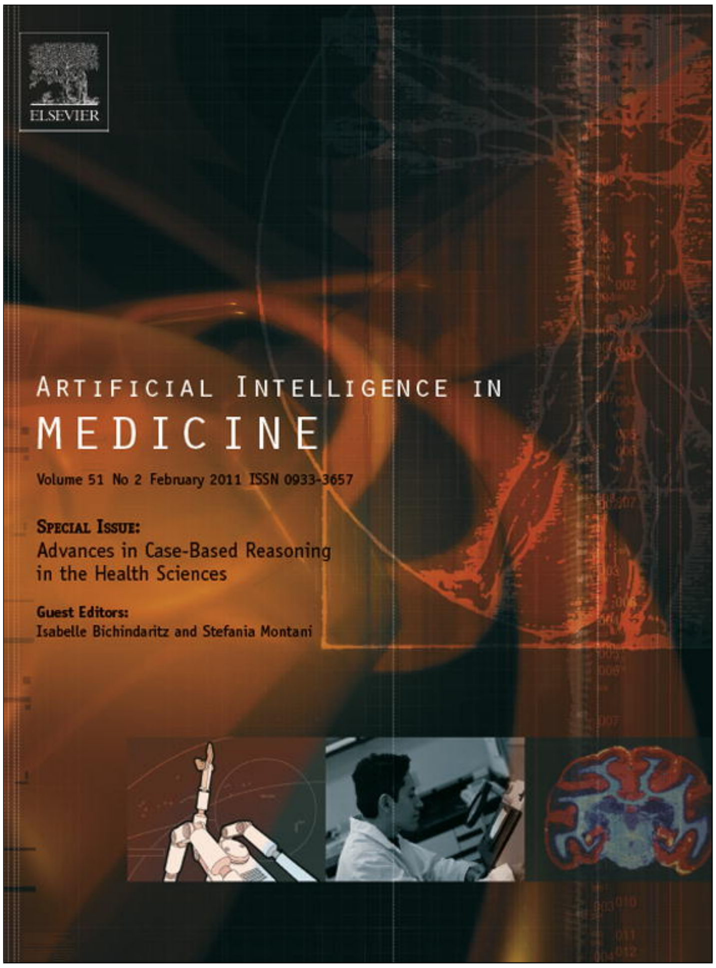
ELSEVIER
ARTIFICIAL INTELLIGENCE IN
MEDICINE
Volume 51 No 2 February 2011 ISSN 0933-3657
SPECIAL ISSUE:
Advances in Case-Based Reasoning
in the Health Sciences
Guest Editors:
Isabelle Bichindaritz and Stefania Montani

Artificial Intelligence in Medicine 51 (2011) 125–131

Contents lists available at ScienceDirect

Artificial Intelligence in Medicine

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

How to use contextual knowledge in medical case-based reasoning systems: A survey on very recent trends

Stefania Montani*

Dipartimento di Informatica, Università del Piemonte Orientale, Via Bellini 25/g, Alessandria, Italy

ARTICLE INFO

Article history:
Received 19 December 2008
Received in revised form 23 April 2010
Accepted 26 April 2010

Keywords:
Case-based reasoning
Context
Retrieval efficiency
Knowledge validity
Knowledge and reasoning adaptation

ABSTRACT

Objectives: This paper aims at systematizing the ways in which the contextual knowledge embedded in the case library can support decision making, within case-based reasoning (CBR) systems. In particular, CBR applications to the medical domain are considered.
Methods and material: After a quick survey on the definition and on the role of context in artificial intelligence research, we have focused on CBR, with a particular emphasis on medical applications. In this field, we have identified a number of very recent contributions, which strongly recognize context *per se* as a major knowledge source. These contributions propose to maintain and to rely on contextual information, in order to support human reasoning in different fashions.
Results: We have distinguished three main directions in which contextual knowledge can be resorted to, in order to optimize physicians' decision making. Such directions can be summarized as follows: (1) to reduce the search space in the case retrieval step; (2) to maintain the overall knowledge content always valid and up to date, and (3) to adapt knowledge application and reasoning to local/personal constraints. We have also properly categorized the surveyed works within these three clusters, and identified the most significant ones, able to exploit contextual knowledge along more than one direction.
Conclusions: Innovative applications of the contextual knowledge recorded in the case library, described and systematized in this paper, can trace promising research directions for the future.

1. Introduction

With its cability of incrementally collecting, reusing and sharing the knowledge implicitly embedded in previously experienced situations, case-based reasoning (CBR) [1] is currently recognized as a very well suited reasoning methodology in medical applications. A significant number of valuable CBR approaches in medicine have been proposed in the literature so far (see e.g. [2,3] for recent surveys on the topic).

