

# AKT-R4: Supporting the Exploration of a Knowledge Web

Andy Aiken, Derek Sleeman

Department of Computing Science,  
King's College, University of Aberdeen,  
Aberdeen, Scotland, AB24 3UE  
{aiken, dsleeman}@abdn.ac.uk

**Abstract.** This paper outlines a proposal for an expert system, AKT-R4, which will be the fourth iteration of the Refiner series. Unlike previous iterations, AKT-R4 will include a non-deterministic feature which will enable it to better handle noisy datasets by generating category descriptions with 'fuzzy' boundaries. The AKT-R4 system will be based on a case-based algorithm focussed on performing a classification task, in particular medical (differential) diagnosis; AKT-R4 will be aimed primarily at medical students and junior doctors (i.e., novice diagnosticians). The AKT-R4 algorithm will incorporate concepts from case-based reasoning, hypothetico-deductive reasoning (HDR) and illness scripts, and will perform diagnosis by means of a diagnostic web, which is an expansion of the illness script concept. This system supports the novice user as they develop and navigate through a medical knowledge web.

## Introduction

### Differential Diagnosis

Differential diagnosis is the name given to the principal method used by clinicians for performing diagnosis. The differential diagnoses for a disease are the collection of other diseases that have symptoms in common with it. To put this in context, consider the following scenario. A clinician sees a patient who presents a set of symptoms which suggest a number of potential diagnoses; the clinician must use his or her experience and judgement to determine which of these potential diagnoses might be the likely correct ones.

The complete candidate set of potential diagnoses includes every disease which has at least one symptom in common with the set of presented symptoms. It is unlikely that any clinician, no matter how experienced, will be able to consider each potential diagnosis in this set. This task is made more difficult by the fact that particularly salient symptoms might not be presented initially.

### Hypothetico-Deductive Reasoning

Hypothetico-deductive reasoning (HDR) is an iterative process by which an expert arrives at a diagnostic hypothesis given a collection of indicators. It incorporates both forward and backward chaining. The process of HDR takes the following general format:

- An initial set of indicators are presented to the expert.
- These indicators give rise to an initial set of hypotheses (forward chaining).

- Each of these hypotheses can be corroborated or refuted by looking for additional indicators (backward chaining).
- The expert decides which of these additional indicators to investigate; this updates the list of indicators and the process begins a new iteration.

The process repeats until all but one hypothesis has been ruled out, or the expert decides that one of the remaining hypotheses is significantly more likely than the others.

### **Illness Scripts**

Beyond HDR there seems to be little homogeneity in clinical reasoning in different situations and domains, and Feltovich [5] suggests that this may be due to the lack of a common problem representation. The suggested general problem representation for medicine, the illness script, contains the following components:

- **Enabling conditions** are features associated with the acquisition of the illness, such as predisposing factors (the patient's case history, for example), boundary conditions (i.e. age and gender), and hereditary factors.
- **Faults** are the underlying physical malfunctions which give rise to the disease.
- **Consequences** are the symptoms of the disease.

### **The Effect of Experience**

How do experts and novices differ? Traditional maxims about expertise include the following observations [6]:

- Experts generally reason from symptoms to a diagnostic class.
- Experts tend to form more abstract problem representations than novices.
- Experts have more knowledge than novices but do not necessarily have more ability to process this knowledge.
- Experts must see many cases over a number of years to become expert, and this expertise needs to be maintained.

This leads us to consider how the process of diagnosis varies with expertise.

- Novices use HDR, but their knowledge tends to be incomplete / inaccurate so the process is incomplete and hence potentially flawed.
- Intermediates perform HDR.
- Experts still use HDR, but at this level of experience it is supported by a corpus of compiled case knowledge (i.e., illness scripts) which provides 'shortcuts' to likely diagnoses; this reduces the amount of processing required.

### **Clinical Decision-Support Systems**

Clinical decision-support (CDS) systems such as Internist-1, the Quick Medical Reference (QMR) [7], DXplain [2-4] and other systems [10, 11] were designed to collate findings from a case and respond with a list of potential diagnoses which would explain these findings. Although these systems tend to perform remarkably well, they are built on large databases which are difficult to develop and maintain.

