

Using Artificial Neural Networks to Predict the Quality and Performance of Oil-Field Cements

P. V. Coveney, P. Fletcher, and T. L. Hughes

■ Inherent batch-to-batch variability, aging, and contamination are major factors contributing to variability in oil-field cement-slurry performance. Of particular concern are problems encountered when a slurry is formulated with one cement sample and used with a batch having different properties. Such variability imposes a heavy burden on performance testing and is often a major factor in operational failure.

We describe methods that allow the identification, characterization, and prediction of the variability of oil-field cements. Our approach involves predicting cement compositions, particle-size distributions, and thickening-time curves from the diffuse reflectance infrared Fourier transform spectrum of neat cement powders. Predictions make use of artificial neural networks. Slurry formulation thickening times can be predicted with uncertainties of less than ± 10 percent. Composition and particle-size distributions can be predicted with uncertainties a little greater than measurement error, but general trends and differences between cements can be determined reliably.

Our research shows that many key cement properties are captured within the Fourier transform infrared spectra of cement powders and can be predicted from these spectra using suitable neural network techniques. Several case studies are given to emphasize the use of these techniques, which provide the basis for a valuable quality control tool now finding commercial use in the oil field.

Cements are among the most widely used and the least well understood of all materials. Although cements are often viewed as simple "low-tech" materials, they are, in fact, inherently complex over many length scales. The starting material, cement

powder, is obtained by grinding cement clinker. The cement clinker is manufactured by firing limestone (providing calcium) and clay (providing silicon, aluminum, and iron). Gypsum (calcium sulfate dihydrate) is then added to moderate the subsequent hydration process. After grinding the clinker and gypsum, the cement powder consists of multisize, multiphase, irregularly shaped particles ranging in size from less than a micrometer to slightly more than 100 micrometers. When this starting material is mixed with water, hydration reactions occur that ultimately convert the water-cement suspension into a rigid porous material, which serves as the matrix phase for concrete, a cement paste-sand-rock composite.

The various chemical phases within the cement powder hydrate at different rates and interact with one another to form various reaction products. Some products deposit on the remaining unhydrated cement particle surfaces, but others form as crystals in the water-filled pore space between cement particles. Moreover, some of the hydration products contain nanometer-sized pores, so that the size range of interest for these materials is from nanometers to hundreds of micrometers, or even centimeters if one includes the rock aggregates used in concrete. Because of these complexities, many questions remain unanswered in the science of cementitious materials. As with most materials of industrial importance, the key relationships between processing and underlying physicochemical properties must be elucidated to obtain better control over the material in use.

The most common application of cement

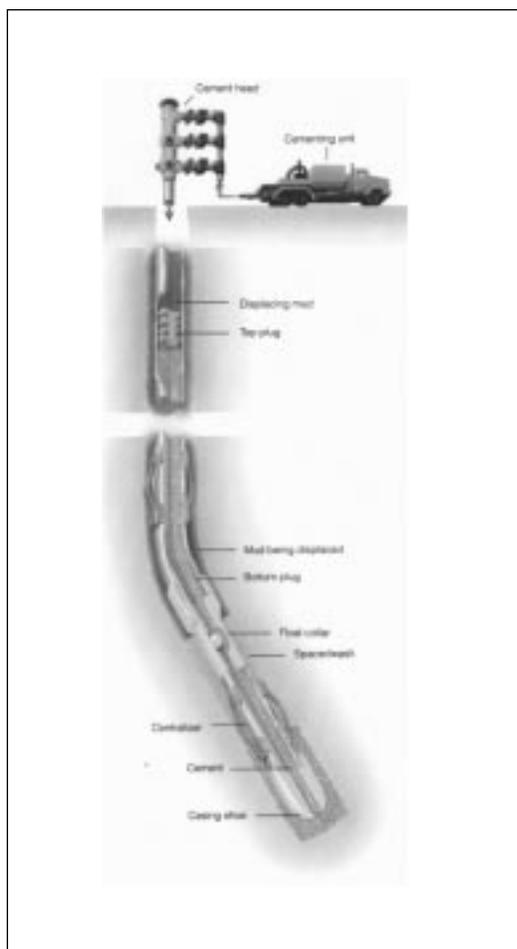


Figure 1. Cementing an Oil Well.

The main objective of such well cementing is to provide complete and permanent isolation of the formation behind the steel casing previously placed in the bore hole. The cement must be mixed to meet appropriate design parameters and is then pumped down hole, displacing all drilling mud from the annulus between casing and formation. Spacers or washes can be used along with top and bottom plugs to separate cement from drilling mud. Centralizers on the outside of the casing are used to keep the annular gap as even as possible.

is, of course, in building construction, where it has been used since at least Roman times. However, the work described here is concerned with another important application of cement—in the oil industry, where about three percent of the world's annual cement output is deployed. Cement is used to line oil and gas wells after drilling by pumping a cement slurry between the well bore and a steel casing inserted into the well, as shown in figure 1. During placement, the cement displaces all the drilling fluid originally present from the drilling operation itself. The cement then sets to form a low-permeability annulus, which isolates the productive hydrocarbon-bearing zones of the well from the rest of the formations, surplus water, and the surface.

Cement is used almost exclusively for oil-field cementing despite the fact that its performance is variable and not completely understood. Cement variability is observed between cements from different manufacturers, different cement batches from the same manufacturer, and samples from the same batch of cement that might have aged differently during storage. Because of all these problems, well cementing has remained until now more a black art than a science.

Various cement-slurry properties, such as compressive strength development, permeability to oil and gas, and flow behavior, need to be specified and controlled, taking into account the high temperature and pressure conditions prevailing down hole. For oil-field cement slurries, the *thickening time* plays a central role during slurry formulation because it is a measure of the time within which the cement is pumpable (American Petroleum 1982). Experimentally, it is the time taken to reach a specified consistency as measured under defined conditions. Longer than required thickening times are a potential waste of drilling time and an inefficient use of expensive chemical additives. Operational problems as a result of short thickening times are especially dramatic because the cement can set prematurely in the casing or pumping equipment. Such major operating failures (MOFs) can necessitate the complete redrilling of a many-thousands-of-feet well bore and can cost from \$1 to \$2 million; less severe MOFs in which a limited amount of redrilling is required typically cost around \$0.6 million. Therefore, considerable, time-consuming experimental effort is devoted to precise control of slurry thickening times.

Application Description: The FTIR Spectra of Cements

In view of the overwhelming complexity of cement hydration, a valuable quality control tool would be a model predicting performance properties of a given cement sample prior to its use. However, the mathematical modeling of cement hydration based on mechanistic understanding is still in its infancy (but see Coveney and Humphries [1996]). The approach taken here is to dispense with detailed physicochemical characterization of the cement particles in favor of methods based on a combination of statistics and AI. With this approach, cement composition and performance properties are correlated with a judiciously chosen measurement that implicitly contains key information on cement composition, particle-size distribution, and surface chemistry. To give the method any chance of commercial success, this measurement also has to be relatively inexpensive and easy to perform on a routine basis.

