Decision-Theoretic Information Pretrieval: A Generalization of Reminding

Michael M. Wagner and Gregory F. Cooper
Section of Medical Informatics, University of Pittsburgh

Reminder systems and clinical medical librarian services often provide information to clinicians without requiring that a clinician actively seek information. This characteristic may explain in part the effectiveness and high clinician acceptance of these systems. We term systems with this characteristic "information pretrieval systems" to distinguish them from information retrieval systems, which require a clinician to articulate an information need in the form of a query. Because of the increasing importance of information pretrieval systems in medical care, we have developed a decision-theoretic model of an ideal information pretrieval system. In this paper, we present this model and suggest its use as an analytic framework for understanding existing approaches, and as a formal basis for a functioning pretrieval system.

INTRODUCTION

Reminding systems and clinical medical librarian (CML) services are examples of medical information systems that have been shown to affect patient outcomes. Reminding increases clinician compliance with selected practice guidelines [6] and improves selected health outcomes [1, 4, 5]. The information given to clinicians by CMLs has been shown to result in changes in patient management [8].

INFORMATION PRETRIEVAL

Reminding systems and CML services are similar approaches in the sense that they may give unsolicited information to a clinician at a time when the clinician needs the information for patient care. It seems likely that these properties contribute significantly to the effectiveness of these approaches. We suggest a new term—*information pretrieval*—to refer to any approach that has these properties. We coin this new term because existing phrases such as surveillance, reminding, or selective dissemination of information do not connote the key idea of anticipating the context-specific information needs of clinicians.[1]

[1] Pretrieval is derived from the phrase *pre-retrieval*. The prefix pre- means before or earlier and it is intended to suggest retrieval of information before it is requested.

A DECISION-THEORETIC MODEL OF INFORMATION PRETRIEVAL

In this paper, we present a decision-theoretic model of an ideal information pretrieval system. We expect that such a model may be useful in two ways: as an analytic framework from which to understand the assumptions underlying existing pretrieval systems and as a computational basis for new systems.

Definition of a document

The model is intended to cover all types of information, hence we shall use the term *document* to denote any information object (e.g., a reminder message, a journal article, a clinical algorithm, or an image).

Problem formulation

Fig. 1 depicts a general model of an ideal information pretrieval system. The system's task is to infer, from available information (the evidence), the best set of documents to present to a clinician. In a decision-theoretic formulation, the best set is the set of documents with the maximum expected utility, which we shall discuss. For the remainder of the paper, we will limit the discussion to automatic pretrieval systems, although the general principles apply to manual pretrieval systems as well.

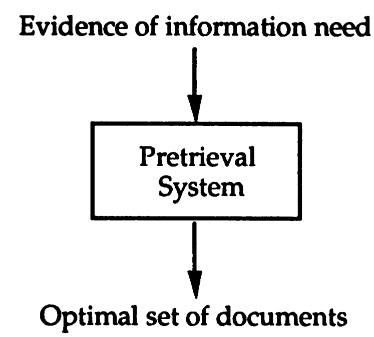

Fig. 1 A general formulation of the information pretrieval task

Assumption of document-document independence

Automatically selecting an optimal set of documents is a difficult problem because documents can interact in complex ways. Without some method to prevent or handle document interactions, a system must know the value of each combination of documents in every clinical circumstance. Therefore, we make an assumption of document-document independence. While this assumption is often invalid (if two information objects are identical, the utility of the pair is not twice the utility of one of them), we make it to simplify the analysis. A model of information *retrieval* making a similar assumption is found in [3]. There are several potential approaches to addressing the limitations imposed by our assumption of document independence. We could prohibit dependent documents, cluster dependent documents manually into independent composite documents, or handle all low-order (e.g., pairwise) interactions by creating a new object for each pair.

