# Artificial Neural Networks for Predicting Failure to Survive Following In-Hospital Cardiopulmonary Resuscitation

Mark H. Ebell, MD
*Detroit, Michigan*

*Background.* Neural networks are an artificial intelligence technique that uses a set of nonlinear equations to mimic the neuronal connections of biological systems. They have been shown to be useful for pattern recognition and outcome prediction applications, and have the potential to bring artificial intelligence techniques to the personal computers of practicing physicians, assisting them with a variety of medical decisions. It is proposed that such an artificial neural network can be trained, using information available at the time of admission to the hospital, to predict failure to survive following in-hospital cardiopulmonary resuscitation (CPR).

*Methods.* The age, sex, heart rate, and 21 other clinical variables were collected on a consecutive series of 218 adult patients undergoing CPR at a 295-bed public acute-care hospital. The data set was divided into two groups. A neural network was trained to predict failure to survive to discharge following CPR, using one group as the training set and the other as the testing set. The procedure was then reversed, and the results of the two networks were combined to form an aggregate network.

*Results.* The trained aggregate neural network had a sensitivity of 52.1% and a positive predictive value of 97% for the prediction of failure to survive following CPR. The relative risk of actually failing to survive to discharge following CPR for a patient predicted not to survive was 11.3 (95% CI 3.3 to 38.2).

*Conclusions.* Predicting failure to survive following CPR is but one possible application of neural network technology. It demonstrates how this technique can assist physicians in medical decision making. Future work should attempt to improve the positive predictive value of the neural network, to consider combining it with an expert system, and to compare it with other predictive tools. Once validated, the network can be distributed as a separate application for use by practicing physicians.

*Key words.* Neural networks (computer); resuscitation orders; cardiopulmonary resuscitation, survival.
*J Fam Pract 1993; 36:297-303.*

The rate of survival to discharge following in-hospital cardiopulmonary resuscitation (CPR) has been reported to vary from 7%[1] to 24%[2] in published studies. In a recent meta-analysis of 14 studies,[3] this author found an average survival rate of 13.5%. While the analysis identified several patient subgroups with a low rate of survival to discharge following CPR, only patients with metastatic cancer were found to have a rate of survival to discharge of less than 1%.[3] The identification at hospital admission of additional patients with a negligible rate of survival to discharge following CPR would provide important prognostic information for discussions about do-not-resuscitate (DNR) orders.

George and colleagues[2] have proposed a Pre-Arrest Morbidity (PAM) index to prospectively identify patients with a negligible rate of survival to discharge following CPR. The index consists of 15 variables assigned 0, 1, or 3 points, with a score greater than 8 associated with failure to survive. Based on the results of the previously mentioned meta-analysis, the author has proposed a modified PAM index, hereafter referred to as the Prognosis After Resuscitation (PAR) score. It is easier to use than the original predictive tool and may be more specific, as it eliminates 7 of the original 15 variables that were found to be unrelated to the survival rate.

*Submitted, revised, October 29, 1993.*

*From the Department of Family Medicine, Wayne State University, Detroit, Michigan. Requests for reprints should be addressed to Mark Ebell, MD, Department of Medicine, Wayne State University, 4201 St Antoine UHC-4J, Detroit, MI 48201.*


ISSN 0094-3509

While neither index has been tested on an independent data set, preliminary studies suggest that each is able to identify as nonresponders a subgroup of approximately 20% to 25% of the patients who eventually undergo CPR.[2,3]

Since the prediction of failure to survive using clinical data is essentially a matter of pattern recognition, artificial intelligence techniques that are able to analyze large amounts of data may be useful. Artificial neural networks are one such technique. First proposed in 1947,[4] neural networks use nonlinear equations to mimic the connections between neurons found in biological systems. Unlike traditional expert systems, they are not rule-based and must be trained on a set of known data.

Artificial neural networks have been successfully applied to a variety of medical questions, such as the diagnosis of myocardial infarctions,[5,6] the identification of cancerous cells in breast aspirates,[7] and the interpretation of chest radiographs.[8,9] They have the potential to bring artificial intelligence to the personal computers of practicing physicians, and assist them with a wide variety of medical decisions. It is the goal of this study to develop an artificial neural network that will, at the time of admission to the hospital, identify patients with a negligible rate of survival to discharge following CPR.