The measurement chosen was based on the use of *infrared spectroscopy*, a common analytic technique used in the chemical sciences. It is well known that every chemical species has its own unique infrared spectrum. Indeed, chemists most commonly use this technique in a qualitative mode by matching up spectral features in an unknown compound with previously recorded spectral data on known compounds available in lookup tables. An experienced chemist, working in a specified area of chemistry, can often identify a chemical by direct visual inspection of its infrared spectrum. A more specialized yet equally well-established application is *quantitative analysis* of chemical mixtures, wherein measured spectra of unknown chemical composition are regressed against linear combinations of infrared spectra either of the pure chemical components or of mixtures of known chemical composition (Beebe and Kowalski 1987).

The particular variant used in this work is that of the Fourier transform infrared (FTIR) spectrum of dry cement powders, sometimes known as diffuse reflectance infrared Fourier transform spectroscopy (DRIFTS). In this technique, white-light radiation from a Michelson interferometer is focused on a compacted sample using a moving mirror. Radiation impinging on the sample undergoes two types of reflection. The first is *specular reflectance*, where the radiation is reflected from the sample surface as if from a mirror. The second is *diffuse reflectance* whereby a proportion of the

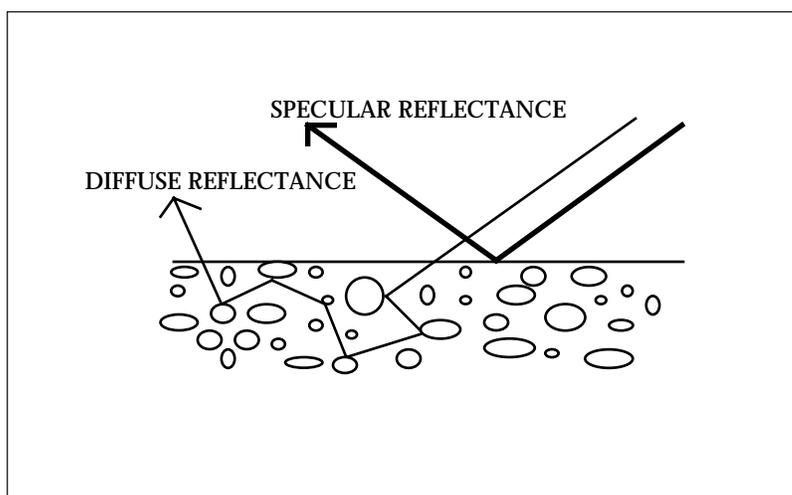


Figure 2. Schematic of the Diffuse Reflectance Process.

radiation penetrates the sample and is reflected from particle surface to particle surface. At each reflection, a degree of energy absorption occurs as indicated in figure 2.

Energy is absorbed because of the vibration and stretching of chemical bonds in the molecules of the powder. The reflected light reemerges from the sample and is collected by a second ellipsoidal mirror. The Fourier transform technique is used to convert the emergent radiation into a spectrum of absorbance versus frequency. The experimental method for collecting FTIR spectra of cement powders has been described elsewhere (Hughes et al. 1995, 1994). The wavelength range of the mid-infrared region of the electromagnetic spectrum is approximately 2.5×10^{-3} cm to 2.5×10^{-4} cm or 4000 to 400 wave numbers, where *wave numbers* are reciprocal wavelength in units of cm^{-1} .

Information Contained within Cement FTIR Spectra

Particle Size: The extent of diffuse reflectance is inherently related to the particle size of the sample but in a generally unknown manner. Large coarse particles allow the incident radiation to penetrate deeply into the sample, thus increasing absorption. However, large particles show greater specular reflectance that distorts the frequency spectrum. As the particle size of a sample is reduced, the depth of penetration and, therefore, absorption is less because more particles are present to reflect and limit the depth of penetration. Spectra are therefore distorted as a function of particle size although sample dilution in KBr minimizes these effects. Accordingly, our spectral mea-

surement is made with samples diluted to a concentration of 10 percent by weight in finely ground, infrared inactive, potassium bromide.

Composition: To a good first approximation, the FTIR spectrum of any multicomponent mineral assembly is a linear superposition of the spectra of the pure mineral components. In the case of oil-field cements, the American Petroleum Institute (API) lays down notional chemical composition specifications based on the so-called Bogue clinker phases: alite (tricalcium silicate), belite (dicalcium silicate), aluminite (tricalcium aluminate), and aluminoferrite (tetracalcium aluminoferrite). These Bogue phases, which themselves provide only an approximate chemical description, are traditionally specified by means of a linear transformation of the chemical composition of the clinker expressed in terms of its major oxides, which can be determined directly by other, more lengthy, noninfrared methods. Spectral features of the major Bogue cement chemical phases in the mid-infrared are dominated by vibrations and stretching modes of water molecules that are located on mineral surfaces, within the sulfate and carbonate minerals, or in calcium hydroxide. Within the mineral phases present, chemical bonds between silicon and oxygen, aluminium and oxygen, and iron and oxygen are also active in the mid-infrared region. Despite the aforementioned complexities that are the result of particle-size distributions and the occlusion of minerals, it is established that linear statistical techniques can be used to correlate spectral characteristics with cement chemical composition, provided due care is taken in the sample preparation (Hughes et al. 1995, 1994).

Other Spectral Attributes: Lack of crystallinity, impurities in minerals, and prehydration have a more subtle effect on spectra, usually broadening absorbance peaks and shifting the frequencies at which absorbance occurs.

We can therefore assert with confidence that diffuse reflectance infrared spectra of cements contain information on the composition, particle-size distribution, and surface chemistry of the material, all of which influence cement hydration.

Methods for Predicting Cement Properties from FTIR Spectra

In this section, we discuss the design and construction of a cement properties database, modeling techniques, and the predictive capabilities of the models.

The Cement Properties Database

The methods we use for making quantitative predictions are based on establishing statistical correlations between cement infrared spectra and selected cement physicochemical parameters. Specifically, we were interested in seeking to establish unambiguous relationships between the infrared spectra and cement properties such as chemical composition, particle-size distribution, and thickening time. To achieve such correlations required the construction of a database containing data on 158 oil well cements collected worldwide. Our database is one of the most comprehensive currently available on oil-field cement properties. It contains the following standard physical and chemical data on each of the 158 cements:

First is the cement mineral composition expressed in weight percent (wt%) of the following minerals: alite, belite, aluminite, ferrite, gypsum, the sulfates bassanite and syngenite, calcium hydroxide, and calcium carbonate.

Second is the cement oxide composition expressed in weight percent of the following oxides: SO_3 , Al_2O_3 , Fe_2O_3 , MgO , Na_2O , CaO , SiO_2 , P_2O_5 , TiO_2 , CrO_2 , MnO_2 , ZnO , and SrO_2 .

