From Adverse Drug Event Detection to Prevention

A Novel Clinical Decision Support Framework for Medication Safety

V. G. Koutkias[1]; P. McNair[2]; V. Kilintzis[1]; K. Skovhus Andersen[2]; J. Niès[3,4]; J.-C. Sarfati[5];
E. Ammenwerth[6]; E. Chazard[4]; S. Jensen[2]; R. Beuscart[4]; N. Maglaveras[1]

[1]Lab of Medical Informatics, Dept. of Medicine, Aristotle University of Thessaloniki, Thessaloniki, Greece;
[2]Capital Region of Denmark, Copenhagen, Denmark;
[3]MEDASYS, Gif-Sur-Yvette, France;
[4]UDSL EA 2694, Univ Lille Nord de France, Lille, France;
[5]ORACLE, Colombes, France;
[6]Institute for Health Informatics, UMIT – University for Health Sciences, Medical Informatics and Technology,
Hall in Tyrol, Austria
[6]UDSL EA 2694, Univ Lille Nord de France, Lille, France

Keywords

Adverse drug events, clinical decision support systems (CDSS), connectivity, contextualization, interoperability, service-oriented architecture (SOA), medication safety

Summary

Background: Errors related to medication seriously affect patient safety and the quality of healthcare. It has been widely argued that various types of such errors may be prevented by introducing Clinical Decision Support Systems (CDSSs) at the point of care.
Objectives: Although significant research has been conducted in the field, still medication safety is a crucial issue, while few research outcomes are mature enough to be considered for use in actual clinical settings.

In this paper, we present a clinical decision support framework targeting medication safety with major focus on adverse drug event (ADE) prevention.
Methods: The novelty of the framework lies in its design that approaches the problem holistically, i.e., starting from knowledge discovery to provide reliable numbers about ADEs per hospital or medical unit to describe their consequences and probable causes, and next employing the acquired knowledge for decision support services development and deployment. Major design features of the framework's services are: a) their adaptation to the context of care (i.e. patient characteristics, place of care, and significance of ADEs), and b) their straightforward integration in the healthcare information technologies (IT) infra-

structure thanks to the adoption of a service-oriented architecture (SOA) and relevant standards.
Results: Our results illustrate the successful interoperability of the framework with two commercially available IT products, i.e., a Computerized Physician Order Entry (CPOE) and an Electronic Health Record (EHR) system, respectively, along with a Web prototype that is independent of existing healthcare IT products. The conducted clinical validation with domain experts and test cases illustrates that the impact of the framework is expected to be major, with respect to patient safety, and towards introducing the CDSS functionality in practical use.
Conclusions: This study illustrates an important potential for the applicability of the presented framework in delivering contextualized decision support services at the point of care and for making a substantial contribution towards ADE prevention. Nonetheless, further research is required in order to quantitatively and thoroughly assess its impact in medication safety.

1. Introduction

Medication constitutes a major clinical intervention that may yield significant benefits for the patients. Yet, it can also cause considerable harm, if inappropriately used/applied. In the US, adverse drug effects are considered as one of the 10 principal causes of death, which has been estimated to harm at least 1.5 million patients/year, with about 25% of the cases being preventable [1]. Hospital admissions, prolonged hospital stays, and serious economic burdens are additional consequences of the problem [2]. A rather common type of medication-related problems constitutes adverse drug events (ADEs) that are typically defined as "injuries due to medication management rather than the underlying condition of the patient" [3]. ADEs are classified as preventable and non-preventable. Preventable

ADEs are assimilated to "medication errors", while non-preventable ADEs are considered adverse drug reactions (ADRs) that may not be avoidable [4, 5].

Clinical Decision Support Systems (CDSSs) have introduced great promise for reducing ADEs and medical errors in general [6]. Although significant research has been conducted on ADE detection by employing various means, e.g., data review methods and/or machine learning and statistical inference techniques applied to patient data repositories [7], there is difficulty in introducing the outcomes at actual clinical practice in order to support clinical decision making in preventing ADEs. This is due to multiple factors, among which are: a) the lack of reliable and accurate knowledge on ADEs, b) the inability to adapt the discovered knowledge at each specific care context resulting in CDSS over-alerting and alert fatigue [8], and c) the architectural limitations of existing/legacy Health Information Systems (HISs) which hamper the interoperability with external CDSSs. Nevertheless, a comparative international study revealed that the attitude of physicians is positive toward automatic alerting, and there is a clear requirement for introducing CDSSs in clinical practice to enhance medication safety [9].

In this paper, we present a decision support framework targeting medication safety with major focus on ADE prevention. The novelty of the framework lies in the holistic approach followed for its design and development. In particular, we approached the problem by: a) starting from knowledge discovery for ADE detection through the analysis of hospital databases to provide reliable numbers about ADEs per country, hospital or medical unit, and describe their consequences and probable causes, and b) employing this knowledge for ADE prevention via decision support services development and deployment at the point of care, with steps (a) and (b) performed in an iterative fashion. Major design features of the framework's services are: a) their adaptation to the context of care (i.e. patient characteristics, place of care, and significance of ADEs), and b) their straightforward integration in the healthcare information technologies (IT)

environment thanks to the adoption of a service-oriented architecture (SOA) and relevant standards.

The paper is structured as follows. In section 2, we present the materials employed for designing the CDSS framework, i.e. an ADE detection ruleset obtained from data mining hospital stays (i.e. the hospitalization periods of patients) in Electronic Health Record (EHR) systems. In section 3, we present the methods employed for the design and development of the framework and, particularly, its overall architecture and the main design principles followed, the means for appropriately exploiting its services at the point of care emphasizing on connectivity and contextualization aspects, and the methods employed for clinical validation. Section 4 presents the obtained results concerning implementation, deployment, technical and clinical validation, while in section 5 we discuss limitations and future directions of the current work. Finally, section 6 concludes the paper.

2. Materials

The basis for developing the proposed decision support framework was a ruleset of ADE prevention rules, obtained through a tentative knowledge discovery phase involving data mining of 155,447 complete hospital stays. These stays included diagnoses, drug administrations, laboratory examination results and patient characteristics (e.g. age) contained in both structured and free-text records, and were obtained from 6 European hospitals (in Denmark, France and Bulgaria). To facilitate data mining, the respective data were unified under a common data model [10].