By using the document independence assumption, the problem of finding the set of documents with the highest expected utility (EU) reduces to the problem of computing the EU of each document. The system can then build the optimal set by including all documents whose EU is greater than zero. We can think of such a pretrieval system as comprising an automated expert for each document whose job is to compute the EU of the document and forward the document to a dispatcher if the EU is greater than zero (see Fig. 2).

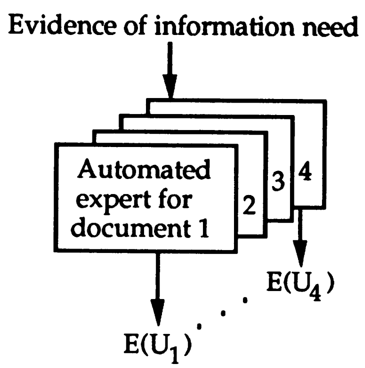

Fig. 2 A model of information pretrieval with the independence assumption

The decision rule that describes the behavior of this system is as follows: *For each document d, pretrieve d if and only if the expected utility of d, $E(U_d)$, given the evidence, is greater than zero.*

Mathematically, *pretrieve document d iff,*

$$E(U_d|evidence) = \int U_d \cdot P(U_d|evidence)dU_d > 0. \quad (1)$$

Evidence of the Information Need

Before we discuss how each document expert computes the EU of its document, let us consider the types of evidence needed for this inference and the types of evidence currently available to a computer-based system.

The value of a document depends on the "information need" of the clinician. This information need is determined, in part, by a patient's problems and the clinician's prior knowledge about the patient's problems. The evidence available in a state-of-the-art clinical information system about a patient's problems includes lists of diagnoses, treatments, and the results of laboratory tests. While we would not expect to find any direct representation of a clinician's prior knowledge, we might find clinician characteristics (e.g., subspecialty, year of graduation, or CME activities) that could provide some evidence of prior knowledge. In any case, we could represent this kind of information, if need be.

Since the evidence available to an automatic pretrieval system is generally not adequate to deterministically deduce the utility of a document, we represent the relation between the evidence of information need and the value of a document probabilistically (see Eq. 1) and use the expected utility to select documents.

Note that the type of evidence that this model can use is not restricted to atomic facts from a database. Like reminder systems, it can use, as evidence, compound statements about a database such as "The patient has a serum bicarbonate less than 18 and yesterday the bicarbonate was greater than 24." We can also imagine variables representing the output of knowledge-based programs such as "The diagnosis of Wilson's disease was established by a diagnostic program." The model can also handle queries by treating them as additional forms of evidence.

Note also that this model can use multiple evidence variables. For example, the automated expert for a document about influenza vaccination might look at three evidence variables: a two-valued variable representing the traditional reminding-rule precondition "The patient is over 65 years old and has not had an influenza vaccination," a variable representing a physician's specialty, and a variable representing whether a physician is academically based.

For notational ease, we shall refer to the set of variables that represent evidence for a particular document, d, as N_d, the model of the information need for that document.

Calculating the Expected Utility of One Document

The main calculation that we must perform in this approach is the expected utility of a single document. Before discussing this calculation, we introduce a distinction between the information value of a document and the time-cost of reading and reacting to the document. We make this distinction because the time that a clinician spends reading and reacting to pretrieved documents is time that could be spent doing something else of value for a patient (e.g., taking a more thorough history). Therefore, we model expected utility as the difference between the expected value of the information in a document, d, and the expected time cost if d is pretrieved.

$$E(U_d | N_d) = E(VDI_d | N_d) - E(TC_d | N_d), \quad (2)$$

where VDI_d denotes the value of document information for d, and TC_d denotes the time-cost for document d. We next discuss these two components of expected utility.

Definition of VDI Depending on the purpose of a pretrieval system (e.g., research or clinical care) we could define the value of a document in different ways. Since we are primarily interested in a pretrieval system for clinical care, we define the value of a document by its ability to improve health outcomes of patients. As suggested by Fig. 3, we measure the value of a document by its ability to improve health states. This type of measurement is called *value-of-information*.

$$\textit{State A} \xrightarrow{\textit{Information}} \textit{State B}$$

Fig. 3 Measuring value of document information (VDI) by the effect of information on the state of the world. VDI = value(State B) – value(State A).