Third is the binned particle-size distribution (PSD bin), in volume fraction occupancy, and mean particle diameter, in microns, as measured by Cilas granulometry.

Fourth is weight loss on ignition, free lime content, and insoluble residue.

Fifth is the surface area (in cm^2g^{-1}) as measured by Blaine's method, which provides an estimate of the total surface area of all the cement particles.

Sixth is the digitized thickening-time curve for a neat cement slurry at 50°C and solid/water ratio of 0.44.

Seventh is the digitized thickening-time curve for a slurry retarded with 0.2-percent D13 (a Schlumberger Dowell proprietary cement-setting retarder) at 85°C and solid/water ratio of 0.44.

Eighth is the diffuse reflectance FTIR spectra recorded at 2 cm^{-1} resolution using a Nicolet 5DX spectrometer.

Modeling Techniques

The primary objective of this research was to construct models to predict cement properties from FTIR spectra as the sole input data. The most important cement information that one would hope to extract from infrared spectra are (1) chemical composition according to the Bogue and oxide representations, (2) particle-size distribution, and (3) thickening-time profiles for neat and retarded ce-

ment slurries. Accordingly, five independent statistical models were constructed for the prediction from FTIR spectra of the following properties selected from the database:

Model A: The concentrations of the four API-specified Bogue minerals plus gypsum, syngenite, bassanite, calcium hydroxide, and calcium carbonate

Model B: The concentrations of the major oxides together with loss on ignition, free lime content, and insoluble residue

Model C: Particle-size distributions and mean particle diameter

Model D: Digitized neat thickening-time curves

Model E: Digitized retarded thickening-time curves

These models were subsequently used independently of one another.

It has previously been demonstrated that cement mineral compositions (model A) can be predicted from FTIR using linear statistical techniques (Hughes et al. 1994; Fierens and Verhagen 1972). A suitable procedure, described elsewhere (Martens and Naes 1989; Beebe and Kowalski 1987; Sharf, Illman, and Kowalski 1986), is based on *partial least squares* (PLS). This technique is a variant on simple multiple linear regression, which has the capacity to filter noise and redundant information from spectra prior to prediction. In this study, PLS is used for the prediction of mineral compositions only. All other models make full use of artificial neural networks (ANNs).

The full relationships between the measurable properties of a cement powder and its slurry performance are not known and are expected to be complex, that is, highly nonlinear (Coveney and Humphries 1996; Fletcher and Coveney 1995; Fletcher et al. 1995; Bensted and Beckett 1993; Billingham and Coveney 1993; Hunt 1986). Therefore, it is best to choose a technique for finding such correlations that makes as few assumptions as possible regarding their nature. ANNs offer the possibility of finding input-output correlations of essentially arbitrary complexity and consequently formed the basis for the AI methods we used in this work. The main feature of the neural network methodology is that input-output information is correlated using a system of interconnected nodes (Hush and Hornee 1993; Lippmann 1987; Rumelhart and McClelland 1986). These nodes, also called *neurons*, are the computational analog of nerve cells in the human brain. A single node is a processing element that combines a set of input to produce a single numeric output (figure 3).

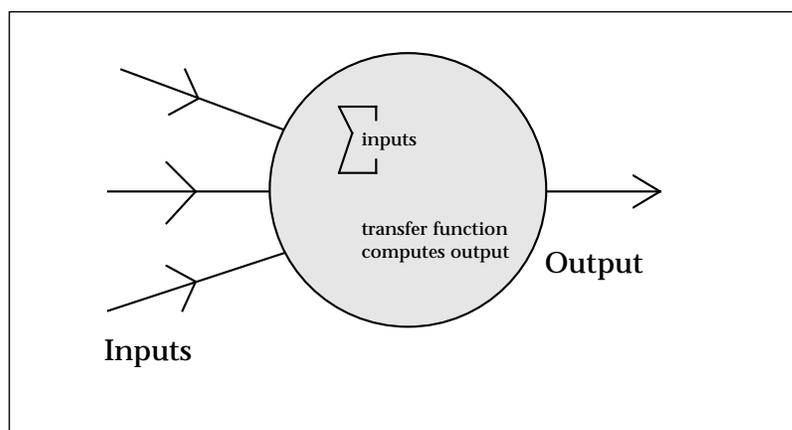


Figure 3. Schematic of a Single Node.

The strength of the output signal is given by a nonlinear function called the *transfer function*. Commonly, the transfer function is based on a weighted sum of the input signals. A complete neural network is constructed from an arrangement of individual neurons that link input data to output data by a network of arbitrary complexity. Within any architecture the strength of the signal received by any one node is a weighted sum of input sent by all the nodes to which it is connected. In the commonly used, supervised, feed-forward, layered networks, nodes in an input layer first receive signals equal to the values of the external input data. This information is passed on in a nonlinearly convolved fashion to nodes in an output layer representing output data (figure 4). The network architectures and nonlinear expressions are modified using a supervised training procedure such that input data are correlated with output data. In some networks, there can be one or more layers of neurons connecting the input and output layers. These (hidden) layers add mathematical features to networks necessary to model complex relationships.

A fully trained ANN is effectively a nonlinear map between specified variables that is capable of filtering noise in the input data and has a predictive capacity; that is, it is capable of making predictions for situations not previously encountered. The procedures for optimizing ANNs are described elsewhere (Masters 1993) and use goodness-of-fit criteria based on minimum residual prediction errors for test data.

Neural networks have the following valuable features: (1) they respond with high speed to input signals, (2) they have generalized mapping capabilities, (3) they filter noise from data, (4) they can perform classification

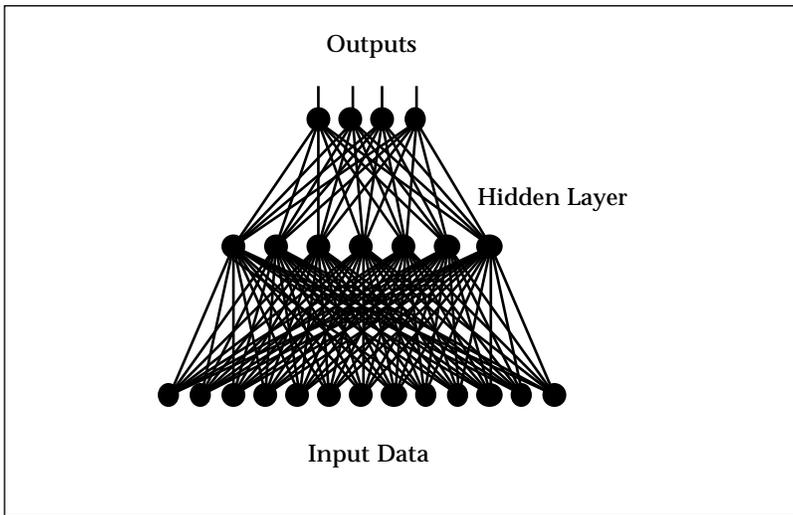


Figure 4. Schematic of an Artificial Neural Network.

as well as function modeling, and (5) they can encode information by regression or iterative supervised learning.

