

Telematics for Clinical Guidelines: A Conceptual Modelling Approach

Colin Gordon¹, Ian Herbert², Peter Johnson³, Peter Nicklin², Philip Reeves¹

1. Royal Brompton Hospital, London

2. NHS Information Management Centre, Birmingham

3. Sowerby Unit for Primary Care Informatics, University of Newcastle

Abstract. *PRESTIGE* is a project for applying telematics to assist the dissemination and application of clinical practice guidelines and protocols. Previous publications have described *PRESTIGE*'s technical approach, including the use of a generic model for representing the knowledge content of clinical guidelines. [1,2,3] This approach offers the possibility of 'plug-and-play' electronic distribution of clinical guidelines produced by multiple authoring bodies for use on multiple healthcare clinical management software platforms. A recent joint workshop held with the Section on Medical Informatics, Stanford University School of Medicine compared the European consensus approach developed in *PRESTIGE* with a parallel series of projects for computer-assisted protocol-based healthcare undertaken at Stanford and other American centres over the past, which confirmed the convergence and complementarity of our approaches, and holds out prospects of world-wide standardization in healthcare protocol knowledge representation. In this paper we summarise the *PRESTIGE* conceptual model set which is the design foundation of our approach.

1. Background: elements of the models.

1.1. Knowledge

Guideline applications are **knowledge-based** applications. Guidelines are modules of clinical knowledge and the function of telematics guideline applications are to assist the availability, accessibility and appropriate use of that knowledge where it is relevant in clinical practice. *PRESTIGE*'s approach, derived from the *DILEMMA* project and other earlier work in clinical knowledge-bases systems, uses an explicit, declarative representation format in knowledge bases which facilitates the adaptation, flexible use and maintenance of clinical knowledge. This approach to knowledge representation requires an explicit knowledge model defining common and standard formats for computerised representation and use of items of knowledge. The *PRESTIGE* models incorporate and refine the *DILEMMA Generic Protocol Model* which offers a common electronic format for representing all clinical guidelines across specialities, user sites and application software platforms.

Clinical guidelines contain knowledge about medical concepts such as diagnoses, therapies and symptoms. They also contain knowledge about how to perform specific healthcare acts and activities such as monitoring consultations and treatment planning. Therefore a model of clinical guideline knowledge needs to call on a **model of healthcare activities and processes**. In *DILEMMA* this need was addressed by use of the UK National Health Service's *Common Basic Specification* (CBS: now retitled the NHS Healthcare Model), a model describing all forms of business performed by a health service. [Add stuff on acts here.]

Types of protocol entities.

The *PRESTIGE* Protocol Model, which forms part of the *PRESTIGE* Conceptual Model, provides rules defining the knowledge statement constructs which may or must be present in a *PRESTIGE* protocol knowledgebase. As an Object model, it defines the kinds of entities (objects) which may appear in the protocol, with their permitted relations and attributes. Several main types of these objects can be identified.

- As protocols are knowledge, they will involve the use of general *concepts* (concepts of activities, acts and case-specific phenomena such as diagnoses and symptoms).
- protocols will have a structure consisting of the overall protocol and its component parts (which are all *individual, named instances* of the concept ‘protocol version’);
- protocols will include *expressions*, built out of concepts, instance names and syntax operations (‘and’, ‘or’, ‘at least...’...) as prescribed in the BNF grammar rules for an expression type (cf. above), which can be used in several roles within the protocol:
 - * conditions governing the appropriate use of a protocol or its parts (such as when a specific act should be performed). A condition is a (sometimes complex) specification of a (simple or composite) ‘phenomenon’ which can be shown to be true or false in the current context of an individual healthcare case. A phenomenon may consist of
 - * individual patient attributes, symptoms and diagnoses (subject characteristics),
 - * facts about things done (recorded acts and their attributes and states),
 - * facts pertaining to a context of care (‘patient already in the care of a chiropodist’; ‘local availability of open access echocardiography’).
 - * ‘method’ expressions to define the specific attributes of an act to be performed (e.g. a drug dosage)
 - * ‘templates’ which define a set of data items which are to be collected

