The Knowledge Engineering Review, No. 5:4, 1990, 251–263

Medical AI systems as appropriate technology for developing countries

KATHLEEN KING

Department of Artificial Intelligence, University of Edinburgh, UK

HOWARD BECK

Artificial Intelligence Appications Institute, University of Edinburgh, UK

Abstract

Expert systems technology has been around for a long time, becoming increasingly easy to use, inexpensive and reliable in recent years. It would seem to provide an ideal vehicle for the dissemination of expertise in developing countries, particularly in the field of medicine, which was the focus of much early work in diagnostic systems. Despite the apparent match of a real problem and a credible solution, however, remarkably few AI systems for medicine in developing countries have been researched, designed and implemented. This paper addresses why this might be the case, reviews some of the extant systems and explores some of the design issues. Particular emphasis is placed on the question of "Appropriate Technology". Various criteria for Appropriate Technology are explored, and an optimal set used to guide principles of design. It is argued that medical AI systems can satisfy these criteria, provided that sufficient care is taken in their design for the country of application.

1 The existing problem

". . . the main social targets of government and WHO in the coming decades should be the attainment of all citizens of the world by the year 2000 a level of health that will permit them to lead a socially and economically productive life." (World Health Organization: resolution WHA30.43)

In rural areas in many developing countries, and in remoter or poorer parts of some developed countries, there is a shortage of doctors and a reliance on nurses, paramedics or individuals with minimal medical knowledge for diagnosis and therapy. This is often a self-perpetuating situation in that when health workers are trained to a higher standard they tend to move elsewhere for employment. The drain of such resources is generally towards urban areas and the distribution of medical expertise becomes progressively unequal. Even in countries where a good health service infrastructure exists or is being set up and the training of health workers is adequate, once these workers are in the field they can have a hard task to keep up with increasing populations, changing patterns of disease and the sheer workload of catering to the health needs of a large rural community.

Towards the goal of "Health for All" by the Year 2000 mentioned above, the WHO defines primary health care as

". . . essential health care based on practical, scientifically sound and socially acceptable methods and technology made universally acceptable to individuals and families in the community through their full participation and at cost that the community and the country can afford to maintain at every stage of their development in the spirit of self-reliance and self determination. It forms an integral part of both the country's health system, of which it is the central function and main focus, and of the overall social and economic development of the community bringing health care as close as possible to where people live and work, and constitutes the first element of a continuing health care process". (WHO, 1978)

This is what the health worker is trying to provide, and this is where AI can make a contribution; not necessarily at the "high-tech" end of medicine where much of the research into AI for medicine has gone on in the past, but at the simple and practical end, where a little support can go a long way.

2 The promise of an AI solution

Many people, when they first encounter the idea of an expert system, and especially when they hear about early experiments with medical diagnostic systems like MYCIN (Clancey and Shortliffe, 1984) and Internist (Miller et al., 1982), think what a boon such an application would be in areas where expertise is thin on the ground. An expert system would seem to have many advantages for such situations; it does not wish for the bright lights, it can be called upon at any time (provided it, and its power supply, are operative), it is consistent, as an *advisor* or backup for a diagnostician it will part with the most elementary information in the same manner as the most obscure, its knowledge can be updated, its pronouncements are distributed equitably and it does not ask for payment.

Expert Systems technology has now become well established. The acceptance of Expert Systems by industry proves the methodology to be cost effective and reliable enough to survive in the harsh world outside of the academic environment. The hardware and software to support development and delivery is becoming cheaper every year. Even hybrid system building environments, previously only for costly workstations, are now available on PCs and many development environments have a cheaper runtime version specifically for delivery of the finished product. It would seem that they would be the ideal vehicle for giving access to medical knowledge for health workers, and thus access to health care for everyone, a step towards the WHO's goal of "Health for all by the year 2000".

Despite the general acknowledgement that one of the most promising characteristics of Expert Systems, particularly diagnostic systems, is their potential for disseminating knowledge and expertise in an accessible manner, we have been able to uncover remarkably few projects concerned with using AI for medicine in developing countries. There are several possible reasons for this, some more credible than others.