Some drawbacks of neural network methods are (1) they are data intensive; (2) training is computationally intensive and requires significant elapsed wall-clock time; (3) they have a tendency to overtrain if the network topology is not optimized, resulting in their mapping training data extremely well but becoming unreliable in dealing with new data; and (4) predictions are unreliable if extrapolated beyond the boundaries of the training data.

Many different types of neural network can be implemented to solve a wide range of complex nonlinear problems. We originally worked with multilayer perceptrons (MLPs), which are composed of three layers in which the number of nodes in the input and output layers were fixed by virtue of the mapping sought. Thus, the number of nodes in the input layer is equal to the number of individual pieces of data in a single-cement *input data record* (also called an *input vector*), but the number of nodes in the output layer is equal to the number of separate parameters being predicted from the input vector. However, there are certain computational drawbacks to these MLP networks: Finding the optimal number of nodes in the hidden layer is time consuming because network training by the back-propagation-of-errors algorithm is slow, and in addition, there is some possibility of the network becoming trapped in a local, rather than the global, error minimum.

The network type we found to be the most

suitable for predicting cement properties employed Gaussian radial basis functions (Moody and Darken 1989, 1988) in a single hidden layer. We preferred these networks because their underlying learning algorithms are fast, and based on linear algebra, they are guaranteed to find global optima. In such *radial basis function networks*, also sometimes referred to as *localized receptive field networks*, the nodes in the hidden layer are different than those in a multilayer perceptron: They are radial-distribution functions that have centers and widths expressed in terms of the n -dimensional space defined by the input data vectors. These Gaussian basis functions produce a nonzero response only when an input vector falls within a small, localized region of this n -dimensional space centered on the mean and within the specified width of the basis function.

The process of constructing and optimizing such networks involves several stages: First, an arbitrary number of Gaussian basis functions have to be selected. Their means and standard deviations (widths) are determined on the basis of the available data vectors to be used for training by a procedure such as n -dimensional K-means clustering. This standard statistical procedure exploits the natural clustering of the input data to locate the means (that is, the centers) of the selected number of nodes such that their average Euclidean distances from all the input data vectors are minimized. The output from this arbitrarily chosen number of radial basis functions are then linearly correlated to the supplied target (output) vectors. The final stage of network optimization is performed by systematically varying the number of clusters and overlap parameters to achieve an optimum fit with the training data.

For all types of ANN architecture employed and models constructed (A to E), the networks were trained using a subset of the full database and their predictive capabilities evaluated using a completely independent test data set—that is, one containing data that had not previously been used by the network during training—selected randomly from the database. The importance of network optimization and training in the construction of reliable and robust ANN models cannot be overstressed.

The extensive computation time for optimizing even radial basis function neural networks becomes an issue when spectral data are used as input variables. A typical mid-infrared FTIR spectrum collected at 2-wave-number resolution of the kind used here has approximately 2000 digitized points. Thus, to use

FTIR spectra as input data for neural networks, it was found necessary to first reduce the number of variables representing any spectrum. This reduction was performed using the principal component method based on spectral eigenvector analysis and led to a reduction in the useful information content in each spectrum to 35 principal components, which allows network training and validation to be performed on PCs and workstations. For each of the models B to E, the spectra were always reduced to 35 principal components, although the optimum architectures were different for each model. More details of the theoretical basis of the modeling procedure are given elsewhere (Fletcher and Coveney 1996).

Predictive Capabilities of the Models

In this section, we discuss the predictive capabilities of the five models.

Model A—Mineral Composition Predictions, PLS Model The expected uncertainties in mineral composition predictions have been described in detail elsewhere (Hughes et al. 1995, 1994). They are summarized in table 1, which lists the various chemical phases present as well as the concentration ranges and associated uncertainties with which the phases are found (in weight percent). As in all the models to be discussed, the quoted uncertainties refer to the imprecision of the model predictions (σ is the standard deviation) compared with the known, experimentally measured values of the same quantities. The predictions of the sulfate minerals (gypsum, bassanite, and syngenite), calcium hydroxide, calcium carbonate, aluminate, and ferrite are generally good and can be used to detect aging of cements, as our later case studies show. The major uncertainties lie in the prediction of the individual silicate phases, although total silicates (alite + belite) is predicted well.

Model B—Major Oxide Analyses, ANN Model Table 2 displays the cement chemical analysis represented more fundamentally in terms of the major oxides present, together with the concentration ranges and associated uncertainties with which these oxides occur.

In all cases, the oxides, weight loss on ignition (LOI), free lime, and insoluble residue variables are predicted well, although the uncertainties are greater than expected errors on the measurements. The major uncertainties lie in the predictions of the concentrations of CaO and MgO. These uncertainties arise from the fact that the variance in the levels of CaO and MgO is known to be small, and the errors in prediction are proportionately large.

Table 1. Uncertainties in Model A Predictions.

Component	Concentration Range (wt%)	Uncertainty/ 2σ (wt%)
Alite	42–70	± 5
Belite	4–35	± 5
Alite + Belite	70–82	± 1.5
Aluminate	0–15	± 1
Ferrite	5–20	± 1
Syngenite	0–3	± 0.6
Gypsum	0–6	± 0.6
Bassanite	0–6	± 0.6
Ca(OH) ₂	0–3	± 0.1
CaCO ₃	0–4	± 0.2

Table 2. Uncertainties in Model B Predictions.

Component	Concentration Range (wt%)	Uncertainty/ 2σ (wt%)
SO ₃	1.5–4.0	± 0.2
Al ₂ O ₃	3.0–7.0	± 0.3
Fe ₂ O ₃	1.8–7.0	± 0.3
MgO	0.5–3.0	± 0.3
Total alkalis	0.2–1.5	± 0.2
CaO	61–67	± 1
SiO ₂	19.5–24	± 0.4
Insoluble residue	0–0.9	± 0.3
Loss on ignition	0.5–2.5	± 0.3
Free lime	0.4–2.1	± 0.3

As with the mineral composition model A, cement-oxide compositions can be estimated with errors a little greater than those of the direct composition measurement itself, but general trends in chemical composition can be determined readily from both models.

Model C—Particle-Size Distribution Bins and Mean Diameter, ANN Model

Figure 5 shows a typical particle-size-distribution prediction for an oilwell cement, and table 3 lists its uncertainties. In all cases, the prediction errors are greater than the expected measurement errors, although general trends are predicted well.

