J. ENVIRONMENTAL SYSTEMS, Vol. 28(2) 157-174, 2000-2001

APPLICATION OF ARTIFICIAL NEURAL NETWORKS TO UPDATE AND MODIFY A DOS-BASED ENVIRONMENTAL EXPERT SYSTEM

MATTHEW R. THOMPSON

WILLIAM F. McTERNAN

Oklahoma State University, Stillwater

ABSTRACT

An Artificial Neural Network (ANN) model was developed to mimic the exact output from a DOS-based Environmental Expert System. Computer codes developed originally for mainframe computers and transported into the DOS environment routinely do not receive modifications necessary to perform under more modern operating systems unless there is sufficient financial incentive. Software written for the environmental market, particularly the classroom market, rarely has this level of incentive, resulting in much previously usable software being rendered obsolete. Much of this software still can play a critical role in the education of future environmental scientists and engineers. The subject research investigated one potential solution to this problem: the development of ANN models capable of producing the exact results of the earlier DOS code while having the capability of ready modification given new information or circumstances. This research illustrated the overall utility of ANN's in this capacity, as a 100 percent compatibility between the underlying Expert System and the ANN was achieved. In addition, the ANN was readily modified to include new information. The ANN developed extends the useful life of the Expert System with minimal developmental costs, without extensive re-programming or retrofitting of the original code.

INTRODUCTION

Computer software written for environmental applications is typically produced in limited quantities for a relatively small market, particularly when compared to that developed for either business or gaming purposes. Environmental codes are often written by governmental or academic researchers and in many cases can trace their lineages to software originally developed for large mainframe systems. Many of these programs having previously made the transition, albeit with some operational difficulties, to DOS-based personal computers are increasingly called upon to operate under other operating systems, particularly Microsoft's Windows.

Another problem faced by the user of some of these codes is that changes either in underlying environmental regulations or in cost or operations databases often render these programs obsolete even when they can still operate efficiently on modern office computer systems. Being a small market with minimal economic incentives, primarily supported by programmers more interested in developing new models for research papers than in updating existing codes, environmental software is often "orphaned" in terms of ongoing maintenance or in supplying critical updates. It is often too time and resource consuming to update models that are only used by hundreds or possibly a few thousand individuals. Greater rewards will come to those whose programs have more universal appeal. This is one of many reasons why computer games (and the machines used to play them) have better graphics, higher speeds, and better audio than many corresponding environmental modeling packages.

Academics, in particular, are most impacted by these trends. College classes use more programs at perhaps lessened depth of application than do engineering firms. Students benefit from testing alternative scenarios within the simulation environment. These simulations allow for investigation of real-world problems, adding greatly to the education of young scientists and engineers. Programs developed either for the mainframe or for early DOS machines, however, are either increasingly off-limits to these users or else perform at reduced levels with the advent of alternative operating systems. Admittedly, programs developed for these new operating systems are generally far better than the original DOS codes but, given the limited market for this type of software, it often becomes more a matter of losing important codes rather than replacing them.

This article presents one alternative available to academics and working scientists and engineers who, because of tight time and resource budgets, are increasingly forced to either employ alternative methods to those previously used or to under-utilize a valued software package. Specifically, an Artificial Neural Network (ANN) was used to test the hypothesis that existing DOS-based systems could be easily and relatively rapidly updated to perform under newer operating systems. The expert system Cost of Remedial Action (CORA) was selected as the DOS-based program for investigation [1].

Expert systems are computer coded decision trees which focus upon developing solutions to various problems. Users respond to a series of questions to generate recommended courses of action including diagnosis, monitoring, planning, design, and interpretation [2, 3]. Routinely used for medical and chemical interpretation and diagnosis, other applications include hazardous waste management, groundwater remediation, and risk assessments [4-7].

USING CORA AS THE MODELED EXPERT SYSTEM

CORA was developed and maintained for the Environmental Protection Agency until the early 1990s. At the time of its last update in 1990 it employed question and answer format to suggest potential remediation alternatives. The questions ranged from what type of chemicals and soil types were present to information concerning the hydrology of the site. The program interpreted each answer and recommended several applicable remediation schemes. Subsequently, the total amount of contaminated materials and the recommended level of clean-up for that site were requested from the user resulting in an approximate cost of the remediation and clean-up.

