Automated quality checks on repeat prescribing

Jeremy E Rogers, Christopher J Wroe, Angus Roberts, Angela Swallow, David Stables, Judith A Cantrill and Alan L Rector

SUMMARY

Background: *Good clinical practice in primary care includes periodic review of repeat prescriptions. Markers of prescriptions that may need review have been described, but manually checking all repeat prescriptions against the markers would be impractical.*

Aim: *To investigate the feasibility of computerising the application of repeat prescribing quality checks to electronic patient records in United Kingdom (UK) primary care.*

Design of study: *Software performance test against benchmark manual analysis of cross-sectional convenience sample of prescribing documentation.*

Setting: *Three general practices in Greater Manchester, in the north west of England, during a 4-month period in 2001.*

Method: *A machine-readable drug information resource, based on the British National Formulary (BNF) as the 'gold standard' for valid drug indications, was installed in three practices. Software raised alerts for each repeat prescribed item where the electronic patient record contained no valid indication for the medication. Alerts raised by the software in two practices were analysed manually. Clinical reaction to the software was assessed by semi-structured interviews in three practices.*

Results: *There was no valid indication in the electronic medical records for 14.8% of repeat prescribed items. Sixty-two per cent of all alerts generated were incorrect. Forty-three per cent of all incorrect alerts were as a result of errors in the drug information resource, 44% to locally idiosyncratic clinical coding, 8% to the use of the BNF without adaptation as a gold standard, and 5% to the inability of the system to infer diagnoses that, although unrecorded, would be 'obvious' to a clinician reading the record. The interviewed clinicians supported the goals of the software.*

Conclusion: *Using electronic records for secondary decision support purposes will benefit from (and may require) both more consistent electronic clinical data collection across multiple sites, and reconciling clinicians' willingness to infer unstated but 'obvious' diagnoses with the machine's inability to do the same.*

Keywords: *automated medical records systems; quality control; software design; clinical decision support systems; medication errors; repeat prescribing.*

J E Rogers, MRCGP, clinical research fellow; C J Wroe, MA, MB BChir, clinical research fellow; A Roberts, MSc, research associate; A L Rector, BA, MD, PhD, professor in medical informatics, Medical Informatics Group, Department of Computer Science; J Cantrill, MSc, MRPharmS, professor of medicines, School of Pharmacy and Pharmaceutical Sciences and National Primary Care Research and Development Centre (NPCRDC); A Swallow, research interviewer/technician, NPCRDC, University of Manchester. D Stables, medical director, Egton Medical Information Systems Ltd, Leeds.

Address for correspondence

Dr Jeremy Rogers, Medical Informatics Group, Department of Computer Science, Oxford Road, University of Manchester, Manchester M13 9PL. E-mail: jeremy@cs.man.ac.uk

Submitted: 2 April 2003; Editor's response: 3 June 2003; final acceptance: 4 August 2003.

Introduction

THE United Kingdom (UK) National Health Service Information Authority (NHSIA) strategy document *Information for Health,* and its subsequent update *Building the Information Core,*[1] states that general practice electronic patient records should become valuable information repositories, guiding clinical decisions and resource planning. Some researchers have measured the quality of the electronic data collected in UK primary care by comparing patient data in aggregate to 'gold standard' disease registers or expected disease incidence figures.[2-6] We report a project using single electronic records on individual patients as the input to software offering support for clinical decisions about them.

The focus of the study was medication review, which is an accepted part of good clinical practice and also now a requirement of the National Service Framework for older people.[7] Management and review of repeat prescriptions issued in UK general practice is currently acknowledged to be poor.[8,9] Detailed manual review of prescriptions for patients over 65 years of age by a pharmacist has recently been shown to be effective.[10] A set of validated indicators,[11] advocated in the National Service Framework as part of repeat prescribing review, has also been developed to identify individual repeat prescribing events for possible clinical reconsideration or review.

