Intrusion-Detection Systems

Daniel E. O'Leary

SYNOPSIS

Researchers recently have begun to develop intrusion-detection systems that detect potential and actual computer system intruders. These systems build and analyze profiles of expected user behavior, which are then compared to actual patterns to ascertain if a user is behaving as expected.

Some of the previous research in intrusion-detection systems is reviewed and a number of the characteristics and assumptions of those systems are discussed. Typically, those systems employ (among other devices) statistical analysis to aid in determining if an attempted intrusion is occurring or when an intrusion has already occurred. Unfortunately, many of the statistical methods make important distributional assumptions or use weak nonparametric methods. If the assumptions are inappropriate or the methods are weak in gathering knowledge from the data, then the results also will be weak. This paper proposes mitigating these problems by using alternative statistical methods.

Key Words: Intrusion-detection systems; Computer-intensive statistics; Continuous auditing; Security.

I. INTRODUCTION

A critical aspect of the security of a computer-based system is the ability to fend off potential intruders and identify actual intruders. Researchers such as Denning (1987) have responded to this problem by developing intrusion-detection systems, based primarily on statistical analysis, and secondarily on computer models of human detection expertise.

INTRUSIONS

Traditionally, passwords and other devices have been used to prevent intrusions. However, frequently passwords are discovered using another computer's own power to penetrate such security, successively iterating through a large volume of passwords.

Once a user has obtained access to a system, there have been few devices or methods developed to ascertain whether or not the system has been violated. In addition, the user or other users appropriately informed, could obtain future access simply by retracing their steps from previous illegal accesses. Thus, intrusions are not limited to single events.

INTRUSION-DETECTION SYSTEMS

The general approach used in intrusion-detection systems is to develop a set of expectations about a user, based on previous interaction with that user or users in general. Those expectations are employed to determine the existence of intruders. Typically, statistical analysis has been used

Daniel E. O'Leary is with the School of Accounting, University of Southern California.

The author would like to acknowledge the helpful comments of the three anonymous referees on two earlier versions of this paper. In addition, the author would like to acknowledge the extensive comments by Bill McCarthy on the previous version of this paper.

both to generate the set of expectations and to compare the expectations to behavior.

For example, the user of a specific password would generate a set of historical behaviors (signs on at a specific terminal, uses a specific printer, etc.). Thus, whenever someone successfully signs on using that password, the current behavior could be compared to the profile. If the behavior is different then there would be at least three options. First, the difference could be reported to a human expert who would then decide what action to take. Second, the system could take action, such as disconnecting the user and that password. Third, the system might take no action, but wait until some costly jeopardy (such as attempted access to a critical database or the generation of a report) is incurred, before choosing the first or second action.

PURPOSES OF THIS PAPER

Intrusion-detection systems have been developed for a number of applications. The next section of this paper summarizes the assumptions and characteristics of those systems, discusses some applications and provides extensions to other accounting and auditing applications.

Further, there are at least two limitations of the statistical analysis in some current intrusion-detection systems. First, statistical tools used in some systems provide very loose bounds. If the statistical tools are not effective in bounding expected behavior, then the systems are in danger of not recognizing an intruder. Second, some of the statistical tools make certain unjustified distributional assumptions about the data, which could provide misleading results. Thus, the following section investigates the statistical basis and limitations of such systems, with particular emphasis on the system discussed in Denning (1987). Then the next section analyzes the use of alternative statistical methods and summarizes their use in two examples. Finally, the last section summarizes the paper and discusses extensions.

II. INTRUSION-DETECTION SYSTEMS

The name "intrusion-detection system" is used since the first systems of the type investigated in this paper were for that purpose. It refers to those systems which are designed to monitor an agent's activity to determine if the agent is acting as expected or if the agent is exhibiting unexpected behavior. In a classic intrusion-detection system the agent is a computer system user. The objective of the intrusion-detection system is to determine if the user is an intruder or not.