Since 1990, CORA has become more difficult to operate under ever more rigorous computer systems as well as being somewhat dated with regard to the costs of the various technologies. CORA, as with any code, could be updated by rewriting and revising the original computer program but, as an alternative, this research developed and applied an artificial neural network (ANN) model to project needed technologies and their attendant costs.

The advantage of updating CORA by using a neural network is that the time needed to update the system was relatively minimal. Access to the source code or even extensive knowledge of CORA's programming language were also not required. Updated information was readily assimilated. Once the modeling was completed, any updates typically took less than five minutes to complete. In addition, little computer information was needed. The publicly available package used in this effort required only a basic knowledge of computer spreadsheets augmented with the neural network "add-on." In this manner, a tool was developed quickly and efficiently which built upon previous engineering knowledge with readily available skills and software. Neural networks have the capacity of learning the complicated decision tree associated with CORA or other expert systems while being able to update the underlying database and decision algorithm.

SCOPE OF THE PROJECT

To facilitate investigation, this project was confined only to sites that contained Volatile Organic Compounds (VOCs) as the contaminant of concern. Although CORA has the capability of giving remediation alternatives for many other categories of chemicals that might exist at a site, it was felt that an initial effort involving a ubiquitous contaminant could readily test the utility of the ANN model development. In this way a problem of sufficient complexity could be evaluated while still proving tractable should difficulties arise.

The ANN model was designed with the same questions as the CORA program for the types of conditions where VOCs may need to be addressed. The ultimate goal of the research project was to develop a tool that updated a very powerful and useful expert system. Further, the effort needed to accomplish this updating should be considerably less than that needed to recode the original program. The neural network program that was used in this research project was Neuralyst [5].

The ultimate goal of the training portion of this project was to achieve a 100 percent recognition rate in an efficient manner. The 100 percent recognition rate occurred when the training of the network subsequently lead to a 100 percent prediction of the test cases. Converting the test cases into training cases and using CORA to obtain additional information for other remediation/containment scenarios eventually achieved the 100 percent success rate. When this criterion was met, the network was ready for almost any type of containment or treatment scenarios dealing with VOCs previously addressed by CORA. The ANN was further modified by development of updated cost estimation methods. Current unit costs replaced the 1990 database.

NEURAL NETWORKS BACKGROUND

ANNs work by pattern recognition linking input data to outputs in a series of training and subsequently testing exercises somewhat analogous to calibration and verification steps of conventional modeling. These exercises are cited as examples of system "learning." Functional forms are developed during training which link the measured inputs and outputs of a problem situation. Testing confirms the model developed. The resulting model can then be used in prediction of different conditions. The capacity of the system to "learn" can then be subsequently utilized whenever technology or costs change.

On the simplest level, a neural network is configured like a human brain with many simple elements (neurons) that work in parallel. The neurons that make-up each layer are connected to each other to create a network of neurons much like the human brain. Input information is processed by a complex array of internal neurons to produce output information. That is, previously measured information is propagated from the input layer used to define the problem through the hidden layers to the output layer which when property configured describes the desired results.

Training consists of manipulating the number of inputs, the number of hidden layers, the functional form applied and the learning rate among other variables to gain conformance with the previously measured outputs. Typically, ANN models have two to six hidden layers which process the inputs while predicting the outputs. Input neurons are what is to be solved while output neurons describe the desired results. The higher the number of hidden layers the more generalizations can be made in the network. The larger number of layers also allows the program to use fewer numbers of neurons in the development of the neural network.

Networks with a lesser number of layers conversely use less computer time. It is reported that most neural networks can be solved with three layer systems [5].

A neural network can solve both linear and non-linear problems. As with any "learning" experience, the network is initially very prone to mistakes. Once calibrated, the neural network can become very precise. The network is ultimately driven by the type of problem to be solved. The neural network begins as a collection of rules and inputs that must be taught to achieve better results. This learning process is termed training.

