Conf. 75116--8

A MULTIVARIATE STATISTICAL PATTERN RECOGNITION
SYSTEM FOR REACTOR NOISE ANALYSIS*

MASTER

R. C. Gonzalez,[#] L. C. Howington,[†]
W. H. Sides Jr., and R. C. Kryter

Oak Ridge National Laboratory
Oak Ridge, Tennessee 37830

By acceptance of this article, the
publisher or recipient acknowledges
the U.S. Government's right to
retain a nonexclusive, royalty-free
license in and to any copyright
covering the article.

NOTICE

This report was prepared as an account of work
sponsored by the United States Government. Neither
the United States nor the United States Energy
Research and Development Administration, nor any of
their employees, nor any of their contractors,
subcontractors, or their employees, makes any
warranty, express or implied, or assumes any legal
liability or responsibility for the accuracy, completeness
or usefulness of any information, apparatus, product or
process disclosed, or represents that its use would not
infringe privately owned rights.

* Research sponsored by the Energy Research and Development
Administration under contract with the Union Carbide Corporation.

[#] Consultant, Electrical Engineering Department, the University
of Tennessee, Knoxville.

[†] Graduate Assistant, Electrical Engineering Department, the
University of Tennessee, Knoxville.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Jeg

R. C. Gonzalez,[#] L. C. Howington,[†] W. H. Sides Jr., and R. C. Kryter

Oak Ridge National Laboratory
Oak Ridge, Tennessee 37830

Abstract

A multivariate statistical pattern recognition system for reactor noise analysis was developed. The basis of the system is a transformation for decoupling correlated variables and algorithms for inferring probability density functions. The system is adaptable to a variety of statistical properties of the data, and it has learning, tracking, and updating capabilities. System design emphasizes control of the false-alarm rate. The ability of the system to learn normal patterns of reactor behavior and to recognize deviations from these patterns was evaluated by experiments at the ORNL High-Flux Isotope Reactor (HFIR). Power perturbations of less than 0.1% of the mean value in selected frequency ranges were detected by the system.

Index Terms: pattern recognition, multivariate, noise analysis, reactor, statistical, algorithms.

1. Introduction

The extraction and use of information from noise signals originating within a nuclear power plant for assessing the operational status of the plant has been advocated for a number of years.¹⁻⁵ Reactor incidents that resulted in a partial loss of core mechanical integrity or of fuel element cooling have been reported where a conspicuous change in the nature of a randomly fluctuating variable preceded the incidents.⁶⁻⁸ Utilization of pattern recognition techniques⁹ in a comprehensive noise analysis system could eliminate the chief impediment to widespread implementation of noise diagnosis in nuclear power plants by providing a tool for on-line, unattended computer monitoring and interpretation of large quantities of operational data needed for detecting anomalous reactor performance. These same pattern recognition techniques also provide a basis for developing learning, adaptive computer procedures that could establish the normal behavior patterns of a reactor and compensate for plant pattern changes introduced by aging, plant alterations, or changed operating procedures. Such a noise analysis system could perform data logging, cataloging, and signature analysis automatically, thus eliminating the need for a trained noise analyst in all but exceptional circumstances.

In the system described in this paper, a linear transformation is used to reduce multidimensional, correlated input observations to a set of uncorrelated samples that can be processed independently and then fed to a learning, adaptive system that estimates (or updates) the probability density function of each variable. Using the resulting densities, the system establishes alarm thresholds in accordance

with false-alarm criteria determined by the plant operator. This ability to set the false-alarm rate to a statistical level that is acceptable to the plant operator is important in assuring reliable operation of the system and its acceptance by industry. Since it must adapt to a changing environment, the system provides the capability for parameter tracking and updating to detect and compensate for slowly occurring changes of the normal operating characteristics of the plant. Because the system algorithms are formulated as recursive relationships, the system parameters can be iteratively updated as new samples become available, thus eliminating the need for storing a complete set of previous data.