However, in this paper, a broad-based view of intrusion-detection systems is taken. The agent might be an accountant providing entries to the system. In that situation, the intrusion-detection system might be concerned with determining if the agent or the entries were legitimate.

Intrusion-detection systems have concentrated on transaction processing applications. There are at least two basic sets of transactions of concern: transactions that occur with use of the system (the use is the transaction); and transactions that occur when an accounting transaction is entered into the system.

Denning (1987) and Tenor (1988) investigated systems where "uses" of the system and variables associated with those uses were transactions. Denning (1987) was concerned with the design of an intrusion-detection system that would be a part of the protection of an operating system. That system was to provide control over users accessing either the computer in general or specific applications. Tenor (1988) was concerned with a system designed to determine the existence of unusual accesses to a credit database. Halper et al. (1989) discussed a system designed to monitor accounting transactions. In their system the transaction base was the set of accounting transactions.

EXPERTISE IN INTRUSION-DETECTION SYSTEMS

Generating a set of expectations and then comparing behavior to expectations is consistent with behavior used by humans in everyday communication. For example, the expression "You are not yourself today" indicates that the initiator of the discussion has some base line behavior they are anticipating.

Typically intrusion-detection systems include the expertise of a human expert. Denning (1987) used expertise regarding the intrusion into computer systems. Tenor (1988) used knowledge acquired from a team of investigators as the basis of an intrusion-detection system for database protection. In both cases, that knowledge was supplemented by using statistical analysis of the data gathered in order to make inferences about the agents using the database.

Expertise is used to choose the behaviors to monitor (e.g., "time spent on the system") and the variables designed to capture those behaviors (e.g., "seconds in application #a"). Together the behaviors and variables provide a "profile" of the user. In addition, expertise is used to develop relationships between variables and to measure those relationships. For example, the correlation between two variables may be an important indicator of behavior.

USER PROFILES AND THEIR USE

The measures of behavior are summarized in user profiles. For example, the random variables investigated by Denning (1987) and Tenor (1988) included:

CPU time
number of accesses
date and time of access
user location
user identification
terminal or port ID
type of access.

A typical use of intrusion-detection systems is determining if a user is the legitimate owner of the password. In order to investigate that issue, the system could analyze the CPU time incurred by the user on a particular application or the difference of CPU time from some mean value. Current times being incurred could be compared to the user profile to determine if the user is the same person identified in the profile.

GENERAL-CASE OR USER-SPECIFIC PROFILES

A comparison of performance and user profiles is usually done in one of two ways. The first approach is to generate a single set of general expectations for all users. For example a "general-case" system designed to detect fraud, Lecot (1988, 19) used the rule "If age-of-account < x number of months then fraud <0.2>." Second, rather than general expectations, a profile could be generated for individuals with "user-specific" profiles.

DATA LIMITATIONS FOR INTRUSION-DETECTION SYSTEMS

There are at least two basic problems in the choice of data that are beyond the scope of this paper that are addressed elsewhere: data designed to mislead the system and changes in the processes being measured. First, there is concern that the system could be manipulated by a user behaving in a manner designed to confuse the system. For example, a user may establish high variance behavior to camouflage a planned intrusion. Thus, e.g., Tenor (1988) integrates a filter into the system in an attempt to ensure that only appropriate data is used to update the profiles. Second, if the data is nonstationary, then when the process changes, only the data relating to the new process should be used. As a result, system designers, such as Tenor (1988) periodically purge collected data from the intrusion-detection systems.

STATISTICAL BASIS OF INTRUSION-DETECTION SYSTEMS

Perhaps the most comprehensive discussion of the underlying statistical models used in intrusion-detection models is included in Denning (1987). In that paper five types of statistical methods are discussed: operational model, mean and standard deviation model, multivariate model, Markov process model, and time series model. In addition, Denning (1987, 226) was open to other alternatives, "Other statistical models can be considered. . . ." The analysis in this paper is concerned with the first three.

