Expert Systems for Clinical Pathology Reporting

Glenn A Edwards
Department of Biochemistry, St John of God Pathology, Wembley WA 6913, Australia.
For correspondence: Dr Glenn Edwards e-mail: Glenn.Edwards@sjog.org.au

Summary

- Conventional automated interpretative reporting systems use standard or "canned" comments for patient reports. These are result-specific and do not generally refer to the patient context.
- Laboratory information systems (LIS) are limited in their application of patient-specific content of reporting.
- Patient-specific interpretation requires extensive cross-referencing to other information contained in the LIS such as previous test results, other related tests, and clinical notes, both current and previous.
- Expert systems have the potential to improve reporting quality by enabling patient-specific reporting in clinical laboratories.

Introduction

After considerable success in various non-health industries, expert system technologies are beginning to gain some traction in medicine, particularly in laboratory medicine. Here I will focus on the use of expert systems to generate interpretative comments for clinical laboratory reports. There are other roles for expert systems in reporting (e.g. managing QC and auto-verification) but I will not consider these further.

It should be noted that this is an emerging field in laboratory medicine. While LIS often allow the creation of rules, which in some cases are used for interpretative reporting, these are usually of limited capability and do not merit the designation "expert system" (indeed, most LIS vendors sensibly do not use this term).

The situation with true expert systems in routine use in Australian clinical laboratories is less clear. The best known of these is the LabWizard system. Developed by an Australian company Pacific Knowledge Systems,[1] LabWizard uses the novel "knowledge acquisition" methodology Ripple Down Rules (RDR) and was specifically designed to overcome some of the technical obstacles facing rules-based systems.[2] Some of the larger commercial laboratories in Australia also operate proprietary software products that support interpretative reporting.

Finally, there have been numerous research reports over the years of specialised, dedicated expert systems for interpreting niche domains in pathology. As very few of these systems have migrated outside the research environment they are not generally available for routine laboratory use.

Current Interpretative Reporting Systems

The typical clinical laboratory handles large volumes of tests and has expectations of rapid turnaround. These constraints preclude manual review of all but a small selection of reports. Most reports will be issued uninterpreted, or with a standard or "canned" comment appended by the laboratory information system.

However, there is growing awareness that doctors and patients need more patient-specific interpretation, in order to improve the clinical outcomes of laboratory testing.[3,4] To provide a more comprehensive interpretative service, some level of automation is required. While LIS rules modules have some role in this, truly individualised patient-specific interpretation of most reports requires the application of expert system technology.

LIS rules

"Standard" Comments

Many biochemists are familiar with the use of programmed "rules", which are used to determine an action. Regardless of the actual programming methodology used, these rules are of the general form:

IF (conditions)
THEN (action)

Most biochemists would not write rules in this programmatic way, however. Instead, they more commonly use an application's tables or forms to manage conditions for actions, for example, criteria for alert or panic levels.

This form of rule management is found in most LIS systems, and more recently in middleware applications. It is sometimes extended to the application of automated comments for patient reports ("standard" comments).

Typically, standard comments apply to a single test result (or a very restricted set of test results). Standard comments, which by their nature are "result-specific", are appropriate *when no reference to the patient context is required*. Common examples include notification of changes to a method or reference interval, or where a specific reference interval applies (e.g. fertility hormones) but where this is not able to be handled by the LIS.

"Canned" Comments

Some laboratories try to extend the use of standard comments to support interpretative reporting - *where reference to the patient context is required*. However, these rules are highly restricted in their ability to cross reference other tests, patient demographics and other information, and they have significant limitations when applied to interpretative reporting.

Seen from the perspective of the referring doctor, "canned" comments may be of little clinical value. Worse, they may potentially be misleading or dangerous.[5,6] The problem of spurious commenting is compounded even further when patients have sequential tests with similar findings, in which case the inappropriate comment will be issued repeatedly.

Problems with LIS rules applications

There are some fundamental issues that prevent more effective use of LIS-type rules. Collectively, these problems have been dubbed the "knowledge engineering bottleneck" and remain the subject of ongoing research.

In practical terms these problems can be summarised as follows:

> Bottleneck 1: Programming environment - demand for access to rules functionality has resulted in many LIS providers creating rule editing environments with more user-friendly, "windows"-type user interfaces.
>
> Bottleneck 2: Rule structure - as collections of rules grow, interactions between the rules rapidly becomes a serious issue. Increasing complexity causes conflicting behaviour that can be very difficult to resolve.
>
> Bottleneck 3: Understanding the expert - coding rules requires that an expert specify all the contexts in which that rule should and should not fire. This apparently simple requirement is in practice a notoriously difficult one. It is virtually impossible for an expert clinical biochemist to specify all of the criteria for appending Comment A, and the circumstances under which Comment A is supplanted by Comment B.
>
> Bottleneck 4: Keeping the rules current - laboratory medicine requires that the rules change to keep pace with growth in knowledge, new tests, altered reference intervals, and updated guidelines.