As a matter of fact, CBR may mitigate the knowledge acquisition effort, since representing a real world situation as a case is often straightforward. Given a set of meaningful features for the domain, it can be sufficient to identify the value they assume in the situation at hand. In classical CBR approaches, the so-obtained set of ⟨*feature*, *value*⟩ pairs provides the *problem description*, which is typically coupled with information about the applied *solution*, completing the situation-action pattern adopted on that occasion. Such data encompasses an amount of domain knowledge, which can be memorized without the need of making it *explicit* in a more abstract and generalized form, as it would be required by other methodologies (e.g. rule-based or model-based reasoning). Maintaining this kind of operative information is extremely useful in the medical domain. In fact it can complement explicit domain knowledge (especially for those diseases that are still not well understood, or for which generalized rules or models do not apply), or it can help make domain knowledge itself immediately usable in real clinical environments (e.g. in the case of clinical protocols interpretation [4]). Operative knowledge can even lead to changes in organizational settings, and improve the overall quality of care provision [5].

Operative knowledge is stored and maintained in the case base without applying filters, systematizations or abstractions. In this way, it does not separate the actual domain knowledge it implicitly embeds from the circumstances and the details of the situation in which that knowledge was exploited [6]. Among the case details, the *context* of application seems to be particularly relevant.

The *context* of a system captures the environment where it operates, including all additional or non-functional aspects that, while not being core to the system's behaviour, nevertheless may affect the way in which that behaviour should be optimized. In medical applications, in particular, the "system" should be interpreted as the set of (human or not) actors and procedures involved in patient care. In this domain, the *context* includes physical characteristics of the clinical setting, as well as many other aspects, such as knowl-

* Tel.: +39 0131 360158; fax: +39 0131 360198.
E-mail address: stefania.montani@unipmn.it.

0933-3657/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.artmed.2010.09.004

edge states (of the patient and of the physician) and emotions. A case then contains a domain knowledge core (i.e. the situation-action pattern, and the rationale behind it), which is surrounded by (and interlaced with) contextual details [6]. These details basically concern: (i) the specific patient being considered; (ii) the specific physician responsible for the described procedure; and (iii) the specific health care provision environment where the case took place.

Abstraction and generalization of domain knowledge contained in the case base, although possible, require the removal of circumstances and details from individual cases, which operatively translates into removing the case context. Removing context may create a gap between the obtained generalized knowledge, and future problem instances [6], giving birth to a knowledge content which is better suited at the population level than at the individual one. While this outcome may be desirable in some situations (e.g. when defining a clinical guideline), it may hamper the decision making activity in others.

Stemming from these considerations, a very recent research trend is now strongly and explicitly recognizing context as a major knowledge source for decision making in many medical applications. Along this line, in this paper we aim at stressing the importance of contextual knowledge *per se*, and at illustrating how it can be relied upon within a CBR framework in order to fulfill several goals.

The paper is organized as follows. In Section 2 we recall the research results on the notion and on the use of context in artificial intelligence (AI) in general, and in CBR in particular. In Section 3, which contains the main methodological contributions of this paper, we systematize the ways in which context can be relied upon in medical CBR applications, identifying and describing three main research directions for the future. Final observations and comments are addressed in Section 4.

2. Context in artificial intelligence

Context has been playing an important role in several domains for many years, since it is seen as a core concept both in the philosophical and linguistic area, and in the scientific one [7].

AI research, in particular, is increasingly interested in the notion and in the possible uses of context, as testified, for instance, by a series of international conferences (see http://mainesail.umcs.maine.edu/Context/, last accessed on March 31st, 2010), which have been taking place since the early 90s.

Within AI, context is used in many different areas [8]: in natural language processing, context is used to assign an interpretation to assertions and resolve ambiguities; in information retrieval, context helps refine the queries made by users; in distributed AI, context is used as a flexible formal tool for the design of systems of autonomous agents; in human–machine interaction, context is used to design context sensitive applications and interfaces. However, in the rest of the section, we will concentrate on the use of context in knowledge representation and reasoning (KRR).

KRR is the AI area whose aim is to devise languages that can represent (1) what (intelligent) programs or agents know about their environment, and (2) the reasoning processes that allow them deriving new knowledge, with the goal of solving problems and to support human decision making. Actually, CBR belongs exactly to this area. As there is a lack of consensus on the meaning of the word "context", we will resort to the definition provided by Brezillon in [7], since it is well agreed upon, at least in the KRR community: "Context is what constraints problem solving without intervening in it explicitly".

Very interestingly, some authors in KRR (see e.g. [9]) identified one of the reasons why the first knowledge-based decision support systems (KBS) failed exactly in the lack of an explicit representation and treatment of context. Actually, the knowledge acquired from humans has a rich contextual component, which was generally not acquired in the knowledge formalization and generalization phase of early KBS. Moreover, in KBS the knowledge was acquired once and for all in the beginning (i.e. in the system design phase), and was therefore unable to take into account misunderstandings and problems that could arise during the system usage [10]. The absence of contextual knowledge also made it very difficult to generate explanations about the reasoning choices of the systems. On the other hand, human decision making can be more properly supported only if the acquired knowledge does *not* undergo a *strong generalization* process, and if its acquisition is *incremental* [7].

