PREDICTIVE MODELING OF CARDIAC ISCHEMIA

Final Report
NASA/ASEE Summer Faculty Fellowship Program-1995
Johnson Space Center

Prepared By:	Gary T. Anderson, Ph.D.
Academic Rank:	Associate Professor
University & Department	University of Arkansas at Little Rock Dept. of Applied Sciences Little Rock, AR 72204

NASA/JSC

Directorate:	Information Systems
Division:	Technology Systems
Branch:	Client/Server
Date Submitted:	8/11/95
Contract Number:	NGT-44-001-800

ABSTRACT

The goal of the Contextual Alarms Management System (CALMS) project is to develop sophisticated models to predict the onset of clinical cardiac ischemia before it occurs. The system will continuously monitor cardiac patients and set off an alarm when they appear about to suffer an ischemic episode. The models take as inputs information from patient history and combine it with continuously updated information extracted from blood pressure, oxygen saturation and ECG lines. Expert system, statistical, neural network and rough set methodologies are then used to forecast the onset of clinical ischemia before it transpires, thus allowing early intervention aimed at preventing morbid complications from occurring. The models will differ from previous attempts by including combinations of continuous and discrete inputs.

A commercial medical instrumentation and software company has invested funds in the project with a goal of commercialization of the technology. The end product will be a system that analyzes physiologic parameters and produces an alarm when myocardial ischemia is present. If proven feasible, a CALMS-based system will be added to existing heart monitoring hardware.

INTRODUCTION

Cardiovascular disease is the leading cause of death in the US, causing about 43% of all mortalities. Each year, more than 5 million patients arrive at Emergency Rooms (ER) with chest pain, with 35-40% of these suffering from acute ischemia [Selker, 1989]. Coronary Care Units (CCUs) have proven to be extremely effective in preventing death from ischemic cardiac events, but the cost of these units limits their presence to only 22% of hospitals. When cardiac patients arrive at a medical facility, a decision must be made as to whether they belong in the CCU or in a less expensive facility such as a Monitored Care Unit (MCU). For patients arriving at a hospital without a CCU, a decision must be made as to whether they can be treated in-house, or should be transported to a tertiary care facility with a CCU.

The cost of wrong triage decisions can be staggering. Estimates of the percentage of patients needlessly admitted to the CCU range from 50% [Rollag, 1992] to 70% [Fineberg, 1984]. Selker [1989] concludes that each year perhaps $4 billion dollars are spent on CCU care for such patients. In addition, many patients who would benefit from CCU services are not admitted. It is estimated that about 11% [Fleming, 1991] of ER patients with acute ischemic disease are inadvertently sent home. Of those admitted, 9 to 12% [Rollag, 1992; Fleming, 1991] who should be admitted to the CCU are sent to the ward or a step down care facility.

Criteria for admission to a CCU can vary, depending on hospital practice [Weingarten, 1993]. It is known that CCU interventions can significantly lower mortality of patients with acute myocardial infarctions. If implemented in the first 6 - 12 hours after an MI, arrhythmia prophylaxis, cardiac monitoring, thrombolytic therapy and resuscitative interventions available in the CCU can all reduce mortality and morbidity rates for cardiac patients. Quick diagnosis and triage decisions are critical for preventing or effectively treating complications of an MI. However, cardiac triage decisions in the emergency room are often made under severe time pressure, making optimal placements difficult. The proposed CALMS technology will assist the ER physician in making difficult triage decisions by giving them an objective, computer-based second opinion on patient prognosis.

The most difficult triage decision concerns patients with unstable angina, chest pain that is non-responsive to drug treatment. 80-90% of these people will respond to medical therapy, while 10-20% will progress to a myocardial infarction (MI). Based on a pilot study of patients at the University of Arkansas for Medical Sciences, about 8% of people in an MCU will later be transferred to the CCU, indicating that the severity of their illness was originally misinterpreted by the attending cardiologist. Emergency room physicians and family practitioners in rural settings could be expected to have a higher misdiagnosis rate. Once in a CCU, very few life-threatening incidents transpire. If surgery patients, catheterization patients, people admitted to the CCU because they are in the midst of a potentially lethal event and co-morbidity patients (who experience chest pain along with another unrelated illness) are excluded, less than 10% of the remaining population will experience life-threatening episodes. One reason for the low event rate is because of interventions available only in a CCU (e.g., administration of intravenous nitroglycerine or dobutamine), which probably prevented morbid incidences that would have occurred otherwise. However, overcautious admission of people to the CCU likely accounts for a large portion of the low event rate [Selker, 1989].

