CONF-930749--41

LA-UR- 93-2537

Title: APPLICATION OF NEURAL NETWORK AND PATTERN RECOGNITION SOFTWARE TO THE AUTOMATED ANALYSIS OF CONTINUOUS NUCLEAR MONITORING OF ON-LOAD REACTORS

Author(s): Jo Ann Howell, George W. Eccleston, James K. Halbig, Shirley F. Klosterbuer, and Ted W. Larson

Submitted to: 34th Annual Meeting of the Institute of Nuclear Materials Management, Scottsdale, Arizona, July 18-21, 1993

Los Alamos
NATIONAL LABORATORY

APPLICATION OF NEURAL NETWORK AND PATTERN RECOGNITION SOFTWARE TO THE AUTOMATED ANALYSIS OF CONTINUOUS NUCLEAR MONITORING OF ON-LOAD REACTORS*

Jo Ann Howell, George W. Eccleston, James K. Halbig, and Shirley F. Klosterbuer
Los Alamos National Laboratory, Los Alamos, NM 87545 USA

Ted W. Larson
California Polytechnic State University, San Luis Obispo, CA 93401 USA

ABSTRACT

Automated analysis using pattern recognition and neural network software can help interpret data, call attention to potential anomalies, and improve safeguards effectiveness. Automated software analysis, based on pattern recognition and neural networks, was applied to data collected from a radiation core discharge monitor system located adjacent to an on-load reactor core. Unattended radiation sensors continuously collect data to monitor on-line refueling operations in the reactor. The huge volume of data collected from a number of radiation channels makes it difficult for a safeguards inspector to review it all, check for consistency among the measurement channels, and find anomalies. Pattern recognition and neural network software can analyze large volumes of data from continuous, unattended measurements, thereby improving and automating the detection of anomalies. We developed a prototype pattern recognition program that determines the reactor power level and identifies the times when fuel bundles are pushed through the core during on-line refueling. Netural network models were also developed to predict fuel bundle burnup to calculate the region on the on-load reactor face from which fuel bundles were discharged based on the radiation signals. In the preliminary data set, which was limited and consisted of four distinct burnup regions, the neural network model correctly predicted the burnup region with an accuracy of 92%.

INTRODUCTION

Nuclear power stations in the United States contain reactor cores, which can be accessed from only one end, usually the top; fuel can be accessed only when the reactor is shut down. One safeguards advantage to this type of reactor is that it is relatively easy for a nuclear safeguarding agency to monitor the fueling process: an inspector can be sent to the site to oversee the fueling procedure. On-load nuclear reactors differ from those in the United States, in that operators may remotely obtain access to the core from both ends, and the reactors can be continuously fueled without shutting them down. Such an operation offer a fuel management advantage, but a safeguards challenge, because it provides a greater opportunity for the diversion of nuclear material.

On-load reactors are well-suited for producing plutonium from their standard fuel bundles. Safeguarding an on-load reactor requires keeping track of fuel as it is pushed through the core. When a fresh fuel bundle is pushed in one side, a spent fuel bundle is simultaneously discharged into a collection mechanism on the other side. Using this fueling scheme, a typical on-load reactor will discharge 55 to 65 fuel bundles per week. Figure 1 shows a conceptual diagram of this fueling cycle. Because this is an ongoing process, it is labor intensive for a safeguarding agency to have an inspector on site to continuously monitor re-fueling.

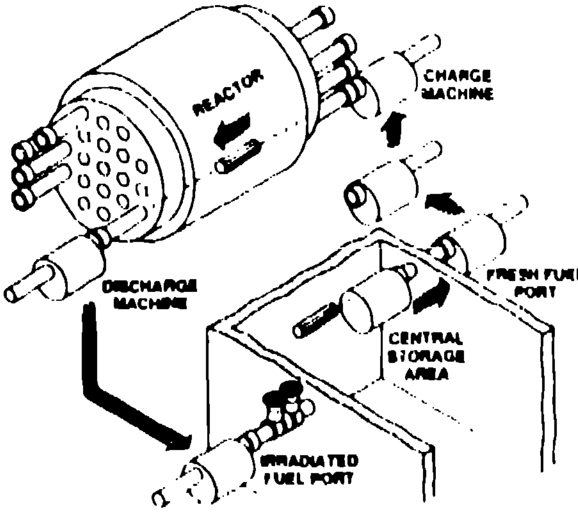

Fig. 1 Conceptual diagram of fueling cycle

* This work supported by the U. S. Department of Energy, Office of Safeguards and Security and Office of Arms Control and Nonproliferation