Several efforts have been completed in order to take into account the observations above. For instance, Walther et al.[11] have defined a context ontology, which includes constraints about how the knowledge terms should be used and combined, while Cyc [12], the largest common sense knowledge base ever built, implements and exploits an explicit notion of context.

From a more technical viewpoint, attempts have been made in order to represent context both using appropriate logic formalisms [13,14], and using rule-based or model- (e.g. graph-)based representations [15]. Some of these methods have provided extremely expressive mechanisms to exploit context in formal theories, as proved e.g. by their recent application to the Semantic Web (see [16]).

However, the acquisition of knowledge in context is very challenging [7]. The contextual information is sometimes implicit, and strongly interlaced with the more objective details about a fact, so that it becomes really hard to separate context, and to provide it in the explicit form, as it is required e.g. by rule-based approaches.

CBR thus appears to be more suited for capturing and maintaining contextual knowledge. Actually, as observed in the Introduction, CBR relies on operative knowledge, which does not require generalizations and abstractions. Therefore, cases implicitly maintain both objective details as well as contextual ones [6]. Moreover, in CBR knowledge acquisition is naturally incremental [1], since new cases are continuously acquired during the system usage.

In CBR, context can serve as a major knowledge source in the various reasoning phases, from retrieval to adaptation. Possible roles of context in CBR applications have been explored since the early years. Leake [17], for instance, has developed a CBR system which exploits contextual knowledge to select the right explanation to be shown to the user (remember that the lack of explanations was one of the weaknesses of the early KBS). Ozturk and Aamodt [18] have relied on contextual information in order to retrieve appropriate cases within the knowledge-intensive CBR system Creek. More recently, Lieber and Napoli [19] have integrated the C-OWL context ontology with CBR: semantic relations between contexts and the associated reasoning mechanisms allow the CBR process in a particular viewpoint reusing and sharing information about the problem and the already found solutions in other viewpoints. Another, very recent line of research is about CBR and context-awareness [20]. In this field, scientists are studying appropriate ways to handle and exploit context for Ambient Intelligence. This objective is particularly challenging, due to the fuzziness of context information in such domain, given – especially in mobile scenarios – the rapidly changing environments and the unstable information sources. They are proposing applications of CBR to pervasive computing, autonomic systems, and ubiquitous computing, relying on sensed and real-world features. How to use context is also being investigated in recommender systems. In this domain, it is progressively being recognized that cases should not store only information about the products to be recommended, otherwise no real learning process can be supported [21]. On the other hand, cases should maintain contextual information too, about the user interests [22], and about the recommendation sessions where the products themselves were selected [23].

The usefulness of contextual knowledge has been recognized early in medical CBR systems too. For instance, the high number of hybrid and multi-modal reasoning contributions (see e.g. [24,25]), integrating CBR with other knowledge sources and/or reasoning methodologies, testifies to the advantages of coupling formalized knowledge with an operative one. It also indirectly testifies to the advantages of exploiting context, since context is incorporated in operative knowledge [6].

Moreover, very interestingly, recent medical CBR applications are now shifting towards an even more explicit and massive exploitation of contextual knowledge. The next section will provide a description and a categorization of these innovative approaches.

3. Using contextual knowledge in medical case-based reasoning systems

Context seems to be particularly relevant in medical applications, where inter-patients variability is extremely high, and where diagnostic and therapeutic decisions always need to be properly tailored to the single patient's peculiar situation. An exhaustive description of the patient thus cannot ignore contextual information. In fact, recent CBR systems applied to the medical domain are relying on context as a major knowledge source for supporting decision making, in various ways. In this section, after some observations about the nature and the representation of contextual knowledge in medical cases (see Section 3.1), we analyze such approaches in detail (see Section 3.2).

3.1. Contextual knowledge in medical cases

Dealing with the medical domain, by *case* we will refer to a given medical procedure applied to a patient, such as a visit, a treatment, or the execution of a whole clinical guideline (GL).

In a medical case, the context is often captured as implicit [6] or difficultly measurable knowledge, deeply interlaced with more objective patient and procedure information. Objective measurements/findings (e.g. the patient's age), can be recorded together with more subjective characteristics of the patient (e.g. moods, emotions) and of the physician (e.g. personal preferences, implicit perceptions), and with the characteristics of the physical environment where the case took place (e.g. available resources and skills, physical or cultural constraints): all these *contextual* elements can influence the medical procedure implementation and its outcome.[1]

It is worth noting that the distinction between the data which are objectively measurable, and the ones that are subjective, or more difficult to express and quantify, is not always clear-cut. For instance, in psychiatry, behaviours and emotions are well codified, and should not be considered as subjective and purely contextual information. Moreover, information about people's behaviour is sometimes strictly paired with more objective measurements, and should not be considered disjointly from them, when the aim is to have a complete description of the situation. For instance, pubertal compliance problems with respect to insulin therapy may be diagnosed by looking at the age of a diabetic patient, at the negative trend in her blood glucose level control, and at a (less objectively quantifiable) behaviour that leads to a reduced care in her self-monitoring activity.