PREDICTIVE MODELS

Predictive models generally depend on information from a patient's medical history and present medical condition. Several physiologic parameters have been shown to be indicators of future cardiac events. For example, factors as varied as age, hypertension, diabetes, length of stay in CCU [Gheorghiade, 1987], ST and T wave changes [Severi, 1988; Bell, 1990], sex, anterior infarction, hypotension at admission, ventricular tachyarrhythmias, diabetes, Killip class III and IV [De Martini, 1990], previous myocardial infarction [Nishi, 1992], and serum urea [Marik, 1990] have all been shown to have short-term prognostic significance. Recently, changes in heart rate variability has also been shown to be a precursor of clinical ischemia [Bianchi, 1993].

Several researchers have developed models to predict which patients could most benefit being in the CCU [Pozen, 1984; Brush, 1985; Weingarten, 1989, Selker, 1991]. Pozen *et. al* developed a model based on seven discrete inputs to the logistic equation. This model worked best at excluding patients from the CCU (rather than predicting who should be admitted), but missed some obvious candidates [Green, 1988]. In addition, two of the criteria can not be reliably found in a patients medical records (nitroglycerine use and history of heart attacks), and another two may have ambiguous interpretations (S-T segment "straightening" and chest pain as the chief complaint). An improved version of the logistic model [Selker, 1991] used twelve discrete inputs and was shown to perform about as well as an ER physician. To be generally accepted by physicians, however, a decision aid must perform significantly better than physician judgment.

Brush [1984] developed a model based on an "ECG score" that predicted complications in cardiac patients, but the model had disappointing performance when used outside the environment it was developed in [Green, 1988]. Other groups have developed practice guidelines based on expert opinions on how to treat cardiac patients [Weingarten, 1993]. These guidelines work best at selecting patients for early transfer from the CCU, rather than choosing patients suitable for admission.

MODELING TECHNIQUES

Neural Networks

Artificial neural network techniques show excellent promise in being able to overcome the limitations of presently used computer methods to predict patient prognosis. This is because these networks can be trained to recognize complex relationships that exist between inputs (i.e., physiologic data) and outputs (i.e., patient outcome) [de Villiers, 1993]. These subtle relationships in the data are automatically recognized by the network, even if they are unknown to clinicians. Because neural networks can learn any arbitrary relationship between a given set of inputs and outputs, they can normally be expected to perform at least as well as and usually better than any other modeling technique. As the complexity of the problem increases, so does the superiority of neural networks over most other methods. Importantly, neural network techniques have previously been shown to be able to handle the inaccuracy and inconsistency associated with patient histories and physical findings [Pike, 1992; Edenbrandt, 1992; Baxt, 1991; Marik, 1990; Gheorghiade, 1988]. Further, the networks appears to be able to deal with the complexities of disease states characterized by several totally differing clinical presentations [Dassen, 1990].

The disadvantage of neural network models is that, while they often have excellent overall results, they do not reveal how a given prediction was made. Physicians sometimes feel uncomfortable with this "black box" approach to patient management in complicated cases because it is difficult to know when to overrule the network prediction. This objection can be overcome by having a model that can demonstratively perform much better than standard physician judgment.

Rough Sets

Rough sets is a new and powerful technique for extracting rules from data [Pawlak, 1984]. Rough sets have been shown to create impressive predictor models and are especially well suited for problems with inconsistent data, as is often the case with medical problems. Like neural networks, rough sets is a completely data driven technique that can find relationships that exist between problem parameters. A major advantage of rough set models is that they can explain the reason a certain decision was made by revealing what rules were fired. This makes it easier for a physician to reject a decision made by the model on the rare occasions when an unusual set of circumstances suggests such action.

In order to create a rough set model, continuous data must be divided into discrete categories, (e.g., high, medium and low). The rough set algorithm will compare the discretized inputs and output, and eliminate redundant inputs. From the remaining data, a set of rules will be generated that indicates what the likely outcome will be for a given combination of inputs. Certain rules are generated from consistent examples and uncertain rules are generated from inconsistent data. For example, an uncertain rule might state that under given conditions the outcome will be positive 80% of the time. Various methods are employed to give strengths to different rules so that when contradictory rules are fired the most important one will determine the decision.