3 Why isn't there much around?

3.1 Appropriate technology

Computers are "high tech" objects, whose place belongs firmly in the first world.

There is something of a feeling around that for anything to be "appropriate technology" it has to be made out of cow-dung and old tractor tyres. Computers are not, so they are inappropriate. This is a particularly common objection, and so we intend to address it here in some detail. In §5 we explore some proposed criteria of Appropriate Technology, and attempt to identify those with specific application to medical AI systems.

3.2 Prioritizing expenditure

If people do not have a proper water supply why spend money giving them computers?

This is a general objection to any expenditure which is not geared towards the widest application at the lowest level. It was also used in the UK against providing computers in schools. There is a point at which it *does* become worthwhile to diversify resources but quite where that point is is difficult to ascertain. The decision to implement technological solutions must be balanced by the existence of enough infrastructure to support them.

3.3 Cost of development and delivery

Both expatriate AI experts and hardware are costly, especially if they have to be bought with foreign currency.

Initial development can be done outside the country of application, as long as enough attention is paid to its *appropriateness. Technology Transfer* at least at user and maintenance level ensures that the costly foreign expertise is not necessary once the system is installed. Hardware is coming down in price and some countries have local assembly plants. It should still be borne in mind, however, that poor exchange rates will always mean that anything imported will be relatively costly and this may heavily influence the choice of delivery machine. The assessment of *real* cost is complex. For example, there may be hidden material costs in the existing health-care setup, which the proposed system might ameliorate, such as overprescription or overuse of expensive laboratory tests.

3.4 Lack of infrastructure

Many third world countries do not have a reliable electricity supply, let alone a "computing policy". To be used, the system must be reliable, and to be reliable, it has to be supported.

Again, technology transfer helps develop local infrastructure. This means training technicians and possibly software engineers. The maintenance of the system must be planned as a fundamental part of the project. In the case of an expert systems project, a particularly important issue in this regard is to allow knowledge bases to be accessible and modifiable by the domain experts; if it is easy to maintain it will be maintained and not become obsolete. However, as mentioned above ("prioritizing expenditure") the choice of what is implemented must be influenced by what facilities are actually available.

3.5 High tech medicine

Expert Systems historically have tended to become geared towards the high-tech or very specialized areas of medicine, and to replicating or improving upon specialist performance. Making a system which deals with, in a sense, all of medicine, appears to be a massive task.

What is really needed here is usually a broad but fairly shallow system. In most areas where this would have application, the bulk of a health worker's caseload is taken up with relatively few and simple diseases. Complex, unknown or dangerous cases are referred up the health-care hierarchy and an important task is to ensure that these cases are not misdiagnosed or underreferred.

4 What there is around

The idea of using medical expert systems in areas where medical experts are thin on the ground is well known, but practical projects to implement it are few. So far, we have uncovered the systems described below. We make no claims that this is an exhaustive list, although we have made a significant effort to find these.

We have indexed these systems (somewhat arbitrarily) by their country of *development* rather than country of *application*. This is because some systems have been applied in more than one country and others have not been field tested. We have tried to describe them systematically, in terms of language, architecture, interface and so on, but often much of this information is unavailable.

4.1 General overview

Given that there have been so few projects of this type it is difficult to decipher overall trends. A few "issues" emerge however:

Filter systems. Mention is often made of the importance of *filtering* patients for referral and sometimes this is said to be the main task of the primary health worker; to be a diagnostic "frontline" to ensure that the next level of the health-care hierarchy is not swamped with inappropriate referrals and the important cases get through. Taking this to an extreme are systems like (Uplekar et al., 1988) which do not actually make a diagnosis at all but merely classify on a "treat or refer" basis.

Flowcharts. Several medical textbooks based on flowcharts (also known as decision trees) exist for use by health workers in developing countries, the most popular probably being Essex, "Diagnostic Pathways in Clinical Medicine" (Essex, 1980). Many people have criticized the use of flowcharts because they force an ordering on gathering information (e.g. Kastner et al., 1984) although others have used precisely this feature as an asset, suggesting that it encourages systematic history taking. Others criticize the strict path on the grounds that it does not correspond to the "hypothesize and test" method of diagnosis (Auvert et al., 1986b). Whatever their merits and demerits, flowcharts are in use in various forms in many areas and some workers have used them as a basis for the structure of their knowledge base.

