Research Paper ■

Predicting Length of Stay for Psychiatric Diagnosis-related Groups Using Neural Networks

WALTER E. LOWELL, EDD, GEORGE E. DAVIS, MD

Abstract **Objective:** To test the effect of diagnosis on training an artificial neural network (ANN) to predict length of stay (LOS) for psychiatric patients involuntarily admitted to a state hospital.

Design: A series of ANNs were trained representing schizophrenia, affective disorders, and diagnosis-related group (DRG) 430. In addition to diagnosis, variables used in training included demographics, severity of illness, and others identified to be significant in predicting LOS.

Results: Depending on diagnosis, ANN predictions compared with actual LOS indicated accuracy rates ranging from 35% to 70%. The validity of ANN predictions was determined by comparing LOS estimates with the treatment team's predictions at 72 hours following admission, with the ANN predicting as well as or better than did the treatment team in all cases.

Conclusions: One problem in traditional approaches to predicting LOS is the inability of a derived predictive model to maintain accuracy in other independently derived samples. The ANN reported here was capable of maintaining the same predictive efficiency in an independently derived cross-validation sample. The results of ANNs in a cross-validation sample are discussed and the application of this tool in augmenting clinical decision is presented.

■ **J Am Med Informatics Assoc.** 1994;1:459–466.

Over the last decade, many states have been seeking means to transfer care of the mentally ill from the state to the private for-profit sector.[1] This effort is driven by basic values that place a premium on providing care in a patient's immediate community, thus reducing the effect of "institutional dependence," and on the perception that the private sector can deliver efficient and high-quality services at a lower cost. There are, however, many barriers to successfully transferring the care of this population,[2,3] not the least of which is frequent or prolonged hospitalization.

Affiliation of the authors: Augusta Mental Health Institute, Augusta, ME.

Correspondence and reprints: Walter E. Lowell, EdD, Director of Information Services, Augusta Mental Health Institute, P.O. Box 724, Augusta, ME 04330.

Received for publication: 1/21/94; accepted for publication: 7/01/94.

In general medicine, prospective payment schemes such as diagnosis-related groups (DRGs) were developed to control hospitalization. The underlying purpose of DRGs is to ensure that resource consumption is sufficient and appropriate, i.e., length of stay (LOS) is limited to statistically derived diagnostic norms. Unfortunately, psychiatric DRGs do not explain enough variation in LOS to be useful. Psychiatric DRGs, for example, explain only 2–15% of the variation in LOS.[4,5] As a result, DRGs have been waived for psychiatry, and without such tools we compromise the ability to effectively plan resource consumption.

While there continues to be considerable work in refining psychiatric DRGs, the results have not been encouraging, leading some to speculate whether they are even feasible.[6,7] This is because variables used in predicting psychiatric LOS, diagnosis included, are weak predictors, fallible, and often inappropriately

Table 1 ■

Demographics of the Patients of the Combined Diagnosis-related Group 430*

	Schizophrenia (n = 399)	Affective Disorder (n = 430)
Gender		
Female	172 (43%)	230 (54%)
Male	227 (57%)	200 (46%)
Nature of residence		
Lives alone	167 (43%)	159 (37%)
Lives with parents	62 (17%)	59 (14%)
Lives with spouse	25 (6%)	60 (14%)
Lives with relatives	12 (2%)	18 (4%)
Other	133 (33%)	134 (31%)
Age—mean	37 years	41 years
No. prior admissions—mean	5	5
Length of stay—mean	50 days	42 days

*This group consisted of psychiatric patients involuntarily admitted to a state hospital during fiscal years 1990 and 1991 who were diagnosed as having schizophrenia or affective disorder.

intercorrelated.[7] In this study, we further explored the emerging technology of neural computing as an alternative approach to predicting LOS.