The neural network has many advantages over conventional models which have been used in the past. Traditional models develop formulae that mimic reality. The data for a particular situation are used to fit the model. A neural network develops functional relationships between the input and outputs. Neural networks can be adapted to the most complex problems where other models may prove limited. The training process teaches the program the importance of every neuron instead of focusing on individual locations such as a maximal or minima. Each neuron is assigned a significance (weight) and is identified by the classification of the connection. This allows the program to establish useful relationships between the neurons in the network. The network can easily be as complex as the original problem without concern for the underlying mathematics. Artificial Neural Network models have been used in groundwater studies as well as in finance and other business applications [8, 9].

NEURAL NETWORK PARAMETERS

There are several parameters that are used by a neural network during training. They include the learning rate and momentum as well as the training and testing tolerances. These parameters are determined iteratively with relative improvement in recognition rate being the evaluation criterion.

The learning rate is used to control the way in which the error corrects the weights in the network for each training case. This error correction is the way the network trains itself. The range of the learning rate is from 0 to 1,where the lower number will reduce unstable behavior. Unstable behavior is when the neural network plateaus during training (number of correct answers does not increase as time progresses). The lower the learning rate the longer the network will take to develop, however.

Momentum deals with the amount of previous error that is applied to the weight adjustment in each training case. If momentum is 0.5, then the weight adjustment will be 50 percent from the current error and 50 percent of the adjustment will be applied to the previous case within the neuron. The neuron then takes the starting value given by the user and, using an exponential decay, reduces the amount of error associated with the next neuron by 5 percent each time.

Training tolerance tells the neural network how much training is needed to consider it trained. This number reflects how precisely the network must come to the desired answer. Tighter training tolerances force the program to numerically approach the desired targets. For instance, if the training tolerance is set at 0.2, then the computer will have to come within $\pm$ 20 percent of the target answers to be considered trained. Testing tolerance is similar to training tolerance when used on the previously segregated testing cases. This tolerance describes the acceptable proximity of the trained network to the target answers to be considered a correct response.

MATERIALS AND METHODS

The research was initiated by accessing CORA for a variety of containment and/or remediation scenario, as illustrated in Table 1. In response to a series of questions posed to the user, CORA produced a suggested alternative and an associated cost estimate for each of these simulations. These data were then used to train the neural network. With some exceptions, the questions CORA asked were in a true-false format. CORA gave different remediation schemes depending on the answers to these questions.

The ANN was established by incorporating the types of questions and answers included in CORA into the training network. Table 2 presents all of the possible questions necessary for a containment evaluation while Table 3 presents the corresponding information from the remediation effort. The ANN input data were collected by selecting true or false responses for these questions. Given the structure of a specific problem, not all questions are asked for during each simulation. For instance, if question 6 from Table 2 was answered "false" and question 7 was answered "true," then questions 8 thru 10 would not be asked. The next question would be 11 (is the site considered a hazard to unauthorized personnel). There were several other instances that involved these if-then questions in CORA and in the subsequent ANN. Similar conditions applied to Table 3 questions. Answering Q1, Q2, and Q4 was required to define the scope of a project [answers: Q1 → Homogeneous Contaminated Unsaturated soils (HCUS) , Q2 → Containment, Q3 → Volatile Organic Carbon].

The input data from CORA for the neural network were initially collected in ten arbitrary training cases completed with five test cases. The 15 cases resulted from the questions and answers associated with CORA data runs. After the first set of data was collected from CORA, the neural network was established. A total of 154 computer simulations were completed for the containment and remediation alternatives. Fifty-four were used for containment with 100 utilized for remediation. Table 4 presents the number of cases and how each case was used in these simulations. This table presents the number of training and test cases that were used for every simulation where a case was defined as an example of the output from CORA used to either train or test the neural network. A run represented a training session of the neural network that attempted to simulate the

Table 1. Containment and Treatment Options Evaluated

Remediation scheme CORA #	Name	Containment	Treatment
105	Surface water diversion/collection	X	X
201	Soil excavation	X	X
301	Onsite incineration		X
302	Offsite incineration		X
305	Soil vapor extraction		X
306	Flaring		X
307	Air stripping		X
308	Vapor phase carbon		X
312	Ion exchange		X
316	Solidification		X
317	In-situ stabilization		X
401	Offsite RCRA landfill	X	X
402	Onsite RCRA landfill—above grade	X	X
403	Offsite RCRA landfill—below grade	X	X
404	Offsite solid waste landfill	X	X
405	Discharge to POTW	X	X
406	Discharge to surface water	X	X
407	Water reinjection		X
503	Groundwater monitoring	X	X
504	Site access restrictions	X	X

Note: "X" denotes when a particular remediation scheme could have been employed by CORA.

