Publets: Clinical Judgement On The Web?

John Fox, PhD, Jon Bury, MB ChB, Michael Humber, MSc,
Ali Rahmanzadeh, MSc, Richard Thomson, MA
The *OpenClinical* Group, Imperial Cancer Research Fund, London, UK

ABSTRACT

The Internet is now a major channel for publishing medical research data and documents, including clinical practice guidelines. It is now possible to capture guidelines in a computer interpretable form opening up the capability of using the internet (and intra/extranets etc.) to deliver patient-specific advice and other services. A development lifecycle and technology for publishing and delivering services at the point of care ("publets") are described. As with all new technologies, however, these new methods entail risks as well as opportunities. The paper closes with a discussion of quality requirements and an argument that publets should include a safety case as an integral part of their content.

INTRODUCTION

The internet is now a major channel for publishing documents and data in many fields. In medicine, for example, it is now used for dissemination of research results (e.g. the Cochrane Library of systematic reviews[1]) and research trials (e.g. the PDQ cancer trials database[2]). More recently it has become a means of disseminating practice guidelines by many organisations, including the Scottish Intercollegiate Guidelines Network[3], the UK's National Institute for Clinical Excellence[4] and the US National Guideline Clearinghouse[5]. This trend is also reflected in initiatives by established medical publishers, such as BMJ Publishing's *Clinical Evidence*.[6]

The creation of national and other large repositories of guidelines has many benefits, including the provision of an accepted authoritative source of information for practitioners, the opportunity to set quality standards for content and presentation etc.

Despite these benefits the limitations of purely document-based dissemination of knowledge about best clinical practice are increasingly recognised:

- Busy clinicians have little time to read guidelines and converting them from paper to electronic pages may not substantially change this reality.
- Even if a clinician has time to read the content of a guideline it may not be reliably memorised and correctly applied in practice.
- Conventional guideline documents do not provide recommendations that are tailored to the needs of individual patients.

There is therefore growing interest in distributing guidelines in an "enactable" form based on standard formats that are machine interpretable such as the Arden Syntax and more recently ASBRU, EON, GLIF, Prestige, Prodigy and PRO*forma* (see www.openclinical.org for summaries of these and other technologies). Technologies based on such formats can provide many services that have been proposed by medical informaticians, including patient-specific prompts and reminders, decision support, monitoring of risks and adverse events, scheduling of clinical tasks and care planning.

PUBLETS

A further refinement of enactable guidelines would be to publish them as web-accessible services that we call "publets" [7]. A publet combines machine-executable knowledge with more conventional media in an encapsulated form for publication on a network. A publet is comparable to a conventional publication in that the preparation may involve familiar disciplines like peer review, but it is distinctive in that the knowledge content is wholly or partially formalised for interpretation by a computer, not just to be read by people. This will

[1] http://hiru.mcmaster.ca/cochrane/cochrane/cdsr.htm
[2] http://cancernet.nci.nih.gov/trialsrch.shtml
[3] http://www.sign.ac.uk/
[4] http://www.nice.org.uk/nice-web/Cat.asp?c=29
[5] http://www.guideline.gov
[6] http://www.evidence.org/

[7] Publet is a contraction of publication and applet.

1067—5027/01/$5.00

permit the computer to apply medical knowledge in supporting clinical practice at the point of need, via the internet or a secure intranet or extranet.

An example of a set of publets is ERA (Early Referrals Application), that is currently being evaluated in a British National Health Service (NHS) trial[8]. ERA has been designed to help primary care physicians comply with practice guidelines published by the NHS[9]. The guidelines set out criteria for deciding whether to refer suspected cancer patients for urgent specialist services. ERA includes textual material from the paper guideline (translated into HTML web pages) augmented with an enactable procedure written in PRO*forma*. This procedure captures the patient history, interprets it, and advises whether or not referral is appropriate under NHS policy. A typical ERA report screen is shown in figure 1 (*full size versions of the figures in this paper can be found at www.acl.icnet.uk/amia*).

