

Development and Implementation of Clinical Guidelines - An Artificial Intelligence Perspective

Tiago Oliveira ¹

Department of Informatics, University of Minho, Braga Portugal

toliveira@di.uminho.pt

University of Minho, CCTC/DI, Campus of Gualtar, 4710 -057, Braga, Portugal

Paulo Novais

Department of Informatics, University of Minho, Braga Portugal

pjon@di.uminho.pt

University of Minho, CCTC/DI, Campus of Gualtar, 4710 -057, Braga, Portugal

José Neves

Department of Informatics, University of Minho, Braga Portugal

jneves@di.uminho.pt

University of Minho, CCTC/DI, Campus of Gualtar, 4710 -057, Braga, Portugal

Abstract: Clinical Practice Guidelines in paper format are still the preferred form of delivery of medical knowledge and recommendations to healthcare professionals. Their current support and development process have well identified limitations to which the healthcare community has been continuously searching solutions. Artificial Intelligence may create the conditions and provide the tools to address many, if not all, of these limitations.. This paper presents a comprehensive and up to date review of Computer-Interpretable Guideline approaches, namely Arden Syntax, GLIF, PROforma, Asbru, GLARE and SAGE. It also provides an assessment of how well these approaches respond to the challenges posed by paper-based guidelines and addresses topics of Artificial Intelligence that could provide a solution to the shortcomings of clinical guidelines. Among the topics addressed by this paper are Expert Systems, Case-Based Reasoning, Medical Ontologies and Reasoning under Uncertainty, with a special focus on methodologies for assessing Quality of Information when managing incomplete information. Finally, an analysis is made of the fundamental requirements of a guideline model and the importance that standard terminologies and models for clinical data have in the semantic and syntactic interoperability between a

¹ Corresponding author

guideline execution engine and the software tools used in clinical settings. It is also proposed a line of research that includes the development of an ontology for Clinical Practice Guidelines and a decision model for a guideline-based Expert System that manages non-compliance with clinical guidelines and uncertainty

Keywords: Computer-Interpretable Guidelines, Ontologies, Decision Support, Quality of Information

1. Introduction

There is an increasing pressure in healthcare professionals to standardize their clinical practice in order to prevent undesired variations. Clinical Practice Guidelines (CPGs) are developed in order to achieve this purpose. In recent years there has been an explosion of interest in CPGs, with initiatives to stimulate guideline development promoted by many countries and healthcare institutions. In fact, CPGs are, currently, the best way to convey information to healthcare professionals, to ensure that their clinical practice follows the rules of medical procedures. This is a very important matter, if one takes into account the consequences that may arise from a poorly conducted clinical process. The prevalence of medical errors is significant in hospitals across the world (Brennan 2000; Kalra 2004). Putting aside the human cost, which is immeasurable, the economic cost from lawsuits and other legal issues resulting from medical error has a deep impact in the budget of healthcare institutions. However, an overzealous practice, like defensive medicine, may have equally undesired consequences (Chawla and Gunderman 2008). The prescription of exams and treatments without scientific proof or basis also has a great economic impact and may seriously undermine the confidence that patients have in their physicians. This has consequences in the mental health of patients as well. The primary objective of CPGs is to provide a scientific support to clinical procedures, thus mitigating the occurrence of these situations.

However, healthcare professionals still show some resistance towards complying with CPGs. The arguments used to justify this behavior are that guidelines stifle change and innovation, and restrain clinical practice, preventing healthcare professionals from adapting their practice to their social, economic and cultural contexts (Thomson et al. 1995). Guidelines evolved in order to address some of this criticism, through the development of mechanisms to smooth updating processes and to accommodate justified variations in clinical practice. Currently, we live in the age of information and, once again, CPGs should evolve to keep up with the rapid growth of scientific knowledge. Research in the field of Computer-Interpretable Guidelines (CIGs) is booming, due to the need to deliver

information to healthcare professionals in a faster way and to support them in decision making.

This paper starts by providing some background information on CPGs and how they are developed. Then some Artificial Intelligence (AI) techniques will be addressed with the objective of determining how AI can improve the current state of the art as well as the development and the execution of CPGs. The final sections of the paper provide a general presentation of the research line that is being followed and how it integrates the conclusions extracted from the analysis made. It goes without saying that a paper format cannot be compared to a computerized guideline, as the first cannot be processed electronically. The perspective this work intends to show is how a digital format can be more advantageous and provide a new set of tools to facilitate the work of healthcare professionals.

2. Clinical Practice Guidelines

2.1. What are Clinical Practice Guidelines?

CPGs are systematically developed statements to assist healthcare professionals and patients about appropriate healthcare in specific clinical circumstances (Miller and Kearney 2004). This is the most widely accepted definition of clinical guideline, provided by the Institute of Medicine, of the United States (US). There are other terms used as synonyms of CPGs such as protocols, practice policies, clinical policies, practice parameters and clinical pathways. Usually, the name given to these documents is a matter of personal preference rather than a reference to a standard nomenclature and it can change across healthcare institutions and countries. Despite these differing nomenclatures, there are common objectives associated (Miller and Kearney 2004) with all of them, such as:

- Help healthcare professionals and patients in decisions about clinical procedures;
- Describe appropriate care based on scientific evidence;
- Act as the focus for quality assessment and activity improvement, including audits.

CPGs are decision tools devised to shorten the distance between real clinical practice and optimal clinical practice (Mead 2000). The potential benefits from

the implementation of CPGs include the reduction of morbidity and mortality, efficiency improvement and cost containment. They also provide their users with a reference by which they guide their clinical practice, and measurable criteria to assess their performance. The evidence contained in CPGs is used, at the same time, to inform healthcare professionals of the latest developments in scientific knowledge and to justify their decisions during the clinical process (Thomson 2000).

The format of these documents is not standardized and shows variations according to the organization producing the guideline and the clinical area it addresses. Since the middle of 1990s, many worldwide organizations started evidence-based CPG development programs, namely the Scottish Intercollegiate Guidelines Network (SIGN), the New Zealand Guidelines Group (NZGG), the Guidelines Advisory Committee (GAC) in Canada, and the National Federation of Cancer Centers (FNCLCC) in France, among others (Rosenbrand et al. 2008). These organizations joined others that paved the way for guideline development like the Institute of Medicine and the Dutch Institute for Healthcare Improvement. In 2002, an international effort towards the dissemination of CPGs culminated in the creation of the Guidelines International Network (G-I-N). (Ollenschläger et al. 2004). Currently, this global network comprises 92 organizations and 127 individual members, representing 48 countries, putting forth efforts in order to standardize guideline development and implementation. In recent years, some online guideline repositories started to appear, among which should be highlighted the National Guideline Clearinghouse (NGC) of the US. NGC² is a public resource for evidence-based CPGs and gathers guidelines from various organizations under different labels that represent their category and medical specialty.

2.2. Development of Clinical Practice Guidelines

Each organization follows its own guideline development process. However, the different development methodologies have common phases and follow similar principles.

Initially, guidelines were only based on the consensus of groups of experts, but with the growth of evidence-based clinical practice, other techniques were

² <http://guideline.gov/index.aspx>

included in guideline development. The Delphi and nominal group techniques are some of the methodologies that were later included in the development process and are still used today (Hutchings et al. 2006). Currently, guideline development is more focused on an extensive research of the literature and thorough analysis of empirical evidence.

The process usually starts with the choice of the guideline topic or subject, based on the problems that motivate the development (Rosenbrand et al. 2008; Rosenfeld and Shiffman 2006). CPGs can be developed to a wide range of subjects and medical areas, including health conditions bound to diseases and economical costs. To choose the topic, it is necessary to do a preliminary check of the available evidence in order to ascertain the validity of the theme.

The composition of the work group is the following step (Rosenbrand et al. 2008; Rosenfeld and Shiffman 2006). The efficiency of the guideline highly depends of the nature of the group producing it. The work group must be multidisciplinary, in a way that includes participants from all the areas affected by the topic of the guideline. Once the group is gathered, the analysis of the underlying problem, to which the guideline must provide a solution, starts. The work group must search for other guidelines concerning the topic, whose existence does not invalidate the creation of a new one as the existing ones may be outdated. The result from the analysis of the problem should be a set of key-questions that clearly identify the population being studied (the group of individuals who will be the target of the diagnosis or intervention), the type of control used and the efficiency measures that will be used to evaluate the interventions.

The objective of the literature research is to find the best available evidence, capable of answering the key-questions formulated in the previous step. The development group has to define some search constraints (e.g., to privilege a published work over an unpublished one) in order to assure the quality of the evidence. Once all the information sources are gathered, the work group does a critical appraisal of the evidence, based on the methodologies used to do the studies that generated them. The reviews are summarized in evidence tables, with a grade being given to the medical trials that were selected (Rosenbrand et al. 2008; Rosenfeld and Shiffman 2006). Healthcare institutions that produce guidelines do not have a common grading system, which is inconvenient when

one has to compare the evidence of similar guidelines. The Grading of Recommendations Assessment, Development and Evaluation (GRADE) (Kavanagh 2009) workgroup was created with the objective of developing an approach to evidence grading that can be used by different organizations. The GRADE system has been adopted by an increasing number of organizations and it is in continuous development.

After the evidence grading, the workgroup must elaborate a sketch of the guideline and submit it to external revision. Usually, this revision is performed by independent entities in conferences or healthcare-related gatherings. This is an iterative process in which the guideline is altered according to the reviews and then proposed for another external revision, until it reaches a stable version. Then the guideline is published and disseminated through conferences and newsletters to healthcare professionals.

2.3. Shortcomings of Clinical Practice Guidelines in the Present

In the final phase of the development process, the development group has to choose suitable means for disseminating guidelines (Thomson et al. 1995). The usual ones are newsletters to healthcare professionals, disclosure at medical conferences and through online PDF repositories of guidelines (Cheater and Closs 1997; Dennis et al. 2004). However, these means do not provide the desired coverage and sometimes fail in the delivery of knowledge to healthcare professionals. This is an important aspect because feedback from the medical community is the best mechanism through which guidelines are improved.

Guideline documents have a structure that makes them difficult to consult. Usually they are long texts and the clinical recommendations are contained in the body of that text. This aspect interferes with the retrieval of relevant information by healthcare professionals and makes the consultation for real time application rather complicated. Moreover, these long documents are difficult to update, which is a great drawback in the evolution of a guideline. They should accompany the development of clinical knowledge in a specific medical area (Rosenbrand et al. 2008).