Widespread and routine manual application of such review processes would be expensive. Although there was an attempt to implement all the indicators in our project as software checks, this paper reports on only one: 'The indication for the drug is recorded in the electronic patient record and upheld in the *British National Formulary* (*BNF*)'.[11]

Method

Although the EMIS system (the existing core software of a major primary care software supplier) allows users to assign a diagnosis as the indication for a drug when it is prescribed, the empirical evidence is that this feature is often unused by clinicians. As a proxy for this missing association, we set out to compare the known indications for each prescribed drug with the list of diagnoses for the same patient recorded in the past 3 months or on their active problem list. If nothing in the diagnosis list matched any of the drug's possible indications, the reason for the prescription was deemed to be unrecorded.

The *BNF*[12] was chosen as a clinically oriented, 'gold standard' source documenting the licensed indications of prescribable medicines in UK general practice. The electronic version of that resource, the *eBNF*, was unsuitable for direct integration into the software. Primarily a human readable resource, it includes many ambiguous expressions. For

HOW THIS FITS IN

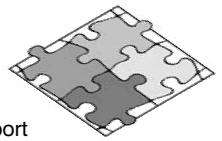

What do we know?

It is often proposed that successful implementation of clinical decision support is a requirement in order to persuade clinicians to make the best use of information technology. Centrally authored decision support rules need coded data that is captured consistently across multiple sites if they are to perform predictably in all sites.

What does this paper add?

Inconsistency in United Kingdom electronic primary care records may be a significant obstruction to successful implementation of decision support software. Apparently simple decision support tools may require surprisingly sophisticated diagnostic inference if they are to avoid being dismissed by clinicians.

example, if it is written that a drug is indicated for 'arrhythmia', how should a computer determine the full set of conditions that may constitute a reason to use the drug? Is the set the same for all drugs carrying this indication? Could there be some kinds of arrhythmia for which the drug is not indicated? Another example is that, if a drug is indicated in 'allergic and vasomotor rhinitis', is this a true logical 'AND', such that the patient must have both conditions simultaneously, or is it a logical 'OR'? Human readers possess contextual and domain knowledge that enables them to infer, with little conscious effort, the correct meaning of such semantically or logically ambiguous expressions. Computers, typically, do not.

Notwithstanding the difficulty a computer might have in interpreting any part of the *eBNF,* it may also struggle even to identify all the relevant information that needs interpreting, as the *eBNF* does not restrict all information about a drug; for example, its side effects, indications, and contraindications, to appearing only in the individual drug monograph. Relevant information may also be found only within section or chapter narratives, or in the appendices. Where a group of similar drugs have the same indications or side effects, individual drug monographs; for example, for atenolol, may refer the reader to the monograph of the index drug; for example, propranol for beta-blockers. Human readers can follow these textual indirections. Computers require the indirections to be formally represented outside the text.

A team of physicians working in medical informatics therefore used knowledge, engineering tools and methodologies, developed within the OpenGALEN programme,[13] to abstract the relevant clinical content from the *eBNF* as an unambiguous computable resource using a prototype drug ontology[14-16] and controlled vocabulary, mapped to an open source reference concept model for the medical domain.[13] Two thousand, four hundred and seventy-three commonly prescribed 'virtual products' were identified, each representing a specific combination of one or more of 1039 pharmacologically active ingredients, given in a specific formulation and by a specific route.

The clinical systems in the study stored prescribing data using a drug dictionary that lists more than 18 000 'actual products' distinguished by dose and manufacturer as well as by ingredient, form, and route. Six thousand, eight hundred and two actual product codes, comprising the majority of repeat prescribed products used in UK primary care, were mapped semi-automatically to the virtual product codes. Codes for different actual products that shared the same active ingredients, form, and route, but different dosages or manufacturers, were mapped to the same virtual product code.