2. Problem Formulation

Figure 1 shows the basic problem considered in this paper. Observations, such as noise measurements, are obtained from a plant and are fed into a pre-processing stage whose function is to convert the raw signal measurements into a form suitable for the recognition stage. In our experiments at the HFIR, these signals were derived from a neutron-flux noise sensor. The preprocessor consists of a set of computer programs based on the fast Fourier transform (FFT).¹⁰ It yields an ensemble-averaged, discrete power spectral density (PSD) function $F(\omega_i)$, $i = 1, 2, \dots, n$, where $\omega_1, \omega_2, \dots, \omega_n$ are frequencies uniformly distributed in the interval 0 to $(n-1)$ Hz. A given PSD function will be represented throughout this paper as an n -dimensional column vector

$$x = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} \quad (1)$$

where $x_i = F(\omega_i)$.

Finally, the output from the preprocessor is fed into a pattern recognition stage whose function is to yield a decision concerning the status of the portion of the plant being monitored. The nature of the HFIR

ORNL-DWG 75-15868

Figure 1. Elements of an automated plant surveillance system.

* Research sponsored by the Energy Research and Development Administration under contract with the Union Carbide Corporation.

[#] Consultant, Electrical Engineering Department, the University of Tennessee, Knoxville.

[†] Graduate Assistant, Electrical Engineering Department, the University of Tennessee, Knoxville.

and preprocessor has been covered in detail elsewhere.^{5,10} The remainder of this paper is devoted to a description of the design and implementation of the pattern-recognition system.

3. Pattern Recognition System Structure

The design of the system was based on the following constraints: (1) adaptability to the statistical properties of the data; (2) learning, tracking, and updating capabilities; (3) time-varying memory; and (4) control over the false-alarm rate.

The first constraint indicates that the system should be independent of specific data characteristics; that is, its applicability should not be limited to data of specific types, such as, for example, Gaussian, log-normal, Rayleigh, etc.

The second set of constraints indicates that the system should be capable of learning (i.e., estimating) the required operational parameters from representative data with as little operator assistance as possible. It should track changes in operating conditions, and update the system parameters to take into account significant changes in the data characteristics.

Time-varying memory means that the system should be capable of "forgetting" operating conditions that no longer represent the plant's status. Finally, the system should have variable alarm thresholds that can be set at statistical levels by the operator.

The pattern recognition stage is expanded in Fig. 2, which illustrates how the functions interact to yield the system output. Each output x of the preprocessor is multiplied by an $n \times n$ decoupling transformation matrix A (box 1) to yield a vector y whose elements are uncorrelated. This transformation (discussed in Sect. 4) allows processing of each element of y independently, which reduces the problem to a set of one-dimensional variables. (The reason for this procedure is that considerable difficulty is encountered in statistically characterizing the original, correlated variables in an n -dimensional space.)

The learning or estimating process (boxes 2 and 3) involves the transformation matrix A and the probability density functions $p_1(y_1)$, $p_2(y_2)$, ..., $p_n(y_n)$ of the elements of y . All learning algorithms require

a set of representative samples, called a "training data set" (box 4), which is derived from the plant sensors during normal operations.

Once the initial learning has been completed, the system parameters are tracked (box 5) to detect significant changes in the data characteristics. These results are used to validate the classification results and to determine whether to update the parameters. The updating criteria are stored in the memory (box 6).

The pattern classifier (box 7) is fed three inputs: decoupled observations, the latest tracking results, and the latest estimates of the probability density functions $p_i(y_i)$, $i = 1, 2, \dots, n$. These functions are used for classification in the following manner. If the i th component y_i of the vector y has the probability density function $p_i(y_i)$, the probability that any y_i is less than or equal to a value a_i is

$$P(y_i \leq a_i) = \int_{-\infty}^{a_i} p_i(y_i) dy_i \quad (2)$$

Similarly, the probability of y_i being equal to or exceeding a value b_i is

$$P(y_i \geq b_i) = \int_{b_i}^{\infty} p_i(y_i) dy_i \quad (3)$$

Any observation y_i that is outside the interval $[a_i, b_i]$ is, by definition, abnormal. The thresholds a_i and b_i are obtained by specifying values for $P(y_i \leq a_i)$ and $P(y_i \geq b_i)$ and then by solving Eqs. (2) and (3) for a_i and b_i by numerical integration methods. Thus, the probability of an alarm (i.e., the probability of y_i outside $[a_i, b_i]$) can be controlled at will by the operator by specifying $P(y_i \leq a_i)$ and $P(y_i \geq b_i)$.