Another issue is the ambiguity of the content of guidelines (Woolf et al. 1999). Ambiguity can be classified into syntactic, semantic and pragmatic (Codish and Shiffman 2005). Syntactic ambiguity occurs when the structure of a statement is

not clear, thus impeding its correct interpretation. Misplaced (or lack of) punctuation and wrongfully applied Boolean connectors are some of the causes of syntactic ambiguity. The classic definition understood generically by people as ambiguity fits the category of semantic ambiguity, characterized by situations in which terms can be interpreted in more than one way. Misuse of abbreviations, such as the case of the word “cold”, which in the context of a guideline can mean “common cold”, “cold sensation” or “Chronic Obstructive Lung Disease”, fall in the spectre of semantic ambiguity. As for pragmatic ambiguity, it happens when the recommendations of CPGs are not consistent or are conflicting with each other.

The vocabulary used in CPGs may also denote vagueness (Codish and Shiffman 2005). Sometimes the boundaries of a term are not completely understood by healthcare professionals. To show an example, temporal vagueness is frequent in guideline recommendations, with the use of terms such as “rare” or “common”. The poor specification of terms is also frequent, with terms like “moderate”, “elderly” and “adequate” being used without sufficient detail for clear interpretation. The texts often have occurrences of probabilistic terms to describe the frequency of events, namely “impossible”, “certain”, “unlikely” and “probable”, whose interpretation falls upon the subjective perception of the reader. The same situation occurs with some of the quantitative terms that are used.

Healthcare professionals often complain that, rather than offering support for clinical practice, CPGs restrain it, the argument being that they do not consider the social, cultural and economic conditions of the context in which they are applied (Woolf et al. 1999). Healthcare professionals may need to adapt their clinical practice according to the origin of their patients, but the steps for doing so are not described in CPGs. This lack of context-awareness is one of the major causes of noncompliance.

Currently, CPGs do not cope with preference-sensitive decisions, for instance, between scientifically valid treatments that may be applied to the same situation. In this case, there should be a group decision that takes into consideration the preferences and goals of the medical team responsible for the clinical case as well as those of the patient (Weidjen et al. 2011). What usually happens, in these cases,

is that the decision is made by one healthcare professional only, without consulting the other parts involved.

The level of uncertainty and incompleteness of the information upon which decisions are made, during the application of guidelines, is also a matter of concern (Logan and Scott 1996). A symptom is a somewhat uncertain indication of a health condition as it may or may not occur together with the disease. Thus, it is necessary a measure of the uncertainty associated to the observation of a symptom and the risk of the occurrence of a disease. During clinical encounters, healthcare professionals have to collect the values of relevant clinical parameters that build the patient's health state. The observations made by healthcare professionals in order to obtain these values have a subjective nature, mainly because a human being is doing them, thus the information they generate may be contradictory/inaccurate and sometimes the values of these parameters may not be obtainable due to the lack of technical means to do so. These cases of contradictory, inaccurate and missing information fall under the designation of incomplete information.

3. Artificial Intelligence and Clinical Guidelines

AI is a field of study that aims to explain and emulate intelligent behaviour in computational processes (Schalkoff, 1990). It is the branch of computer science that is concerned with the automation of intelligence. The ability to make machines think like human beings creates new possibilities in many areas. Research in AI helped the development of new technologies that nowadays are the basis of many big systems. These technologies are primarily used to automate tasks and improve knowledge-based processes, such as decision making.

The application of AI in medicine can be traced back to the middle of the 1970s and early 1980s, and led to the appearance of a subarea in AI, called Artificial Intelligence in Medicine (AIM). Research in this new field was pioneered by research groups in the US. An early definition of AIM was provided by Shortliffe (1993), stating that the primary concern of this research area was the construction of AI programs that perform diagnosis and make therapy recommendations. This definition reflects the primary focus of AIM at that time, which was the understanding and automation of the clinical encounter. Nowadays, AIM is more focused on giving support to healthcare workers rather than trying to replace

them. As so the identification of the right areas of medicine in which this support can be given is the key aspect that dictates the acceptance of AI technologies by clinicians.

The variety of roles AI programs may play in medicine is very wide. The use of medical knowledge is one of such roles, namely the support to human cognition that can be implemented, for instance, as reminder systems that alert healthcare professionals of clinical events or contradictions in treatment plans. AI programs can also be used to create new knowledge by discovering new phenomena through data analysis, pattern discovery and associations. Machine learning is the subfield of AI that deals with the generation of new knowledge and includes different techniques to produce systems capable of providing a description of clinical features. Case-based Reasoning (CBR) is one of such techniques. Based on past clinical cases, CBR is able to generate recommendations to new ones. The form in which these recommendations are provided varies, but rules and decision trees are among the most commonly used. An example of this type of system is KARDIO, for interpreting ECGs (Bratko et al. 1989). Another application of machine learning in medicine is the use of data-mining in the construction of pathophysiological models and drug discovery. AI systems containing medical knowledge, usually about a specific domain, are capable of reasoning and reaching conclusions based on data. The array of functions AI programs can perform includes: alerts and reminders, diagnostic support, agents for information retrieval and image recognition/interpretation. DXplain (Barnett et al. 1992) and HELP (Gardner et al. 1999) are examples of these knowledge-based decision support systems and are among the first ones to be developed.

In the remainder of this section, the focus will be placed on some topics and technologies of AI that may provide effective responses to the shortcomings of CPGs and help the development of clinical practice.

3.1. Group Decision Making

Group decision is a common phenomenon in human decision making activities. It is an arduous task because it implies the aggregation of individual alternatives to yield a decision that is acceptable to the group as a whole (John et al. 2008). The group explores a number of alternative solutions, answering *what-if* questions and

the participants may have different roles in the decision process, according to pre-established criteria by the organization.

During the application of CPGs, there are moments when this type of decision is required. The selection among scientifically valid options during the clinical process must be done based on the opinions of the parts involved (healthcare professionals and patients). Technology-assisted decision making may help the generation of ideas and actions, the choice of alternatives and the negotiation of solutions. The existence of CIG models and a tool for execution of CPGs enables the implementation of automated group decision making.

The work of Karacapilidis and Pappis (1997) summarizes some of the aspects that must be taken into account when developing a framework for group decision. The first one is the spatial distance between decision makers and the electronic communication facilities that enable them to communicate with each other. In the clinical setting, it is not uncommon for a clinical case to be treated by a medical team whose members are from different healthcare institutions, so the development of a virtual environment that enables the communication between them may be an advantage to the discussion of guideline recommendations, as shown in previous works with successful knowledge exchanges (Anogianakis et al. 1998; Househ et al. 2011). The type of environment influences the goals of decision makers. The goals are different in an environment where the group wants to solve a common problem cooperatively from another in which bargaining takes place. Typically, in a clinical environment both situations can occur, a medical team may be discussing the diagnosis of a patient and their members may have different opinions based on different evidence. The implementation of techniques in a virtual clinical decision environment to extract information about the actors (e.g., stress level) (Novais et al. 2012), would assist the definition of the type of interaction between the group members and consequently the selection of a suitable decision model. However, the development of mechanisms that enable them to express their preferences is necessary. The type of control over the decision process is also important. The group members may follow a democratic process in order to reach a solution (e.g., voting) or they may follow a hierarchical model in which the system is supported by a mediator, capable or not of imposing decisions to the other members.

A group decision environment with a decision model will help healthcare professionals and patients clarify their position in the decision making process and assure not only that their perspectives and preferences are heard, but also that they conform with the recommendations of CPGs,

3.2. Expert Systems

An Expert System (ES) is a computer program capable of performing at the level of a human expert, or above it, in some knowledge domain (Nikolopoulos 1997). This type of systems uses knowledge and inference procedures to solve difficult problems. They have to mimic the adaptation capabilities of human beings in order to find solutions to new problems (Jackson 1990). In this sense, there is four fundamental aspects of the construction of ESs (Nikolopoulos 1997): the knowledge acquisition module, the knowledge base, the inference engine and the interface. The knowledge representation in an ES applies concepts of logics to create structured formalisms, inference rules as well as ontologies to define the context of the domain. The knowledge itself may be introduced in the system by human experts as rules, obtained from past experience through learning algorithms or both.

Expert systems not only apply knowledge to situations but also generate new knowledge for new situations. The advantage of these systems, namely in healthcare, is that they are able to justify the decisions they make and provide confidence measures in their decisions. One of the problems healthcare professionals are faced with is the efficient use of all the information concerning clinical cases that they have. ESs provide means to treat large amounts of information and extract knowledge to be used in the future.

The applicability of ESs in healthcare has been proven through cases such as those of MYCIN (Melle 1978), for the diagnosis of infectious diseases; and NED (Zhou et al. 2002), which is used for the detection of lung cancer cells in the images of the specimens from needle biopsies. The usefulness of ESs in healthcare is evidenced in the work of Seto et al. (2012) that comprises the development of a rule-based ES for the monitoring of heart failure.

The development of an ES based on CPGs will enable the implementation of guideline acquisition tools based on domain ontologies that represent the different aspects of the clinical process. Such a system will also enable referencing the

evidence and clinical trials that endorse a clinical recommendation, in order to provide healthcare professionals the support they need to justify their actions. Inference rules are a fundamental feature of the system. Initially, they must be based on the available evidence, researched during the guideline development process by human experts, but afterwards, learning techniques, such as CBR, may be used to reinforce the rules of the system or offer alternatives to the recommendations of guidelines. Such an ES will also enable healthcare professionals to give their feedback of guideline recommendations, according to the outcomes of their application, producing data that may be used to improve them. The issue of guideline contextualization may also be addressed by the ES through the use of information retrieval techniques to search for news and articles that fit the scope of the health conditions addressed by a guideline. Such a feature may be useful when dealing, for instance, with flu outbreaks because the healthcare professional may consider relevant the information about new virus strains that are currently active and characteristics of the population that make them particularly vulnerable to those strains, adapting his clinical practice accordingly.

Currently, web applications are growing fast. They present some advantages to their desktop counterparts that make them the ideal support for ESs. They require no installation and updating and are accessible from anywhere on the internet. The data is stored remotely and they do not require high specs from the devices in which they run. This portability makes them accessible to low spec PCs, smartphones and tablets. The coming of age of cloud computing and mobile cloud computing will have a positive impact in the way e-health services are made available to healthcare professionals, mainly due to the pervasive access to information granted by these technologies (Dinh et al., 2011). Moreover, a cloud-based health information system eases the integration of different services from different service providers through the internet to meet user demands. A web-based ES for the application of CPGs would allow healthcare professionals to access the information they need when they are in contact with their patients, filling in any knowledge gaps they might have. It would also provide decision support during the clinical process and solve the problem of guideline delivery to healthcare professionals.