Table 2. CORA Containment Questions[a]

What waste types apply to the site (Q1)
What response action do you wish to consider (Q2)
What types of contaminants are in the soil (Q4)
Will excavation of the contaminants cause environmental or public impacts (Q5)
Is the contaminated soil a hazardous waste (Q6)
Is the contaminated soil concentration above land disposal restrictions (Q7)
Is an onsite landfill reasonable (Q8)
Select all types of contaminants in leachate from landfill (Q8-a)
Are contaminated soils located in a 100-year flood plain (Q8-b)
Is a shallow aquifer present that would not allow a below grade landfill (Q9)
Type of discharge option either water reinjection (Q10)—water infiltration (Q10-a)—discharge to POTW (Q-10b)—discharge to surface water (Q10-c)
Could site conditions threaten health or safety of unauthorized visitors (Q11)
Are exposed soils on the site exposed to erosion (Q12)
Pick the location of the site: above floodplain (Q12-a)—at base of hill above floodplain (Q12-b)—in floodplain (Q12-c)

[a]The questions are not sequentially numbered due to other questions that resulted in the treatment section of CORA.

CORA outputs. The test data had not previously been placed into the neural network as a training aid. If the test case had been used to train, the neural network would already know the answer from previous training.

ANN Updating of CORA's Technical and Cost Bases

As new remediation technologies become available, the ANN produced in this effort can be updated by including the developing information into the network as a new training data set. The original network has the capacity to "learn" the correct responses when modified with new input data sets. In this way, the ANN developed can be more readily adapted to changing conditions than can a conventionally coded model.

The ANN completed for VOC containment and treatment options were capable of producing CORA compatible answers. Updating the costs portion of the expert system proved necessary. Several alternative approaches were evaluated before selecting an approach based on unit costs. This allowed every site to be economically evaluated by eventually knowing the unit cost from the neural network and the volume of contaminated material from the engineer's studies. The costs of the remediation schemes came from several sources. They include the Environmental Protection Agency [10], Ground-Water Remediation and Analysis

Table 3. CORA Treatment Questions[a]

What waste types apply to the site (Q1)
What response action do you wish to consider (Q2)
What is the hydraulic conductivity of the soil (Q3)
What types of contaminants are in the soil (Q4)
Will excavation of the contaminants cause environmental or public health impacts (Q5)
Is onsite incineration precluded based on space or local considerations (Q5-a)
Type of discharge option either water reinjection (Q10)—water infiltration (Q10-a)—discharge to POTW (Q10-b)—discharge to surface water (Q10-c)
What is the hydraulic conductivity of saturated zone (Q10-d)
Is the water table greater than 5 feet below surface (Q10-e)
Is the ash a hazardous waste (Q10-f)
Is concentration of contaminant in ash above land disposal (Q10-g)
Is an onsite RCRA landfill for solidified ash reasonable (Q10-h)
What types of contamination are in the leachate from landfill (Q10-i)
Would a shallow aquifer preclude a below grade landfill for solidified ash (Q10-j)
Is the solidified material landfill footprint in a 100-year floodplain (Q10-k)
Could site conditions threaten health or safety of unauthorized visitors (Q11)
Are exposed soils on the site exposed to erosion (Q12)
Pick the location of the site above floodplain (Q12-a)—at base of hill above floodplain (Q12-b)—in floodplain (Q12-c)

[a]The questions are not sequentially numbered due to other questions that resulted in the treatment section of CORA.

Table 4. Training and Test Cases from the Neural Network

Run #'s	New training cases	Total training cases	Test cases
1, 2, 3, 4, 5, 6, 7, & 8	10	10	5
9, 10, 11, & 12	9	24	5
13, 14, 15, & 16	10	39	5
17	10	64	5
18	10	79	5
19	10	94	5
20	10	109	5
21	10	124	5
22	10	139	15

Center [11], environmental design engineers, city officials, and landfill designers. Table 5 presents these cost figures.

