Contributions from the 2016 Literature on Clinical Decision Support

V. Koutkias¹, J. Bouaud^{2,3}, Section Editors of the IMIA Yearbook Section on Decision Support

¹ Institute of Applied Biosciences, Centre for Research & Technology Hellas, Thessaloniki, Greece

² AP-HP, Department of Clinical Research and Innovation, Paris, France

³ INSERM, Sorbonne Université, UPMC Univ Paris 06, Université Paris 13, Sorbonne Paris Cité, UMRs 1142, LIMICS, Paris, France

Summary

Objectives: To summarize recent research and select the best papers published in 2016 in the field of computerized clinical decision support for the Decision Support section of the IMIA yearbook.

Methods: A literature review was performed by searching two bibliographic databases for papers related to clinical decision support systems (CDSSs). The aim was to identify a list of candidate best papers from the retrieved papers that were then peer-reviewed by external reviewers. A consensus meeting of the IMIA editorial team finally selected the best papers on the basis of all reviews and section editor evaluation.

Results: Among the 1,145 retrieved papers, the entire review process resulted in the selection of four best papers. The first paper describes machine learning models used to predict breast cancer multidisciplinary team decisions and compares them with two predictors based on guideline knowledge. The second paper introduces a linked-data approach for publication, discovery, and interoperability of CDSSs. The third paper assessed the variation in high-priority drug-drug interaction (DDI) alerts across 14 Electronic Health Record systems, operating in different institutions in the US. The fourth paper proposes a generic framework for modeling multiple concurrent guidelines and detecting their recommendation interactions using semantic web technologies.

Conclusions: The process of identifying and selecting best papers in the domain of CDSSs demonstrated that the research in this field is very active concerning diverse dimensions, such as the types of CDSSs, e.g. guideline-based, machine-learning-based, knowledge-fusion-based, etc., and addresses challenging areas, such as the concurrent application of multiple guidelines for comorbid patients, the resolution of interoperability issues, and the evaluation of CDSSs. Nevertheless, this process also showed that CDSSs are not yet fully part of the digitalized healthcare ecosystem. Many challenges remain to be faced with regard to the evidence of their output, the dissemination of their technologies, as well as their adoption for better and safer healthcare delivery.

Keywords

Medical informatics; International Medical Informatics Association; Yearbook; Decision Support Systems

Yarb Med Inform 2017;133-8

<http://dx.doi.org/10.15265/IY-2017-031>

Published online August 18, 2017

Introduction

Since the inception of the Yearbook of the International Medical Informatics Association (IMIA) in 1992, a section has been dedicated to decision support, a landmark topic in Medical Informatics. The goal of this synopsis is to summarize recent research on decision support and to select the best papers published in this field during the previous year. This literature review of 2016 publications was still targeted toward clinical decision support systems (CDSSs) and computerized provider order entry (CPOE) systems. Of note is this year's survey paper of the decision support section [1], in which research advances in the period 2015-2016 are outlined.

The synopsis is organized as follows: the next section briefly describes the method used to select the best papers on the topic as well as the slight changes made to the review process compared to the one used previously; results of this year's selection process are presented in the following section and, finally, the last section presents in detail the main contributions of the four best papers as well as noticeable research works in the decision support domain, which were identified during the selection process.

Paper Selection Method

A comprehensive literature search on topics related to CDSSs and CPOE systems was performed to identify candidate best papers following the established protocol applied the past years [2]. PubMed/Medline (from NCBI, National Center for Biotechnology Information) and Web of Science® (WoS,

from Thomson Reuters) were queried with similar queries (taking into account the query formulation syntax of each source), for journal papers published in 2016, written in English, and on these topics. From the retrieved citations, a preliminary review was performed by the two section editors to select 15 candidate best papers. These candidate best papers were then individually reviewed and rated by external reviewers from the international medical informatics community. Based on reviewers' ratings and comments, the Yearbook editorial committee kept a handful of them as best papers of the year in the decision support domain.