Although there is a widespread research in the field of healthcare ESs, their application in real life is not so widespread. Among the reasons for this situation is the fact that, in many of them, developers do not consider the cognitive necessities of healthcare professionals when designing their interfaces (Johnson and Turley 2006). Intelligent interfaces reflect users' goals, tasks and processes in order to make human-machine interaction a collaborative experience. As such, they provide an abstraction level of the processes that occur in the internal structure of the system that resembles the cognitive process of the users. This is beneficial to the implementation of CPGs, it enables the development of user-friendly tools consisting of graphical interfaces that support primitives for drawing the control information within the guideline, windows for acquiring the internal properties of the objects, facilities for browsing CPGs and an environment for consistency checking of clinical recommendations. To achieve these purposes, the interface needs specific data about the clinical domain that is being addressed as well as models for the representation of the knowledge of CPGs, their rules and processes.

3.3. Case-Based Reasoning

CBR is an AI approach that makes use of past experience to solve current problems (Aamodt and Plaza 1994). The applicability of CBR in health sciences is vast, given the similarities this research method has with the cognitive process of healthcare professionals: it is a natural process for them (Bichindaritz and Marling 2006). Case histories are the main training tool for clinicians and the medical literature is filled with accounts of treatments of individual patients. Moreover, some diseases still remain a mystery to the medical community, which impedes the definition of generic models to manage them. The approach to these clinical cases requires background knowledge recorded in practice cases. These background cases complement guidelines and help to interpret them. The human body is a complex biological system that is difficult to describe and even in well-known health conditions (e.g., hypertension and heart disease) several diagnoses interact to produce a given set of symptoms.

Typically, a CBR process is composed of four sequential phases: retrieve, reuse, revise and retain (Aamodt and Plaza 1994). The first phase consists in retrieving one or more previously experienced cases that are relevant. The

relevance of the cases corresponds to a similarity measure, (e.g., the difference of the sums of the different attributes that build the case). During the reuse phase, the solutions of the retrieved cases are mapped to the new case, which may involve adapting the solution in order for it to fit specific requirements of the new problem because it is unlikely that an exact match of the new case exists in the case memory. In the third phase, revise, the best matching solution is tested in order to predict the results of its application. If the result does not meet the expectations, the action taken is revised. In the last phase, retain, the solution of the new case is stored in the case memory, contributing to its enrichment.

Among the applications of CBR systems in the health sciences domain, CASEY is one of the earliest (Koton 1988). This system diagnosed heart failure patients by comparing them to earlier patients whose diagnoses were known. CASEY also integrated an earlier model-based system and pioneered the combined use of CBR and another reasoning methodology. PROTOS (Bareiss 1989) is another early CBR system that assigned patients to pre-defined diagnoses based on past cases (Bichindarits and Marling 2006). Since the debut of these systems, CBR has been used for other tasks, such as nursing diagnosis (e.g., FLORENCE system (Bradburn and Zeleznikow 1994)), radiation therapy design (e.g., ROENTGEN system (Berger 1994)) and diagnosis of degenerative brain diseases through image segmentation of CT and MR brain images (e.g., HPISIS system (Perner 1999)), to name a few.

CBR may be used to manage non-compliance with CPGs. When executing a clinical guideline in an ES, the healthcare professional may have to face a situation that was not predicted by the guideline or in which his professional opinion is different from the recommendations provided by it. Moreover, the unavailability of relevant patient data or resources, and the existence of data that is outside the range foreseen by the guideline may also require a deviation from the protocol by healthcare professionals. When faced with these situations, the ES may allow the healthcare professional to change the guideline in order to fit the current case. With the help of CBR, the system may construct a case memory of these deviations where the description of the cases (the pair attribute/value of the clinical parameters) and their solution (the alteration made to the guideline) are stored for later retrieval to solve similar cases. This way, the system could grasp the constraints (social, economic and cultural) of the medical practice of

physicians and provide useful feedback of the applicability of a certain guideline. An elevated number of cases in the memory case for a specific guideline are an indicator that a step of the guideline or the guideline itself is no longer fit for medical application.

3.4. Medical Ontologies

In the context of AI, an ontology is a formal representation of knowledge as sets of concepts and the relations between them within a domain (Gruber 1993). An ontology defines a vocabulary that contains all the concepts that may be used to model the domain and how they relate to each other. This conceptualization is achieved through the definition of classes and subclasses of individuals along with the properties of the individuals in a class.

Ontologies have a key role in the Semantic Web (Berners-Lee et al 2001), since they structure underlying data for the purpose of comprehensive and transportable knowledge and machine understanding. Besides allowing machines to read and interpret information, ontologies present other advantages to knowledge engineering such as automated validation and consistency checking.

In a complex domain such as the clinical one, ontologies provide significant advantages in the formalization of CPGs. The vagueness and ambiguity that, sometimes, is present in guidelines can be removed through the usage of controlled vocabularies, thus eliminating fuzzy relations between the concepts of the domain. It would also allow the extraction of rich patterns, that would go unnoticed otherwise, and the construction of inference mechanisms in the domain. The guideline ontology can be shared in ontology repositories for widespread use and dissemination. Currently there is a growing interest of clinical guideline researchers in ontology-driven execution of CPGs (Isern et al. 2012).

The Unified Medical Language System (UMLS) (Bodenreider 2004) reflects the efforts of the US National Library of Medicine to remove ambiguity and vagueness from the clinical setting. It is an ontology that aggregates terms used to describe the same concept from existent knowledge sources (e.g., SNOMED CT, LOINC, ICD-10, MeSH) under the same identifier. The UMLS has three main components: the Metathesaurus, the Semantic Network and the SPECIALIST Lexicon. The integration of this ontology in a CPG ontology would effectively improve the understanding of clinical recommendations and provide an easy

access to semantic networks that could provide precise definitions of medical terms.

3.5. Reasoning under Uncertainty

AI provides some techniques that deal with uncertainty and incomplete information in decision making scenarios. They can be classified in qualitative and symbolic methods (Clark 1990). The advantage of symbolic methods is that they bring some common sense validity that can be found in approaches such as non-monotonic logics, default logic and defeasible reasoning. There are also other symbolic methods, often called *reason-based* (Fox et al. 2001), that use informal endorsements for multiple options and formalizations of everyday strategies for reasoning about competing beliefs, *argumentation* being one of these techniques. However, the health sciences are, currently, more interested in the numeric methods such as Bayesian Networks, Dempster-Shafer Theory or Fuzzy Logic (Clark, 1990).

Bayesian Networks were derived from probability theory and appeared for the first time in the middle of the 1980s (Pearl 1986). It is a statistical model defined by two components: a qualitative component and a quantitative component (Clark 1990). The qualitative component is an acyclic orientated graph whose nodes represent a random variable that may assume any value from a given set and is associated with a probability distribution. The existence of an arch between two variables means that they are statistically dependent. The quantitative component is a conditional probability distribution. The essence of this approach is the representation of hypotheses and relations in the domain under consideration. In the medical domain, the relation of causality between clinical parameters and diseases may be effectively represented through Bayesian Networks and it is possible to obtain these relations from CPGs (van Gerven et al. 2008). Moreover, the prior probabilities for the different variables and the conditional probabilities may be gathered from the empirical evidence displayed in the guideline. Thus, the combined use of Bayesian Networks and CPGs adds value to the clinical process and provides quantitative measures that enable healthcare professionals to assert the solidity of their decisions. The work of Lucas (2004) is heavily focused on the combined use of clinical guidelines and Bayesian networks for clinical decision support systems.

The Dempster-Shafer Theory of evidence was initially developed by Dempster in 1967 (Dempster 1967) and later extended by Shafer in 1976 (Shenoy and Shafer 1986). It relies in degrees of belief to represent uncertainty. This approach allows the assignment of degrees of belief to sets of hypotheses (e.g., {gastric cancer, gastric ulcer}, i.e. gastric cancer is caused by a gastric ulcer) instead of individual clinical parameters, like Bayesian networks. For this reason, it is considered that Dempster-Shafer Theory is better able to represent the process of narrowing hypotheses with the accumulation of evidence. This process is claimed to mimic diagnostic reasoning. Since Dempster-Shafer Theory identifies a set of solutions that reflect any other options that are not discretized, it can deal with ignorance and non-exhaustiveness (not pointing out all the existing solutions). However, it receives some criticism concerning the computational complexity that generates for large sets of hypotheses (Clark 1990). Despite these shortcomings, Dempster-Shaffer Theory has been used efficiently for the representation of medical knowledge and uncertainty in some critical areas (Straszecka 2004).

The Fuzzy Sets approach was initially developed with the objective of quantifying imprecise classes in natural language (Zadeh 1975). It is most useful when sets are defined by vague concepts and variables are continuous (e.g., height, warmth, age). Natural language is full of examples of fuzzy classifiers, like predicates (e.g., small, large, young), quantifiers (e.g., most, many, few), probabilities (likely, unlikely) or truth values (e.g., very true, quite true, mostly true). The quantification in this method is provided by membership functions that attribute a value in the interval $[0,1]$ to the relevant elements. Fuzzy Logic was derived from fuzzy sets and is based on the notion of truth degree of a preposition. It defines operators that express the disjunction and conjunction of prepositions, independently of their meaning. Just as it is difficult to estimate the prior probabilities of a Bayesian Network, the production of membership functions is complex (Clark 1990). Many disciplines of medicine already use Fuzzy Sets in ESs, for tasks such as diagnostic and imaging analysis (Abbod et al. 2001). Fuzzy Sets are being researched for the representation of operational guideline knowledge and the definition of threshold values for clinical parameters. In fact, currently, Fuzzy Logic is being integrated with the Arden Syntax

(Vetterlein 2010) guideline model in order to produce a continuously graded applicability of statements.

All these approaches deal with uncertainty from a perspective of causality and interdependence, but do not address the aspects of incomplete information and the different forms it can assume. Further in this paper we will present a methodology called Quality of Information (QoI) (Neves et al. 2012) that provides ways of dealing with this information and making it useful to the decision making process.

4. Computer-Interpretable Guidelines

4.1. Living Guidelines

As a response to the challenges presented by CPGs, the concept of Computer-Interpretable Guideline emerged (De Clercq et al. 2008). CIGs are representations of CPGs in a digital format, suitable for being interpreted by machines. A digital format of CPGs may be a game changer in all the aspects that revolve around them, namely development, dissemination, implementation and execution.

