Construction of a Literature Database and its Use to Provide Probabilities for Decision-Analytic Models of Thrombolytic Therapy

John F. Murphy, M.D.,[1] Nilesh L. Jain, M.S.,[2,3] Calixto A. Romero, Jr., M.D.,[1] and Michael G. Kahn, M.D., Ph.D.[2]

[1]Cardiovascular Division; [2]Division of Medical Informatics; [3]Department of Computer Science; Washington University School of Medicine, St. Louis MO 63110

Probabilities for decision-analytic models are routinely obtained from the medical literature. This study describes development and use of a literature database to facilitate obtaining probabilities for decision-analytic models of thrombolytic therapy for acute myocardial infarction. Implementation demonstrates the concept of a literature database to be both feasible and effective. Specific difficulties encountered in the evaluation of continuous variables, the potential storage of actual probabilities, and the advantage for easy growth with the literature are discussed.

INTRODUCTION

The ability to develop a complete decision-analytic model and the merit that will be assigned to its recommendations rests heavily on the data used to provide probabilities for the model. Medical experts can be consulted to estimate the probabilities, but there are numerous pitfalls with this approach, including representativeness and the availability heuristic [1]. Probabilities for decision analysis are therefore usually obtained from the medical literature. However, extracting the required probabilities from the literature is a difficult and laborious task. No one source can appropriately include all the references required to attain the knowledge base necessary to develop the appropriate model and to furnish all the needed probabilities. A combination of electronic searches on MEDLINE and manual searches of bibliographies from papers and textbooks is necessary [2]. Once all the appropriate papers have been located, keeping track of all the clinical variables and the status of each variable in every paper quickly becomes unwieldy as the number of papers increases.

Precedent for a literature database to facilitate storage and retrieval of probabilities for decision-analytic models does not exist. Morris *et al.* have constructed a clinical trials database to support meta-analyses [3]. As discussed later it is a step forward but cannot serve an additional purpose of supporting decision analysis.

This study describes development of a literature database to provide appropriate clinical information on thrombolytic therapy for treatment of acute myocardial infarction. The database tabulates the papers so that papers pertaining to any clinical variable or outcome can easily be retrieved. The database is being used to provide information for decision-analytic models used to resolve current controversies in thrombolytic therapy for acute myocardial infarction. Several insights were gained from this process. The ability of the database to provide probabilities is directly dependent on the type of clinical variable to be evaluated in the decision-analytic model. Storage of actual probabilities would require assessment of each paper by database developers. The literature database is a way of organizing a specialized body of literature that allows physicians to master their paper collections with minimal effort. These issues and other specific advantages and disadvantages of the literature database are discussed.

LITERATURE DATABASE

There are over sixty clinical variables and outcomes that can influence the decision to administer thrombolytic therapy for acute myocardial infarction. Papers on thrombolytic therapy were assembled from MEDLINE in both broad searches and searches focused on individual variables. Manual searches of the bibliographies from these papers and from standard cardiology textbook bibliographies were used to increase the number of papers collected.

A database was then created with a data schema consisting of two tables. The first table stored information about every variable in every paper as a boolean value. Clinical variables included information on demographics, physical exam, myocardial infarction characteristics, medications, and outcomes. Papers that pertained directly to thrombolytic therapy and at least one variable were entered into this table. Citation information was recorded in separate text fields. The second table provided background papers for any