## Methods

### *Artificial Neural Networks*

The basis of learning for artificial neural networks and living systems alike is the theory that when two neurons are simultaneously excited, the connection between them should be strengthened. For example, if an input neuron fires when it sees something red, and an output neuron fires to signal that an apple has been identified, repeatedly observing that most apples are red should strengthen the connection between these two neurons.[10]

The actual architecture of an artificial neural network consists, at a minimum, of a layer of input units or neurons and at least one output neuron. Optionally, there can also be one or more intermediate or "hidden" layers of neurons. In the most common schema, each neuron in one layer is connected to each neuron in the layer above it (Figure 1). Neural networks can employ a feed-back or feed-forward architecture, a term describing the flow of information in the trained network. Most networks in use today are the feed-forward type, as they may have a better capacity for learning and are faster to use once trained.

Each neuron in such a network receives input from all of the neurons in layers below it, with each input

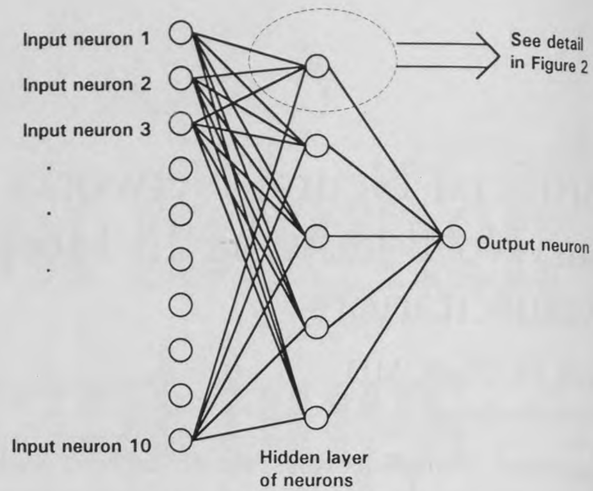


Figure 1. Sample neural network. This network has 10 input neurons, one hidden layer with 5 neurons, and 1 output neuron. For the sake of clarity, only the connections for 4 of the input neurons are shown.

consisting of a signal multiplied by a weight. All of the inputs (excitatory and inhibitory) to a given neuron are summed, and a transfer function determines the strength of the resulting output signal (Figure 2). This signal is then passed to all of the neurons in the next layer of the network. The most common neuron transfer functions are linear, threshold, and sigmoid. If linear, the strength of the output signal is proportional to the sum of the input signals, whereas a threshold function exhibits all-or-nothing behavior. A sigmoid function behaves in a nonlinear fashion and is the most widely used. As in most living systems, a small increase in the input signal may lead to a large increase in the output signal, making the sigmoid function especially appropriate for biomedical applications.[11]

At the start of the training process, random weights are assigned to the connections between neurons. Next, the network tries to predict the desired outcome, and compares the network output with the known result. Based on the difference between the calculated and actual results, the weights are adjusted, and the network tries again. This technique is known as back-propagation, and is the most commonly used training paradigm. After hundreds or thousands of such iterations, a set of weights is found that accurately predicts all or most of the outcomes in the training group. This set of weights, and the equations connecting the individual neurons, make up the neural network. Once created, such a network must be tested on an independent set of data to assess its utility. While training a network may take hours on a desktop computer, once trained, the network can make a prediction based on a set of unknown data in milliseconds.