The methods for developing the ruleset have been presented thoroughly by Chazard et al. [11]. In summary, the knowledge discovery focused on important drugs, such as those used in coagulation disorders and renal disorders, in order to explore both severe and frequent ADEs and rely on sufficient data for analysis. Various supervised rule induction methods were employed, such as decision tree and association rule mining, in order to trace different kinds of outcomes, e.g., hyponat-

remia and hyperkalemia. The result was a set of production rules of the general form:

$$C_1 \text{ AND } C_2 \text{ AND } \dots \text{ AND } C_n \to E, \quad (1)$$

where C_i, i = [1, n] denote the Boolean conditions of the rule comprising of a variable, an operator and a value. The variables may correspond to: a) groups of drug codes expressed in the ATC (Anatomical Therapeutic Chemical) classification system, b) groups of laboratory examination results expressed in C-NPU/IUPAC (Nomenclature, Properties and Units/International Union of Pure and Applied Chemistry), c) groups of diagnosis codes encoded in ICD-10 (International Classification of Diseases), or (d) patient parameters, e.g. age and gender. In ▶Equation 1, E denotes the potential ADE.

As presented in [11], the discovered rules were filtered and validated from a pharmacological viewpoint, resulting in a final set of 236 rules. These rules have been also evaluated in the context of retrospective ADE detection and reached a positive predictive value of 53.5% (details are provided in section 4). According to a recent review study of ADE detection methods [12], this performance is superior compared to various other methods, since the study concluded that "rule accuracy was variable and often poor with positive predictive value 0.9% – 64%". The rules contain also several statistical features (e.g., the confidence, relative risk, and median delay of outcome appearance) that are meant to be computed per rule and for the medical department of interest. Such features, and especially the confidence of the rule that provides the probability to observe the outcome when the conditions of the rule are met, are employed as contextualization parameters for ADE prevention in the targeted clinical setting.

It is important to highlight the novelty of the discovered knowledge, which lies in the fact that data-mining techniques are able to associate more complex patterns of conditions compared to the rules that are typically found in pharmacovigilance Knowledge Bases (KB). Examples of such conditions are drug discontinuation or absence, and laboratory examination results, based on which the probability of the

ADE occurrence is specified more accurately.

The discovered rules along with the standard terminologies employed for expressing their conditions and outcomes, i.e., ATC, C-NPU/IUPAC, and ICD-10, constituted the core of a KB [13], specifically designed to support the delivery of alerts and recommendations to the clinical personnel for ADE prevention. These shall be generated upon the conduction of a drug-related clinical task (e.g., a new drug prescription or a drug discontinuation), which can be registered/captured in a HIS, such as a Computerized Physician Order Entry (CPOE) or an EHR system.

In order to reduce over-alerting and alert fatigue [8], the alerts and recommen-dations may be adapted according to the care context, i.e., specialized for each patient case and the department/hospital with respect to the significance of the ADE (e.g., based on the local statistical features). To this end, the KB additionally encapsulates higher-level knowledge in the form of processes that indicate and may control whether rules should be considered or not in specific circumstances. This involves, for example, threshold-based filtering as regards the statistical significance of the corresponding triggered rules, in order to determine the most significant alerts or recommendations that will reach the CDSS end-user. A detailed description of the knowledge representation and reasoning scheme is provided in [13].