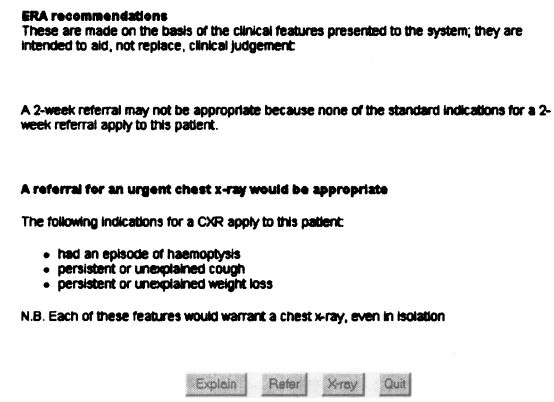

Figure 1: Example decision report for lung cancer referral by ERA (demonstration version can be accessed at www.infermed.com/wap/era)

GUIDELINE PUBLISHING METHODS

The National Guideline Clearinghouse™ "invites organizations, societies, and other developers of clinical practice guidelines to submit completed guidelines and related background information [with] two paper copies of each guideline ...together with electronic copies on disk (if available) for each guideline submitted". While the NGC *publishes* on the web its submission and quality review processes are conventional.

The BioMed Central[10] site also "publishes peer reviewed research across all areas of biology and medicine with immediate, barrier-free access for all". While the peer review process is open and on the web the content is otherwise standard.

Recognising the greater opportunities and needs for future web based publishing, Ida Sim has pointed out that the natural next step in clinical-trials reporting will be to describe trials not just in text but also using knowledge engineering and database technologies to facilitate reporting and retrieval of trial data [1]. For trial banks to be truly integrated they must share a common conceptual model of clinical trials supported by an appropriately adapted publishing and review process.

With new and powerful technologies becoming available for creating and enacting decision support and guideline systems we may expect medical service providers, publishers and other organisations to increasingly adopt such methods. If this is to become routine, however, we shall require appropriate tools and processes for authoring, reviewing and disseminating machine interpretable knowledge.

PUBLET DEVELOPMENT CYCLE

We propose an extended publishing cycle for publets, as shown in figure 2.

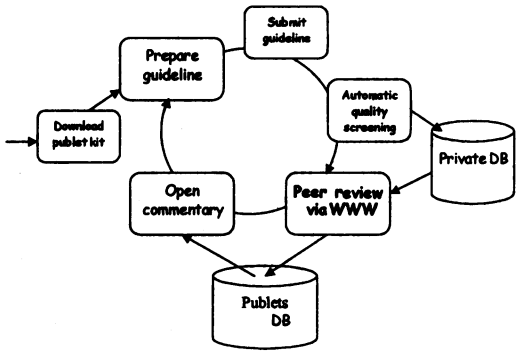

Figure 2: A publet life-cycle, in which part of the publishing process is conventional (based on peer review) and part automated by software.

The publishing cycle is a natural development of a conventional authoring and editorial process.

1. The author prepares the knowledge content that is to be embedded in the publet (guideline, protocol, care pathway, clinical algorithm or other knowledge) using appropriate tools. These may be downloaded from publisher's sites or

[8] www.infermed.com/wap/era
[9] http://www.doh.gov.uk/pub/docs/doh/guidelines.pdf
[10] www.biomedcentral.com

used directly over the web; the latter would be preferable where publet development is a group project.

2. As part of the preparation process, the knowledge content of the publet is automatically checked for syntactic correctness, and for logical adequacy by running the application against patient data.
3. Authors next submit the publet to a publisher's web site (e.g. by email) together with key words reflecting the content. Authors include the patient data against which the application has been validated.
4. A software agent receives the publet and screens it according to technical and quality standards (e.g. syntactic rules, use of standard clinical terms and, if a standardised conceptual model is available, semantic checking).
5. Submissions that fail screening are returned to the author with an automatically generated report. Acceptable submissions are posted on the publisher's site, and an acknowledgement, URL and password are returned to the author.
6. The next phase of publication is an adapted form of peer review. The publishing agent maintains a database of reviewers and their areas of clinical expertise and interest. It creates a short-list of reviewers by matching to the keywords provided by the authors. The responsible (human) editor makes a selection and sends the publet's URL to the preferred reviewers.
7. Reviewers critically evaluate the submission, running the application (over the web) against the authors' cases, and their own test cases if required. An anonymous report is prepared in the usual way and sent to the editor and authors (along with any additional test data that have been used).
8. If the publisher follows the standard peer review process for conventional publications, a decision to accept, reject or request revisions may be made at this point. However the web also offers further possibilities, e.g. an "open commentary" step in which publets are made publically available for a limited period together with the anonymous reports, permitting open discussion between authors, referees and the wider community.
9. Accepted publets appear with appropriate ancillary documentation (see discussion below) on the publisher's web site.

TOOLS FOR CREATING PUBLETS

We have described elsewhere a set of software tools to support the creation of decision support systems, interactive guidelines, smart care pathways etc. based on the PRO*forma* language, a formalism for representing clinical processes. PRO*forma* processes are modelled in terms of four basic classes of "task" (decisions and plans, clinical acts and data requests) together with control information describing their enactment over time and under uncertainty. The PRO*forma* language and technology have been successfully applied in a wide range of medical applications ([9]).