4. Decoupling Transformation

From a set of random vector variables of the form $x = (x_1, x_2, \dots, x_n)^T$, we shall obtain corresponding vectors $y = (y_1, y_2, \dots, y_n)^T$ by means of a linear transformation

$$y = Ax, \quad (4)$$

where A is an $n \times n$ matrix. This transformation must produce a set of y vectors whose elements are uncorrelated.¹¹

The mean or average vector for the x 's is

$$m_x = E\{x\}, \quad (5)$$

where E indicates the expected value with respect to x . Similarly, the covariance matrix is

$$C_x = E\{(x - m_x)(x - m_x)^T\}. \quad (6)$$

Since $y = Ax$, we have

$$m_y = E\{Ax\} = AE\{x\} = Am_x, \quad (7)$$

where $m_y = AE\{x\}$ is true because A is constant.

ORNL-DWG 75-15869

Figure 2. Diagram of the recognition stage.

The covariance matrix of the y vectors is

$$C_y = E\{(y - m_y)(y - m_y)^T\}. \quad (8)$$

Using $y = Ax$ and $m_y = Am_x$ results in the following expression for C_y in terms of C_x :

$$\begin{aligned} C_y &= E\{(Ax - Am_x)(Ax - Am_x)^T\}, \\ &= AE\{(x - m_x)(x - m_x)^T\}A^T, \\ &= AC_xA^T. \end{aligned} \quad (9)$$

Since C_x is a symmetric matrix, its eigenvectors corresponding to distinct eigenvalues are orthogonal.¹² If the eigenvalues of C_x are distinct (which is normally the case in practice), and if the rows of A are chosen as the normalized eigenvectors of C_x , the resulting vectors y will have components that are uncorrelated and their variances will be given by the eigenvalues of C_x .¹¹ This implies that the covariance matrix C_y has the following properties: (1) the off-diagonal elements c_{ij} ($i \neq j$) are zero; and (2) the element c_{ii} is equal to the variance of the i th component of the transformed vectors.

By using Eq. (4) to decorrelate the components of x and by working in the transformed space, each element of y can be treated independently, as far as first- and second-order statistics are concerned. This allows utilization of several one-dimensional algorithms for testing and estimating one-dimensional probability density functions. (This approach is analogous to the second-order approximation used frequently in the design of control systems.)

5. Estimation of the Probability Density Functions

In Section 3 a signal vector y was defined as abnormal if any of its components y_i is outside its corresponding interval $[a_i, b_i]$. Three approaches will be discussed for obtaining the threshold values a_i and b_i from the density function $p_i(y_i)$, $i = 1, 2, \dots, n$. Since the densities $p_i(y_i)$ are estimated independently of each other, we simplify the notation considerably by dropping the subscripts and letting y represent any component of y and $p(y)$ its corresponding probability density function.

Histogram Approach

A histogram of each component y can be used to determine the threshold parameters a and b . Suppose the i th component of a set of N vector observations varies between 0 and $y_{\max}$. The interval $[0, y_{\max}]$ is divided into I subintervals, and each of the N observations of y is assigned to the subinterval into which it falls. Dividing the total number of observations that fall in each subinterval by N yields a histogram of the values of y . Since a histogram is a discrete approximation to $p(y)$, the values of a and b can be obtained directly from the histogram by replacing the integral with a summation sign in Eqs. (2) and (3).

Statistical Inference Approach

Suppose that a certain form of the densities $p(y)$ is postulated. It is possible to test this hypothesis statistically by using the Kolmogorov-Smirnov (K-S) goodness-of-fit test.¹³ Let $Y = \{y_k\}$ be a set of N sample values of an element of the vector y , and define

$$u(y_k) = \begin{cases} 1 & \text{if } y_k > 0 \\ 0 & \text{if } y_k < 0 \end{cases}, \quad (10)$$

where $k = 1, 2, \dots, N$. The empirical distribution function $F_N(y)$ is given by

$$F_N(y) = \frac{1}{N} \sum_{k=1}^N u(y - y_k). \quad (11)$$

The continuous distribution function is, by definition,

$$F(y) = \int_{-\infty}^y p(y) dy, \quad (12)$$

where $p(y)$ is the assumed probability density function of the component y . The K-S one-sample statistic is defined as

$$D_N = \max_y |F_N(y) - F(y)|, \quad (13)$$

which is the maximum of the differences between the postulated and empirical distribution functions over the entire sample set Y .