3. Methods

3.1 Decision Support Framework Architecture

From an architectural viewpoint, the major principle in the design of the proposed framework was the construction of an extensible, sustainable, adaptable, and interoperable solution. The architecture discriminates knowledge engineering and CDSS development from knowledge exploitation and CDSS deployment, respectively. In particular, as depicted in ▶Figure 1, the Global Knowledge Management Platform (GKMP, upper half of ▶Figure 1) is devoted to the construction of a generic/blueprint CDSS as a subsequent step of the knowledge discovery phase, while the
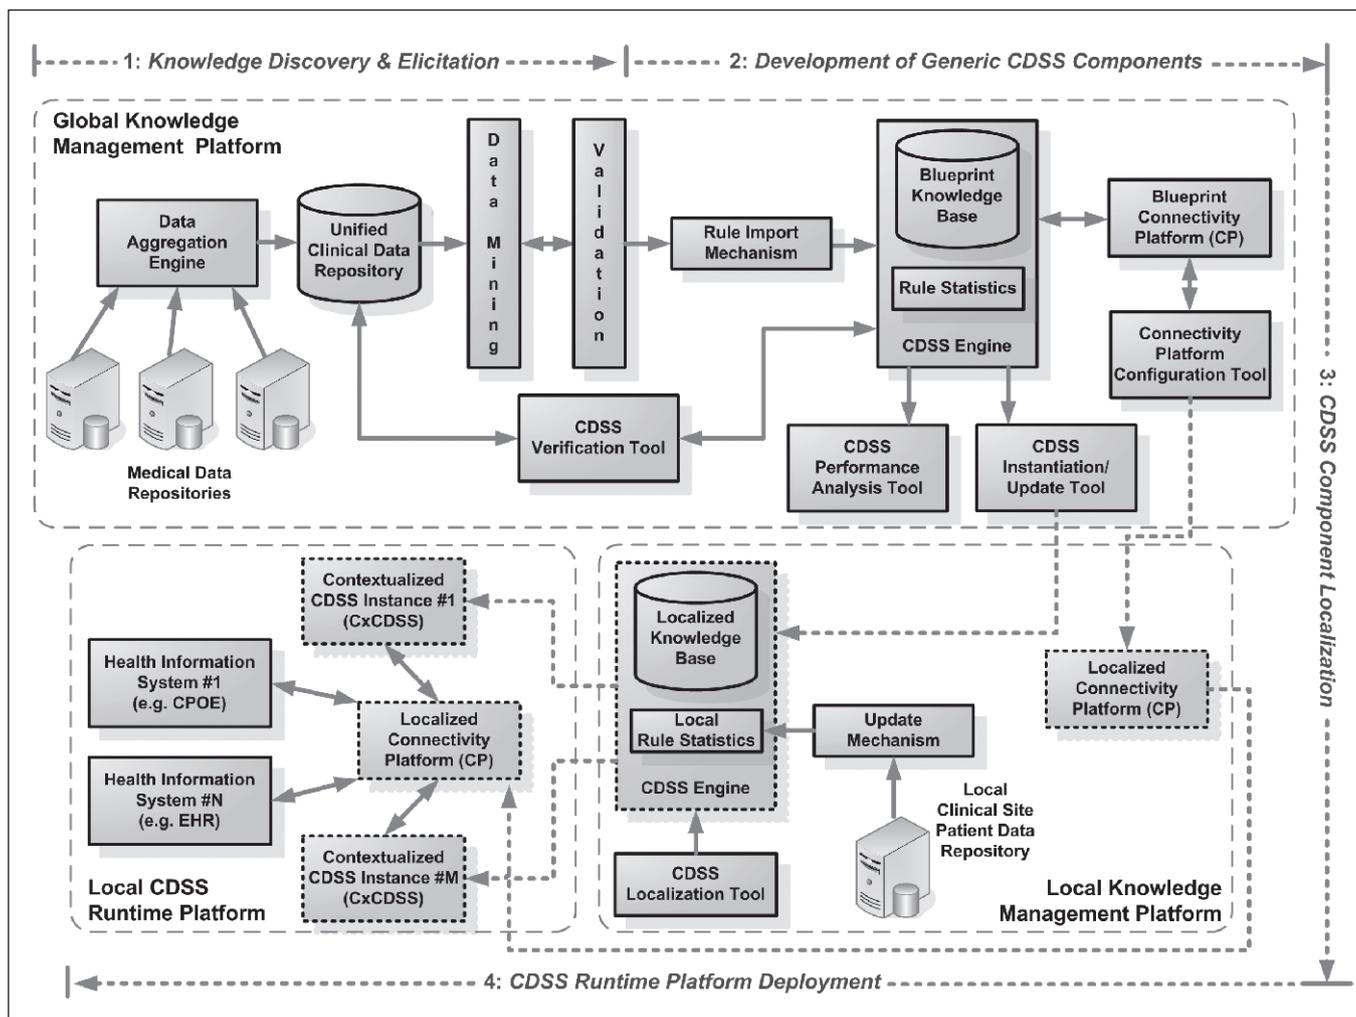
Figure 1 The CDSS framework architecture (linkage among the Global Knowledge Management Platform, the Local Knowledge Management Platform, and the Local CDSS Runtime Platform and their subsequent development is depicted in a clock-wise order, starting from the upper left side)
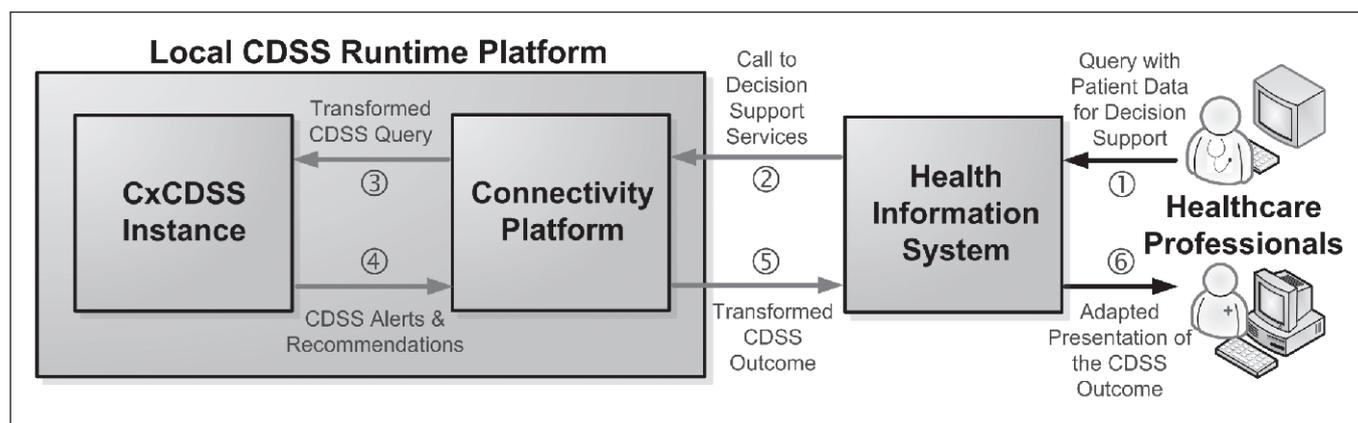
Figure 2   Architecture and communication flow for exploiting the decision support services in the clinical environment

Local Knowledge Management Platform (LKMP, down-right side of ►Figure 1), constitutes an intermediary step to construct a Local CDSS Runtime Platform (LRP, down-left side of ►Figure 1), specifically adapted for the current context, i.e., hospital, department/clinic, along with relevant management tools.

►Figure 1 illustrates the lifecycle for developing the framework (cf. external annotated arrows labeled from 1 to 4 in clockwise order), as well as major components of each platform. In more detail, in the GKMP, data-mining originated rules derived from the formulated Unified Clinical Data Repository (after applying both clinical and statistical validation) are imported in the Blueprint KB of the CDSS Engine. The CDSS Engine performs its reasoning based on the imported rules, meta-rules, and filtering mechanisms encapsulated in the KB, while Rule Statistics denote the statistical significance of rules [13]. Support tools such as the CDSS Verification Tool and the CDSS Performance Tool are available through the GKMP, enabling consistency checking and benchmarking, respectively, before operationalizing the decision support services in the targeted clinical setting. A major part of the GKMP constitutes the Blueprint Connectivity Platform (CP), a software module enabling communication between the CDSS Engine and HISs, as presented in section 3.2.

Appropriate instantiation-update and configuration mechanisms are available for both the CDSS Engine and the CP, in order to construct specialized versions of these components for the local clinical setting in the respective LKMP. The LKMP offers also the CDSS Localization Tool, which enables configuring the operation of the Localized CDSS Engine, according to the particular requirements/preferences that may be applicable (e.g., selection of rules to be deployed). In addition, the statistics of the incorporated rules are updated according to the Local Clinical Site Patient Data Repository through a relevant mechanism.

In the context of LRP, using the tools offered by the LKMP, Contextualized CDSS Instances (CxCDSS) may be constructed, which enable medication-related IT systems of the local environment to access the decision support services offered through the Localized CP.

For the overall design of the framework components, the Reference Model of Open Distributed Processing (RM-ODP) was employed as a modeling methodology [14], as it constitutes a well-established standard in the field offering several advantages compared to alternative approaches [15].

3.2 CDSS Connectivity with the Healthcare Enterprise

Considering the technical heterogeneity among HISs, the development of a generic CP offers an interoperable solution for exploiting the services offered by the CDSS. ►Figure 2 provides a detailed view of the LRP part (►Figure 1) of the overall architecture. The CP acts as the mediator between the LRP and HISs, providing a systematic, controlled, and robust linkage with the decision support services offered. The main design principles for realizing the CP lie in its generalization and openness, so as to support the connectivity among a variety of HISs and (potential) CDSS Engines, independently of their underlying technical characteristics. In addition, the CP is capable of managing potential communication failures among the interacting parties.

The above features that are offered by the CP are achieved through the adoption of communication standards and technologies that are employed to realize a robust SOA [16]. SOA offers a loosely-coupled computing paradigm that has become a key ingredient of modern business applications and IT infrastructures. For the healthcare IT sector, the added value and the potential of SOA have been highlighted in several studies [17, 18].