Recently we have developed an extended version of this technology to support the publet concept. This provides an intuitive set of software engineering tools for designing clinical processes in terms of the four classes of tasks supported by the PRO*forma* language. The environment provides visual programming and CASE tools for sketching the clinical process and scheduling component tasks, populating task details, defining decision criteria and testing correct enactment. The system also includes a mechanism for submitting the resulting publet to a selected website. The publet authoring environment is shown in figure 3.

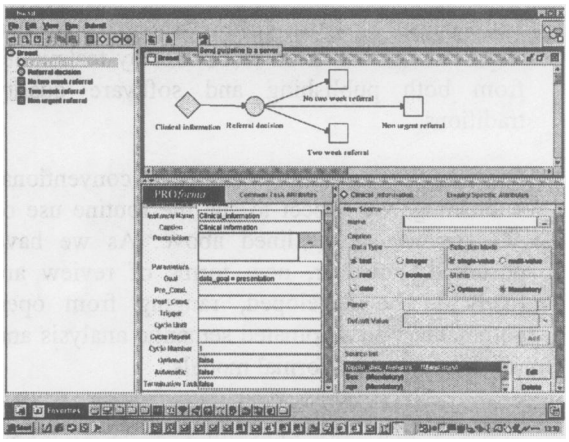

Figure 3: Publet authoring environment. This development and testing environment provides tools for developing a PRO*forma* specification of the tasks required in clinical processes. The top window shows the main "plan" that is used to represent a referral process for suspected breast cancer (see description of ERA guidelines in the text). The guideline consists of five tasks: first a request for data (diamond), followed by the referral decision (circle) based on the data obtained, and the three alternative actions that could be selected by the decision (squares). The left panel shows a thumbnail view of the task structure (useful for more complex applications) while the panels at bottom right are CASE tools for populating the definition of any selected task. A button near the centre of the tool bar near the top of

the display submits the application to the publisher's web site.

At the publishing site a software agent receives, processes and acknowledges submissions. The technology is currently oriented towards PRO*forma* applications, but we believe that the publet approach could also be used with other formats, such as Arden, Asbru or GLIF.

The final component of the technology is *Solo,* a task enactment engine that runs PRO*forma* applications over the web, and permits access to publet services using a standard web browser. Following installation, a publet is immediately available for access and *Solo* automatically creates generic HTML pages as user interfaces for each task as it executes [2]. These generic pages can be tailored to provide a specific "look and feel" for the application and/or to incorporate the imprimatur of the publisher.

PUBLET QUALITY MANAGEMENT

In order to maintain quality standards for this new kind of publishing, we must develop an appropriate methodology. Some of the quality issues for guideline systems generally are reviewed elsewhere [3] but publets raise additional issues. Since they bring together conventional documents and executable knowledge systems, we can expect that requirements will need to be met by techniques from both publishing and software quality traditions.

Drawing on techniques used in conventional publishing, we expect to see the routine use of peer review, as outlined above. As we have observed, however, new forms of review are likely to be developed, ranging from open commentary to automated semantic analysis and validation against formal models.

Publets must obviously be *documented* to a high standard. The designers of the Arden Syntax [4] set out a comprehensive documentation scheme for decision support modules. Fourteen general requirements are identified including documentation of *applicability* (e.g. when was the module produced? Which version is this? To what extent has it been validated?) and *library* information, including citations that may support or question the assumptions of the application and its clinical role (e.g. what is its purpose? In what context is it appropriate?). Publets should include at least this level of documentation, and all of it should be accessible at the point of use.

The enactment of codified medical knowledge to give patient-specific advice is a type of clinical intervention and so is subject to the same *provenance* requirements as any other clinical procedure. Ideally, publets are based on reliable and reproducible evidence of their efficacy and safety. If the content cannot be evidence-based, it should be consensus based.

As a web-based publishing method, publets offer the obvious further possibility of including links to research and other relevant background documents (such as the evidence base published in the Cochrane Library and TrialBank). In addition, however, as with expert systems and other kinds of decision support software, publets should provide *patient-specific justifications* for recommended actions and decisions. Peer review and screening processes should reflect these requirements

Since a publet is a piece of software as well as a medical publication, the authoring process should support good engineering practice as well as good science and good scholarship. For example, as the range of available publets develops, there will be an increasing need for a *component-based* approach to design and documentation, both to support reusability and reduce development costs. A publet that offers advice on the management of heart failure, for example, may include a component publet dealing with the use of ACE inhibitors published by an independent source. Quality requirements will therefore include making provision for reuse in the component as well as ensuring access to the component's own documentation from within the main application.