Interface. All of these systems have a very simple user-interface and none appear to have a knowledge-base maintenance interface. The simplicity of the interfaces is such that often the entirety of the users' interactions are to select "yes" or "no" in response to a displayed textual question. We feel that concentration on the user interface is one of the most crucial parts of the system if it is to be used in developing countries by people who do not necessarily have *any* familiarity with computers at all. Of course sophisticated graphics are only available on more expensive machines, but surprisingly effective results can be obtained in a simple way, with a little thought (see for example Trappl and Horn's (1983) interface based around a simple anatomical screen model).

4.2 UK

Some ten years ago now Pankhurst (1980) worked on a project to generate diagnostic flowcharts, as in Essex (1980), using an adapted program originally designed to generate biological taxonomic keys. The diagnostic method described here is not itself computational, although the generation of the flowcharts uses a method algorithmically similar to rule induction, as in Quinlan (1979). A novel application suggestion is the use of "polyclaves" for diagnosis in developing countries. A polyclave is basically a set of hollerith cards, one for each observed symptom and each symptom card having a hole for every disease manifesting that symptom. These are combined until only one hole remains common. The polyclave has an advantage over the flowchart in that questions, or symptom considerations, do not have to be strictly ordered.

language: N/A
architecture: classification mechanism for generating flowcharts
interface: N/A
development machine/environment: N/A
delivery machine: N/A
testing: flowcharts trialed in Tanzania with one medical student (2030 cases, 94% agreement with experts) and Guinea Bisau with one doctor (1618 cases, 87% agreement with experts) polyclaves not trialed.
current status: flowchart generation software available for PCs.

The paper describing this work is most interesting for its exposition of the construction and use of flowcharts for use in a developing country (Tanzania).

4.3 France

TROPICAID (Auvert et al. (1986a) and Auvert et al. (1986b)) is a small medical expert system on a hand-held computer. The machine runs on rechargeable (with a solar cell) batteries and each charge can run for 20 continuous hours. The system has a database of drugs (60 essential drugs) with information on indications and contraindications, posology, pharmacological properties etc, diseases (210) with information on clinical description, whether referral is necessary, appropriate drugs etc, and bases its diagnosis on the recognition of symptoms (460). It is flowchart based, the whole system having 464 nodes, 412 questions and 210 diagnoses. Users enter data primarily by choosing from menus or giving yes/no answers. It has been used and tested successfully in Chad in 1984.

language: Pascal
architecture: based on Essex's (Essex, 1980) flowcharts
interface: menus and yes/no questions
development machine/environment: IBM PC
delivery machine: "Basic" hand-held computer and IBM PC
testing: in Chad with 50 doctors and paramedics
current status: Currently being tested in Zaire. Has also been turned into a teaching system, ConsultEAO (Aegerter et al., 1987).

The authors point to several problems, mainly due to the choice of knowledge representation. Firstly, the path is of rigid structure and cannot take into account regional or seasonal variation or avoid questions that are beyond the user's medical knowledge. Secondly, they feel that the diagnostic process should be more closely linked with the treatment module. Thirdly, the graph-type diagnosis does not correspond at all to the "hypothesis-testing" method of diagnosis. Fourthly, the representation is not compact and a larger application might require unwieldy amounts of memory.

4.4 India

SYMPMED (Uplekar et al., 1988) is essentially a *filter* system based on the recognition of common *symptoms* which helps the user decide whether a patient should be sent to a doctor. For each presenting symptom, the aim is to extract the relevant information from the user to decide on a prescription and to make the user aware of associated symptoms or signs which might require referral. There is a symptom file and a drug file, and for each symptom there is a list of drugs which can and cannot be prescribed. On comparison with GP performance, it has been shown to give no false negatives, and some false positives, i.e. to err on the side of safety.

language: unknown
architecture: not clear
interface: not clear, probably yes/no questions
development machine/environment: unknown
delivery machine: IBM compatible PC
testing: with a junior doctor (500 cases) and an engineer (400 cases)
current status: unknown

One of SYMPMED's biggest drawbacks seems to be that it does not actually give a diagnosis. Instead, the user proceeds directly from symptoms to prescription. The authors feel that this is not a problem as they claim

> ". . . diagnosis could be a pointless diversion in the therapeutic process, which is after all what interests the patient."