The K-S goodness-of-fit test is to reject, with significance level α , the hypothesis that the sample set is governed by the specified density function if $D_N > D_{N,\alpha}$. Values of $D_{N,\alpha}$, such that $P(D_N > D_{N,\alpha}) = \alpha$, have been tabulated for numerous values of N and α .^{13,14,19} As an illustration, suppose that $\alpha = 0.05$. For $N = 50$, the tables in ref. 13 give $D_{N,\alpha} = 0.1884$. Then the hypothesis that a postulated $p(y)$ characterizes the given set is rejected if $D_N > 0.1884$. What this implies is that the probability of rejecting the hypothesis when it is true is 0.05. Another important property of the K-S test is that it can be used to determine the number of samples (N) needed to approximate $F(y)$ by the empirical distribution $F_N(y)$, such that the error D_N does not exceed an acceptable value $D_{N,\alpha}$.¹³

Direct Estimation by the Potential Function Method

Use of the K-S test requires that a form be specified for the density $p(y)$. In many cases it may be desirable to estimate $p(y)$ directly from the sample set Y . The method of potential functions^{9,15,16} developed in the following discussion can be used to formulate an iterative procedure for the estimation of these densities.

Let the unknown probability density function be represented by the M -term series expansion

$$p(y) = \sum_{m=1}^M c_m \phi_m(y), \quad (14)$$

where the $\phi_m(y)$'s, $m = 1, 2, \dots, M$, are a set of orthonormal functions (e.g., Hermite, Legendre, or trigonometric polynomials).⁹

Given any sample y_k from the set Y , the potential function of this sample is defined by the expression

$$K(y, y_k) = \sum_{m=1}^M \phi_m(y) \phi_m(y_k). \quad (15)$$

After the appearance of L samples from Y , an estimate of $p(y)$, denoted by $\hat{p}(y)$, is formed, as follows:

$$\hat{p}(y) = \frac{1}{L} \sum_{k=1}^L K(y, y_k). \quad (16)$$

It can be shown¹⁶ that the average squared error

$$e = \int_Y [p(y) - \hat{p}(y)]^2 dy \quad (17)$$

decreases as a function of increasing L . Equation (16) can be expressed in a variety of recursive forms that are useful for updating purposes.^{9,15}

Each of the approaches previously described has inherent advantages and disadvantages in its implementation. The histogram is one of the simplest methods for obtaining an estimate of a probability density function. If the histogram requires periodic updating, however, there must be a way of storing a histogram value for each specified subdivision of every variable in the vector y . If y is n -dimensional, and L divisions are specified for each component, then nL histogram values must be kept either in the system memory or in some peripheral storage device.

The K-S approach is suitable for applications where a form of each density can be specified. It can also be used in conjunction with a method that estimates the densities directly, such as the potential function technique, in which case the K-S test can be used to evaluate the quality of the estimates. The advantages of being able to statistically evaluate a specified (or calculated) density are many. For example, if a Gaussian or a log-normal density assumption passes the K-S test, a significant reduction in storage over the histogram approach can be realized because only two parameters--the mean and the variance--need to be stored for each component of y .

Direct estimation using potential functions is, in principle, the most powerful of the three approaches. Application of this method requires specification of a set of orthonormal functions and selection of specific values for M and L . Once the orthonormal functions have been specified, the potential function approach can be implemented in a recursive framework that uses the K-S test to establish the quality of the approximation for any values of M and L .

6. Classification of Multivariate Gaussian Data in the Transformed Space

Implementation of the pattern recognition system requires estimation of the mean vector m_x and covariance matrix C_x of the vectors x . These parameters completely define the multivariate Gaussian probability density function

$$p_x(x) = \frac{1}{(2\pi)^{n/2} (\det[C_x])^{1/2}} \exp[-1/2(x - m_x)^T C_x^{-1} (x - m_x)] \quad (18)$$

Since m_x and C_x are available, it is of interest to compare the classification capabilities in the original and transformed space if the data are (or are assumed to be) Gaussian.