Models D and E—Thickening-Time Curve Predictions, ANN Model Figure 6 shows predictions of the full digitized thickening-time curves for the retarded and neat slurries for a typical oil-field cement. In this

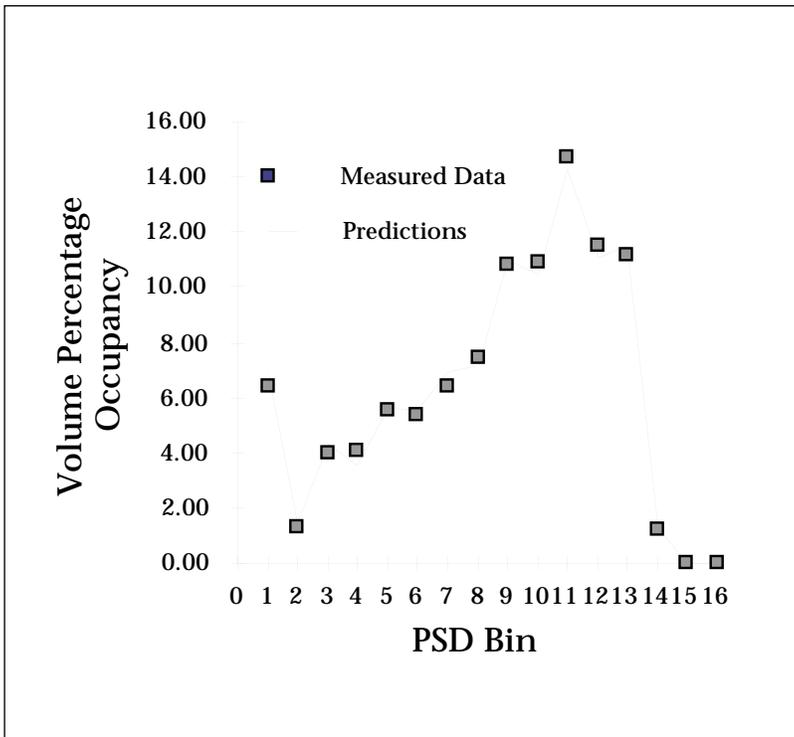


Figure 5. Measured and Predicted Particle Size for a Typical Oil-Field Cement.

Table 3. Uncertainties in Model C Predictions.

Bin	Diameter Range/ μm	Uncertainty/ 2σ (%) ¹
1	0-1	± 1
2	1-1.5	± 1
3	1.5-3	± 1
4	2-3	± 1
5	3-4	± 1
6	4-6	± 1
7	6-8	± 1.2
8	8-1	± 1.2
10	16-24	± 1.2
11	24-32	± 1
12	32-48	± 1
13	48-64	± 1
15	96-128	± 1
16	128-192	± 1

1. Error on mean particle diameter $2\sigma = \pm 1\mu\text{m}$.

example, the digitization simplifies the curve, yet the general trends, including the point of departure for cement setting and the actual thickening time, can be seen clearly.

The expected error limits for thickening-time predictions for both neat and retarded formulations are shown in table 4. These uncertainties are typically less than experimental thickening-time measurement errors and, thus, support the use of FTIR as a rapid, quantitative predictor of cement-slurry thickening times.

Application Case Studies

The application offers two basic levels of interpretation: One is the qualitative assessment of cement FTIR spectral features. The other is the interpretation of quantitative predictions from the models. In this section, we provide a few examples of how this AI-based application works in commercial operations where it has been deployed for almost two years.

Qualitative Interpretation

Qualitative interpretation involves identifying and comparing relevant features of cement spectra with no reference to the AI system. The simplest qualitative method is direct visual inspection of the spectra. This inspection involves identifying the presence of specific components by the presence of characteristic absorbance bands and reference to the spectra of pure components provided in preexisting lookup tables (Hughes et al. 1995, 1994). The relative spectral changes can give indicators to changes in performance between batches. An alternative qualitative technique is spectral subtraction where one FTIR spectrum is subtracted from a second to leave a so-called residual spectrum that is used to identify differences in the spectra. The residual spectrum is particularly useful in the detection of contaminants.

Quantitative Interpretation

Quantitative interpretation involves predicting the composition, particle-size distribution, and performance properties of the cement using the ANN-based prediction modules. Most applications involve comparing the properties of one cement with another possibly suspect batch using the cement FTIR spectra passed through our predictive models. Any statistically significant differences in composition or particle-size distribution will indicate differences in performance. In some cases, such as the detection of aged cements, the changes in mineral compositions can indicate

changes in performance. The performance predictions themselves can be used to support or confirm the qualitative interpretations. In some cases, the compositional differences between cements are subtle and not easily interpreted. In these cases, direct prediction of performance is informative.

Case A—Detecting a Barite-Contaminated Cement A cement sample was observed to yield an unexpectedly long thickening time compared with a normal cement taken from a different storage silo. A residual FTIR spectrum was obtained by subtracting the spectrum of the normal cement from that of the rogue cement. Figure 7 shows the residual spectrum compared with the spectrum of pure barite; the barite characteristics are confirmed by table lookup from existing databases. (Use of the AI system is unnecessary for this application). The correspondence of spectral features confirmed the presence of barite in the rogue sample. Barite contamination leads to the slurry being overretarded when the cement is used in a slurry formulated on the basis of an uncontaminated cement.

Case B: Detecting an Aged Cement A cement from one storage silo was observed to show mixing and pumping problems and yield a short thickening time compared to cement samples from other silos. The spectrum of the problem cement is shown in figure 8, where it is compared with the spectrum of a normal cement. Enhanced syngenite features are visible in the spectrum of the problem cement; as in the case of the barite-contaminated sample, the spectral features characteristic of syngenite are confirmed by table lookup from existing databases.

The linear partial least squares composition model A predicted the syngenite content of the problem cement to be 2.7 wt% compared with 0.9 wt% for the normal cement. Aging to form syngenite is consistent with the observed shortening of thickening times and pumping problems. The retarded slurry performance ANN model predicted the thickening time for the aged cement to be 50 minutes shorter than for the normal cement.

Case C: Identification of a Rogue Cement Figure 9 shows a retarded thickening-time curve for an oilwell cement as predicted from its FTIR spectrum using our ANN model. The predicted data are compared to an average thickening-time curve obtained from experimental measurements on five different batches of the same cement. An indication of the normal batch-to-batch variation because of storage is given by the two standard-devia-

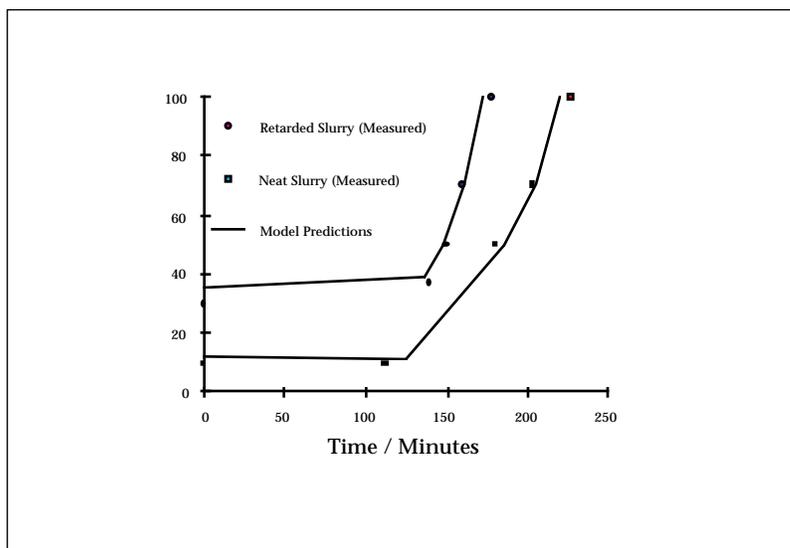


Figure 6. Measured and Predicted Thickening-Time Curves for a Typical Oil-Field Cement.