There is a set of features that guideline researchers would like to see guidelines acquire (Rosenbrand et al. 2008). Features such as modularity, dynamism and interactivity are gathered under the concept of *living guidelines* (Seyfang et al. 2007), which basically is translated into guidelines that are easy to update and modify and have an active role in providing knowledge to healthcare professionals. The objective of researchers is to change the static and passive nature of guidelines. CIGs are, currently, the best way to achieve this purpose.

The development of a standard model of CIGs may provide a deeper understanding of the clinical process and may have significant benefits. A depiction model for CPGs can be used to identify the different requirements that must be accomplished before making a decision, to establish decision criteria and thus helping healthcare professionals in this critical moment of the clinical process (Elkin et al. 2000). Having a model also enables the definition of methods to verify the semantic and syntactic structures of guidelines, providing a way to distinguish a well formed guideline from a poorly made one (Elkin et al. 2000). If the model enables the definition of modular components, like for instance clinical

procedures that are common to different guidelines, it may be possible to reuse this knowledge when building a new digital guideline (Elkin et al. 2000).

The creation and use of CIGs offer a better description and recording of patient states and may provide selective access to background knowledge to be used in specific circumstances. The use of automatic reminders according to the recommendations of guidelines may also be implemented (Fox et al. 2008).

4.2. The Document-centric and Model-centric Approaches

Decision support systems based on CPGs may support healthcare professionals in following the best clinical practice in a consistent way. Formalization of guidelines in a guideline representation language may follow two different approaches (Sonnenberg and Hagerty 2006): document-centric and model-centric.

The document-centric approach (Kaiser and Miksch 2009) consists in using mark-up tools on the original guideline documents. The original document is either marked up or annotated to produce a more structured format with defined semantic elements. Usually, this process is carried out in stages. First, the mark-up is used to identify elements in the text of the guideline. Then, using a specialized tool, a semantic tag is assigned to the elements and the connections between them are made. The advantage of this approach lies in enabling the encoding of CIGs without the need to have a profound knowledge of a specialized computer language. However, the current tools that perform this process are not perfected yet and it is difficult to construct long and complex guidelines using this method.

On the other hand, in the model-centric approach (De Clercq et al. 2004), a depiction model is formulated by domain experts and the relationship between the new model and the original document is indirect. The acquisition of guidelines in the model-centric approach is done directly by healthcare professionals into the new model. Through this process, it is possible to develop friendlier interfaces for healthcare professionals to encode their guidelines and, at the same time, they become more knowledgeable of the different steps that compose the clinical process.

GEM Cutter (Shiffman et al. 2000), Stepper (Ruzicka and Svatek 2004) and DELT/A (Votruba et al. 2004) are some of the most relevant mark-up-based tools that generate semi-formal models of marked guideline texts. GEM Cutter was one of the first tools to apply a document-centric approach and transform guideline

information to an *ad hoc* format, called GEM. Stepper is a tool that segments the initial text in multiple user definable steps coded in XML. DELT/A stands for Document Exploration and Linking Tool/Add-ons and, as its name indicates, it supports the translation of HTML documents into any XML language, among which is the Asbru guideline representation model. There are methodologies (e.g., LASSIE (Kaiser and Miksch 2009)) for document-centric approaches that use information extraction techniques that rely on databases of medical terminologies to acquire guidelines in a semi-automatic way, thus eliminating the requirement of having an healthcare professional manually tagging the terms in the original document.

Some applications for model-centric acquisition of CPGs will be presented when the different representation models for CPGs are addressed, further in this paper.

4.3. Aspects of CIG-based Systems

In the conception and development of CIG-based decision support systems, researchers identified four aspects that must be taken into consideration in the development process (De Clercq et al. 2004): guideline representation and modelling, guideline acquisition, guideline validation and testing, and guideline execution.

The model is the fundamental feature of a CIG-based decision support system (Peleg et al. 2003). It has to provide enough expressivity in order to accommodate every step of a guideline. Normally, the models created exclusively for guideline representation have a set of construction units (e.g., tasks or steps) that are used to build a guideline (De Clercq et al. 2004). These building blocks are given the designation of representation primitives (e.g., decisions, actions) and are used according to a Task Network Model (TNM). Some works consider the adaption of business process models, such as Petri Nets (Quaglini et al. 2000), to the modelling of CPGs. However, these approaches do not have enough expressivity due to them being developed to support business organizations and processes rather than medical organizations and processes. The possibility of using them in clinical settings is being actively studied in order to define higher abstraction layers, capable of expressing the different steps of the clinical process, on top of the basic model.

Whichever model is chosen, the degree of complexity the representation is able to accommodate is an important factor. Different models may differ in terms of the abstraction levels they allow, for instance, in the nesting of guidelines inside other guidelines. CPGs possess two different types of knowledge (Rosenbrand et al. 2008), the declarative (scientific knowledge about the domain) and procedural (the inference methods and the decision model), which have to be formalized through a language in the representation model. The language should provide an objective vocabulary, syntax and semantics, so that an inference engine can be developed. In a complete representation, there should also be triggering criteria, which include initial screening to assess if a patient should enter a protocol or not and connect the different elements of the guideline according to the output of decisions. Another indispensable feature is temporal patterns because guideline recommendations depend mostly of the state of the patient in a given moment. Knowing this, it is essential for a guideline model to provide mechanisms to define durations, repetitions and cycles of tasks.

An acquisition tool must be developed in order to help healthcare professionals structure the knowledge according to the guideline model that was defined (De Clercq et al. 2004). The tool must take into account the approach followed for guideline acquisition, if it is either document-centric or model-centric.

The precision, the syntactic correctness and the semantic coherence are extremely important in the acceptance of CIGs by healthcare professionals and in their integration in daily practice (De Clercq et al. 2004). As such, the inclusion of mechanisms for guideline validation and testing in the guideline execution engine is necessary. During the execution, the guideline execution engine should have access to a database containing the values for the clinical parameters that build the patient's state in order to apply CPGs in real time.

4.4. Current Approaches to Guideline Modelling and Execution

Currently there are few CIG systems available and they lack application in real clinical settings. This review addresses them by their depiction models and mentions the execution engines available for each one as well as the underlying platforms. The selection of the approaches was based on opinions collected from the literature that deemed them the most relevant. Table 1 shows a summary of the available software tools and models as along with their main features.

4.4.1. Arden Syntax

Arden Syntax (Hripcsack 1994) was developed in 1989 and is now a standard of Health Level 7 (HL7). The current version of Arden Syntax is Arden Syntax 2.8. The primary aim of this approach is the sharing of simple and independent guidelines as modules. Each clinical guideline is modelled as a Medical Logic Module (MLM), which comprises relevant knowledge for only one judgment. Initially, each MLM was an ACII file divided in three partitions: *maintenance*, *library* and *knowledge*.

The *maintenance* and *library* partitions possess administrative information about the guideline, namely authoring and version number. The constructs of the *maintenance* partition are *title*, *(file)name*, *author*, *version*, *institution*, *date of last modification* and *validation status*. The *validation status* contains information about the approval of the guideline in a local institution and it may have three possible values: *testing*, *research*, *production* and *expired*. The transition from *testing* to *production* means a shift of responsibility from the institution that developed the MLM to the local institution where the guideline will be applied. The *library* compartment contains constructs used for a detailed description of the guideline and among them the attribute *purpose* enables the expression of the clinical objective of the MLM.

The main constructs of the *knowledge* compartment are *data*, *evoke*, *logic* and *action*. The *data* construct is used to obtain the values of the concepts referred in the MLM from the information system of the healthcare institution. The *evoke* construct contains the events that trigger the execution of the MLM and these events are related with the clinical parameters in *data*. The decision criteria are expressed in the *logic* construct through *if-then-else* rules and sets of logical, mathematical and temporal operators. When a rule is assessed to the value *true*, a given procedure of the construct *action* is proposed. These procedures may include messages/alerts or the execution of other MLMs. This approach reveals great modularity and gives transparency to the decision making process, but given its simplicity, the ability to capture the full content of a clinical guideline is compromised. Arden Syntax is mainly used in alert and monitoring systems, like the ones provided by the Regenstrief Institute (Anand et al. 2004). Initially it was defined in Backus-Naur Form (BNF), a notation technique used to describe the

syntax of computation languages. Currently the development of Arden Syntax by HL7 is based on XML (Kim et al. 2008).

There is a myriad of tools to acquire and execute guidelines in Arden Syntax. We will highlight the Arden Syntax IDE (Samwald et al. 2012), which is a simple development environment that provides syntax highlighting and testing functionalities for MLMs. The Arden Syntax IDE contains a compiler that generates java classes from MLM code. These classes are then executed by an *Arden Syntax Rule Engine* that works together with another component, the *MLM manager*, which gives the rule engine the access to the available MLMs in the system. Arden Syntax is a highly portable format, conceived to be integrated in Clinical Management Systems (CMSs).

4.4.2. Guideline Interchange Format (GLIF)

The Guideline Interchange Format (GLIF) (Ohno-Machado et al. 1998) represents an effort of Intermed Collaboratory for the development of a sharable clinical guideline representation model. The first published version of GLIF was released in 1998 and its current version is GLIF3 (Boxwala et al. 2003). This approach was developed in order to reflect a flowchart of steps and consists of a set of classes that describe the fundamental characteristics of a guideline and constructs that contain the clinical parameters. Through this flowchart representation, GLIF3 provides a better understanding of the clinical process to healthcare professionals.

A guideline in GLIF3 is an object that contains different steps, namely: *decision steps*, *patient state steps*, *branch steps*, *synchronization steps* or *action steps*. This approach follows the Task Network Model (TNM), so that every moment of the clinical process is labelled as a step.

Decision steps model decision points in a guideline and direct the careflow from one to alternative steps. There are two subclasses in *decision*: *case step* and *choice step*. A *case step* contains a set of logical expressions that initially corresponded to an excerpt of Arden Syntax. Currently, GLIF3 uses an OCL (Object Constraint Language) expression language called GELLO (Sordo et al. 2003) that has more expressive power than the previous. As for *choice steps*, they contain only a set of options for the next step in the clinical process and the selection is made by an external agent (e.g., the user).

Patient state steps function as labels that have constructs used to describe the patient's health condition. These steps are used as data entry points in the system. When the state of the patient is updated, the guideline that possesses the corresponding patient state is executed.

The tasks of the clinical process are modelled in the *action steps* through three distinct constructs: *medical actions*, *activity oriented actions* (e.g., messaging, retrieving of patient data) and *control actions* (invocation of structures such as sub-guidelines).