There is ample evidence of these bottlenecks in clinical laboratories, although the problem is rarely articulated in this way. Instead the bottlenecks are embedded implicitly in our choice of IT tools, statements that IT staff have other priorities to attend to, and our widespread acceptance of canned comments.

Expert Systems

Expert systems are computer programs that "represent and reason with knowledge of some specialist subject with a view to solving problems or giving advice".[7] They have potential in clinical laboratory reporting by supporting automation of the process, improving accuracy and consistency, and enhancing the quality of work.

As can be seen from the example below, patient-specific interpretation requires extensive cross-reference to other information contained in the LIS such as previous test results, other related tests and clinical notes, both current and previous. An expert system therefore needs to be adequately populated with a large number of rules that capture this knowledge.

Note again that this is not just a matter of *quantity* of rules – many rules makes for much conflict, as above. Therefore, an expert system that can find routine application in clinical laboratory practice *must have its own internal mechanisms for minimising or eliminating this conflict.*

"Embedded" Expert Systems

There are many examples of embedded expert systems in daily use, from motor vehicle diagnostics to domestic appliances such as washing machines. In these cases, the human operator has no interaction with the expert system.

At first glance, embedded expert systems suggest some promise for clinical laboratory reporting. Indeed, many such applications in clinical pathology have been described in research settings, such as interpreting renal function tests[8] and hepatitis serology.[9] Despite many years of work, however, very few of these have made their way into routine operation outside the laboratory environment in which they were developed.

There are several major obstacles:

- These systems are highly specialised.
- They lack mechanisms for interfacing to LIS systems, integration with laboratory workflow, or scaling up to deal with the volumes of tests done in large laboratories.
- They require the end user to accept the knowledge built into the systems by the developer's experts.
- Expert knowledge changes regularly and there needs to be a mechanism for updating the rules to keep them current.

Consequently, embedded knowledge-rich systems are not generally available for laboratory reporting applications.

Configurable Expert Systems ("Rules Engines")

These modules allow the customer to build and maintain their own rules. This flexibility overcomes many of the limitations of narrow and static embedded systems.

There are likely to be significant training and workload implications for IT, pathologists and laboratory staff when implementing configurable rules engines. In particular high volume laboratories, which rely on extensive use of auto-validation to manage turnaround times, need to work closely with the system provider to manage auto-validation and therefore the impact on laboratory staff and pathologists' time.

As well there needs to be an ongoing commitment to resource the project so that the system's behaviour remains current and appropriate.

LabWizard

At this time, LabWizard may be the only commercially available system for clinical laboratory report commenting. It supports the management of patient-specific laboratory reports and is designed around routine laboratory workflows.[10] LabWizard is available as a rules engine package (where local biochemists supply the content) or expert content may be supplied.

LabWizard supports ongoing, incremental deployment of rules. This flows from RDR's automated internal mechanisms for managing rule conflicts, thus overcoming most of the bottlenecks cited above.[11]

Bare ("Shell") Rules Engines

These systems are highly generic, in that they are not industry-specific. They allow input of data and creation of rules, but they do not address laboratory-specific demands such as result and report workflow, test result representation (including reference intervals), nor is there specification for interfacing to LIS systems or workflow integration.

Potentially, they might find some application in the clinical laboratory. However this would require a significant collaborative engineering project, including expert system engine provider (often this level of support is not available), LIS vendor and laboratory staff, to customise the generic tool to the specific requirements of real world laboratory reporting.

Key Steps to Implementing an Expert System for Interpretative Reporting

Before looking at technology solutions, a careful review of the laboratory's current reporting strategies is required. In particular, identify the key stakeholders (clinical, lab, IT, executive) and ascertain their needs and the contribution they will make to the project.

A formal project should be established. The project team will need to establish the project's objectives. What are the clinical and commercial drivers for the project? Where is the clinical area of need that is most likely to deliver results? What is the time frame to work to? What investment in technology and project management does the business case require?

All stakeholders need to understand that an expert system will remain a "live" project, requiring regular review and input to keep it current and effective.

In choosing a technology platform, some key issues to consider include:

- What level of interfacing to LIS is required?
- Will the product adapt to our current laboratory reporting workflow? Do we need to change how we do things? Does the product need to be customised to fit our lab?
- What impact will the product have on core service levels e.g. turnaround times?
- What level of user training is required?
- Does the product support ongoing evaluation against clinical, quality and commercial indicators?
- Does the product support true patient-specific interpretation?
- Does the product include internal mechanisms for managing rule context?
- Does the product support ongoing evolution to remain current?
- What resourcing is required for rule management tasks? Will this change as the system matures?