Information, including the indications for use of each virtual product, was expressed within the new computable resource (Box 1). The indications for the top 80% of prescribed virtual products by volume were collated, comprising 529 different conditions. A mapping was declared from each indicated condition to a set of 5-byte Read codes (Box 2).

The combined mappings database supported the following functionality: for each 'actual product' that a patient was prescribed, a 'virtual product' could be identified and, through that, a list of indications. Each indication mapped to a list of Read codes. The presence of any member of that Read code set (or any of their descendent codes), when encountered in the electronic patient record, was accepted as evidence that the patient had at least one documented possible reason for taking the original 'actual product'. If no such Read code was encountered, then a 'no indication' alert was triggered.

The completed knowledge base and execution software, known as Tool 1, was integrated with EMIS and installed in three practices for a 4-month period in 2001. The software was configured to run daily, examining all repeat prescriptions as they were printed and adding the item to a job list for later examination by the clinician if an alert was raised.

The original project plan had been for the pilot clinicians to record whether each alert was correct or not, and for the new drug information resource to be improved iteratively as required. Early user feedback, however, suggested that 'no indication' alerts were being incorrectly triggered much more frequently than had been expected, and that efforts to improve this were not working. The planned systematic end user evaluation of the validity of alerts, and later project stages investigating any clinical management changes prompted by the alerts, were therefore abandoned.

To investigate the extent and cause of the incorrect alerts further, the software was run prospectively in two of the practices, covering 10 433 patients between them. The records of 3707 fully anonymised patients who were scheduled for a repeat prescription of at least one item within 6 months, beginning in July 2001, were processed as a batch. In total, 8794 repeat prescription items were identified for further study (Figure 1).

Of the 8794 items examined prospectively by Tool 1, a 'no indication' alarm was raised by 3428 (38.9%) because the computer could not find a coded patient record entry to match the known drug indications. An exhaustive manual study of all 3428 alerts and their associated anonymised disease entry records was undertaken. Four different types of error leading to incorrect alerts were identified (Box 3), and each alert was assigned to one of the four error cate-

***BNF* text (abstract)**	**After rewrite**
4.3.1 Tricyclic and related antidepressant drugs These drugs are most effective for treating moderate to severe *endogenous depression* associated with psychomotor and physiological changes	MAIN drug CONTAINS tricyclic group HAS_DRUG_FEATURE indication FOR treating ACTS_ON depressive illness
AMITRIPTYLINE HYDROCHLORIDE ***Indications***: depressive illness, particularly where sedation is required; nocturnal enuresis in children (see section 7.4.2)	MAIN amitriptyline hydrochloride PROPERTIES HAS_DRUG_FEATURE indication FOR treating HAS_PATIENT child ACTS_ON nocturnal enuresis

Box 1. Sample of rewritten BNF entry showing extract of original BNF text and the corresponding computable representation after knowledge engineering.

Drug	**Indication**	**Read code mapping set**	**Text of Read code**
Digoxin	Atrial fibrillation	14AN.	H/O: atrial fibrillation
Amiodarone		3272.	ECG: atrial fibrillation
Esmolol		3273.	ECG: atrial flutter
		7936A	IV pacer control atrial fibrillation
		G573.	Atrial fibrillation/flutter

H/O = history of; ECG = electrocardiograph; IV = intravenous.

Box 2. Example of indication mapping showing Read code mappings for 'atrial fibrillation'.

- The electronic patient record contains a code referring to an indicated condition, but the code used is idiosyncratic and would not normally be mapped to that indication
- The electronic patient record contains a code that is very suggestive of the (unrecorded) existence of a condition that is an indication, although it is not an accepted *BNF* indication
- The electronic patient record contains a widely used code referring to an indicated condition, but the mapping between that indication and the code was not declared to the algorithm
- The electronic patient record contains a code that is almost certainly the indication, but the *BNF* does not include the condition as a licensed indication

Box 3. Four causes of error — categories used to analyse cause for alarm.

gories or to a fifth category of 'correct alert' when the record contained no indication for an item.