At Columbia University, GLIF is being integrated with the Clinical Event Monitor and the Computerized Physician Order Entry systems to provide clinical decision support (Peleg and Wang 2006) for post-CABG (Coronary Artery Bypass Grafting) (Zheng et al. 2010). Encoded GLIF guidelines are also being used in Israeli clinics for the management of feet injuries of diabetics. The GLIF3 Guideline Execution Engine (GLEE) (Wang et al. 2004) is a tool for executing guidelines in this format. It defines three layers of abstraction: *data*, *business logic* and *user interface*. The *data* level contains the Electronic Medical Record (EMR) and a guideline repository. The execution engine is in the *business logic layer* and includes a server that interacts with the data layer and clients that interact with the users. Applications exchange data with the other two layers through the *user interface* layer. GLEE may be linked with a clinical event monitor, thus enabling event-driven execution of CPGs, responding to alterations in the state of the patient. This software tool also defines a set of methods to connect it to CMSs and uses representations like the Resource Description Framework (RDF) and HL7 as a general patient data model (Schadow et al. 2006) to support CPGs and encode medical data in order to share information across different institutions.

Table 1 Software tools for guideline development (adapted from Isern and Moreno (2008)).

Tool	CG Repository	CG Editor	CG representation language	CG basic elements	Run-time engine	Access to EMR	Access to CMS	Standards used
Arden Syntax IDE	Yes	Yes	Arden Syntax	MLMs	Rule-based	No	No	XML
GLEE	Yes	Yes	GLIF3	Decision, action, patient state, branch, synchronization steps	Event-based and rule-based	Yes	Yes	RDF,HL7
Arezzo	Yes	Yes	PROforma	Plan, action, decision, enquiry	Rule-based	Yes	Yes	No
DeGeL	Yes	Yes	Asbru	Preferences, intentions, conditions, effects, plan body	Rule-based	Yes	Yes	XML, ICD-9, SNOMED-CT, CPT, LOINC
GLARE	Yes	Yes	Graph-like	Query actions, work actions, decision actions, conclusions	Rule-based	Yes	Yes	XML, ICD-9
SAGE	Yes	Yes	SAGE model	Context, action, decision, routing nodes	Event-based	Yes	Yes	HL7, UMLS

4.4.3. PROforma

In 1998, the Advanced Computation Laboratory of Cancer Research of the United Kingdom initiated the development of the PROforma (Fox et al. 1998; Sutton and Fox 2003) depiction model. The objective of this model is the development of guidelines as flowcharts where the nodes are instances of pre-defined classes of tasks. The main classes are *plans*, *actions*, *decisions* and *enquiries*. Each class has a set of attributes that reflect its information needs. The syntax of PROforma was also initially defined in BNF in an ASCII file.

Every task of a guideline derives from a common task called *root task*. A *root task* contains several guidelines encoded as sets of tasks called *plans*. On the other hand, a *plan* contains any number of instances of atomic tasks, such as *actions*, *decisions* and *enquiries*. A *plan* also has constructs that enable the definition of clinical objectives (that reflect the objective of a guideline), abort or termination conditions and scheduling constraints on the atomic tasks. It is also possible to define temporal constraints on *plans*, such as cycles, durations and number of repetitions. One of the core features of PROforma is the possibility of nesting plans inside other plans.

An *action* in PROforma is a task whose execution has to be performed by an external agent. Typically, these tasks consist in sending messages and calling external programs or clinical procedures.

The *enquiry* task defines data entry points in the guideline as questions to the patients or internal procedures to retrieve the relevant information from the patient's EMR. This class contains *data definition* constructs that specify how a value for a clinical parameter must be stored (e.g., data type, unit).

Perhaps the most important class in PROforma is *decision*. Among all the CIG formalisms, PROforma was the first to offer a support to deal with uncertainty during the decision process. A *decision* contains constructs to express candidate solutions to the decision problem as well as logical expressions that endorse or refute each candidate. Each expression, in favour or against a candidate, is associated with positive signs (represented by a plus sign +) and negative signs (represented by a minus sign -). The weight of an argument in the overall score of its candidate depends of the number of positive and negative signs it has. This is a

symbolic method of endorsement that uses a mathematical function to convert the signs in numeric values for calculating the scores of each option. Then the options are presented by descending order of scores. According to the results of a *decision* task, the next task in the careflow is the one that has a construct called *trigger condition* matching the output of the *decision*. PROforma has been used in a few prototypes for clinical management, namely CAPSULE for general practice and Bloedlink for advice on laboratory tests and management of chronic diseases (dyspepsia, asthma and depression) (Fox and Thomson 1998).

Among the software tools for PROforma, Arezzo (Fox et al. 2006) is arguably the most disseminated. It has an architecture composed of three elements: a *composer*, a *tester* and a *performer*. The *composer* is responsible for providing an acquisition suite of guidelines in PROforma. The *tester* verifies the syntactic integrity of the PROforma guidelines before deployment by the *performer*, which is an inference engine. The *performer* can be linked to existing EMRs and CMSs to acquire data related to patients and also defines different states of guideline execution (e.g., *waiting for data*, *suspended*, *finished*).

4.4.4. Asbru

Asbru (Shahar et al. 1998) is the result of collaboration between Stanford and Vienna Technology Universities. This formalism presents a notion of *plan* similar to PROforma in the sense that it represents a collection of items. The knowledge required to perform a *plan* is defined by its *knowledge roles*, which include *preferences*, *intentions*, *conditions*, *effects* and *plan body*.

The content of a *plan body* is composed of other *plans* until they are no longer decomposable, reflecting a parent-child structure. The plans that cannot be decomposed are called *actions*. The functionalities of *plans* and *actions* are defined by the remaining *knowledge roles*. The *plan body* is the layout of a given *plan*.

The restrictions on the execution of a *plan*, in order to achieve a given objective are defined by *preferences*. The categories in preferences that define these restrictions are *select-method*, *resources* and *strategy*.

The objectives of *plan* are represented in the *intentions knowledge role*. The definition of *intentions* helps the selection of an adequate *plan* and is crucial in decision support. Intentions are defined as temporal constraints on the actions of

healthcare professionals. There are four types of intentions: *intermediate state*, *intermediate action*, *overall state pattern* and *overall patient pattern*. *Intermediate state* refers to patient states that must be maintained, reached or avoided (e.g., the control of levels of substances in the blood) during the execution of the *plan*. On the other hand, *intermediate actions* define the actions the healthcare professionals must perform during the *plan*. The *overall state pattern* is the state of the patient that must be verified at the end of the execution of the *plan* and the *overall action pattern* is the pattern of actions of the healthcare professional that should result from the *plan*.

There are different types of *conditions*, namely *filter-preconditions* and *setup conditions*, *suspend conditions*, and *abort conditions*, that are used to express the respective following situations: conditions that must hold for a *plan* to be considered applicable, conditions that determine the suspension of a *plan* and conditions that determine the abortion of a *plan*.

Effects describe the expected behaviour of the execution of a plan. It is composed of the following two constructs: *argument-dependency* and *plan-effect*. The first is used to describe the functional relationship between the *plan* arguments and the measurable parameters, describing how they influence each other. The second describes the relationship between the overall plan and its expected effect.

Asbru is heavily focused on temporal aspects of CPGs and that is evident in its temporal annotations, which specify four points in time for the execution of plans and verification of conditions, with the particularity of allowing the expression of uncertainty in starting time, ending time and duration of a time interval. The temporal annotations of Asbru are *earliest starting shift* (ESS), *latest starting shift* (LSS), *earliest finishing shift* (EFS) and *latest finishing shift* (LFS). It is also possible to specify two types of durations: minimum duration (MinDu) and maximum duration (MaxDu).

This model has been used in the Asgaard project in the development of prototype applications for the management of diabetes, jaundice and breast cancer (Zheng et al. 2010).

Acquisition and execution of Asbru guidelines is possible through DeGeL (Shahar et al. 2004), a tool in development at Ben Gurion University, in Israel, and is a web-based architecture that facilitates the conversion of textual guidelines

to Asbru guidelines. This distributed architecture has some key components responsible for the creation of new guidelines, guideline retrieval from an XML repository as well as testing and execution of guidelines. The execution module is called *Spock* and it incorporates an inference engine that can retrieve data from the patient's EMR. It is a modular client-server application that consists of a set of classes to store guidelines, a parser to interpret their content and a synchronizer that establishes the communication with external systems. DeGeL also has a vocabulary server for supporting guideline specification and establishing mappings between the standardized terms and each clinical database vocabulary. The set of standard terminologies that is used includes International Classification of Diseases (ICD-9), Standard Nomenclature of Medicine (SNOMED-CT), Current Procedural Terminology (CPT) and Logical Observation Identifiers, Names, and Codes (LOINC).

4.4.5. GuideLine Acquisition, Representation and Execution (GLARE)

The GuideLine Acquisition, Representation and Execution (GLARE) (Bottrighi et al. 2006) project includes a guideline depiction model and a system to acquire and execute CPGs. It was developed by the Computer Science Department of the University of Piemonte Orientale, Alessandria, Italy.

The depiction model does not use a standard representation. Instead, it defines a proprietary graph-based structure for displaying CPGs, where a clinical action is represented by a node. It is possible to define atomic actions that represent simple tasks like *queries* to obtain external information, *work actions* that represent medical procedures, *decision actions* with sets of conditions to select alternatives and *conclusions* that describe the output of a *decision*. *Decision actions* are specific types of actions that contain the criteria used to select from alternative paths from a guideline. These criteria are represented as sets of triplets in the form <diagnosis, parameter, score> and, in turn, a parameter is another triplet <data, attribute, value>. It is also possible, in GLARE, to define composite actions, which are collections of atomic actions or other composite actions. GLARE was designed to cope with different types of temporal constraints and implements specialized temporal reasoning algorithms.

The GLARE execution engine (Bottrighi et al. 2006) distinguishes between the acquisition phase and the execution phase of guidelines. Similarly to GLEE,

GLARE defines three architecture levels, namely *System*, *XML* and *DBMS*. The *System* level encompasses the acquisition and execution modules. The *XML* level is responsible for the data exchanges between the *System* level and the *DBMS* level. The *DBMS* is the bottom level, responsible for establishing a physical connection between the top levels and the databases where the information for creating and executing guidelines is stored. This information includes open instances of guidelines, a repository of guidelines and medical records of patients. GLARE uses ICD-9 as a terminology standard.