Rough sets have a few minor disadvantages that have to do with the requirement for discretization of continuous data. If a problem has more than a few inputs, a large amount of data is required to extract rules for all possible combinations of input categories. If a rule has not been generated for a particular combination during training (i.e., rule extraction from a training set of example cases), then no decision can be made when this particular combination occurs during model use. Also, several examples of each combination of categories are desirable to ensure the rules work for a majority of cases. Therefore, a large number of training examples are necessary for the rough set model to generate reliable rules for all possible scenarios.

A second slight disadvantage of rough sets has to do with the crispness of the categories defined for continuous data. For example, a heart rate of 40 - 60 might be considered low, 61 - 80 medium and 81-120 high. Two people may have nearly identical physiologic signs, but one has a heart rate of 80 and the other a heart rate of 81. These people would be considered as being in different categories (80 = Medium, 81 = High), even though they are nearly identical. If a large set of examples is available to extract rules from, this disadvantage can be overcome by using a large number of categories for important variables.

Logistic Regression

Logistic regression is a standard statistical tool that has been used for predictive models with some success [Pozen, 1984; Selker, 1991]. Logistic regression assumes the desired output (usually a "yes" or a "no") fits the sigmoid-shaped logistic equation. The technique has advantages over discriminant analysis in that it can accept combinations of categorical and normal or non-normal continuous data. Data is fit to the equation:

$$Y = \frac{1}{1 + \exp(-u)} \tag{1}$$

where Y is the desired outcome, X are the inputs, b_n are the coefficients of X and $u = b_0 + b_1X_1 + b_2X_2 + \ldots + b_pX_p$. Logistic regression has been shown to work well with categorical and non-normal inputs. Its major disadvantage is that it assumes the data fits a rigid form of equation that may not reflect the subtle interactions actually present between factors in the problem.

DATA ANALYSIS

A pilot study, based on an NSF/Whitaker Foundation planning grant, was conducted to determine the feasibility of developing neural network and rough set predictive models from CCU data. A total of 118 records from patient who had gone through the CCU of the University of Arkansas for Medical Sciences' University Hospital in the past five years was input into a database. Surgery patients, catheterization patients and people admitted to the CCU because they are in the midst of a potentially lethal event were excluded. Thirty seven physiologic parameters from the patients charts were recorded, with 28 model inputs recorded at admission and 9 upon admission to the CCU (see Table I). Four possible adverse outcomes were noted: 1. Type II 2nd degree AV block or 3rd degree AV block; 2. More than 15 seconds ventricular tachycardia; 3. Blood pressure less than 85 with the use of pressors; 4. Death. A total of 44 of the patients suffered serious events while in the CCU. Due to the small number of total events, all four adverse outcomes were combined into a single outcome that was positive if any of the four complications occurred.

Model Input Selection

Data from 118 cases was collected, but only 40 of these had a complete set of inputs. The type of data collected creates special problems for model development for several reasons: 1) there are too few training cases for the number of inputs present; 2) the inputs are correlated; and 3) bad data points probably exist in both the inputs and outputs. A set of predictive model inputs was chosen in a two step process. First, data was divided into two groups based on the outcome (yes = event and no = no event). Student t-tests are a method of testing whether the mean of two groups are equal. t-tests were run to look for differences in each variable between the two groups. Afterwards, stepwise logistic regression was run on the variables selected by the t-tests to choose the final set of model variables. The t-tests were necessary because stepwise regression is performed only on cases that have a full set of all inputs. If a single input is missing from a example, then the entire case is removed from the procedure. This, when applied over the entire dataset, then leaves very few complete cases for model development. On the other hand, t-tests can be performed on all cases where the variable under consideration is present, irrespective of whether any of the other inputs are missing. This allows each candidate input to be evaluated over a larger sample size, thus giving a more solid basis for elimination of parameters that show no difference between outcome groups. After candidate inputs are selected by the t-tests, stepwise regression is performed to eliminate redundancies in the inputs caused by correlations between variables.

Eighteen variables were chosen by the t-tests ($p<0.1$ using either the yes and no groups pooled or separated for calculation of variances) as being possible candidates for the model inputs. The eighteen were: sex, age, weight, diabetes, chest pain, systolic pressure, respiration rate, white blood count, ventricular arrhythmias, ST segment depression, rales, syncopy, S3 heart sound, temperature in CCU, diastolic pressure in CCU, respiration in CCU, aspirin use, class III drug use, class IV drug use, and change in body temperature between ER and CCU. After running stepwise logistic regression, seven inputs were chosen for model development: sex, age, weight, diabetes, ST segment

TABLE I. - INPUT PARAMETERS FOR THE PREDICTIVE MODELS.

