Care pathway records and variance data: Enabling research through the use of ontologies

Mark Olive
CSCT, University of the West
of England, Bristol, UK
Mark2.Olive@uwe.ac.uk

Alison Lashwood
Clinical Genetics, Guy's
Hospital, London, UK
Alison.Lashwood@gstt.nhs.uk

Tony Solomonides
CSCT, University of the West
of England, Bristol, UK
Tony.Solomonides@uwe.ac.uk

Abstract

Integrated care pathways (ICPs), a fine-grained form of medical guideline including the explicit recording of any deviation, or 'variance', could serve as a rich source of data for research. Not only do they incorporate a wealth of operational knowledge, but feeding the results of the analysis of variance into the development of a pathway could be an effective way of capturing evidence from practice. In our principal case study we propose a system for extracting data from care pathways with the aid of ontologies, and a method for inferring ICPs from other patient records, combining these with data collected for retrospective and prospective studies in preimplantation genetic diagnosis (PGD) for assisted reproduction.

1. Introduction

Our previous papers summarised research into the development and use of Integrated Care Pathways (ICPs), and the potential value of variance data – and practice-based evidence in general – for healthcare research ^[1]; this paper expands on that work, and proposes a practical application for research purposes of a general purpose ontology for modelling care pathways. In part, this work arose out of our involvement in the EC funded SHARE project. The SHARE project aimed at identifying the most important steps and significant milestones towards the deployment and adoption of grid technology for medical research and healthcare in Europe. Part of the project's final report explored the future use of this technology for general healthcare, taking Chronic obstructive pulmonary disease (COPD) as an example. COPD refers to an airway obstruction caused by chronic inflammation. It is usually progressive, not fully reversible, and often occurs as a result of smoking

but other factors such as air pollution can also contribute to the development of COPD. In the UK almost 900,000 people have been diagnosed with the disease, and the true number of people suffering from the condition is estimated to be around 1.5 million.

According to NHS guidance, the management of the disease should be tailored to the individual, with adjustments being made based on responses to treatment. The guidance includes a large number of drugs, including some 'off-label' drugs (drugs either not certified for the treatment of this condition, or for the particular group of patients being treated). In the UK, the National Institute for Health and Clinical Excellence (NICE) has issued national guidelines for the treatment of the disease, but such guidelines are frequently modified to account for local variations and priorities. As a result, the procedure for assessing and treating the disease can vary around the country. The evidence on which this guidance is based comes from a variety of sources, such as national studies by NICE and systematic reviews with an international scope.

Two main concerns were expressed in the SHARE report for general healthcare; supporting the travelling patient, such as migrating elderly populations, and enabling decision support systems that can account for local variations in best practice and clinical evidence.

For example, the evidence supporting the use of an off-label drug might not be available to a doctor from a particular locale. A patient moving from the US to the UK, for example, may well be taking a drug certified for them in the former, but not the latter. In order to continue treating the patient concerned, the doctor must be aware of the evidence and guideline/pathway that informed the plan of care for that patient, and any deviations from that plan that have occurred to date. This may not be a trivial matter of simply retaining a link to the relevant material, as there may be language barriers (when patients travel between European states, for example), and local reasons why the guidance

followed in one locale would not be appropriate in another.

The report also noted that not only can various forms of evidence feed into guidelines, but in the healthcare domain, the knowledge can be both declarative and operational in nature. Declarative knowledge is probably most often disseminated through peer-reviewed publication. On the other hand, the ‘best ways’ of treating patients – in a particular context – are often solely described in care pathways and guidelines. In SHARE’s vision of future healthcare, these are brought together for the better treatment of patients and at the same time to improve research, through appropriately controlled secondary use.

Clearly, guidelines are important, and failure to adhere to them can have disastrous results. High profile enquiries into healthcare failings, such as the investigation of Mid Staffordshire NHS Foundation Trust^[5], inevitably mention inadequate adherence to accepted national clinical guidelines. The fact that particular guidelines had not been adopted or others were not being followed with sufficient care is either seen as a failing in itself, or as an indication that the overall quality of care is not what it should be. As a result of this and other initiatives to improve the quality of care provided by the NHS, following guidelines (and the more fine-grained care pathways) is increasingly viewed as necessary in order to meet the required standard of care, and those who are unable to follow them must be able to provide an adequate reason.

