The ADE Scorecards: a tool for Adverse Drug Event detection in Electronic Health Records

Emmanuel CHAZARD [a,1] Adrian BĂCEANU [b] Laurie FERRET [c,d], Grégoire FICHEUR [a]
[a] *Department of Medical Information and Archives, CHRU Lille; Univ Lille Nord de France; EA 2694; Lille, France*
[b] *IDEEA Advertising, Romania*
[c] *Central dispensary, CH Denain, France*
[d] *CERIM, Univ Lille Nord de France; EA 2694; Lille, France*

Abstract. Although several methods exist for Adverse Drug events (ADE) detection due to past hospitalizations, a tool that could display those ADEs to the physicians does not exist yet. This article presents the ADE Scorecards, a Web tool that enables to screen past hospitalizations extracted from Electronic Health Records (EHR), using a set of ADE detection rules, presently rules discovered by data mining. The tool enables the physicians to (1) get contextualized statistics about the ADEs that happen in their medical department, (2) see the rules that are useful in their department, i.e. the rules that could have enabled to prevent those ADEs and (3) review in detail the ADE cases, through a comprehensive interface displaying the diagnoses, procedures, lab results, administered drugs and anonymized records. The article shows a demonstration of the tool through a use case.

Keywords. Adverse Drug Events, Adverse Drug Reactions, data mining, Electronic Health Records

Introduction

The Institute Of Medicine defines ADEs as "injuries due to medication management rather than the underlying condition of the patient" [1]. That definition emphasizes that ADEs are due to a combination of causes, including drugs (drug administration, dose variations, and drug discontinuations) and characteristics of the patient (such as the age, diseases, renal and hepatic functions) [2].

When computerized provider order entries (CPOEs) are used to prescribe drugs, it is possible to detect situations at risk of ADE via prevention rules, such as *"Heparin & age>70 → increased bleeding risk"*. Those rules enable to detect risky situations and to prevent from an ADE by alerting the prescriber. The ADE is still not observed when the alert fires: that can be called *prospective ADE prevention.*

Another subject of research is *retrospective ADE detection*. It aims at analyzing past hospital stays to discover cases where ADEs really occurred. An ADE case is a hospital stay where an outcome occurred, and where that outcome is explained by a set

[1] Corresponding Author: Dr Emmanuel Chazard ; Clinique de Santé Publique, CHRU de Lille, 2 avenue Oscar Lambret 59037 Lille Cedex ; emmanuel.chazard@univ-lille2.fr

of causes related to drug administration or discontinuation, possibly combined with characteristics of the patient. Several approaches have been developed in the field of retrospective ADE detection [3-4]. They can be classified into expert-operated methods, or automated methods. The expert-operated methods suppose that an expert explicitly identifies the ADE cases. Those methods consist of retrospective medical chart reviews and reporting systems. The development of automated methods is more recent. Those methods include natural language processing of discharge summaries [5-8], and data mining of electronic health records [9].

Whatever the method used for ADE retrospective detection, a tool that could display the detected ADE cases and related statistics to the physicians of medical units does not exist yet. As a consequence, the physicians are not aware of how many ADEs occur in their medical unit, and they cannot improve their medication management.

The objective of the present work is to develop and deploy a tool that can be installed in any hospital to automatically detect past ADE cases and to display those cases to the physicians. The tool must take as input records of past hospitalizations extracted from the Electronic Health Records (EHR) of the hospital, and a set of ADE detection rules. The tool must run the rules, and provide the physicians of the hospital with comprehensive statistics about ADEs in the current department, the ADE detection rules that are interesting in the current department, and the ability to review the ADE cases that are detected by the system.

1. Material

The material consists of data that correspond to past hospital stays, and a set of ADE detection rules obtained by means of data mining.

1.1. Records of past hospital stays

As the objective is to mine past hospital stays to discover ADE cases, the Scorecards must be provided with structured description of the stays extracted from the EHR of the hospital where it is installed. This description fits a data model that has been designed previously within the PSIP Project [10]. It only uses routinely-collected data: no data have to be specifically recorded or computed for the Scorecards. The data model includes:

- Medical and administrative information (e.g., age, gender, admission date)
- Diagnoses encoded using the International Classification of Diseases (ICD10)
- Medical procedures
- Drugs administered daily to the patient, encoded using the Anatomical Therapeutic Chemical classification (ATC)
- Laboratory results encoded using the International Union of Pure and Applied Chemistry classification (IUPAC) or local terminologies
- Anonymized free-text records, such as discharge letters.