To provide data that are useful to inspectors, a core discharge monitor (CDM) system[1] has been installed on the on-load reactor. The CDM collects data continuously and automatically from radiation sensors that monitor the reactor core and the fueling process of the on-load reactor. Currently, the CDM data are manually examined by a safeguards inspector using graphical review software to determine when on-line fueling activity occurred. Because this system has the potential to generate massive quantities of data, efficient automatic algorithms would help make interpretations. These algorithms could extract information from the data, reduce analysis times, and relieve inspectors from time-consuming manual data reviews. Automated quantitative analysis programs could help safeguarding agencies gain a better perspective on the complete picture of the fueling activity of an on-load nuclear reactor. These programs could provide a cost-effective solution for automated monitoring of on-load reactors, significantly reducing personnel time and effort. In this paper we discuss prototype pattern recognition and neural network software developed to test automated data analysis and provide a tool for inspectors. The pattern recognition program was developed to test the feasibility of analyzing CDM data to identify when fuel bundle pushes occurred during on line refueling and to monitor the power level of the reactor. The neural network model was developed to test the feasibility of determining the region on the reactor face from which each fuel bundle set was discharged and to try to predict the burnup of fuel bundles. These programs were tested using preliminary start up data collected from a CDM system installed on an on load reactor.

CORE DISCHARGE MONITOR (CDM) SYSTEM

The CDM system used in this study consists of four gamma ray and neutron detectors (GRANDs) located near the nuclear core: two on each reactor face. The faces of the reactor core are on the east and west sides of the building. Fueling takes place from east to west or west to east and each GRAND detector array is designated by its location in relationship to the core, either the southeast (SE), northeast (NE), southwest (SW), or northwest (NW) corner as shown in Fig. 2. The GRAND operates continuously, collecting data at discrete time intervals from the detector arrays. These arrays monitor radiation signals from the reactor that show the discharge of spent fuel from the reactor core. The data are transmitted to an MS-DOS computer for permanent recording, archiving, and analysis by inspectors.

Each GRAND collects nuclear radiation data from the detector enclosure, filters it, time stamps it, and temporarily stores it. The data are then fed to the collection computer upon request for more permanent storage. At a later time, data can be off-loaded from the collection computer for off-line review. The detector data fed from the GRAND consist of five channels of information. The channels are labeled as follows: fission chamber A, fission chamber B, fission chamber C, ion chamber 1, and ion chamber 2. Fission chamber A corresponds to the first neutron detector in the detector enclosure. Fission chamber B is another view of the first neutron detector, which can be used for tamper detection. The second neutron detector in each detector enclosure is labeled as fission chamber C. This neutron detector is not wired to its corresponding GRAND, but rather to the GRAND on the opposing face. For example, the NE fission chamber C is wired into the NW GRAND, and the NW fission chamber C is wired into the NE GRAND. This provides the overall system with a backup, in case the GRAND for one of the detectors fails. This cross wiring is shown in Fig. 2 as the splice box between the two GRANDs on each side of the reactor core.

Finally, the two gamma ray detectors correspond to the ion chamber 1 and 2 channels, respectively. Figure 3 shows the layout of a detector enclosure. An in depth discussion of the detector assemblies and the GRAND electronics package can be found in Ref. 2