2. Integrated Care Pathways (ICPs)

ICPs are, according to the European pathways association, a tool to improve the quality and efficiency of evidence-based care, and are a means of operationalizing guidelines. They provide a patient-centred multidisciplinary ‘journey’ through the planned course of care, documenting the actual care received by the patient. Clinicians are not only reminded of the best practice for the condition concerned, but are actively guided through the steps required. ICPs are typically developed from one or more guidelines.

In order to prevent ICPs from becoming a form of ‘cookbook medicine’, they incorporate a mechanism called variance tracking. A doctor or nurse records any variations from planned care as it occurs, with the reason and any resolution. To quote Middleton et al., “while the ICP acts as a template of the care to be provided to the chosen group of patients, it is not intended to compromise clinical judgment. Any member of the clinical team can deviate from the

pathway if there is a valid reason for doing so. In essence, the pathway asks each clinician to determine whether each defined intervention is appropriate for a given patient, thereby promoting clinical freedom based on the needs of the individual”^[4]. Therefore variance is to be expected, merely ‘an alternative plan of care’ as opposed to a more negative title such as ‘non-compliance’. The tracking of variances from the plan of care is the single most important difference between guidelines and ICPs. The facility to depart from the pathway, provided that such departures are recorded and explained, should address the common criticism of guidelines and protocols being too prescriptive.

Variance tracking forms a record of “things that have been tried, and what happened in order to take the next steps”^[3]. The standard method of recording variance involves ‘triggers’ (typically a particular goal, milestone or outcome being/not being met). Where the outcome was not a successful resolution of the problem (where the variance was as a result of a problem to begin with), a further entry will be added and so on, building a chronological record.

An ICP is therefore also a way of documenting the care received by a patient, complementary to or even a replacement for the patient record for the period of care concerned.

In earlier papers, we examined the idea of using variance to collect ‘evidence from practice’; the analysis of variance records could indicate where a pathway should be modified, such as creating a new branch or alternative pathway. We argued that while such evidence could probably not be used to challenge ‘gold standard’ evidence, it was useful for more than simply ‘filling the gaps’, and could serve as a useful complementary resource. Some care pathways, such as the Liverpool Care Pathway for the Dying Patient (LCP) have achieved a fairly high degree of harmonization across organizations, allowing ‘tertiary’ analysis of variances to take place – this has been recently highlighted by the Healthcare Quality Improvement Partnership (HQIP). However, such harmonization currently relies upon a low level of local customisation; with too many differences introduced to account for local priorities and resource issues, this tertiary analysis would not be possible. The system outlined later in this paper offers a potential solution to this problem.

We have also identified two further examples of where branching or competing pathways would be beneficial to compare and understand processes requiring more formalisation. The first was mentioned by a consultant gastroenterologist. A number of their

patients have a certain amount of malnourishment due to having 'nil by mouth' prior to surgical procedures. A number of tests are conducted to check for signs of refeeding syndrome (where potentially fatal metabolic disturbances occur as a result of introducing nutrition too rapidly), but ultimately the consultant makes a judgement of how the patient is doing, and decides how aggressively nutrition should be reintroduced. In this case, the consultant and a dietician frequently disagreed on the correct course of action; tests indicated possible early signs of refeeding syndrome, but the outcome for the patients concerned was always positive. In order to determine the correct course of action for any given patient, the process of determining which patients should return to full nutrition would need to be formalised, with set criteria that should be satisfied; once in place, a branching pathway could be employed to refine these. The second example was very similar, but involved the use of thrombolytic therapy for patients that had had a stroke within a certain period of time. In this case, the decision whether or not thrombolytic therapy would be beneficial was related to the period of time that had elapsed since the onset of stroke; the assessment appeared to involve more than just a fixed cut-off time, but again the doctors concerned felt that the process had not been sufficiently formalised.

3. Gathering data for research into PGD

With electronic care pathways (eICPs or ECPs) and the electronic health record (EHR) both being advanced as enhanced sources of data for research purposes, an area of interest here would be how data from both sources could be intelligently integrated. An example of how such integration might be performed has recently emerged through our work on the EuroPGDCode project, which involved a pilot data collection study.

Preimplantation genetic diagnosis (PGD) is a treatment for patients that have or are carriers of an inherited genetic disorder, which might be a chromosomal abnormality or genetic mutation, which could imply a serious health risk for any children that they might have. PGD is conducted alongside assisted reproduction, such as in-vitro fertilisation (IVF) or intracytoplasmic sperm injections (ICSI). A number of embryos are produced, and the DNA and chromosomes of each are tested for the suspected disorder by removing one or two cells for analysis. Embryos that are unaffected by the disorder may then be transferred into the uterus of the mother, where they would develop normally. A licence to practise PGD is

required from the Human Fertilisation and Embryology Authority (HFEA), and there are other legal issues, such as those relating to PGD for specific conditions, that must also be considered.