How parts of the Conceptual Model are interlinked

Several of the same types of object occur (either as uses of concepts or as named instances of a concept) in the other Conceptual Models, so that the different segments of the models are interlinked. Clinical concepts are used in the patient record to describe individual (instance) patient characteristics; act concepts are used to record things done; ‘protocol use’ instance objects will reference the individual (instance) protocol version, defined in the knowledge base, which is being used in an individual patient case. Patient records will identify specific agents as observing data or performing acts; an enterprise database will identify the agent type (concept) to which an individual agent belongs, and the roles and accountabilities which it can take on; a protocol knowledgebase may identify the agent types needed to execute a part of a protocol, or as recipients of a referral action specified in a protocol. The use of clinical terms used to express all concepts mentioned here will be constrained, in each application site, by a common terminology/semantics model.

For example: if a protocol is to specify the activity in which it is providing support to a clinician, the term it uses must exist in the application site terminology/semantics model, and must be classified in that model as a subtype of the concept ‘activity’.

2. Patient Data

To implement telematics support for clinical guidelines, we need not only knowledge bases and knowledge models, but a mature electronic patient record. Clinical guidelines define patient-specific recommendations where the option selected is a function of the patient's medical condition and other factors. A useful computerised aid will be able to detect which options are relevant for an individual patient by directly interrogating the patient record. (patient-specific prompts have been shown in randomised controlled trials to enhance the effectiveness of guideline implementation). Moreover a record should be capable of keeping track of how a guideline is used over time, since later options are also determined by earlier choices made. An important part of the *PRESTIGE* models is therefore the Electronic Patient Record [EPR] model specification which defines what facts and data a *PRESTIGE* application EPR must be able to represent.