The Scorecards are installed in four hospitals (in Denmark, France and Bulgaria) and provided with about 90,000 records over 3 years (2007-2010). In some of those hospitals, the data are updated monthly.

1.2. Adverse Drug Events detection rules

The knowledge about ADEs can be summarized by means of ADE detection rules. An ADE detection rule is made of one or several Boolean conditions that lead to an outcome, with a given probability, such as $Cause_1$ & $Cause_2$ & $Cause_3$ $\rightarrow$ *Outcome*. That representation is widely used either for prospective ADE prevention or retrospective ADE detection [11]. Generally, the conditions are simple: two drugs, a drug and a lab result, a drug alone, a drug and a patient's characteristic, or a drug and a drug allergy [4, 12-23]. In this work we use a set of 236 rules that have been discovered in a previous work by data mining of EHRs [9]. Those rules involve 1 to 4 conditions that lead to an outcome. The conditions can be of demographic characteristics of the patients, drug administrations or discontinuations, laboratory results, or diagnoses. The number and the kind of the conditions were not constrained by the methods but were optimized by the use of statistical procedures. The rules enable to discover 56 kinds of outcomes, displayed in Table 1.

Table 1. Number of ADE detection rules per outcome

Outcome	Rules
Coagulation disorders	
Hemorrhage (detected by the administration of haemostatic)	7
Heparin overdose (activated partial thromboplastin time>1.23)	5
VKA overdose (INR>4.9 or administration of vitamin K)	59
Thrombopenia (count<75,000)	24
Other coagulation disorders	23
Ionic and renal disorders	
Hyperkalemia (K^+>5.3 mmol/l)	63
Renal failure (creatinine>135 µmol/l or urea>8 mmol/l)	8
Other ionic disorders	4
Miscellaneous	
Anemia (Hb<10g/dl)	2
Bacterial infection (detected by the administration of antibiotic)	4
Diarrhea (detected by the administration of an anti-diarrheal)	2
Fungal infection (detected by the administration of an antifungal)	10
Hepatic cholestasis (alk. Phos.>240 UI/l or bilirubins>22 µmol/l)	3
Hepatic cytolysis (ala. trans.>110 UI/l or asp. trans.>110 UI/l)	4
Hypereosinophilia (eosinophilocytes>10^9/l)	4
High level of pancreatic enzymes (amylase>90 UI/l or lipase>90 UI/l)	7
Neutropenia (count<1,500/mm^3)	2
Others	5
Total	*236*

The rules are described as a set of structured XML files [24]. Those files include:

- Mappings, that enable to transform the raw data into Boolean variables, e.g. *potassium≥5.3* $\rightarrow$ *hyperkalemia=1*.
- The set of rules, identified as set of conditions linked to outcomes.
- A lexicon that enables to automatically replace the names of the variables by understandable English, French or Danish labels.
- A set of free-text explanations that describe each rule and provide with bibliographic references. Those explanations are available in three languages for several uses (short label, long label, "what to do" label) and for several users (physicians, nurses and patients).

2. Methods

The display of statistics on ADEs and ADE cases relies on two steps (Figure 1). The first step, *computation step*, consists in applying the ADE detection rules to the hospital stays in order to detect ADE cases and to compute statistics about ADEs. The second step, *Web-based display tool*, consists in displaying the statistics and the ADE cases.

The Method section mainly deals with the computation step. The conception of the display tool is briefly explained in this section, and then the Web-based interface is illustrated in the Results section.