RESULTS

Network Training

Twenty-two different runs were required by the neural network to emulate CORA's responses. These runs are detailed in Table 6. A lesser number would have sufficed if the project goals were less stringent. That is, this effort was constrained to achieve a 100 percent recognition rate between the Expert System and the subsequent ANN for VOC containment and/or treatment. The containment option required 16 training runs, while the treatment portion utilized six. The containment scenarios needed more runs because they were used to develop the number of neural network layers and the network parameters, which were then utilized by the treatment simulations. The number of runs and data samples was determined iteratively. If the first trial of the data achieved the desired goal, then the number of data samples would be sufficient for that particular case. In both cases, ten additional cases were included to further test the precision of the network model.

Table 6 displays the run number, number of samples, number of layers, neurons, parameters, and results for each of the 22 runs. These data show that a training rate of over 90 percent could be achieved readily; requiring only one run with two layers and 32 and 21 neurons per layer, respectively. Further, this single run relied extensively upon default levels included in the ANN code for most of the operating parameters. Should the scientist or engineer find 93 percent satisfactory, ANN models can be very readily developed for this problem.

Training to achieve a 100 percent recognition level required that the training tolerance be varied. While all of the final network parameters were set to code recommended default values, the training tolerance was set to 0.05. This forced the neural network to be within 0.05 of the actual number during the training procedure. This tighter control during training forced the training numbers closer to the target values. The other network variables were not modified because upon further review the default values were in the appropriate range to achieve the best results in a short period of time.

The final networks became very large so that three and four layer systems could not handle the number of variables and information in a routine manner. The run time for three training sessions was stopped at two hours without 100 percent training rate. This was arbitrarily deemed an unacceptable time limit. The 2-, 3-, and 4-layer systems were eliminated after the completion of containment section. The desired goal of 100 percent was achieved with the five-layer system in the containment section. It was then used for the training of the treatment model. If a smaller (or larger) layer system is ever needed, the ANN can be readily altered and the network reloaded instantly.

Table 5. Updated Costs for CORA

CORA #	Name	Unit costs
105	Surface water diversion/collection	$10,963/Ac[a]
201	Soil excavation	Approximately $2-$5/yd^3
301	Onsite incineration	$164–$730/ton
302	Offsite incineration	$200–$1,000/ton
305	Soil vapor extraction	$100/ton
306	Flaring	$300/hole for pipes
307	Air stripping	Depends on electricity costs—requires 1.5 hp/foot of stripping
308	Vapor phase carbon	$1,000-$40,000 for the machine carbon = $2–$3/lb
312	Ion exchange	$0.30–$0.80/1000 gal treated
316	Solidification	$100/ton including excavation
317	In-situ stabilization	Shallow—$40–$60/yd^3 Deep—$150–$250/yd^3
401	Offsite RCRA landfill	$15/yard—excludes transportation
402	Onsite RCRA landfill—above grade	$500–$1140/cy range from 7000–220000cy[a]
403	Offsite RCRA landfill—below grade	$490–$1121/cy range from 7000–220000cy[a]
404	Offsite solid waste landfill	$4.00/cy plus $1.50/ton
405	Discharge to POTW	$5.25/gal first 1000 gal, $2.00/gal after
406	Discharge to surface water	NPDES permit = $7,000[b]
407	Water reinjection	$1.00/gal haz, $0.55/gal non haz—excluding pump truck
503	Groundwater monitoring	$2,000/well/month plus quarterly monitoring
504	Site access restrictions	$28.50/ft includes fencing and signs[a]

[a]These costs came from CORA and were updated from 1990 dollars to current dollars with a factor of inflation of 3 percent. [b]This price depends on the city that issues the NPDES permit (this price is for Sand Springs, Oklahoma).