►Figure 2 illustrates the entire workflow (steps (1) to (6)) that takes place upon the submission of requests for decision support by healthcare professionals, posed implicitly/explicitly through their interaction with a HIS (step (1)). Specifically, the HIS formulates a request encapsulating the patient data for assessing the clinical case that is transmitted to the CP via a Web service invocation (step (2)) expressed in its local format. The CP identifies the "appropriate" CxCDSS instance, transforms the request posed to the respective CxCDSS format (step (3)), and routes the message to the CxCDSS instance. The CxCDSS instance, in turn, assesses the clinical case and generates potential alerts (accompanied with recommendations for actions), according to the clinical context, e.g., the statistical significance of the alerts in the

local setting, and runtime configuration options. These alerts are returned to the CP in the CxCDSS response format (step (4)), and the CP routes the message to the HIS, after its transformation to the HIS response format (step (5)). Finally, the HIS provides the CxCDSS outcome to the healthcare professionals by handling the relevant information presentation issues (step (6)).

Interestingly, the "many-to-many" relationship among CxCDSSs and HISs that the CP supports (▶ Figure 1, LRP part) has been considered for a hospital or healthcare region in a scenario where different CxCDSS instances are exploited by HISs operating in different clinics. Overall, the proposed approach provides a scalable and extensible solution, which may be capable of growing sustainably over time.

3.3 Approach for CDSS Contextualization

Contextualization is considered as a key attribute for eliminating CDSS over-alerting, and reinforcing user-acceptance in the field of medication safety [9, 19]. In general, contextualization of decision support services aims to address the delivery of "the right information, at the right time, to the right person, at the right place, and in the right format". A model for information contextualization of medication-related decision support services has been introduced by Niès et al. [20], proposing four major categories that have to be elaborated, namely *Environment*, *Tasks*, *Users*, and *Temporal Aspects*. In the scope of this work this model was applied and its dimensions involve:

- *Environment*: country and language characteristics, significance of the ADEs in the specific clinical setting, and epi-

demiological features such as risk factors of the particular population.
- *Tasks*: the tasks of prescription, dispensation, administration, and compliance (PDAC) of the medication chain, performed by healthcare professionals.
- *Users*: presentation of decision support information according to their expertise and role (i.e., junior and senior physicians, nurses, pharmacists).
- *Temporal Aspects*: Description of sequences involving the above three descriptors that could lead/contribute to ADEs.

The above dimensions were analyzed to identify how they should be elaborated for the realization of the proposed decision support framework by considering both technical (i.e., architecture, performance, and technology-related) and organizational aspects (i.e., procedures taking place in clinical settings, as well as human factors). The analysis outcome reflects the following principles that were adopted in the implementation phase: a) the CDSS runtime handles *Environment* and *Time*-related aspects of contextualization, and b) the HIS component requesting access to the CDSS services handles the *Tasks* and *Users* contextualization dimensions. To keep its connectivity role intact, the CP handles routing and data transformation; however, it allows contextual parameters (e.g., statistical thresholds for rule filtering) to be passed to the CDSS through the respective request.

An approach for the contextualization of the adopted rules for ADE prevention is illustrated via the following example. Let us consider the 182nd rule from the ruleset:

R_{182}: *CRI & NSAID & no K^+ sparing diuretic → hyperkalemia,*   (2)

where *CRI* stands for chronic renal insufficiency, *NSAID* for a non-steroidal anti-inflammatory drug and K^+ for potassium. According to this rule, patients who suffer from renal insufficiency and who receive an NSAID but no potassium sparing diuretics could experience hyperkalemia, with a probability given by the "confidence" of the rule. The first condition is related to the medical condition of the patient, the second is related to a drug the patient is taking, and the third to a drug the patient does not take.

The confidence C_{182} of the above rule is defined as the probability of facing hyperkalemia given that the conditions of the rule are met, i.e.:

C_{182} = *P(hyperkalemia | CRI & NSAID & no K^+ sparing diuretic).*   (3)

C_{182} is computed using historical clinical data for the clinical setting of interest. In particular, ▶ Table 1 provides the value of C_{182} across six European hospitals providing data in our study and for the medical departments of hospital #4. In this example, the confidence differs significantly between the hospitals, but it is not significantly different among the units of hospital #4.

In order to cope with over-alerting, a threshold T can be applied on C_{182} to filter the alerts that R_{182} may generate when its conditions are met. This is performed according to the following simple logic: *If C_{182} exceeds the value of T ($C_{182} \geq T$) or if the situation never occurred before (i.e. C_{182} = 0/0) in the current clinical setting, then alerts from R_{182} are generated by the respective CxCDSS; otherwise the CxCDSS*

Table 1   Comparison of C_{182} estimations: (a) in six hospitals, (b) in the departments of hospital #4

(a)	**Hospital:**	**#1**	**#2**	**#3**	**#4**	**#5**	**#6**	*p value*
	C182 value:	0/56 = 0%	3/174 = 1.7%	3/34 = 8.8%	68/703 = 9.7%	5/43 = 11.6%	0/2 = 0%	0.0061
(b)	**Department:**	**Cardiology**	**Geriatric**	**Obstetrics & Gynecology**	**Internal Medicine**	**Pulmonology**	**Surgery**	*p value*
	C182 value:	27/255 = 10.6%	6/43 = 14%	0/0	23/253 = 9.1%	15/115 = 13%	3/47 = 6.4%	0.818

Remark: If a patient is hospitalized in multiple departments of the same hospital, then for the computation of the confidence the stay is taken into account in all the respective departments. However, such a case is counted only once for the computation of the confidence for the entire hospital.

Table 2   CxCDSS output when applying an 8% threshold to C_{182} (Table 1), if all the conditions of R_{182} are met: (a) in six hospitals, (b) in the departments of hospital #4

(a)	Hospital:	#1	#2	#3	#4	#5	#6
	Output of R_{182}	No alert	No alert	Alert	Alert	Alert	No alert
(b)	Department:	Cardiology	Geriatric	Obstetrics & Gynecology	Internal Medicine	Pulmonology	Surgery
	Output of R_{182}	Alert	Alert	Alert	Alert	Alert	No alert

suppresses the alerts concerning R_{182}. For instance, if the conditions of R_{182} are matched and T is set to 8% for all the hospitals and departments referred in ►Table 1, then the outcomes (i.e. alerting or not for R_{182}) are those presented in ►Table 2. Of course the selection of T for each rule may be set according to the local clinical setting.