Reusability was one of the original objectives of the Arden Syntax though success has been somewhat limited by procedural and other features of the model [5]. More recent languages like Asbru, EON, GLIF and PRO*forma* are more declarative, though experience is needed to determine how much reusability is possible in practice.

More critically, as with many drugs and medical procedures, publets may have safety implications if they are used improperly or in settings that are not anticipated by the designers. We follow Wyatt and Spiegelhalter [6] and others in emphasizing the importance of *clinical validation* of publets. Again, this should be routinely supported as part of a systematic publishing lifecycle [3].

We would go further by also emphasizing the need for supporting the development of an explicit *safety case* for applications where there may be significant

morbidity or mortality issues. Techniques for developing software safety cases, including HAZOP (Hazard and Operability Analysis) [7] and safety lifecycles [8] are well established in other areas of software engineering and need to be brought into this field. Several techniques are reviewed by Fox and Das [9] in the context of medical decision support and guideline systems.

Finally, since publets are designed for use in practical clinical settings we need to consider *usability criteria*. Much is known about good design of screen layouts, command languages, menu design and other interface functions (e.g. [10]). http://world.std.com/~uieweb/biblio.htm is a very valuable collection of resources. This body of knowledge is as relevant for web browsers and other user interfaces for publets as for any other class of software (e.g. [11]). The body of human factors research aimed specifically at medicine is somewhat scanty, and the demands and constraints of clinical environments seem particularly challenging (see [12] for a discussion of some general principles).

CONCLUSIONS

The internet, and also hospital and corporate intranets, are becoming important channels for publishing medical knowledge. We have demonstrated the feasibility of integrating decision support systems and enactable guidelines into such publications, and it seems likely that this will be an increasing feature of medical publishing in the future.

The combination of conventional electronic publishing and knowledge engineering will require new quality methods. A major component of this should be the development of an explicit publishing lifecycle and associated tools that support the creation of medical content, peer review of the behaviour, evidence base and usability of the applications, and publication should include comprehensive on-line documentation.

Although some version of conventional peer review may be sufficient for assessing the efficacy and usability of a publet we doubt that any form of peer review alone will be sufficient to ensure technical reliability and safety. For this we believe that it will be necessary to introduce the concept of a "safety case" from the software engineering community which documents the safety issues associated with the application, and the design and testing methods used to address them.

Acknowledgement: We would like to thank Jeremy Wyatt for his valuable comments on an earlier draft of the paper.

REFERENCES

[1] Sim I, Owens D K, Lavori P W, Rennels GD. Electronic Trial Banks: A Complementary Method for Reporting Randomized Trials. *Medical Decision Making* 2000; 20(4):440-450.

[2] Humber M, Butterworth H, Fox J, Thomson R. Medical decision support via the Internet: PRO*forma* and Solo. *Proceedings of Medinfo 2001.*

[3] Fox J, Bury J. A quality and safety framework for point of care clinical guidelines. In: *Proceedings of AMIA Annual Symposium* 2000, Los Angeles, 2000.

[4] Arden Syntax for Medical Logic Systems Version 2.0 (1992). Draft revision by Clinical Decision Support & Arden Syntax Technical Committee of HL7.

[5] Musen MA, Gennari JH, Eriksson H, Tu SW, & Puerta AR. PROTEGE II: Computer

Proceedings of MEDINFO 95, The Eighth World Congress on Medical Informatics, Vancouver, B.C., Canada, 766-770. 1995.

[6] Wyatt JC & Spiegelhalter D. Evaluating medical expert systems: what to test and how. *Med. Inform.* 1990, vol 15, no. 3, 205-217.

[7] Redmill F, Chudleigh M, Catmur. *System Safety: HAZOP and Software HAZOP.* John Wiley & Sons, 1999.

[8] Leveson N., *Safeware: system safety and computers.* Addison Wesley, Reading MA, 1995.

[9] Fox J & Das S. *Safe and Sound: Artificial Intelligence in Hazardous Applications.* AAAI and MIT Press, 2000.

[10] Wickens C D, Gordon S E & Liu Y. *An Introduction to Human Factors Engineering.* Addison-Wesley (Longman Imprint), New York, NY, 1998.

[11] Fleming J. Web *Navigation: Designing the User Experience.* O'Reilly & Associates: Sebastopol, CA, 1998.

[12] Cortes L. Designing a Graphical User Interface. *Medical Computing Today*, May 1997.