Table 6. Results of the Training Runs

Run number	Number of samples	Number of layers	Neurons per layer	Network parameters[a]	Results (correct)	Comments
1	10 train 5 test	2	32,21	Default values	93%	Need more data
2	10 train 5 test	2	32,21	Training Tol. = 0.05 Testing Tol. = 0.15	88%	Need more data
3	10 train 5 test	3	32,30,21	Default values	94%	Need more data
4	10 train 5 test	3	32,30,21	Training Tol. = 0.05 Testing Tol. = 0.15	90%	Need more data
5	10 train 5 test	4	32,30,45,21	Default values	93%	Need more data
6	10 train 5 test	4	32,30,45,21	Training Tol. = 0.05 Testing Tol. = 0.15	89%	Need more data
7	10 train 5 test	5	32,30,45,30,21	Default values	91%	Need more data
8	10 train 5 test	5	32,30,45,30,21	Training Tol. = 0.05 Testing Tol. = 0.15	90%	Need more data
9	24 train 5 test	2	32,21	Training Tol. = 0.05 Testing Tol. = 0.15	96%	Need more data
10	24 train 5 test	3	32,30,21	Training Tol. = 0.05 Testing Tol. = 0.15	93%	Need more data
11	24 train 5 test	4	32,30,45,21	Training Tol. = 0.05 Testing Tol. = 0.15	92%	Need more data
12	24 train 5 test	5	32,30,45,30,21	Training Tol. = 0.05 Testing Tol. = 0.15	95%	Need more data
13	39 train 5 test	2	32,21	Training Tol. = 0.05 Testing Tol. = 0.15	—	Would not train eliminated 2 layer networks
14	39 train 5 test	3	32,30,21	Training Tol. = 0.05 Testing Tol. = 0.15	96%	Need more data
15	39 train 5 test	4	32,30,45,21	Training Tol. = 0.05 Testing Tol. = 0.15	99%	Need more data
16	39 train 5 test	5	32,30,45,30,21	Training Tol. = 0.05 Testing Tol. = 0.15	100%	100% training for containment—10 more cases to test

Table 6. (Cont'd.)

Run number	Number of samples	Number of layers	Neurons per layer	Network parameters[a]	Results (correct)	Comments
17	64 train 5 test	5	32,30,45,30,21	Training Tol. = 0.05	59%	Need more data
18	79 train 5 test	5	32,30,45,30,21	Training Tol. = 0.05	90%	Need more data
19	94 train 5 test	5	32,30,45,30,21	Training Tol. = 0.05	94%	Need more data
20	109 train 5 test	5	32,30,45,30,21	Training Tol. = 0.05	90%	Need more data
21	124 train 5 test	5	32,30,45,30,21	Training Tol. = 0.05	96%	Need more data
22	139 train 5 test	5	32,30,45,30,21	Training Tol. = 0.05	100%	100% training for treatment—10 more cases to test

[a]Default network parameters were as follows: Learning Rate = 1, Momentum = 0.9, Input Noise = 0, Testing Tolerance = 0.3, Training Tolerance = 0.1, and 1 Epochs per update These values were used unless otherwise stated. In this case, the default values were used except for the changes noted.

Example Comparisons

The following example presents a comparison of the output from the neural network model with those generated by CORA. This hypothetical site is of the type commonly encountered. The example was further divided into five cases which, although similar, had subtle differences to test the resolution of the ANN model to evaluate data and suggest alternative renovation approaches. Two cases (1 and 2) involved a containment scenario, while the remaining three cases (3-5) utilized treatment.

The following information was used as background for the site for all five of the test cases:

- A homogeneous contaminated unsaturated soil;
- Non-hazardous volatile organic carbons (VOCs) as the chemical of concern;
- Clay soil;
- A site where the physical nature was dangerous to trespassers;
- Exposed soils that may erode;
- A site that can be excavated;
- A small site that restricts any type of onsite landfill;

- A site that is located above the flood plain; and
- A site that is located near surface water.

The information that varied among the cases:

- The type of remediation scheme employed (containment/treatment);
- Concentration of the soil (above/below land disposal restrictions);
- Incineration (yes/no); and
- Concentration of ash (above/below land disposal restrictions).

Case 1 looked at a containment scheme where the concentrations of the soil were not above land disposal restrictions while Case 2 investigated the same scenario except that the contaminant concentrations in the soils were above these restrictions. The first two treatment schemes, Cases 3 and 4, compared on-site incineration as a treatment alternative. Case 3 did not use this option while Case 4 produced an ash that was not above the landfill disposal restrictions. This compared to the similarly configured Case 5 where the ash was above the land disposal restrictions.