By using Eqs. (4), (7), and (9), it can be shown that $p_y(y) = p_x(x)$. Consequently, the linear transformation $y = Ax$ does not affect the form of the Gaussian density. Since the elements of y are uncorrelated, it follows in the Gaussian case that

$$p(y) = p_1(y_1) p_2(y_2) \dots p_n(y_n), \quad (19)$$

where $p_i(y_i)$, $i = 1, 2, \dots, n$, are one-dimensional Gaussian densities. The importance of Eq. (19) is that each function $p_i(y_i)$ can be evaluated separately, using the K-S test to assess how closely it approximates the Gaussian assumption. To the authors' knowledge, no test exists to do this in the original multivariate space.

The decoupled variables can be classified by using the techniques described in the previous sections. If the data pass the Gaussian test, however, an alternative classification method based on the Mahalanobis distance concept may be used.^{9,17} This distance measure is defined as

$$D(x) = (x - m_x)^T C_x^{-1} (x - m_x) \quad (20)$$

for the original variables, and as

$$D(y) = (y - m_y)^T C_y^{-1} (y - m_y) \quad (21)$$

for the transformed variables.

By using Eqs. (4), (7), and (9), it can be shown that $D(y) = D(x)$, indicating that identical results are obtained in both spaces. However, it is a simpler procedure to implement Eq. (21) because C_y is a diagonal matrix whose inverse can be obtained with little computational effort.

Using the Mahalanobis distance, an observation y is said to be abnormal if

$$D(y) > T, \quad (22)$$

where T is a nonnegative threshold. Setting $D(y) = T$ defines the equation of a hyperellipse in n -dimensional space. Use of Eq. (21), therefore, can be interpreted as enclosing the region of normality by this surface. Then, any point y outside the hyperellipse (i.e., $D(y) > T$) is treated as abnormal.

To set T at a statistical level, we note that Eq. (21) is a chi-square distribution with n degrees of freedom;¹⁷ that is,

$$\chi_n^2 = (y - m_y)^T C_y^{-1} (y - m_y). \quad (23)$$

Values of χ_n^2 are available in standard statistical tables. For example, if $n = 30$ and a 0.99 interval of confidence is desired, the result is that $\chi_{30}^2 = 50.89$. It follows that by letting $T = \chi_n^2$ a statistical level of confidence can be associated with the chosen threshold. The tightness of the fit can be controlled by using

$$T = \chi_{cn}^2, \quad (c > 0), \quad (24)$$

since χ_{cn}^2 increases with increasing n . The decoupled-variable approach also offers an alternative method for setting the threshold T based on the training data. This technique is to specify thresholds a_1 and b_1 for each density $p_1(y_i)$ using Gaussian densities in Eqs. (2) and (3). The training patterns are classified using these thresholds, and, for each pattern that is classified as normal, the distance $D(y)$ is computed. At the end of the procedure, T is set equal to $cD_{\max}(y)$, where c is a positive number and $D_{\max}(y)$ is the maximum of $D(y)$ computed during the learning phase. Factor c controls the degree of tightness of the surface that encloses the region defined as normal.

Although this discussion is limited to a Gaussian assumption, a similar argument could be developed for a larger class of ellipsoidally symmetric probability density functions that would include the Gaussian density as a special case.⁹

7. Parameter Tracking and Data Labeling

The principal parameters required to implement the pattern recognition system are the mean vector, the covariance matrix, and the probability density functions. These parameters are estimated initially with a block of training data. During normal operation, however, changes in the initial estimates over long periods should be expected, and there must be a way to track and update the system parameters.

Development of an unattended monitoring system must accommodate changes produced in signals being monitored by routine changes in the operating conditions of the plant. For example, neutron noise PSD measurements at the HFIR show that there is a strong correlation between PSD amplitude at some frequencies and the withdrawal position of the control rods. Unlike the tracking problem, these effects are generally predictable and should be taken into account in the design of the system.