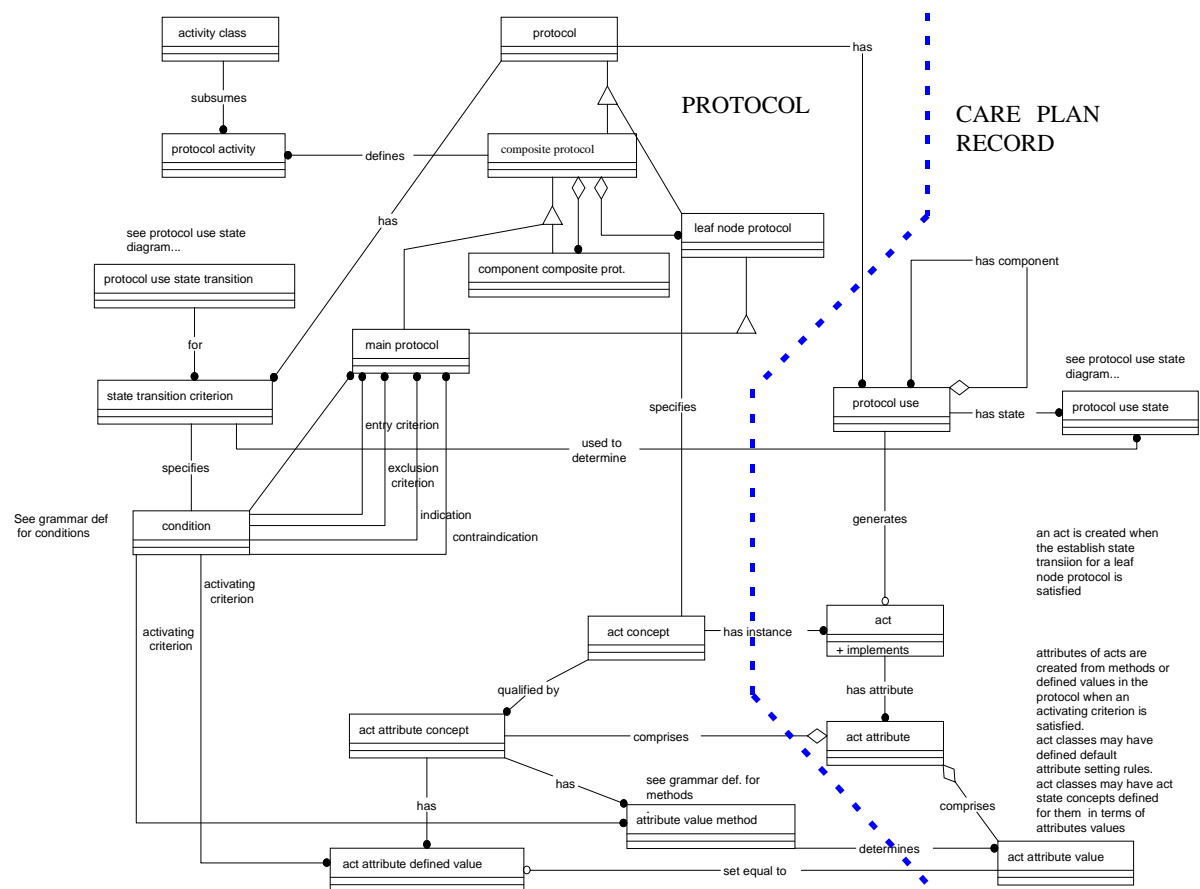


Figure 1. Protocol and care plan record: syntax and semantics of the protocol model

3. The Healthcare Enterprise

As part of the planning and designing *PRESTIGE* applications, we analyse and document the existing clinical business process which a guideline is intended to assist, and which form the context where it will be used. This activity uses a methodology for business process re-engineering developed during the 3rd Framework AIM project SHINE. This approach identifies the organizations, professionals, data items and transactions involved in the cooperative delivery of healthcare by health services to patients. The models produced also allow a definition of how where guidelines and associated technology will be inserted

into modified versions of these processes (and, in some cases, how use of the guideline is intended to modify an existing process).

This type of analysis is not directly dealt with in the *PRESTIGE* conceptual model set since the models produced do not have to be expressed in a formalised notation and, although a vital step for application requirements definition, are not an immediate basis for software design. The project conceptual model set does, however, need to provide a formal representation of some healthcare organizational and administrative concepts which need to be explicitly addressed within software design. These include concepts about the healthcare **enterprise** (healthcare professionals, service units and services provided; types of resources required in providing specific services), together with administrative, demographic and **identification** data about individual patients. (Identification data about patients do not often play a significant role in clinical guideline themselves, but a consistent overall model of such information is nevertheless important in guideline applications, because cooperative healthcare professionals involved in applying a shared guideline (e.g. hospital staff and GP) may use different means to identify the same patient.)

Enterprise information enters in guideline applications in two ways. Often clinical guidelines need to be tailored to fit the specific constraints and circumstances of the local healthcare enterprise (such as the availability of specific skills and services): the local version of a guideline knowledgebase may then need to be extended or modified to refer explicitly to some local resource criterion; in applying a guideline, it may be necessary also to test such criteria by querying the current availability of local resources or services in an enterprise database. For both these services, a standard vocabulary and syntax must be provided for each enterprise to maintain its own such database and express enterprise constraints or criteria in its locally adapted guidelines.