Two other researchers independently analysed a randomly selected subset of the alarms (n = 326; n = 299). Overall agreement for each researcher with the original categorisation was 76% (κ = 0.67) and 77% (κ = 0.69), although specific category agreement (a measure of agreement on the application of any one category) was higher (range = 84–96%, κ range = 0.81–0.96). These κ-values indicate that moderate inter-rater variability remained. The differences between the three human raters were examined, revealing disagreement centred mainly on how to categorise certain commonly occurring patterns and inconsistent application of agreed rules both within and between raters.

A new algorithm (Tool 2) was devised to achieve a more consistent categorisation. All 'no indication' items were automatically categorised into one of 43 therapeutic groups (Box 4), according to the lexical properties of the chemical or brand name of the drug. For example, ACE-I inhibitors were identified as a group by searching for all prescription items whose text included the substring elements '*pril*', '*arace*' or '*novace*'.

The set of Read codes that, if encountered in the record, had prompted each of the four categorisations in the manual analysis, was collated for each therapeutic group. Database scripts and queries were written to recategorise all alerts according to these formalised rules.

All 3428 positive alerts reported by Tool 1 were categorised automatically using Tool 2. Overall agreement of the result with the original manual categorisation was 85% (κ = 0.8) and specific category agreement varied between 92% and 96% (κ = 0.91–0.95). Overall agreement with the other two raters was 74% (κ = 0.63) and 79% (κ = 0.70).

The qualitative reaction to the software by 10 of the clinicians was later assessed by audiotaped semi-structured interviews. The interviews were transcribed verbatim and analysed to identify common themes.

A research ethics committee was approached before the project started but advised that ethical approval was not necessary.

Further technical material relating to this study is available from http://www.cs.man.ac.uk/mig/.

Results

The results of the automatic categorisation of the 3428 'no indication' alerts by Tool 2 were as follows:

Correct alert: indication not stated in the record

Out of 3428 items examined, 1301 (38%) of the 'no indication' alerts studied were correct. This indicates that at least 14.8% of all 8794 repeat prescribed items examined by Tool 1 did not include a valid coded indication in the associated electronic record (Box 5). It should be noted that the study did not measure the number of items where the record con-

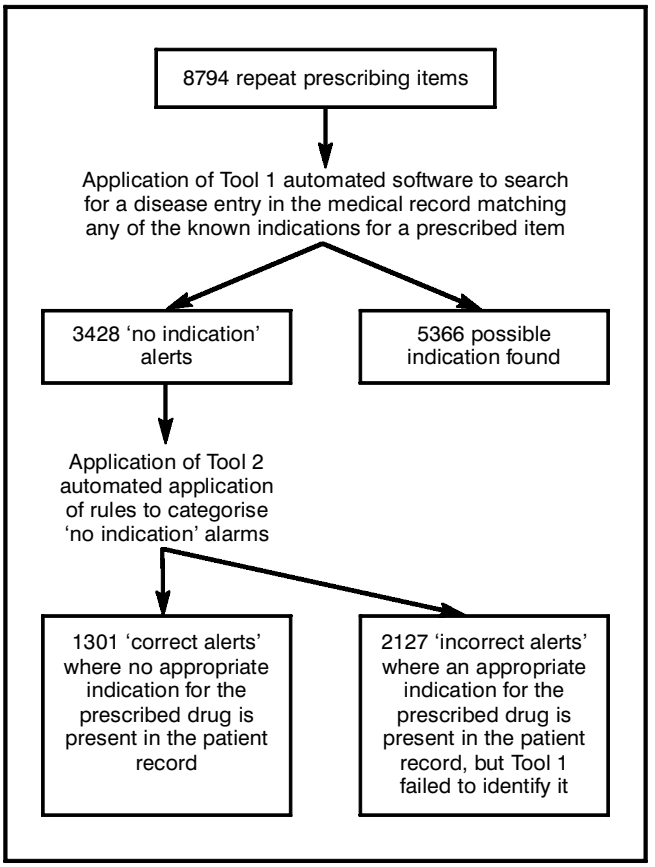

Figure 1. Flow chart of study.

tained no indication but where the algorithm incorrectly reported that it did: the true incidence of 'no recorded indication' prescribing in the study practices may therefore be higher.