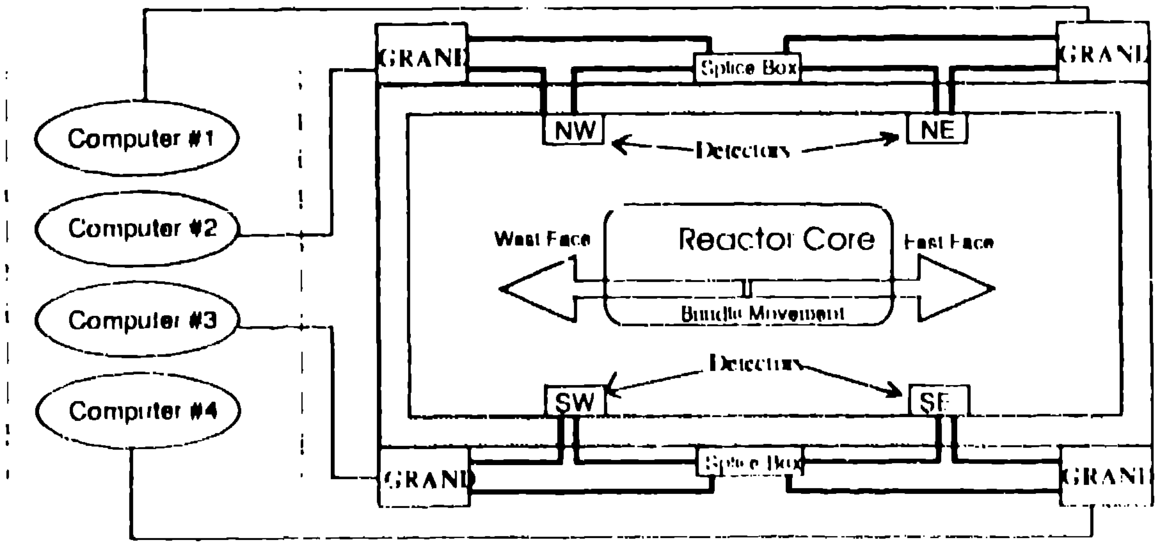

Fig. 2 Sample layout of a typical on load reactor

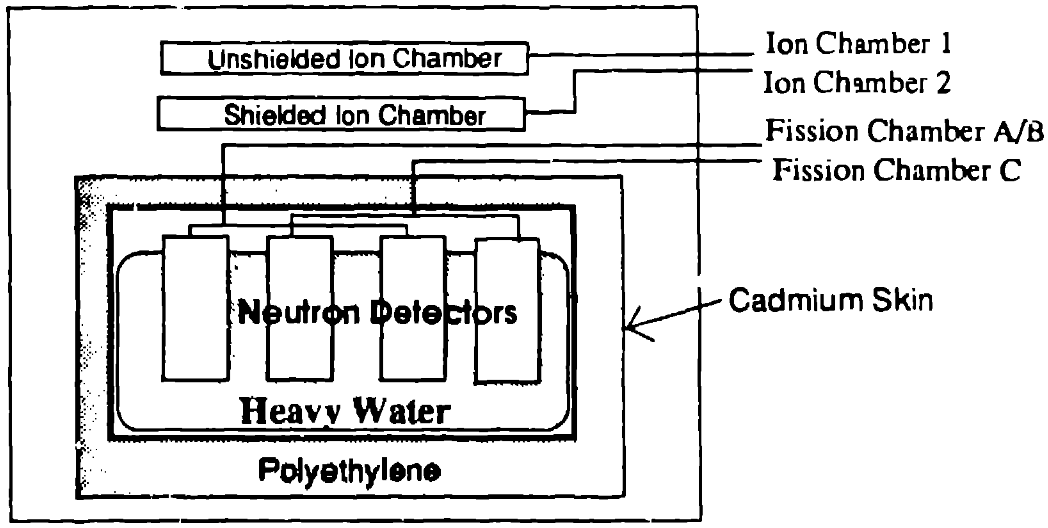

Fig. 3. A typical detector enclosure.

The GRAND records data every 10 to 11 seconds requiring around 100 megabytes to store all the data points collected from one reactor for 90 days; normally statistically insignificant data are filtered so the actual data amount stored is closer to 10 to 20 megabytes per 90 days. It is impractical for inspectors to quantitatively analyze this much data. Shown in Fig. 4 are graphs of data from two detectors during one particular day. Each large spike on the graph corresponds to a pair of fuel bundles being discharged from the reactor. Smaller spikes or decay curves or both on the graph may correspond to other activities such as the rotation of the fueling machine or the radioactive decay of the spent fuel being held in the fueling machine during a refueling operation. Reactor power level can also be determined from the data because the background level the detectors are sensing corresponds to the current power level of the reactor. The background in this context is considered to be the amount of radiation the reactor emits when no fuel is present outside of the core. A safeguards inspector counts the number of spikes on the graph to determine the total number of fuel pushes the reactor made in a particular day. The counted number of fuel pushes is then compared to facility declarations for safeguards verification. An automated process can considerably reduce the analysis time and help a safeguards inspector review the large volume of CDM data.