Artificial Neural Networks

A complete discussion of artificial neural networks (ANNs) is beyond the scope of this article, but it would suffice to say that these systems are really computers that learn. An ANN is a form of regression for modeling nonlinear relationships between variables. Regression is a statistical technique that estimates the best-fitting curve to a continuously varying parameter. Traditional approaches to regression assume a function of some form. Linear regression, for example, assumes a straight line to fit the data. Regression analysis derives an explicit model, or equation, with known parameters. This equation can then be used to predict values of a dependent variable from known values of independent variables. A key test of a regression equation is its ability to accurately predict outcome in an independently derived data set. This has been a major stumbling block for regression equations designed to predict LOS for psychiatric inpatients.[7]

The problem with regression analysis is that it requires an *a priori* assumption of the particular form (linear or nonlinear) of the equation. If the wrong assumption is made, the regression leads to a poor fit between the data and the curve, which results in poor predictive ability. ANNs, on the other hand, require no *a priori* assumption about the data, nor do they derive a known function or equation.[8] ANNs are robust in their ability to identify patterns through training in complex and incomplete data sets and emulate the biological characteristics of learning, association, memory, and even forgetfulness.[9–14]

In a previous study we demonstrated that an ANN could accurately predict LOS for psychiatric patients as well as or better than could the treatment teams.[15] A trained ANN learned patterns of patient-specific information and was able to discriminate between short- and long-stay patients correctly with an accuracy of 70%. Moreover, the ANN maintained this accuracy in an independently derived cross-validation data set.

Since diagnosis has been shown to account for a small but significant amount of variance in the LOS for psychiatric patients, the question was raised whether controlling for diagnosis would improve predictive accuracy.[7] Unlike multiple regression techniques, ANNs have no method to control for specific variables such as diagnosis. Control is accomplished by physically manipulating the data (i.e., training the data on one diagnosis group at a time). From a neural computing perspective, we expected that multiple input variables for different diagnostic groups would have uniquely different patterns related to LOS. This is a reasonable clinical assumption to make since diagnosis varies on a continuum with a variety of demographic, social, and severity of illness factors.

The purpose of this study was to test the effect of diagnosis on improving the accuracy of ANNs in predicting LOS for involuntarily admitted psychiatric patients. We selected schizophrenia and affective disorder as the primary focus for the study. When combined, they represent DRG 430. These two diagnosis categories were selected because 1) they represent approximately 54% of all admissions to our facility and thus provide a large sample for training and testing; 2) as a result of being involuntarily admitted, these patients often have wide variations in severity of illness and LOS; and 3) when combined, these diagnoses represent not only a major prospective payment category, DRG 430, but also some of the most expensive categories to treat.

Practically, utilization review studies at our hospital have shown that approximately 20% of these patients have short stays (≤7 days). Since this short-stay group is discharged before commitment proceedings take place, its members represent good candidates to be treated in the community rather than in the state hospital. Accurate predictive tools to identify this population prior to admission would improve triage, with the ultimate result of lower admission rates and shorter LOSs.

Method

We used two strategies to evaluate the accuracy of the ANN in predicting LOS. First, data sets were collected for all patients admitted during the study period who had schizophrenia or affective disorders. These two diagnosis groups were trained separately and the ANN predictions for each group were compared with actual LOS and evaluated for accuracy. These diagnosis groups were then combined to represent DRG 430. These combined data (n = 829) were then trained and similarly evaluated. Accuracy of the trained ANN is evaluated as follows. When the data set is entered into the ANN for training, the experimenter selects a random number of records to be withheld from training. We routinely withhold at least 10% of the total sample to be used as a test set. This subset is not used in training and we refer to it as an internal test set. We use the term "internal" because the data are a subset of the training data as distinguished from an externally derived data set. The sizes of the test sets were: schizophrenia, 39; affective disorder, 40; and DRG 430, 80.

The second strategy was to go back to the original data and randomly select a 60% training sample (n = 568) for DRG 430. The 60% was an arbitrary number based on other reported studies, and it provided us with a large training sample.[8] The randomization strategy was used to attenuate fluctuations in the data caused by events independent of treatment. These events include changes in factors such as resources, policy, and staff that are known to affect LOS.[7] A cross-validation data set consisting of 167 records was also randomly selected independent of the training set. The ANN trained on DRG 430 data was thus evaluated using two different test sets: the first was the internally derived sample (n = 56) and the second was the externally derived cross-validation sample (n = 167). In training, some records are identified as "bad facts." The ANN cannot determine a consistent output category for these records because they contain data that are contradictory. These bad facts were eliminated from training. Generally, the amounts of facts eliminated ranged from 1% to 7% of the total sample. Tables 1 and 2 display demographics for data used in the study.