0195-4210/94/$5.00

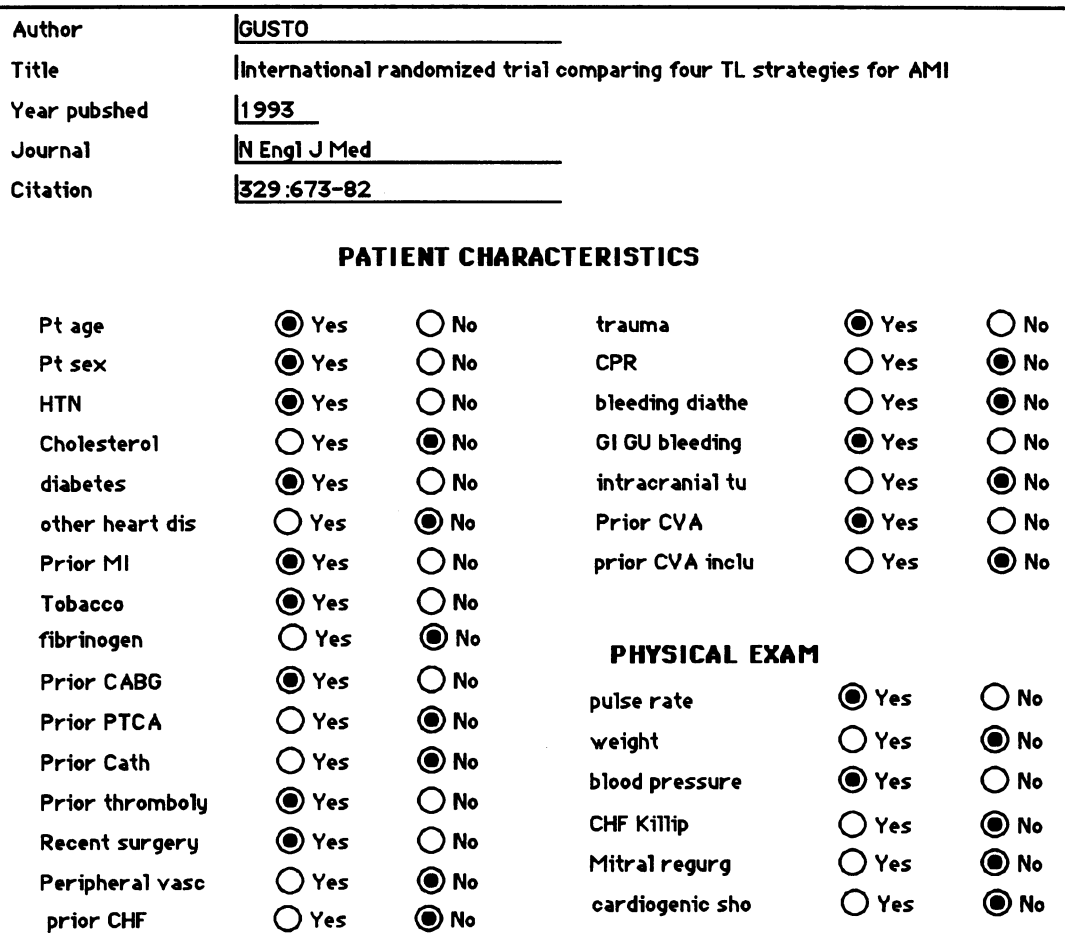

Figure 1: View of one of the screens used for data entry. Shown at the top are the fields for the citation, and at the bottom is a partial list of the boolean clinical variables used for data entry.

given variable. These papers had supplemental medical knowledge to model the clinical situation and to represent the consequences of not giving thrombolytic therapy. These papers did not have any information about thrombolytic therapy, and hence were designated as background papers. These references were listed as citations and one text field for a one-word subject and were not further indexed. Actual probabilities were not stored in the database. The database was built using a commercial relational database, 4th Dimension (ACI US, Inc., Cupertino, CA).

Data entry was straightforward. Citation information of each paper was entered, then the presence of a clinical variable was recorded as the boolean value for that variable. The first data-entry screen is in Figure 1; it contains almost half of the boolean clinical variables stored for every paper and displays actual information from a clinical trial publication.

Papers were then retrieved by selecting any single variable or a combination of clinical variables, including basic citation information. For example, all papers with information on streptokinase, heparin, and mortality can be selected. The search for conditional probabilities can also easily be narrowed. For example, citations for papers which dealt with both location of infarct and post-thrombolytic ejection fraction can be retrieved; a manual search of the papers is then necessary to see if the ejection fraction outcomes are stratified by location of infarct.

DATABASE USE TO SUPPORT A DECISION-ANALYTIC MODEL

A total of 447 key publications have been collected and entered into the database thus far. There are 293 which directly pertain to thrombolytic therapy, and 154 which provide background information.