4.4.6. Standards-Based Sharable Active Guideline Environment (SAGE)

The Standards-Based Sharable Active Guideline Environment (SAGE) (Ram et al. 2004; Tu et al. 2007) project is a collaboration of six research groups (IDX Systems, University of Nebraska Medical Center, Intermountain Health Care, Apelon, Inc., Stanford Medical Informatics and the Mayo Clinic). SAGE includes a guideline depiction model and a guideline authoring and execution environment and is perceived as an evolution of GLIF3 and EON. Its objective is to establish an infrastructure to enable sharing guidelines in heterogeneous clinical information systems. SAGE is involved with standards organizations to bridge the gap between guideline logic and real life implementations.

In this depiction model a guideline is a *recommendation set*, which is composed as a graph of nodes. These nodes can be *context*, *action*, *decision* and *routing nodes*. *Context nodes* describe the environment in which the guidelines are applied (e.g., a physician in an emergency room). The *action nodes* represent activities of the information system that support the execution of a guideline. The control of the careflow is performed by the *decision* and *routing nodes*. The patient state is retrieved directly from the electronic health record of the healthcare entity.

So far, application of SAGE in practice is very limited. However the Mayo Clinic has plans to apply it in the implementation of guidelines for immunization, diabetes and pneumonia in controlled environments (Zheng et al. 2010).

The SAGE system consists of an execution engine, an event listener and a set of services (terminology, patient record and general applications) (Ram et al. 2004; Tu et al. 2007). The execution engine is called *SAGEDesktop* and it interprets the content of the *context*, *action decision* and *routing nodes*. The event-listener

communicates with the CMS and the EMR with the objective of detecting sudden alterations in a patient's state and notifies the execution engine if that is the case. There is also a terminology server that was added to customize terms used in local applications. The communication between the CMS and the execution engine is facilitated by an Application Programming Interface (API) developed specifically for this purpose, which hints to the main focus of SAGE, interoperability. Semantic and syntactic interoperability of clinical data requires the use of standard data types, terminology, information models and conventions for expressing clinical statements. SAGE puts this to practice through the use of HL7 version three and the UMLS.

5. Active Guidelines in a Clinical Setting

After the analysis of some of the existing projects in the field of CIGs, there are some issues that leave room for improvement and innovation. This work focuses mainly on guideline modeling and decision support during guideline execution.

Concerning guideline modelling, there are some issues that may be pointed out, namely the fact that the available models lack real life application outside the academic environment and are still in the development phase. As such, there is not a reference standard for CIG representation that can be used when developing a system for acquisition and execution of CPGs. None of the models was largely adopted by the health informatics community. The degree of complexity the different models can accommodate is also a matter of discussion: the model cannot be too simple because it may not be able to represent all the information contained in a guideline. A paradigmatic example of this case is Arden Syntax, arguably the model with the highest number of implementations, with its MLMs capable of only representing a decision point in a guideline. At the other end is PROforma, which defines a number of proprietary specifications for data that may be difficult to use and apply to real practice. The challenge is to develop a model capable of representing complex guidelines, yet simple enough to do it with a minimal set of components. Most models and tools do not use terminology and data model standards, which makes the transition to clinical applications in a clinical setting difficult and impairs interoperability with other software tools already used in such environments. Moreover, some of the models require some proficiency in languages to formalize logical rules, numerical expressions and

temporal constraints (e.g., BNF, GELLO), which non-expert clinicians do not have, thus precluding the correct acquisition of CPGs. Furthermore, the software tools provided for editing, visualizing and executing guidelines are often too complex and not user-friendly. The ontology and related tools should be developed in order to allow clinicians with no advanced programming knowledge to revise and customize the guideline representations. Most of the existing models are specialized in certain disease domains which limits their capability to represent other knowledge domains and their applicability to other areas of medicine.

Decision support has also been neglected in the current CIG approaches. The current systems do not complement the decision schemes proposed by their models with techniques to infer the confidence in the outputs of the decision process. Furthermore, the problem of incomplete and uncertain information mentioned in previous sections of this paper remains unaddressed.

Next one will mention the efforts that are underway towards the development and implementation of active CPGs, by extracting elements from the current CIG representations and applying the above mentioned techniques of AI. Ultimately, the aim is to create truly interactive and *living* guidelines by building upon the work done so far and through the introduction of new technologies, models and methods.

5.1. Ontology for Clinical Practice Guidelines

The approach to guideline modelling that one intends to develop presents an abstract view of decision making processes and task management during a clinical procedure (Oliveira et al. 2012). The model is depicted in Fig. 1. The main objective is the development of an ontology capable of accommodating any clinical guideline. To do so, the model should be generic in order to adapt to the context and necessities of different guidelines and, at the same time, allow the definition of constraints characteristic of clinical workflows. There are certain aspects to take into account when developing the model, namely scheduling constraints of the recommendations, time constraints, clinical parameter constraints, terminology and the modularity of the model.

Every task described in a clinical guideline is modelled as a task displayed in an oriented graph. The task is the basic unit of the model. As so, a guideline is

viewed as a *plan*, which is a collection of tasks represented by the following constructs: *action*, *question*, *decision*, and *aggregation module*. A *plan* has any number of instances of these constructs and their ordering and sequence inside a *plan* will be expressed in the form of a linked list that connects the different instances.

Actions represent tasks that must be performed during the execution of a guideline. They can either be clinical procedures, exams, medication prescription, simple recommendations or internal data operations (e.g., calculation of the body mass index from the available clinical parameters).

To feed inputs to the system one uses the *question* task. A *question* is a task to obtain data about the patient's state, in the context of a guideline. It is a data entry point that acts as the substratum for the execution of the other tasks, it is the mechanism through which one feeds information to the CIG system. It also contains a series of constructs to describe the clinical parameters and the data types for their values.

When a decision point is reached in the guideline workflow, the decision task is used. This task contains rules that associate options to the parameters of the patient's state. It has constructs to express the conjunction and disjunction of conditions.

The *aggregation module* aims at controlling special cases in guidelines and groups tasks that are part of a cycle or iteration, creating the conditions for the user to define their *periodicity*, *duration* and *objective*. It also enables the representation of tasks that belong to alternative pathways of the clinical workflow, like the ones that follow a *decision* task, in which the system chooses the next step of the clinical process according to the conclusion reached at the *decision*. One more requirement is the representation of simultaneous tasks that should be executed in parallel.

Another relevant aspect of the model is the *terminology* construct of a *plan*. *Terminology* comprehends the terms used in all the tasks of the plan along with their Concept Unique Identifier (CUI) (Bodenreider 2006), which is a code used in the UMLS Metathesaurus to identify a concept and associate the different terms that can be used as synonymous. This controlled vocabulary is an answer to the ambiguity and vagueness of CPGs.

The *version* construct contains administrative information about the guideline and its authoring, as displayed in Fig. 1. Also in Fig. 1, it is visible a construct called *plan reference* whose function is to make a reference to other *plan* that must be executed in the context of the current one.

To capture the knowledge of the domain and thus create the guideline ontology one will resort to the Ontology Web Language (OWL) (Antoniou and Hamerlen, 2009) because it is the emerging modelling paradigm of the Semantic Web and a standard of the World Wide Web Consortium (W3C). More specifically, the OWL undertaking to be used is OWL-DL, which uses description logics to formalize its classes, individuals and properties. There are a number of reasoners developed to verify the semantic correctness of OWL ontologies (e.g., HermiT, FaCT++), which is an advantage of using this language for modelling guidelines. Moreover, the underlying support for OWL is provided by RDF and XML, which are well known standards.

The set of tools that support the ontology are crucial and they greatly determine the adoption (or not) of the model by the medical community. Knowing this, the ontology will be integrated in a web-based ES, whose advantages were already mentioned in this paper. The importance of the interface in such a system is paramount, mainly because the interface is often the factor of exclusion of a system by clinicians. The system must possess a patient data model and it is essential that the data model is compliant with HL7, namely HL7 version 3 (Dolin and Alschuler 2011).




Fig. 1 Representation of the ontology proposed for guideline modelling.

5.2. Clinical Decision Model

Before applying a clinical decision model that contemplates incomplete information, it is necessary to represent this information in an appropriate way. The Extension to Logic Programming (ELP) (Neves et al. 2012; Novais et al. 2010) is one of the few techniques that enable this representation, using Mathematical Logic. ELP uses two types of negation: default (weak) negation and classic (strong) negation. The use of these two types of negation is the core feature that enables the association of ELP programs to sets of abducibles, represented as exceptions to the extensions of the predicates that represent the clinical parameters. This representation technique augments the usual truth values that are assigned to information (*true* and *false*), by adding the truth value *unknown*, and allows one to represent explicitly negative information. For instance, in cases of inexactitude where there are different possibilities for the value of a clinical parameter, these possibilities are represented as abducibles or exceptions. In cases of uncertainty, if the value of the clinical parameter is unknown, this is represented as a null value.

Decision making in these situations requires the use of an information quantification method. The Quality of Information (QoI) (Neves et al. 2012;

Novais et al. 2010) is a methodology associated with ELP. It is defined in terms of truth values taken in the interval $[0, 1]$ that are attributed to the clinical parameters of the patient according to their number of abducibles and null values. Knowing this, it is possible to calculate the QoI of each condition in a *decision* and calculate scores for each option with the relative weights of its conditions.

By assimilating the concepts of CBR and contextual information with the ELP and the QoI in the context of a runnable clinical guideline, it is possible to devise a decision model that focuses on preeminent matters of guideline execution, non-compliance and inadequacy (Oliveira et al. 2012). Such a decision model is depicted in Fig. 2. Starting with the retrieval of relevant information about the clinical parameters of the patient, this data is presented to the healthcare professional along with a feed of contextual information. This newsfeed is composed of recent news and articles retrieved by an agent from relevant online sources (e.g., the website of the Center for Disease Control and Prevention). Then, based on this information, the healthcare professional assesses the adequacy of the guideline to the case in hand and defines if it is a compliance situation with the guideline he is following or, on the contrary it is a non-compliance situation. In the compliance situation, the decision process moves to the core stages of the decision model. These core stages start with the *Formulation of Clinical Hypotheses*, where the system carries out a survey on the available options in a *decision* task of the clinical guideline. The following stage is *Voting*, where, for each option and consequently for the rules that dictate their choice, the system performs an *Evaluation of Conditions*, to see if they hold true. Next, in the *Evaluation of the QoI*, the system assesses the state of the information responsible for validating each rule and assigns a score to each option. In the following stage, the *Selection of the Clinical Option*, the output of the *decision* is generated. The selected option will be used as a trigger condition for the following tasks in the clinical process. In the *Clinical Task Assignment*, the next task of the clinical process is selected according to its trigger condition. On the other hand, before a non-compliance situation the system may perform one of two things: retrieve a similar case from the case memory or suggest that the healthcare professional alters the current guideline in order to fit the current case. The case memory contains the previous alterations made to the guideline, as well as the clinical parameters of the patient that made him alter it and the output of the process

generated by the alteration. If a similar case does not exist, the healthcare professional alters the guideline accordingly and this alteration will enter the memory case as a new case. Once selected the case or made the alteration, the system shifts from a non-compliance situation to a compliance one and enters the core stages of the model.