Data

Data for individuals admitted during fiscal years (July–June) 1990 and 1991 who had a Diagnostic and Statistical Manual of Mental Disorders (DSM) III-R diagnosis of schizophrenia, schizoaffective, or schizophreniform (n = 399) and those who had a DSM III-R diagnosis of bipolar depression, major depression, or dysthymia (n = 430) were collected and used for training the networks. The data were obtained from the hospital's computer databases. These databases include the following patient information: age, sex, county of admission, date of admission, nature of residence prior to admission, prior admission history, multi-axis diagnosis information, employment history, family support system, severity of illness, dangerousness, ability to care for self, assaultiveness, community LOS, and hospital LOS.

The data were recorded on admission and at various points during hospitalization. Date of admission is used to establish quarterly data reflecting seasons of the year. Diagnosis data represent the primary working diagnosis for the hospitalization, and include DSM

Table 2 ■

Demographics of the Patients of the Randomized Diagnosis-related Group 430*

	Training Set (n = 568)	Cross-Validation Set (n = 167)
Diagnosis		
Schizophrenia	265 (47%)	76 (46%)
Affective disorder	303 (53%)	91 (54%)
Gender		
Female	271 (48%)	81 (49%)
Male	297 (52%)	86 (51%)
Nature of residence		
Lives alone	197 (35%)	59 (36%)
Lives with parents	82 (14%)	30 (19%)
Lives with spouse	64 (11%)	16 (10%)
Lives with relatives	8 (1%)	1 (1%)
Other	217 (38%)	55 (34%)
Age—mean	39 years	38 years
No. prior admissions—mean	4	4
Length of stay—mean	50 days	41 days

*This group consisted of psychiatric patients involuntarily admitted to a state hospital during fiscal years 1990 and 1991 who were diagnosed as having schizophrenia or affective disorder. These patients were randomly selected to be in a training set or a cross-validation set.

Table 3 ■

Data Domains, Data Elements, and Coding Examples Used in an Artificial Neural Network to Predict Length of Stay for Psychiatric Patients Involuntarily Admitted to a State Hospital

Domain*	Data Element†	Example	No. of Elements
Input variables			
Patient identifier	Unique number	UNIQNO (0,1)	1
Season	Admission date	QUARTER1 (0,1)	4
Gender	Gender	MALE (0,1)	1
Geography	Admission catchment	CUMB (0,1)	5
Age	Age in years	ADOL (0,1)	5
Admission history	Total prior admissions	FIRSTAD (0,1)	3
Community LOS	LDPP–admission date	COMMON1 (0,1)	4
Living arrangements	Nature of residence	ALONE (0,1)	5
Marital status	Marital status	MARRIED (0,1)	1
Family support	Nature of family support	SUPPORTIVE (0,1)	1
Diagnosis	DSM III, Axis 1a, 1b, II, IV	PSYCH (0,1)	9
Severity of illness	GRS psych, social, substance abuse	PSYSEV (0,0.5,1)	3
Dangerousness to self or others	Reason for admission	DANGERA (0,0.5,1)	1
Organization and ability to care for self	Reason for admission	CARE (0,0.5,1)	1
Assaultiveness	SRC as a result of physical assault (actual or threatened)	ASSAULT (0,0.5,1)	1
Output variables			
Hospital LOS	Actual inpatient days	LOSWEEK (0,1)	4
TOTAL			49

*LOS = length of stay.
†DSM = Diagnostic and Statistical Manual of Mental Disorders.