The operational model compares a new observation, x, against fixed limits. Those fixed limits are used for those statistical events where, based on prior experience with the system, they can be established. For example, it may be a policy that jobs that require a CPU time of y are investigated.

The mean and standard deviation model and the multivariate model both employ traditional parametric and nonparametric methods. As noted by Denning (1987, 225), "A new observation x_{n+1} is defined to be abnormal if it falls outside a confidence interval that is d standard deviations from the mean for some parameter d . . . ", that is, mean ± d * standard deviation. By assuming a normal distribution, confidence intervals can be developed. In addition, nonparametric approaches are used. For example, Chebyschev's inequality is used to establish confidence intervals, where (Denning 1987, 225) ". . . the probability of a value falling outside this interval is at most $1/d^2$ . . . ."

"The multivariate model is similar to the mean and standard deviation model except that it is based on correlations between two or more metrics" (Denning 1987, 225). Such an approach may provide "better discriminating power." Denning suggests that the relationship between CPU time and input/output (I/O) units is a typical example.

LIMITATIONS OF STATISTICAL APPROACHES

There are a number of problems with these approaches. The assumption of a normal distribution may not meet the needs of some situations. For example, it is likely most variables are truncated at 0 and skewed in one direction or the other. As a result, a priori distribution assumptions are likely to be incorrect. Although a normal distribution often is a robust substitute for the actual distribution, unfortunately, the user has no a priori idea how robust, unless the actual distribution is known. Further, the normal distribution may lead to entirely inappropriate conclusions.

In addition, as noted by a number of authors, for example, Kaplan (1982, 330), the nonparametric approach provided by Chebyschev's inequality ". . . is much too conservative, greatly reducing the power of the statistical procedure." Further, as Kaplan (1982, 330) also notes:

> There is little reduction in the information required to use Chebyschev's inequality, since the mean and the standard deviation of the error term still need to be estimated. Just how one would estimate the standard deviation of a distribution without having any knowledge of the underlying distribution is never made very clear by researchers advocating this procedure. It seems more reasonable to assume a simple parametric form, such as the normal, unless one has specific knowledge to the contrary. Chebyschev's inequality is mathematically interesting, but its practical significance . . . is questionable.

In addition, the evaluation of the statistical significance of the correlation for the multivariate model depends on a distribution assumption. Unfortunately, that assumption may be inappropriate and misleading.

Because of the limitations of assuming, e.g., a normal distribution, and because of the criticisms of Kaplan (1982) noted above, it is prudent to examine alternative approaches that do not make use of Chebyschev's inequality and do not require normality. If they assume normality, the reasonableness of a normality assumption is examined.

ALPHA LEVEL AND BETA LEVELS

Development of intrusion-detection systems requires addressing the potential for error by the system. These errors are commonly called "alpha" and "beta" levels or "type I" and "type II" errors. Typically, there is a trade-off of the costs and benefits associated with the different error types. In the case of the system for credit files discussed in Tenor (1988), there was little cost to have an intruder browse a set of files, since no changes could be made to the files and users were charged for use of the system. Thus, even if there was an intruder, the firm would receive revenues for the use of the system. Similarly, there was substantial cost to the firm if a legitimate user was unable to access the system since that would lead to lost revenues from nonuse of the database.

III. ALTERNATIVE STATISTICAL APPROACHES

Two examples are used to illustrate alternative statistical methodologies that minimize the limitations of the statistical approaches employed in intrusion-detection systems. In the first example there is some initial data on usage (base case) by a particular user, and later, additional data (recent data) is generated from system use. The example data is given in Table 1. The data are to be evaluated from at least two perspectives:

a. "In the base case data, are any of the observations of behavior not consistent with the other data, and thus, possibly an indication of an intrusion?" or "Are there any outliers of behavior?"

b. "Did the same user generate both sets of data?" or "Is the data in the base set from the same distribu-