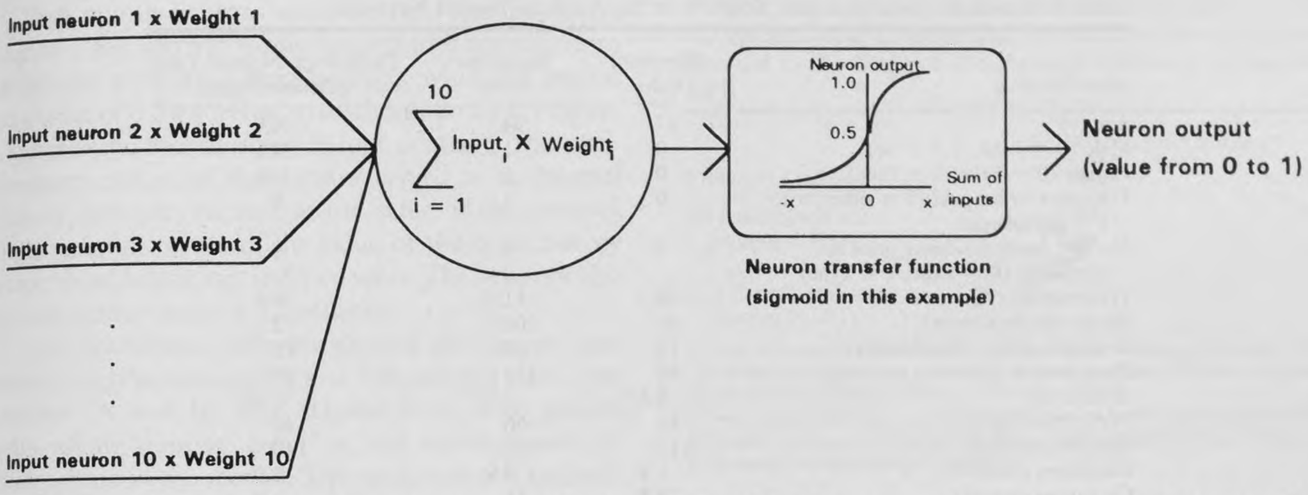


Figure 2. Function of an individual neuron in a neural network. The value of each input neuron is usually normalized to a range from −1 to 1, then multiplied by a weight. These products are then summed, and passed to a transfer function that determines the value of the neuron's output.

As a clinical example of a neural network application, the training set could be clinical data about a group of 100 nursing home patients and their rate of survival following surgery. What is traditionally called the dependent variable (survival after surgery in this case) becomes the output neuron, while the clinical variables are assigned to input neurons. If desired, one or more hidden layers can be added to the network before training. Once trained, the network could be tested on a second set of 100 patients to determine its sensitivity and positive predictive value.

### *Subjects*

A retrospective chart review was conducted using the medical records of all inpatients over age 16 years who underwent CPR on the general wards or in the intensive care units at Athens Regional Medical Center between January 1, 1989, and July 31, 1991. A retrospective design was chosen because only approximately 1% of all patients admitted to Athens Regional Medical Center undergo CPR. A prospective study would require the enrollment of over 20,000 patients in the study in order to identify 218 who had undergone CPR, which is clearly impractical.

Athens Regional Medical Center is a public acute-care hospital with 295 beds, including 24 intensive care unit beds. Cardiopulmonary resuscitation was minimally defined as the application of chest compressions and rescue breathing by the "Code 99" team; CPR was directed by the attending physician or emergency room physician, and assisted by ancillary staff on the unit. Data for CPR performed in the emergency department, catheterization laboratory, or operating rooms were not available for inclusion.

A review of the CPR log forms completed by nursing staff following each Code 99 event identified 184 potential subjects. Review of the medical records of these subjects revealed that 6 had either not received chest compressions or rescue breathing or were under 17 years of age. These 6 subjects were excluded, leaving 178 subjects identified from CPR logs. The records of 57 additional patients who were not identified on a CPR log form but who were charged for use of the "code cart" were reviewed. Forty of these patients had undergone CPR and were also included in the study, resulting in a final study population of 218 patients. As the outcome measured by the study was survival to discharge from the hospital, patients undergoing multiple resuscitative efforts were included only once in the data set.

### *Variables Measured*

The outcome variable measured in this study was each patient's discharge disposition (dead or alive). Since the purpose of the study was development of a predictive tool that could be used at the time of hospital admission, only data that were available within the first 24 hours of hospitalization were used as input variables for the neural network. The input variables assigned to input neurons in the neural network are summarized in Table 1. The patient's age, sex, and reason for admission to the hospital were obtained from a review of the admitting physician's history-taking and physical examination of the patient. The presence of an active diagnosis of cancer or metastatic cancer was noted; a diagnosis of basal cell