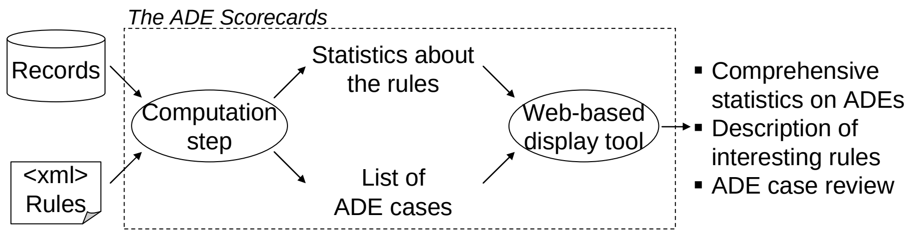

Figure 1. The ADE Scorecards rely on a computation step and a Web-based display tool.

2.1. Computation step

The computation step consists in applying the rules on the hospital stays that are extracted from the EHR. A rule is a set of conditions leading to an outcome, such as C_1 *&* … *&* $C_k \rightarrow O$. A stay that "matches the conditions of the rule" is a stay that belongs to the set $C_1 \cap \ldots \cap C_k$, if in addition the conditions are compatible regarding time: $max(startTimeC_1, \ldots, startTimeC_k) \leq min(stopTimeC_1, \ldots, stopTimeC_k)$. A stay that "matches the rule" is a stay that belongs to the set $C_1 \cap \ldots \cap C_k \cap O$, if in addition the conditions and the outcome are compatible regarding time: $max(startTimeC_1,\ldots, startTimeC_k) \leq startTimeO \leq min(stopTimeC_1,\ldots, stopTimeC_k)$. This enables to compute several statistics for each rule in the hospital. The same statistics are also computed separately in each medical department, we call them "contextualized statistics". The statistics are:

- Support = $P(O \cap C_1 \cap \ldots \cap C_k)$
- Confidence = $P(O \mid C_1 \cap \ldots \cap C_k)$
- Relative risk $RR = \dfrac{P(O / C_1 \cap ... \cap C_k)}{P(O / (\overline{C_1 \cap ... \cap C_k}))}$
- P value of the Fisher's exact test for independency between the outcome (O) and the set of conditions($C_1 \cap \ldots \cap C_k$)
- Median delay between t_1 (the conditions are met) and t_2 (the outcome occurs)
- Description of the background of the patients: average age, sex ratio, prevalence of renal insufficiency, hepatic insufficiency, and alcoholism.
- Description of what happens to the patients thereafter: average length of stay, death rate, etc.

2.2. Conception of the Web-based display tool

A Web-based tool is developed to display the statistics described above, the rules that are interesting, and the ADE cases. The following constraints are taken into account.

The Scorecards must be easily accessible: they are developed in PHP as a Web-based application and made available through an Apache HTTP server connected to a MySQL relational database. Any member of the staff equipped with a Web browser can use the Scorecards, assuming he has valid credentials.

The Scorecards must preserve the anonymity of the patients: the data used in the Scorecards concern patients who have already been discharged. The knowledge brought by the Scorecards is generic and there is no need to connect the data to the original records by name. The free-text records (e.g. discharge summaries) are automatically anonymized. The structured data do not contain any directly or indirectly nominative data (identifiers, names, birth date, dates of the stay, precise age, ZIP code…). Finally, the Scorecards are deployed in the intranet of each hospital.

The Scorecards must be easy to use: the must be able to quickly and simply find the relevant information, and not to be flooded by too much useless information. The scorecards have been developed using a Human-centered design process [25].

The Scorecards must provide the users with contextualized information: the information displayed to the user must depend on the user's characteristics and requirements. The statistics that are displayed are computed especially in the medical department of the user, and the cases that are displayed really occurred in his department. In addition, the Scorecards are fully multilingual. For the moment, the following languages are supported: English, French and Danish.

The Scorecards must be easy to deploy: The Scorecards are developed as a bootable ISO image, so that it requires a few time to deploy them into a new hospital, assuming that the data extraction are available in the form of tabulated text tables.

3. Results

This section describes the ADE Scorecards. The main features are described in the first section, and the second part consists of a use-case that demonstrates the tool.

3.1. Main features

The ADE Scorecards are a Web tool for ADE detection and ADE-related knowledge visualization. The basic course of events consists of 3 steps (Figure 2). Once logged in, the user can visualize global statistics about ADEs in his department. On a comprehensive page, it is possible to know how many ADEs occurred with respect to their kind. Then, by choosing a type of ADE, the user accesses the list of rules that are interesting in his department, i.e. the rules that would have enabled to prevent some ADEs in the department. Those rules are complemented by contextualized statistics. There is a hypertext link to the ADE cases, which allows the user to visualize all the anonymized data, including demographics, diagnoses, procedures, lab results (in tabular or graphical form), drugs administered to the patient (in tabular or graphical form), and anonymized free-text reports. This helps the user making his opinion about the case.