Besides employing the local statistical features of rules as illustrated in the above example, contextualization of the alerting scheme is also based on a number of meta-rules and filtering mechanisms (details provided in [13]).

3.4 Clinical Validation

The clinical perspective was considered of outmost importance for the design and the acceptance of the proposed decision support framework. As illustrated in [9], clinicians really need automatic medication alerting systems, but they are skeptical concerning the appropriateness of the alerts produced.

To assess the positive predictive value of the ADE detection rules, a sample of 24,753 patient records were explored to automatically detect possible ADEs and especially cases of possible drug-associated hyperkalemia. Hyperkalemia was selected as it presents a severe and potentially life-threatening situation, with around three cases per month noticed in the departments of the participating hospitals. A manual expert review by clinicians and pharmacists was then performed to identify whether the automatically detected hyperkalemia cases were in fact actual ADEs. From this, positive predictive value and sensitivity were calculated.

Furthermore, realistic test cases were systematically developed, aiming to assess correctness, completeness, and under-standability of the CDSS alerting functionality per se. Nine clinicians from hospitals participating in the study from three countries were invited to develop test cases using a predefined template and a Web-based repository. The clinicians were not informed of the architecture, functionality and capabilities of the CDSS and were not directly involved in the development of its Knowledge Base. The approach used to identify test cases varied between clinicians and included analysis of national guidelines on medication use, official drug information, reports of actual ADEs from error reporting systems, complex patient cases and review of scientific literature on typical problems.

The test cases covered especially anticoagulation treatment due to the high prevalence of this problem especially in elderly populations. Each test case contained a description of the patient case, including age, medical history, allergies, recent or planned treatment and medications, and the alert/warning the underlying clinical case should generate (e.g., warning on increased bleeding risk). All the proposed test cases were then reviewed by five other clinicians not being involved in the test cases definition to assess clarity and correctness, with a test case at least reviewed by two experts. In case of disagreement that could not be resolved through discussion, the test case was excluded from the repository. Validation via the test cases was launched as soon as the first versions of the CDSS runtime were made available, in order to obtain feedback as early as possible concerning the adopted approach and the exploitation of the discovered knowledge. The generated output was assessed by two expert clinicians from two different countries to assess whether the expected alert (as defined in the test case description) was generated, and whether the alert was clinically correct.

Disagreements were resolved by discussion in joint workshops.

4. Results

4.1 Implementation

The CDSS runtime was implemented using Gaston [21], an IT tool developed for building tailor-made DSSs for the healthcare domain, which was selected by analyzing various candidates with respect to the requirements defined. The CDSS Engine operates as a server receiving requests for data analysis encoded in a specifically-defined XML request format, in which the actual clinical data are encapsulated following the common data model [10], and generating an XML response message.

The CP was developed using Oracle® SOA Suite [22], and Oracle® BPEL (Business Process Execution Language) technology. As the CP aims to provide transformation services, it has been employed to map the format of the requests received from various HISs into the CDSS XML request. While the HIS requests are product-specific, the CDSS responses are transformed uniformly by the CP into the CAP (Common Alerting Protocol) format [23], a standard not specifically targeting the healthcare IT domain, but adequate for our aim to enable generalization of the proposed approach. ►Figure 3 illustrates an excerpt of an example CDSS response message encoded in CAP. As remarked in section 3.2, the CP does not handle presentation issues for the end-user, as regards the information obtained from the CDSS, i.e., the messages' language, content, and level of detail for describing the alerts, as well as the way of providing the alerts to the user. This remains a responsibility of the decision support requester HIS

(e.g., CPOE or EHR, step (6) depicted in ▶Figure 2).

The CDSS response messages incorporate the following information per alert (an example is presented in ▶Figure 3 for the rule presented in section 3.3):

- the (unique) identifier of the respective rule;

```
<alert>
  <cap:identifier>R182</cap:identifier>
  <cap:sender>PSIP</cap:sender>
  <cap:sent>2012-02-02T09:29:51+01:00</cap:sent>
  <cap:status>Actual</cap:status>
  <cap:msgType>Alert</cap:msgType>
  <cap:source>CxCDSS-3</cap:source>
  <cap:scope>Public</cap:scope>
  [...]
  <cap:info>
   [...]
   <cap:headline>Hyperkalemia (K+>5.5) </cap:headline>
   <cap:description>Diag: chronic renal insufficiency is
True;ATC_DRUG_ M01AA01 is True;ATC_DRUG_ C03D is False
</cap:description>
    [...]
    <cap:parameter>
     <cap:valueName>Short_Physician</cap:valueName>
     <cap:value>Nonsteroidal Anti-Inflammatory drugs can cause a
hyperkalemia.
     </cap:value>
    </cap:parameter>
    <cap:parameter>
     <cap:valueName>Long_Physician</cap:valueName>
     <cap:value>Non Steroid Anti-Inflammatory drugs are nephrotoxic by
inhibiting prostaglandin I2, which regulate renin-angiotensin system. The
probability of hyperkalemia is altered by renal status and by potassium
sparing drugs.
In case of treatment by NSAIDs, clinical and biological monitoring should be
enhanced.
Ref.: Martindale - The Complete Drug Reference – 36th Ed., The
Pharmaceutical Press, London 2009:1-134.  </cap:value>
    </cap:parameter>
    <cap:parameter>
     <cap:valueName>Confidence</cap:valueName>
     <cap:value>0.097</cap:value>
    </cap:parameter>
    [...]
  </cap:info>
  <cap:info>
   [...]
  </cap:info>
</alert>
```

Figure 3   Sample alert encoded in the CAP format, corresponding to the example rule presented in section 3.3

- the effect that the rule is associated with, e.g., hyperkalemia;
- textual descriptions of the alert in various levels of detail (short/long description) for the involved healthcare professional types (physician, nurse, etc.), and in different languages (currently, Bulgarian, Danish, English, Greek, and French are supported, while a straightforward procedure has been established to support further languages);
- statistical features denoting the significance of the alert in the specific context, e.g., the confidence value (probability of having the ADE effect knowing that the conditions of its occurrence are met), the support (probability of having the ADE effect and matching the conditions of its occurrence at the same time), the Fisher test p-value, etc., and
- the conditions that triggered the rule.

The various tools of the CDSS platform, e.g., the CDSS Localization Tool, the CDSS Verification and Performance Tool, etc., have been developed as Java™ applications. ▶Figure 4 presents a screenshot of the CDSS Localization Tool, illustrating the wizard-like procedure that is offered for generating a particular CxCDSS instance for a specific clinical setting. Configuration options involve setting up thresholds to define the statistical significance of the ADE prevention rules, manually enabling/disabling rules no matter their statistical significance is, etc.