III-R Axis Ia and Axis Ib, Axis II, and Axis IV.[16] Axis Ia and Axis Ib are used to identity patients who have dual diagnoses (e.g., mental illness coupled with substance abuse). In this study, DSM III-Axis Ia was used to identify specific diagnostic groups for training. Severity-of-illness data are also collected on admission, including severity of psychopathology, social impairment, and substance abuse. Scales used to collect these data were designed and developed by the hospital's Psychological Services Department and are available upon request. Reliability estimates range between 0.7 and 0.8 for these scales. Data about dangerousness and ability to care for self are also collected at the time of admission from an emergency involuntary referral form. Employment history and family support system are assessed at the time of discharge. Length of hospitalization is based on actual inpatient days from admission date to last date physically present, exclusive of leave days.

Data Transformation

ANNs can handle a wide variety of data, including character, numeric, and image data. Research suggests that ANN performance improves when data are presented in a modified binary format and our previous work indicated that data in this format could be trained with accuracies as high as 70%.[15] Binary data were used because the network does not discriminate as well between small variations in data such as three or four as it does between larger variations of 50 or 100. Binary data enable the network to know whether a category is present, read 1, or absent, read 0. Networks can also be taught when information is unknown by assigning a value of 0.5 to a variable.

Since our application uses a combination of numeric data (LOS) and character data (nature of residence), all data were transformed into a modified binary format. In our application, we experimented with continuous (analog) and binary data sets and found that performance improved when binary data were used. We therefore transformed all data to binary format. This procedure simplifies data acquisition and conforms more consistently to the general "fuzzy" nature of psychological variables. Table 3 displays data domains, data elements, and coding examples used in this study.

Network Characteristics

We used a commercially available ANN software program to organize the data into a format suitable for the computer to learn.[11] We used a back-propagation algorithm in our application. Neural networks are composed of at least three neural layers: an input back-propagation, an intermediate or "hidden" layer, and an output. Intermediate neurons may have more than one layer, but most applications find that no more than two intermediate layers serve any useful purpose. The number of neurons in the input layer is equal to the number of input variables plus a "threshold neuron" valued at unity to prevent a zero sum.[10] The number of neurons in the intermediate or hidden layer is often determined by experimentation, but may be at least twice the number of input neurons. The number of output neurons is a function of the output purpose of the network. Detailed discussions of the structure of networks have been published.[10,12,17]

Our networks evolved over months of experimentation. There were 48 input neurons, from 100 to 200 "hidden" neurons in a single intermediate layer, and four LOS output neurons (less than one week, more than one week but less than 30 days, more than 30 days but less than six months, and more than six months but less than one year).

In our tertiary-care facility, over half of the patients stayed longer than 20 days. In other hospitals with shorter LOSs, different output categories could be used. We selected the above categories because they represented important LOS distributions at our hospital.

Training the Network

Training the network involves establishing the level of tolerance of accuracy in prediction, which is a function of the degree of variation within the training fact output set. We have previously discussed training issues.[15] The neural network automatically randomly removes approximately 10% of the data to be used as an internal test set, reducing the total number of cases available for training. Basically, multiple variables for each patient are compared as an input data vector with similar vectors (i.e., patients) in the training set.

This process is iterated until all training facts are learned. Depending on the size of the data set, training can take up to several hours. The network outputs a probability estimate for each of the four LOS categories. For example, an output pattern of

I	II	III	IV
0.89	0.10	0.01	0.003

indicates that the patient would probably stay less than a week, category I. In another case

I	II	III	IV
0.19	0.61	0.63	0

the output may be more problematic, implying that the network is having difficulty deciding between categories II and III. We used the rule that the largest probability determined the "answer," even though values might have been close. If all four categories have low probabilities, i.e., less than 0.2, then the network probably does not have enough experience to make a firm estimate. The output categories are independent, therefore the sum of the categories will not add up to unity.

We report two sets of results: one using an exact category, where the predicted and the actual categories are the same, and the other distinguishing between long and short stays. We defined short stay as 30 days or less, since this is a common dividing point used by current reimbursement schemes. We used the term "extended tolerance" (ET) if the network could distinguish between long- and short-stay patients, which in our application would be between categories I–II and III–IV.