It may be the case that patients are more interested in *treatment* than diagnosis, but a system like this which in effect skips the diagnostic element of the consultation (a) is giving the impression that it is merely *symptoms* that are being treated rather than any underlying causal disease and (b) will not help the paramedical worker, or whoever else uses the system, to learn anything from using it.

In fact it encourages *dependence* on the machine to hand out a prescription, with no explanation, rather than helping the user towards making diagnoses themselves.

4.5 Thailand

"Medical Tutor" is an ICAI system for teaching the differential diagnosis of anaemia. It helps a student through the diagnostic procedure for one of a number of case histories according to the student's answers. This paper (Eamsiri et al., 1987) is rather difficult to follow but it appears that the student is given a textual case history, and prompted "What do you want to ask?" A keyword search of the responses decides how much the student knows (enough or not enough) by use of a few simple rules like "if the student selects 2 or more strong criteria they understand". The system then prompts "What is your evidence?" and the student must highlight suitable points in the casenotes. The system's responses to this are generated from the scripts dependent on how much it decided the student understands.

language: Turbo-Prolog
architecture: semantic net and scripts
interface: highlighting text, keyword matching of typed input
development machine/environment: Same as delivery machine (we think)
delivery machine: PC
testing: with 20 medical students, 10–20 minutes per student
current status: being extended to include more areas (genetics and haematology)

This is not a diagnostic system but was included here for completeness. It is aimed at *medical students* so is probably not really appropriate for those in primary health care. It is also based on a very small subdomain of medicine.

4.6 USA

Kastner et al. (1984) and Ostroff et al. (1986) developed a small, simple system for primary eye care, again acting as a filter system for referrals to secondary health centres, again designed for eventual use of a hand-held machine. The system was implemented in a rule based shell (EXPERT). It is based on WHO decision strategy guidelines, in which "findings" and rules with certainty factors lead directly to recommendations for treatment. The prototype in Ostroff et al. (1986) has 49 observations, 26 treatments or recommendations and 77 reasoning rules. The authors considered flowcharts but felt that the rule based approach was better because the system builder can

> ". . . experiment with the wide variety of ways in which rules can be combined until the best performance level is achieved. The rule sequences that lead to the desired expert performance can then be translated into a decision algorithm." (Kastner et al., 1984)

language: BASIC
architecture: rule based
interface: numerical, yes/no, checklist and multiple choice questions
development machine/environment: rule-based shell, EXPERT, developed on Vax and Dec mainframes
delivery machine: Apple IIe and Epson HX20
testing: tested with hypothetical cases and real ones collected from California and Yemen. To be trialed in Tunisia by public health nurses and primary physicians, and in Egypt by primary care physicians.
current status: unknown

This system again bypasses the actual diagnosis stage, partly because the authors feel that treatment and prognosis are more important, and partly to simplify the model and presumably make it easier for the limited capacities of the hand-held delivery machine to cope with.

4.7 South Africa

Because of its political and social situation, South Africa presents a unique mixture of first and third world medicine. SASEP is a production-rule system written in BASIC and implemented on Hewlett Packard touchscreen machines. Rules forward chain until either a hypothesis is true with a probability within acceptable confidence limits or no conclusion can be reached within these limits. There is also an associated therapeutic system (CATT: Computerized Aid To Therapy). Again there appears to be no diagnosis of disease, but rather three referral categories, 1-treat, 2-refer and 3-hospitalize (Briggs et al.).

language: BASIC
architecture: rule based
interface: touchscreen with yes/no questions and menus
development machine/environment: same as delivery machine (we think)
delivery machine: Hewlett Packard touchscreen machines
testing: piloted in a provincial hospital with nurses
current status: unknown

Our information on this system is scanty and limited to an unpublished paper (Briggs et al., 198x). It is not an in-depth description of the system and we are still unsure about several issues such as how cumulative probability is calculated, how an explanation facility which the authors allude to works or how the diagnostic system interfaces with the therapeutic.