Preimplantation genetic screening (PGS) is similar, and employs the same techniques as PGD to identify embryos that may have a serious genetic disorder, but does not look for a specific condition. Typically, the reasons for a patient having PGS would be related to their age and/or a history of miscarriages, or problems with embryo implantation.

In order to gather more information about the efficacy of PGD and PGS, members of the project based at Guy's Hospital, London (the largest and most successful centre in the UK for PGD) produced a paediatric follow-up questionnaire, which was circulated to centres around the world. This questionnaire was completed by mothers that had successful PGD/PGS cycles, with the centres including additional information about the cycles and children involved.

This data has now been collected, and staff at Guy's are in the process of entering it into a database designed by project members at the University of the West of England, Bristol. The web interface for this database includes the ability to query the data held. Although this is at an early stage, already several questions have emerged.

Some of these are as follows:

- ◆ Did the parents have had any medical conditions either related or unrelated to the genetic condition in their family? Where an abnormality was picked up in a child, is this something that has been seen in either parent previously? Could it be prevalent in that family (i.e. asthma)?
- ◆ Where there were obstetric complications during the birth of a PGD baby, had this occurred in previous pregnancies for that couple? For example, in relation to prematurity or intra-uterine growth retardation, is this something common to the family?
- ◆ Did the patients have a choice of suitable embryos at the point where embryos were transferred? If so, then the chance of pregnancy can be greater as the healthiest are selected. If there is no choice, the embryo that is available is transferred. But does this have any impact on the outcome of children born? Where there was no choice of embryo, to what extent is it more likely that the children born from these PGD cycles would have abnormalities or health problems?

While other questions could be answered by using data collected via the questionnaire, these questions require more information about the patients and periods of care they have had previously. What would be needed is a way of extracting the relevant pieces of data from records that correspond to those previous episodes of care, and this is what we will now propose.

As with the questions mentioned above, the centralised PGD database serves as a starting point, allowing the researcher to formulate their questions, perform exploratory queries and identify patients of interest. All records held in the database are

pseudonymised. The next step is to identify which ICPs are involved and contain data of interest to the researcher.

The patient identifiers and a list of required ICP records are then sent to the distributed centres. Each ICP would need to be described by an ontology previously defined. Hurley and Abidi^[8] (later extended by Zhen et al.^[7]) have created a generic ontology for care pathways, incorporating temporal aspects, outcomes and variance, which serves as our starting point; Zhen et al. also demonstrated how this could be embedded into an EHR.

Figure 1. An example of relevant variance that might occur following the delivery of a PGD baby, using the generic care pathway ontology defined by Hurley & Abidi^[8]

Where the ICP in use is recognised, data must now be extracted from the ICP record. Using an ontology to aid in this extraction allows for some variation between ICPs, rather than a formal model approach which would be less flexible. For example, Daniyal et al.^[6] demonstrated how location-specific pathways could be unified by using an ontology for prostate cancer pathways. When extracted, the data is then pseudonymised and sent back.

Where the data is held in an EHR rather than a recognised ICP, an attempt must be made to match the record to an ICP. A branching care pathway or guideline for a specific diagnosed condition could be expressed as a series of possible paths, each with anchor points such as referrals or key test results. The EHR could be mined for these anchor points, which may correspond to periods where a patient's care followed a particular guideline or pathway. The details between the anchor points can then be analysed to determine if any variance was likely; at a basic level, the length of time between anchor points could be a

good indication that the course of care may not have been as predicted. The original pathway would then be augmented with comments to indicate where complications are likely, further branching might be beneficial, etc. The formalisation of the guideline or pathway would be a prerequisite for this approach, with *PROforma* being one example of a formal guideline modelling language; Seyfang et al. have reported some success in easing the conversion via an intermediate ontology^[9]. This approach also has other applications, such as monitoring compliance for example. Map of Medicine, as used throughout the NHS, is interested in knowing when and why clinicians are consulting and deviating from their pathways, but this is still currently limited to monitoring access.

Where a match is successfully made, the care pathway ontology would be employed to extract the relevant data. In order to conduct a case-control study, records for controls could also be selected in a similar way at each site.