Results

The accuracy of predicting LOS was determined by selecting random test sets from the schizophrenic, affective, and combined data sets. For the schizophrenic group we were able to predict 60% by ET and 35% perfectly. For the affective group the comparable figures were 70% and 53%. For DRG 430, the results were 59% by ET and 40% perfectly. These results are based on the internal test sets.

To test the practical utility of each network, its output was compared with the clinical treatment teams' assessment of LOS. These data were collected from the 72-hour team conference notes in the medical record for each case in the test sets. This was considered an important test for the networks since the clinical team was more knowledgeable of the patients, had worked with the patients for a longer time, and had the benefit of a multidisciplinary estimate of LOS. The networks, on the other hand, were limited to a small number of variables. The results of these analyses indicated that for schizophrenia and affective disorder, there was no statistically significant difference between the network and the treatment team for predicting exact LOS category. However, the ET comparisons indicated that the network predicted more accurately than did the treatment team. These differences were statistically significant for both the schizophrenia and the affective disorder groups (χ^2

Table 4 ■

Comparison of the Treatment Team With the Artificial Neural Network for Accuracy of Predicting Length of Stay for Three Groups of Psychiatric Patients

Group*	Treatment Team	Neural Network	χ^2	p Value†
Schizophrenia				
Perfect	8 (21%)	13 (35%)	1.6	NS
ET	10 (26%)	24 (60%)	11.3	0.001
Affective disorder				
Perfect	14 (35%)	21/40 (53%)	2.5	NS
ET	16 (40%)	28/40 (70%)	7.3	0.01
DRG 430 (schizophrenia and affective disorder combined)				
Perfect	26 (33%)	32 (40%)		NS
ET	33 (41%)	48 (59%)	5.6	0.01

*ET = extended tolerance; DRG = diagnosis-related group.
†NS = not significant.

= 11.2, $p \leq 0.001$ and $\chi^2 = 7.2$, $p \leq 0.01$, respectively).

The comparisons for the DRG category indicated that the network performed as well as or better than the treatment team at 72 hours. The results reported above and displayed in Table 4 assumed that the treatment team attempted to predict LOSs for all cases. The treatment team did not make LOS estimates for all cases. We designated these cases with missing clinical predictions as errors made by the clinical team. The value of knowing LOS early in an admission is to enable timely treatment and discharge planning. Failure to estimate may be due to a lack of time or information, or to the complexity of the case. Not knowing a LOS estimate implies a lack of a sense of how ill the patient is, which in turn implies a higher probability of error in treatment and disposition. However, this strategy does inflate the number of apparent errors made by the clinical team. For example, the team estimated LOSs for only 54% of the cases in the schizophrenia group, 63% of those in the affective disorder group, 69% of those in the DRG. To compare the results between the network and the actual number of cases estimated by the treatment team, the distribution was reanalyzed for only those cases estimated by the treatment team. The results indicated no significant difference for any comparison. While these samples were attenuated, bias was not an issue because clinical estimates had no effect on the outcome of the ANN predictions. This indicates that the ANN performed as well as the treatment team for all three groups.

Finally, we trained the randomized DRG 430 training set (n = 568). The results for the internal test set (n = 56) were 47% perfect and 66% correct for ET. For the cross-validation data set (n = 167), the results were 47% perfect and 64% correct for ET. The random data set results were better than the results for the first DRG training set. We attributed this improvement to randomizing non-patient factors such as changes in policy, resources, and staffing, which allowed actual treatment results to emerge. A random sample of 42 cases was selected from the cross-validation data set and compared with predictions of the treatment team at 72 hours. These results indicated no statistically significant difference in predicting exact category between the treatment team at 72 hours and the ANN. There were statistically significant results ($\chi^2 = 8.05$, 1, $p \leq 0.01$) for the ET category favoring the neural network. These findings indicate that the ANN predicted as well as or better than did the treatment team at 72 hours.

Discussion

The purpose of this study was to test the accuracy of ANNs in predicting LOS for two major diagnostic groups separately and combined, the latter group representing DRG 430. A second DRG 430 group was created by randomly selecting a sample consisting of both schizophrenic and affective disorder patients. Randomization was used to minimize the effects of factors unrelated to the etiologies of these disorders but known to affect LOS.