Parameter Tracking

As an illustration of the parameter tracking problem, let N denote the number of samples in the training set. The mean vector is given by Eq. (5), which can be approximated by the relation

$$\bar{x}_x(N) = \frac{1}{N} \sum_{i=1}^N x_i, \quad (25)$$

where $\bar{x}_x(N)$ denotes the estimate obtained with N samples. If one sample is added to the set, Eq. (25) becomes

$$\bar{x}_x(N+1) = \frac{1}{N+1} \sum_{i=1}^{N+1} x_i,$$

which can be expressed in terms of $\bar{x}_x(N)$ as follows:⁹

$$\bar{x}_x(N+1) = \frac{1}{N+1} (N\bar{x}_x(N) + x_{N+1}). \quad (26)$$

This recursive expression is a convenient method for updating the mean vector as new samples become available.

When Eq. (26) is used over a long updating sequence, the accuracy of the estimates is affected by round-off error. To minimize this effect, the parameters are updated in "blocks" rather than a sample at a time. Thus, a mean vector $\bar{x}_x(M)$ computed from Eq. (25) with M samples can be added to the original N -sample estimate by using the relation

$$\bar{x}_x(N+M) = \frac{1}{N+M} [N\bar{x}_x(N) + M\bar{x}_x(M)], \quad (27)$$

which can be easily derived from Eq. (26). A similar expression can be obtained for the covariance matrix.⁹

A more sophisticated tracking method is described in ref. 18 where the tracking algorithms developed account for the nature of the parameter variations and potentially can yield better estimates of the mean vector and covariance matrix in a variety of dynamic conditions.

Data Labeling Approach

Thus far, the discussion has been limited to one transformation matrix for all observations, which assumes that the data are reasonably clustered about some region in the vector space of x . In practice, there may be more than one such cluster because of changes in operating conditions of the plant. In the HFIR, for example, it is not unusual to have several distinct data clusters during the course of a fuel cycle.⁵ This precludes the use of a single transformation matrix to describe the behavior of the plant for all operating conditions.

The basic idea behind data labeling is to associate with each observed sample a set of identifying parameters, or labels, $L = \{L_1, L_2, \dots, L_R\}$ which depend on the operating conditions of the plant when the sample is taken. To minimize the number of different elements in L , each component L_i , $i = 1, 2, \dots, R$, can be made to correspond to a range of operating conditions. In the HFIR a minimum set of two parameters can be used, where L_1 represents increments of stationary rod position, and L_2 denotes increments of operating power level.

Once a set of labels has been specified for a particular plant, the training samples are divided into groups according to their label specification. The design procedure previously discussed is applied to each of the data groups. Sample x with label set L is classified by using the transformation function and decision thresholds associated with this label set.

Our experience with labeling has been limited to a formulation of the concept. As the recognition system is exercised in practical applications, however, we expect that labeling techniques will have considerable impact on classification performance.

8. Experimental Results

The pattern recognition system was implemented on a minicomputer with 32K words of core memory, two 1.25 M-word disks, and analog-to-digital (A/D) converters. The computer is also equipped with a video graphics display unit which, although not essential for on-line operation, does facilitate the presentation of recognition results. A smaller configuration consisting of 28K words of core, one disk, and A/D converters would be sufficient to implement all the necessary algorithms.

The data for testing recognition performance were recorded at the HFIR during rod-perturbation experiments (Fig. 3) in which a 4 ± 0.5 -Hz noise signal was injected into the control rod servo demand system. (The noise signal was obtained by bandpass filtering a white noise signal from a generator.) The control rod moved in response to the fluctuating demand, causing perturbations of less than 0.1% about a 98-MW mean level in the signal from the neutron detector. This signal was amplified and recorded, and its PSD was computed with and without the 4-Hz perturbation. The PSDs without the perturbation constituted the training set; those with the perturbation were used as "abnormal" observations to test the recognition sensitivity of the system after the training phase was completed. The training set contained 357 PSDs from a 12-hr continuous learning period. The abnormal set had 51 PSDs.

ORNL-DWG 75-19870

Figure 3. Diagram of the HFIR experimental setup.

Each pattern x was formed as the $\log_{10}$ of a 30-dimensional vector: components 2-19 contained PSDs from 0.5 to 5 Hz at 0.25-Hz intervals, components 20-30 contained PSDs from 5.5 to 10 Hz at 0.5-Hz intervals, and component 1 was the square root of the average of the other 29 components.

The system was first trained using histograms with probabilities of 0.005 to establish thresholds a and b for each of the 30 components of the pattern vectors. After training was completed, each of the 51 abnormal patterns was input for classification, and all were flagged as falling outside the bounds of normal operation.