Payne [8] suggests that successful clinical decision support systems must satisfy three conditions:

- They should give patient-specific recommendations.

- They should save time for the user.
- They should be incorporated into the organisation's workflow.

Some common issues and concerns when implementing a diagnostic CDS system are:

- Using the system should not require additional time, although this may be the case initially as the users become familiar with it.
- Particularly in medical domains, decision support systems should be designed to maintain patient confidentiality.
- The system's diagnostic accuracy is obviously of key importance. The correct diagnosis should be suggested for all cases, and other potential diagnoses, if they are presented at all, should not be ranked higher than the correct diagnosis.

The Quick Medical Reference (QMR) has an array of features useful for clinical diagnosis and computer-assisted learning (CAL):

- It can generate a differential diagnosis, given a set of clinical findings. It presents the user with a list of potential diagnoses, each with a score suggesting the likelihood of the diagnosis.
- It can offer information on a wide range of diseases (over 600 in 1999), including associated diseases and common complications.
- It can offer strategies to confirm or exclude disorders (by suggesting salient tests to confirm or exclude the diagnosis, for example), list other related diseases, and highlight the findings which are consistent or inconsistent with the diagnosis.
- It can provide simulated cases for educational purposes.

DXplain [2-4] is a system which provides access to a database of differential diagnoses. Its goals are to remind the physician of diseases that should be considered, and to provide information about diseases that can be eliminated from consideration. It is not designed to determine the single 'correct' diagnosis, as a clinician will always have a more complete picture of the patient than will be entered into the system.

DXplain allows the clinician to enter a set of findings, displaying the most likely diagnostic hypotheses (separated into 'common' and 'rare' diseases) and also suggesting a list of additional findings which might be used to discriminate between them, thus enabling an HDR-style interaction.

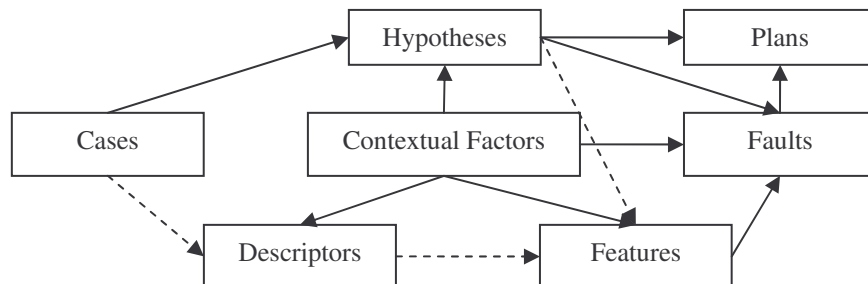
### The Refiner Series

The basic philosophy of the original Refiner system [9] is that it is better to detect errors and inconsistencies in sets of cases before a case-based reasoning (CBR) session than during one. Refiner reads in a set of labelled cases (i.e. cases for which the domain expert had specified the class to which it belongs) and background knowledge if it is available, and the system generates a generalised description for each of the classes. The underlying assumption of the Refiner series is that each category should have a unique and unambiguous description – that is, there should be no overlap between categories, and each case should only match the category it has been assigned to by the expert. This disambiguation process is performed incrementally, i.e. each new case causes each of the class descriptions to be recalculated. Moreover, at each stage if the case base is inconsistent e.g., a case can be classified as belonging to more than one class, this fact is flagged and Refiner suggests a number of ways in which that latest inconsistency could be removed. The system requires the domain expert to choose the appropriate modification.

As Refiner asked the domain expert to remove an inconsistency as soon as it arose, this led to changes being made at one “cycle” which would sometimes be “undone” during a later “cycle”, as the domain expert did not have a good overview of the data-set. Therefore, it was decided to re-implement the algorithm as a batch algorithm, which was named Refiner+ [12]. Refiner+ reads all the cases before it creates any of the class descriptions or detects any inconsistencies. The great advantage of this approach is that the algorithm now has an overview of the data sets and so can suggest an ordered set of changes (that is, changes which are likely to remove a sizable number of inconsistencies are listed before those which only remove a single inconsistency). The third iteration of this series, Refiner++ [1], came about when the system was reimplemented in Java (as opposed to the original LISP). A number of additional changes were made to the algorithm at this time. The principal strength of the Refiner++ system seems to be in domain elucidation, but it does not tend to deal well with real-life datasets as it has a limited number of ways to deal with noisy data.