The database was then used to extract probabilities from the literature for decision-analytic models. The first decision-analytic model developed using the database sought to find the threshold of blood pressure above which thrombolytic therapy should be withheld because the risk of a hemorrhagic stroke and concomitant morbidity and mortality outweighed the benefit of decreased cardiac mortality. For all clinical events, the model required a probability of an event after thrombolytic therapy and the probability of the same event in patients not treated with thrombolytic therapy.

Papers were retrieved using the blood pressure, post-thrombolytic stroke, post-thrombolytic hemorrhage, and mortality variables. There were 9 background papers on hypertension and stroke. There were 107 papers which contained blood pressures. The blood pressures were all reported in very different fashions. One method was to report the mean of a population but not stratify the mortality [4]. Another method was to report ranges only, usually in the form of an upper boundary for exclusion criteria [5]. Some studies did some stratification with broad intervals and unspecified ranges for the first and last intervals [6]. No publications specifically addressed the question of what is the probability of a hemorrhagic stroke at a given specific blood pressure. The only way to combine these pressures to develop probabilities was to combine studies using exclusion criteria. For example, studies with patients with systolic blood pressures less than 200 mm Hg could be pooled and the needed probabilities extracted. The patient population and protocol used were not the same across studies, but the acute MI population entered into studies was relatively homogeneous with regard to risk factors for hemorrhagic stroke, such as female sex: roughly 80% of patients enrolled in virtually every study were male. A decision-analytic model evaluating patients with systolic blood pressure less than 200 mm Hg found they should receive thrombolytic therapy.

However, finding decision-analytic support for a recommendation to give thrombolytic therapy to this group as a whole will not provide adequate basis for the recommendation for the subgroup of patients with blood pressure between 190-200 mm Hg since their results are diluted by the entire population. The risk of hemorrhagic stroke stratified by blood pressure is required to answer the question about this narrow range of blood pressure, and the probabilities were not found in the literature. A sensitivity analysis will locate the threshold stroke rate at which the decision to give thrombolytic therapy changes, but the hemorrhagic stroke rate at a given blood pressure is essential because its comparison to the threshold is the basis for the recommended decision. Obtaining raw patient data from a clinical trial resolved this dilemma by providing the necessary stratification with the data, allowing attainment of the initial model goal of providing thrombolytic therapy recommendations for specific blood pressures [7].

DISCUSSION

The ability of the database to support decision-analytic models depends strikingly on the type of variable to be analyzed by the model. The medical literature, and hence this database, is generally well suited for probabilistic models involving ordinal or nominal variables, but a significant disadvantage is the inability to fully support decision-analytic models for continuous variables. This is readily apparent given the numerous ways in which an objective clinical finding, blood pressure, is reported in the literature. It can be presented as a demographic mean without outcome stratification by blood pressure, as a range with an upper limit given as an exclusion criterion, or with a minimal level of stratification but without a range. The various ways of reporting blood pressure are not a result of inattention to blood pressure in thrombolytic trials, but rather are part of a larger problem of reporting information on continuous variables. Even in trials for evaluation of the ability of a medication to lower blood pressure the blood pressures are not reported in a standardized way to facilitate the most accurate comparison [8]. This problem is easily extrapolated to other continuous clinical variables such as age, heart rate, respiratory rate, and weight.

One remedy to this problem is to use aggregate patient data to supplement the literature. This aggregate data can be in the form of an electronic medical record, the raw data from a clinical trial as was used in this study, a clinical data repository, or observational patient databases. The advantage of these sources is the ability to provide stratified data. These sources also provide conditional probabilities; several papers from a single trial do not have to be analyzed together to arrive at the conditional probabilities. The single source can ensure homogeneity of the patient population and a standard level of evidence, eliminating major steps requiring expert review. Probabilities are easier to retrieve, and no paper review is necessary. An electronic medical record would have the additional advantage of a robust amount of data on each subject whereas clinical trials often record only select clinical variables and outcomes. A second

advantage of the electronic medical record would be the inclusion of all patients, not just those enrolled in trials where a statistically significant result is obtained; in short, publication bias would be eliminated [9]. Obtaining probabilities via a query instead of reviewing many papers would also save time.