AUTOMATED SOFTWARE ANALYSIS

We developed prototype analysis software to investigate the feasibility of the following objectives:

1. Identifying sections in the CDM data for an inspector to examine in greater detail,
2. Locating and counting fuel bundle pushes and determining when they occurred,
3. Determining reactor power level as a percentage of full power,

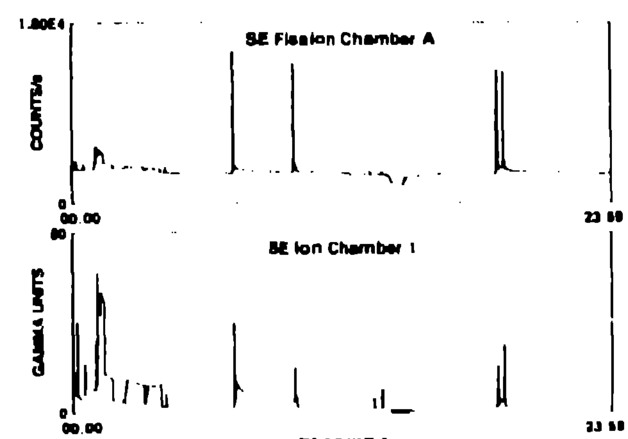

Fig. 4. Sample CDM data from an on-load reactor.

4. Correlating events between detector channels to assure the channels are operating correctly and to check for possible tampering,
5. Identifying the fueling channel from which the spent fuel was discharged, and
6. Predicting the burnup of discharged spent fuel bundles.

A prototype pattern recognition software tool, CDM Analysis, was developed to test objectives 1 through 4. A neural network model was developed to test the feasibility of predicting fuel burnup and location of fuel discharged from the reactor. To fully test CDM Analysis and the neural network models, a considerable amount of data is needed. For this study, only about [illegible] days of data were available. Although the total amount of data used was sparse, the analysis software still performed well suggesting this approach could be developed into a useful tool for inspectors.

CDM Analysis makes two passes over the CDM data during its search for areas of interest. In the first pass, it slides an average along the signal looking for significant changes. When the slope of the signal jumps above or below the sliding average by more than 10%, the data points are flagged for later examination. In the first pass, a large quantity of data may be flagged as interesting. To reduce the clutter, a second pass is made over just the areas that were flagged. Areas near each other in the time series are clustered together with the maximum data point being marked as the middle of the event. From the resulting list, a report can be generated to alert the safeguards investigator to specific areas of the data. Radiation spikes caused by refueling are found by setting the search threshold very high (50%). This technique provides all the fueling spikes for a given data set.

MONITORING REACTOR POWER LEVEL AND POWER LEVEL CHANGES

Once the areas of interest are identified, power level monitoring is straightforward. When no events are occurring, the background radiation sensed corresponds to the reactor power level. The average of the background can be used to compute the power level by establishing a baseline reading of what is considered to be full power. This baseline is computed by examining data from a reactor that is operating at a fixed power without fuel outside the core. The average value recorded by each detector is used as the baseline. This baseline is marked on the graph in Fig. 5 by a horizontal line. If the average value of the background moves from this baseline, then the power level is changing. The data have shown that most power changes occurred in a step wise fashion. CDM Analysis evaluates the power changes in the following manner. If the reactor power is raised or lowered, the slope of the average background starts to become very steep. This is marked as the beginning of a power change. When this slope flattens out again, the end of the power change is marked. The

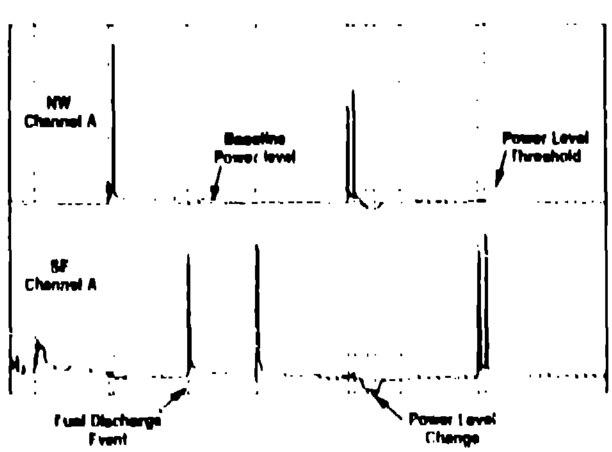

Fig. 5. Sample output from the CDM analysis program.

new value at which the average background comes to rest is considered the new power level of the reactor. The average background as a percentage of the pre-defined baseline is the percentage of full power at which the reactor is running.