Table 1. Variables Used as Input Neurons in the Artificial Neural Network

| Input Neurons | Minimum Value | Maximum Value | Physiologic Normal Value (Default Value) |
|---|---|---|---|
| Age (y) | 17 | 94 | NA |
| Male (0 = False, 1 = True) | 0 | 1 | NA |
| Female (0 = False, 1 = True) | 0 | 1 | NA |
| Transport to hospital (0 = other, 1 = ambulance) | 0 | 1 | 0 |
| Nursing home residence prior to admission (0 = False, 1 = True) | 0 | 1 | 0 |
| Temperature (°C) | 28.3 | 41.8 | 38.6 |
| Heart rate (beats/min) | 40 | 200 | 72 |
| Respiratory rate (breaths/min) | 10 | 60 | 16 |
| Mean arterial pressure (mm Hg) | 30 | 173 | 90 |
| Arterial pH | 6.67 | 7.59 | 7.40 |
| $Pco_2$ (mm Hg) | 23 | 66 | 40 |
| Sodium (mmol/L) | 112 | 165 | 140 |
| Potassium (mmol/L) | 1.4 | 6.9 | 4.0 |
| Creatinine (mg/dL) | 0.3 | 14 | 1.0 |
| Bicarbonate (mmol/L) | 12 | 46 | 24 |
| Hematocrit (%) | 9.5 | 66.1 | 45 |
| White blood count (1000 cells/$mm^3$) | 0.3 | 50 | 6.0 |
| Glasgow Coma Scale score (0–15) | 3 | 15 | 15 |
| Chronic illness (0 = False, 1 = True)* | 0 | 1 | 0 |
| Primary admitting diagnoses (0 = False, 1 = True): | | | |
| Myocardial infarction | 0 | 1 | 0 |
| Coronary artery disease† | 0 | 1 | 0 |
| Cancer | 0 | 1 | 0 |
| Metastatic cancer | 0 | 1 | 0 |
| Pneumonia | 0 | 1 | 0 |

**See Appendix for a detailed definition.*
*†Coronary artery disease is defined as an admitting diagnosis of myocardial infarction, unstable angina, rule out myocardial infarction, arrhythmia or congestive heart failure.*
*NA denotes not applicable.*

carcinoma or an inactive diagnosis (ie, colectomy for colon cancer 10 years ago, now asymptomatic) was not included.

The patient's place of residence before admission and mode of transport to the hospital were noted. The Glasgow Coma Scale (GCS) score was recorded on arrival in the emergency department in 58.7% of the patients in the study. For the remaining patients, the GCS score was derived from the neurologic examination, if one had been performed by the patient's physician at admission. For example, a description of a patient as "alert, appropriate and oriented times three" was considered consistent with a GCS score of 15. An estimate of chronic health status adapted from the Acute Physiologic and Chronic Health Evaluation (APACHE) II[12] was recorded (Appendix).

Physical and laboratory values were collected and recorded at the time of admission for the majority of patients. The first value recorded within the first 24 hours of admission was used. Three patients did not have a white blood count or hematocrit documented in their records, and two others had not had a sodium, potassium, or creatinine recorded within 24 hours of admission.

### *Neural Network Architecture*

An artificial neural network was created using the BrainMaker version 2.3 software.[13] The network utilized a sigmoid transfer function, a learning rate of 1.0, and a training tolerance of 0.1. The learning rate determines how large a correction should be made when the network detects an error, while the training tolerance determines how close the network output must be to the correct answer to be considered correct. A value of .1 means that the output must be within 10% of the correct answer to be considered correct during training.

The prearrest variables were mapped to 25 input neurons as shown in Table 1. In addition, the minimum and maximum possible input values are shown, along with the physiologic normal value. In the event that data were missing for a particular variable, the physiologic normal value was used as a default value. This is consistent with the practice used in other types of predictive indices such as the APACHE II and III scores, where missing data are also assumed to be normal. An alternative technique would substitute a null value (ie, 0) if data were missing, but was believed to be more likely to skew the weights calculated for that neuron.

Two output neurons were created. The first had a value of 1 for survival to discharge and 0 for failure to survive (the network output neuron), while the second had a value of 0 for survival to discharge and 1 for failure to survive (the confirmatory output neuron). The complementary nature of the outputs served as an internal check of certainty, because as the value of the network output neuron increased, the value of the confirmatory neuron should decrease, and vice versa. The network also had one hidden layer of 12 neurons.