After detection of an abnormality, the system displays a set of data analysis options (Fig. 4a) to aid the operator in interpreting the abnormality. The following descriptions of these options are based on the first pattern in the abnormal set.

Option 1 ("code 1" of Fig. 4a) displays graphs on the same coordinate system of the average PSD (as determined from the training set) and the abnormal PSD input pattern. Option 1(A) plots the patterns before their transformation, and flags abnormal individual components (Fig. 4b). Option 1(B) plots the mean and abnormal pattern after their transformation (Fig. 4c). Exercise of option 1(C) yields a plot

of the abnormal components of the transformed pattern, such that their distances from the corresponding component of the average plot are relative to the number of standard deviations that each component lies outside the thresholds of abnormality (Fig. 4d).

Option 2(A) lists the limits (a_i, b_i) for each abnormal component, the value of each abnormal component, and the relative distance outside the normal limits in units of standard deviations. Although data classification is carried out in the transformed space, interpretation is made easier by knowledge of which components in the untransformed space contributed most to a given abnormality. One way to obtain this information is the following. After training is completed, a variable vector x is set equal to the mean of the training patterns. Then, a single component x_i is incremented by some Δx_i , transformed, and tested for abnormality by using the limits in the transformed space. When the pattern is sufficiently distorted to be classified as abnormal, that value of x_i becomes the upper limit of normality for the component in the untransformed space. The lower limit is determined in the same manner by decrementing x_i from its mean value. This procedure can be used to determine the limits of normality for every component similar to the limits (a_i, b_i) calculated in the transformed space. However, since the limits in the untransformed space do not account for correlation effects, they are not used for classification. The results of this procedure constitute option 2(B) (Fig. 4f). Components 13-16, corresponding to the frequency range from 3.25 to 4.0 Hz, are labeled as abnormal, which corresponds approximately to the 4-Hz noise perturbation. Option 2(C) is analogous to option 2(B), except that the technique for determining abnormal components in the original space is based on the Mahalanobis distance (Fig. 4g).

Option 3 displays both the Mahalanobis distance of the observed pattern and the base Mahalanobis distance $D_{\max}(y)$. The Mahalanobis distance of the observed, abnormal pattern is much greater than the maximum Mahalanobis distance for any of the training patterns, indicating that, in this case, the Mahalanobis distance and the histogram approaches agree that the observation is abnormal. This agreement held for all patterns in the abnormal set.

Option 4 displays a histogram of any component of the training set in the transformed space, along with the corresponding limits (a_i, b_i). The histogram of component 30 (Fig. 4i), for example, shows that this component is well outside the upper limit.

Options 5-8 give the user latitude in obtaining other observations. Option 5, which allows the user to recall a pattern that has been processed, might be desirable for checking the conditions just prior to the occurrence of an abnormality. With option 6, the user can observe the next pattern without losing control of the surveillance program, and option 7 returns control to the surveillance program to continue monitoring. Finally, option 8 terminates the monitoring process.

As a second experiment, the system was trained by statistical inference. The K-S test was applied to the training data, assuming that each transformed component was characterized by a log-normal probability density function. Since the training data were formed from the logarithms of PSD values, the problem was to test these data against a Gaussian assumption. With $\alpha = 0.01$ and $N = 357$, the tables of ref. 19 gave 0.055 for $D_{N,\alpha}$. Only component 30 failed the test at the $\alpha = 0.01$ level of significance. This overall

(a)

(b)

(c)

(d)

(e)

(f)

(g)

(h)

(i)

Figure 4. Analysis options: (a) system output following detection of an abnormality, (b) option 1(A), (c) option 1(B), (d) option 1(C), (e) option 2(A), (f) option 2(B), (g) option 2(C), (h) option 3, and (i) option 4.

agreement with the log-normal assumption explains the similarity of classification results obtained with the histogram and Mahalanobis distance approaches.

The limits for each component were determined by integrating the corresponding log-normal densities [Eqs. (2) and (3)], and the classification experiment was repeated. Although the limits were generally different than those in the histogram experiment, each of the abnormal patterns in the set was also flagged by the system when log-normal densities were used. Figure 5 shows the transformed limits for the same abnormal pattern described earlier. Comparison of this figure with Fig. 4(e) shows that the largest discrepancy in classification is in component 30, as expected.