This decision model leaves the door open to further research on the complementarity that other techniques that manage uncertainty in different ways, namely Bayesian Networks, Dempster-Shafer Theory and Fuzzy Logic, may offer to the QoI.

The implementation of such a decision model is necessary in order to capture the context of the execution of guidelines and provide measures of confidence in the outputs.




Fig. 2 Clinical decision model for the execution of Clinical Practice Guidelines in an Expert System.

The development of such a decision model is but a step in the construction of a wider decision platform, represented in Fig. 2, where healthcare professionals, members of the same medical team, possibly dispersed across different locations, can discuss the case of a patient in the context of an intelligent environment. Through the use of AI techniques, it is possible to perceive information about the state of stakeholders, namely their attitudes and emotions and thus determine the type of interaction they are developing. If one throws into the equation relevant

knowledge, from exterior sources, concerning the health condition (that is the object of the discussion) and guideline recommendations, a group decision environment is established for healthcare professionals to discuss if a guideline is suitable for the situation at hand and mediate/negotiate solutions. Having all this information enables the medical team, in cases of non-compliance of guidelines, to build new strategies and adapt their content to maximize the probability of a successful treatment.




Fig. 3 Characterization of an intelligent clinical environment where a group decision framework is established using Clinical Practice Guidelines.

6. Conclusions and Future Work

It is widely accepted that the adoption of CIGs would greatly improve the efficiency of healthcare, both in clinical and in economic aspects. This is an on-going research line with numerous people working on the implementation of useful models and the development of execution engines. However, after perceiving the main necessities of paper-based CPGs and analysing the current CIG approaches, one may conclude that they do not solve completely the shortcomings of guidelines, as it is evident by the fact that the available models and systems are not widely implemented.

The line of research proposed in this paper focuses on the development of an ontology for the representation of CPGs that effectively encompasses different clinical domains and, at the same time, shows portability, for implementation in heterogeneous systems. The requirements to achieve this purpose include the conception of structures to accommodate different types of clinical tasks, temporal and scheduling constraints, logical rules, triggering criteria and shows conformance with data and terminology standards. The current CIG models are not complete in the way that they do not have a transversal approach to all of these issues.

It is also viable to conclude that there is a need for a decision model that addresses the aspects of the contextualization of guideline execution and the handling of incomplete and uncertain information. However, healthcare professionals and their opinions should not be excluded of the process because one of the current criticisms to guidelines is that they are too rigid and do not give space for innovation and change..

Acknowledgements

This work is funded by national funds through the FCT – Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) within project PEst-OE/EEI/UI0752/2011".

References

- Aamodt A, Plaza E (1994) Case-based reasoning: foundational issues, methodological variations, and system approaches. *AI Commun* 7(1): 39-59
- Abbod M, von Keyserlingk D, Linkens D, Mahfoul M (2001) Survey of utilization of fuzzy technology in Medicine and Healthcare. *Fuzzy Set Syst* 120:331-349
- Anand V, Biondich P G, Liu G, Rosenman M, Downs S M (2004) Child Health Improvement through Computer Automation: The CHICA System. *Medinfo* 11(1): 187-191
- Antoniou G, Harmelen F (2009) Web Ontology Language: OWL. In: Staab S, Studer R (eds) *Handbook on Ontologies*. Springer Heidelberg, Berlin, pp 91-110
- Anogianakis G, Maglavera S, Pomportsis A, Bountzioukas S, Beltrame F, Orsi G (1998) Medical emergency aid through telematics: design, implementation guidelines and analysis of user requirements for the MERMAID project. *Inter J Med Inform* 52: 93-103
- Bareiss R (1989) Exemplar-based knowledge acquisition: a unified approach to concept representation, classification, and learning. Academic Press Professional, Inc., , San Diego

- Barnett G, Hoffer E, Packer M (1992) DXplain-demonstration and discussion of a diagnostic decision support system. *Proc Annu Symp Comput Appl Med Care*: 822
- Berger J (1994) Roentgen: radiation therapy and case-based reasoning. In: O'Leary D, Selfridge P (eds) *Proc 10th Conf Artif Intell Appl*. IEEE Computer Society Press, Los Alamitos, pp 171-177
- Berners-Lee T, Hendler J, Lassila O (2001) The Semantic Web. *Sci Am* 284(5): 34-43
- Bichindarits I, Marling C (2006) Case-based reasoning in health sciences: What's next? *Artif Intell Med* 36: 127-135
- Bodenreider O (2006) The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Res* 32: 267-270
- Bottrighi A, Terenziani P, Montani S, Torchio M, Molino G (2006) Clinical guidelines contextualization in GLARE. *AMIA Annu Symp Proc*:860
- Boxwala A, Peleg M, Tu S, Ogunyemi O, Zeng Q, Wang D, Patel V, Greenes R, Shortliffe E (2004) GLIF3: a representation format for sharable computer-interpretable clinical practice guidelines. *J Biomed Inform* 37: 147-161
- Brennan A (2000) The Institute of Medicine report on medical errors--could it do harm? *N Engl J Med* 342 (15): 1123-1125
- Bradburn C, Zeleznikow J (1994) The application of case-based reasoning to the tasks of health care planning. In: Wess S, Althoff K, Richter M (eds) *Topics in Case-based Reasoning*, Springer Berlin Heidelberg, Berlin, pp 365-78
- Bratko I, Mozetic I, Lavac N (1989) KARDIO: A study in deep and qualitative knowledge for Expert Systems. MIT Press, Cambridge
- Chawla A, Gunderman R (2008) Defensive medicine: prevalence, implications, and recommendations. *Acad Rad* 15 (7): 948-9
- Cheater F, Closs S (1997) The effectiveness of methods of dissemination and implementation of clinical guidelines: a selective review. *Clin Eff Nurs* 1: 4-15
- Clark D (1990) Numerical and symbolic approaches to uncertainty management in AI. *Artif Intell Rev* 4(2): 109-146
- Codish, S, Shiffman R (2005) A Model of Ambiguity and Vagueness in Clinical Practice Guideline Recommendations. *AMIA Annu Symp Proc*:146-150
- De Clercq P, Blom J, Korsten H, Hasman A (2004) Approaches for creating computer-interpretable guidelines that facilitate decision support. *Artif Intell Med* 31(1): 1-27
- Dempster A (1967) Upper and lower probability inferences based on a sample from a finite univariate population. *Biometrika* 54(3-4): 515
- Dennis S, Edwards S, Partridge M, Pinnock H, Qureshi S (2004) The dissemination of the British guideline on the management of Asthma. *Respir Med* 98: 832-837
- Dinh H, Lee C, Niyato D, Wang P (2011) A survey of mobile cloud computing: architecture, applications, and approaches. *Wirel Commun Mob Comput*. doi:10.1002/wcm.1203/pdf
- Dolin R, Alschuler L (2011) Approaching semantic interoperability in Health Level Seven. *J Am Med Inform Assoc* 18(1): 99-103

- Elkin P, Peleg M, Lacson R, Bernstam E, Tu S, Boxwala A (2000) Toward standardization of electronic guideline representation. *MD Comput* 17(6): 39-44
- Fox J, Alabassi A, Patkar V, Rose T, Black E (2006) An ontological approach to modelling tasks and goals. *Comput Biol Med* 36: 837-856
- Fox J, Black R, Chronakis I, Dunlop R, Patkar V, South M, Thomson R (2008) From Guidelines to Careflows: Modelling and Supporting Complex Clinical Processes. In: Teije A, Miksch S, Lucas P (eds) *Computer-based Medical Guidelines and Protocols: A Primer and Current Trends*. IOS Press, Amsterdam, pp 44-62
- Fox J, Glasspool D, Bury J (2001) Quantitative and Qualitative Approaches to Reasoning under Uncertainty in Medical Decision Making. In: Quaglini S, Barahona P, and Andreassen S (eds) *Artificial Intelligence in Medicine- Lecture Notes in Computer Science*, vol. 2101,, Springer-Verlag, Berlin, pp 272-282
- Fox J, Johns N, Rahmzadeh A (1998) Disseminating medical knowledge: the PROforma approach. *Artif Intell Med* 14(2): 157-182
- Fox J, Thomson R (1998) Decision Support and Disease Management: A Logic Engineering Approach. *IEEE T Inf Technol B* 2(4): 217-228
- Gardner R, Pryor T, Warner H (1999). The HELP hospital information system: update 1998. *Int J of Med Inform*, 54(3), 169-182
- Gruber T (1993). A translation approach to portable ontology specifications. *Knowl Acquis* 5(2):199-220
- Househ M, Kushniruk A, Maclure M, Carleton B, Cloutier-Fisher D (2011) The use of conferencing technologies to support drug policy group knowledge exchange processes: An action case approach. *Inter J Med Inform* 80: 251-261
- Hripcsak G (1994) Writing Arden Syntax medical logic modules. *Comput Biol Med* 24(5): 331-363
- Hutchings A, Raine R, Sanderson C, Black N (2006) A comparison of formal consensus methods used for developing clinical guidelines. *J Health Serv Res Policy* 11 (4): 218-24
- Isern D, Sánchez D, Moreno A (2012) Ontology-driven execution of clinical guidelines. *Comput Methods Prog Biomed* 107(2):122-139
- Isern D, Moreno A (2008) Computer-based execution of clinical. *Int J Med Inform* 77: 787-808
- Jackson P (1990) *Introduction to Expert Systems*. Addison-Wesley Longman Publishing, Boston
- Johnson C, Turley J (2006) The significance of cognitive modeling in building healthcare interfaces. *Int J Med Inform* 75: 163-172
- Kaiser K, Miksch S (2009) Versioning computer-interpretable guidelines: Semi-automatic modeling of 'Living Guidelines' using an information extraction method. *Artif Intell Med* 46(1): 55-66
- Kalra J (2004) Medical errors: an introduction to concepts. *Clinical biochemistry* 37 (12): 1043-51
- Karacapilidis N, Pappis C (1997) A framework for group decision support systems: Combining AI tools and OR techniques. *Eur J Oper Res* 103(2): 373-388
- Kavanagh B P (2009) The GRADE System for Rating Clinical Guidelines. *PLoS Med* 6 (9): 5