The BrainMaker software divided the data set into two randomly selected groups of 109 patients each. Two networks, A and B, were trained using each patient group as the training group in one instance and the testing group in the second. The same network architecture (number of neurons, transfer function, and so forth) was used for each network. Finally, the results of networks A and B were summed to assess the performance of the aggregate network. Such a "crossover" design provides a more accurate assessment of network performance, as it prevents the training of a network that is designed to perform especially well on the specific testing group used in the study.

A microcomputer based on the Intel 386DX chip operating at 25 MHz (without a math co-processor) was used to train the network. Training was judged to be complete when all of the patients in the training group were correctly segregated by the network. A value of <.001 for the network output neuron was used to identify patients predicted not to survive to discharge following CPR. This figure was chosen because it had the best ability to discriminate between survivors and nonsurvivors. The sensitivity, positive predictive value, and relative risk for a network output value of <.001 as a predictor of failure to survive were calculated.

Receiver-operating characteristic (ROC) curves were calculated for both the neural network and the PAR score described earlier, using the data collected above. An ROC curve plots the false-positive rate (1 − specificity) on the *x*-axis and the true-positive rate (sensitivity) on the *y*-axis. The closer an ROC curve is to the upper left corner of the graph and the larger the area under the curve, the more accurate the predictive instrument, because the true-positive rate approaches 1.0 as the false-positive rate approaches 0.[14]

## Results

Neural network A required only 876 iterations, while network B required 1127 iterations. For each case in which the value of the network output neuron was <.001, the value for the confirmatory neuron was >.900. For network A, of the 64 patients who were identified as having a network output <.001, 63 did not survive to discharge. The sensitivity for the prediction of failure to survive of patients in network A was 68.5%, and the positive predictive value was 98.4%. In network B, of the 35 patients who were identified as having a network output <.001, 33 did not survive to discharge. For network B, the sensitivity was 35.9%, while the positive predictive value was 94.3%.

The output from neural networks A and B was summed to create the aggregate network output. The results are shown in Table 2 in the form of a 2 × 2 table. A total of 99 patients had a network output <.001 (predicted not to survive), of whom 96 actually failed to survive to discharge following CPR. Thus, the aggregate network had a sensitivity of 52.2% and a positive predictive value of 97.0% for the prediction of failure to survive. The relative risk of failure to survive to discharge following CPR for a patient with a neural network output of <.001 was 11.3 (95% CI 3.3 to 38.2).

Table 2. Actual Outcomes and Predicted Outcomes Compared Using an Aggregate Neural Network to Predict Survival to Discharge of 218 Patients Following In-Hospital CPR

| | Actual Outcome* | |
|---|---|---|
| Predicted Outcome† | Died | Survived |
| 99 Patients will die (output <.001) | 96 | 3‡ |
| 119 Patients will survive (output ≥.001) | 88 | 31 |

**Actual outcome refers to whether the patient survived to discharge following CPR.*
*†Predicted outcome refers to the neural network's prediction of whether the patient would survive to discharge following CPR.*
*‡Three patients who survived were incorrectly classified by the neural network. Overall the network had a sensitivity of 52.2% and a positive predictive value of 97%.*
*CPR denotes cardiopulmonary resuscitation.*

The ROC curves for the aggregate neural network and the PAR score are shown in Figure 3. The neural network is closer to the upper left corner of the graph over most of its course, indicating that it is a better predictive instrument than the PAR score. Using the Wilcoxon statistic, the area under the neural network curve was estimated to be .765 (standard error [SE] = .048), as opposed to an area of .717 (SE = .051) under the PAR score curve.[15] This larger area under the curve is consistent with better performance for the neural network.

Three of the patients who underwent CPR and survived to discharge were misclassified by the neural network. These three patients were a 47-year-old man with diabetic ketoacidosis and a witnessed episode of symptomatic bradycardia; a 55-year-old man with pneumonia and a witnessed respiratory arrest; and a 71-year-