TABLE 1
DATA FROM CONTROL INVESTIGATION

Set A—Base Data	
1	3.0
2	0.0
3	−0.5
4	−1.0
5	−1.0
6	4.0
7	2.0
8	−0.5
9	0.0
10	0.0
11	2.0
12	−0.5
13	−1.0
14	−1.0
15	0.0
16	0.0
17	−1.0
18	−0.5
19	−1.0
20	0.0
21	0.0
22	0.0
23	3.0
24	0.0
25	0.0
26	−1.0
27	−2.0
28	3.0
29	3.0
30	3.0
31	0.0
32	0.0
33	−1.0
34	0.0
35	0.0
36	−1.0
37	0.0

Set B—New Data	
1	2.0
2	0.0
3	0.0
4	2.0
5	1.5
6	0.0
7	0.0
8	−1.0
9	−1.0
10	1.0
11	1.0
12	2.0
13	−1.0

Source: Dungan (1983).

> tion as the data in the recent observations?''

In the second example, it is assumed that there are data for CPU time and I/O units, for a user, and for a number of system uses. The question of concern is whether there is a relationship between those two variables. The correlation is used as the basis of analysis. Statistical significance is used to determine if there is a relationship that might be investigated. Table 2 shows the example data.

Consider the ''base case'' data in the first example. Assume there is concern about an observation of 4. The mean of the data is .27 and the standard deviation is 1.473. As a result, Chebyschev's inequality becomes

$$\Pr(|4 - .27| > 3.73) < 1.472/3.732. \tag{1}$$

Thus,

$$\Pr(|4 - .27| > 3.73) < .16. \tag{2}$$

Alternatively, if we assume a normal distribution, we can develop a .95 confidence interval

$$.27 + 1.96 * 1.473/(\sqrt{37}) = .27 + .476 \tag{3}$$

Another approach is to use nonparametric, univariate outlier analysis (Velleman and Hoaglin 1981). Outliers are those observations that stand apart from the rest of the observations.

Typically, outliers are observations that do not derive from the same distribution as the other observations. Thus, outlier analysis is consistent with the interest in identifying those observations that seem to be different than the others.

One of the primary tools of outlier analysis are boxplots. Boxplots focus attention on extreme values by providing a means to determine those observations that are a ''sufficient distance'' from the median. If an observation is a sufficient distance from the middle of the data set, then it is likely that it does not come from the same distribution as the rest of the data.

TABLE 2
EXAMPLE: CPU TIME MEASURE AND I/O UNITS

CPU Time Measure	I/O Units
54	15
55	5
15	2
41	17
54	4
27	3
28	1
37	1
28	0
32	0
18	0
19	0
17	0
53	0
21	0
23	20
42	9
20	0
19	0
40	0
13	1
20	6
28	1
27	3
30	1
52	0
102	6
14	0
19	0
54	0
26	0
35	1
62	14
19	0
46	7
59	0
23	2
38	9
23	4
16	0
15	0
16	0
13	9
15	10

The boxplot defines that distance as follows.

The median of the set of observations splits the data into two sets. The median of each of those two data sets forms what is called the "hinge." The difference between each of those hinges is called the H-spread. The hinges and the H-spread are used to calculate inner and outer fences that are used to point to those values that are likely to be outliers:

Lower Inner Fence = Lower Hinge − 1.5 * H-Spread
Upper Inner Fence = Upper Hinge + 1.5 * H-Spread
Lower Outer Fence = Lower Hinge − 3.0 * H-Spread
Upper Outer Fence = Upper Hinge + 3.0 * H-Spread

Values between, respectively, the lower inner and lower outer and the upper inner and upper outer fences are "possible outliers." Values outside, respectively, the lower outer and the upper outer fences are "possible far outliers."

This approach eliminates the direct dependence on a distribution assumption. It replaces reliance on the mean and standard deviation with use of the median and H-spread. The amounts 1.5 and 3.0 are roughly analogous to the z values used in the analysis of normal distributions.