In conclusion, it can be difficult, or unuseful, to make a clear distinction between objective data and contextual information in many medical domains. The strength of CBR is exactly that of allowing the storage, retrieval and reuse both knowledge types, without requiring their explicit separation [6].

3.2. Using contextual knowledge: recent trends

By analyzing recently published system descriptions, we have identified three main directions in which contextual knowledge can be profitably resorted to within a CBR framework, namely:

1. to reduce the search space in the retrieval step;
2. to help maintain the overall knowledge content always valid and up to date, and
3. to adapt knowledge application and reasoning to local/personal constraints.

Details of these lines of research will be provided in the next sections.

3.2.1. Reducing the search space

Contextual information may be relied upon to select a subset of the cases contained in the case base, thus reducing the search space for the retrieval step in a CBR system. This choice clearly can make retrieval computationally faster, and hopefully more meaningful, since only cases taken under comparable circumstances are retrieved.

The approach has been applied e.g. in the field of diabetes [27,28], and in the field of hemodialysis [29]. In both domains, the patient's behavioural or emotional situation may strongly influence the therapy outcome. For instance patients in the pubertal age may be less compliant with insulin therapy, or elderly patients may be less compliant in completing a (long) hemodialysis session, especially if they habitually feel bad during it. In these domains, it is not easy to objectively quantify and code behaviours and emotions, which represent contextual information, deeply interlaced with other patient's features. However, contextual knowledge can guide search space reduction, e.g. using a clustering technique [29]. Clustering does not aim at labelling the cases in a group with a specific tag (as it happens in classification), where the tag represents a piece of generalized domain knowledge, extracted from the subsumed cases. In clustering contextual knowledge remains in the implicit form; however, the most similar cluster(s) allow the identification of the cases collected under similar circumstances, and the limitation of retrieval just to them.

3.2.2. Keeping knowledge up to date

In medical practice, past diagnostic or therapeutic recommendations tend to become quickly obsolete, due to the development of new technologies and to the availability of new scientific evidences. Therefore, when a CBR approach is resorted to, the case base should be maintained and updated in a life-long learning perspective [2], where recent developments and findings are integrated, while old ones are carefully evaluated before being reused.

Anyway, while highly abstracted evidence loses its significance when becoming obsolete (and its re-application may even be harmful), the knowledge implicitly embedded in past cases can always be re-interpreted in the future. Re-interpretation allows the verification of its validity, possibly leading to different conclusions in the light of recent scientific discoveries [30]. This observation further corroborates the hypothesis that memorizing operative knowledge and relying upon it is a desirable task. In fact, since operative knowledge is context-dependent, it allows rebuilding theories *dynamically*. On the other hand, explicit domain knowledge is *static*: it cannot be re-interpreted, since the original context from which it was abstracted is no longer available [6]. The heterogeneous, deeply interlaced information encoded in cases is thus more capable to

[1] For instance, non-compliance to GL may emerge due to the need to adapt the GL itself to the local reality ([26], see also Section 3.2.2).

retain validity and relevance over time and space [6], at least to a reasonable extent.

However, at the same time, the need to properly manage contextual knowledge in order to preserve recency and to continuously verify validity clearly makes the acquisition, adaptation and reuse activities more challenging.

These challenges are approached e.g. in [30], by resorting to the definition and exploitation of prototypes [31]. Prototypes are a generalization from single to clustered typical cases. Their main purposes are to structure the case base and to guide and speed-up the retrieval process. However, prototypes can also help reuse (by possibly adapting) past cases. In the case of prototypes the definition of a reuse/adaptation strategy becomes easier, since the specific details of ground cases leave space to a more generalized kind of knowledge. If a hierarchy of prototypes exists, reuse/adaptation can also be seen as a top down search to find the most specific case that fits for the current problem. Moreover, despite the fact that prototypes do summarize a set of (very similar) cases, and generalize them to some extent, they do not constitute highly abstracted evidence, and they do not remove contextual knowledge (at least, not completely). Therefore, knowledge in prototypes may still be reviewed and re-interpreted to reach possible new conclusions. The work in [30] proposes to mine some prototypes from the biomedical literature, as well as to learn others by applying conceptual clustering on existing cases in the case base. In particular, the new prototypes automatically built from the literature take into account the flow of biomedical advances. On the other hand, the prototypes learned by clustering organize and represent clinical practice ground cases. Prototypes then guide retrieval and reuse. In the retrieval phase, both prototypes and ground cases are extracted; ground cases are retrieved also when they are not indexed under any prototype in the taxonomy. Prototypes are therefore “add-ons” to the case-base, i.e. forms of organizing the cases allowing improving search and reuse/adaptation.