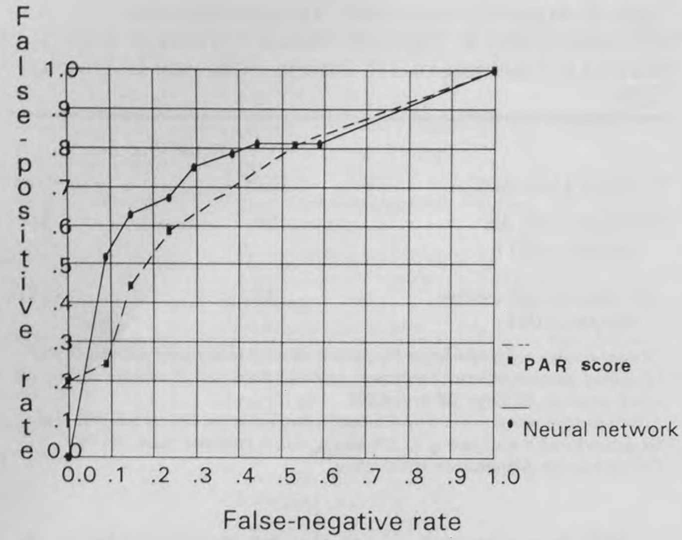


Figure 3. Receiver-operating characteristic curves for the aggregate neural network and the PAR score are shown. Better predictive instruments are characterized by curves that are closer to the upper left corner of the graph (greater sensitivity or true-positive rate, and a lower false-positive rate).

old woman recovering from peripheral vascular surgery with an unwitnessed episode of ventricular fibrillation. Each of these patients was documented as having received full resuscitative efforts, including chest compressions and artificial ventilation.

## Discussion

The trained aggregate neural network was able to identify 52% of the patients who would eventually undergo CPR and not survive to discharge based on information at the time of admission. This reflects a positive predictive value of 97%. While not a perfect predictor of failure to survive, an artificial neural network can provide important prognostic information for clinicians and their patients when they discuss DNR orders. Such information will not only help guide decision making, but in the event that nonsurvival is predicted, it will reassure patients and their families that a decision to write a DNR order is appropriate. In addition, the identification of poor responders to CPR at the time of admission to the hospital may stimulate discussion of DNR orders with these high-risk patients, and prevent needless morbidity and the misapplication of medical resources.

In addition, comparison of the ROC curves shows that the neural network is a better instrument for the prediction of failure to survive to discharge following CPR than the PAR score. Future work should focus on improving the positive predictive value of the neural network. To be useful in a clinical setting, the positive predictive value should approach 100%. Since neural networks learn by identifying patterns in the training set, a larger and more representative training set generally leads to the development of a more accurate neural network model.[4] In this study, two training groups of 109 patients each were used. Future studies should use larger training groups consisting of 200 to 300 patients.

Network accuracy could also be enhanced by including variables that reflect additional elements of the history and physical examination, such as the presence of cardiomegaly or a third heart sound, and including patients from a variety of acute-care facilities. Once an accurate network has been developed, it must be tested in a prospective fashion in a variety of patient populations to ensure its broad applicability. In addition, research is underway that directly compares artificial neural networks with other predictive tools such as the PAM index, the modified PAR score proposed by the author,[3] and the APACHE III score,[16] using the same set of patient data.

Another approach to improving the predictive power of the neural network could involve hybrids between rule-based expert systems and neural networks, known as "expert networks."[17] Such a system could, for example, classify all patients with metastatic cancer as nonresponders to CPR, and young trauma patients as persons who should always undergo CPR. A neural network could then attempt to sort out the remaining patients.

Clinicians may hesitate to place much confidence in "black-box" decision-making aids.[18] The knowledge of the neural network is distributed across the network rather than in a set of rules, and cannot be readily extracted and quantified into a series of rules. In defense of neural networks, it has been argued that the clinical decision making of a human expert, based on years of experience with thousands of patients, can be equally difficult to quantify or codify. Nevertheless, society does not hesitate to accept that expertise.[19] The use of ROC curves to compare neural networks with each other and to compare their performance with other predictive tools will help researchers objectify that task.

Once validated, the neural network can be distributed as a stand-alone application for use by practicing physicians on their personal computers. Because trained neural networks require minimal computer power, they can easily run on existing laptop or palmtop computers. Different networks can be designed to assist physicians with a variety of medical decision-making tasks, generating answers within a fraction of a second.

The accurate prediction of failure to survive following CPR will help physicians and their patients decide when it is appropriate to forgo this medical intervention.