Questions and *answers* employed by CORA as inputs to the neural network include:

Case 1

Q1: What waste types apply to the site? *Homogeneous contaminated unsaturated soils*
Q2: What response action do you wish to consider? *Containment*
Q4: What types of contaminants are in the soil? *VOC's*
Q5: Will excavation of the contaminants not cause environmental or public health impacts? *True*
Q6: Is the contaminated soil a hazardous substance? *False*
Q7: Is the contaminated soil concentration above land disposal restrictions? *False*
Q11: Could site conditions threaten health or safety of unauthorized visitors? *True*
Q12: Are exposed soils on the site exposed to erosion? *True*
Q12-a.b.c: Pick the location of the site: *Above Floodplain (Q12-a)*

The actual questions instead of question numbers are included in Table 2.

The independent CORA and neural network simulations produced the same outputs. Recommendations from both models included:

- site restrictions;
- groundwater monitoring;
- surface water diversion;
- soil excavation; and
- offsite solid waste landfill.

Case 2

Featured the same questions and with the exception of Q7, the same *answers* as Case 1.

The results again showed complete agreement between the two modeling approaches. Further, Case 2 results were the same as Case 1 with the exception of the exclusion of an offsite solid waste landfill. The contaminant concentrations in this case exceeded the landfill restrictions.

Case 3

Questions that were asked by CORA (and *answers* generated) for this treatment case were:

Q1: What waste types apply to the site? *Homogeneous contaminated unsaturated soils*
Q2: What response action do you wish to consider? *Treatment*
Q3: What is the hydraulic conductivity of the soil? *Clay*
Q4: What types of contaminants are in the soil? *VOC's*
Q5: Will excavation of the contaminants not cause environmental or public health impacts? *True*
Q5: Is on-site incineration option precluded based on space or local considerations? *True*
Q11: Could site conditions threaten health or safety of unauthorized visitors? *True*
Q12: Are exposed soils on the site exposed to erosion? *True*
Q12-a.b.c.: Pick the location of the site: *Above Floodplain (Q12-a)*
The actual questions instead of question numbers are located in Table 3.

The results from Case 3 included:

- site access restrictions,
- groundwater monitoring,
- surface water diversion,
- soil excavation, and
- offsite incineration.

As before, the ANN results were in complete agreement with CORA. Differences were noted however, between the Case 1 and 2 containment evaluations by the addition of offsite incineration.

Case 4

This case addressed a treatment scheme that allowed onsite incineration. The questions asked by CORA and the answers provided were the same as for Case 3 until question 5a, "Is on-site incineration option precluded based on space or local considerations?" was answered "false." This generated an alternative path through both models resulting in the following, additional questions and *answers*:

Q10-a,b,c,d: Type of discharge option? *Discharge to surface water (Q10-d)*
Q10-f: Is the ash a hazardous waste? *False*
Q10-g: Is the concentration in ash above land disposal requirements? *False*
The actual questions instead of question numbers are located in Table 3.

This modified the results generated by both models such that for Case 4

- onsite incineration,
- ion exchange,
- discharge to surface water, and
- offsite solid waste landfill

were added to the recommended treatment sequence. This modification illustrated the sensitivity of these models to changes in field conditions. While CORA has long shown this capability, it was considered very positive that the ANN performed in an equivalent manner. The addition of the ion exchange unit to the recommended processes and activities illustrated more than a simple additive property of including just onsite incineration. The discharge to surface waters of treated effluent and the utilization of an offsite solid waste landfill further supported this observation.

Case 5

Case 5 was the last example completed. While similar to Case 4, it differed in that contamination concentration of the ash exceeded the land disposal restrictions resulting in the following different questions and/or *answers*:

Q10-g: Is the concentration in ash above land disposal requirements? *True*
Q10-h: Is an onsite RCRA landfill for solidified ash reasonable? *False*
The actual questions instead of question numbers are located in Table 3.

Case 5 involved more and different remediation alternatives than did the other cases. When compared to Case 4, the following processes or activities, resulting from the contaminant concentration being above the land disposal requirement, were different:

- in-situ stabilization,
- Solidification, and
- offsite RCRA landfill.

