 1987) where the authors deal with systems that have limited computing and memory resources and must cope with the world in real-time, however they do not touch the problem of continuous learning.

3. Incremental learning with INFLUENCE

INFLUENCE is a system that was designed to experiment a different approach to knowledge maintenance from the more classical TMS like ones. Basically, the model relies on "spontaneous" changes in the knowledge base under local control when new information is added or under the pressure of internal perturbations. A good image is one of a chemical molecule that re-organizes its configuration in face of new constraints without the help of an external supervisor. This allows us to see the process of knowledge acquisition more as an *adaptive* one than one of creation of a knowledge base from scratch, and leads naturally to the study of sequential and incremental learning.

3.1. The domain

One domain of application that seems particularly appropriate for the study of the problems of incrementally building a model of the world is Natural Language Processing. Indeed, it presents us with *the typical constraints one encounters in incremental learning*. Data are provided and analysed sequentially. They can be ambiguous and their interpretation necessitates to take into account the context, that is the informations currently known to the system. The integration of these data to the memory implies a change in their representation. And, finally, the global interpretation or model of the world rests generally on incomplete and uncertain informations and therefore on hypothetical inferences that may be contradicted by subsequent data. Hence the system must be able to detect resulting inconsistencies and correct them in some way. All these problems are to be found in one form or another in any sequential learning system.

3.2. The INFLUENCE system²

INFLUENCE is a realisation for artificial intelligence purposes of a dynamical system builded out of numerous basic autonomous entities. It is known that these systems can

² This section intends to give only a brief overview of the INFLUENCE system, a detailed account of the system and of the corresponding experiments can be found in (Cornuéjols, 1989) and (Cornuéjols, 1987).

exhibit quite interesting self-organized and adaptive behaviours when immersed in environments rich in flows of energy or matter. How can one use this metaphor in information systems ?

It would be nice to take a collection of very primitive interacting devices, such as cells in a cellular automaton, and let them self-organized when confronted with a given environment. The problem is to be able to bridge the gap between what we would like the system to realize and what is realized at the level of the cells. In Physics or in Chemistry this is still an open problem in its generality, in Artificial Intelligence it appears sometime under the guise of the credit assignment problem, and except for very constrained systems (such as back-propagation schemes in neural networks) it is immensely difficult to solve. This is due to the difference of description between the two levels: global and local. In part for this reason, and in part because there existed already a way of representing knowledge in AI with objects that could be made autonomous, we choose to found our system on frame like entities.

There are three advantages to this choice. First, the gap between the description of the behaviour of the global system and the description of the local representation is leveled out. Second, frames can conveniently be considered as self-contained agents. Third, it is easy to cast existing frame-based Natural Language Processing systems (such as BORIS (Dyer, 1983)) into the INFLUENCE framework.

These considerations are important since they determine somewhat *the type of learning realised*. Here, the system is supposed to be endowed from the start with the adequate concepts necessary to perceive and interpret the environment, these correspond to the available library of frames and the system does not have to learn them, however the task is to progressively build a model of the world (a kind of mega-concept) in term of which incoming messages can make sense, and decisions and actions can be made. We call this type of learning *episodic* or *contextual learning*. A system that learns to make diagnostics in certain types of situations falls into this category.

In frame-based systems, interpreting, understanding and inferring mean instantiating schemas or frames and establishing links between them. Conversely, retrieving informations mean identifying some relevant structure in memory and "reading" from there the appropriate schemes and links. This is what is routinely offered as facilities in existing

frame-based systems. These systems however are generally unable to allow knowledge revision except in some cases under the supervision of a heavy supervisory system à la TMS (Doyle, 1979)(Granger, 1980)(O'Rorke, 1983) and with the need to have kept most informations from the raw data in the first place, two unpleasant requirements.

Yet, it suffices to destroy some links and frames and to construct new ones to change the content of the memory. This is what does INFLUENCE. In this model, the knowledge base is conceived as a self-organized organism where the *links are autonomous agents* that can decide for themselves if and when to detach from their current target to look for a new one. This process is guided by *local informations* on the adequation between constraints inscribed in the slot corresponding to the link and the current target. The better this adequation is, the stronger is the corresponding link and the more unlikely is its detachment from the target. The uncertainties of the system are thus translated into the instabilities of the links in memory.

Meanwhile, the choice of a new target is determined by the respective "*influences*" or activities of the potential targets. These influence factors represent the importance of the attached scheme or concept in the current context. They are directly depending on the connections maintained by each frame in memory. Specifically, the use of these coefficients allows to attract links preferably towards those concepts that are the most central in the current situation rather than towards more peripheral ones. This implements in a way the parcimony principle that dictates to prefer simple and economical explanatory schemes rather than complicated ones.