Currently, CDM Analysis does not examine more than one channel on one detector when making its power level computations. In a production-quality analysis package, this percentage should be an average of all the percentages computed from all channels on all detectors. By taking power level measurements from all sides of the reactor core and averaging them, we could obtain a more accurate power level reading. Even though examining just one channel gives a fairly accurate reading, within 5%, examining all channels is a much better strategy because it provides a redundancy check. Figure 6 is an example of the power level of a reactor being raised from startup to full power. Notice that the power changes occur in multiple steps. CDM Analysis is also capable of printing a report that details each step of the power level change and the power level to which the reactor moved.

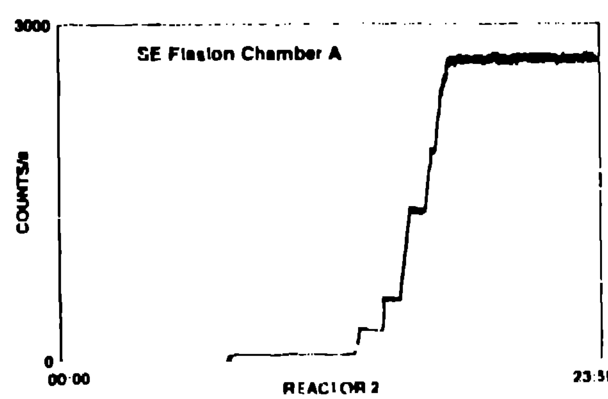

Fig. 6. The multiple steps of a power level change.

STATISTICAL PHENOMENA OF CDM DATA

We statistically analyzed the available CDM data to test the correlation of the height of the radiation spikes from fuel discharge events with the fuel burnup. We also cross correlated radiation signals from detectors in different positions to try to determine the location of the fuel channel during a refueling event.

Determining burnup is a difficult, multivariate problem. The CDM data showed that detectors on one face of the reactor are insignificantly affected by refueling events occurring on the opposite face. A significant correlation does exist between detector arrays located on the same face of the reactor. The variance in the ion chamber data was found to be pronounced. This effect was traced to data sampling with insufficient integration time to provide accurate ion channel currents. The neutron channels were not affected and provided stable readings that were used for the analysis.[1]

NEURAL NETWORKS FOR SOLVING THE FUEL GEOMETRY PROBLEM

Neural networks are based on a mathematical model that is derived from cell biology.[4] These networks are organized into layers consisting of several neurons (nodes) connected with adjustable weights. Each layer performs a particular function. The input layer processes the data being presented to the network, one or more hidden layers encode "features" in the data, and the output layer holds the response of the network to a given input.

Two phases of operation are required: the learning phase and the testing and recall phase. Learning consists of presenting a stimulus (an input vector) to the input layer together with a desired response. The network then calculates a result using the current weights and given input values. This "answer" is next compared with the desired response. If a difference of sufficient magnitude exists, the weight values are adjusted. Over time, as this learning process is repeated with more vectors, the weights will converge, and the network is said to be trained. During the testing/recall phase, similar examples are presented to the network to test whether the training was adequate. The difference between the desired and actual output is a measure of success, with differences of smaller magnitude representing greater success than those of larger magnitude.

When using neural networks, one must obtain an adequate set of training data. It is difficult to quantify the amount of training data required for good results because the quantity depends on the complexity of the records and the number of "features" embedded in the data. The 30 days of available reactor data yielded only 170 examples of fuel discharge events, which we consider minimal for adequate training and testing. In addition, these events came from only 90 of the 460 available fuel channels in the reactor core and represented a start-up activity rather than normal refueling. Even with these limitations, we were still able to train a neural network to classify the data into different regions on the face of the reactor.