4.2 Deployment and Technical Validation

The decision support framework has been incrementally developed following an iterative procedure, according to data availability and the subsequent releases of source knowledge [11]. The procedure involved rapid prototyping, deployment, and validation via test cases concerning primarily communication and performance issues, as well as the accuracy of the CDSS outcome.

In particular, an extensive evaluation of the LRP response time has been performed, taking into account a variety of requests with respect to the amount of data

Figure 4   Using the step-by-step wizard of the CDSS Localization Tool to configure a specific Contextualized CDSS (CxCDSS) instance: the step of reviewing ADE prevention rules and setting up statistical significance filtering levels is depicted.

considered and the underlying complexity (e.g., number of rules included in the ruleset, data payload, etc.). ►Figure 5 presents a scatter-plot of the response time with respect to the data request size, along with a histogram depicting the distribution of classifications across response time. As remarked in [13], the response time varies according to the inference procedure applied by the CDSS Engine, e.g., whether the encapsulated filtering mechanisms are enabled or not. Equally important, the response time depends critically on the CPU processing power. The outcomes presented in ►Figure 5 (response time 6 – 8 sec in the majority of the classifications) correspond to the full-feature scenario, in which the CDSS Engine employs all its internal post-processing and filtering mechanisms on

the entire ruleset, with the entire LRP deployed in a medium-class notebook computer (CPU Intel® i3).

An automated verification procedure for the CDSS runtime has been also iteratively applied aiming to identify whether the implemented rules are consistent with the input provided by the knowledge discovery techniques. For this purpose, the CDSS responses for the same clinical database used in the knowledge discovery phase, as well as a list of hospital stays fitting each rule conditions were used as input. The outcome of this procedure has been assessed via contingency tables [24], comparing the identities of the CDSS rules fired on each hospital stay with its counterpart from the data mining. This procedure resulted in the optimization of the ruleset,

in order to address more adequately the requirements of a prospective decision support operation.

The generic design of the CP makes it independent of the CDSS Engine and the HIS side (►Figure 2). Due to this approach, the specifications defined enable vendors of either CDSS Engines or HISs to plug-in their systems in the CP. As a proof-of-concept, DxCare® (a commercial EHR system from Medasys®, France), EPM® (a CPOE system from IBM®, Denmark), and a Web-based prototype (independent of any existing healthcare IT product) have been successfully employed to communicate with the LRP [25], by deploying various CxCDSS instances simultaneously. In this respect, we have successfully implemented the mapping from HL7 CDA (Clinical

Document Architecture) format [26], and a proprietary format that the above products require/support to the CP internal request format.

Links to demonstrators (live or videos) of the above prototypes are accessible in [27], in which different perspectives were adopted to handle presentation issues of the decision support outcome.

4.3 Clinical Validation

In the sample of 24,753 patient records that was analyzed, 997 possible ADEs (i.e., 4% of all hospitalizations) were detected by the ADE rules, including 507 cases of possible drug-associated hyperkalemia. The manual expert review confirmed that 271 of these hyperkalemia cases were in fact actual ADEs. Thus, the positive predictive value of the ADE rules in this sub-study was found to be 53.5% and the sensitivity was 95.1% [20]. These figures for ADE detection are superior to many other methods presented in the literature [12]. We need to remark though that, since the performance of ADE detection methods varies significantly according to the data being analyzed and the events that are being investigated, comparisons among detection methods without employing a common basis are indicative.

In addition, as remarked in section 3.4, 24 test cases were analyzed by the CDSS and the generated output was assessed by two expert clinicians. Inter-rater agreement before discussing and resolving all disagreements was good (0.74). 71% of the test cases showed the expected alerts; 40% of the alerts were judged as clinically correct; in 31% of the alerts, the experts could not judge the clinical correctness due to missing clinical information. The difficulties faced by the experts in the review procedure were mainly due to the fact that the alerts produced were based on data-mining methods which are able to statistically associate quite complex patterns of conditions compared to what can be typically met in pharmacovigilance KBs.

Overall, based on the detailed analysis of the evaluation output, a comprehensive set of recommendations for improvement of the decision support output were made, including wording of alerts, grouping of alerts, adding more explanations for the user, and providing recommendations for actions in response to the alert. The majority of these recommendations have been addressed in ongoing iterations of the CDSS development.

5. Discussion

The potential of IT in the field of patient safety has been illustrated in various studies [28]. In particular, decision support for medication safety has been recognized as a significant research topic [1], and various systematic efforts have been reported in the literature especially focusing on integrating decision support in CPOE systems [7]. However, due to the complexity of medication management and the diverse types of associated safety risks, medication-related decision support needs to be integrated in various HISs, beyond CPOE sys-
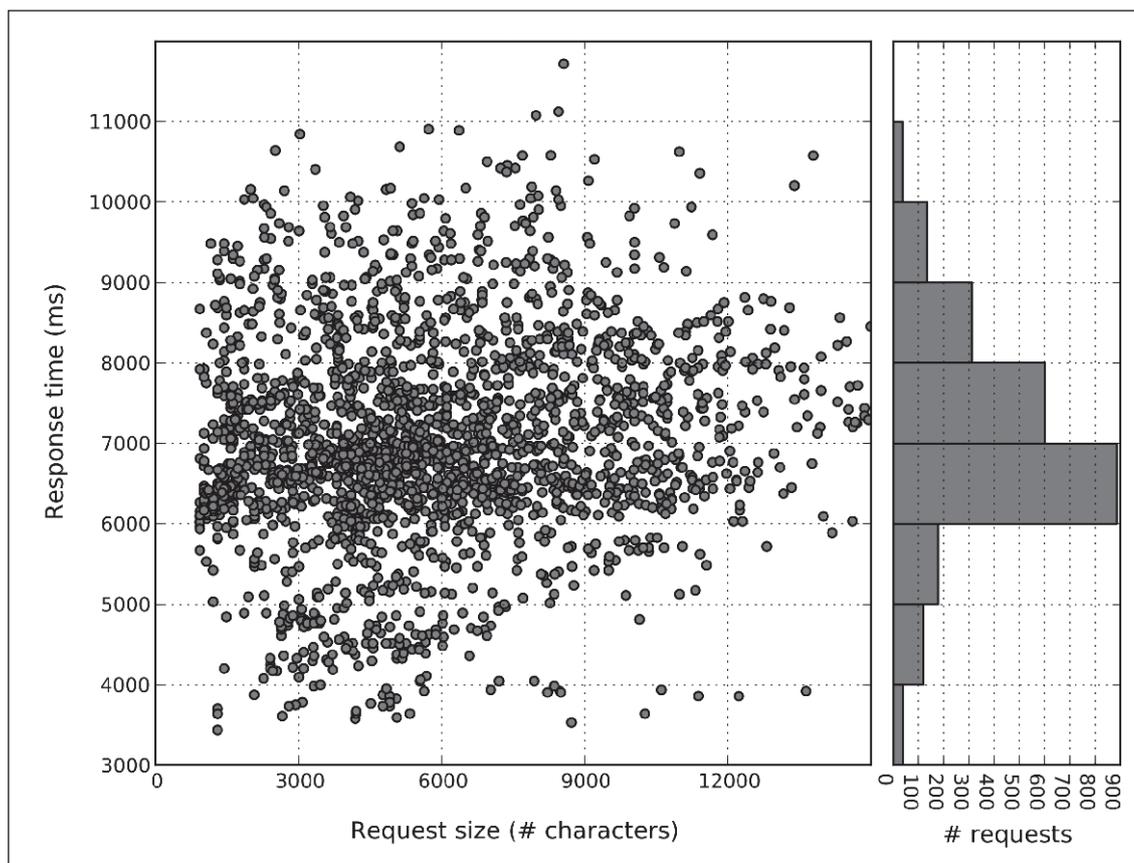
Figure 5
Left-side: Scatter-plot depicting the CDSS response time with respect to the data request size; right-side: histogram depicting the distribution of classifications across the response time