Exhibit 1 illustrates a boxplot. Pictorially, the median is denoted with a "+." The hinges form the vertical outer edges of the "box." The box contains fifty percent of the data. Possible outliers are denoted with a "*," while possible far outliers are denoted with a "0." It can be seen that the value of 4 is classified as a "far-outlier." This result is consistent with making a normal distribution assumption. However, Chebyschev's inequality provided a bound of a probability of .16. Thus, this provides evidence that if we had used Chebyschev's inequality, then that bound provides a loose estimate. As a result, actual intrusions might be missed because of the looseness of the bound.

TESTING THE NORMALITY OF THE DATA

Another assumption that is made in intrusion-detection systems is that of normality or of a t distribution. In some cases there may be reason to suspect that the data is from a normal distribution or other distribution. Such an assertion may be based on previous experience with the data or from an understanding of the process that generated the data. However, in any case, it is desirable to be able to test for normality.

Normal probability plots are a computer-based statistical approach to identifying whether a set of data are likely to be from a normal distribution. As noted by Filliben (1975, 111), "The . . . test statistic is conceptually simple, is computationally convenient, and is readily extendable to testing non-normal distributional hypotheses." Filliben (1975) empirically generated normal probability correlation coefficients for various percent points and sample sizes, using random normal deviates as input. A similar approach can be used for other distributions.

The results of the normal probability plot for the base data (normal z scores) in the example are given in Exhibit 2. The correlation of those sample normal scores and the normal distribution is 0.934, which falls between the 5th and 0th percentile of the null distribution generated by Filliben (1975). Elimination of the far outlier does not change that result significantly. The results indicate that there is no evidence that the data comes from a normal distribution. Thus, use of a normal distribution could result in misleading inferences.

COMPUTER-INTENSIVE APPROACHES

Since it is apparent that the use of a normal distribution can be inappropriate, there is a need for an alternative approach. Computer-intensive approaches use the computer itself to generate a distribution for the test statistic, such as the mean or standard deviation. Computer-intensive

EXHIBIT 1

Boxplot

```
             ---------
     --------I       +--------       *       *       O
             --------

  +---------+---------+---------+---------+---------+------
-2.4      -1.2       0.0       1.2       2.4       3.6
```

Note:
* = possible outliers
O = possible far outliers
+ = median

methods get their name because as noted by Noreen (1986, 5) ". . . they require recomputing the test statistic for . . . data sets many (typically 100 to 1,000) times." The basic assumption of computer-intensive methods is that the sample contains all the information that is known about the distribution (Efron 1979). As noted by Noreen (1986, 5), "The significance of virtually any well-defined test statistic can be accessed using one of these methods."

Further, some computer-intensive approaches, such as the primary approach discussed here, are valid for nonrandom samples, as well as for random samples (Noreen 1986). In addition, as noted by Noreen (1986, 7), ". . . if a conventional test is feasible it will often yield virtually the same significance level as the computer-intensive test. . . ." As a result, there is little to lose and much to gain with the use of a computer-intensive approach.

RANDOMIZATION: A TEST OF DIFFERENCES IN MEANS

The computer-intensive approach discussed in this paper is randomization (Edgington 1980 and Noreen 1986). Generally, randomization is used to test that one set of random variables is distributed independently of some other set of random variables (Noreen 1986). Thus, randomization can be used to examine the possibility that recent user behavior, characterized in the observations of some random variable (x_{n+1}, ..., x_{n+m}), is not the same as previous behavior (say a user profile), characterized by the observations (x_1, ..., x_n). Differences in the two distributions might signal the presence of an intruder.

Computationally, randomization would be used in the following manner.