The upper continuous curve displays the artificial neural network predictions for a neat slurry, and the continuous lower curve shows similar predictions for a retarded slurry.

Table 4. Uncertainties for Thickening-Time Predictions.

	Neat Slurry	Retarded Slurry ¹
Mean errors/minutes	±16.0	±19.5
2σ /minutes*	±31	±37
Mean % error	±9.5	±10.6
Range of fit/minutes	50–350	50–50

1. 2σ indicates upper maximum expected errors.

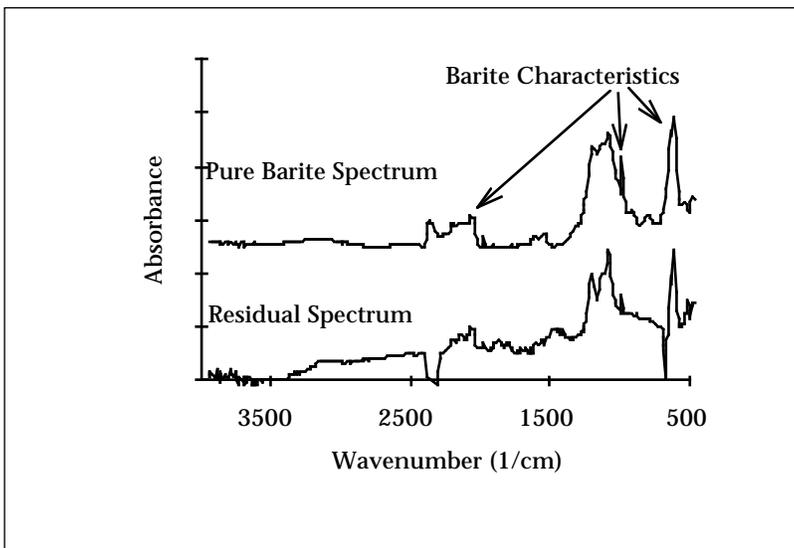


Figure 7. A Barite-Contaminated Cement.

Pure-barite FTIR spectrum and residual spectrum obtained by subtracting the FTIR spectrum of a barite-contaminated cement from the spectrum of a normal cement.

tions limit. The rogue batch is identified as having a very short thickening time compared to the expected range for this cement and an unusually high initial consistency. These predictions were subsequently confirmed experimentally.

This example makes critical use of the ANN-based performance-prediction capability to identify a rogue cement without recourse to interpreting cement composition or particle-size distribution, which, on its own, is likely to provide ambiguous results. Recall that premature setting of the slurry, as is the case here, would likely lead to a costly major operating failure if such a cement were pumped in the field.

These case studies indicate the scope and power of prediction afforded by ANNs. However, the intriguing issue remains of how these ANNs actually succeed in making correct composition and performance predictions from the compressed and convolved representation of cement FTIR spectra.

ANNs have been criticized at times because they appear to work like black boxes. However, the substantial quantity of knowledge they encode is available for more detailed interrogation and can be used effectively in its own right. To illustrate, we mention in passing that we have constructed other ANN models that map chemical compositions and particle-size distributions directly onto slurry thickening-time curves (Fletcher and Cov-

eney 1996). These networks can be used to investigate the sensitivity of, for example, changes in thickening time to changes in the values of input parameters, such as the amount of aluminum or iron in the cement, and so on. In some cases, these models have confirmed previously established qualitative trends known within the cementing community, such as the observation that increasing the iron content increases thickening times, but at the same time have made these relationships more quantitative. However, in many other instances, including, for example, the dependence on composition variables of the kickoff time ("point of departure") and the subsequent rapidity of thickening following the usual quiescent period during which there is essentially constant slurry consistency (see figures 6 and 9), no previous knowledge—either qualitative or quantitative—existed. Moreover, we have shown that a genetic algorithm can be used to invert the nonlinear forward mapping provided by such an ANN to furnish the precise physicochemical composition of a cement needed to deliver specified performance properties. This capability is particularly remarkable because it implies that in principle at least, it might one day be possible to tailor-make a cement to suit any particular application.

The cement quality-assurance tool that we have described here is the result of the powerful combination of a modern AI technique (ANNs) and the established laboratory measurement technique of FTIR spectroscopy. The integration of these two methodologies is achieved in routine use by passing the digitized output from an FTIR spectrometer into a 486 PC on which the trained ANNs reside. In this way, a single cement powder FTIR spectrum provides information simultaneously on cement chemical composition, particle-size distribution, and setting profile (including thickening time), together with a flag indicating the degree of statistical reliability to be expected from the predictions emanating from the AI device. This flag indicates whether or not a cement being analyzed lies within the part of infrared parameter space on which the ANNs have been trained: If the former, the predictions are classified as reliable; if the latter, they are described as unreliable. As a consequence, one can record the FTIR spectrum of a cement powder and predict its setting time in about 15 minutes of real time rather than wait for more than 4 hours to observe when the slurry will actually set.

Our radial basis function neural network and other codes were homemade and were

developed on UNIX platforms. At the time when the method was transferred from Schlumberger Cambridge Research, where it had been developed, to Schlumberger Dowell's Europe-Africa Technology Center in Aberdeen, a decision was taken to port all codes to the MATLAB commercial package—which is platform independent and, thus, immediately accessible on a PC available in the field. The porting of codes to MATLAB reduced the coding requirements of the commercial product to a minimum. It should be noted that our ANN codes made no special use of MATLAB's intrinsic features, nor did MATLAB influence in any way our choice of network architecture; our work was completed prior to the availability of the neural network toolbox within this package.

Application Use and Payoff

The cement quality control technique described here proved so successful in our research laboratories that a decision was made to turn it into a commercial product, called CEMQUEST (cement quality estimation). The technique is now being used in our Aberdeen regional field laboratories to detect and avoid cementing problems normally associated with cement quality and variability. CEMQUEST can predict, directly from the FTIR spectrum, composition, particle-size distribution, and thickening times for certain cement-slurry formulations.

The advantages of using CEMQUEST compared with previous cementing practice are manifold. Obviously, there is the large time and personnel savings that accrues from predicting cement-setting properties in this way. Other benefits include the avoidance of operational cementing failures because of batch-to-batch variation, aging or cement contamination, and improved efficiency of cement-slurry formulation design through the identification of important slurry performance characteristics.