4.8 Austria

Porenta et al. (1988) "Village Health Worker" is a decision support system built using a hybrid knowledge-engineering tool (VIE-KET, said to ressemble KEE), based on five different problem (disease) areas. Two different knowledge representation architectures are used for these five subdomains. In expanded decision networks, branching at each node is determined by the relation between a threshold value and the sum of weights of symptoms e.g. one "is present" or two "maybe-presents" cause a yes branch. Rules with certainty factors are also employed for different subdomains (it is not clear why the different representations were used for the different domains). Designed from limited experience with village health work in Northern Ethiopia but not yet field tested.

language: IQlisp
architecture: decision trees and rule based with certainty factors
interface: yes/no questions or choice from numbered menu
development machine/environment: knowledge engineering tool, VIE/KET
delivery machine: IBM compatible PC
testing: untested
current status: prototype

Funding ceased on the project in 1987 and despite promising beginnings the workers were unable to secure support for field testing.

5 Appropriate Technology: what is it anyway?

As mentioned above in §3, we feel that one of the main reasons for the dearth of medical expert systems for developing countries is that they are considered to be "inappropriate technology". This view is widespread. The governments of developing countries may be understandably unwilling to become involved in projects which, although largely aid-funded, could use up scarce local expertise resources if the projects turn out to be insupportable. Aid agencies are also sensibly unwilling to fund enterprises which seem inappropriate. The academic community, more familiar with the high-tech applications of expert systems, find it difficult to visualize their use in less "technological" environments.

We here review the notion of appropriate technology, looking to various sources for principled definitions. The criteria range from the very general; what is appropriate technology at all (Eckhaus, 1977), through what is appropriate technology for medicine (Committee, 1988) and (Association, 1985), to computer-specific criteria (Machanik, 1988). From these we have produced a general list of criteria for appropriate technology which can be applied to computing systems, and to medical expert systems in particular. These criteria are than used as principles of *design* in §6.

5.1 NAS report

Richard Eckhaus in the National Academy of Sciences' report "Appropriate Technologies for Developing Countries" (Eckhaus, 1977) notes that a number of different criteria for appropriate technology in general have been proposed including

1. maximization of output, or cost minimization
2. maximization of the availability of consumption goods
3. maximization of the rate of economic growth
4. reduction in unemployment, or maximizing employment
5. regional development
6. reduction in balance of payments deficits
7. greater equity in the distribution of income
8. promotion of political development (including national self-reliance)
9. improvement in the quality of life

Many of these criteria are only partially relevant to our domain as they relate primarily to economic and political issues such as employment and means of production. Those which seem most relevant to health care are:

Cost Minimization: although any cost restriction must be balanced against resultant benefits. Obviously, choosing the cheaper of two equally effective methodologies is a rational strategy, hence the proliferation of generic drug programs. Note that below other authors qualify the "cost restriction" criterion into "cost effectiveness".

Improvement in Quality of Life: Health care is about improving quality of life, and maintaining life itself. As a criterion for appropriateness in health care this is really a requirement that the method be *effective*.

5.2 WHO report on appropriate technology for cardiovascular diseases

The report of a WHO expert committee on appropriate technology in the management of cardiovascular diseases (Committee, 1988) suggests that appropriate health technology means methods procedures techniques and equipment that are

1. scientifically valid
2. adapted to local needs
3. acceptable both to those who use them and to those for whom they are used
4. can be acquired, maintained and utilized with resources the community or country can afford

This is very similar to the WHO description of primary health care (WHO, 1978) (see quotation in §1), suggesting that *primary* health care and *appropriate* health technology are in fact very similar creatures, and indeed they note

"Whether or not a particular diagnostic technology is appropriate is thus dependent partly upon the health care level—primary, secondary or tertiary—of the proposed setting and partly upon the degree of development of, and the availability of resources for, the overall health care system"