Table 7 summarizes the results of all of these five cases for both CORA and the ANN model. This table shows that the developed ANN recommended the same technology as CORA. In all cases the results were directly comparable with a 100

Table 7. Comparison of Results between CORA and the Calibrated ANN Model

Case		Remediation number (from CORA)[b]											
number	Results	105	201	301	302	312	316	317	401	404	406	503	504
1	CORA	x	x							x		x	x
1	NN[a]	x	x							x		x	x
2	CORA	x	x									x	x
2	NN	x	x									x	x
3	CORA	x	x		x							x	x
3	NN	x	x		x							x	x
4	CORA	x	x	x	x	x				x	x	x	x
4	NN	x	x	x	x	x				x	x	x	x
5	CORA	x	x	x	x	x	x	x	x		x	x	x
5	NN	x	x	x	x	x	x	x	x		x	x	x

[a]NN = Neural Network. [b]The actual remediation names instead of the numbers are located in Table 5.

percent matching rate between these two programs. This indicated that the neural network was ready for use on similar projects. Model calibration and verification had been achieved.

The final phase of this project was to update the original cost information for select CORA suggested remediation approaches. Data from Table 5 were used to determine subproject costs for work items recommended by the ANN in conjunction with estimates of the contaminant mass suggested by the user. These updated costs allow other users of this neural network to have a base cost for twenty types of remediation schemes for both containment and treatment remediation. As with CORA generated cost projections, the costs are subject to change over time and with actual site conditions. These cost projections however, can be readily modified within the ANN platform as newer data become available.

SUMMARY AND CONCLUSIONS

A DOS-based environmental expert system capable of identifying alternative groundwater and soil remediation options and their attendant costs was converted to artificial neural network (ANN) modeling platform as a means of updating the original code. The original expert system upon which this research was based was still capable of performing these tasks but had become more difficult to operate given current computer systems. In addition, the underlying expert system had become somewhat dated in terms of cost basis used for some technologies.

The effort ultimately required 54 data cases and 16 training runs to accurately produce the 100 percent pattern recognition rate deemed necessary. The trained neural network can replace the original expert system to precisely suggest remediation alternatives for VOCs in a homogeneous contaminated saturated soil. While updated cost data are included in the revised model, the approach selected affords future users the opportunity to also readily update the ANN with different types of chemicals and new or innovative technologies.

REFERENCES

1. CH2M Hill, *Cost of Remedial Action User Manual Version 3.0,* Environmental Protection Agency, Washington, D.C., 1990.
2. M. Stefik, Basic Concepts for Building Expert Systems, in *Building Expert Systems,* F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (eds.), Addison-Wesley Publishing, Reading, Massachusetts, pp. 59-86, 1983.
3. E. Castillo, J. Gutierrez, and A. Hadi, *Expert Systems and Probabilistic Network Models,* Springer-Verlag, New York, 1997.
4. E. Fisher, *The Potential Role of Expert Systems I the New Chemical Hazard Identification Process,* American Association for the Advancement of Science, Environmental Science and Engineering Fellow Report, Washington, D.C., 1984.
5. Y. Shih, *Neuralyst User's Guide (v1.4),* Cheshire Engineering Corporation, Pasadena, California, 1994.
6. L. Rogers, F. Dowla, and V. Johnson, Optimal Field-Scale Groundwater Remediation Using Neural Networks and the Genetic Algorithm, *Environmental Science and Technology, 29*:5, pp. 1145-1155, 1995.
7. P. A. Arora and W. F. McTernan, An Expert System to Determine the Probability of Pesticide Leaching, *Agricultural Water Management, 25,* pp. 57-70, 1994.
8. A. Aziz, A. Rashid, and K. V. Wong, A Neural Network Approach to the Determination of Aquifer Parameters, *Ground Water, 30*:2, pp. 164-166, March-April 1992.
9. R. Wilson and R. Sharda, Neural Networks: "Brain Metaphors" of Information Processing Turns Heads in OR/MS Community, *OR/MS Today,* pp. 36-42, August 1992.
10. United States Environmental Protection Agency, 2000, http://clu-in.com.
11. Ground-Water Remediation Technology Analysis Center, 1998, http://www.gwrtac.org.

Direct reprint requests to:

William F. McTernan
Professor
School of Civil and Environmental Engineering
207 Engineering South
Oklahoma State University
Stillwater, OK 74078-5033