Within this framework, possible interpretations of a given set of informations are alike local equilibrium states corresponding to different basins of attraction. Revision or modifications to the current model of the world occur when perturbations, either external (tied to incoming data) or internal (due to the instabilities of the uncertain links) push the system towards other attractors. The traditional phases of inconsistency detection and belief revision are thus tightly mixed, and cannot be in fact distinguished.

Typically, at any given time, there are some links in the memory network that are "uncommitted" from their current targets and are looking for new ones. The probability that any one of them attach to a given target depends on the relative influence of the latter. If the current network corresponds to a satisfactory state of the memory, links are rather stable

(local adequations are good) and current targets have large influences making high the probability that the rare uncommitted links re-attach to the same target, leaving the memory unchanged. The network is then quite stable. On the other hand, if the current memory model is unsatisfactory, links tend to be unstable, and frames that were previously rather poorly influent may get a chance to become more central yielding in turn other modifications and leading to a new organization of the network and to a change of the overall interpretation.

The same may happen when new data arrive. Either the frame representation of these new informations fits well within the existing network resulting in an extension of it and a corresponding elaboration of the current world model, or they do not fit, leading to one of two possible outcomes. One is that the new data are simply forgotten because of their being not significant enough to maintain themselves in memory in isolation, the other is that a more or less important re-organization of the memory has to take place to accomodate the new data.

For a casual observer, the whole process of data integration into memory is quite analogous to the way a chemical molecule re-arranges spontaneously its structure when coping with new stresses.

One noteworthy point besides is that the system never reach rest except in the rare cases when memory is empty (an hopefully extreme situation indeed), or when all links represent perfectly certain assertions and are thus totally stable. That means in particular that the dynamics of the system must be taken into account when controlling interactions with the environment. If incoming data are too sparse, the risk is high that too much forgetting takes place between incoming messages to maintain a faithful model of the world. If, on the other hand, they are coming too fast, the risk exists that the system be overwhelmed and not be able to accomodate for all of them. This is at least the intuitive view.

To sum up, because INFLUENCE is a model of adaptive learning taking place over time, and because its out of equilibrium mode of functioning distinguishes it from classical quasi-static system where one quietly waits for the system reaching a stable state, one can expect interesting experiments and results about all aspects of incremental learning. The next section provides a glimpse of such experiments.

3.3. Experiments

Numerous and various experiments have been realized with the INFLUENCE³ system applied to Natural Language Processing tasks. The typical protocol was the following: various entries corresponding to a given text were translated by hand⁴ under the form of their canonical representation in frame-based structures, and then these fragments of semantic networks were entered into the system in various orders or/and various rates (with respect to the internal inferences cycles of the system).

One text for instance read as follow :

"When John learnt in the news that he had won the lottery, he bought immediately an airplane ticket for Indonesia. He was very excited. He rushed to make his luggages, and left his apartment without closing the gas. Three days later in Djakarta, he learnt that two mafiosi had died in the explosion of his apartment. He was relieved."

Note that this text presents an interesting property since most of us will at first think that John went abroad for vacation before choosing, upon reading the last sentences, the less mundane interpretation that he was fleeing the mafia. The questions are, will the system be able to do the same, and will its understanding process be more involved than if the data had been in a different order with the last items of informations appearing first ?

One would thus compare the learning and interpretation processes when in situation 1:

1st input : *"When John learnt in the newspaper"*

2nd input : *"he had won the lottery"*

3rd input : *"he bought an airplane ticket"*

4th input : *"he was very excited"*

5th input : *"he made his luggages"*

³ INFLUENCE has been implemented with Allegro Common Lisp on a Mac Iix micro-computer.

⁴ But existing parsers could be used as well with minor modifications so as to attribute a certainty factor to each link depending on syntactical constraints and heuristics used during the parsing.

6th input : *"he left his apartment without closing the gas"*

7th input : *"John knew that two mafiosi were looking for him"*


8th input : *"John learnt they had died in his apartment explosion"*

and in situation 2 where the last two items are entered in the system before the others.

In other experiments, one would vary the time delays between inputs, say 0 cycle (all data provided at once), 10 cycles, 20, 50 and so on, and with various parameters combinations that limited space prevent us from discussing here.

The first interesting result is that the system is indeed able to change correctly its interpretation of the world when subsequent data are in conflict with the current one, as for instance when input 7 and 8 are taken into account in situation 1 above.

The following picture gives an idea of the progress of the learning task as inputs are provided to the system (x axis). The y axis represents the completion (in a crude approximation) of the current memory model measured in percentage against the desired final model. One can observe that the completion measure increases with each new input, but also that there is a jump when inputs 7 and 8 are provided. Note too that the curve is going up between inputs as well, reflecting the fact that the system continues to make inferences and to build its memory network.