Of course, any predictive tool provides only prognostic information, and should never be the only resource used in decision making. It is only through careful discussions with patients or their surrogate decision makers that physicians can understand the patients' values and assist them in making the appropriate decisions regarding DNR orders.


References

1. Taffet GE, Teasdale TA, Luchi RJ. In-hospital cardiopulmonary resuscitation. JAMA 1988; 260:2069–72.
2. George AL, Folk BP, Crecelius PL, Campbell WB. Pre-arrest morbidity and other correlates of survival after in-hospital cardiopulmonary arrest. Am J Med 1989; 87:28–34.
3. Ebell MH. Survival following in-hospital cardiopulmonary resuscitation: a meta-analysis. J Fam Pract 1992; 34:551–8.
4. Lawrence J. Introduction to neural networks and expert systems. 4th ed. Nevada City, Calif: California Scientific Software, 1992.
5. Furlong JW, Dupuy ME, Heinsimer JA. Neural network analysis of serial cardiac enzyme data: a clinical application of artificial machine intelligence. Am J Clin Pathol 1991; 96:134–41.
6. Baxt WG. Use of an artificial neural network for the diagnosis of myocardial infarction. Ann Intern Med 1991; 115:843–8.
7. Wolberg WH, Mmangasarian OL. Computer-aided diagnosis of breast aspirates via expert systems. Anal Quant Cytol Histol 1990; 12:314–20.
8. Gross GW, Boone JM, Greco-Hunt V, Greenberg B. Neural networks in radiologic diagnosis. II: interpretation of neonatal chest radiographs. Invest Radiol 1990; 25:1017–23.
9. Asada N, Doi K, MacMahon H, et al. Potential usefulness of an artificial neural network for differential diagnosis of interstitial lung diseases: pilot study. Radiology 1990; 177:857–60.
10. Hebb D. The organization of behavior. New York: Wiley Publications, 1949.
11. Firth WJ. Chaos—predicting the unpredictable. BMJ 1991; 303: 1565–8.
12. Knaus WA, Draper EA, Wagner DP, Zimmerman JE. APACHE II: a severity of disease classification system. Crit Care Med 1985; 13:818–29.
13. BrainMaker [computer software]. Grass Valley, Calif: California Scientific Software, 1992.
14. Meistrell ML. Evaluation of neural network performance by receiver operating characteristic (ROC) analysis: examples from the biotechnology domain. Comp Prog Biomed 1990; 32:73–80.
15. Hanley JA, McNeil BJ. The meaning and use of the area under a receiver operating characteristic (ROC) curve. Radiology 1982; 143:29–36.
16. Knaus WA, Wagner DP, Draper EA, et al. The APACHE III prognostic system: risk prediction of hospital mortality for critically ill hospitalized adults. Chest 1991; 100:1619–36.
17. Caudill M. Expert networks. Byte 1991; Oct:108–16.
18. Hart A, Wyatt J. Evaluating black-boxes as medical decision aids: issues arising from a study of neural networks. Med Inf (Lond) 1990; 15(3):229–36.
19. Guerriere MR, Detsky AS. Neural networks: what are they? Ann Intern Med 1991; 115:906–7.


# Appendix

## Definition of Chronic Disease Variable (0 = no chronic disease, 1 = chronic disease)

Organ insufficiency or immunocompromised state must have been evident prior to this hospital admission and conform to any of the following criteria:

*Liver.* Biopsy-proven cirrhosis and documented portal hypertension; episodes of past upper gastrointestinal bleeding attributed to portal hypertension; or prior episodes of hepatic failure/encephalopathy/coma.

*Cardiovascular.* New York Heart Association Class IV.

*Respiratory.* Chronic restrictive, obstructive, or vascular disease resulting in severe exercise restriction, ie, unable to climb stairs or perform household duties; or documented chronic hypoxia, hypercapnia, secondary polycythemia, severe pulmonary hypertension (>40 mm Hg), or respiratory dependency.

*Renal.* Receiving chronic dialysis.

*Immunocompromised.* The patient has received therapy that suppresses resistance to infection (eg, immunosuppression, chemotherapy, radiation, long-term or recent high-dose steroids), or has a disease that is sufficiently advanced to suppress resistance to infection (eg, leukemia, lymphoma, AIDS).