ISABEL: a novel Internet-delivered clinical decision support system

P Ramnarayan, G Kulkarni, A Tomlinson, J Britto

ISABEL Clinical Decision Support Systems
Acrow Building
St. Mary's Hospital
South Wharf Road
London W2 1NY

Keywords: ISABEL, clinical decision support, diagnosis, Internet

Abstract

The development, and the results of clinical evaluation of ISABEL, an Internet-delivered integrated clinical decision-support system developed over the past 3 years, is described. Utilising a novel technique, pattern-recognition software from Autonomy™ is used to search standard medical text. The system consists of an innovative differential diagnosis tool to provide a list of potential diagnoses for a patient's clinical features. Further clinical corroboration for each of these diagnostic suggestions is provided by the associated text description, relevant images annotated by experts, national guidelines for management, and pertinent clinical experiences. Initial evaluation of ISABEL showed that in around 85% of the cases tested, ISABEL presented the final diagnosis in the list of differential diagnoses. In an assessment of impact on clinical decisions made by doctors, both in a simulated setting and in real life, ISABEL reminded the user to consider a significant diagnosis in 1 out of every 8 cases.

Introduction

Biomedical knowledge has grown exponentially in the past few years, resulting in severe information overload for clinicians[1]; it is estimated that this problem will double every 20 years[2]. It is clear though that information related to newer medical tests and treatments is constantly evolving, whereas traditional knowledge pertaining to clinical diagnosis has changed relatively little. This suggests that all domains of medical

knowledge have not been equally affected by this rapid growth. Recent techniques that attempt to summarize latest treatment recommendations to comply with medical evidence are now available, and are popular with clinicians[3,4]. Systems that might assist in routine clinical diagnosis remain scarce. We know that medical information relevant to making clinical diagnoses is constantly needed, and that this information need is fulfilled mainly by consulting textbooks[5].

We also know that errors related to misdiagnoses, or missed diagnoses, constitute a significant proportion of the preventable burden of medical error[6]. Medical error is reportedly common, yet there are no reliable estimates of the incidence of medical error within the National Health Service (NHS). Previous studies in the USA and Australia have estimated adverse events to occur in between 4 and 16% of all hospital admissions[7,8]; a recent study from the UK reported[9] an estimate of medical error in the UK. In this latter study, 10.8% of all inpatients suffered an adverse event in hospital, either related to diagnostic evaluation, operative procedure, ward management, or drug prescription. Some patients suffered more than one adverse medical event. Underlying reasons that predispose the occurrence of medical error are often systemic, not related to individuals. As the authors of the British study conclude, these reasons may include dependence on diagnoses made by inexperienced clinicians, poor medical records, poor communication and inadequate consultant input[10].

Diagnostic errors were recorded in a fifth of the cases. These errors may be due to errors of omission – failure to consider all clinically relevant diagnoses during initial workup. Additional contributors to diagnostic error may include incorrect formulation of the clinical problem, as well as difficulty in extracting relevant information from textbooks quickly, whether paper-based or electronic. The current structure of healthcare delivery in the NHS may also contribute to the occurrence of medical error. The current 'inverted knowledge pyramid' of clinical care, junior doctors with relatively little knowledge and experience at the bottom and experienced, knowledgeable consultants at the top, may be one main reason[11]. Unfortunately, the critically patient is initially assessed by a junior doctor at the frontline, deprived of specialist advice out of working hours and unarmed with clinical experience. Financially as well as in practice, it seems clear that inverting this knowledge pyramid is difficult to achieve without incorporating the use of techniques to manage knowledge and deliver it in usable form to the clinicians at the bedside.