Reviewing hundreds of papers (beyond one thousand) is a fastidious task and, during the past years, we tried to optimize our queries to increase the specificity of the retrieved papers [3]. Basically, PubMed was searched with two queries, one based on plaintext fragments, noted as $Q_{\text{Pub_plain}}$, and the other based on MeSH terms, noted as $Q_{\text{Pub_indexed}}$. Differently, but with the same objective, searching WoS involved two queries. The first query, noted as $Q_{\text{WoS_restricted}}$, used plaintext fragments, but it was restricted to two subject areas (a WoS-specific indexing schema) of our domain: "Medical Informatics" and "Health Care Sciences & Services". The second WoS query, noted as $Q_{\text{WoS_filtered}}$, used the same plaintext fragments, but excluded subject areas which were not relevant to the domain, such as "Architecture", "Dance", "Zoology", etc. In the past years, the citations returned by these four queries were merged into one unique list, which underwent the preliminary selection review. However, when gaining some experience, we noticed that each query had not the same specificity and, consequently, the same proportion of noise, i.e. non-relevant

articles. When merging the four citation lists obtained by each query, we lost the origin of citations. So, this year, we proceeded to some changes of the reviewing process: the four queries were made disjoint and their hits were reviewed and discussed by the two section editors in separate batches.

Review Results

Database extraction on the 2016 literature with the four queries was performed on January 9, 2017. A total of 1,145 references was obtained, distributed as follows: 860 for $Q_{\text{Pub_plain}}$, 51 for $Q_{\text{Pub_indexed}}$, 69 for $Q_{\text{WoS_restricted}}$, and 169 for $Q_{\text{WoS_filtered}}$ yielding sub-totals of 911 references from PubMed and 234 from WoS. Compared to the previous year, there were 171 papers more in total. The four batches of references were reviewed separately by the two section editors. It must be noticed that both editors did not consider any candidate in both $Q_{\text{Pub_indexed}}$ and $Q_{\text{WoS_restricted}}$. Then, only the $Q_{\text{Pub_plain}}$ and $Q_{\text{WoS_filtered}}$ lists were merged, among which 62 references were kept aside by at least one reviewer. These 62 references, 56 from $Q_{\text{Pub_plain}}$ and 6 from $Q_{\text{WoS_filtered}}$, were reviewed again by the two section editors to finally select 15 candidate best papers. Following the IMIA Yearbook best paper selection process, these papers were then peer-reviewed by external reviewers and the Yearbook editors. Four papers were finally selected as best papers for 2016. Three papers [4–6] were selected

from PubMed and one paper [7] from WoS in a journal which is not indexed in PubMed. The four papers are listed in Table 1 in the alphabetical order of the first author's surname, and they are discussed in the next section. The summaries of their contents are available in the Appendix of this synopsis.

Discussion and Outlook

The study reported by Lin et al. [4] in the first paper addresses the prediction of breast cancer therapeutic decisions made during multidisciplinary team (MDT) meetings and regarding the choice of adjuvant drug therapy, i.e. after surgery was completed. At this step of the patient management, drug options may be split into three categories, chemotherapy, endocrine therapy, and targeted therapy, and for each category, the outcome can be recommended, non-recommended, or discussable. To predict MDT decision outcomes, the authors principally adopted a machine learning approach to capture the decision-making expertise of the MDT from clinicopathologic criteria. The authors implemented a total of 18 classifiers, based on variants of 10 supervised machine learning methods. Classifiers were trained using stratified ten-fold cross-validation on a dataset of 1,065 retrospective MDT decisions of a unique hospital site over an 8-year period. For each decision category, the best classifiers accurately predicted MDT decisions with areas under the receiver op-