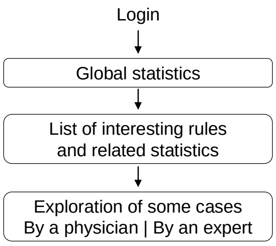

Figure 2. Basic course of events

From a technical viewpoint, the Scorecards are distributed as a bootable ISO image that contains a Web server and a set of PHP scripts. It is to be installed onto the intranet of a hospital; the installation is immediate. The hospital records have to be extracted in tabulated text according to the data model, and are automatically loaded into the database. If they are available, the free-text reports have to be anonymized first. The rules are stored as a set of XML files that can be easily updated or replaced by a customized rule set. The users have to be registered into a specific table. Then, the tool is available from the intranet through a HTTP connection.

3.2. Use case

The features of the Scorecards are presented through a sequence of commented screenshots that correspond to the following possible scenarios: “A physician working in a hospital, from which the ADE Scorecards are available, uses the Scorecards for various purposes. (Scenario 1) He wants to have a comprehensive overview of the ADEs that have been detected in his medical department during the last 6 months. (Scenario 2) Among those kinds of ADEs, he wants to explore the rules that lead to hyperkalemia (Scenario 3). Then he wants to explore one of the probable ADE cases to form his own opinion.”

3.2.1. Scenario 1: comprehensive overview about ADEs in a department

The user has to use a computer connected to the intranet and equipped with a Web browser. Once logged in, he has access to the synthesis page (Figure 3). The language select box allows for choosing the language: French, English or Danish. The synthesis page (Figure 3) consists of 3 zones. The table (part 1 of Figure 3) displays the number of ADEs detected month per month. Each line of the table is a kind of ADE; each column is a month of the current year. The line chart displays the same information using a chart (part 2 of Figure 3). In the third zone (part 3 of Figure 3), the user can chose a period of the analysis, from 2007 to 2010. He is also able to choose some kinds of ADEs and validate the form in order to generate the scorecards per kind of ADE.

3.2.2. Scenario 2: exploration of the interesting rules in a department

Once the user has chosen one or several types of ADEs and validated the form, he is displayed one page per kind of outcome chosen in the previous list. In this use case, the user focuses on the cases of hyperkalemia. The potassium is an electrolyte; its level in the plasma is regulated by the kidneys and might be influenced by some drugs and diseases. In case the potassium level raises up to 5.3 mmol/l, there is a hyperkalemia: this kind of anomaly could lead to lethal cardiac rhythm troubles.