A breakdown of the frequency of correct 'no indication' alerts by therapeutic category is shown in Table 1. Significant differences were observed across therapeutic groups: generated alerts were usually correct when they concerned prescriptions for eye or ear drops, antihistamines, contraceptives, hormone replacement preparations, H_2 antagonists, antibiotics, psychotropic agents, and anticoagulants. Conversely, alerts were usually incorrect when they concerned drugs for cardiovascular disease or asthma. However, because the greater volume of all prescribing relates to those therapeutic categories where the alert generation system performed most poorly (drugs for cardiovascular disease and asthma), the overall performance of Tool 1 was skewed further towards lower accuracy.

Incorrect alert: missing indication-to-EPR (electronic patient record) mappings in the knowledge base

Nine hundred and twenty (27%) of all alerts raised were incorrect because the system failed to identify records that contained an entry corresponding to a valid indication for the drug, and where that entry was made using a commonly used code for the condition. These incorrect alerts arose because the indication-to-Read code mapping (Box 2) contained omissions.

Incorrect alert: idiosyncratic local coding

Nine hundred and thirty-nine (27%) of all alerts raised were incorrect because, although the patient record contained an entry coding for an indicated condition, the code used was idiosyncratic and normally had a different interpretation. The code was therefore not mapped to the relevant indicated condition; for example, when considering a drug licensed for use in asthma, the software at all sites searched the record only for any occurrence of the 5-byte Read code 'H33.. Asthma' or any of its more specific descendent codes such as 'H333. Acute exacerbation of asthma'. One of the two study practices, however, routinely recorded asthma only using the Read code '14B1. History of asthma'. Other

ACE I inhibitor	Antiplatelet	Constipation	Nicorandil
ACE II inhibitor	Antipsychotic	Cystic fibrosis	Nasal disease
Alcoholism	Aspirin	Diabetes	Non-steroidal anti-inflammatory
Alpha-blocker	Asthma and chronic obstructive pulmonary disease	Digoxin	Oral contraceptives
Anaemia	Benzodiazepine	Diuretic	Osteoporosis
Antibiotic	Beta-blockers	Ear conditions	Peptic ulceration
Anticoagulant	Breast cancer	Ocular disease	Prostate cancer
Antidepressant	Calcium channel blockers	Gout	Skin preparations
Anti-emetic	Claudication	Hormone replacement therapy	Statins
Antihistamine	Colitis	Isosorbide	Transplant medicine
Antimigraine		Menorrhagia	Verapamil

Box 4. Therapeutic groupings used by Tool 2.

Drug	Valid indications	Actual contents of electronic patient record
Atenolol	Hypertension	Repeat prescription monitoring
	Angina	Adult health exam
	Arryhthmia	Diagnostic gastroscopy NEC
		Injection of steroid for local act NEC
		Cervical spondylosis — no myelopathy
		Polymyalgia rheumatica

NEC = not elsewhere classified.

Box 5. Example of 'no indication' record.

Table 1. Subanalysis of rates of correct alert by class of drug.