Since early 1995, CEMQUEST has been in routine use within Schlumberger Dowell, where it is part of the overall set of techniques employed for achieving improved cement-slurry design and reliability on a daily basis. It has also attracted the attention of cement manufacturers and clients (oil companies) for whom cement quality control work is also now being done on a regular basis. CEMQUEST is able to save about \$3 to \$5 million a year for each client through its ability to detect potential major operating failures before they arise. The costs of slurry formulation are also reduced by CEMQUEST: Rapid screening and

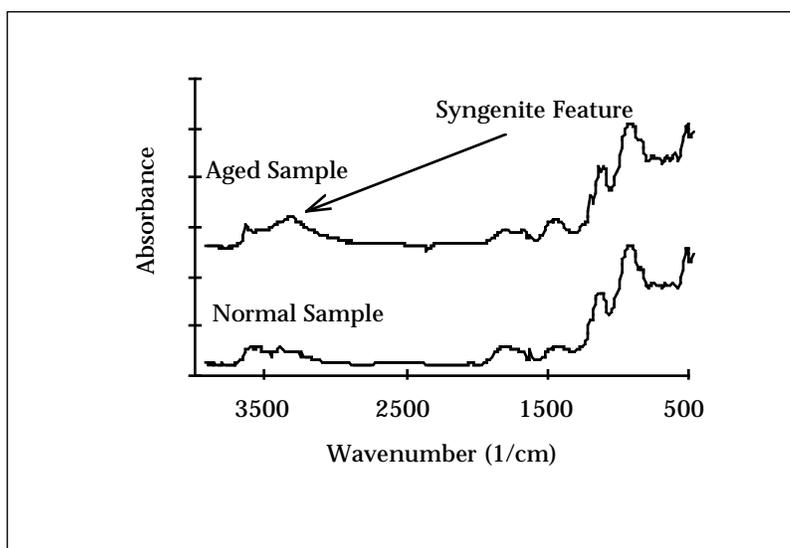


Figure 8. A Syngenite Aged Cement.

elimination of bad cements saves about 10 percent of the time taken by the lengthy process of formulation optimization. This savings translates to about \$1000 a week for each formulation in routine laboratory testing.

We expect additional benefits to arise in time for at least two reasons: The first reason will be as a result of the buildup of a larger cement database, extending the domain of validity of the existing neural network models (which will require periodic retraining). A second reason will be as a result of an enhanced reputation for Schlumberger Dowell based on the increasing reliability of its cementing jobs through the use of the current product on a day-to-day basis.

Application Development and Deployment

The development of the CEMQUEST prototype at Schlumberger Cambridge Research was the result of about 12 person-years of effort, beginning in 1991 and ending in mid-1993. The work involved coordinating a vast cement data-collection exercise, with samples sent from all areas of the world in which Dowell has cementing operations. Thus, about 160 distinct cements, with various physicochemical properties—chemical composition, particle-size distribution, FTIR spectra, slurry thickening curves, and so on—had to be recorded. The reproducibility of all these measurements had to be investigated. All this work required the cooperation of

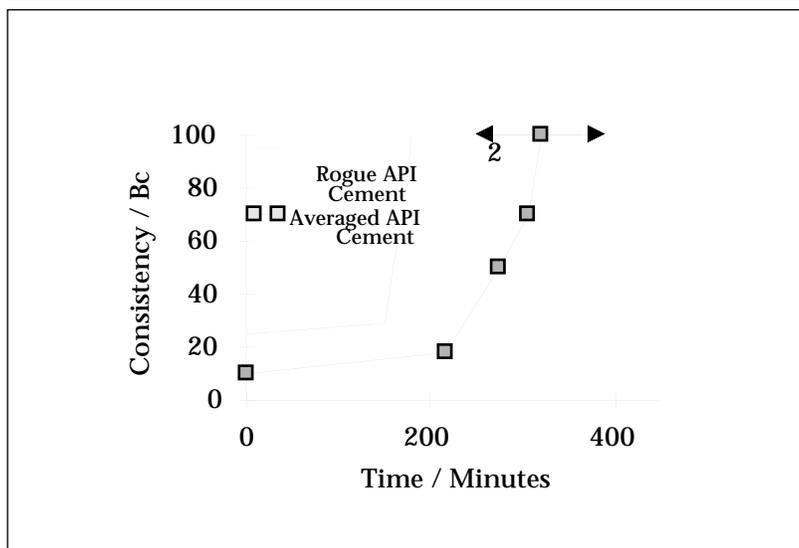


Figure 9. Thickening-Time Curve Predictions for a Rogue Oil-Field Cement.

many colleagues in our product center (then in St Etienne, France) and in Aberdeen. In addition, some of the chemical-analysis work was performed externally at low cost. The samples needed careful storage in the absence of moisture and carbon dioxide to prevent alteration of cement properties with time because these substances are readily absorbed by cement powders. The end result was a substantial cement database that was used for developing the final neural network models.

While data were being acquired, approximately three person-years of effort were devoted to an investigation of the feasibility of cement-quality estimation using FTIR spectra linked to thickening-time curves. The initial aim was to establish whether any of the cement data could reliably be used for such predictive purposes. When it was decided possible in late 1991 and early 1992, the target was to demonstrate that the same could be achieved on the basis of the single and easily performed FTIR powder measurement. The feasibility of this process was fully confirmed in late 1992 and opened the way to a commercially viable product. During 1993, about one person-year's effort was assigned to the development of the basic MATLAB code for transfer to Aberdeen in mid-1993. One of the authors was transferred to Aberdeen to ensure correct technical implementation of the product and prepare for its commercialization. This transfer was seen as important to guarantee a successful future for the product because at that site, there was previously only limited expertise in the recording and interpreting of FTIR spectra.

Maintenance

Experience in the transfer of this product from research to operations showed that the key limitation is associated with the recording of FTIR spectra in field laboratories. Owing to the large size of the database and the technical issues involved in producing accurate predictive models, all these models were developed in our research laboratory during 1992 to 1993, using FTIR spectra recorded there. Thus, it was of paramount importance to ensure that FTIR cement spectra recorded on different spectrometers in the field center were closely coincident with the database spectra recorded in research. Clear guidelines for ensuring reproducible spectra had to be laid down by the research group.

Maintenance of the software and the database is now the responsibility of Aberdeen. To date, it has not proved necessary to update this knowledge base, owing to the rather wide representative coverage of the original cement data. However, data are being kept on all significant outlier cements detected by the reliability flag within the current AI system. Predictions of the physicochemical and performance properties of such outlier cements cannot be made reliably using the existing database; so, at a future stage, their measured FTIR spectra and performance properties will be added to supplement the data set. When an addition is made, new ANN and other models will also need to be constructed and validated. This activity will be carried out entirely in Aberdeen on a periodic basis.