## Diagnostic Web

After discussions with our domain experts, and studying the various existing CDS systems, we have concluded that the concept of the illness script represents a simplified model of the cognitive knowledge sources which experts utilise when performing diagnosis. Therefore, we have formulated a more inclusive concept: the Diagnostic Web.



We can summarise this diagram as follows:

- An expert’s knowledge is derived largely from **cases** (and also from book learning, but this is often presented in terms of real or hypothetical cases). This is perhaps the most significant difference between a novice and an expert; experts have a larger, more complete casebase to use when performing diagnosis.
- Each case has values for a set of **descriptors** (or fields). These values can be of a range of data types. Cases are unlikely to have values for all descriptors; more likely, a case will have values for a small subset of particularly significant descriptors. Potential descriptors in a medical diagnosis application might include blood pressure, heart rate, and respiratory rate.
- Each case is associated with one or more **hypotheses** (or categories). Hypotheses are represented in a hierarchy; when a hypothesis is associated with a case, we assume that the ancestors of the hypothesis are also applicable to the case. In the case of medical diagnosis, hypotheses represent the diseases and conditions that patients might suffer from.

- Descriptor values are combined to form **features**, which are analogous to illness script consequences. A feature, then, is a compilation (or generalisation) of all the values the expert has encountered for a particular descriptor for a particular hypothesis. The features for each descriptor for a hypothesis make up the model of that hypothesis (in other words, the hypothesis is described by these features); new cases can be compared to this model to determine their degree of similarity to other cases with that hypothesis, and hence whether the descriptor values in the new case are indicative of that hypothesis.
- Features illuminate **faults**, which are the underlying physical malfunctions which give rise to the problem to be diagnosed, such as bacteria invading an organ or a blood vessel rupturing. These are analogous to illness script faults. Each hypothesis can be associated with one or more faults.
- Hypothesis and faults can be associated with one or more **plans**. This allows the expert to develop a management plan for the case, once a diagnostic hypothesis has been confirmed. For example, in medical diagnosis a course of action might specify the type of drug to administer.
- **Contextual factors**, which are analogous to illness script enabling conditions, allow the expert to modify certain aspects of the diagnostic process (for example, the relevance of certain features to a hypothesis, or the probability of a particular hypothesis or fault being present) given certain background knowledge about the case. In the medical domain, these contextual factors might include age, gender, current medication, and case history. For example, if the patient has been in contact with someone suffering from meningitis, the chance that the patient has also contracted the disease is much greater than it would otherwise be.

New cases can be diagnosed in the following manner:

- Given the initial presenting symptoms, case history, etc., the clinician creates a number of diagnostic hypotheses, which might be influenced by contextual factors. The expert then chooses one of these hypotheses to confirm or rule out.
- The new case contains values for some descriptors. Each value is compared to that descriptor's feature for the hypothesis. The result of these comparisons allows the expert to determine how good a match there is between the case and the hypothesis. This process is indicated by the dashed path on the diagram.
- This comparison can be repeated for all the hypotheses in the diagnostic web (of particular interest to the expert might be the hypotheses that have faults in common). The hypothesis showing the best degree of fit with the new case is deemed the most likely diagnosis.
- The user can then engage in a HDR process by obtaining values for other descriptors, repeating the process until no significant differences are found from iteration to iteration.

## Proposed Work

During interviews with experts [1], it was noted that the Refiner++ system did not cope well with noisy data as the algorithm classified cases absolutely; there is no facility for incorporating any measure of uncertainty in this regard. The only way in which Refiner++ is able to capture uncertainty is in its category descriptions; features are

described by a conjunction of case values, producing a range of values. We therefore intend to develop a new system, AKT-R4, to address this issue.

Adding the facility to handle noisy datasets means that we lose the terse feature descriptions of Refiner++ (e.g. 'Age: 50 – 56') in favour of less easily-described generalisations. This in turn leads to more complex category descriptions which, although more useful for performing classification, are much more difficult to display to a user succinctly. Whereas the aim of the previous Refiner systems was the creation of unambiguous category descriptions from case data, this addition of uncertainty changes the focus of AKT-R4 to the accurate classification of new cases.