As seems reasonable afterwards, with a system presenting good abilities for re-interpretation the effect of the differing orders of input streams were not as strong as expected. Overall, with different sets of data and after numerous experiments, *it was found that learning and interpreting was difficult mostly in those cases where data with high but ambiguous information content were given at first without rapidly providing keys that would allow disambiguation.* Otherwise it was found that the best results in terms of learning speed and performance were attained when giving first key informations followed by complementary ones that allow the system to develop its memory network and to confort its first inferences.

Of course, these observations would deserve to be more formalised. The *information content* of each piece of data was measured with respect to the number of frames and links necessary to represent them, and in term of their likeliness. For instance you can arrange your background knowledge (expressed in the form of the frames and the procedures attached to their slots) so that it is likely that you learn on the radio that there was a coup d'état somewhere in the world, but unlikely that you learn on the same radio that your reseacher salary has been raised. In fact, in this last case, the information content depends also on the context. Whether you are yourself a chief of state or a reseacher can totally change the impact of the data on your view of the world. This is why *the learning system itself should be the judge in the last resort of what is informative or not.* This has been done using some simple heuristics in INFLUENCE, but space limitations prevent us to elaborate on this here.

As for the effects of the rate of input data on the learning process, results show that the highest rates are the best ones in terms of the correctness and completion of the final memory model obtained. On the other hand, intermediate rates appear better with respect to the speed of learning each individual data. Lowest rates lead to irremediable loss of informations (note that the forgetting phenomenon is interesting in itself: unimportant details are the first to disappear).

These results are somewhat surprising since one would expect that giving all data at once would result into great difficulties for the system to analyze them all. It must be noted however that the INFLUENCE model is misleading in some sense with respect to this problem because it implements massive parallelism and does not have, in principle, limited

computing resources. This is why the system can absorb, and reflect upon, large batches of data.

In this way direct learning seems better than incremental one. But as we have seen, even if this rule was to be confirmed, realistic learning situations would prevent its use.

4. Conclusion

As the uses of Artificial Intelligence systems increasingly involve more automatic learning extending over the life of the systems from incomplete and uncertain informations, sequential learning (where data are considered as they come) and adaptation of the knowledge bases will become the norm rather than the exception. It thus happens that for all but some very constrained cases, sequential learning implies a sensitivity on the "history" of the learning process, a phenomenon I call incremental learning. It is then crucial to ask oneself how to optimize the education of such systems so as to reach the desired performance levels.

Likewise, the growing complexity of Artificial Intelligence systems and of the tasks to them devolved, will render untenable the assumption of quasi-staticity required nowadays in their operating. Future systems will have to be able to function out of equilibrium. Now, how are we to control and guide such systems ? Answers to this question remain largely in the realms of fantasy.

We advocate that more research is needed, and urgently needed, to explore the temporal or dynamic dimension of learning. Large theoretical studies are necessary. They will undoubtedly be difficult and involved. In the meantime, experimental studies can prove themselves worthwhile. The system INFLUENCE was designed to have a stab at such an inquiry. Results obtained so far are still modest (I think that one reason is that there are too many problems in dealing with open domains like Natural Language Processing, particularly at the level of knowledge representation. Further experiments should be realised in more constrained and better formalised domains like fault diagnosis). I hope however that these experiments point towards interesting directions.

REFERENCES

- Antoine CORNUEJOLS (1987) : "**Revision of "interpretation" in Episodic Memory by using chemistry instead of reason maintenance systems.**" *Proc. MARI-87*, Paris, may 18-24,1987.
- Antoine CORNUEJOLS (1989) : "**De l'Apprentissage Incrémental par Adaptation Dynamique: le système INFLUENCE**" (Incremental Learning through Dynamical Adaptation: the INFLUENCE system). Ph.D. thesis. University of Paris-Sud, Orsay, France, january 6th 1989.
- Jon DOYLE (1979) : "**A Truth Maintenance System**", *Artificial Intelligence Journal*, 12, 1979, pp.231-272.
- Michael DYER (1983) : **In-Depth Understanding**, MIT Press, 1983.
- J. ELGOT-DRAPKIN, M. MILLER & D. PERLIS (1987) : "**Life on a Desert Island: Ongoing Work on Real-Time Reasoning**", in the *Proc. of the 1987 Workshop on The Frame Problem in Artificial Intelligence*, Frank Brown (ed.), Lawrence, Kansas, avril 12-15, 1987.
- Richard GRANGER (1980) : "**When expectation fails: Toward a self-correcting inference system**" in *Proc. of the First National Conference on Artificial Intelligence*, Stanford, California, 1980.
- James N. MacGREGOR (1988) : "**The Effects of Order on Learning Classifications by Example: Heuristics for finding the Optimal Order**", *Artificial Intelligence*,34, pp.361-370.
- Paul O'RORKE (1983) : "**Reasons for beliefs in understanding: applications of non-monotonic dependencies to story processing**", *Proc. of the NCAI-83*, pp.306-309.