

Artificial Intelligence and Large Scale Scientific Experiments

Scott H. Clearwater
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, California 94304
clearwat@xerox.com

Abstract

As ever larger scientific endeavors are made, it becomes even more crucial that they be designed, performed, and analyzed very carefully. Artificial intelligence (AI) techniques are seen as a methodology for automating some of the tasks involved in large scale scientific experimentation. Scientific experimentation, in turn, is seen as a methodology for improving the capabilities of AI through the special demands it makes. This paper discusses how large scale scientific experimentation can be helped by AI and vice-versa.

Large Scale Scientific Experiments

Scientific experimentation has challenges that no other domain has to offer. It differs in large part by its exploratory nature which means that the theories and tools used to test the theory must be state of the art and so are not as reliable as traditional methods, such as those used in industrial applications. This is exemplified by the fact that an experiment is usually only performed one time, as opposed to an industrial application that are often performed many times. This means that expertise in an experiment is constantly evolving and puts extra demands on the management of it. Also, a scientific result demands quantitative verification to a level that exceeds most other domains.

As scientific experiments become more and more complex they also become much more costly and it becomes imperative that the experiments be carefully designed to insure they will be able to meet their goals. High Energy Physics (HEP) is the biggest of "Big Science" endeavors and paradigmatic of large scale experimental science. As such, HEP can serve as a testbed for applications of intelligent computation technology to scientific experimentation. (In this paper, intelligent computation will refer to AI methods such as expert systems or learning programs.) At the same time, the special demands of scientific experimentation allow AI researchers a chance to improve their methodologies to handle more of the scenarios presented by experiments.

Specifically, AI has a role in large experimental settings as a means of automating, at least to some extent,

the design and data analysis of an experiment. In the design phase, AI can be used to verify the design of an experiment more thoroughly and quickly than can be done by scientists directly. The efforts of the scientists can then be focussed on the more abstract characteristics of the experiment while a program handles the computational details. This gives the scientist more time to conceptualize the experiment as well as spend time on other aspects of the experiment. In the data analysis phase, AI can be used to explore many more possible analyses of the data so that the most effective one can be found.

AI has something to learn from experimental science as well. The Scientific Method has not been codified and any part of it that can be will improve the ability to do science.

Critical Issues

Critical issues of large scale scientific experimentation are related to design, operation, and data analysis. Of these, the first and last are related. In particular, they involve comparisons between the data collected by the experiment (simulated data in the case of design) and that predicted by a theory. The design must be robust enough to insure that the experiment will be able to measure the desired phenomena even if the parameters that went into the design are not quite correct. Further, large scientific projects usually involve enormous amounts of data that must be filtered carefully. For the design process to be efficient, artificial intelligence techniques such as machine learning may be used to automate the search for robust designs.

It is important to remember that the scientists' goal is to do a good experiment and that a wide margin for error in the experiment design, execution and analysis must be included. This desire for safety outweighs "cute" solutions that may seem attractive in the AI lab but have not demonstrated any utility in the real world. It is therefore important to identify which AI techniques will be most relevant to science experiments.

The wide variety of search algorithms used in AI systems are often novelties to scientists who rely, often un-

knowingly, exclusively on hill-climbing programs. The time saved and confidence gained by performing an automated search are significant. Even if the results of the search are not themselves conclusive they do provide a much more informed result in a shorter time than would have been otherwise possible. It is important to realize that a scientist does not need to have everything automated, just an improvement over what is currently done.

Further, representation schemes used in AI are much more sophisticated than typically used in scientific programming and may also be of use. Scientists used to working with purely numerical arrays will find AI's many data structures of use.

One of the ubiquitous problems in large scale scientific experiments is that of classification. This applies whether the domain is high energy physics, astronomy, or molecular biology. All three of these seemingly diverse areas have similar requirements. They are very large, often the largest experiments or projects in their respective disciplines. They take many people and many years to complete. They also rely heavily on computerized datalogging, databases and numerical simulations. However, the tools for designing and analyzing the experiments are often not very automated. These characteristics indicate that there is a need for automation in large scale experimentation.

High Energy Physics: An Example of "Big Science"

HEP is carried out at a few large facilities scattered over the world. A facility has one or more particle accelerators to generate very high energy particles. Each experimental apparatus, called a detector, weighing hundreds or thousands of tons is used to record the outgoing particles that result from a collision of particles produced by the accelerator. Both the accelerator and the detector are very complicated and take many scientists and technicians to operate. The success of the facility is entirely dependent on the ability of the accelerator to deliver beams of particles to the experiment and the ability of the experiment to measure the results of collisions. An accelerator facility can also be thought of as similar in some ways to a factory in that it produces a product (the beam or scientific results) in a complex setting much like a more familiar kind of factory that produces automobiles or electronic components. Both scientific and industrial factories are concerned with saving money and producing a quality product. Thus, some of the kinds of applications of AI used in HEP will find analogs in industrial settings as well.