One reason why computerized diagnostic aids are scarce may be related to the difficulty of converting traditional medical knowledge into computer-readable form. Expert systems developed for internal medicine, such as DXplain, QMR and ILIAD, were developed over many years, involving the input of multiple experts to provide semi-

probabilistic relationships between thousands of clinical features and hundreds of diseases. These tools were also developed to assist the clinician primarily during the rare entity of a diagnostic dilemma (clinical dead-end), by acting as oracles. As a result of this design, they often required the user to expend a considerable amount of time interrogating the system. To regularly use such a stand-alone system in practice, a clinician had to be highly motivated. This may be one reason why diagnostic decision support as a concept has not captured clinicians' interest.

Aims

To develop an easy to develop and user-friendly, integrated clinical decision-support system meant to be used for point-of-care decision-support in hospital, and to evaluate the clinical performance and impact of this system.

Methods

Underlying knowledge base

Standard electronic medical text was used to populate a pre-designed diagnostic tree comprising around 3500 diagnoses. Text pertaining to each diagnosis within the tree was copied into the ISABEL database, without any modifications, by one research nurse. Thus, unformatted text relating to the same disease, from different sources, was collated under one disease label. Where new disease labels were necessary to accommodate text from a new source, they were created within the same overall diagnostic tree model. Since there was no modification of the text involved, this entire process took less than 4 months for one researcher.

Inference engine

Commercially available software was used as the inference engine (Autonomy™, Autonomy Corporation, Cambridge, UK, www.autonomy.com). This software utilizes Bayesian inference and Shannon's principles of information theory to generate pattern-matching algorithms to enable sophisticated concept extraction from documents. By collating text related to one specific diagnosis under a single diagnostic label within the diagnostic tree, it was possible for the software to generate a unique signature of key concepts for each diagnosis, by using its concept extraction techniques. This signature was constantly modified with the addition of text from each additional source used to populate the ISABEL database.

Search methodology

In response to a set of key clinical features for a patient (concepts),

ISABEL utilized Autonomy™ to search the underlying database of text, and return all documents (diagnostic labels) whose concept signature matched that generated from the clinical features.

System architecture and delivery platform

In order to maximize the use of the system, and eliminate inequalities of regional distribution, the intention was to deliver ISABEL on the World Wide Web to all medical practitioners, after a short registration process. To this end, a website was created by DynamicWeb, UK using Javascript. The Autonomy Dynamic Reasoning Engine™ (DRE) was hosted on a dedicated ISABEL server, as was the database comprising the diagnostic tree, which could be construed as being the equivalent of multiple documents of text in html. At the front end, a free text box, into which the user could enter the clinical features of a patient, was created on a dedicated diagnostic tool webpage on the ISABEL website. On searching the database with these features, a list of all matching diagnostic labels (diseases) was returned. The maximum and minimum number of the diagnoses displayed on the results page could be varied by the developers of the system.

Figures 1 & 2 show how a set of clinical features entered in free text into the search box produce a set of diagnostic labels (with the preceding text: have you considered?) for the user's attention.

