

# A platform and infrastructure for creating and sharing knowledge of best clinical practice and evidence-based decision making

## OpenClinical.net: White Paper 1, version 1, 2013

"Imagine the stimulus that a well-researched, evidence-based repository of standardized medical knowledge, with tools for delivering patient-specific advice at the time of need, would have on the ability to deliver computerised decision support services" Robert A Greenes (Ed) *Clinical Decision Support: The Road Ahead*, New York: Academic Press, 2007.



OpenClinical has not yet been formally launched and its goals, the services it offers, and governance and IP arrangements are still in discussion.

## Executive summary

1. The OpenClinical.net project is a radical reconfiguration and repositioning of the well-known [www.openclinical.org](http://www.openclinical.org) information service established in 2002.
2. The central objective of the new project is to create and maintain an **open access and open source repository** of medical knowledge, specified in a machine readable format, which can be used to support best clinical practice, and address the recognised demands of achieving quality and safety alongside efficient yet individualised care.
3. The specific goals of OpenClinical are to provide a:
  - a. Library of reusable decision, workflow and other machine-interpretable models for delivering high quality health and social care;
  - b. Software platform for creating and sharing machine readable knowledge and deploying applications and components at the point of care<sup>1</sup>;
  - c. Hosting service for members of the OpenClinical community to trial *Repertoire* content and for quality assessment and peer review;
  - d. Resource for researchers wishing to develop and test techniques and methodologies for using advanced IT to support best practice;
  - e. Data warehouse for information collected in c and d, available for data mining and other tools to support clinical research and further application development;
  - f. Practical infrastructure for rapid translation of new clinical research and new scientific findings into routine clinical practice.
4. The feasibility of the concept has been demonstrated and a complete software platform is operational;
  - a. The project has been showcased in the reports *Big Ideas for the future*<sup>2</sup> from Research Councils UK (2012), and *Doctor Know: A knowledge Commons for Healthcare* (NESTA 2013)
  - b. However, the service has not yet been officially launched at the date of this white paper as there are a number of key tasks still to be completed;
5. We have nevertheless reached the point where expressions of interest in the project, offers of collaboration or support, commentaries and criticisms on technical matters or governance issues would be very welcome indeed;
6. Please feel free to contact [john.fox@eng.ox.ac.uk](mailto:john.fox@eng.ox.ac.uk) about any aspect of the paper.

---

<sup>1</sup> We aim to launch with a range of applications that demonstrate a variety of decision-support, workflow and other knowledge-based services, including: clinical data capture and interpretation; physician order entry; risk assessment; trial eligibility; quality alerts and reminders; diagnosis and prognosis; test selection and pre-authorisation; treatment selection and planning; workflow management; coordination of care.

<sup>2</sup> <http://www.rcuk.ac.uk/Publications/reports/Pages/BigIdeas.aspx>

## Imagine...

A standard of care is developed and published as a clinical guideline by a healthcare quality organisation (e.g. National Guideline Clearing House in the US, NICE in the UK) or a healthcare provider (such as a hospital or clinical research group) in order to promote best practice in the diagnosis and treatment of cancer or some other condition. The authors now go on to prepare decision-making and workflow models that are described in the guideline, using the open logic framework supported by OpenClinical.net. The original guideline is enhanced, linking the recommendations for care to computer interpretable models of data, rules, decisions, evidence and treatment pathways.

The result of this process is a hybrid of human readable material and machine interpretable knowledge that translates general clinical guidance into patient specific recommendations and plans of care. This content can be relevant to risk assessment, investigation, diagnosis, treatment and many other decisions to assist healthcare professionals, and potentially their patients, by providing quality monitoring and safety alerts, adaptive workflow management and so on.

The key idea is that the informal statement of the medical knowledge in the guideline is enhanced with formal, computer executable representations of the knowledge. The enhanced guideline can be stored on an open access, open source repository from where it can be used by clinicians, researchers, technologists and others in developing IT services which provide patient-specific guidance at the point of need.