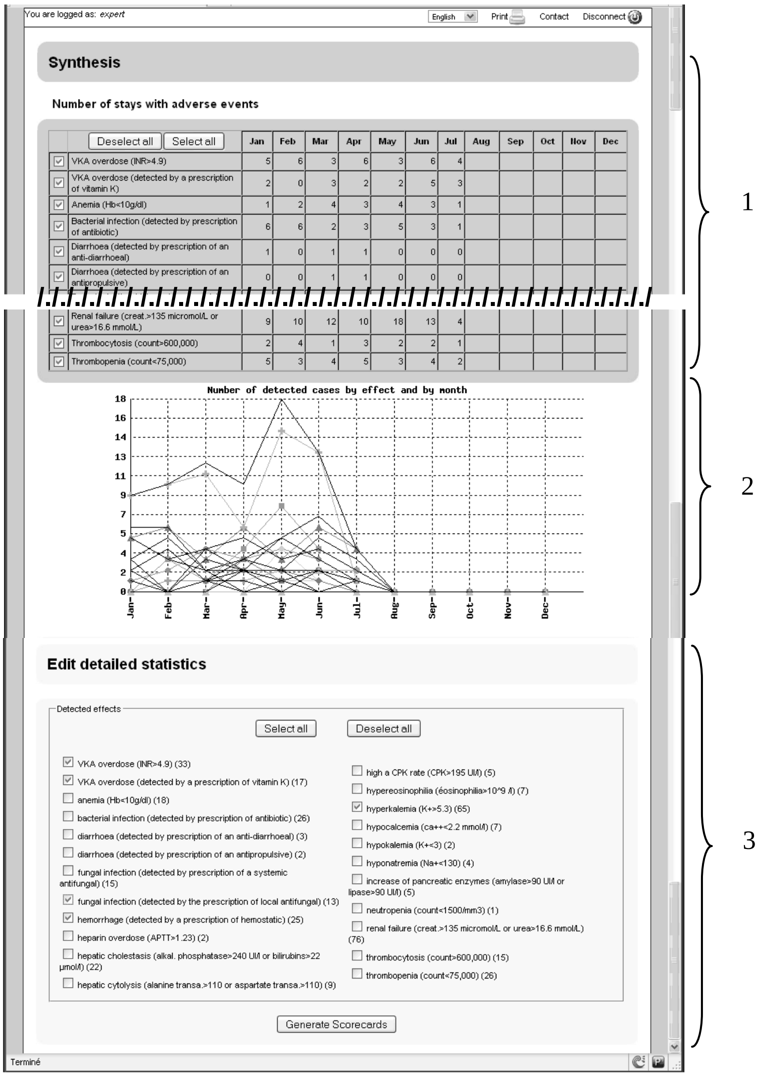

Figure 3. Synthesis page of the Scorecards

The complete scorecard is displayed (Figure 4). The page contains 4 zones, and is conceived to be either explored on the screen or printed on paper. At the top of the page, the user can read the period, the place, and the outcome that is traced (part 1 of Figure 4). In the second area, descriptive statistics are computed (part 2 of Figure 4); they describe the stays that have been detected within all the rules. In the third area, all (and only) the rules that enable to detect potential ADE cases in the current department are displayed (part 3 of Figure 4). For instance, the user can read that Low Molecular Weight Heparins (LMWH) can induce hyperkalemia especially for patients suffering from renal insufficiency (rule N°1). In the current department, 17% of patients with LMWH and renal failure encountered a hyperkalemia in a median delay of 4.5 days. At the bottom of the page (part 4 of Figure 4), more detailed explanations are provided for each rule. They can be reached by clicking on the internal hypertext links placed on the number of each rule. If the user wants to check one of those stays, he just has to click on the number of stays beside a given rule, on the right. Doing this, a popup displays the different cases that match the rule. The user can reach the corresponding stay by clicking on its identifier.