For example, suppose that *State A* represents a 65 year old patient who has not been vaccinated for influenza. Suppose further, that a pretrieval system gives the patient's doctor a document about influenza vaccination that causes the doctor to immunize the patient. If *State B* is the patient's health state after vaccination, then we define the value of the information for this document as the difference between the value of *state A* and value of *state B*.

While measuring the value of the health state of individuals is an open problem, there is general agreement that the attributes of a health state that people value are length-of-life and quality-of-life. Therefore, we measure VDI in units of quality-adjusted life-years (QALYs).

In the influenza vaccination example, the expected benefit due to influenza vaccination for patients of different ages has been established by cost-benefit analysis. Thus, if we send a document to a clinician about influenza vaccination, we know that the expected value to the patient will either be zero (if vaccination is not administered), or approximately 26 quality-adjusted life-days.

Obviously, documents can lead to changes in patient management whose expected value in QALYs has not been established. For this decision-theoretic model to handle such documents, we must ask the involved clinician or an expert to estimate the value of such changes.

Calculating Expected VDI We calculate the expected value of information of a document from the conditional probability distribution over the value of the document, given the evidence

$$E(VDI_d | N_d) = \sum_{VDI_d} VDI_d \cdot P(VDI_d | N_d), \quad (3)$$

where the sum is taken over all possible values of VDI_d. Note that, although the utility of a document is continuous in the general model, here we discretize the VDI variable because many methods for estimating probability distributions from data are limited to discrete variables.

Definition of Time When a clinician receives a document, he or she may spend time reading the document, discussing any proposed actions with a patient, implementing the actions, and reviewing and reacting to their results. We define time as the *total* time that a clinician spends in all of these activities.

Time-Cost As with VDI, the definition of a time-cost utility function depends on the objective of a pretrieval system. For example, if our objective is to optimize patient care within the time-constraints of existing medical practice we might define time-cost as the improvement in patient health-state that would have occurred if the physician had used the "document time" for the next best non-document activity for the patient (e.g., elucidating the history). This is the economic concept of *opportunity cost* [7]. In theory, we could estimate this cost-function by measuring the visit duration and the average change in health-states of similar patients resulting from encounters with similar clinicians.

Alternatively, if our objective is to optimize medical care from the prevailing economic perspective, we could define time-cost as the value, in units of QALYs, of T units of clinician time. To derive this function, for example, we could assume that the economic value of one QALY is at most $100,000 (this is an upper limit value sometimes mentioned by medical consensus committees in medical policy making). This means that we are willing to pay up to $100,000 for medical care that we expect will save one QALY. If we further suppose that the economic value of clinician-time is equal to the price that insurers are willing to pay for non-specialty office visits (approximately $50 for 15 minutes), then one hour of clinician time is equal to 17.51 hours of quality-adjusted life-expectancy. Thus, our time-cost function would be $TC=17.51*T$ This cost function implies that we should pretrieve a document that will take an expected one minute of clinician time only if we expect that the benefit to the patient will be at least 17.51 minutes improvement in quality-adjusted life-expectancy.

Calculating expected time cost With a linear time-cost function, we determine the expected time-cost for a document by applying the time-cost function to the expected time.

$$E(TC_d | N_d) = TC\left(\sum_{T_d} T_d \cdot P(T_d | N_d)\right), \quad (4)$$

where T_d is the time for document d, $TC()$ is the time-cost utility function, and the expression in brackets represents the expected time.