OpenClinical will provide a complete set of tools to achieve this vision. The goal is to help in the drive to improve quality and safety of patient care and reduce waste and costs common to health services everywhere.

### *The knowledge cascade*

Countless hospitals around the world develop and maintain collections of clinical guidelines that establish local standards of care. Such guidelines are often of variable quality but even when prepared to a high standard they are simply text documents and are frequently difficult to access, keep up to date and rarely consulted. OpenClinical.net aims to offer a new option, allowing healthcare providers to download guidelines developed by national, international or other authoritative bodies, adapt them to meet local priorities and constraints and integrate them with their local electronic records and IT infrastructure. If a user organisation chooses to create a localised version this can now be *republished* on the OpenClinical repository for others to adopt, or modify further and share again.

Whenever an application is used data about the patient can be recorded on an EHR (EMR) or research database, providing a powerful resource which can be used for clinical audit and peer review by healthcare providers, and to feed back front line experience to the original quality organisation and into research.

## Background - openclinical.org

The OpenClinical project was established in 2001<sup>3</sup>, initially as an information service and web portal for raising awareness and promoting safe design and use of clinical decision support and other clinical knowledge based services.

*“Goal: To promote awareness and adoption of advanced decision support and knowledge dissemination technologies and to inform and guide the proper use of these tools to improve quality, safety and efficiency of patient care based on scientific knowledge and rigorous engineering” ([www.openclinical.org](http://www.openclinical.org)).*

Until recently the OpenClinical web site was under continuous development and by 2011 it had about 600 pages of material<sup>4</sup> covering technical, clinical, and policy aspects of the field and countless links to other sources, projects and demonstrations. The site was widely used for reference and teaching and at peak was receiving more than 350,000 visitors per annum from across the English speaking world (about 15% UK; 50% USA). OpenClinical achieved the status of a “trusted brand”.

The dramatic growth of activity in eHealth over the last ten years, and the practical development of clinical decision support and other knowledge-based systems in particular, has meant that keeping information on the site up to date has become difficult for a single tiny organisation. However we believe that the project achieved its initial objectives so we are now taking a new and more sustainable but more radical direction to build on the success of the first phase.

## From OpenClinical.org to OpenClinical.net

From the outset we saw OpenClinical as an opportunity to be engaged in the creation and use of knowledge based services as well as providing an information portal. Our goal was to distribute practical tools and techniques for building applications for supporting quality and safety of patient care. OpenClinical.net was established in 2005 and has been developing steadily in a collaboration between University College London (Clinical Informatics Incubator, Department of Academic Oncology) and Oxford University (Department of Engineering Science)<sup>5</sup> with help from Deontics Ltd (a commercial spinout from Oxford, UCL and Cancer Research UK).

OpenClinical.net is based on two key foundations: the open standard *PROforma* process modelling language<sup>6</sup>, and the *Tallis* software suite, a complete “end to end” platform for authoring and publishing clinical decision support, workflow and other applications. *Tallis* makes it possible to design and implement a wide range of applications in *PROforma*, deploy them simply and quickly via a web client, and provide others with the ability to do the same. Members of the OpenClinical.net community will not only be able to publish their applications on an open access repository on the web they will also be able to download applications of other authors and reuse and adapt them for their own purposes.

<sup>3</sup> Cancer Research UK, J Fox and R Thomson

<sup>4</sup> Thanks largely to the sustained efforts of Richard Thomson, the managing editor of OpenClinical.org throughout this period.

<sup>5</sup> <http://www.ibme.ox.ac.uk/news-events/news/breast-cancer-patients-and-clinicians-benefit-from-technological-developments>

<sup>6</sup> Awarded the European Federation of Medical Informatics 20<sup>th</sup> Anniversary gold medal (“Laureate Prize” 1996)

The OpenClinical.net publishing system and knowledge repository *Repertoire* are fully operational; OpenClinical.net members can upload applications to *Repertoire* and download existing applications from it. They can then be imported into the *Tallis* authoring system in order to adapt and redeploy them. *Repertoire* presently carries about 30 example applications to demonstrate a range of applications for use in different clinical specialties and different sectors across healthcare. The publishing platform is functionally complete and we believe that the project is ready to be scaled up.