Summary and Conclusions

Using ANNs and conventional statistical methods, we showed that the information in the FTIR powder spectra of cements can be used to predict composition, particle-size distributions, and thickening-time curves for simple slurries. This discovery has established the FTIR measurement as a signature for cement performance. The measurement can be used as a rapid technique to estimate cement quality and detect batch-to-batch variability in cements. Specific case studies have demonstrated that the product can detect batch-to-batch variability between manufacturers as well as aging and contamination of a given cement. Thus, it is capable of preventing the occurrence of costly major operating failures in oil-field cementing operations. Under the name of CEMQUEST, the application is finding successful commercial application use in the oil field.

Acknowledgments

PVC is grateful to Reid Smith for helpful comments and advice made during the preparation of this article.

References

- American Petroleum. 1982. API Spec 10: Materials and Testing for Well Cements. Dallas: American Petroleum Institute.
- Beebe, K. R., and Kowalski, B. R. 1987. An Introduction to Multivariate Calibration and Analysis. *Analytical Chemistry* 59:1007A.
- Bensted, J., and Beckett, S. J. 1993. *Advances in Cement Research* 5:111.
- Billingham, J., and Coveney, P. V. 1993. *Journal of the Chemical Society: Faraday Transactions* 89:3021.
- Coveney, P. V., and Humphries, W. 1996. *Journal of the Chemical Society: Faraday Transactions* 92:831.
- Fierens, P., and Verhagen, J. P. 1972. *Journal of the American Ceramic Society* 55:306.
- Fletcher, P., and Coveney, P. V. 1996. Predicting Cement Composition and Performance with Artificial Neural Networks and FTIR Spectroscopy. *The American Institute of Chemical Engineering Journal*. Forthcoming.
- Fletcher P., and Coveney, P. V. 1995. *Advanced Cement Based Materials* 2:21.
- Fletcher P.; Coveney, P. V.; Hughes, T.; and Methven, C. M. 1995. *Journal of Petroleum Technology* 47:129.
- Hughes, T. L.; Methven, C. M.; Jones, T. G. J.; Pelham, S. E.; Fletcher, P.; and Hall, C. 1995. *Advanced Cement-Based Materials* 2:91.
- Hughes, T. L.; Methven, C. M.; Jones, T. L. H.; Pelham, S. E.; Vidick, B.; and Fletcher, P. 1994. Rapid Cement Quality Control Method for Improved Oilfield Cementing. Paper presented at the Offshore Technology Conference, 2-5 May, Houston, Texas.
- Hunt, L. P. 1986. *Cement and Concrete Research* 16:190.
- Hush, D. R., and Hornee, B. G. 1993. Progress in Supervised Neural Networks. *IEEE Signal Processing Magazine* 10:8.
- Lippmann, R. P. 1987. An Introduction to Computing with Neural Nets. *IEEE Signal Processing Magazine* 4:4.
- Martens, H., and Naes, T. 1989. *Multivariate Calibration*. Chichester, U.K.: Wiley.
- Masters, T. 1993. *Practical Neural Network Recipes in C++*. San Diego, Calif.: Academic.
- Moody J., and Darken, C. J. 1989. *Neural Computation* 1:281.
- Moody, J., and Darken, C. J. 1988. Learning with Localized Receptive Fields. In *Proceedings of the 1988 Connectionist Models Summer School*, eds. A. N. Touretzky, G. Hinton, and T. Sejnowski, 133. San Francisco, Calif.: Morgan Kaufmann.
- Rumelhart, D. E., and McClelland, J. J., eds. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1*. Cambridge, Mass.: MIT Press.
- Sharf, M. A.; Illman, D. L.; and Kowalski, B. R. 1986. *Chemometrics*. New York: Wiley Interscience.



Peter Coveney is a senior scientist at the Schlumberger Cambridge Research Laboratory, United Kingdom. He holds a B.A., an M.A., and a D.Phil. from Oxford University. Before joining Schlumberger Cambridge Research, he was a lecturer in physical chemistry at the University of

Wales, a junior research fellow at Oxford University, and a visiting fellow at Princeton University. He also currently holds a visiting fellowship at Wolfson College, Oxford University, in theoretical physics. Coveney is the coauthor of two books, *The Arrow of Time* (Ballantine, 1991) and *Frontiers of Complexity* (Ballantine, 1995), both with Roger Highfield.



Philip Fletcher is a senior scientist and section head at Schlumberger Dowell Europe-Africa Technology Center in Aberdeen, Scotland. He holds a B.Sc. in chemistry and a Ph.D. in physical chemistry. Before joining Schlumberger Cambridge Research in 1984, he held research fellowships at Oxford University, the University of California, and the City University London. He has researched in theoretical physical chemistry, solution chemistry, mineral chemistry, colloid chemistry, and applications of statistics. His current work is in product development and field support. Fletcher is also the author of an undergraduate textbook in geochemical thermodynamics.



Trevor Hughes is a research scientist at the Schlumberger Cambridge Research Laboratory, United Kingdom, where he has worked since 1985. He obtained a B.Sc. in geochemistry from the University of Liverpool and an M.Sc. in mineral chemistry from the University of Birmingham.

Hughes began his career in the precious metals mining and refining industry. He holds 5 U.S. patents concerning techniques for monitoring the composition of oil-field drilling fluids and has published 17 scientific papers.



Diagrammatic Reasoning

Cognitive & Computational Perspectives

Edited by Janice Glasgow, N. Hari Narayanan, and B. Chandrasekaran

Foreword by Herbert Simon

“Understanding diagrammatic thinking will be of special importance to those who design human-computer interfaces, where the diagrams presented on computer screens must find their way to the Mind’s Eye.... In a society that is preoccupied with ‘Information Superhighways,’ a deep understanding of diagrammatic reasoning will be essential to keep the traffic moving.” – *Herbert Simon*

Diagrammatic reasoning—the understanding of concepts and ideas by the use of diagrams and imagery, as opposed to linguistic or algebraic representations—not only allows us to gain insight into the way we think but is a potential base for constructing representations of diagrammatic information that can be stored and processed by computers.

Diagrammatic Reasoning brings together nearly two dozen recent investigations into the cognitive, the logical, and particularly the computational characteristics of diagrammatic representations and the reasoning that can be done with them. Following a foreword by Herbert Simon (coauthor of one of the most influential papers on reasoning with diagrams, which is included here) and an introduction by the editors, chapters provide an overview of the recent history of the subject, survey and extend the underlying theory of diagrammatic representation, and provide numerous examples of diagrammatic reasoning (human and mechanical) that illustrate both its powers and its limitations. Each of the book’s four sections begins with an introduction by an eminent researcher who provides an interesting personal perspective while he or she places the work in proper context.

ISBN 0-262-57112-9 800 pp., index. \$50.00 softcover

The AAAI Press • Distributed by The MIT Press

Massachusetts Institute of Technology, Cambridge, Massachusetts 02142

To order, call toll free: (800) 356-0343 or (617) 625-8569. MasterCard and VISA accepted.