Figure 1: Differential diagnosis tool

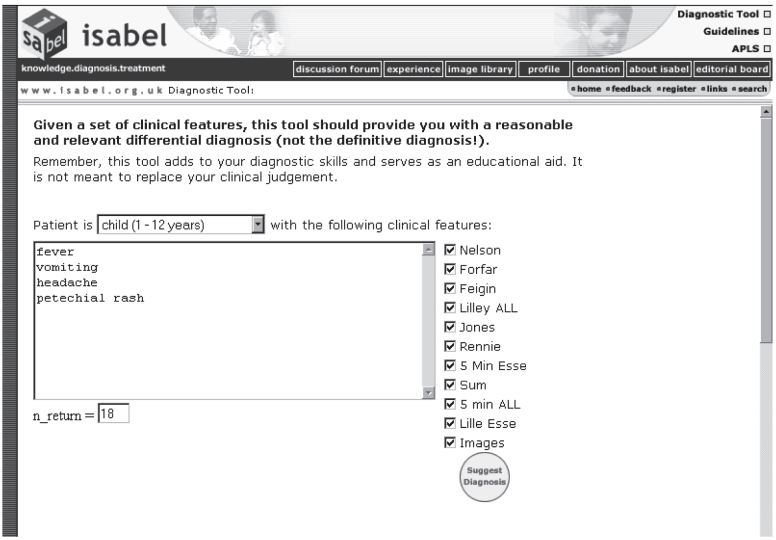

Figure 2: Differential diagnosis tool output, ranked in order of textual match

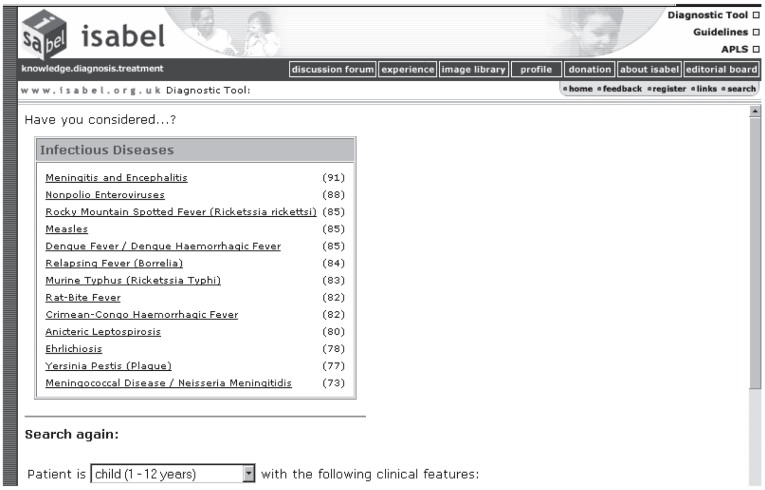

Figure 3: Guideline page

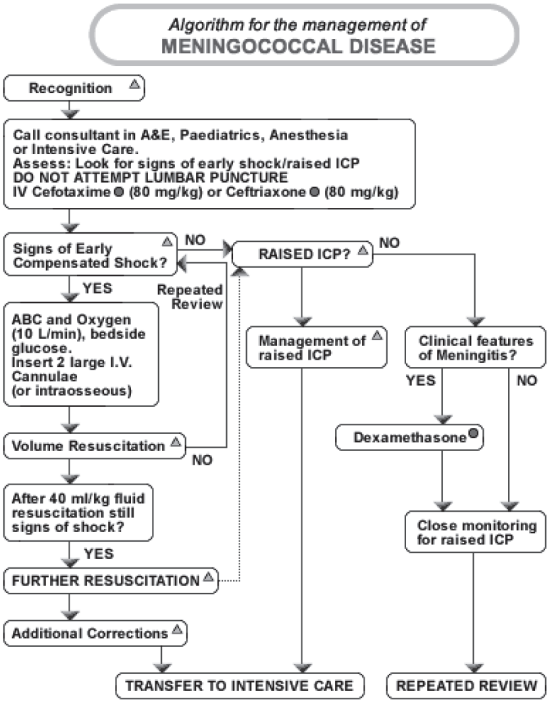

Guidelines
National guidelines for the management of various conditions were collected from numerous sources. These included the Advanced Paediatric Life Support group (APLS) and other specialists. The most appropriate and evidence-based guideline in use for each particular condition was chosen. This was then adapted for ease of access and uniformity on ISABEL. All guidelines were represented in the same format, as algorithms with clear decision nodes and action nodes. Hyperlinks were created for each drug name in the guideline to link to the appropriate page in the British National Formulary (BNF), with description of the dosage, interactions and side effects. Bare detail was presented to the user in a first layer, and the evidence and textual description for each node hidden in a second more complex layer for use only if chosen by the user.

Images
Clinical images from experts in various specialities were collected and digitised for use on the website. Once annotated with textual description by the experts online, each image was then linked to a diagnostic category as was previously described for the DDT. This annotation included a description as well as a management plan for each image. Images included clinical photographs, radiographs, blood pictures and other clinical pictures. These images could be searched in multiple ways – by image type, body system involved as well as by keywords.

