The NewGuide Project: guidelines, information sharing and learning from exceptions

P. Ciccarese[1], E. Caffi[2], L. Boiocchi[1], A. Halevy[3], S. Quaglini[1], A. Kumar[1], M. Stefanelli[1]

[1]Dipartimento di Informatica e Sistemistica, University of Pavia, Italy; [2]Consorzio di Bioingegneria e Informatica medica, Pavia, Italy; [3]Ness-ISI Ltd., Beer Sheva, Israel
paolo @aim.unipv.it

Abstract. Among the well agreed-on benefits of a guideline computerisation, with respect to the traditional text format, there are the disambiguation, the possibility of looking at the guideline at different levels of detail and the possibility of generating patient-tailored suggestions. Nevertheless, the connection of guidelines with patient records is still a challenging problem, as well as their effective integration into the clinical workflow. In this paper, we describe the evolution of our environment for representing and running guidelines. The main new features concern the choice of a commercial product as the middle layer with the electronic patient record, the consequent possibility of gathering information from different legacy systems, and the extension of this "virtual medical record" to the storage of process data. This last feature allows managing exceptions, i.e. decisions that do not comply with guidelines.

1 Introduction

In the past years, we developed a tool for clinical practice guidelines (GLs) implementation [1]. Other research groups had put efforts in this field, and recently a comprehensive paper comparing different tools has been published [2]. Compared to the past research in this area, our focus is now shifting towards a different representation of data, information, and knowledge, according to their source and degree of generality/specificity. The patient goes through many healthcare organisations, and may be enrolled in more than one GL. Each organisation has its own legacy system, and the granularity of information is, in general, different from the one required by the GL. Moreover, while it is quite normal to store patient's clinical data (more or less detailed), the process data are rarely registered (workflow technology is not widely adopted in the healthcare setting). We argue that implementing GL without accessing process data allows only a partial exploitation of the GLs potential for care delivery improvement. As a matter of fact, to evaluate the physicians' behavior, and in particular their compliance to the GLs, it is necessary to know when, why and how a certain task has been performed: the rough datum in the Electronic Patient Record (EPR), representing the result of the task, often is not sufficient. It is clear that the two contexts, medical and organisational, are to be taken into account. The NewGuide Project allows the integration of these two fields, while maintaining their own specificity: it puts together the experience gained on Medical Knowledge Formalization [1, 2] and Workflow Management Systems (WfMS) [3]. The scenario that we devise is the following: the NewGuide inference engine suggests the action to be performed; if the physician accepts the suggestion, the action control is transferred to a WfMS; the latter, according to the action type, facilitates its execution, by looking for operators with the appropriate roles, by advising them through the most suitable communication systems, etc. It also stores the action execution details. The best performance is achieved when the WfMS is integrated with the EPR, because as soon as the WfMS "finds" the operator that will perform the action, it will provide him, when necessary, with the correct electronic form to be filled.

This paper focuses on the medical knowledge, but with particular attention to the influence that process data may have on a GL inference engine. Section 2 briefly illustrates the graphical formalism, Section 3 shows the solution adopted for the VMR, and Section 4 is about the VMR extension to process data.

2 The NewGuide representation formalism

Concerning the GL representation, our approach is still flow-chart like, with a strong connection to Petri Nets, that are a good theoretical basis for process management [4]. Given the health care profiles complexity we use a multi-level representation where a sublevel is the detail of a concept expressed in the higher level. The health care process is therefore composed by a sequence of blocks, on different levels, each of them addressing a medical task or a flow management function. For the specification of both the rules associated to arcs and the criteria for defining abstractions, we have implemented an object-oriented language that can also manage qualitative and temporal abstractions [6].

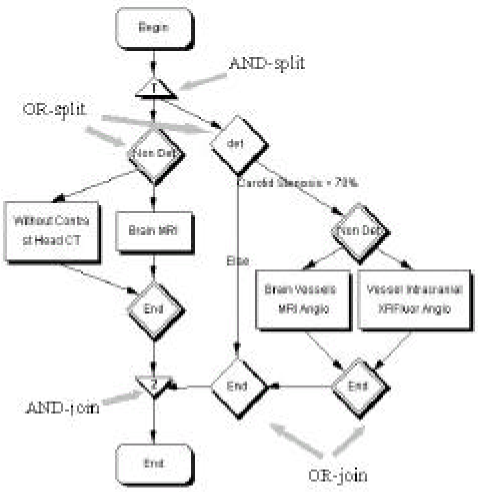

Fig. 1. A page of the stroke GL: a diagnostic strategy where choice is among different image-based examinations. Analogies with Petri Nets are shown.