4. Clinical concepts and terms.

PRESTIGE provides a set of generic models for guidelines, patient records and enterprise information which are independent of language and clinical domain. Within the language and clinical context of a given application site, it remains essential that a common and consistent clinical language is used in representing the site's clinical guideline knowledge, patient data and enterprise information. Without this common vocabulary, the telematics application will be unable to make proper combined use of these information sources in order to assist clinical professionals using a guideline. Standard clinical coding systems such as ICD, ICPC, SNOMED and (in the UK) Read Codes are currently able to meet part of this requirement, but are insufficient to express the knowledge or data entities used in typical clinical guidelines. The 3rd Framework project GALEN developed a compositional and generative approach to modelling medical concepts which provides a consistent means to create clinical termsets of any required scope and complexity ('acute chest pain with sudden onset precipitated by cold and exercise').[6] In *PRESTIGE*, this model (the GRAIL Model) will be utilised to create supplements to standard codesets in current use (ICD, Read, ICPC) for each specific guideline application. GRAIL is not itself a part of the *PRESTIGE* Conceptual Model, but a harmonization process is being undertaken to ensure that the system of concepts used within GRAIL to model clinical concepts is consistent with (or mappable to) the assumptions and vocabulary of the Conceptual Model.

Note that this Model will not be a simple list of permitted terms. It will be a dictionary and (simplified) grammar for constructing noun phrases including legal qualifier phrases,

together with a built-in classification of terms and term components which can be used to constrain the legal uses of a term within a knowledgebase or a patient record (and to guide the selection and browsing of terms by a patient record user interface or a knowledge authoring tool).

5. Discussion.

The models of the Protocol and the Patient also comprise specifications of the dynamic behaviour of various objects in the model. For example we describe formally how the dynamics of a **protocol use** object in the EPR [see figure] will involve creating objects (?) to monitor the conditions specified in the protocol to govern its stages of execution, and creating **act** objects where specific protocol recommendations are being generated. This dynamics modelling effectively provides the semantics of the protocol model: it specifies the operational meaning of a guideline in terms of inferences performed on the basis of a protocol and a patient record, and its consequences in terms of recommended additions to a Care Plan within the EPR.

Different parts of the model will be used to constrain the final implementation of various components of the *PRESTIGE* technology:

- * Both the Protocol Model and Terminology/Semantics Model will constrain the Guideline Authoring and Dissemination Tool (GAUDI) which is used to create and edit protocol knowledgebases.
- * The EPR model will determine an API interface specification for all *PRESTIGE* application EPR implementations, to ensure that they deliver the care plan recording capabilities necessary for guideline applications.
- * The Protocol Model (including dynamics) and EPR model will jointly directly determine the engineering implementation of the Protocol and Act Manager which performs the inferencing functions necessary to support protocol use.

See [3] and [4] for discussions of the requirements addressed by these model sets and the development methodologies for their utilisation in specific healthcare applications. The EON and INTERMED projects at the Section on Medical Informatics, Stanford University School of Medicine and associated American centres are pursuing similar objectives to *PRESTIGE* and have adopted a generic protocol modelling approach partly influenced by study the results of the DILEMMA project which preceded *PRESTIGE*. [5] Comparisons conducted by SMI and *PRESTIGE* at a recent joint workshop indicates that the models and their operational semantics are broadly convergent; a full cross-mapping to demonstrate their formal equivalence is planned.

References

- [1] C Gordon and JP Christensen (Eds.), Health Telematics for Clinical Guidelines and Protocols. IOS Press, 1995.
- [2] C Gordon and M Veloso, The *PRESTIGE* Project: Implementing Guidelines in Healthcare. Medical Informatics Europe '96, IOS Press 1996, 887-891.
- [3] C Gordon, SI Herbert, P Johnson, Knowledge Representation and Clinical Practice Guidelines: the DILEMMA and *PRESTIGE* projects. Medical Informatics Europe '96, IOS Press 1996, 511-515.
- [4] C Gordon, P Johnson, C Waite, M Veloso. Algorithm and care pathway: clinical guidelines and healthcare processes. Proceeding of AIME97, Grenoble 1997.
- [5] M A Musen, S W Tu, A K Das, Y Shahar, A Component-Based Approach to Automation of Protocol-Directed Therapy. JAMIA, forthcoming.
- [6] A L Rector, J E Rogers, P Pole. The GALEN High Level Ontology. MIE 96, 174-178..