Two additional steps could greatly expedite using the literature database to obtain probabilities. This first step is inclusion of study size; searches then could be limited to studies with a specific minimum number of patients. Second, the study design of each paper could be included as a set of boolean variables with the appropriate study design selected; the boolean format would in turn optimize searching. The decision-analytic modelers could then choose a level of evidence for their model, and then limit their search to only those papers with a study design which meets the required level of evidence.

The next step of actually storing probabilities and the fractions used to derive the probabilities would eliminate paper review by the decision-analyst and hence save a considerable amount of time. However, the paper review is not eliminated altogether but shifted to the literature database developers. Selecting papers to provide a probability requires identification of homogenous populations and a standard level of evidence to which each paper must adhere if its probabilities are to be included in the database. These are not small tasks. Identification of a homogeneous population requires homogeneous demographics, risk factors, medical illnesses, concomitant procedures and medications, and a uniform disease and disease stage. Identifying a standard level of evidence goes beyond study size and study design and includes all the quality assessment issues facing meta-analysis. Meta-analysis by definition seeks to combine quantitative data, and it is distinct from decision analysis in that it seeks to use published clinical trial data and arrive at a new standard of care [10]. Decision-analytic results aim at serving as a clinical guideline pending a clinical trial to establish a standard of care. The pitfalls in selecting publications and combining the results apply to both fields. The pitfalls are numerous and include specification of protocol, treatment assignment, selection bias, data-extraction bias, financial bias, statistical methods used, subgroup analysis, quality assessment and quality assessment methods, publication bias, economic impact, outcome definitions, confounding, and misclassification [11,12,13].

One proposal to address the obstacles to meta-analysis is to develop meta-analysis registers which will promote collaboration and provide a basis for methodologic research [14]. Another approach to these problems is the development of the aforementioned clinical trials database. Papers receive a quality score based on clinical protocol and experimental design [15]. Studies deemed to have sufficient quality are included in the list provided by the database which also stores study size, year of publication, agents used with doses, duration of therapy, additional therapy, and the percent of patients in whom therapy is effective. Modifications would be necessary to provide the quantified data necessary for decision analysis. First, all trial endpoints should be included and evaluated separately. Second, papers that include analysis of a different endpoint measured in the same patients should be included and incorporated to facilitate calculation of conditional probabilities. Third, database developers review papers focusing on protocol and design with little attention given to disparity among patients and outcomes. A more comprehensive review is necessary to furnish the probabilities needed for clinical decisions for specific patients. For its stated purposes of producing quick access to an updated list of clinical trials and selection of papers for meta-analysis, the clinical trials database should be lauded. The points raised here are meant to discuss its suitability for a different use, providing probabilities for decision-analytic models.

Once papers have been deemed worthy of use in calculating a probability, shortcomings in papers should be included by adjusting probabilities to account for them. An existing system which contains such corrective properties is THOMAS, a bayesian statistical expert system which takes a single paper, adjusts probabilities for statistical shortcomings, and when given prior probabilities and utilities packages the results into a clinical decision recommendation [16]. The incorporation of the entire above process in storing probabilities would transform the database into a knowledge base analogous to the creation of a disease profile in QMR, previously known as Internist-1 [17]. QMR investigators have attempted to standardize disease profile creation with a knowledge acquisition tool (QMR-KAT) to provide probabilities for diagnostic use [18]. QMR is an excellent educational and consultative resource, but selecting probabilities in this way may remove some of the probabilities essential to sensitivity analysis which validates all decision-analytic models. The probabilities are less likely to be excluded if the decision-analysts themselves review the papers. Conversely, including all published probabilities including those that are not applicable to the population considered and those that do not meet an

essential standard of level of evidence will also give inaccurate probability ranges for sensitivity analyses. Hence there is a trade-off between reduced time for decision-analytic model development and careful selection of papers.

A separate objective of providing a source of background knowledge for decision-analytic model construction, distinct from finding probabilities to use in the model, is easily accomplished with a comprehensive literature database. Searching both tables of the database for references to a clinical variable will provide this information. Consequently, the literature database is an rich source of data for a single decision-analytic model, and it most appropriately serves as a source of data to evaluate comprehensively a medical intervention through decision analysis.