5.3 Appropriate technology: articles from the British Medical Journal

In 1985 the British Medical Journal published a collection of articles on appropriate technology for medicine in developing countries (Association, 1985) and in the introduction by Katherine Elliott, six criteria are proposed:

1. Effectiveness: it must work and fulfil its purpose in the circumstances in which it needs to be used.
2. Cultural acceptability: it must fit in with the users (although new ideas may stimulate local enterprise and promote self sufficiency).
3. Affordability: it doesn't have to be *cheap* but tradeoffs between cost and effectiveness must be considered.
4. Local sustainability: it shouldn't be over-dependent on imported skill for function, maintenance and repair.
5. Measurability: impact and performance needs proper evaluation.
6. Political responsibility: it shouldn't alter an existing balance in ways that might be counterproductive.

5.4 Criteria for medical education software

Machanik (1988) narrows things down somewhat by defining computer-specific criteria for appropriate technology, hardware and software, based on some originally general criteria proposed by Jequier and Blanc (1979). This analysis is as a prologue to a project designing medical education software

1. hardware should be low in cost
2. multiuser or networked configurations should be avoided
3. the software should be easy to use
4. minimal training should be required
5. the user interface should readily adapt to new languages
6. the functionality should easily adapt to new contexts
7. the approach should adapt to a wide range of people and localities
8. the cost per user should be low

5.5 A collation

These concerns seem to be mainly about effectiveness, cost effectiveness, usability, adaptability, supportability and local relevance. Our feeling is that for an advanced technology such as AI even more emphasis should be placed on the idiosyncratic requirements of local application and maintenance. We have added the criterion of *usefulness* and expanded *relevance* into relevance both of form and content. We propose the following as a usable set of criteria for Appropriate Technology, which we will define and attempt to apply to the use of medical AI systems for developing countries.

1. Effectiveness: it should function adequately and measurably in the environment for which it is designed.
2. Usefulness: it should answer an existing need.
3. Usability: it should be "easy to use" and should make the process it is designed to assist less complex or time consuming.

4. Supportability: it should be supportable within the existing infrastructure and economy.
5. Adaptability: it should be adaptable to changing requirements and circumstances.
6. Cost effectiveness: the cost of the system should be offset by measurable improvements in health care.
7. Relevance of content: the knowledge base should be locally relevant and rooted in familiar facts (e.g. disease profiles, drug availability).
8. Relevance of form: the interface should be similarly related to familiar things (e.g. local training manuals).
9. Minimal training: it should be as "obvious" as possible and such training as is required should fit in with existing e.g. nurse training.

6 Design Issues: making the tool fit the job

In this section we enlarge upon each of the criteria defined in §5.5 and try to show how they should directly influence the design of not only the software/system itself, but also the surrounding project to implement it (projects need to be designed too).

6.1 Effectiveness

It should function adequately in the environment for which it is designed
This means that there must be a substantial evaluation component in the project, to see if the thing does actually work. Evaluation is of course not all done after the event (summative evaluation). Testing throughout the design process with potential users (formative evaluation) gives valuable early clues about where things are going wrong, and right. It may be worth investigating, even investing in, an expert systems *development environment* to do fast *prototyping* on.

6.2 Usefulness

It should answer an existing need
The application of AI in developing countries should not be undertaken merely to satisfy a researcher's whim. Too many projects in developing countries have the bad image of being developed exclusively in the first world for application in the third. In medicine particularly it is important not to have useless and expensive equipment "plonked" into overworked health centres. To avoid this developers should consider the *functionality* of the proposed system in its setting; is it for epidemiological data-gathering, record keeping, education and training, diagnosis, decision support, access to medical databases (e.g. drugs, diseases)? Which of these (or other) applications would be most useful?