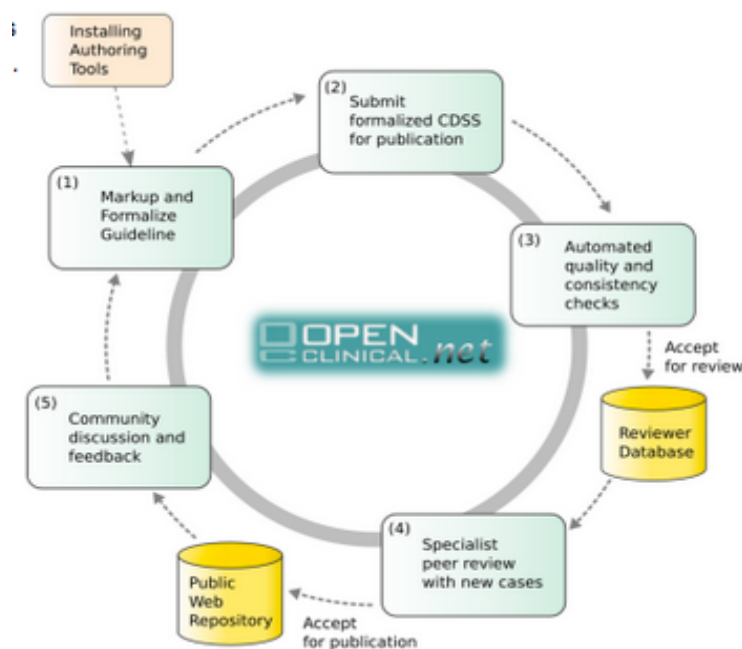
Our target now is to establish an international community of authors, developers, researchers and users who have the common aim of creating and sharing knowledge of best medical practice and helping to improve current methods for building high quality and safe knowledge based systems in healthcare. We are developing a systematic approach to ensuring quality of content on *Repertoire*, including technical methods and peer review, and discussing how to ensure safe and ethical use of the content.

## Current status

### Lifecycle and tools

OpenClinical.net has a fully operational suite of software with which authors can create, test and publish applications, based on the lifecycle diagram shown here.<sup>7</sup>

Members of the OpenClinical community can download the suite of tools. These are designed to be intuitive and flexible. Experience shows that they are easy to use, even for non-programmers (some of the best applications on OpenClinical were designed and implemented by clinicians without experience in software development).



### Repository of applications

The heart of OpenClinical.net is *Repertoire*, an open access and open source repository of “CDS-ready” content that perfectly fits Bob Greenes’ vision quoted above. OpenClinical.net aims to build on the reputation established by OpenClinical.org by supporting an international network of collaborating clinicians and scientists who wish to promote good practice in their field of special interest, either as individuals or under the auspices of professional bodies and other non-profit organisations.

<sup>7</sup> Deontics Ltd. permits use of the *Tallis* version 1.7 tool set at no charge for non-commercial use by registered members of the OpenClinical community, and we anticipate that this arrangement will continue indefinitely.

When an application is complete it can be submitted for peer review and possible publication on *Repertoire*. Once a “publet”<sup>8</sup> has been accepted on *Repertoire* other members of OpenClinical.net community, anywhere in the world, can download it and adapt it for their own organisations and circumstances.

At the time of writing *Repertoire* contains about 30 applications; at launch we hope there will be about 50. If OpenClinical.net achieves its goals, however, there will eventually be hundreds or even thousands of reusable applications on the repository that have been made available by their authors for others to use and adapt. We hope that *Repertoire* will evolve into a comprehensive library of applications and reusable templates, covering much of modern healthcare and regularly updated as new treatments and new research become available.