The use of a grid or cloud could be a logical approach here, given in part that the ICP data is distributed, ICPs (and control over them) vary at a local level, and the potentially computationally intensive task of extracting the required data. As the data involved is in part from ICP records, the subsequent analysis could also include the identification of missing data of interest that could be collected (either in the form of coded variance, or additional items that should be recorded), and any branches in the pathway that might be beneficial.

4. Resolving drug interactions

We are also examining the potential use of an ontology for resolving problematic drug interactions that can occur when patients are following multiple care pathways in parallel, as is often the case. Several current decision support systems based on formal guideline modelling languages such as GLIF and PROforma include the facility for alerting the user when a problematic drug interaction may occur, but it seems to be left to the user to determine from the referenced literature what the severity and nature of that interaction might be.

Figure 2. The main classes and selected properties from an ontology of drug interactions

Far more detailed information exists from Cerner, the British National Formulary and other sources on drug interactions that could be exploited. Alongside a modified version of Hurley and Abidi's ontology, this could provide useful information on alternate drugs for substitutions, required modification of pathways to incorporate drug management tasks, and allow variance data to feed into a ranking process for alternate drugs, as well as an ongoing evaluation of the prescribing of off-label drugs.

Our next step will be using this ontology to develop a decision support system for pathway orchestration within a specific domain. Although contraindications

are included, identifying which would apply to the specific patient would be a manual process here as the necessary data would not necessarily be present in a pathway, so the clinician would need to refer to the patient record.

5. Conclusion

The system outlined above involves integrating information in a variety of forms; retrospective and audit data collected via the questionnaire and stored in a central repository, highly structured and fine grained data from care pathways, stored at remote sites, and the

semi-structured notes of an EHR (albeit mapped to a care pathway). The sensitive nature of the data requires extensive pseudonymisation, which may not be trivial given the level of detail involved. The use of ontologies allows data from pathways with local customisations to be aggregated, and aids in the extraction of the data. The ‘fuzzy matching’ approach of mapping EHR records to one of several paths through an ICP is a particularly interesting problem, and one that would also be of interest for monitoring compliance with pathways. We hope to pursue this soon.

Our previous work involved examining ICPs as vehicles for conducting research and collecting ‘evidence from practice’ in the form of variance records. This paper highlights another example of how ICPs might be exploited for research purposes, and how data from both ICPs and EHRs might be intelligently combined through a process involving the mapping of EHR data onto an ICP model, and the use of an ontology to aid in the subsequent extraction of that data. The drug interaction ontology being developed will also demonstrate the potential of using ontologies with ICPs for resolving such difficulties in healthcare. As a generic care pathway ontology is used rather than a specific ontology for PGD, this approach has fairly wide applicability.

5. References

- [1] Olive M & Solomonides T (2009). Variance analysis as practice-based evidence. In Healthgrid Research, Innovation and Business Case: Proceedings of Healthgrid 2009, Studies in Health Technology and Informatics, Volume 147, IOS Press.
- [2] Andoulsi I, Blanquer I, Breton V et al. (2008). SHARE the Journey: A European Healthgrid Roadmap. European Commission, Directorate General Information Society and Media. ISBN: 978-92-79-09668-6.
- [3] Proceedings of Integrated Care Pathways 2007, London, June 2007.
- [4] Middleton S & Roberts A (1998). Clinical Pathways Workbook, Clinical Pathways Reference Centre, VFM Unit.
- [5] Wood H, Fontaine P, Harvey D et al. (2009). Investigation into Mid Staffordshire NHS Foundation Trust. Commission for Healthcare Audit and Inspection. ISBN: 978-1-84562-220-6.
- [6] Daniyal A, Abidi S R & Abidi S S R (2009). Computerizing Clinical Pathways: Ontology-Based Modelling and Execution. Medical Informatics in a United and Healthy Europe, IOS Press.
- [7] Zhen H, Jing-song L, Hai-yan Y et al. (2009). Modelling of Clinical Pathways Based on Ontology. IEEE International Symposium on IT in Medicine and Education (ITIME).
- [8] Hurley K & Abidi S S R (2007). Ontology Engineering to Model Clinical Pathways: Towards the Computerization and Execution of Clinical Pathways. IEEE International Symposium on Computer-Based Medical Systems (CBMS).
- [9] Seyfang A, Miksch S, Marcos M et al. (2006). Bridging the Gap between Informal and Formal Guideline Representations. European Conference on Artificial Intelligence (ECAI-2006).