Therapeutic category	All alerts	Correct	Percentage (%)
Cardiovascular system and asthma			
Anticoagulants and antiplatelet	82	48	59
Other treatment for ischaemic heart disease	251	127	51
Calcium channel blockers	185	68	37
Diuretics	373	132	35
Inhalers and anti-asthmatics	485	159	33
ACE inhibitors	285	74	26
Beta-blockers	279	60	22
Aspirin 75 mg for prophylaxis of myocardial infarction	332	69	21
Statins and cholesterol lowering	290	17	6
Subtotal	2562	754	29
All other medicines			
Eye and ear preparations	22	20	91
Antihistamines	154	129	84
Contraceptives	18	15	83
Hormone replacement therapies	87	65	75
H_2-antagonists	94	71	76
Antibiotics	51	36	71
Antidepressants and psychotropics	203	123	61
Skin preparations	18	9	50
Treatments for osteoporosis	37	16	43
Gout treatments	47	16	34
NSAIDs	60	12	20
Other	75	35	53
Subtotal	866	547	63
Total	3428	1301	38

NSAIDs = non-steroidal anti-inflammatory drugs.

major diagnoses were similarly recorded using 'history of …' codes. The algorithm therefore did not recognise these patients as having asthma, hypertension, or other major conditions.

Incorrect alert: omissions in the'gold standard' drug information corpus

One hundred and sixty-three (5%) of the alerts raised were incorrect because the *BNF*, as the reference corpus used in this study, did not always list all current clinically accepted indications for a drug. For example, the condition 'benzodiazepine dependence' was present in the record of a number of patients receiving such drugs, but it is not recognised in the *BNF* as an accepted indication.

Incorrect alert: indication only inferable by a clinician

One hundred and five (3%) of the incorrect alerts arose because an 'obvious' diagnosis had not been entered explicitly in the record. For example, a clinician might reasonably infer an unstated diagnosis of breast cancer in a patient known only to have had a 'breast lump symptom' and a 'total mastectomy', and so would justify a repeat prescription of tamoxifen. Similarly, hypertension is not itself an indication for aspirin, but it is a risk factor for ischaemic heart disease and myocardial infarction, both of which are listed as indications for antiplatelet therapy using aspirin. While clinicians can make such inferences, the software was not designed to perform such sophisticated operations.

Results of interviews

The semi-structured interviews revealed that most clinicians remained enthusiastic regarding the software's aims. Several interviewees described proper documentation of the reason for a prescription as an important goal in itself:

'I think it is very important, if any medication is prescribed I think it has to be documented because if it is not documented then nobody else knows.' (GP3.)

'The indication for taking the tablets, essential … if we prescribe we must always know for what we are prescribing for.' (GP7.)

'If the patient is not one of yours and they are on repeats and you cannot understand why, it is vital that there is a reason.' (GP10.)

'Very important. I have just remembered why we signed up to do the project in the first place, as I was hoping the system would help us with that.' (GP5.)

Discussion

Summary of main findings

Sixty-two per cent of alerts raised automatically were found to be incorrect. The users stated a desire to have the functionality of the software, but rejected it until the unacceptably high number of incorrect alerts could be reduced.

Scope to reduce the incorrect alert rate is, however, limited without simultaneous change in clinical coding prac-

tices. Even taking into account all possible drug knowledge base and mapping improvements, this analysis suggests that a residual minimum rate of around 44% incorrect alerts will persist: this is the number of alerts still attributable to idiosyncratic records or those requiring inference (= 939 + 105), as a percentage of all alerts that would still be raised after all formulary omissions and mapping errors were corrected (= 3428 - 920 - 163).

Such a system would most likely still be sufficiently irritating to clinicians that acceptance of the system would be limited and generated alerts routinely ignored. In the event that idiosyncratic coding was also abolished, but records requiring inference remained, the incorrect alert rate would be 7.5%.

Uniform algorithms for machine interpretation of the record, intended for application across multiple sites, require standardised input, not only in terms of which coding scheme they use, but in how those schemes are used locally. Idiosyncratic coding conventions within and between different practices, or between different primary care organisations, cut across this standardisation and will be a significant continuing contributor to incorrect alerts in this kind of decision support. Unless coding can be made more consistent between sites, supposedly general algorithms to interpret individual electronic patient records may require considerable local 'tuning'.