Finally, we see *Repertoire* as a research resource for medical informaticians, clinical and health service researchers, even social scientists and others. It will offer a collection of applications for members of the OpenClinical community to develop and test techniques and technologies, and a vehicle for clinical researchers to translate new biomedical science into new clinical practice. In this latter role OpenClinical will “close the loop”, feeding back information about the impacts and outcomes of new treatments and practices into research.

### *Kinds of content*

The kinds of knowledge content that we are able to publish on *Repertoire* are currently limited to computer-interpretable content that can be represented in the PROforma process modelling language, including rules, decision models, workflows etc. This limitation is for the practical reason that this is a language for which there is a published standard<sup>9</sup> and a complete suite of software to create and deploy applications, the Tallis suite. PROforma is a versatile language that has been shown to be able to support a wide range of the decision support and knowledge based services that the eHealth community is currently seeking to implement. However, it is anticipated that other types of content will be supported in due course.

The kinds of formats we would like to be able to support through OpenClinical.net are illustrated by the following list.

- Data standards, terminologies and ontologies (e.g. OpenEHR, RDF/OWL, SNOMED CT...).
- Rule sets and logic databases (e.g. Arden Syntax, GELLO, Prolog, SPARQL, OWL, AIF)
- Guideline models (e.g. ASBRU, GLIF, PROforma)
- Statistical inference models (e.g. XBN)
- Task networks and workflow models (e.g. ASBRU, GLIF, PROforma)
- Document models (e.g. CDA, GEM)

General requirements for inclusion of content on *Repertoire* will include: (a) the content is defined in a format that is either a *de facto* standard or (b) is used and advocated by a number of members of

---

<sup>8</sup> Contraction of “publication” and “applet”.

<sup>9</sup> For practical reasons OpenClinical will launch using the PROforma modeling language which has been published as an open standard (Sutton and Fox, *J Am. Medical Informatics*, 2003) and with which we now have a great deal of experience. Our aim is to broaden the base of formalisms and tools as soon as we can, to ensure that OpenClinical has as wide a constituency as possible and to support the community in comparing different options and learning how they can be best used.

the OpenClinical community, (c) it has a set of associated software tools for creating content that complies with the standard, and (d) there is an existing capability for, or practical prospect of, interoperability between this content and the publets already published by OpenClinical.net.

We do not currently expect to publish much conventional content (text, pdf, xml etc) except for documentation of apps, and articles that have been accepted for publication in the *Point of Care* research journal to be jointly published by OpenClinical and an established medical publisher.

The OpenClinical project also aims to support internet-based collaborations, where there is a wish to establish and maintain a consensus about best practice in specialist areas. Consensus groups may be local, national or even international. OpenClinical intends to have a purely facilitating role in the creation and operation of these special interest groups; broad governance principles will be applied but such groups should be otherwise independent of OpenClinical and pursue their own research and development programmes as they see fit.

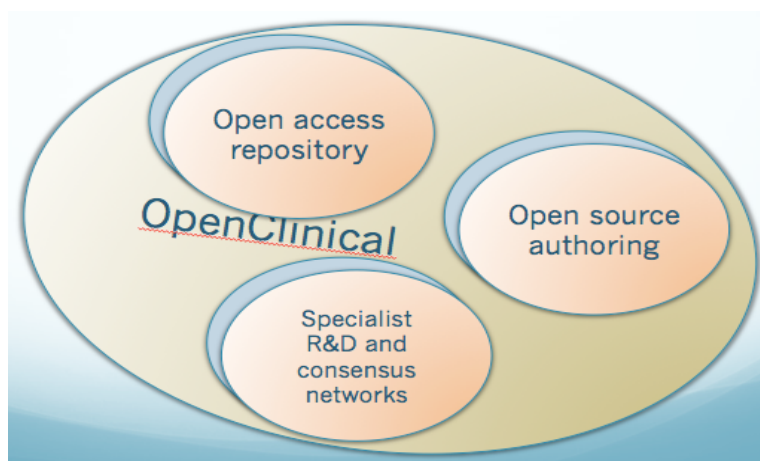
## The “business model”

Nowadays we all know about many different kinds of internet organisations with diverse business models; OpenClinical.net can be viewed as a fusion of three well-known types of organisation.