Note that while we are assuming a linear time-cost function here, this is not a restriction of the model. With a non-linear function, however, we lose the independence condition between documents for the time-cost component of expected utility. To handle this, we can change the model so that each document expert returns the expected time and VDI, not the expected utility. The system would use these values to compute the expected utility for each *set* of documents. The decision rule that describes the behavior of this system is *pretrieve the set of documents with the highest expected utility.*

Estimating the Probability Distributions We can estimate the distributions $P(VDI_d | N_d)$ and $P(T_d | N_d)$ from data collected from clinician users of a system. For VDI, we could elicit the VDI directly from a sample of recipients of a document, or elicit the clinical action taken in response to the document and have experts make the value determination at a later time. For time distributions, time studies would be required.

The need for the above distributions may be obviated if documents suggest a limited number of actions (e.g., Pap smear, or no Pap smear) whose value or time requirements are independent of the evidence. For example, we may know that a Pap smear can be considered, performed, and interpreted by a clinician in about five minutes in any patient. Similarly, perhaps on average 30 seconds is spent reacting to the document if the Pap smear is not done. With this independence condition we can compute expected time as

$$\sum_{T_d,\ action} T_d \cdot P(T_d | action) \cdot P(action | N_d). \quad (5)$$

$P(action | N_d)$ may be relatively easy to estimate (e.g., from compliance data collected automatically). For example, using such data from [6], the expected time for a Pap smear reminder is (300 s)(1)(0.15) + (300 s)(0)(0.85) + (30 s)(0)(0.15) + (30 s)(1)(0.85) = 70.5 s.

EXAMPLE: A MAMMOGRAPHY REMINDER

In this section, we show how to implement one practice guideline statement in this model, and we demonstrate the calculation of expected utility of this document to a particular patient (a 60 year old female) seeing a particular clinician (a medical resident).

The practice guideline is that females in the age range 50-70 should have annual mammography. For evidence of information need for this document, we use three variables. The variable *MDL* (for medical decision logic) is a compound statement about the patient "Age 50-70 and no record in the database of mammography this year" which takes the values *true* and *false*. The variables *PGY* and *SPE* are also two-valued (for expository simplicity) and represent, respectively, whether the clinician is a resident (res) or attending (att), and whether the clinician is an internist (IM) or general practitioner (GP). Also for simplicity, we restrict VDI to two values: 2 life-days and 0. Two days is our estimate, based on a published cost-benefit analysis, of the expected increase in life expectancy attributable to a single screening mammogram in this age group [2].

Table 1 is the probability distribution required by Eq. 3. We have used data published by McDonald and colleagues in 1984 [6] to estimate the probability that a reminder will result in a mammogram being performed in such a patient. Note that we only show four of the sixteen rows of the table to save space. We do not have estimates for many of the elements in this distribution, however, we do not need them for this example. In principle, they could be collected automatically as described in [6]. Using this distribution and Eq. 3, we compute an expected VDI of 2(0.06) + 0(0.94) = 0.12 days.

Table 1 $P(VDI_d \mid MDL, PGY, SPE)$

MDL	*PGY*	*SPE*	*VDI*	$P(VDI_d \mid N_d)$
true	res	IM	2	0.06
true	res	IM	0	0.94
true	res	FP	2	–
.	.	.	.	.
false	att	FP	0	–

Although we do not know the expected time for this mammography reminder, we can work backwards from the expected VDI of 0.12 days and our previously derived cost-function of *17.51*T* to derive a time threshold of 0.12 days/17.5=10 minutes. This is an upper limit of the expected time to process this reminder to net a positive expected utility. This means that, if we are willing to pay $100,000 to gain one QALY, then we should be willing to have our clinician spend up to 10 minutes to read and react to this reminder (and to the results of the mammogram, if ordered) which has only a 6% chance of a benefit. Ten minutes far exceeds any reasonable estimate of expected processing time for this message. Note, however, that if the cost of physician time was an order of magnitude greater, or our cost-to-benefit threshold was $10,000/QALY, then the expected time threshold would be 1 minute, which is probably close to the expected time for this message. In that case, the reminder should not be forwarded to the clinician.