Worked Example

Consider a patient with the following results:

Sex:	M	
Age:	54	
Total cholesterol:	8.0	mmol/L
Total triglyceride:	0.9	mmol/L
HDL-C	1.3	mmol/L
LDL-C	6.3	mmol/L

LIS rules would append a *result-specific* comment to this report. For example:

> Increased LDL can be seen in hypothyroidism and nephrotic syndrome. Primary causes include polygenic hypercholesterolaemia, Familial Hypercholesterolaemia, Familial Defective ApoB-100 and Familial Combined Hyperlipoproteinaemia. Increased LDL is a risk factor for CVD. Assess overall risk and consider treatment. LDL-C <2.0 mmol/L is the treatment target for very high risk individuals (NHFA 2005).

An expert system would append a *patient-specific* comment. This should reflect *all that is known about, and relevant to, this patient from the laboratory's perspective*. An example would be:

> Interpretation: Hyperlipidaemia persists and has not improved despite statin therapy. Patient remains at HIGH RISK for CHD in view of previous CHD and raised LDL.
>
> Recommendation: Review medication regime if no response to statin. Suggest TSH to exclude subclinical hypothyroidism.

Implicit in this patient-specific comment is knowledge of:

- Current and previous lipid results
- Trend in these results (i.e. they are not improving)
- Clinical notes provided currently and previously (CHD history, statin therapy)
- Any tests done to exclude secondary conditions (hypothyroidism – no available TSH results)
- Other information indicating that the patient is already known to be hypothyroid or have kidney disease e.g. clinical notes, test results

The Table below contrasts result-specific (canned) comments with those generated by an expert system (patient-specific).

Summary

Implementing an expert system to improve laboratory reporting can be a very rewarding project. The right solution should deliver tremendous clinical value for patients and doctors, enhance the standing of the laboratory in the healthcare community and support greater job satisfaction for clinical and laboratory staff. However, success requires a significant investment, and key stakeholders need to understand the clinical, quality and commercial benefits and costs of implementing a system.

Care with product selection is critical. In some cases, standard comments may be sufficient. If so, most rule packages within LIS systems, or perhaps middleware products, will suffice. Be careful to understand the limitations of standard comments, and if you find yourself heading in the "canned comment" direction, give careful consideration to an expert system.

If considering an expert system, careful research is vital. Many promises are made by software providers, especially regarding the level of skill required to use the product. This needs careful evaluation, with special reference to issues of ongoing maintenance. Most importantly, be clear that the product you are considering can indeed provide the level of detailed, patient-specific interpretation that you require.

Comment attribute	Result-specific	Patient-specific	Example
Identifies key results	Yes	Yes	Raised LDL
Provides generic interpretation	Yes	If required	List of possible causes
Acknowledges previous results	No	Yes	"…persists.."
Acknowledges trend in results	No	Yes	"…has not improved.."
Considers clinical history	No	Yes	Statin, previous CHD
Considers other related tests	No	Yes	TSH not done previously
Provides specific interpretation	No	Yes	"Patient remains at high risk"
Provides generic advice, which may be redundant	Yes	No	"consider treatment"
Provides specific advice	No	Yes	"Suggest TSH" since this has not been done

Competing Interests: The author has a financial interest in Pacific Knowlege Systems.

References

1. Pacific Knowledge Systems. www.pks.com.au (Accessed 1 March 2008).
2. Compton P, Peters L, Edwards G, Lavers T. Experience with Ripple-Down Rules. Knowledge-Based Systems 2006;19:356-62.
3. Barlow IM. Are biochemistry interpretative comments helpful? Results of a general practitioner and nurse practitioner survey. Ann Clin Biochem 2008;45:88-90.
4. Laposata ME, Laposata M, Van Cott EM, Buchner DS, Kashalo MS, Dighe AS. Physician survey of a laboratory medicine interpretive service and evaluation of the influence of interpretations on laboratory test ordering. Arch Pathol Lab Med 2004;128:1424-7.
5. Lim EM, Sikaris KA, Gill J, Calleja J, Hickman PE, Beilby J, et al. Quality assessment of interpretative commenting in clinical chemistry. Clin Chem 2004;50:632-7.
6. Laposata M. Patient-specific narrative interpretations of complex clinical laboratory evaluations: who is competent to provide them? Clin Chem 2004;50: 471-2.
7. Jackson P. Introduction to Expert Systems. 2nd ed. St Louis, USA: Addison-Wesley; 1990.
8. Fugleberg S, Greulich A, Stenver D. Computer-assisted diagnosis of acute azotemia: diagnostic strategy and diagnostic criteria. Comput Bio Med 1991;21:399-406.
9. Adlassnig K, Horak W. Development and retrospective evaluation of Hepaxpert-I: a routinely-used expert system for interpretive analysis of hepatitis A and B serologic findings. Artif Intell Med 1995;7:1-24.
10. Edwards G, Compton P, Malor R, Srinivasan A, Lazarus L. PEIRS: a pathologist-maintained expert system for the interpretation of chemical pathology reports. Pathology 1993;25:27-34.
11. Edwards G, Kang BH, Preston P, Compton P. Prudent expert systems with credentials: managing the expertise of decision support systems. Int J Biomed Comput 1995;40:125-32.