tems. Various methods have been developed for detecting potential drug risks [6], but still there is a gap in employing the discovered knowledge in actual clinical settings for harm prevention through CDSSs. This shortcoming stems from multiple factors, among which the architectural limitations of existing HISs hampering integration aspects, the inability to adapt the discovered knowledge at the local context, and the overall inherent complexity of the clinical environment.

The aim of the current work was the design and development of a CDSS framework targeting medication safety through: a) identification of drug safety risks and provision of reliable numbers regarding their prevalence per hospital or medical unit, b) contextualization of the decision support services, in order to address the requirements of the local clinical setting that these shall be offered, c) seamless connectivity with medication-related HISs through an interoperable design, and d) maintenance and extension of knowledge and system components. The application focus of the presented framework was preventable ADEs, an important subset of medication safety risks [1–5].

Overall, the design of the presented framework illustrates a holistic approach for the construction and delivery of decision support services (exceeding the application scope of medication safety), starting from the knowledge discovery phase (GKMP), continuing with the adaptation of the services for the local environment through contextualization (LKMP), and resulting at the actual deployment and use of the services (LRP). The knowledge discovery phase that was elaborated discriminates our work from other CDSSs for ADE prevention which rely only on expert knowledge obtained from clinical guidelines and medical reference books like ADEAS [29].

The presented framework is aligned with other SOAs that have been proposed for clinical decision support, such as SANDS [30]. SANDS defines a set of interfaces that a decision support service should make available, leaving the choice of knowledge representation up to the implementer. As we focused on ADE prevention, we had to first seek for reliable knowl-

edge concerning the prevalence of ADEs and then proceed with the construction of decision support services for exploiting this knowledge at the point of care. Thus, our framework addresses also the issue of knowledge discovery and management towards contextualized decision support. Equally important, we elaborated on knowledge maintenance and update by providing dedicated tools and mechanisms to support the iterative process from knowledge discovery to decision support deployment, enabling this way the overall sustainability of the framework. On the other hand, SANDS supports scenarios where several disparate decision support services are needed to synthesize a decision, which have not been elaborated in the current work, although the proposed architecture enables multiple HISs to connect to multiple CxCDSS instances through the CP.

For a successful framework implementation, we share the same experience reported for SANDS [30], concerning the importance of prototype development, in order to identify and address challenges and special requirements which are not anticipated at the design level. Our prototypes enabled us to develop a fully functional and technically efficient architecture.

The accuracy of the employed methods for ADE detection, along with the fact that the adopted rules are able to associate more complex patterns of conditions (e.g. drug discontinuation or absence, and laboratory examination results) compared to the rules that are typically found in pharmacovigilance KBs, provides the basis for developing novel and valuable decision support services.

Furthermore, contextualization of decision support has been introduced as a mechanism to confront (besides the overall adaptation of the provided services and) the issue of over-alerting. Over-alerting may be defined as "sending an alert to prevent an outcome although the probability of its occurrence is too low or null", and it is crucial for the effective operation and user acceptance of medication-related CDSSs [8, 9]. In the current implementation, through the CxCDSS instances that may be deployed, this issue is being dealt with: a) the selection of the exact rules of interest for the local setting, e.g., related with co-

agulation disorders, b) the calculation and use of the local statistical features for the selected rules to assess their statistical significance, and c) a number of meta-rules and filtering mechanisms [13].

In the scope of this work, obtaining true alerts was the main priority. As the empirical confidence of the rules has been estimated in real data in the clinical settings considered (hospital or department), although a rule is always applicable from a pharmacological viewpoint, the empirical probability of experiencing the outcome knowing that the conditions of the rule are met, enables to "silent" the CDSS in case the probability is too low. Although the effect on alert fatigue has not been empirically evaluated, the adopted contextualized approach reduces over-alerting and could contribute in reducing alert fatigue. The CDSS Localization Tool that has been developed as part of the LKMP enables the flexible configuration of the relevant thresholds as required in the respective application setting. Machine learning strategies could be employed to support/automate the selection of such thresholds. Further experimentation on the direction of handling contextualization aspects through the CP, so as to extend its transformation and routing scope, constitutes an additional future work direction.

Concerning connectivity, the CP enables "plugging" new systems in the architecture without redeveloping the entire process. Following this approach, each stakeholder (either CDSS or HIS vendor) interested in participating in the deployment of the proposed architecture needs to be aware of how to communicate with the CP, rather than how to communicate with the other systems/parties. This is possible thanks to the well-defined Web service interfaces that are available through the adopted SOA, and standards-based data communication protocols.

The CDSS response time depends heavily on the inference procedure employed, e.g., whether the various processing mechanisms offered by the CDSS Engine are enabled or not [13], as well as on the available computing power. The significance of the delay recorded in our benchmark test relies on the mode employed for operating the decision support services,

i.e., reactive in the case of the EPM® prototype, or proactive in the case of the DxCare® prototype [25]. Nevertheless, performance improvement is an issue that needs to be further elaborated (especially for the reactive operation), concerning the inference part as well as the deployment of the LRP in more advanced computational infrastructures such as cloud-based.