Experience section
Clinical experience is a result of years of learning and an accumulation of clinical lessons learnt. These experiences are commonly reported as case reports in journals to alert the readers of potential pitfalls in clinical medicine. The aim of this section of ISABEL was to be able to collate and categorise lessons gathered in everyday clinical practice, as contributed by users and experts, as well as from journal reports. For users, the facility to anonymously record clinical errors and lessons was created, so that a short description of the clinical setting, the associated lesson learnt and the keywords to be assigned to the particular experience could be recorded. An editor for the section vetted the experience for utility as well as for significance. This was then assigned a category as described for the DDT. Users could search for individual experiences by keywords or by the system involved. In addition, experiences were integrated into ISABEL, as described later.

Users
ISABEL was primarily intended for junior doctors and primary care physicians, but other medical professionals could also use the system.

While the initial version of the website was restricted to paediatrics, it is being extended to adult medicine and surgery. The user group has also extended to other medical professionals including nursing staff, physiotherapists and others. Members of the general public can view only selected sections of the system. This was ensured by registration and password access, by confirming credentials using the unique identification number assigned by each professional body eg General Medical Council, United Kingdom Central Council for Nursing, Midwifery and Health Visiting (UKCC) etc.

Clinical evaluation

Validation of the DDT has been performed extensively, and reported in the peer-reviewed medical literature. The scheme followed for the evaluation process is shown.

Figure 4: Scheme of clinical evaluation process

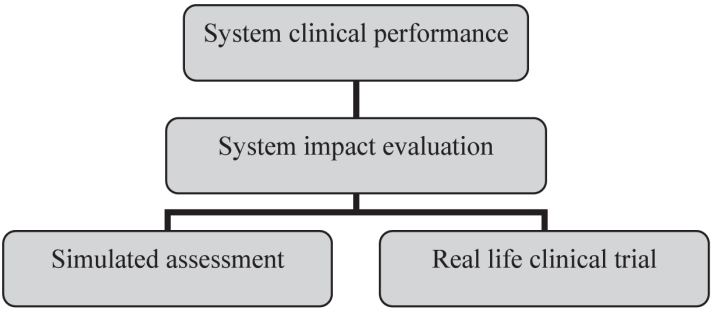

Results

Usage statistics

25 GB of data was transferred in the period from July 2001-August 2003 (average/day: 32 MB). There were 9,168,147 successful page requests (average/day: 11,641); over 12,000 users registered to use the site in the specified period. Over a fifth of users accessed the system more than five times since registration. The entire National Health Service (NHS), UK was provided log-in free access via IP address recognition in mid-2002; it then proved difficult to estimate the true number of users.

Evaluation results of the DDT have been encouraging. They were reported[11] in the Archives of Disease of Childhood in 2003. The initial system performance, using hypothetical cases and data from real cases, is summarised in the table.

Table 1: Initial performance evaluation summary[12]

	Stage 1 validation	**Stage 2 validation**
Objectives	To assess the number of case histories in which the expected final diagnoses was presented in ISABEL DDT	To assess the number of cases for which the final diagnosis was presented in ISABEL DDT
Source of cases/features	Junior doctors in one paediatric centre	Real children in 4 paediatric centres
Nature of cases	Mainly hypothetical	Real
Source of final diagnosis	Junior doctors' expected diagnoses for a set of clinical features	Discharge summary diagnosis
Number of diagnoses in ISABEL List	10	10-15
Accuracy rate*	87%	83%

***Accuracy**: defined as proportion of cases in which the final/expected diagnosis was present in the DDT output list.

The results of the impact evaluation, both for the simulated experiment, as well as the clinical trial conducted at four NHS hospitals is summarised below.