6.3 Usability

It should make the process it is designed to assist less complex or time consuming
In a clinical diagnostic situation any computational support must be *usable*, otherwise it won't be used. Clinical decisions have to be made rapidly, and anything which hinders/interferes with the clinical decision making process rather than supporting it will be swiftly rejected. Usability is notoriously difficult to define, but in the design of the *interface* general HCI principles obtain:

- displays should be clear, effective and consistent
- there should be a simple means of *inputting* information
- feedback should be rapid
- failure, if it happens, should be graceful
- there should be help and recovery support

6.4 Supportability

It should be supportable within the existing infrastructure and economy
This refers to both hardware and software. The most obvious comment to make about hardware is that if there is a national computing policy this should be adhered to as standardization of hardware makes support and maintenance easier. Other considerations would be is there a reliable local electricity supply? If so, it's probably best to use a mains driven rather than battery-operated machine. Batteries are often expensive and difficult to get hold of, especially if they are not manufactured in the country. Consideration should be given to the environment in which the machine will be used. Is there enough space for, say, something the size of a PC? Will it be needed, or used, for outreach work, in which case it will have to be portable? Are there extremes of temperature, humidity or dust? What sort of machines are available in the country anyway? If you use something like everybody else does there are more likely to be local people who can mend them; demand creates supply.

With regard to software, maintenance of the system is very much bound up with *adaptability* (below), and as mentioned in §3.4, the maintenance should be "planned into" the design.

6.5 Adaptability

It should be adaptable to changing requirements and circumstances
The knowledge base of an Expert System, particularly in medicine, should not be a static entity. The disease profile in the country may change, new diseases can emerge (witness the emergence and spread of AIDS), others be eradicated (e.g. smallpox). Drugs are particularly transient, new drugs are developed (although new drugs tend to take a long time to find their way into developing countries) and others become obsolete, drug resistant organisms develop, requiring a change in prescribing practice. If these changes were not reflected in the diagnostic system it would cease to reflect current diagnostic and prescribing practice and so become obsolete. There should be some fairly accessible method of updating the knowledge base available to relatively untrained domain experts. Knowledge base updating should not require the skills of the developer.

6.6 Cost effectiveness

The cost of the system should be offset by measurable improvements in health care
Cost effectiveness, as previously suggested, is a very difficult entity to measure. For a system to be cost effective, planners would be looking for one or more of the following:

- Improvement in diagnostic accuracy.
- Improvement in referral rate.
- Reduction in drug overprescription, and, related, prevention of emergence of drug resistant organisms.
- Recommendation of therapy using generic drugs.
- More efficient use of lab facilities.
- Reduce administrative load on health workers.

6.7 Relevance of content and relevance of form

The knowledge base should be locally relevant and rooted in familiar facts (e.g. disease profiles, drug availability). The interface should be similarly related to familiar things (e.g. local training manuals)
In planning a project of this type, perhaps the major effort should go into ensuring the appropriateness of the proposed system by tailoring to local requirements. A preliminary investigation should attempt to answer such questions as What's local training like? What are local

disease profiles? What drugs are available? How are people coping at the moment? (they may be doing very well thank you) What are people's attitudes to computers? Would the machine make the gap between diagnostician and patient even greater?

The results of this survey should feed into the construction of the knowledge base and the design of the interface. If a generic drug list is available this can be used as a database (although some investigation should also be conducted into how many of these drugs are actually available in the dispensaries). Attention should be paid to existing practices in the clinical setting so that, for example, the system does not ask for data which is not available or beyond the scope of extant facilities. Local training manuals will give clues as to how history taking is taught, so that any "history taking" required by the machine can be prompted for in the same manner.

6.8 Minimal training

It should be as "obvious" as possible and such training as is required should fit in with existing e.g. nurse training

This is a catch-all. Satisfying the requirements of Usability, Relevance of Content and Relevance of Form should ensure satisfaction of this criterion.

7 Conclusion

We feel that AI *does* hold out promise of a solution to the problem of providing "health for all" in remote, poor or developing areas. Expert Systems technology has proven itself to be reliable and useful in other than academic environments. The lack of medical AI development projects can be attributed to various factors, one of the most important being that people feel AI to be "inappropriate technology" for the developing world. On the contrary, we feel that it is very appropriate, providing expertise and support in areas where it is most needed. However, it is crucial for the success of any such projects that the circumstances of the environment for which the system is to be produced are taken into account, in other words, that its designers build *appropriateness* into their product.