An important feature of the database is the ability to grow with minimal effort. As a new paper is published, it can be entered into the database as it is being read. Hence it adds only a few clicks of the mouse button to the usual "keeping up with the literature" physicians undertake. The number of variables to be considered also helps the researcher focus and read more critically and objectively without additional time spent reading. Aside from its purpose of providing probabilities, the database can be used by physicians to master their own journal collections.

In conclusion, the concept of a literature database to support decision analysis by providing background medical knowledge and probabilities is feasible, effective, and can grow with the literature. Decision-analytic evaluation of continuous clinical variables may require supplemental stratified data from another data source such as raw data from clinical trials. The storage of actual probabilities in the database would save time for model developers but would require extreme care to ensure accurate probability ranges for sensitivity analysis.

Acknowledgments

Dr. Murphy is a cardiology fellow supported by a training grant from the NHLBI. Dr. Kahn is supported in part by Grant #5-R29-LM05387 from the NLM.

References

[1] Dawson N, Arkes H. Systemic errors in medical decision making: Judgment limitations. J Gen Intern Med 1987; 2: 183-7.

[2] Dickersin K, Hewitt P, Mutch L, et al. Comparison of MEDLINE searching with a perinatal trials database. Controlled Clin Trials 1985; 6: 306-17.

[3] Morris RD, Lau J, Arena NJ, et al. A clinical trials database as a research tool in health care. Online J Curr Clin Trials 1992; Jul 17: Doc No 14.

[4] Topol EJ, Califf RM, George BS, et al. Coronary arterial thrombolysis with combined infusion of recombinant tissue-type plasminogen activator and urokinase in patients with acute myocardial infarction. Circulation 1988; 77: 1100-7.

[5] GUSTO Investigators. An international randomized trial comparing four thrombolytic strategies for acute myocardial infarction. N Engl J Med 1993; 329: 673-82.

[6] ISIS-2 Collaborative Group. Randomised trial of intravenous streptokinase, oral aspirin, both, or neither among 17,187 cases of suspected myocardial infarction: ISIS-2. Lancet 1988; 2(8607): 349-60.

[7] Murphy JF, Jain NL, Kahn MG, et al. The impact of hypertension on the risk of hemorrhagic stroke in thrombolytic therapy: An appraisal and hypothesis. 1994; Submitted for publication.

[8] Chalmers TC. Clinical trial quality needs to be improved to facilitate meta-analyses. Online J Curr Clin Trials 1993; Sep 11: Doc No 89.

[9] Levy G. Publication bias: Its implications for clinical pharmacology. Clin Pharm Ther 1992; 52: 115-9.

[10] Cook DJ, Guyatt GH, Laupacis A, et al. Rules of evidence and clinical recommendations on the use of antithrombotic evidence. Chest 1992; 102: 305S-11S.

[11] Sacks S, Berrier J, Reitman D, et al. Meta-analyses of randomized clinical trials. N Engl J Med 1987; 316: 450-5.

[12] Thompson SG, Pocock SJ. Can meta-analyses be trusted? Lancet 1991; 338: 1127-30.

[13] Spector TD, Thompson SG. The potential and limitations of meta-analysis. J Epidemiol Comm Health 1991; 45: 89-92.

[14] Dickersin K, Chalmers TC, Simes RJ, et al. Report from the panel on the case for registers for clinical trials. Controlled Clin Trials 1988; 9: 76-81.

[15] Chalmers TC, Smith, Jr. H, Blackburn B, et al. A method for assessing the quality of a randomized control trial. Controlled Clin Trials 1981; 2: 31-49.

[16] Lehmann HP, Shortliffe EH. THOMAS: Building Bayesian statistical expert systems to aid in clinical decision making. Comp Meth Prog Biomed 1991; 35: 251-60.

[17] Miller RA, Pople, Jr. HE, Myers JD. INTERNIST-1: An experimental computer-based diagnostic consultant for general internal medicine. N Engl J Med 1982; 307: 468-76.

[18] Giuse DA, Giuse NB, Bankowitz RA, et al. Heuristic determination of quantitative data for knowledge acquisition in medicine. Comp Biomed Res 1991; 24: 261-72.