	Simulated field trial	**Real life clinical trial**
Objectives	To assess the impact of ISABEL-delivered diagnostic suggestions in a simulated environment	To assess the impact of ISABEL-delivered diagnostic suggestions in a natural NHS clinical environment
Subjects	76 clinicians, comprising consultants, medical students, registrars and SHOs	60 junior paediatricians in training at 4 NHS hospitals in the South-East of England
Case material	24 simulated cases, balanced to represent different specialities and levels of difficulty	Real patients assessed by the participating junior doctors, and thought to warrant the use of ISABEL during the trial period

	Simulated field trial	Real life clinical trial
Gold standard	Decisions judged by an expert panel to be clinically significant Final diagnoses for patients	Decisions judged by an expert panel to be clinically significant
Outcomes	Number of significant diagnoses reminded by ISABEL	Number of significant diagnoses reminded by ISABEL Number of final diagnoses reminded by ISABEL
Results	In 1 out every 7 cases, a significant diagnosis was reminded	In 1 out every 8 cases, a significant diagnosis was reminded

Discussion

This project aimed to produce a clinically useful Internet-delivered decision-support system for the use of the busy clinician at the bedside. The separate components of the system integrate with each other, so that when the user is provided a differential diagnosis, textual description of that diagnosis is made available to the clinician, images are provided to support the diagnosis, guidelines are provided to aid management of the child's condition, and finally clinical experience from a number of experts and journals is brought to the point-of-care at the appropriate point to improve the junior doctor's clinical reasoning. Such integration is crucial to quality decision-making in a modern healthcare system and attempts to invert the knowledge pyramid.

Extension of the DDT to adult medicine as well as surgery and obstetrics is underway. The addition of various other components of the site is also ongoing. Clearly, the use of knowledge management techniques is crucial to a profession that is knowledge-intensive. Clinicians regularly use up to 2 million pieces of information to make decisions. The use of clinical decision-making support systems on the Internet is the way of the future, considering the need for universal access and the problems with geographically disparate medical care settings. ISABEL is one attempt to address this need.

References

1. Wyatt JC. Clinical Knowledge and Practice in the Information Age: a handbook for health professionals. The Royal Society of Medicine Press. 2001.
2. Wyatt J. Uses and sources of medical knowledge. Lancet. 1991, 338:1368-72.
3. Brassey J, Elwyn G, Price C, Kinnersley P. Just in time information for

clinicians: a questionnaire evaluation of the ATTRACT project. BMJ. 2001, 322(7285):529-30.

4. Godlee F, Smith R, Goldmann D. Clinical evidence. BMJ. 1999, 318(7198):1570-1.
5. Ely JW, Osheroff JA, Ebell MH, Bergus GR et al. Analysis of questions asked by family doctors regarding patient care. BMJ. 1999, 319(7206):358-61.
6. Leape LL, Brennan TA, Laird N, Lawthers AG et al. The nature of adverse events in hospitalized patients. Results of the Harvard Medical Practice Study II. N Engl J Med. 1991, 324(6):377-84.
7. Brennan TA, Leape LL, Laird NM, Hebert L et al. Incidence of adverse events and negligence in hospitalised patients. N Engl J Med. 1991, 324:370-376.
8. Wilson RM, Runciman WB, Gibberd RW, Harrison BT et al. The quality in Australian health care study. Med J Aust. 1995, 163:458-471.
9. Vincent C, Neale G, Woloshynowych M. Adverse events in British hospitals: preliminary retrospective record review. BMJ. 2001, 322(7285):517-9
10. Neale G, Vincent C, Woloshynowych M. Exploring the causes of adverse events in NHS hospital practice. J R Soc Med. 2001, 94(7):322-30.
11. Ramnarayan P, Britto J. Paediatric clinical decision support systems. Arch Dis Child. 2002, 87(5):361-2.
12. Ramnarayan P, Tomlinson A, Rao A, Coren M et al. ISABEL: a web-based differential diagnostic aid for paediatrics: results from an initial performance evaluation. Arch Dis Child. 2003, 88:408-13.