> ALGORITHM: The absolute value of the difference in the means between two samples is computed. Then the original n observations and the subsequent m observations would be stacked together. Then the entire set of observations would be shuffled. The absolute value of the difference between the means of the shuffled data for the first n and the last m observations would be computed and stored as an observation in a distribution of such absolute differences. This would be done a large number

EXHIBIT 2

```
                    Normal Probability Plot of Data

         -
      4.0+                                                     *
         -
  Data   -                                           5
         -
      2.0+                                      2
         -
         -
         -
         -
      0.0+                               +
                               4
         -              9
         -
         -
     -2.0+   *
         ---------+---------+---------+---------+---------+--------Normal
              -1.60     -0.80      0.00      0.80      1.60        Scores

            Correlation of Data and Normal Scores = 0.934
```

> of times, e.g., one hundred or more. The observations would be ranked, based on the size (smallest to largest) of the absolute differences so that using this process, an entire distribution of absolute differences could be computed. Then the absolute value of the difference between the mean for the original and subsequent observations would be placed in distribution in the appropriate position to determine its relative probability.

The results for the first example, with 200 shuffles, are contained in Exhibit 3. The actual absolute value of the difference in the means of the base case and subsequent case was .2297. This is between the 59th and 60th ordered observations, which is at the 70% level (1-(60/200)). Since it is likely to have obtained such an absolute value, there is no reason to reject the null hypothesis that the two sets of observations are not from different distributions.

MULTIVARIATE ANALYSIS: DISTRIBUTIONS OF CORRELATIONS

Multivariate investigations could be employed for those cases where the system's behavior is a function of more than one variable. For example, Denning (1987) suggested a relationship, presumably positive, existing between CPU time and I/O

EXHIBIT 3

Absolute Differences Between Means

```
  :
  :
  :             .
. : . .(actual)  :
: : : :0.2297: . : :
: : : :*     : : : : . :
: : : : .  . .:: : : : : :    :
: : : : .. : ::: : : : : : : : : .
: : : : :..: ::: : : : : : : : : : .
: : : : ::::.::: : : : : : : : : :.   :  .    . .
+---------+---------+---------+---------+---------+-------diff
0.00      0.25      0.50      0.75      1.00      1.25
```

Number of randomizations = 200

(diff is the difference between the means of two sets of randomly shuffled data, one of size 37 the other of size 13)

units be investigated using correlation analysis.

The relationship between CPU time and I/O can be investigated in a number of different ways. However, for purposes of exposition assume that over time, for a particular user and particular software, a set of data point pairs (CPU,I/O) is available. One null hypothesis is that CPU is independent of I/O.

Randomization could be used to develop a distribution of correlations to determine if there is an unusually high relationship between the two variables. Assume that there are n sets of pairs of, e.g., CPU times (x vector) and I/O Units (y vector), (x_i, y_i), for $i = 1, \ldots, n$. Randomization can be used to generate a distribution of correlations to determine if the original correlation is unusually large.

ALGORITHM: Compute the correlation between the x and y vectors, while in their original order—this is the original correlation. Then the randomization process can be used. Shuffle the y vector while holding the x vector stationary. Compute the absolute value of the correlation coefficient. Do that randomization process a large number of times. Rank the correlation values from smallest to largest. The location on the distribution will indicate whether the original relationship is in the tail, and thus unusual and deserving of further investigation.

Using the second example, one hundred correlation coefficients were developed. It was found that the original correlation coefficient was significant at the .04 level. If a normal distribution were used it would be significant at about the same level.

IMPLICATIONS OF COMPUTER-INTENSIVE STATISTICAL TECHNIQUES FOR THE DESIGN OF ACCOUNTING INFORMATION SYSTEMS

This section has found that some data sets that would be processed by accounting information systems do not necessarily conform to existing distributions. Further, there can be substantial effort in testing to determine whether or not data conform to any of a large set of distributions (e.g., normal).