INPUT #	PHYSIOLOGIC PARAMETER	RANGE
1	sex	male or female
2	age	continuous
3	weight	continuous
4	smoking	yes or no
5	history of diabetes	yes or no
6	previous MI	yes or no
7	chest pain	yes or no
8	heart rate	continuous
9	systolic blood pressure	continuous
10	diastolic blood pressure	continuous
11	body temperature	continuous
12	respiration rate	continuous
13	hematocrit	continuous
14	serum K	continuous
15	white blood count	continuous
16	creatine	continuous
17	current MI	yes or no
18	anterior MI	yes or no
19	atrial arrhythmias	yes or no
20	ventricular arrhythmias	yes or no
21	ST segment depression	yes or no
22	ST segment elevation	yes or no
23	# of ventricular ectopics in a run	continuous
24	rales greater than 1/3 up	yes or no
25	syncope	yes or no
26	height	continuous
27	S3	yes or no
28	history of congestive heart failure	yes or no
29	heart rate in unit	continuous
30	systolic blood pressure in unit	continuous
31	diastolic blood pressure in unit	continuous
32	respiration in unit	continuous
33	aspirin	yes or no
34	class I drugs	yes or no
35	class II drugs	yes or no
36	class III drugs	yes or no
37	class IV drugs	yes or no

depression, respiration rate in CCU and aspirin use. A total of 95 out of the original 118 cases had all seven of these inputs present.

Factor analysis by principle component decomposition was performed on these seven inputs plus an additional input, presence or absence of atrial arrhythmias, to try to eliminate correlations in the inputs. Three factors were chosen by this method: factor 1 was a combination of sex, respiration rate in the CCU, ST segment depression and diabetes. Factor 2 combined weight, diabetes and atrial arrhythmias, while factor 3 combined aspirin usage and atrial arrhythmias. The resulting factors were fed into a stepwise logistic regression model. The logistic model selected only a constant term, indicating that these three factors have little, if any, predictive power. It was therefore concluded that factor analysis was not an effective means of reducing this particular dataset.

Training and Testing Set Selection

Model development and validation were performed by dividing the database into two categories, one for model training and the other for model testing. Ideally, a training set should capture the important features in the data. The training set should normally be unbiased (i.e., have an equal number of yes and no outcomes), or be intentionally biased to favor a particular result. It is also desirable to have the testing set representative of the data as a whole, so as to get a true idea of model performance. To accomplish these, the data set was clustered by cases, using a nearest neighbor algorithm. Six clusters were visually identified, with between 2 and 31 members in each cluster. Four cases were far from all others, and these were placed in the test set. Two training sets were developed, one with 61 cases and the other with 40. The set with 40 cases was nearly equally balanced between yes and no answers, while the other one had 24 extra no outcomes. The test set, which contained 33 cases, had all clusters represented and contained 13 positive and 20 negative outcomes.

Neural Network Results

The models created were evaluated by using sensitivity and specificity:

$$\text{sensitivity} = \frac{tp}{tp+fn}$$

$$\text{specificity} = \frac{tn}{tn+fp}$$

where tp is true positives, tn is true negatives, fp is false positives and fn is false negatives. Sensitivity is a measure of how likely a model will predict a condition if it is actually present, while specificity indicates how likely a condition is to be present if the model results are positive. Several neural network architectures were investigated, with the best results shown in Table 2.

TABLE 2.- NEURAL NETWORK RESULTS FOR 7 INPUT MODELS.

	Number of Hidden Nodes		
	3	**4**	**3-1**
Average % Correct	58	57	55.5
Sensitivity	0.62	0.54	0.46
Specificity	0.50	0.60	0.65

In Table 2, average % correct is the average of the sensitivity and specificity x 100, while 3-1 indicates a four layer network with three nodes in the first hidden layer and one node in the second hidden layer. The results for three hidden nodes used the training set with 40 cases, while the others used the set with 61 cases.

It was thought that the test set results may have suffered from too many inputs for the number of training cases, so a reduced set of inputs was chosen for further model generation. The new inputs were age, weight, ST segment depression and respiration rate in the CCU. The training set for this network had 61 cases. A network with 2 hidden nodes had the following results:

Average % Correct = 70.5, sensitivity = 0.46, specificity = 0.95

The results are significant. While the model only correctly predicted about one half of all the cardiac events, when it did forecast an event the patient was extremely likely to suffer one (19 out of 20 cases). This network can therefore be used as a screening tool to help decide to place patients in the CCU or, if they are already in the CCU, to keep them there.