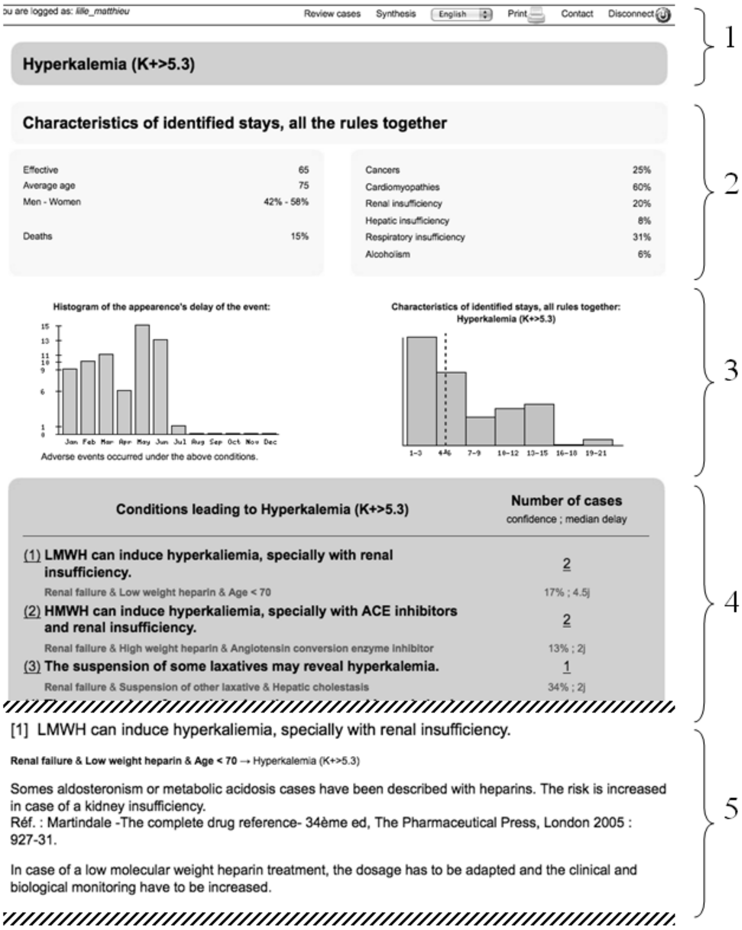

Figure 4. Scorecard of hyperkalemia (K^+>5.3)

3.2.3. Scenario 3: review of an ADE case

By clicking on its identifier, it is possible to review a potential ADE case. The user can reach several pages that display all the available information according to the data model. A page also provides comprehensive information about the stay; we present here only this screen (Figure 5). This screen is made up of 3 main parts. The top frame contains several buttons that will be described later. The left panel enables to review all the drugs that have been administered to the patient. The right panel enables to review all the laboratory results.

In the lab panel (right), by clicking on the "Potassium" checkbox (label 1 on Figure 5), the user makes the Potassium chart appear on the screen. The Potassium checkbox has a colored background because it is identified as the outcome within the rules that fire on the present stay. Several charts can appear on the same page if necessary. In this case the Potassium ion reaches a value of 5.7 on the seventh day (label 2 on Figure 5).

If the user wants to see the rules that fire for that stay, he just has to click on the button “rule info” in the head panel. A popup appears as displayed in Figure 6.

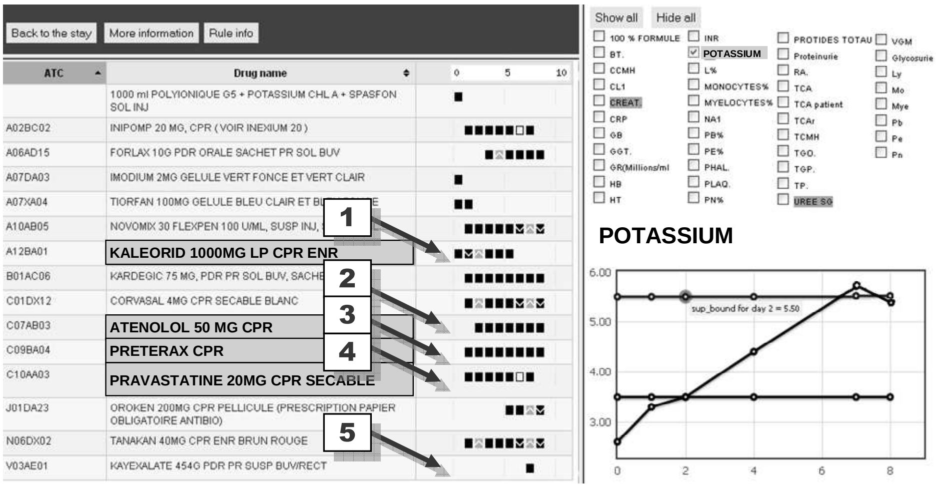

Figure 5. Main screen of the stay review facility

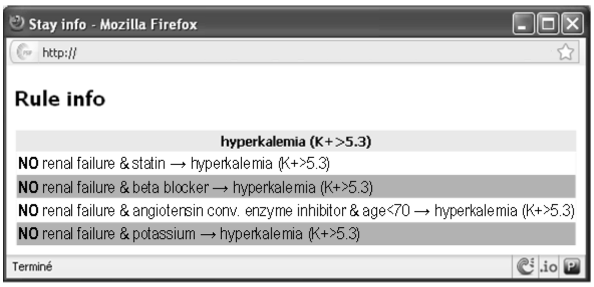

Figure 6. Popup displaying information about the rules of the current stay

In the present case, according to the rules, the drugs involved are statins, beta blockers, angiotensin conversion enzyme inhibitor, and potassium.