We therefore intend to develop AKT-R4 as a user-directed expert system to aid (primarily) junior doctors and medical students with the task of diagnosis. We intend to design the system such that it could be used on a desktop machine, over the web, or on a handheld device such as a PDA. In keeping with the Refiner series, AKT-R4 is intended to be directly usable by the domain experts themselves.

The proposed AKT-R4 diagnosis algorithm extends part of the existing Refiner++ algorithm and also employs useful concepts from both HDR and illness scripts. We have chosen to integrate these algorithms into AKT-R4 because we feel that if the program closely mirrors the way in which experts approach this class of tasks, we will see a number of benefits:

- Most obviously, the program should be easier for them to use.
  - The experts will accept the program more readily than they otherwise might.
  - The program's results are likely to more accurately reflect the experts' expectations.
- The Refiner++ algorithm lends itself well to HDR as it shares its iterative component. In Refiner++, the category descriptions are iteratively refined, each time resulting in a more accurate model and a new set of suggested refinement strategies. In HDR, a set of initial observed findings leads to a set of potential classifications, which in turn gives rise to a set of potential findings (findings that would support or preclude the potential classifications), which update the list of observed findings for the next iteration. It is our intention in developing AKT-R4 that the diagnosis profiles generated by the algorithm can be seen as being analogous to illness scripts. The three components of an illness script, and their AKT-R4 equivalents, are:

- **Enabling conditions**, or predisposing factors, which are used in illness scripts as a form of background knowledge, adding information about the context in which each case is encountered and providing shortcuts to particular diagnoses when certain symptom 'triggers' are met. The issue of implementing enabling conditions is discussed in more detail below.
- **Faults**, the physical malfunctions giving rise to the symptoms, which are used in illness scripts to link two or more potential diseases together; in AKT-R4 this role is filled by the diagnosis hierarchy, which shows the relationships between diagnoses.
- **Consequences**, which are analogous to AKT-R4's descriptor values when they are combined to form features.

To properly mirror the illness script concept, we must implement enabling conditions in some way. The difficulty with this is that these are difficult to come by without a large KA effort (getting experts to come up with the rules, fine-tuning the values given by experts, etc). On the other hand, having these enabling conditions would make the diagnosis (a) more accurate, as they might take into account factors which aren't

present in the dataset and (b) potentially faster, as they might allow us to discount certain diagnoses and thus reduce the algorithm's search space.

Some enabling conditions are well known for certain conditions. For example, alcoholism increases blood pressure, so a high value in this descriptor might not be as indicative for patients with alcoholism. For a simpler example, the male gender precludes pregnancy.

Our current plan is to implement enabling conditions as a collection of *contextual factors* which would be able to perform a variety of modification functions, such as:

- Changing the probability of a diagnosis, or discounting a diagnosis completely.
- Discounting the contribution of a descriptor to the diagnosis, or making its contribution more significant.

The user would then be able to switch these modifiers on for individual cases. The difficulty lies in the question of how the system would populate its collection of these modifiers. This body of knowledge clearly influences diagnostic decisions and could comprise a system-independent knowledge source, once formulated.

#### AKT-R4 and Other Systems

In terms of exposed functionality, AKT-R4 and other clinical decision-support systems, perhaps especially QMR and DXplain, are superficially very similar: HDR-style reasoning with constantly updated lists of likely diagnoses and suggested discriminatory and confirmatory features to investigate.

The algorithms for these systems all rely on the existence of a large, complex knowledge base, which has been built over time at presumably great expense of money and effort, and requires constant modification and updating. In contrast, the AKT-R4 algorithm uses a set of cases which require no knowledge engineer to maintain.

Although these systems could potentially be used to great effect in other domains, such as the diagnosis of faults in electrical or mechanical equipment, the creation of a knowledge base to perform this work would be costly and time-consuming. By contrast, AKT-R4 requires no explicit domain knowledge, and could presumably perform equally accurately on problems in any domain, given an appropriate set of cases and associated knowledge of contextual factors.