Another technique tried to improve model performance was to combine the outputs of the best networks for sensitivity and specificity. These were used as the inputs for a second neural network, with the idea that if each of the original models searched a different area of the solution space then combining them will produce results better than either alone. The output from the network that had a sensitivity of 0.62 (see Table 2) and the one that had a specificity of 0.95 (described above) were combined. The best architecture had four nodes in a single hidden layer:

Average % Correct = 64.5, sensitivity = 0.54, specificity = 0.75

The results are in between the original networks for sensitivity and specificity, thus indicating that the networks were probably keying in on the same features.

The final method tried was to add simulated training cases in order to increase the allowable degrees of freedom in the problem. This procedure also forces the network to learn relationships between inputs. The procedure is as follows:

1. Calculate an average value over all the cases in the training set for each input.
2. For each case, the number of new exemplars created will equal the number of inputs to the model.
3. Each new exemplar replaces a single input with its mean, so that the number of simulated cases created equals the original number of cases times the number of model inputs.

The procedure described above allows a network to be trained with a larger number of hidden nodes without overtraining the network. The inputs for this model were: sex, age, weight, diabetes, ST segment depression, respiration rate in CCU and aspirin use. The original training set had 41 cases, 19 of which were positive outcomes and 22 negative. The new training set had 328 cases with 152 positive outcomes and 176 negative ones. The best network had a single hidden layer with four hidden nodes:

Average % Correct = 66, sensitivity = 0.57, specificity = 0.75

The results improve upon those shown in Table 2, but are slightly worse (66% vs. 70.5% average correct) and not as useful as those from the network with a reduced set of inputs. The limited number of training cases and the combining of four disparate events into a single outcome probably preclude better model performance on this dataset.

Logistic Regression and Rough Set Results

A logistic model was also developed from the same dataset. The training set with 61 inputs was used for coefficient determination (see Equation 1), and the standard 33 case test set was used for model validation. The best validation results were obtained with the following inputs: age, weight, ST segment depression, respiration rate in CCU plus the interactions age x ST segment depression, and weight x respiration rate in CCU:

Average % Correct = 64.5, sensitivity = 0.54, specificity = 0.75

The results are not as good as the best neural networks, but better than many of the networks developed. The logistic model therefore is probably a good benchmark to compare the neural network models to, because it gives an indication if the optimal neural network architecture has been developed for a given problem.

A rough set model was developed from four inputs: age, weight, ST segment depression and respiration rate in the CCU. Continuous inputs were divided into four equally spaced categories that spanned their range. Twelve rules were extracted from the 61 case training set, five for negative decisions and seven for positive decisions. The rule certainty was 100% for eleven rules, and 96% for the twelfth. Each negative rule had between four and twenty-five cases supporting it, with positive rules having between one and six cases supporting them. Decisions were made in 31 of 33 cases in the test set. Model results were:

Average % Correct = 73.5, sensitivity = 0..58, specificity = 0.89

These results are excellent compared to logistic regression and neural network techniques. Although the specificity was slightly less than the best neural network model, its overall performance was better. Moreover, the rough set model made no decision in cases that were not similar to those it was developed on, whereas neural networks will always give an output for all cases.

CONCLUSIONS

Rough sets, neural networks and logistic regression have all proven to be effective tools for predicting the outcome of cardiac patients in a CCU. The rough set model gave the best overall results, and has the advantage of being able to explain how a decision was made. Also, rough set models will not make decisions on cases that are far from the ones they were developed on, adding a degree of confidence to the results. The best neural network model proved to be the most practical, with a specificity of 0.95, although overall results were not quite as good as with rough sets. Logistic regression proved useful as a benchmark against which other methods could be tested.

The key to developing these prognostic models is to choose a good set of predictor variables. This was done in a two step process, using student t-tests and stepwise logistic regression. Selection of cases for training and testing models is also crucial for model creation and validation. A clustering algorithm that measures the distance between cases, while requiring subjective decisions, has shown itself to be useful.