In our earlier work we trained an ANN using multiple diagnostic categories as part of the ANN input.[15] The results reported here indicate that controlling for diagnosis can improve predictability for affective illness, but degrade predictability for schizophrenia. The ANN accuracy for affective disorder was higher than that for schizophrenia and also higher than that of our earlier reported work using multiple diagnostic categories. We attribute this finding to the intrinsic complexity of schizophrenia as opposed to affective disorders. For example, affective disorder has a more predictable treatment response in the hospital, as shown by these data.

To put these results in perspective, the probability of discriminating between a short-stay patient ($\leq$30 days) and a long-stay patient (>30 days), assuming that 70% of a caseload is short stay, in five consecutive admissions is $(0.7)^5 = 0.17$. The randomized

DRG network, as an example, predicts at a rate of 64%, which is 3.8 times better than chance alone (0.64/0.17 = 3.8).

Of interest are the results of the cross-validation study. Studies of predicting LOS for psychiatric inpatients using linear regression approaches have not been impressive.[8] This is due to the shrinkage in predictive power or efficiency when an equation derived from a sample is applied to a second or cross-validation sample. This problem has led some to argue that a system developed in one sample cannot be successfully applied to another, even when the two samples are nearly identical.[7]

We have shown that using a different technology (namely, neural networks) eliminates the problem of shrinkage. A network trained with one set of data was used to predict LOS in a separate randomly derived data set with minimal differences in outcomes. Comparisons between the internal and the cross-validation data sets for the DRG 430 data can be found in Table 4.

The reason that the results of the ET category were better than those of the exact category was due to effectively reducing the complexity of the output categories by half, from four to two. This makes a simpler decision tree for the network to learn. By simplifying output categories to four and then to two, we were able to demonstrate increasing improvement. Whether traditional regression approaches can yield similar results is open to question since no direct comparison between ANNs and regression using psychiatric data has been done.[18] These results suggest that neural networks offer a viable alternative to current statistical approaches to predicting psychiatric LOS. We should add that our primary interest is in developing a predictive tool that has practical clinical applications and not in explaining or determining whether any one variable or combination of variables is best.

As a final test of the accuracy of the ANN, we compared network results with those of the clinical team at 72 hours after admission. Clearly, for any predictive tool to be useful in a clinical setting, the results of that tool must be at least as good as, if not better than, those of the clinician. The results indicate that even when clinicians are using diagnosis to estimate LOS, the ANN predictions are equal to or better than those of the clinicians at 72 hours.

The findings from the cross-validation sample and comparisons between the ANN and the clinical team suggest that ANNs offer a significant step forward in predicting LOS for psychiatric inpatients. These findings are particularly impressive when compared with those reported elsewhere in the psychiatric literature. Of equal importance is that this simple tool can provide accurate estimates for patients who have complex presenting problems. Moreover, the ability of the ANN to find consistent patterns in these complex data suggests that psychiatric treatment practice has consistent patterns that are recognizable by this technology.

An important feature of neural networks is that they can be used easily by paraprofessionals as well as professionals, and do not suffer from the steep learning curve of many expert decision support systems. In addition, when new knowledge becomes available, ANNs can be retrained without the complex reprogramming inherent in Boolean or Bayesian systems. This means that emerging patterns within the data can be learned, thereby changing prediction patterns that are direct reflections of actual clinical practice. This suggests that ANNs will be adaptable to different settings with retraining. Physicians have been reluctant to pursue computer decision enhancement primarily because of cumbersome or tedious processes, outdated information, specificity concerns, and a fear that a machine rather than themselves would be determining the course of therapy.[19] ANNs reflect only what clinicians are collectively doing themselves. This may make this technology more appealing to professionals in the field and may conceivably lead to a more efficient use of our precious resources.

The authors thank Linda Breslin, CSW, Superintendent of Augusta Mental Health Institute, for her ongoing support of this work and John Clark, Librarian, for his assistance in obtaining reference material.