The first neural network model divided the channel map into eight regions. This channel map and the eight regions are shown in Fig. 7. Almost all the regions were chosen because of the distribution of the points in the available data. Because detectors on one face do not reliably see events on the opposing face, only 10 channels from the same face out of the 20 total channels were used in the neural network model. The ion chambers act as noise during the training process to help separate the input vectors into appropriate categories. Back propagation was chosen as the modeling paradigm because of its ability to use real-valued inputs.[5] The neural networks used in this proof-of-principle were created using NeuralWorks Professional II/Plus,[6] a commercial neural network development tool manufactured by NeuralWare, Inc.

NEURAL NETWORKS FOR FUEL BURNUP PREDICTION

Because it may be important to determine if a facility is discharging low-burnup fuel from the reactor, we built a neural network model similar to the one described above to predict fuel burnup. It is difficult to compute an actual value for the burnup of each individual fuel bundle because the spike recorded by the CDM is an additive value of two bundles being discharged simultaneously. In this data set, burnup fell into one of four distinct regions. Therefore, we built a neural network to classify burnup into one of the four categories based upon the CDM data, as shown in Fig. 8.

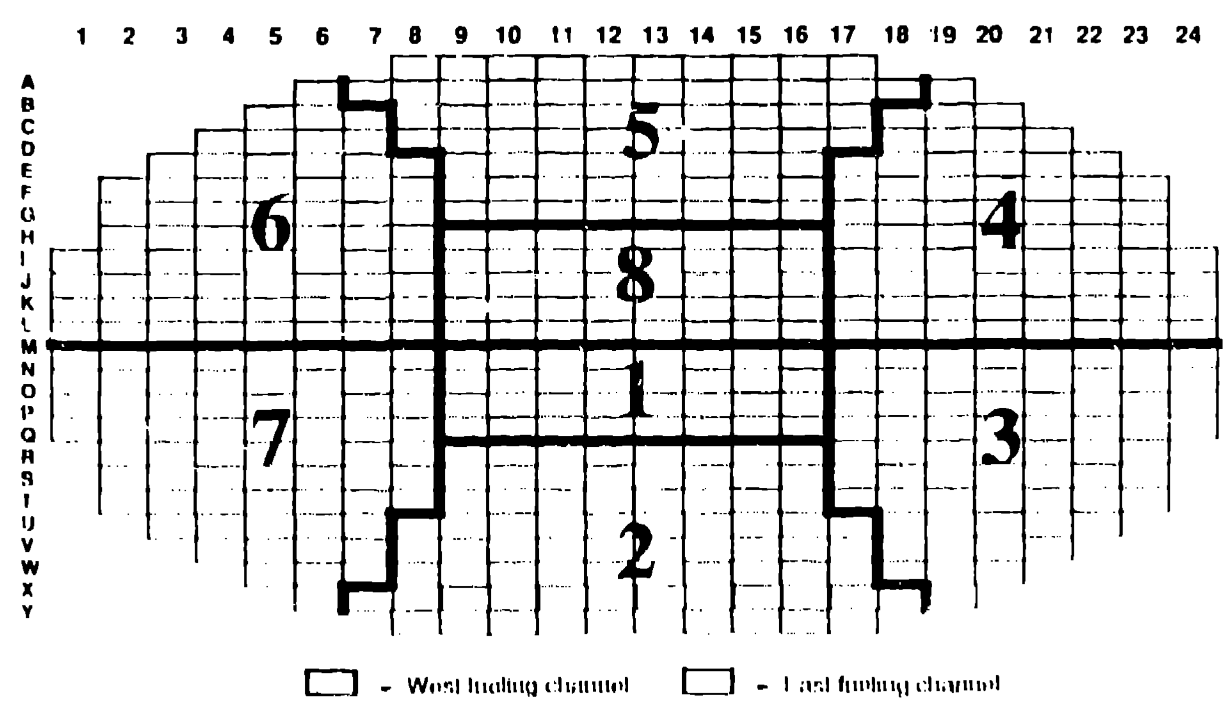

Fig. 7. Eight-region map of reactor face

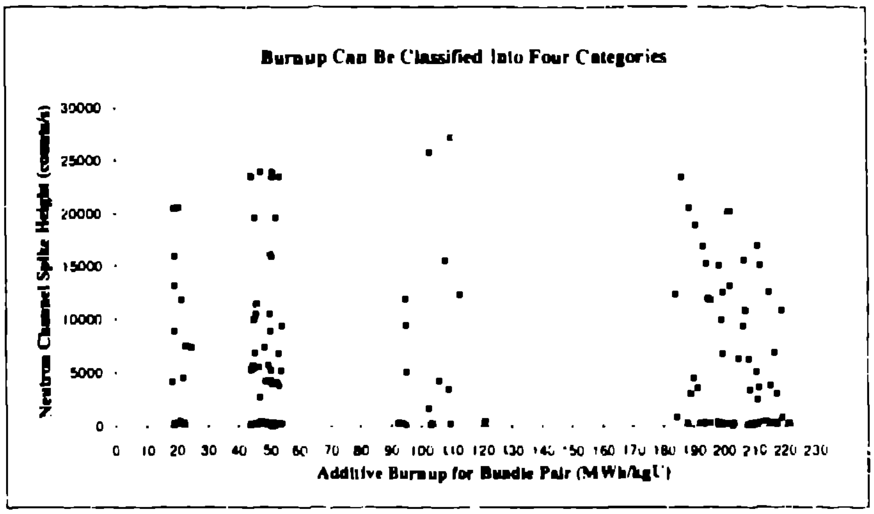

Fig. 8. Four categories of recorded burnups in sample data.

RESULTS OF NEURAL NETWORKS FOR SOLVING GEOMETRY AND BURNUP

The neural networks used for solving the geometry problem were trained and tested on data from the east face of the reactor, although the west face could have been used just as well. The training set consisted of 63 patterns and the test set of 72 patterns. After 50,000 training iterations, the network correctly classified the region of the fuel discharge in 82% of the patterns in the test set. For the fuel burnup problem, the network performed better, with an accuracy of 92% in predicting fuel bundle burnup. In spite of the very small data set, the networks performed remarkably well. The result was a neural network model of reactor geometry that correlates power level, burnup, and the number of fuel bundles pushed through the reactor.

CONCLUSIONS