The positive response by the study clinicians to the idea of the software suggests a desire to improve the quality of their prescribing documentation. Notwithstanding existing data quality initiatives in the study practices, at least 14.8% of all their repeat prescriptions were found on manual analysis to have no documented indication and this supports the case for better tools that would make it possible for clinicians to achieve this improvement.

Strengths and limitations of this study

Other researchers have reported successful statistical analysis over populations of electronic UK general practice records, typically measuring the local prevalence of certain common diseases. Our experiment was a preliminary test of the feasibility of applying a uniform interpretive algorithm to individual electronic patient records, and across the full range of conditions. As a small study limited to three practices with detailed analysis in only two, the results must, however, be taken with caution.

Recording the indication for every prescription may initially appear of limited clinical value. However, implementing the other more clinically valuable checks, such as whether the dose is too high or too low, or whether a medication has been prescribed for too long, often requires knowledge of the indication; dosage and regime parameters are often indication dependent. This exposes an unstated tension between what is required of the medical record in fulfilling its primary purpose of supporting direct clinical care by a human, and meeting the needs of secondary analysis, possibly by external agencies, including computers.

We acknowledge that the second automated categorisation of the positive alerts performed significantly better than the very much more complex algorithm whose original output it was applied to. This result is not unexpected: heuristic approaches, covering post hoc only the cases encountered in a specific dataset, generally perform very well. However, they are typically much harder to maintain or scale, and perform less well and unpredictably on other datasets. The complexity of the approach taken for the original algorithm (Tool 1) was specifically devised to address issues of scale, maintenance and generalisability.

Implications for future research

A detailed characterisation of inter-practice and intra-practice recording variations was not possible within this study. Some researchers have measured coding agreement in the laboratory across different medical coding schemes, reporting that agreement is higher using Clinical Terms Version 3 compared to 5-byte Read.[17] Results from PRIMIS[18] suggest that coding inconsistency is widespread in real practices, but a more detailed investigation may be particularly relevant now that the new UK primary care contract seeks to introduce a data-driven quality framework.

Thiru *et al* commented that no standard measure of electronic patient records data quality exists.[19] Further work is needed, in particular to determine what level of inter-rater coding agreement is required before automated analysis of the individual record, such as described here, becomes viable.

A significant implementation cost was incurred in mapping from the EMIS drug dictionary to the virtual product code set, even though the dictionary substructure allowed for a semi-automatic mapping process. No special problems were encountered using that dictionary; a similar semi-automatic mapping was authored to a different dictionary for a different project. However, the cost of maintaining mappings to multiple-product dictionaries for any system intended to work across multiple sites is significant. The NHSIA goal of a single UK clinical product reference source (UKCPRS) may address this issue, particularly if it includes sufficient structure and content to facilitate fully automatic mapping.

This study suggests that before apparently simple prescribing quality assurance and documentation checks can be routinely automated, two criteria must be met. First, the inconsistency of recording in electronic medical records must be improved. PRIMIS[18] is an important NHS data quality project providing software toolsets to detect omissions in what is recorded post hoc, but a different and complementary approach and toolset will be required pre hoc to change how data is recorded, if it is recorded at all.

Second, both doctors in practice and authors writing professional information; for example, drug formulary information, must make allowances for limitations in machine reasoning. The natural desire for speed and conciseness of expression must be reconciled with the machines' need for completeness and explicitness, at least while they cannot mimic human heuristic or expert reasoning. There is an obvious irony that, in the absence of pedantic record keeping, successful implementation of apparently simple reminders may require relatively sophisticated diagnostic inference to avoid generating warnings that doctors consider 'obviously' inappropriate.