**Guideline Repositories.** The evidence-based medicine movement (notably the *Cochrane Collaboration*) has made huge strides in assessing and distilling new research on patient care, and promoting the importance of taking into account objective research in clinical decision making. Countless medical organisations now draw on this research to develop practical summaries of care in specialist areas in the form of clinical guidelines for healthcare professionals and patients and publish them on open repositories. UK readers will be familiar with the NICE collection; internationally the National Guideline Clearing House (<http://www.guideline.gov/>) has been established as a central repository of these summaries. More than 300 organisations in the English-speaking world deposit copies of their practice guidelines in the NGCH with the intention that they should be used to inform the development of high quality and cost-effective healthcare services. However these guidelines are only intended for people to read, and their real impact in changing practice is known to be limited. *Repertoire* aims to become a repository of knowledge of best practice which is analogous to the NGCH but whose content is *machine readable* and can therefore be used to provide active guidance at the point of care.

**Wikipedia.** In Wikipedia authors can prepare articles on any topic in which they believe that they have specialist knowledge and they wish to publish as part of this well-known *open access* web encyclopaedia. Once the article is published it can also be improved and extended by other knowledgeable authors because Wikipedia content is also *open source*. This model means the content can reflect different points of view and knowledge of the topic, and can foster the establishment of a trustworthy consensus. It also addresses the problem of ensuring that content is maintained, particularly where the domain of the article is changing over time, as in science, engineering, medicine and so forth. OpenClinical.net builds on the Wikipedia model by providing tools for medical experts to author and disseminate their knowledge in a machine readable format that can be used for research or to provide guidance in the care of individual patients.

**LinkedIn.** Many of us who may or may not be Facebook or Twitter enthusiasts increasingly see *LinkedIn, ResearchGate, Mendeley* and other professional networking tools as important for maintaining awareness of current developments, monitoring who is doing what in our fields, and achieving visibility for our own contributions. These platforms offer many functions, and are still evolving, but they clearly point the way to how science and engineering are going to achieve more effective dissemination of new knowledge and practice in the future, and promote more effective interactions between academia, industry, public policy organisations, and so on. OpenClinical.net aims to bring these kinds of benefits to the world of medical research and clinical practice, to encourage collaboration and consensus building and promote compliance with best practice within and across clinical specialties.



The content created by **NGC, Wikipedia** and **LinkedIn** and other web communities is in natural language, and therefore only intended to be written and read by people. This seriously limits its value. The unique feature of OpenClinical is that *Repertoire* content is written using specialised computer languages that were developed for describing medical expertise in formal models, including logic and rules, decisions and plans, and the specialist knowledge needed to implement care pathways and protocols. By formalising medical knowledge in this way it becomes possible for a computer to apply current knowledge to electronically recorded data about individual patients and assist healthcare professionals follow evidence-based recommendations while complying with local organisational constraints and policy requirements.

## Governance

Four principles will govern the activities of the OpenClinical project, and the rights of its members and sponsors. The CAFÉ principles are still under discussion and should therefore be considered only as indicative at this stage.

**Community:** the primary goal is to support the ability to create, document and share models of clinical expertise (e.g. rules, decisions, plans and pathways) and other kinds of computer readable knowledge for the benefit of healthcare professionals and their patients and to promote engagement with interested biomedical and clinical research groups.



**Accessibility:** we will provide *open access* to this content for clinical and research use and wherever possible make this content *open source* to permit reuse of the content by others under licensing arrangements to be agreed by the community.

**Fairness:** we will seek to ensure fair recognition of the efforts and value created by all stakeholders in the OpenClinical project (including content authors, OpenClinical itself and its sponsors) while ensuring that OpenClinical remains independent of commercial and other interests.