On the drug panel (left part of Figure 5), the user can review the drugs. The drugs that correspond to the various rules appear on a colored background (labels 3-6). The user can check that the potassium (label 3), the beta blocker (label 4), the association of the angiotensin conversion enzyme inhibitor and potassium sparing diuretic (label 5) and the statin (label 6) have been administered before day 7, the date of the outcome. All those drugs are known to increase the potassium blood level. In the present example, the user can also notice the reactions of the physicians. Hopefully the potassium is suspended before the hyperkalemia occurs (label 3). But as the potassium level reaches a very high level, a potassium lowering drug is administered during the seventh day (label 7).

The user can also access additional information by clicking on the “more information” button of the head panel. A popup appears and displays the age, the gender, the length of stay, the exit mode, and the diagnoses. In the present case, the hypokalemia is encoded (it was probably the admission ground), but the hyperkalemia is not. Finally, the Scorecards also enable the user to read the anonymized letters and reports that are previously anonymized. In that precise case, the hypokalemia is mentioned in the report but not the hyperkalemia. The physician concludes “woman admitted for a hypokalemia in relation to a gastro-enteritis (…) after correction, the potassium level is normal (…)”.

4. Discussion & Conclusion

The ADE Scorecards are an innovative tool that enables to automatically detect occurred ADE cases, by screening anonymized data extracted from an EHR with a set of rules. The detection is automated and doesn’t need any expert review, contrary to chart reviews or voluntary declarations. The rules used here have been obtained by data mining of EHRs but, as the rules consist of a set of XML files, it is simply possible to use a custom set of rules instead. Occurred ADE cases are detected, and several statistics are automatically computed, allowing the physicians to get quantitative knowledge about ADEs. The physicians are also provided with contextualized knowledge about ADEs, in the form of the set of rules that are interesting for them in their own department. This feature is important, as the knowledge about ADEs is very profuse, and not sorted by probability. Using the Scorecards, the physician can get a reasonable amount of qualitative knowledge: that knowledge is contextualized and describes their own medical unit. Moreover, the users are more responsive to that knowledge because it concerns ADEs that really occurred in their own medical unit, and they are able to review the cases in detail.

The ADE Scorecards can very easily be deployed in any hospital, as they consist of a Web server that is distributed as a bootable ISO image. The hospital has to be able to provide the Scorecards with structured extraction of data from the EHR, including administrative data, diagnoses, lab results and drug administration. If the hospital is able to provide the Scorecards with anonymized reports, then the users will benefit from them.

The Scorecards are currently being evaluated through three aspects. (1) The accuracy and the reliability of the set of rules are evaluated by medical experts who are reviewing the ADE cases detected by the tool. (2) A team of ergonomists and psychologists is evaluating the usability of the tool. (3) A prospective impact assessment is performed, to assess if the tool could help reducing the incidence of ADEs in a French hospital.

Acknowledgements

The authors would like to acknowledge every healthcare professionals (physicians, nurses, head nurses and pharmacists) and technicians whose involvement made completion of this study possible. The research leading to these results has received funding from the European Community's Seventh Framework Program (FP7/2007-2013) under Grant Agreement n° 216130 - the PSIP project.

References

[1] Aspden P. Preventing medication errors: Natl Academy Pr; 2007.

[2] Nebeker JR, Barach P, Samore MH. Clarifying adverse drug events: a clinician's guide to terminology, documentation, and reporting. Annals of internal medicine. 2004;140(10):795.

[3] Bates DW, Evans RS, Murff H, Stetson PD, Pizziferri L, Hripcsak G. Detecting adverse events using information technology. J Am Med Inform Assoc. 2003 Mar-Apr;10(2):115-28.

[4] 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 Aug;13(4):306-14.

[5] Cantor MN, Feldman HJ, Triola MM. Using trigger phrases to detect adverse drug reactions in ambulatory care notes. Quality and Safety in Health Care. 2007;16(2):132.

[6] Gysbers M, Reichley R, Kilbridge P, Noirot L, Nagarajan R, Dunagan W, et al., editors. Natural language processing to identify adverse drug events2008.

[7] Melton GB, Hripcsak G. Automated detection of adverse events using natural language processing of discharge summaries. Journal of the American Medical Informatics Association. 2005;12(4):448-57.