To see how different documents compete with each other in terms of expected utilities, suppose that reference 2 is a second document. Reference 2 may have a greater expected time-cost than a terse reminder about mammography. For clinicians who do not accept that mammography is beneficial in this class of patients, however, it may have a greater expected utility than the reminder. This would be the case if the clinicians ordered mammography when given reference 2 sufficiently more often. Once a clinician is familiar with reference 2 (which we could model with a variable representing whether this document had been seen before), the probability that the clinician will heed the reminder may increase to the point that the expected utility of a short reminder exceeds the expected utility of reference 2.

DISCUSSION

There are three key ideas in the decision-theoretic information pretrieval (DTIP) model. First, the decision to give a document to a clinician is an inference made under uncertainty, which is based on evidence of information need. Second, the decision should be made on the basis of the expected utility of the document. Third, in time-constrained domains such as medicine, the time-cost of information should be explicitly represented and reasoned about.

Since it is a model of an ideal document pretrieval system, the DTIP model may be useful as analytical framework for understanding existing pretrieval approaches. For example, a typical rule-based reminding system is a DTIP system in which the *evidence* is limited to a single statement about the medical characteristics of a patient, the *time-cost* of reminding is zero, and the *expected value of information* is assumed to be greater than zero. In reminding systems, the implicit assumption that expected time-cost is always less than value of information means that if real tension exists between information value and time cost, it is being addressed implicitly by the rule authors.

The usefulness of this model as the basis of an actual pretrieval system requires empirical investigation. Some key issues are what types of evidence predict value of information, how to efficiently measure the value of document information and time, how to define time-cost utility functions for clinical medicine, defining the range of document types for which this model can work, and relaxing the document-document independence assumption. Note that the DTIP model reduces to a traditional reminding system when initialized with suitable probabilities and utilities. Thus, it can be viewed as a traditional reminder system that adjusts its parameters and improves its performance as feedback is obtained from clinician users.

One reason to pursue the research problems associated with the implementation of this model is that clinician-time may be a basic constraint that all medical information systems will have to respect. There may be many information systems in the clinical environment that manage information or that need information from a clinician (computer-based diagnostic systems often need information from a clinician in order to produce a differential diagnosis list that may ultimately be of value to a patient). These systems are competing for a clinician's time and value-of-information and time-cost provide a basis for coordinating the activities of these systems.

ACKNOWLEDGMENTS- Support was provided by the National Library of Medicine under grant 5 F37 LM00013 and contract N01-LM-1-3535 and the National Science Foundation under grant IRI-9111590. We thank Gil Kuperman for suggesting the term *pretrieval*.

References

1. Bradshaw KE. *A Computerized Laboratory Alerting System to Warn of Life Threatening Events.* Doctoral dissertation, University of Utah, June ,1988.
2. Eddy DM. Screening for breast cancer. *Ann Int Med* 111:389-99, 1989.
3. Kraft DH. A decision theory view of the information retrieval situation: an operations research approach. *JASIS* 24(5):368-76, 1973.
4. Larsen RA, Evans RS, Burke JP, Pestotnik SL, Gardner RM, Classen DC. Improved perioperative antibiotic use and reduced surgical wound infections through the use of computer decision analysis. *Infect Cont Hosp Epidemiol* 10:316-20, 1989.
5. McDonald CJ, Siu LH, Tierney WM. Effects of computer reminders for influenza vaccination on morbidity during influenza epidemics. *MD Comput* 9(5):304-12, 1992.
6. McDonald CJ, Siu LH, Smith DM, Tierney WM, Cohen SJ, Weinberger M, McCabe GP. Reminders to physicians from an introspective computer medical record. *Ann Intern Med* 100(1):130-8, 1984.
7. Russell LB. Opportunity costs in modern medicine. *Health Affairs* 11(2):162-9, 1992.
8. Scura G, Davidoff F. Case-related use of the medical literature. Clinical librarian services for improving patient care. *JAMA* 245:50-52, 1981.