References

Aegerter, P, Auvert, B, Gilbos, V, Desve, G, Benillouche, E, Landre, MF and Bos, D, 1987. "A computer-assisted instruction system designed for rural health workers in developing countries". In: *Proceedings of Conference on Medical Informatics in Europe 1987.* Rome.

Auvert, B, Aegerter, P, Gilbos, V, Benillouche, E, Boutin, P, Desve, G, Landre, MF and Bos, D, 1986. "TROPICAID: a portable expert system for medical decision-aid in developing countries". In: *MEDINFO 86.* Amsterdam: Elsevier Science Publishers.

Auvert, B, Aegerter, P, van Look, F, Huong-Du, LT, Boutin, P, Monier, JL, Emmanuelli, X, Gilbos, V and Benillouche, E, 1986. "A hand-held decision aid system designed for rural health workers. *Computers and Biomedical Research* 19.

Briggs, BR, Lavett, SL, Pistorius, JG, van den Berg, ADP, Prax, M and de Klerk, JN. "SASEP a computerised diagnostic program". Unpublished paper.

British Medical Association, 1985. *Appropriate Technology: Articles from the British Medical Journal.* London: British Medical Association.

Clancey, WJ and Shortliffe, EH, 1984. *Readings in Medical Artificial Intelligence: The First Decade.* Addison Wesley.

Eamsiri, J, Malasit, P, Songsivilia, S and Chongstitvatana, P, 1987. "Intelligent tutor for medical teaching". In: *Proceedings of the Regional Symposium on Computer Science and its Applications.* Bangkok.

Eckhaus, RS, 1977. *Applied Technology for Developing Countries.* Washington DC: National Academy of Sciences. Prepared for the panel on Appropriate Technologies for Developing Countries.

Essex, BJ, 1980. *Diagnostic Pathways in Clinical Medicine.* Edinburgh: Churchill Livingstone.

Jequier, N and Blanc, G, 1979. *Appropriate Technology Directory.* Paris: OECD.

Kastner, JK, Dawson, CR, Weiss, SM, Kern, KB and Kulikowski, CA, 1984. "An expert consultation system for frontline health workers in primary eye care". *Journal of Medical Systems* **8**(5) 389–397.

Machanik, P, 1988. "Design of medical education software as appropriate technology using artificial intelligence and software engineering". *Technical Report 88-01*. Johannesburg: University of the Witwatersrand.

Miller, RA, Pople, HE and Myers, JD, 1982. "INTERNIST-I: an experimental computer-based diagnostic consultant for general internal medicine". *New England Journal of Medicine* 307(8).

Ostroff, JH, Dawson, CR, Kastner, JK, Weiss, SM, Kulikowski, CA and Kern, KB, 1986. "Preliminary results from field testing of an expert advisory system for primary eye care in developing countries". In: *AI and Advanced Computer Technology, Proceedings of the 2nd Conference*. Illinois: Wheaton.

Pankhurst, RJ, 1980. "Medical diagnosis in developing countries". *Computers in Biology and Medicine* **10**(69) 82.

Porenta, G, Pfahringer, P, Hoberstorfer, M and Trappl, R, 1988. "A decision support system for village health workers in developing countries". *Applied Artificial Intelligence* 2.

Quinlan, R, 1979. "Discovering rules from large sets of examples: a case study". In: D Michie (ed.), *Expert Systems in the Microelectronic Age*. Edinburgh: Elsevier Science Publishers.

Trappl, R and Horn, W, 1983. "Making interaction with a medical expert system easier". In: *MEDINFO 83*. Amsterdam: Elsevier Science Publishers.

Uplekar, MW, Antia, NH and Dhumale, PS, 1988. "Sympmed1: computer program for primary health care". *British Medical Journal* 297.

WHO, 1978. "Primary health care: report on the international conference on primary health care". Alma Ata USSR 1978. *Health for All Series 1*.

WHO Expert Committee, 1988. "Appropriate technology in the management of cardiovascular diseases". *Technical Report 772*. Geneva: World Health Organisation.