As mentioned above, any schema built with New-Guide can be translated into a Petri Net (see the analogies in Fig.1). Maruster et al [7] showed that Petri Nets allow performing a “process mining”, thus discovering a workflow model starting from the workflow logs. Having a Petri Net-compliant GL representation, allows comparing the learned process with the theoretical one.

3. Computerised Guidelines and Legacy Systems

On top of the above illustrated formalism, we built an inference engine. We must consider that, usually, patients are treated in different settings, with different information systems. In general these are not shared between different health care professionals and it is difficult to retrieve information of different kind and nature. Moreover the same organisation may implement different GLs. Thus, the simple creation of a GL-oriented middle layer between the GL engine and the EPR is not sufficient. A more general level is needed. The International Medical Informatics community is tackling this problem since many years. The idea of a middle layer, called “virtual patient record” (VPR), or “virtual medical record” (VMR) has also been discussed by different authors [8]. The Decision Support Technical Committee of HL7 is actually working at its VMR specification. We are also working on a similar concept, and we believe that this middle layer must carry not only the clinical patient data representation but also all the information related to the decisional processes along the health care pathway.

The Medical Case Study- We considered the stroke patients admitted to the Stroke Unit of the "IRCCS C. Mondino" Hospital (from here on, SU). Two of the major risk factors for stroke are hypertension and hypercholesterolemia. Thus, it is not unusual that patients admitted to the SU were already enrolled in the outpatient departments devoted to those chronic pathological conditions. These two departments belong to another hospital, the "IRCCS Policlinico San Matteo" (from here on SM), with a different information system. SU also implements a GL for the management of the acute/subacute stroke phases, while SM implements a GL for hypertension treatement. Since stroke is a very severe condition, and patients (often unable to provide information) need to be treated in a very short time, it is essential to retrieve all the possible information from whatever data source. Thus, in the acute stroke phase, SU physicians will benefit from receiving the patient's data from the SM database. Moreover, in the post-acute phase, SM physicians will benefit from receiving the stroke history from the SU, in order to assess the best strategy for secondary prevention.

The Adopted Solution: *dbMotion*[1] - Since we do not want to modify the legacy systems, we need to collect the information in different formats and than re-organize it for a homogeneous management at the GL engine level. This target can be reached with a database connectivity product. We chose the ***dbMotion*** platform, a commercial technology for planning, establishing, operating and managing an Internet-based Virtual medical information Community. The ***dbMotion*** platform enables on-line collection of medical data components from decentralized databases and transferal of the data to authorized users according to parameters pre-defined by user profiles. Most of its infrastructure is transparent to the end-user, who retrieves the information on his workstation by means of an online viewer or, as in our case, of a data server. In fact, the system administrators are the ones who configure ***dbMotion*** according to the arising needs and requirements of the organization. To our purposes, ***dbMotion*** is the

[1] DbMotion is a product of Ness-ISI Ltd. Beer Sheva, Israel (www.dbMotion.com)

way to map all the different data sources into a unique structure. It is based on a VMR, called the International Clinical Information Schema (ICIS). It is object-oriented and HL7 compliant[5].

```
<ObservationObject
   PatientId="34533332"
   PatientSystemId="00534532"
   ProcedureId="C4589385"
   ObservationCode="8309-7"
   ObservationType="C4844893"
   ObservationClass= "BDYTMP.ATOM"
   ObsName="BODY TEMPERATURE"
   ObservationDate="12/02/2003"
   ObservationUnits="°C"
   ObservationValue="36"
>
```

Fig.2. The ***dbMotion*** object

NewGuide and *dbMotion* VMR- The ICIS provides the objects for the representation of patients, healthcare structures, observations, diagnoses, drug therapies. These objects are useful not only for defining the information entity but also for sharing it with the external world, through the use of well-accepted medical terminology systems. In fact the tool requires a preliminary mapping of the observations into LOINC terminology [9] and ICD9-CM disease classification, and then it guarantees a homogeneous management of the different EPR data. For example, Figure 2 shows the object representing the body temperature value, as it is returned by ***dbMotion***.