References ■

1. Dorwart R, Schlesinger M. Privatization of psychiatric services. Am J Psychiatry. 1988;145:543–53.
2. Dorwart R, Schlesinger M, Davidson H, Epstein S, Hoover C. A national study of psychiatric hospital care. Am J Psychiatry. 1991;148:204–10.
3. Isaac R, Armat J. Madness in the Streets: How Psychiatry and the Law Abandoned the Mentally Ill. New York: The Free Press, 1990.
4. Stoskopf C, Horn S. The computerized psychiatric severity index as a predictor of inpatient length of stay for psychoses. Med Care. 1991;29:179–95.
5. Taube C, Lee ES, Forthofer R. Diagnosis related groups for mental disorders, alcoholism, and drug abuse: evaluation and alternatives. Hosp Community Psychiatry. 1984;55:452–5.
6. Taube C, Lee ES, Forthofer R. DRG's in psychiatry: an empirical evaluation. Med Care. 1984;22:597–610.

7. Choca JP, Peterson CA. Problems in using statistical models to predict psychiatric length of stay: an illustration. Hosp Community Psychiatry. 1988;39:195–7.
8. Klimasauskas C. Applying neural networks—Part 2. PC AI. Mar-Apr 1991;27–34.
9. Caudill M. GRNN and bear it. AI Expert. 1993;8:28–33.
10. Johnson R, Brown C. Cognizers—Neural Networks and Machines that Think. New York: John Wiley & Sons, 1988.
11. Lawrence J. Introduction to Neural Networks, 3rd ed. Grass Valley, CA: California Scientific Software, 1991.
12. BRAINMAKER PROFESSIONAL—Neural Network Simulation Software. Version 2.0. Grass Valley, CA: California Scientific Software, 1991.
13. Caudill M, Butler C. Naturally Intelligent Systems. Boston: MIT Press, 1990.
14. Admit DJ. Modeling Brain Function: The World of Attractor Neural Networks. Boston: Cambridge University Press, 1989.
15. Davis G, Lowell W, Davis G. A neural network that predicts length-of-stay. MD Comput. 1993;10:87–92.
16. Diagnostic and Statistical Manual of Mental Disorders, 3rd ed. Washington, DC: American Psychiatric Association, 1987.
17. Mulsant BH. A neural network as an approach to clinical diagnosis. MD Comput. 1990;7:25–36.
18. Cohen J, Servan-Screiber D. Introduction to neural network models in psychiatry. Psychiatr Ann. 1992;22:113.
19. Tu J, Guerriere R. Use of a neural network as a predictive instrument for length of stay in the intensive care unit following cardiac surgery. Comput Biomed Res. 1993;26:220–9.

Fellowships in Health Care Informatics
University of Missouri-Columbia

Two- to three-year fellowships are available to train postdoctoral and predoctoral trainees for an academic career in health care informatics. The training program and curriculum prepares individuals to develop, use, and evaluate applications of innovative information methods and computers in the health care environment. Core faculty have active informatics research projects in biomedical and health care fields in the Schools of Medicine, Nursing, Health Services Management, Veterinary Medicine, Health Sciences Library and Information Sciences. NIH stipends, tuition and fees, travel expenses, and computer support are provided. Applicants must be U.S. citizens or permanent residents. Minority and women candidates strongly encouraged. Applications are considered until positions are filled. For priority consideration, send application by December 15, 1994.

Contact: Joyce A. Mitchell, Ph.D., Professor and Director, Medical Informatics Group, University of Missouri, 605 Lewis Hall, Columbia, MO 65211. Telephone: 314/882-6966; fax: 314/884-4270; e-mail: migjoyce@muccmail.missouri.edu

classified ads should be sent to Carol Emery, Hanley & Belfus, Inc., 210 South 13th Street, Philadelphia, PA 19107; FAX (215)790-9330. They should be received 45 days prior to the date of issue. **Advance payment of $2.00 per word must accompany the order.** Add $20 for a boxed ad. Checks should be made out to Hanley & Belfus, Inc.