The developed proof-of-concept prototypes, their integration in real hospital environments, as well as the technical and clinical validation that we conducted, illustrate the potential and applicability of the presented framework in delivering contextualized decision support services at the point of care to contribute in ADE prevention. However, it is important to conduct long-term evaluation of the decision support services through appropriate field studies, in order to quantitatively and thoroughly assess their impact in clinical practice and medication safety.

6. Conclusion

This paper presented a systematic approach towards the design, development, and deployment of decision support services for ADE prevention. It approached the problem by first seeking for reliable knowledge concerning the prevalence of ADEs and then proceeded with the construction of knowledge-based CDSSs contextualized and exploitable at the point of care. Our results concerning our prototypes, i.e., their performance and integration with real healthcare IT products, introduce an important potential for the applicability of the presented framework in delivering contextualized decision support services at the point of care to contribute in ADE prevention. Nonetheless, further research is required in order to assess its use in clinical practice and, ultimately, its impact in medication safety.

Acknowledgments

References

1. Institute of Medicine. Preventing Medication Errors. The National Academic Press, US, 2007.
2. Schachter M. The epidemiology of medication errors: how many, how serious? Br J Clin Pharmacol 2009; 67 (6): 621–623.
3. Morimoto T, Gandhi TK, Seger AC, Hsieh TC, Bates DW. Adverse drug events and medication errors: detection and classification methods. Qual Saf Health Care 2004; 13 (4): 306–314.
4. Aronson JK. Medication errors: definitions and classification. Br J Clin Pharmacol 2009; 67 (6): 599–604.
5. Nebeker JR, Hurdle JF, Hoffman J, Roth B, Weir CR, Samore MH. Developing a taxonomy for research in adverse drug events: potholes and signposts. J Am Med Inform Assoc 2002; 9 (Nov–Dec suppl): S80–S85.
6. Murff HJ, Patel VL, Hripcsak G, Bates DW. Detecting adverse events for patient safety research: a review of current methodologies. J Biomed Inform 2003; 36 (1–2): 131–143.
7. Kuperman GJ, et al. Medication-related clinical decision support in computerized provider order entry systems: a review. J Am Med Inform Assoc 2007; 14 (1): 29–40.
8. van der Sijs H, van Gelder T, Vulto A, Berg M, Aarts J. Understanding handling of drug safety alerts: a simulation study. Int J Med Inf 2010; 79 (5): 361–369.
9. Jung M, Hoerbst A, Hackl WO, Kirrane F, Borbolla D, Jaspers MW, Oertle M, Koutkias V, Ferret L, Massari P, Lawton K, Riedmann D, Darmoni S, Maglaveras N, Lovis C, Ammenwerth E. Attitude of physicians towards automatic alerting in Computerized Physician Order Entry systems: A comparative international survey. Methods Inf Med 2013; 52 (2): 99–108.
10. Chazard E, et al. Detection of adverse drug events: proposal of a data model. Stud Health Technol Inform 2009; 148: 63–74.
11. Chazard E, Ficheur G, Bernonville S, Luyckx M, Beuscart R. Data mining to generate adverse drug events detection rules. IEEE Trans Inf Technol Biomed 2011; 15 (6): 823–830.
12. Forster AJ, Jennings A, Chow C, Leeder C, van Walraven C. A systematic review to evaluate the accuracy of electronic adverse drug event detection. J Am Med Inform Assoc 2012; 19(1): 31–38.
13. Koutkias V, et al. Knowledge engineering for adverse drug event prevention: on the design and development of a uniform, contextualized and sustainable knowledge-based framework. J Biomed Inform 2012; 45 (3): 495–506.
14. ISO/IEC. 10746-1:1998 ITU-T Recommendation X.901, Open Distributed Processing – Reference Model – Part 1: Overview.
15. Chen D, Doumeingts G, Vernadat F. Architectures for enterprise integration and interoperability: Past, present and future. Comput Ind 2008; 59 (7): 647–659.
16. Singh MP, Huhns MN. Service-Oriented Computing: Semantics, Processes, Agents. New York: Wiley, 2005.
17. Kawamoto K, Lobach DF. Proposal for fulfilling strategic objectives of the U.S. roadmap for national action on decision support through a service-oriented architecture leveraging HL7 services. J Am Med Inform Assoc 2007; 14 (2): 146–155.
18. Katehakis DG, Sfakianakis SG, Kavlentakis G, Anthoulakis DN, Tsiknakis M. Delivering a lifelong integrated electronic health record based on a service oriented architecture. IEEE Trans Inf Technol Biomed 2007; 11(6): 639–650.
19. Ammenwerth E, Hackl WO, Riedmann D, Jung M. Contextualization of automatic alerts during electronic prescription: researchers' and users' opinions on useful context factors. Stud Health Technol Inform 2011; 169: 920–924.
20. Niès J, et al. Information contextualization in decision support modules for adverse drug event prevention. Stud Health Technol Inform 2011; 166: 95–104.
21. de Clercq PA, Hasman A, Blom JA, Korsten HHM. Design and implementation of a framework to support the development of clinical guidelines. Int J Med Inform 2001; 64 (2–3): 285–318.
22. Oracle SOA Suite, Available at: http://www.oracle.com/us/products/middleware/soa/ Last access: February 19, 2014.
23. OASIS Emergency Management Technical Committee. Common Alerting Protocol Standard, v. 1.2 (CAP-V1.2). July 2010. Available at: http://docs.oasis-open.org/emergency/cap/v1.2/CAP-v1.2-os.pdf, Last access: February 19, 2014.
24. Brender J, Talmon J, Egmont-Petersen M, McNair P. Measuring quality of medical knowledge. Proceedings of MIE, 1994. pp 69–74.
25. Bernonville S, et al. Three different cases of exploiting decision support services for adverse drug event prevention. Stud Health Technol Inform 2011; 166: 180–188.
26. Dolin RH, et al. HL7 clinical document architecture, release 2. J Am Med Inform Assoc 2006; 13 (1): 30–39.
27. The PSIP (Patient Safety through Intelligent Procedures in medication) Prototypes Portal. Available at: http://psip.univ-lille2.fr/prototypes/public/ Last access: February 19, 2014.
28. Koutkias VG, Niès J, Jensen S, Maglaveras N, Beuscart R. Patient Safety Informatics. Stud Health Technol Inform 2011; 166, IOS Press.
29. Rommers MK, et al. Development of a computerised alert system, ADEAS, to identify patients at risk for an adverse drug event. Qual Saf Health Care 2010; 19 (6): e35.
30. Wright A, Sittig DF. SANDS: a service-oriented architecture for clinical decision support in a National Health Information Network. J Biomed Inform 2008; 41 (6): 962–981.