9. Concluding Remarks

The foregoing experimental results show that it is feasible to implement a recognition system that will (1) learn the characteristics of normal operation in a reactor, and (2) detect small variations from the normal pattern. Future work will test a larger data base to determine with greater certitude the recognition capabilities of the system and to refine its tracking and labeling procedures.

Figure 5. Transformed limits of abnormality by statistical inference approach.

Acknowledgments

The authors wish to thank J. C. Robinson and K. R. Piety for providing the HFIR data, and F. L. Miller, Jr., for his assistance in implementing the Kolmogorov-Smirnov test.

References

1. J. A. Thie, *Reactor Noise*, Rowman and Littlefield, New York, 1963.
2. R. E. Uhrig, *Random Noise Techniques in Nuclear Reactor Systems*, The Ronald Press Co., New York, 1970.
3. D. N. Fry and J. C. Robinson, "Neutron Density Fluctuations as a Reactor Diagnostic Tool," in *Incipient Failure Diagnosis for Assuring Safety and Availability of Nuclear Power Plants*, CONF-671011, USAEC/DTIC (January 1968).
4. R. E. Uhrig, "State of the Art of Noise Analysis in Power Reactors," in *Amer. Nucl. Soc. Topical Meeting Water Reactor Safety, Salt Lake City, 1973*.
5. R. C. Gonzalez, D. N. Fry, and R. C. Kryter, "Results in the Application of Pattern Recognition Methods to Nuclear Reactor Core Component Surveillance," *IEEE Trans. Nucl. Sci.* 21(1), 750-756 (1974).
6. J. A. Thie, "Reactor-Noise Monitoring for Malfunctions," *Reactor Technol.* 14(4), 354-65 (1972).
7. D. N. Fry, "Experience in Reactor Malfunction Diagnosis Using On-Line Noise Analysis," *Nucl. Technol.* 10, 273-282 (March 1971).
8. D. N. Fry, R. C. Kryter, and J. C. Robinson, *Analysis of Neutron-Density Oscillations Resulting from Core Barrel Motion in the Palisades Nuclear Power Plant*, ORNL-TM-4570 (May 1974).
9. J. T. Tou and R. C. Gonzalez, *Pattern Recognition Principles*, Addison-Wesley Publishing Co., Reading, Mass., 1974.
10. R. C. Kryter, "Application of the Fast Fourier Transform Algorithm to On-Line Reactor Diagnosis," *IEEE Trans. Nucl. Sci.* 16(1), 210 (1969).
11. D. N. Lawley and A. E. Maxwell, *Factor Analysis as a Statistical Method*, Butterworth, London, England, 1963.
12. B. Noble, *Applied Linear Algebra*, Prentice-Hall, Englewood Cliffs, New Jersey, 1969.
13. J. D. Gibbons, *Nonparametric Statistical Inference*, McGraw-Hill Book Co., New York, 1971.
14. Z. W. Birnbaum, "Numerical Tabulation of the Distribution of Kolmogorov's Statistic for Finite Sample Size," *J. Amer. Stat. Assoc.* 47, 431 (1952).
15. R. R. Lemke, "On the Application of the Potential Function Method to Pattern Recognition and System Identification," Ph.D. Dissertation, Purdue University, Lafayette, Indiana, 1968.
16. M. A. Aizerman, E. M. Braverman, and L. I. Rozonner, "The Probability Problem of Pattern Recognition Learning and the Method of Potential Functions," *Automatika i Telemekhanika* 25(9), 1307-1323 (1964).
17. W. W. Cooley and P. R. Lohnes, *Multivariate Data Analysis*, John Wiley and Sons, New York, 1971.
18. H. S. Raulston, M. O. Pace and R. C. Gonzalez, "Adaptive Bayes Classifier for Remotely Sensed Data," *Proc Conf. Machine Processing of Remotely Sensed Data*, Purdue University, West Lafayette, Indiana, June, 1975, pp. 1A1-1A7.
19. H. W. Lilliefors, "On the Kolmogorov-Smirnov Test for Normality with Mean and Variance Unknown," *J. Amer. Stat. Assoc.* 62, 399-402 (1967).