4 VMR Extensions to process data

To trace and support the effective usage of a GL we need some extensions to the VMR. They will allow storing information about "task substitution", "task abortion" and so on. Indeed these data objects are predisposed to contain a motivation provided by the user. These motivations concern four levels of compliance/non compliance:

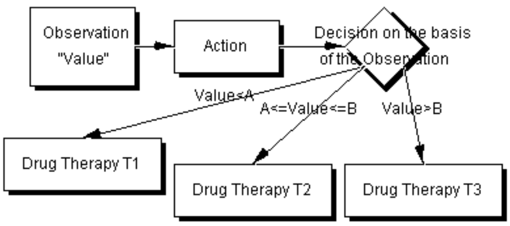

Fig.3. A GL flow that can be transformed as a consequence of a non compliance (see text)

Compliance (Level 0): the GL flow is completely preserved as well as the original intention or meaning; **Weak Non Compliance (Level 1)**: identifies a substitution of an action with another one that is similar in terms of medical goals or finding (for example substitution of an RMI with a CT or with an RX); the flow suggested by the GL is preserved; **Definite Non Compliance (Level 2)**: identifies a suspension, abortion or delay; the flow of actions suggested by the GL is modified in terms of medical intention and paths but no new actions are added; **Strong Non Compliance (Level 3)**: identifies the rejection of a GL suggestion or the insertion of new actions. This is both a flow and intention modification. In order to better understand why the "non compliance management" is necessary for the GL inference engine to run appropriately, we consider a simple flow, shown in Figure 3. We can have different exceptions: **Level 1**- the Observation as requested by the GL is not available, and the physician substitutes the Observation with a similar one that can provide the same information but, for example, with a different confidence (i.e.: RMI substituted by a CT scan). The process will continue but, at the rule-based decision block, the system will advice the user that the confidence of the Value is less than the one expected by the GL; **Level 2 or 3-** the Observation cannot be done (i.e.: lack of resources) or it is rejected (i.e.: the physician does not agree with the GL), the system will continue through the flow until the rule-based decision; at this point it will convert automatically the decision block into a "non rule-based" decision in which it is up to the user to choose the next task, if any: alternative options are to go directly to the end of the decisional block, or to leave the GL.

6 Conclusion

NewGuide is intended to support the physician in the whole patient management. It is particularly difficult to find a clear cut between "pure medical actions" and "organisational actions", i.e. those actions that could be managed by a workflow management system. We are trying to assess such a distinction, implementing a communication layer able to propagate the effect of the exceptions that can arise during the GL-based process.

Acknowledgments. This work has been partially funded by the FIRB project, Italian Ministry of University and Research. We thank Prof. E. Marchesi, Dr. G. Micieli and their staff for the contribution on the medical side.

References

1. Quaglini S, Stefanelli M, Cavallini A, Micieli G, Fassino C, Mossa C. Guideline-based Careflow Systems. Artificial Intelligence in Medicine 2000; 20(1) 5-22
2. Peleg M, Tu S, Ciccarese P, Kumar A, Quaglini S, Stefanelli M et al. Comparing models of decision and action for guideline-based decision support: a case-study approach. JAMIA 2003; vol. 1 ,n.10, 52-68
3. Panzarasa S, Maddè S, Quaglini S, Pistarini C, Stefanelli M. Evidence-based careflow management systems. Journal of Biomedical Informatics 2002; vol. 35, 123-139
4. van der Aalst W, van Hee K. Workflow Management. Models, Methods, and Systems. The MIT Press 2002
5. Jenders RA, Sujansky W, Broverman CA, Chadwick M. Towards improved knowledge sharing: assessment of the HL7 Reference Information Model to support medical logic module queries. Proc AMIA Annu Fall Symp 1997, 308-12
6. Bellazzi R, Larizza C, Lanzola G. An http-based server for temporal abstractions. IDAMAP '99 working notes ,pag. 52 - 62
7. Maruster L, van der Aalst W, Weijters A, van den Bosch A, Daelemans W. Automated Discovery of Workflow Models from Hospital Data. Proceedings of ECAI 2002; 32-36
8. Johnson P, Tu S, Musen MA, Purves I. A Virtual Medical record for Guideline-Based Decision Support. AMIA Annual Symposium, Washington, DC, . 2001
9. LOINC Users Guide release 1.0N, 02/04/ 2000.