Computer-intensive approaches mitigate the search process for matching existing distributions to data and provide an approach that allows the development of distributions that meet the specific needs of the application. It was seen that conventional tests may yield the same significance level. As a result, the computer-intensive approach can be quite helpful in those situations where the systems designer/developer needs to have the system perform statistical analysis.

IV. SUMMARY AND EXTENSIONS

This paper investigates the basic nature of intrusion-detection systems. Many of the implicit characteristics of these systems were elicited and discussed in detail. Intrusion-detection systems have many applications in accounting, including monitoring accounting agents and transactions for expected behavior.

Intrusion-detection systems employ profiles of user behavior in transaction processing systems. Profiles are compared to actual performance to determine whether the user is exhibiting expected behavior. Intrusion-detection systems developed to-date employ statistical methods or results to develop profiles and to determine when behavior is deviating from those profiles. Computer-intensive statistics were suggested as a means to tighten bounds and eliminate use of inappropriate distributions.

Other approaches can be integrated into those presented in this paper and the techniques discussed in the original paper by Denning (1987). For example, if the system designer is concerned with choosing a distribution that meets the needs of the data, then there are other techniques to assess the possibility of, e.g., the normality of the data. For example, Ewart et al. (1978) discuss the Kolmogorov-Smirnov test which can be used to investigate whether the data is any of a number of distributions. Further, the probability plot approach can be extended to other distributions (Filliben 1975).

In addition, bootstrap methods (Efron 1979; Efron and Gong 1983) may be used to develop other tests for intrusion-detection investigations. For example, rather than using the outlier analysis approach or confidence intervals discussed above, bootstrap resampling (e.g., Noreen 1986) could be used to analyze whether a number of recent observations come from the same distribution as the base case data.

Finally, control charts and cost variance approaches (e.g., Kaplan 1982) could be used to assess some of these same problems. Control limits would be used to provide bounds on normal behavior. If an observation exceeded those bounds then that observation would be regarded as abnormal.

REFERENCES

Denning, D. 1987. An intrusion-detection model. *IEEE Transactions on Software Engineering* SE 13 (No. 2, February): 222–232.

Dungan, C. 1983. A model of audit judgment in the form of an expert system. Ph.D. Dissertation, University of Illinois.

Edgington, E. 1980. *Randomization Tests.* New York: Marcel Dekker, Inc.

Efron, B. 1979. Bootstrap methods: Another look at the jackknife. *Annals of Statistics* 7: 1–26.

———, and G. Gong, 1983. A leisurely look at the bootstrap, jackknife and cross-validation. *The American Statistician* 37 (February): 36–48.

Ewart, P., J. Ford, and C. Y. Lin. 1978. *Applied Managerial Statistics*. Englewood Cliffs, New Jersey: Prentice-Hall.

Filliben, J. 1975. The probability plot correlation coefficient test for normality. *Technometrics* 17 (No. 1, February): 111–117.

Halper, F., J. Snively, and M. Vasarhelyi. 1989. The continuous audit of on-line systems. Unpublished paper presented at the Second International Symposium on Expert Systems in Business, Finance and Accounting (November).

Kaplan, R. 1982. *Advanced Management Accounting*. Englewood Cliffs, Prentice-Hall.

Lecot, K. 1988. Using expert systems in banking: The case of fraud detection and prevention. *Expert Systems Review* 1 (3, June): 17–20.

Neter, J., W. Wasserman, and M. Kutner. 1985. *Applied Linear Statistical Models*. Homewood, Illinois: Irwin.

Noreen, E. 1986. An introduction to testing hypotheses using computer intensive methods. Unpublished paper, Graduate School of Business, University of Washington (November). Published as *Computer Intensive Methods for Testing Hypotheses: An Introduction*. New York: John Wiley, 1989.

Tenor, W. 1988. Expert systems for computer security. *Expert Systems Review* 1 (2): 3–6.

Velleman, P. and D. Hoaglin. 1981. *Applications, Basics and Computing of Exploratory Data Analysis*. Belmont, CA: Duxbury Press.