Knowledge-Based Systems 98 (2016) 95–116

Contents lists available at ScienceDirect

Knowledge-Based Systems

journal homepage: www.elsevier.com/locate/knosys

A fuzzy framework for encoding uncertainty in clinical decision-making

Aniello Minutolo, Massimo Esposito*, Giuseppe De Pietro

National Research Council of Italy, Institute for High Performance Computing and Networking (ICAR), Via P. Castellino 111, 80131 Naples, Italy

ARTICLE INFO

Article history:
Received 27 October 2015
Revised 16 December 2015
Accepted 20 January 2016
Available online 27 January 2016

Keywords:
Clinical Decision Support System
Clinical guidelines
Fuzzy Logic
Approximate reasoning

ABSTRACT

In recent decades, technological advances coupled with research efforts have made possible to develop very complex Decision Support Systems (DSSs) able to exhibit highly sophisticated reasoning capabilities in order to improve clinical decision-making, and, thus, promote more efficient care practices. One of the most significant factors influencing, and in particular limiting, the adoption of clinical DSSs is represented by the modality of representation and computerization of clinical guidelines in form of patient-specific recommendations. Until now, many knowledge representation formalisms have been developed, mainly focused on time-oriented guidelines. However, they can generate an unrealistic over-simplification of reality, since they are not able to completely handle uncertainty and imprecision typically affecting clinical guidelines. In this respect, this paper proposes a novel fuzzy framework expressly thought for building guideline-based DSSs, by efficiently modelling and handling the peculiarities of clinical knowledge affected by uncertainty and imprecision and encoded in the form of guidelines. This framework has been devised with the aim of: (i) offering a set of patterns for easily inserting and editing clinical recommendations belonging to a guideline as a group of one or more fuzzy rules expressing positive evidence and one fuzzy ELSE rule including negative evidence; (ii) defining a set of Fuzzy Guideline Systems (FGSs), one for each guideline encoded, characterized by ad-hoc configurations for the mathematical operators necessary to evaluate rules and generate the outcome expected; (iii) implementing a multi-level inference scheme able to treat different FGSs as a whole and efficiently enable their interconnection, i.e. the chaining among the groups of fuzzy rules belonging to each FGS; (iv) exposing a set of graphical facilities for guiding the definition of fuzzy rules to be embedded into a clinical DSS and enabling their automatic encoding and execution by using an XML-based machine executable language. A usability evaluation has been performed, showing a good satisfaction of medical users with respect to the framework implemented, and, thus, proving both its feasibility and usefulness.

1. Introduction

The twenty-first century has seen incessant changes in health care practices, due to an ever-expanding knowledge base in clinical medicine, and a growing clinical data set describing every patient characteristic from phenotype to genotype [26]. In this constantly evolving scenario, clinical decision-making has become an inordinately complex and multifaceted process, where health professionals are required to make decisions with multiple foci (e.g. diagnosis, intervention, interaction and evaluation), in dynamic contexts, using a diverse knowledge base (including an increasing body of evidence-based literature), with multiple variables and individuals involved [49]. Such decisions can differently affect clinical practices, producing substantial variations due to a host of issues, such as the comprehension of medical information by health care personnel, the development of skill and models of competency for complex tasks, the coordination of knowledge among persons with different backgrounds [42].

In recent decades, technological advances coupled with research efforts have made possible to develop very complex Decision Support Systems (hereafter, DSSs) able to exhibit highly sophisticated reasoning capabilities in order to improve clinical decision-making, and, thus, promote more efficient care practices. Although, in the past, the academic world have provided strong motivation for working in the area of clinical DSSs, only in the last years these systems have been markedly recognized as practical tools able to proficiently face the inexorable growth in healthcare complexity and cost [40].

* Corresponding author. Tel.:+39 0816139512; fax: +39 0816139531.
E-mail addresses: aniello.minutolo@na.icar.cnr.it (A. Minutolo), massimo.esposito@na.icar.cnr.it (M. Esposito), giuseppe.depietro@na.icar.cnr.it (G. De Pietro).

http://dx.doi.org/10.1016/j.knosys.2016.01.020

As a result, the role of computer systems in clinical practice has changed tremendously and clinical DSSs have become an essential element of modern clinical practice, whose usage has been requested by a virtuous desire to support physicians in dealing with their daily activities and contribute to their continuing medical education with "just in time" clinical information. Among the different typologies of DSSs that can be realized, knowledge-based DSSs have been evaluated as having the most efficient impact, since they combine a computer-interpretable formalization of clinical guidelines with electronic medical records in order to concretely bring evidence-based medicine into clinical practice. Indeed, knowledge-based DSSs enable clinicians to access, at the time and place of a consultation, not just the clinical decision tailored to the specific patient considered, but also the set of computerized clinical guidelines from which it has been derived, and the literature that explains its scientific evidence.