[8] Aramaki E, Miura Y, Tonoike M, Ohkuma T, Masuichi H, Waki K, et al. Extraction of adverse drug effects from clinical records. Studies in health technology and informatics. 2010;160:739.

[9] Chazard E, Ficheur G, Merlin B, Genin M, Preda C, Beuscart R. Detection of adverse drug events detection: data agregation and data mining. Stud Health Technol Inform. 2009;148:75-84.

[10] Chazard E, Merlin B, Ficheur G, Sarfati JC, Beuscart R. Detection of adverse drug events: proposal of a data model. Stud Health Technol Inform. 2009;148:63-74.

[11] Koutkias V, Kilintzis V, Stalidis G, Lazou K, Collyda C, Chazard E, et al. Constructing Clinical Decision Support Systems for Adverse Drug Event Prevention: A Knowledge-based Approach. AMIA Annu Symp Proc. 2010;2010:402-6.

[12] Bates DW, O'Neil AC, Boyle D, Teich J, Chertow GM, Komaroff AL, et al. Potential identifiability and preventability of adverse events using information systems. Journal of the American Medical Informatics Association. 1994;1(5):404.

[13] Jha AK, Kuperman GJ, Rittenberg E, Teich JM, Bates DW. Identifying hospital admissions due to adverse drug events using a computer-based monitor. Pharmacoepidemiol Drug Saf. 2001 Mar-Apr;10(2):113-9.

[14] Del Fiol G, Rocha B, Kuperman G, Bates D, Nohama P, editors. Comparison of two knowledge bases on the detection of drug-drug interactions2000: American Medical Informatics Association.

[15] Gandhi TK, Weingart SN, Seger AC, Borus J, Burdick E, Poon EG, et al. Outpatient prescribing errors and the impact of computerized prescribing. J Gen Intern Med. 2005 Sep;20(9):837-41.

[16] Judge J, Field TS, DeFlorio M, Laprino J, Auger J, Rochon P, et al. Prescribers' responses to alerts during medication ordering in the long term care setting. Journal of the American Medical Informatics Association. 2006;13(4):385-90.

[17] Paterno MD, Maviglia SM, Gorman PN, Seger DL, Yoshida E, Seger AC, et al. Tiering drug–drug interaction alerts by severity increases compliance rates. Journal of the American Medical Informatics Association. 2009;16(1):40.

[18] Schedlbauer A, Prasad V, Mulvaney C, Phansalkar S, Stanton W, Bates DW, et al. What Evidence Supports the Use of Computerized Alerts and Prompts to Improve Clinicians' Prescribing Behavior? Journal of the American Medical Informatics Association. 2009;16(4):531-8.

[19] Teich JM, Glaser JP, Beckley RF, Aranow M, Bates DW, Kuperman GJ, et al. The Brigham integrated computing system (BICS): advanced clinical systems in an academic hospital environment. International journal of medical informatics. 1999;54(3):197.

[20] Kuperman G, Teich J, Bates D, Hiltz F, Hurley J, Lee R, et al., editors. Detecting alerts, notifying the physician, and offering action items: a comprehensive alerting system1996: American Medical Informatics Association.

[21] Honigman B, Lee J, Rothschild J, Light P, Pulling RM, Yu T, et al. Using computerized data to identify adverse drug events in outpatients. J Am Med Inform Assoc. 2001 May-Jun;8(3):254-66.

[22] Field TS, Gurwitz JH, Harrold LR, Rothschild JM, Debellis K, Seger AC, et al. Strategies for detecting adverse drug events among older persons in the ambulatory setting. Journal of the American Medical Informatics Association. 2004;11(6):492.

[23] Handler SM, Altman RL, Perera S, Hanlon JT, Studenski SA, Bost JE, et al. A systematic review of the performance characteristics of clinical event monitor signals used to detect adverse drug events in the hospital setting. J Am Med Inform Assoc. 2007 Jul-Aug;14(4):451-8.

[24] Chazard E, Ficheur G, Merlin B, Serrot E, Beuscart R. Adverse Drug Events Prevention Rules: Multi-site Evaluation of Rules from Various Sources. Stud Health Technol Inform. 2009;148:102-11.

[25] ISO 13407:1999, Human-centered design processes for interactive systems.