The work in [32], on the other hand, suggests that human perceptions and feelings can be exploited to make experience reuse more reliable and up to date in CBR systems. In particular, the authors observe that physicians' comments and notes appended to clinical practice cases, and usually reported as text, may be valuable to capture data that are not visible in more objective measurements (e.g. in sensor readings). Moreover, such notes provide additional information to better interpret measurements themselves. As a matter of fact, textual notes implicitly record human perceptions, whose exploitation may enhance objective features understanding and reuse/adaptation. The work is still in its implementation phase; however, it provides an analysis of the major issues which have to be afforded to use contextual information to maintain knowledge up to date, especially when supporting case reuse. From a technical viewpoint, methodologies from the textual CBR research area will be relied upon to deal with physicians' notes interpretation and exploitation.

The work in [33] can be categorized in the recency maintenance line of research as well. In particular, this paper describes a CBR tool which our research group is developing for managing non-compliance with clinical guidelines (GL). GL can be defined as a means for specifying the most appropriate clinical procedures and for standardizing them. Despite their proved efficacy in improving patient care, examples of non-compliance with their prescriptions are often reported. Reasons for this can be identified in an improper or weak GL definition, e.g. due to the presence of biases, changes in evidence, or, more frequently and more interestingly for the purposes of our discussion, obsolescence of data and/or procedures. In [33], non-compliance episodes are represented as cases. In particular, the problem description consists of the reasons for non-compliance (e.g. abnormal patient's parameters, unavailable resources, different medical opinions). When possible, these data are recorded as explicit features (e.g. in the case of abnormal patient's parameters). In other situations, they are stored as additional textual comments, or remain hidden within the explicit features themselves, as contextual information. The case solution consists of the changes in the GL procedure that were implemented. In front of a new non-compliance episode, the tool is able to retrieve past cases, and to show them to physicians, for further analysis. Retrieval entails the treatment of several issues, such as the management of heterogeneous features (i.e. numeric, symbolic and textual), and the identification of similar cases by applying appropriate metrics.

The tool is also able to extract more general indications from the ground cases. In particular, after some non-compliance cases have been stored in the case base, it compares them, aiming at discovering: (i) frequent modifications, confirmed by several examples; (ii) atomic modifications, identified by separating heterogeneous information within a single case; (iii) semantic relations between different situations. By iterating this process, (components of) two or more different cases can be merged in a new structure, namely a prototype [31], confirming and generalizing the content of the elements from which it was built. In the system, prototypes can then guide and support further retrieval sessions, and can be considered for a more complete GL revision by a committee of medical experts.

As already observed, prototypes support the evolution of the case base content, but do not lead to an explicit formalization of domain knowledge, nor to an abstraction from contextual information, thus preserving the possibility of knowledge re-interpretation in time.

3.2.3. Adapting knowledge application and reasoning

Context can also support decision making in many medical applications by enabling an adaptation of procedures and reasoning strategies to the local or personal constraints described by contextual information itself (see also Brezillon's definition in Section 2). A really applicable and usable system for care provision should therefore be context-aware.

The issues of reaching context-awareness and of realizing context-based reasoning are being routinely addressed in mobile and ubiquitous computing (see also Section 2), but are now being recognized as key aspects also for a wide range of other areas, among which patient care.

In particular, how to implement a context-aware CBR tool in the medical domain has been explored in [34], a work which abstractly analyzes the problem of changing medical procedures (i.e. hospital workflow steps) on the basis of the constraints imposed by the local contextual reality, and explores possible solutions for case acquisition and case maintenance.

On the other hand, the work in [32] underlines how context-awareness is important for decision support in medical diagnosis as well as in treatment plans, and suggests resorting to the environmental situation to design appropriate human–computer interaction strategies and interfaces.

Remarkably, also the work in [33] can be interpreted in the perspective of designing a context-aware system, able to adapt domain knowledge and to tailor reasoning strategies to local constraints. Actually, non-compliance episodes with respect to GL can also be due to the need to adapt the GL itself to the local reality. Environmental constraints (i.e. human skills, technical and physical resources, or funds locally available at a given hospital) can motivate non-compliances, and legitimate changes, able to deal with local/cultural characteristics, can be made in recommendations, even when the evidence they are based on is the same [26]. Our tool, able to retrieve and reorganize past non-compliance cases, can thus support the design of a locally executable version of the GL, and make it really available at the point of care (see also [4]), by