The CDM Analysis tool has shown the potential for automated analysis of CDM data to determine refueling activity and to monitor the reactor power level. Neural network implementations for determining the location of fuel discharge and the burnup of fuel bundles appear successful enough to warrant further research. It appears that neural network models could be developed to provide close to 100% accuracy in predicting position and burnup if a complete set of representative data from an operating on-load reactor were available. The data needed to achieve this capability should include fuel pushes from all 160 channels of the reactor face and a complete cycle of fuel through all 13 positions in every channel.

Future work should include devising a more accurate technique for determining areas of interest in the CDM data, rather than a sliding average. Power level monitoring using an average over all 20 channels will also yield a more accurate power level calculation. Deficiencies in the collection of quantitative data should be corrected. We need more samples of data per unit time and a gamma channel reading more representative of the measurement period. In addition, different types of neural network models should be tried once a representative amount of data has been obtained. The portability of neural network models to other reactors of the same type should also be investigated. Neural network models hold great promise for future work in the area of core discharge monitoring and automated examination of large volumes of continuously collected data to improve nuclear safeguards. We firmly believe that a commercial-grade tool for monitoring power and counting fuel bundles from CDM data should be developed.

REFERENCES

1. J. K. Halbig and A. C. Monticone, "Proof of Principle Measurements for an NDA Based Core Discharge Monitor," *Nucl. Mater. Manage.* **XIX** (Proc. Issue) 847-852 (1990).

2. J. K. Halbig, A. C. Monticone, L. Ksiezak, and V. Smilaneks, "The Design and Installation of a Core Discharge Monitor for CANDU-type Reactors," *Nucl. Mater. Manage.* **XIX** (Proc. Issue) 839-846 (1990).

3. Ted W. Larson, James K. Halbig, Jo Ann Howell, George W. Eccleston, and Shirley F. Klosterbuer, "Automated Software Analysis of Nuclear Core Discharge Data," Los Alamos National Laboratory report LA-12516-MS (March 1993).

4. David E. Rumelhart and James L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (The MIT Press, Cambridge, Massachusetts, 1986).

5. Philip D. Wasserman, *Neural Computing Theory and Practice* (Van Nostrand Reinhold, New York, 1991).

6. NeuralWorks Professional II/Plus Reference Guide (NeuralWare, Inc., Pittsburgh, Pennsylvania, 1991).