**Empowerment:** to empower stakeholders to achieve any objective that is concerned with improving the quality and safety of healthcare, including empowerment of authors to disseminate their expert knowledge; organisations to promote best practice; clinicians to ensure they offer best care; healthcare providers to manage resource use; scientists to implement trials ... and patients who use services derived from OpenClinical.net content).

The OpenClinical membership rules will let anyone participate who wishes to contribute to any aspect of the work of the project who has the relevant skills (content authors, software developers, project managers, peer reviewers etc). Membership of the OpenClinical community will in the first instance be limited to individuals and organisations that have been nominated by existing members, but we hope to lift this restriction when the project is more mature and we are confident that we are capable of coping with a larger and wider membership.

Now that the technology and a first version of the [www.openclinical.net](http://www.openclinical.net) web site are in place the primary objective is to establish OpenClinical as a recognised tool for creating and sharing CDS-ready content and to build the *Repertoire* knowledge base into a credible resource for technical and medical professionals. At launch we expect to have 50+ applications on *Repertoire*, of many types and levels of sophistication, including a diverse range of clinical demonstrations, tutorial examples and reusable templates and components as well as clinical-strength applications. Launching at a significant scale will create credibility for the project and encourage a growing range of organisations and individuals to begin to use OpenClinical as a new kind of engine for driving evidence-based medicine and promoting best clinical practice.

## Success criteria

The value and timeliness of the OpenClinical.net project are based on proven clinical need, demonstrated technical feasibility, and a critical mass of researchers who are keen to “get going”.

Direct impact measures will be (1) the establishment of an international community of practitioners who create, share and exploit CDS-ready content on an open licence basis; (2) the number and range of applications accepted for publication on *Repertoire*; (3) the number of applications adopted for research or clinical use.

Indirect measures of success will include (1) publication of results from collaborative research in peer-reviewed journals; (2) demonstration of effective and scalable processes for editing and curating *Repertoire* content, (3) 3<sup>rd</sup> party adoption of the OpenClinical method and (4) successful foundation of the *Point of Care* journal.

## Sustainability

We aim to establish OpenClinical as an independent non-profit organisation that will raise funds from a number of sources to pursue its objectives. It will be established initially in the UK as a social enterprise or charity<sup>10</sup>, but once the methods and tools have matured we hope to develop the project internationally, perhaps within a year following the UK launch. We will not seek to establish a commercial operation but must aim to become an efficient and sustainable operation.

### *Potential sources of funding*

In the academic sector we will work with collaborators in the OpenClinical community to help them obtain research grants (from the NHS, NIHR, DoH, Research Councils and so forth in the UK). OC.net is a highly innovative research project in its own right, but it also offers a platform for clinical, technical and social science projects in eHealth which can be used for independent research by collaborating groups. We are working with several UK and non-UK teams and our aim is to bring additional groups into the project as early as possible and support them in developing their research programmes.

Given the reputation and international user base achieved by [www.OpenClinical.org](http://www.OpenClinical.org), we believe that we have an excellent foundation for attracting general sponsorship and sponsorship in specific clinical areas.

An important component of the OpenClinical project will be *Point of Care*, a research journal focused on all aspects of the design, development, deployment, evaluation and use of CDS and other knowledge based services. *Point of care* will adopt the business model of open access journals, attracting publication fees to cover publication costs and *Repertoire* hosting costs. We expect to develop and run the journal through a partnership with an established open access medical publisher.

By making *Repertoire* open access and open source we do not wish to imply that we wish to prevent commercial use of the content on OpenClinical. On the contrary, authors who wish to licence their applications to third parties will be able to do so on an appropriately adapted open source contract. A substantial fraction of the income from commercial use of an application will go to the authors but a portion of the income will be retained to cover the running costs of the OpenClinical service.

---

<sup>10</sup> The pros and cons of each option are currently the topic of consultations with others having experience with both models.