As a specific example I now discuss how AI can help in both the accelerator operation and the experiment design by discussing implemented work. Of course it is also possible to design an accelerator and operate an experiment with AI as should be evident after reading the following discussion.

AI for Diagnosis

Once an accelerator has been designed and built there is still the formidable task of making it work properly. Since the cost of operating the accelerator is so large it is important that the diagnosis be done correctly as well as reasonably quickly. It should be pointed out that the commissioning phase of a new accelerator is typically a year so that some problems may last for weeks or even months before being solved.

Expert systems for diagnosis of accelerator problems in conjunction with numerical simulations of the accelerator can be used to quickly and efficiently localize a fault (Lee et al. 1987). Because of the deep knowledge required to diagnose a fault, only relatively easy problems can be handled at present. However, even this is highly useful because it frees experts to concentrate on the bigger problems and not have to worry so much about the smaller ones.

Using an AI-based diagnosis tool that uses physicists deep knowledge of particle beam transport replaces manual knob-twiddling done by non-technical operators. The deep knowledge is in part contained in precise mathematical models. The expert system functions as a smart driver of existing mathematical tools including numerical optimization programs. Typical errors that plague every accelerator are mis-calibrations or mis-alignments of the magnets that focus and steer the beam particles. Similar problems plague the beam monitors that measure various attributes of the beam such as where it is in relation to an ideal trajectory.

Machine Learning for Diagnostic Rules

In addition to expert systems where physicists' knowledge is explicitly coded into the system, it is possible to automate some of this knowledge acquisition with machine learning. For example, learning rules to distinguish between magnet and monitor errors generates a reproducible set of predictions whose efficiencies can be easily tested (Buchanan et al. 1988). Setting up a set of attributes for learning and then critiquing the results of learning is much less time consuming of the expert's time than traditional knowledge acquisition.

Machine Learning and Experiment Design

A HEP experiment is basically a counting experiment. This means that the data consists of collecting discrete "events" (collisions of one beam of particles with another beam). The events consist of new particles generated as a result of the collision. The problem is that the rate of uninteresting collisions is many orders of magnitude greater than the interesting event rate. The problem, then, is one of classification. This process is performed on-line by a "trigger system" that selects the interesting collisions and stores them on permanent media for later more detailed analysis.

Machine learning provides a possible means of automating the classification triggering task (Clearwater

& Stern 1991). As in the case of the accelerator, a complex numerical model is used. In this case domain knowledge is embodied in a model that generates simulated events according to the currently held physics model as well as a simulation of the response of the detector to the newly created particles. This simulates what the trigger system sees. The physicist supplies the attributes used in classification and the minimum selection efficiencies for the various processes.

The triggering process is one of the most crucial parts of the experiment and a great deal of human toil is spent checking which conditions will allow the experiment to be performed in a successful manner. The traditional method employed by physicists is to guess which attributes of an event will be useful for distinguishing between interesting and uninteresting events and then plot them on histograms or two dimensional plots and choose a value that provides the best separation. The advantage of a rule learning program is that it, in effect, automatically searches the enormous space of all plots (up to the dimension of the rule complexity supplied by the physicist) and only returns those that satisfy the desired separation criteria. The physicist's job has been reduced to supplying the list of event attributes and the separation criteria and thresholds. If the list of attributes is ordered by preference then the search can be performed even more efficiently. The physicist then can choose from among the rules, if any, that are found which are the best for the experiment.

Our experience with rule learning has been that the results have several uses, not all initially planned. First, the program does return rules that agree with the physicists' intuition and provides confidence in the program. Also, in cases where the program does not find any rules the physicist can confidently know that at least a certain space has been thoroughly tested. Finally, the program is also able to find rules that are not obvious to the physicist and require some thought to understand. This acts to enhance the credibility of the program even more.

Neural networks have also been used to classify event types as shown in (Denby and Linn 1990), (Hubert 1990), (Hubert 1990a), and (Lonblad et al. 1990). These systems work similarly to induction programs in that they use pre-classified examples from an experiment simulator and learn to discriminate patterns. Neural nets have the advantage that they can more easily use low level data than a knowledge-based rule system. This means they can be used for more pattern recognition type applications. On the other hand, neural nets have the disadvantage that their black box nature is an undesired uncertainty in the experiment.

Other applications of AI in large experiments can be found in (Nuclear Instruments and Methods 1990) and (Expert Systems for High Energy and Nuclear Physics 1990 and 1992). AI has also been used in small scale experimentation where the problems of data volume and

extreme cost are not present (Zytkow et al. 1990).

In summary, High Energy Physics facilities provide a fertile ground for developing AI techniques because of their need for high productivity and existing heavy reliance on computers in all phases of operation.