References

1. NHS Information Authority. *Information for health*. Section 2.8. NHSIA, 2001. http://www.nhsia.nhs.uk/def/pages/info4health/contents.asp (accessed 16 September 2003).
2. Hassey A, Gerrett D, Wilson A. A survey of validity and utility of electronic patient records in a general practice. *BMJ* 2001; **322:** 1401-1405.
3. Hogan WR, Wagner MM. Accuracy of data in computer-based patient records. *J Am Med Inform Assoc* 1997; **4:** 342-355.
4. Brown PJB, Sönksen P. Evaluation of the quality of information retrieval of clinical findings from a computerised patient database using a semantic terminological model. *J Am Med Inform Assoc* 2000; **7:** 392-403.
5. Gray J, Majeed A, Kerry S, Rowlands G. Identifying patients with ischaemic heart disease in general practice: cross-sectional study of paper and computerised medical records. *BMJ* 2000; **321:** 548-550.
6. Gray J, Orr D, Majeed A. Use of Read codes in diabetes management in a south London primary care group: implications for establishing disease registers. *BMJ* 2003; **326:** 1130.
7. Department of Health. *Medicines and older people: implementing medicines-related aspects of the NSF for older people.* London: Department of Health, 2001. www.doh.gov.uk/nsf/medicinesop/interventions.htm (accessed 17 September 2003).
8. Audit Commission. *A prescription for improvement. Towards more rational prescribing in general practice*. London: HMSO, 1994.
9. NHS R&D Directorate and Sowerby Centre for Health Informatics at Newcastle. *NHS repeat prescribing project.* 2000. www.schin.ncl.ac.uk/RepeatPrescribing (accessed 17 September 2003).
10. Zermansky AG, Petty DR, Raynor DK, *et al.* Randomised controlled trial of clinical medication review by a pharmacist of elderly patients receiving repeat prescriptions in general practice. *BMJ* 2001; **323:** 1340-1343.
11. Cantrill JA, Sibbald B, Buetow S. Indicators of the appropriateness of long-term prescribing in general practice in the United Kingdom: consensus development, face and content validity, feasibility and reliability. *Qual Health Care* 1998; **7:** 130-135.
12. Joint Formulary Committee. *British National Formulary.* London: British Medical Association and Royal Pharmaceutical Society of Great Britain, 2003. www.bnf.org (accessed 17 September 2003).
13. Rogers JE, Roberts A, Solomon WD, *et al.* GALEN ten years on: tasks and supporting tools. In: Haux R, Rogers R, Patel V (eds). *Studies in health technology and informatics*. Vol. 84*. Proceedings of the 10th World Congress on medical informatics*. Amsterdam: IOS Press, 2001: 256-260. www.opengalen.org/ (accessed 17 September 2003).
14. Solomon WD, Wroe CJ, Rector AL, *et al.* A reference terminology for drugs. *Proc AMIA Symp* 1999; 152-156.
15. Wroe CJ, Solomon WD, Rector AL, Rogers JE. *Inheritance of drug information*. 2000 Annual fall symposium of American Medical Informatics Association, Los Angeles. Philadelphia: Hanley & Belfus Inc., 2000: 1158.
16. Wroe CJ, Solomon WD, Rector AL, Rogers JE. DOPAMINE — A tool for visualising clinical properties of generic drugs. In: Lavrac N, Miksch S, Kavsek B (eds). *Proceedings of the 5th International Workshop on Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP) a workshop at the 14th European Conference on Artificial Intelligence*. Berlin: ECAI, 2000: 61-65. http://ai.ijs.si/Branax/idamap-2000_AcceptedPapers/All.pdf (Accessed 8 Oct 2003)
17. Brown PJB, Warmington V, Laurence M, Prevost AT. Randomised crossover trial comparing the performance of Clinical Terms Version 3 and Read Codes 5 byte set coding schemes in general practice. *BMJ* 2003; **326:** 1127-1120.
18. NHS Information Authority. Primary Care Information Services (PRIMIS). www.primis.nhs.uk (accessed 18 September 2003).
19. Thiru K, Hassey A, Sullivan F. Systematic review of scope and quality of electronic patient record data in primary care. *BMJ* 2003; **326:** 1070.

Acknowledgements

This work was funded by a grant from the UK Department of Health.