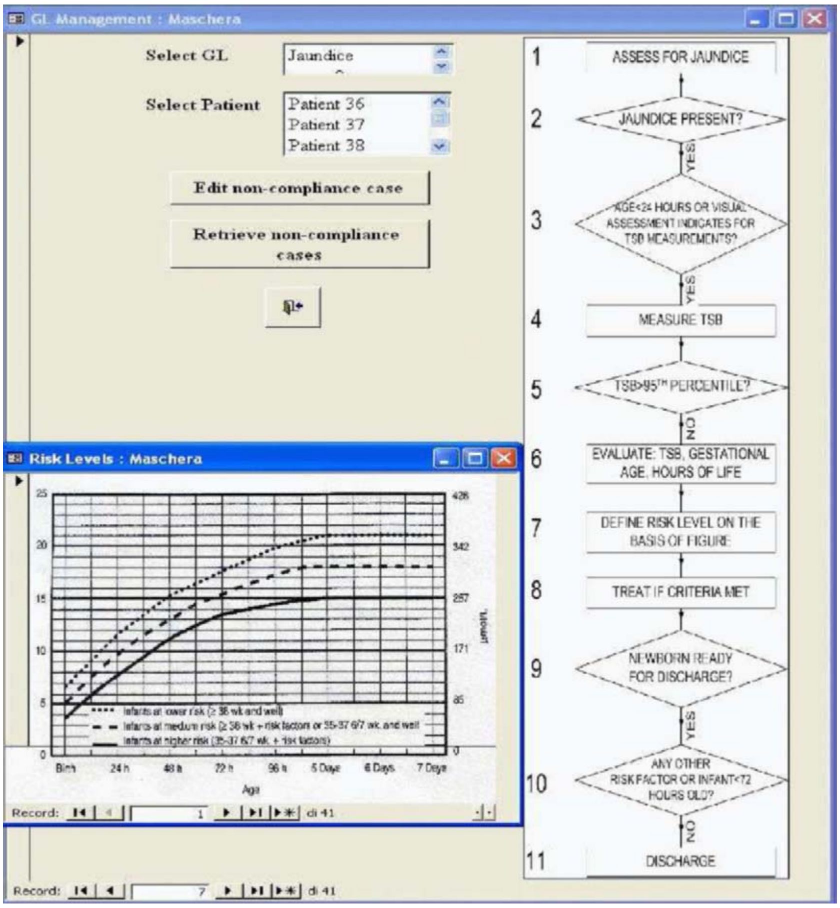

Fig. 1. Part of the American Academy of Pediatrics guideline for managing jaundice, visualized using the tool in [33].

helping the persons in charge of patient care better understand the needs of the environment where the GL itself is meant to be applied. We thus operate somehow similarly to [34], which, although very preliminarily, approaches the problem in the even larger perspective of medical workflow management (note that GL application is typically just one aspect of workflow execution, which also includes resource assignment and activity coordination).

While the contributions in [34,32], which testify to a growing interest in the topic of context-awareness in medical CBR, are not fully implemented, the tool in [33] is in its testing phase, showing how the use of contextual knowledge for adaptation to local constraints is not only a research speculation, but can also become a concrete advance in medical practice. As an example, in the next section we provide a description of an interesting case study.

3.2.4. Adapting knowledge in practice: a case study

Our tool for managing non-compliance to GL has recently been made available at the Obstetrics Department of Policlinico S. Matteo hospital in Pavia, Italy. In this section, we report on a significant experience collected at this test site.

Problem statement. The Obstetrics Department of Policlinico S. Matteo suffers from an insufficient availability of beds for mothers, in front of a growing number of women who choose this site for the birth of their babies. Therefore, if a baby is discharged, but later on needs to be re-hospitalized (e.g. due to jaundice problems), it may become very difficult to re-admit the mother as well, thus creating troubles with breast feeding. Pavia implements the American Academy of Pediatrics GL for managing jaundice [35]. Fig. 1 shows a snapshot of the tool, being applied to this GL. By following the American Academy of Pediatrics GL, at 48 h of age Total Serum Bilirubin (TSB) is measured in all babies in whom jaundice appears (by visual assessment) to be severe enough. High risk (see the pop-up window in Fig. 1) infants are treated with phototherapy; medium risk ones are re-examined before discharge (typically at 72 h of age).

Non-compliance cases generation. Motivated by bed unavailability in Pavia, pediatricians generated two sets of non-compliance cases, where low risk level babies were re-evaluated before dismissal, and medium risk level babies were treated with phototherapy, respectively.

System usage. The system properly clustered the cases, and linked them under two distinct prototypes. However, it also highlighted the semantic relation between the two prototypes themselves. Actually, the two situations were not identical, since the suggested interventions were different. Nevertheless, in all the subsumed cases the patients were treated as if their risk level was higher than the actual one. Consequently more serious/invasive procedures were applied, with respect to the ones indicated by the default GL thresholds. On the basis of this semantic similarity, the tool finally grouped the two prototypes under a single, higher level prototype, giving birth to a multi-level hierarchy in the case base (see also [36]), which facilitates navigation and comprehension by end users.[2]

[2] Additional details about automatic case base structuring in our system are outside the scope of this paper, but can be found in [37].