However, even if the new central role of clinical DSSs has been widely recognized, the examples of knowledge-based DSSs concretely used in daily practice are extremely reduced due to a host of reasons [40]. One of the most significant factors influencing, and in particular limiting, the adoption of clinical DSSs is represented by the modality of representation and computerization of clinical guidelines in form of patient-specific recommendations. Indeed, most guidelines are represented in the form of text documents (paper or electronic) which contain narrative sections regarding background clinical issues, the methodology of guideline development and supporting evidence, and a summary of standard care recommendations for diagnostic or therapeutic decision-making [30,48]. As a result, the thrust of a guideline can be distilled into a series of condition-action recommendations, expressed in the form "if the conditions are verified, then one should perform the recommended actions" [48].

Guidelines may include 'clinical algorithms' in the form of flowcharts to be followed in appropriate situations, but, despite their computer-inspired notation, they are usually intended for humans [18]. Moreover, their textual nature makes them complex to be computerized since, first, they are not expressed according to a computer-inspired formulation and, more importantly, they are intrinsically pervaded by uncertainty and imprecision.

Until now, many knowledge representation formalisms have been developed to fulfil the mismatch existing between the unstructured narrative form of guidelines and the formality that is necessary for the operationalization of clinical knowledge in their computerized versions for DSSs. These solutions have been focused on time-oriented guidelines, and represent each single recommendation, arranged into a control flow structure, in the form of one or more if-then rules, that are simple statements that link a logical combination of conditions to a set of actions [10,43]. However, they can generate an unrealistic over-simplification of reality, since they are not able to completely handle uncertainty and imprecision typically affecting clinical guidelines and pertaining, for instance, the diagnosis, the accuracy of available diagnostic tests, the natural history of the disease, the effects of treatment in an individual patient or the effects of an intervention in a group or population as a whole [22]. For instance, hypertension guidelines state that patients with systolic blood pressure greater than or equal to 160 mmHg or diastolic blood pressure greater than or equal to 95 mmHg could be recommended for the hypertension treatment. Thus, completely different recommendations could be suggested for patients with values of blood pressure that are close but placed around the thresholds (e.g. for patients with values of systolic blood pressure equal to 159 mmHg and 161 mmHg, respectively), so as to lead to possible wrong interpretations with respect to a direct evaluation guided by common sense or by heuristics.

In the past, Fuzzy Logic [59] (hereafter, FL) has widely demonstrated its capability to overcome such critical issues in encoding clinical guidelines, as proved by numerous examples proposed in literature [7,12,25,32]. This is due to the fact that FL formalism is suitable to deal with uncertainty intrinsic to many kinds of guidelines, by offering a more realistic representation of clinical knowledge in the form of fuzzy "if–then" rules to be included into a DSS, and also providing an understandable language for describing them in a natural manner close to the human perception.

The most relevant drawback of the existing approaches for formalizing clinical guidelines in terms of fuzzy rules within a clinical DSS relies on the absence of any type of vertical arrangement for the specific type of knowledge modelled.

First, typically, a guideline can be formalized as a set of one or more fuzzy rules, where different clinical information is combined into the premises in order to infer, as a consequence, the strongest recommendation among the various possible alternatives. However, in contrast to classical logic, the mathematical expression for evaluating fuzzy rules is not uniquely defined and different operators could be used to aggregate premises or infer conclusions, leading, of course, to different results. Indeed, Fuzzy Logic does not foresee any abstract theoretical criteria that make some operators better than the others, enabling to treat many different situations within the same semantic framework [57]. The price for such universality is the lack of straightforward selection rules for those mathematical operators. As a result, every particular guideline should be encoded into a collection of fuzzy rules and examined separately by using the most appropriate mathematical operators, chosen in order to make the behaviour of the rules similar to that of the original guideline.

Secondly, guidelines typically use positive evidence to guide decision-making, in accordance with the assumption that the presence of a causally relevant factor is always of maximum value to confirm a hypothesis. In particular, physicians tend to seek and use predominantly only evidence to confirm, rather than to weaken, a medical hypothesis, for instance related to the diagnosis of a disease [13]. However, when a guideline is codified in terms of fuzzy rules, also the absence of a causally relevant factor should be appropriately taken into account in order to produce a final decision that is coherent with the one expected by the physician. For instance, negative evidence should be modelled in terms of rules and weighted accordingly in case when a sign or a symptom is always expected in order to produce a diagnosis, and it has not been observed at all. More generally, since FL states that each rule whose premise has a non-zero matching degree with respect to its inputs will contribute to the final decision with strength equal to the matching degree of its premise, thus, both positive and negative evidences should be modelled at the same way. However, the definition of fuzzy rules for expressing the negative evidence is a very thorny task, due to the lack of clinical knowledge explicitly reported in the original guidelines. Thus, clinicians are forced to formalize expressly a complementary set of rules to address this issue.

Thirdly, recommendations contained into clinical guidelines describe both the dependencies between one or more premises involving clinical information, and the outcomes regarding a decision to suggest. The premises can be referred, for instance, to the state of the patient, describing his/her clinical situation in terms of physiological parameters, symptoms, and execution stages of therapies or medications, whereas the outcomes can regard, for instance, a diagnosis or a new therapy to perform [56]. However, these recommendations should not be simply considered as separate silos to be evaluated and executed independently, because they can be potentially chained among them. More specifically, guideline chaining occurs when one guideline recommendation has a premise that matches the outcome of another guideline recommendation. For instance, a recommendation suggesting an adjustment of a therapy can include, as one of its premises, the