- Kim S, Haug P, Rocha R, Choi I (2008) Modeling the Arden Syntax for medical decisions in XML. *Intern J Med Inform* 77(10): 650-656.
- Koton P (1988) Reasoning about evidence in causal explanations. In: Mitchell T, Smith R (eds). *Proc National Conference on Artificial Intelligence AAAI-88*. AAAI Press, Menlo Park, pp 256-263
- Logan R, Scott P (1996) Uncertainty in clinical practice: implications for quality and costs of health care. *The Lancet* 347: 595-98
- Lucas P (2004) Bayesian networks in biomedicine and health-care. *Artif Intell Med* 30(3): 201-214
- Mead, P. (2000) Clinical guidelines: promoting clinical effectiveness or a professional minefield? *J Adv Nurs* 31 (1): 110-116
- Melle W (1978) MYCIN: a knowledge-based consultation program for infectious disease diagnosis. *Int J Man Mach Stud* 10(3): 313-322
- Miller M, Kearney N (2004) Guidelines for clinical practice: development, dissemination and implementation. *Int J Nurs Stud* 41 (7): 813-821
- Neves J, Ribeiro J, Pereira P, Alves V, Machado J, Abelha A, Novais P, Analide C, Santos M, Fernández-Delgado M (2012) Evolutionary intelligence in asphalt pavement modeling and quality-of-information. *Prog Artif Intel* 1(1): 119-135.
- Nikolopoulos C (1997) *Expert Systems: Introduction to First and Second Generation and Hybrid Knowledge Based Systems*. Marcel Dekker, Inc., New York
- Novais P, Carneiro D, Gomes M, Neves J (2012) Non-invasive Estimation of Stress in Conflict Resolution Environments. In Demazeau Y, Müller J, Corchado J, Bajo J (eds.) *Advances on Practical Applications of Agents and Multi-Agent Systems - 10th International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS 2012)* Springer - Series Advances in Intelligent and Soft Computing, vol 155, pp 153-160
- Novais P, Salazar M, Ribeiro J, Analide C, Neves J (2010) Decision Making and Quality-of-Information. In: Corchado E, Novais P, Analide C, Sedano J (eds) *Soft Computing Models in Industrial and Environmental Applications, 5th International Workshop (SOCO 2010)*. Springer - Series Advances in Intelligent and Soft Computing vol. 73, pp 187-195
- Oliveira T, Costa A, Novais P, Neves J (2012) An Interpretable Guideline Model to Handle Incomplete Information. In: Omatu S, Paz Santana J, González S, Jose M, Bernardos A, Corchado J (eds) *Distributed Computing and Artificial Intelligence, 9th International Conference (DCAI 2012)*, Springer - Series Advances in Intelligent and Soft Computing, pp 437-444
- Ohno-Machado L, Gennari J, Murphy S, Jain N, Tu S, Oliver D et al. (1998) The guideline interchange format. *J Am Med Inform Assoc*, 5(4): 357-372
- Ollenschläger G, Marshall C, Qureshi S, Rosenbrand K, Burgers J et al. (2004) Improving the quality of health care: using international collaboration to inform guideline programmes by founding the Guidelines International Network (GIN). *Qual Saf Health Care* 13(6): 455-460
- Pearl J (1986) Fusion, propagation, and structuring in belief networks. *Artif Intell* 29(3): 241-288

- Peleg M, Tu S, Bury J, Ciccarese P, Fox J, Greenes RA, Hall R, Johnson PD, Jones N, Kumar A, Miksch S, Quaglini S, Seyfang A, Shortliffe EH, Stefanelli M (2003) Comparing computer-interpretable guideline models: a case-study approach. *J Am Med Inform Assoc* 10(1):52-68
- Peleg M, Wang D (2006) The Guideline Interchange Format. *Open Clinical*. http://www.openclinical.org/gmm_glif.html. Accessed 23 May 2012
- Perner P (1999) An architecture for a CBR image segmentation system. *Eng Appl Artif Intell* 12(6): 749-59
- Quaglini S, Stefanelli M, Cavallini A, Micieli G, Facino C, Mossa C (2000) Guideline-based careflow systems. *Artif Intell Med*. 20(1): 5-22
- Ram P, Berg D, Tu S, Mansfield G, Ye Q, Abarbanel R, Beard N (2004) Executing clinical practice guidelines using the SAGE execution engine. *Stud Health Technol Inform* 107(1):251-255.
- John R, Zhou S-M, Garibaldi J, Chiclana F (2008) Automated Group Decision Support Systems Under Uncertainty: Trends and Future Research. *Int J Comp Intell Res* 4(4): 357-371
- Rosenbrand K, Croonenborg J, Wittenberg J (2008) Guideline Development. In: Teije A, Miksch S, Lucas P (eds) *Computer-based Medical Guidelines and Protocols: A Primer and Current Trends*. IOS Press, Amsterdam, pp 3-21
- Rosenfeld R, Shiffman R (2006) *Clinical Practice Guidelines: A Manual for Developing Evidence-Based Guidelines to Facilitate Performance Measurement and Quality Improvement*. Otolaryng. *Head Neck* 135:1-28
- Ruzicka M, Svatek V (2004) Mark-up based analysis of narrative guidelines with the Stepper tool. *Stud Health Technol Inform* 101:132-136
- Samwald M, Fehre K, Bruin J, Adlassnig K-P (2012) The Arden Syntax standard for clinical decision support: Experiences and directions. *J Biomed Inform*. doi: 10.1016/j.jbi.2012.02.001
- Schadow G, Mead C, Walker D (2006) The HL7 reference information model under scrutiny. *Stud Health Technol Inform* 124: 151-156
- Schalkoff R (1990) *Artificial Intelligence: an engineering approach*. McGraw-Hill, New York
- Seto E, Leonard K J, Cafazzo J A, Barnsley J, Masino C, Ross H J (2012) Developing healthcare rule-based expert systems: Case study of a heart failure telemonitoring system. *Int J Med Inform* 81(8): 556-565
- Seyfang A, Martinez-Salvador B, Serban R, Wittenberg A, Miksch S, Marcos M, Teije A T, Rosenbrand K (2007) Maintaining formal models of living guidelines efficiently. In: Bellazzi R, Abu-Hanna A, Hunter J (eds) *Proc 11th Conf Artif Intell Med*. Springer-Verlag, Berlin, pp 441-445
- Shahar Y, Miksch S, Johnson P (1998) The Asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines. *Artif Intell Med* 14(1-2): 29-51
- Shahar, Y, Young O, Shalom E, Galperin M, Mayaffit A, Moskovitch R, Hessing A (2004) A framework for a distributed, hybrid, multiple-ontology clinical-guideline library, and automated guideline-support tools. *J Biomed Inform* 37(5): 325-344

- Shenoy P, Shafer G (1986) Propagating belief functions with local computations. *IEEE Expert* 1(3): 43-52
- Shiffman R, Karras B, Agrawal A, Chen R, Marengo L, Nath S (2000) GEM: a proposal for a more comprehensive guideline document model using XML. *J Am Med Inform Assoc* (5):488-98
- Shortliffe E (1993) The adolescence of AI in medicine: will the field come of age in the '90s? *Artif Intell Med* 5(2): 93-106
- Sonnenberg F, Hagerty C (2006) Computer-Interpretable Clinical Practice Guidelines: Where are we and where are we going ?. *Method Inform Med* 45(1): 145-58
- Sordo M, Ogunyemi O, Boxwala A, Greenes R (2004) Description and Status Update on GELLO: a Proposed Standardized Object-Oriented Expression Language for Clinical Decision Support. *Stud Health Technol Inform* 107(1): 164-168
- Sutton D, Fox J (2003) The syntax and semantics of the PROforma guideline modeling language. *J Am Med Inform Assoc* 10: 433–443.
- Straszeka E (2004) Medical Knowledge Representation in Terms of IF-THEN Rules and the Dempster-Shafer Theory. In: Rutwoski L, Siekmann J, Tadeusiewicz R, Zadeh L (eds) *Artificial Intelligence and Soft Computing - ICAISC 2004*, vol. 3070. Springer, Berlin, pp 1056- 1061
- Thomson P (2000) Implementing evidence-based health care: the nurse teachers' role in supporting the dissemination and implementation of SIGN clinical guidelines. *Nurs Educ Today* 20 (3): 207-17
- Thomson R, Lavender M, Madhok R (1995) How to ensure that guidelines are effective. *BMJ* 311: 237-242
- Tu S, Campbell J, Glasgow J, Nyman M, McClure R, McClay J, et al. (2007) The SAGE Guideline Model: Achievements and Overview. *J Am Med Inform Assoc* 14(5): 589–598
- van Gerven M, Taal B, Lucas P (2008) Dynamic Bayesian networks as prognostic models for clinical patient management. *J Biomed Inform* 41: 515-529
- Vetterlein T, Mandl H, Adlassnig K-P (2010) Fuzzy Arden Syntax: A fuzzy programming language for medicine. *Artif Intell Med* 49: 1-19
- Votruba P, Miksch S, Kosara R (2004) Facilitating knowledge maintenance of clinical guidelines and protocols. *Stud Health Technol Inform* 107(1): 57-61
- Wang D, Peleg M, Tu S, Boxwala A, Ogunyemi O, Zeng Q, et al. (2004) Design and implementation of the GLIF3 guideline execution engine. *J Biomed Inform* 37(5): 305-318
- Weijden T, Boivin A, Burgers J, Schunemann H, Elwyn G (2011) Clinical practice guidelines and patient decision aids: An inevitable relationship. *J Clin Epidemiol* 65(6): 584-589
- Woolf S, Grol R, Hutchinson A, Eccles M, Grimshaw J (1999) Potential benefits, limitations and harms of clinical guidelines. *Br Med J* 318 (7182): 527-530
- Zadeh L (1975) Fuzzy logic and approximate reasoning. *Synthese* 30(3): 407-428
- Zheng K, Padman R, Johnson M, Hasan S (2010) Guideline Representation Ontologies for Evidence – based Medicine Practice. In: Khoubati K, Dwivedi Y, Srivastava A, Lal B (eds) *Handbook of Research on Advances in Health Informatics and Electronic Healthcare*

Applications: Global Adoption and Impact of Information Communication Technologies.
IGI Global, New York, pp 234-254

Zhou Z, Jiang Y, Yang Y, Chen S (2002) Lung cancer cell identification based on artificial neural network ensembles. *Artif Intell Med* 24(1): 25–36