Evaluation. The physicians working with us were notified to introduce the changes indicated in the prototypes hierarchy, supported by the set of concrete implementations from which the prototypes were derived. Physicians formally reviewed this result, and finally decided to routinely adopt a modified GL version. In this version all babies with a medium (or high) TSB value at 48 h of age must undergo phototherapy. All babies with (even mild) jaundice, must undergo an additional TSB control/treatment before discharge.

Observe that GL adaptation was strongly motivated by context-awareness, and by the limits of the local situation: as a matter of fact, the new GL appears to be a cost-effective solution in Pavia, but it might not be acceptable in a different setting, due to the increased costs.

In synthesis, the work in [33] appears to be particularly significant in the recent literature panorama, since it deals with the double aim of exploiting contextual information to maintain medical knowledge up to date, and of being context-aware in order to enhance domain knowledge application in practice.

4. Conclusions

Contextual knowledge is naturally interlaced with domain knowledge in a case library. Recent medical informatics literature is increasingly recognizing the importance of contextual knowledge *per* se in supporting human decision making [6]. In this paper, we have analyzed this trend focusing on CBR systems.

In particular, we have identified three main directions (see Section 3) in which contextual knowledge can be explicitly and profitably resorted to: (i) firstly, contextual knowledge can be helpful in reducing the retrieval search space, thus making retrieval itself faster and more meaningful; (ii) secondly, while highly abstracted knowledge cannot be re-interpreted, since the original context from which it was obtained is no longer available, the context-rich knowledge embedded in cases can help revise conclusions and rebuild theories dynamically [6], in order to maintain medical knowledge itself always up to date [30]; (iii) thirdly, contextual information may be very helpful to adapt medical knowledge and reasoning strategies to specific local/personal constraints before applying them in practice.

The work in [33], which proposes a CBR approach for managing non-compliances to GL, can be interpreted as a contribution which aims both at maintaining medical knowledge up to date, and at adapting and enhancing knowledge application on the basis of local constraints. Together with the other works presented in this paper, it thus traces new directions for CBR in medicine research, which, in our opinion, will stimulate significant investigations from now on.

Acknowledgements

The author is grateful to the anonymous reviewers, whose very helpful comments significantly helped her in restructuring and improving the paper.