Whereas the terms used by these systems are general, AKT-R4 will be able to deal with specific values – that is, where QMR or DXplain might talk about a patient with the 'high blood pressure' finding, AKT-R4 would be able to talk about a case with a value of 100mmHg for the 'diastolic blood pressure' descriptor.

#### Unique Contribution of AKT-R4

The unique contribution of AKT-R4 to decision-support systems and intelligent tutoring systems can be summed up as follows:

- AKT-R4 will be the only diagnostic decision-support system to use a case base rather than a manually created knowledge base, which allows it to be used on multiple domains without an expensive knowledge acquisition effort. In addition, this approach should make comorbidity much easier to diagnose.
- AKT-R4 represents a merging of the concepts of HDR and illness scripts, allowing it to perform the task of differential diagnosis analogously to an expert. Mirroring the cognitive process of diagnosis in this way should result in a useful diagnosis

tutoring system. The addition of a robust problem representation (the illness script) and a weak PSM (HDR) should result in a much stronger problem solver.

- AKT-R4 is intended to contain information about the recommended course of action for each disease, which will allow the user to create management plans for patients.
- Unlike other clinical decision support systems, AKT-R4 will be aimed at novice diagnosticians, and is intended to scaffold their understanding of how diagnosis should be performed, rather than as a support for experienced clinicians. It would be possible, for example, to introduce a series of graded cases which gradually increase the students' exposure to various aspects of diagnosis.
- To complement its set of cases, AKT-R4 will also make use of a unique collection of contextual factors to modify the diagnosis process. Once elicited, this knowledge source could be used in conjunction with other diagnostic programs.

## References

1. Aiken, A., Sleeman, D., "Refiner++: A Knowledge Acquisition and Refinement Tool", in *Capturing Knowledge from Domain Experts*, KCAP-03 (2003).
2. Barnett, G. O., Famiglietti, K. T., Kim, R. J., Hoffer, E. P., Feldman. M. J., "DXplain on the Internet", *Proceedings of the 1998 AMIA Annual Fall Symposium*, Nov 7-11 1998, Orlando, FL. 607-611 (1998).
3. Barnett, G.O., Cimino, J. J, Hupp, J. A., "DXplain – An Evolving Diagnostic Decision Support System", *JAMA* 258, 67 – 74 (1987).
4. Barnett, G. O., Cimino, J. J., Hupp, J. A., Hoffer, E. P., "DXplain: Experience with Knowledge Acquisition and Program Evaluation", in Stead, W. W. (ed), *Proceedings of the Eleventh Annual Symposium on Computer Applications in Medical Care*, Washington, DC, 150 – 154 (1987),
5. Feltovich, P. J., and Barrows, H. S., "Issues of generality in medical problem solving", in H. G. Schmidt and M. L. De Volder (eds) *Tutorials in Problem-Based Learning*, the Netherlands, Van Gorcum (1984)
6. Gilhooly, K.J., "Cognitive Psychology and Medical Diagnosis", *Applied Cognitive Psychol*, 4:261-272 (1990).
7. Lemaire, J. B., Schaefer, J. P., Martin, L. A., et al, "Effectiveness of the Quick Medical Reference as a Diagnostic Tool", *CMAJ*, 61(6), 725 – 8, (1999).
8. Payne, T. H., "Computer Decision Support Systems", *Chest*, 118(90020), 47S – 52 (2000).
9. Sharma, S., Sleeman, D., "Refiner: A Case-Based Differential Diagnosis Aide for Knowledge Acquisition and Knowledge Refinement", *EWSL* 1988, 201-210.
10. Warner, H., Haug, B., Bouhaddou, O., et al, "ILIAD as an Expert Consultant to Teach Differential Diagnosis", in Greenes, R. (ed), *Proceedings of the Twelfth Annual Symposium on Computer Applications in Medical Care*, Washington DC, IEEE (1998).
11. Waxman, H., S., Worley, W., E., "Computer-Assisted Medical Diagnosis: Subject Review and Evaluation of a New Microcomputer-Based System", *Medicine (Baltimore)*, 69 (3), 125 – 136 (1990).
12. Winter, M., Sleeman, D., "Refiner+: An Efficient System for Detecting and Removing Inconsistencies in Example Sets", *Research & Development in Expert Systems XII*, 115-132 (1995).