What Experimental Science can Contribute to AI

This section briefly discusses what type of contributions scientific methodology can make to AI. In experimental science it is often as important to know the sensitivity of the answer as it is to know the answer itself. In particular, up to now, machine learning has all but ignored this issue. Performance Sensitive Learning (PSL) (Clearwater and Lee 1992) is an attempt to add sensitivity information to the hypothesis formation in machine learning. In a rule-based learning system, this means requiring rules to be less sensitive to the training examples, while still maintaining a high level of coverage. This also has an advantage in incremental learning because the rules learned at one time will more likely still be good at a later time.

Experimental science can also be used as a paradigm to learn a more robust set of rules (Lee and Clearwater 1992). For example, many rule-learning systems have a requirement that each rule cover at least one unique example. While this leads to a concise concept description, it does not take into account the realities of an experiment. For instance, if a large set of the training examples are covered using a rule that relies on an attribute that corresponds to a part of the apparatus that fails, then all those examples will fail to be covered. Thus, in an experiment, uniqueness is not important, but redundancy is. Redundancy can be obtained by requiring that each training example be covered by a minimum number of rules. Learning can be made more efficient by biasing the search towards using certain attributes over others to insure that a proper set of rules will be found.

While the above two techniques do not explain sources of error they provide a means of either estimating them or creating an experiment that is less sensitive to them.

Discussion

Large scale experimental science involves elements of classification (of phenomena) as well as production (of data). Machine learning techniques with their suite of representations and search strategies are ideal candidates for automating classification tasks. Some examples of the efficacy of learning techniques have been given and the early results are promising. Expert systems can be used for the production type tasks that go with a long term experiment. These systems must be able to make recommendations that can reliably reduce the downtime of an experiment. Expert systems in sci-

ence carry the extra burden of having to perform in an environment that is so novel that the expertise is still developing.

A more intriguing use of AI in experimental science is in the area of discovery. Up to now, though, all discovery programs have been used to discover things that already known (Langley and Zytkow 1989). Discovery could take the form of a new phenomenon, or even the creation of a new descriptive feature of an existing phenomenon.

Because AI techniques are so new to large scale experimentation there needs to be a thorough "shakedown" so that scientists can gain confidence and become more familiar with them before they see more widespread use. However, we have seen that the interaction between experimental science and AI is a two-way street where both sides benefit from the methodologies of the other. Continued interaction can only add to the progress.

References

Buchanan, B. G., J. Sullivan, T. Cheng and S. Clearwater 1988. *Simulation Assisted Inductive Learning*. In Proceedings of the Seventh National Conference on Artificial Intelligence, 552-557. Menlo Park, CA.

Clearwater, S., Papcun, G., and Clark, D. 1986. Developing an Expert System to Control a Beam Line at the Los Alamos Meson Physics Facility. *Nuclear Instruments and Methods A247*: 193-196.

Clearwater, S., Cleland, W., Provost, F., Stern, E., and Zhang, Z. 1989. A Real-time Expert System for Trigger Logic Monitoring. *Nuclear Instruments and Methods A293*: 491-495.

Clearwater, S. H. and Stern, E. G. 1991. A rule-learning program in high energy physics event classification. *Computer Physics Communications* 67: 159-182.

Clearwater, S. and Lee, Y. 1992. Sensitivity Constraints in Learning. In Proceedings of the Second Workshop on Computational Learning Theory and Natural Learning Systems. In press.

Denby, B. and Linn, S. 1990. *Computer Physics Communications* 56: 293-.

Humbert, B. 1990. On the Use of Neural Networks in High-Energy Physics Experiments. *Computer Physics Communications* 56: 299-.

Humbert, B. 1990a. *Computer Physics Communications* 58: 223-.

Langley, P. and Zytkow, J. 1989. Data-Driven Approaches to Empirical Discovery. *Artificial Intelligence* 40: 283-312.

Lee, M. J., Clearwater, S. H., Kleban, S. D., and Selig, L. J. 1987. Error-finding and Error-correcting Methods for the Start-up of the SLC. In Proceedings of the 1987 Particle Accelerator Conference, 1334-1336.

Lee, Y. and Clearwater, S. 1992. Tools for Automating Experiment Design: A Machine Learning Approach.

In Proceedings of Tools for Artificial Intelligence 92. In press.

Lonnblad, L., Peterson, C. and Rognvaldsson, T. 1990. Finding Gluon Jet with a Neural Trigger. *Physical Review Letters* 65: 1321-1324.

Nuclear Instruments and Methods in Physics Research A293, 1990. (Entire issue is devoted to the International Conference on Accelerator and Large Experimental Physics Control Systems and includes numerous AI articles.)

Proceedings of the International Workshop on Software Engineering, Artificial Intelligence and Expert Systems for High Energy and Nuclear Physics, 1990 and 1992. (As the title indicates, includes many articles on AI related to HEP.)

Zytkow, J., Zhu, J., and Hussam, A. 1990. Automated Discovery in a Chemistry Laboratory. In Proceedings of the Eight National Conference on Artificial Intelligence, 889-894. Menlo Park, CA: AAI Press.