References

[1] Kolodner JL. Case-based reasoning. San Mateo, CA: Morgan Kaufmann; 1993.
[2] Bichindaritz I, Marling C. Case-based reasoning in the health sciences: what's next? Artificial Intelligence in Medicine 2006;36:127–35.
[3] Montani S. Exploring new roles for case-based reasoning in heterogeneous ai systems for medical decision support. Applied Intelligence 2008;28:275–85.
[4] Bichindaritz I, Kansu E, Sullivan K. Case-based reasoning in care-partner: gathering evidence for evidence-based medical practice. In: Smyth B, Cunningham P, editors. Proc. European workshop on case-based reasoning, LNCS 1488. Berlin: Springer; 1998. p. 334–45.
[5] Stefanelli M. The socio-organizational age of artificial intelligence in medicine. Artificial Intelligence in Medicine 2001;23(1):25–47.
[6] Pantazi SV, Arocha JF, Moehr JR. Case-based medical informatics. BMC Medical Informatics and Decision Making 2004;4:1–23.
[7] Brezillon P. Context in problem solving: a survey. The Knowledge Engineering Review 1999;14:1–34.
[8] Bouquet P, Ghidini C, Giunchiglia F, Blanzieri E. Theories and uses of context in knowledge representation and reasoning. Journal of Pragmatics 2003;35:455–84.
[9] Brezillon P, Pomerol JC. Misuse and nonuse of knowledge-based systems: the past experiences revisited. In: Humphreys P, editor. Implementing systems for supporting management decisions. London: Chapman and Hall; 1996. p. 44–60.
[10] Fischer G. Communication requirements for cooperative problem solving systems. Information Systems 1990;15:21–36.
[11] Walther E, Eriksson H, Musen M. Plug and play: construction of task-specific expert system shells using sharable context ontologies. Technical report. Stanford Education; 1992.
[12] Guha RV. Contexts: a formalization and some applications. Technical report ACT-CYC-423-91. Austin, TX: MCC; 1991.
[13] Giunchiglia F. Contextual reasoning. Epistemologia 1992;345:345–64.
[14] Serafini L, Bouquet P. Comparing formal theories of context in AI. Artificial Intelligence 2004;155:41–67.
[15] Brezillon P. Context-based formalisms and representation levels in decision making. In: Adam F, Brezillon P, Carlsson S, Humphreys P, editors. Creativity and innovation in decision making and decision support. France: Ludic Publishing Ltd.; 2006. p. 523–45.
[16] Bouquet P, Giunchiglia F, VanHarmelen F, Serafini L, Stuckenschmidt H. C-OWL: contextualizing ontologies. In: Fensel D, Sycara K, Mylopoulos J, editors. Proc. international Semantic Web conference, LNCS 2870. Berlin: Springer; 2003. p. 164–79.
[17] Leake DB. Evaluating explanations. USA: Lawrence Erlbaum Associates Inc.; 1992.
[18] Ozturk P, Aamodt A. A context model for knowledge-intensive case-based reasoning. International Journal on Human-Computer Studies 1998;48: 331–55.
[19] Lieber J, Napoli A. Decentralized case-based reasoning for the semantic web. In: Gil Y, Motta E, Benjamins VR, Musen MA, editors. Proc. international Semantic Web conference, LNCS 3729. Berlin: Springer; 2005. p. 142–55.
[20] Third workshop on case based reasoning and context awareness, European conference on case based reasoning (ECCBR); 2008. http://www.iis.uni-hildesheim.de/cacoa08/ [last accessed on March 31st, 2010].
[21] Lorenzi F, Ricci F. Case-based recommender systems: a unifying view. In: Mobasher B, Anand S, editors. Intelligent techniques for web personalization. Springer Verlag; 2005. p. 89–113.
[22] Montaner M, Lopez B, deLaRosa JL. Improving case representation and case base maintenance in recommender systems. In: Craw S, Preece A, editors. Proc. European conference on case based reasoning, LNAI 2416. Berlin: Springer-Verlag; 2002. p. 234–48.
[23] Fesenmaier DR, Ricci F, Schaumlechner E, Wober K, Zanella C. DIETORECS travel advisory for multiple decision styles. In: Frew AJ, Hitz M, OConnors P, editors. Proc. information and communication technology in tourism. Berlin: Springer-Verlag; 2003. p. 232–41.
[24] Bradburn C, Zeleznikow J. The application of case-based reasoning to the tasks of health care planning. In: Wess S, Althoff KD, Richter MM, editors. Proc. topics in case-based reasoning: first European workshop, EWCBR-93. Berlin: Springer-Verlag; 1994. p. 365–78.
[25] Marling C, Whitehouse D. Case-based reasoning in the care of alzheimer's disease patients. In: Aha DW, Watson I, editors. Proc. international conference on case based reasoning, LNCS 1488. Berlin: Springer; 2001. p. 702–15.
[26] Burgers JS, Fervers B, Haugh M, Brouwers M, Browman G, Philip T, et al. International assessment of the quality of clinical practice guidelines in oncology using the appraisal of guidelines and research and evaluation instrument. Journal of Clinical Oncology 2004;22(10):2000–7.
[27] Montani S, Bellazzi R. Supporting decisions in medical applications: the knowledge management perspective. International Journal of Medical Informatics 2002;68(1–3):79–90.
[28] Marling C, Shubrook J, Schwartz F. Case-based decision supportfor patients with type 1 diabetes on insulin pump therapy. In: Althoff KD, Bergman R, editors. Proc. European conference on case-based reasoning, LNAI 5239. Berlin: Springer; 2008. p. 325–39.
[29] Montani S, Portinale L, Leonardi G, Bellazzi R, Bellazzi R. Case-based retrieval to support the treatment of end stage renal failure patients. Artificial Intelligence in Medicine 2006;37:31–42.
[30] Bichindaritz I. Prototypical cases for knowledge maintenance for biomedical CBR. In: Weber R, Richter M, editors. Proc. international conference on case based reasoning, LNCS 4626. Berlin: Springer; 2007. p. 492–506.
[31] Gierl L, Stengel-Rutkowski S. Integrating consultation and semi-automatic knowledge acquisition in a prototype-based architecture: experiences with dysmorphic syndromes. Artificial Intelligence in Medicine 1994;6: 29–49.
[32] Ahmed MU, Begum S, Funk P, Xiong N, VonScheele B. Case-based reasoning for diagnosis of stress using enhanced cosine and fuzzy similarity. Transactions on Case-Based Reasoning on Multimedia Data 2008;1:3–19.
[33] Montani S. Case-based reasoning for managing non-compliance with clinical guidelines. Computational Intelligence 2009;25(3):196–213.

[34] Kofod-Petersen A, Aamodt A. Case-based reasoning for situation-aware ambient intelligence: A hospital ward evaluation study. In: McGinty L, Wilson D, editors. International conference on case based reasoning, LNCS 5650. Berlin: Springer; 2009. p. 450–64.

[35] American Academy of Pediatrics. Clinical practice guideline: management of hyperbilirubinemia in the newborn infant 35 or more weeks og gestation. Pediatrics 2004;114(1):297–316.

[36] Lieber J. Strong, fuzzy and smooth hierarchical classification for case-based problem solving. In: van Harmelen F, editor. Proc. European conference on artificial intelligence. Amsterdam: IOS Press; 2002. p. 81–5.

[37] Montani S. Prototype-based management of business process exception cases. Applied Intelligence. doi:10.1007/s10489-009-0165-z [published online in February 2009].