Future work invcludes applying the data analysis techniques described above to the Contextual Alarms Management System (CALMS) project. The goal of CALMS is to develop sophisticated models to predict the onset of clinical cardiac ischemia before it occurs. The system will continuously monitor cardiac patients and set off an alarm when they appear about to suffer an ischemic episode. The models take as inputs information from patient history and combine it with continuously updated information extracted from blood pressure readings, oxygen saturation measurements and five ECG leads. Data is now being collected on twenty patients at the cardiac catheterization laboratory at

Cooper Hospital in New Jersey. Raw data is read into specialized analysis software developed by Po-Ne-Mah. A total of 110 physiologic parameters are written to a text file, which is updated every 1 second. Episodes of ischemia are annotated by physician during the procedure. Since there are too many parameters for the number of patients, each patient will be compared with themselves, with data taken during ischemic episodes compared with data taken when the patient is not suffering ischemia. Student t-tests and logistic regression will be used to choose indicators of ischemia. These will be input into logistic regression, neural network, rough set and expert system models to diagnose and predict future onset of ischemic conditions. One problem that needs to be addressed is drift in these physiologic conditions with time. One possibility for addressing this problem is to look at changes in parameters when ischemia begins, as opposed to absolute readings. Another possibility is to look at inputs in the frequency domain to examine parameters such as heart rate variability and QRS frequency components.

REFERENCES

Baxt, WG. Use of an Artificial Neural Network for the Diagnosis of Myocardial Infarction. Annals of Internal Medicine. vol. 115, number 11, pp. 843, 1991.

Bianchi AM, Mainardi L, Petrucci E, Signorini MG, Mainardi M, and Cerutti S, Time-variant power spectrum analysis for the detection of transient episodes in HRV signal, IEEE Transactions on Biomedical Engineering, Vol 40, 2:136-144, 1993.

De Martini M, Valentini R, Cesana B, Massari FM, Lettino M, Pupilella T, Ambrosini F, Eriano G, La Marchesina U, Lotto A. Early and late prognosis in acute myocardial infarct. A retrospective study in patients admitted to the coronary care unit in the past 10 years. Giornale Italiano Di Cardiologia 20:215-26, 1990.

De Villiers J, Barnard B. Backpropagation neural nets with one and two hidden layers. IEEE Transactions on Neural Networks. vol 4 no 1. pp 136-141, 1993.

Fineberg HV, Scadden D and Goldman L, Care of patients with a low probability of acute myocardial infarction, New England Journal of Medicine, Vol. 310 20:1301-1307.

Gheorghiade M, Anderson J, Rosman H, Lakier J, Velardo B, Goldberg D, Friedman A, Schultz L, Tilley B, Goldstein S. Risk identification at the time of admission to coronary care unit in patients with suspected myocardial infarction. American Heart Journal. 116:1212-17, 1987.

Green L, and Smith M., Evaluation of two acute cardiac ischemia decision-support tools in a rural family practice, Journal of Family Practice, Vol 26, No. 6:627-632, 1988.

Marik PE. Lipman J. Eidelman IJ. Erskine PJ. Clinical predictors of early death in acute myocardial infarction. A prospective study of 233 patients.South African Medical Journal. J77:179-82, 1990.

Pawlak Z. Rough sets. International Journal of Information and Computer Science 11-2. pp. 341-356, 1982.

Rollag, AJ, Jonsbu, Aase, O, and Erikssen, J, Standardized use of simple criteria from case history improves selection of patients for cardiac care unit admission, Journal of Internal Medicine, 232:299-304, 1992.

Selker, HP, "Coronary unit triage decision aids: how do we know they work?", American Journal of Medicine, 87:491-493, 1989.

Selker HP, Griffith JL and D'Agostino RB, A tool for judging coronary care unit admission appropriateness, valid for both real-time and retrospective use, Medical Care, Vol 29, 7:610-627, 1991.

Weingarten SR, Ermann B, Reidinger MS, Shah PK and Ellrodt AG, Selecting the best triage rule for patients hospitalized with chest pain, American Journal of Medicine, 87:494-500, 1989.

Weingarten SR, Ermann B, Bolus R, Reidinger MS, Rubin H, Green A, Karns K and Ellrodt AG, Early "step-down" transfer of low risk patients with chest pain, Annals of Internal Medicine 113:283-289, 1990.

Weingarten SR, Agocs L, Tankel N, Sheng A and Ellrodt AG, Reducing lengths of stay for patients hospitalized with chest pain using medical practice guidelines and opinion leaders, American Journal of Cardiology, 71:259-262, 1993.