erating characteristic (ROC) curve close to or greater than 0.9. The overall approach is somewhat classical, apart from the fact that the predicted outcome is a decision and not a patient-related event. Classically, these predictive models, based on past decision-making activity, might provide decision support as claimed by the authors, following thus the mantra of big data promoters. These models were in fact not tested or validated prospectively. However, one very interesting aspect of this work, beyond the machine learning part, is the use of other predictors based on external, qualitative knowledge and not on data. In particular, the authors implemented two guideline-based predictors relying on two widely and internationally recognized guidelines addressing best practice for adjuvant therapies in breast cancer, one from the European Society for Medical Oncology (ESMO), the other from the National Comprehensive Cancer Network (NCCN) of cancer institutions in the USA. All predictors were compared. Results showed that, on the same retrospective sample, machine learning predicted MDT decisions more accurately than guideline-based predictors. However, if there was no difference between the best classifiers and both guidelines for endocrine and targeted therapies, the difference was significant for the prediction of chemotherapy. This suggests that MDT decisions were guideline-compliant for endocrine and targeted therapies, whereas they were not compliant for chemotherapies. Authors considered that some external non-clinicopathologic criteria, taken into account in local MDT practice and captured by learning models, but not embedded within guidelines, explain guideline non-adherence. In fact, chemotherapies are highly toxic drug treatments that may be contraindicated. Our comment is that this assumption is somehow strong and should be tested considering both process and outcome measures; the opposite, iconoclastic position would indeed suggest that adhering to guidelines would have been better. Though, by explicitly comparing data-based and knowledge-based predictions, this work implicitly raises the topical question of their relative level of evidence and questions the kind of knowledge that should be incorporated into decision support tools that should drive actual practice.

Table 1 Best paper selection of articles for the IMIA Yearbook of Medical Informatics 2017 in the section 'Decision Support'. The articles are listed in alphabetical order of the first author's surname.

Section
Decision Support
<ul style="list-style-type: none"> Lin FP, Pokorny A, Teng C, Dear R, Epstein RJ. Computational prediction of multidisciplinary team decision-making for adjuvant breast cancer drug therapies: a machine learning approach. <i>BMC Cancer</i> 2016 Dec 1;16(1):929. Marco-Ruiz L, Pedrinaci C, Maldonado JA, Panziera L, Chen R, Bellika JG. Publication, discovery and interoperability of Clinical Decision Support Systems: A Linked Data approach. <i>J Biomed Inform</i> 2016 Aug;62:243–64. McEvoy DS, Sittig DF, Hickman TT, Aaron S, Ai A, Amato M, Bauer DW, Fraser GM, Harper J, Kennemer A, Krall MA, Lehmann CU, Malhotra S, Murphy DR, O'Kelley B, Samal L, Schreiber R, Singh H, Thomas EJ, Vartian CV, Westmorland J, McCoy AB, Wright A. Variation in high-priority drug-drug interaction alerts across institutions and electronic health records. <i>J Am Med Inform Assoc</i> 2017 Mar 1;24(2):331–8. Zamborlini V, Hoekstra R, Da Silva M, Pruski C, ten Teije A, van Harmelen F. Inferring recommendation interactions in clinical guidelines. <i>Semantic Web</i> 2016;7(4):421–46.

The second paper, authored by Marco-Ruiz et al. [5], introduced a linked-data approach for publication, discovery, and interoperability of CDSSs. The study was motivated by the high costs involved in the development and maintenance of CDSSs, as well as the difficulties to share them. Advocating the potential of sharing CDSSs across different systems and organizations, the authors proposed to overcome interoperability-related issues via “linked services”. This is an evolution of the Semantic Web Services paradigm to process linked data over the web, aiming to address functional data execution and non-functional semantics. Considering CDS services as linked services and exposing their interface specification as linked data would facilitate their publication, discovery, and interoperability. The authors developed models compliant with the linked-data principles to create a semantic representation of the components that make up interoperable CDS services. Overall, the developed methods and models offer a semantic layer based on machine interpretable ontologies that facilitates their interoperability and reuse. These ontologies provide unambiguous descriptions of CDS service properties to expose them to the Web of Data. The description of Web services was based on the Minimal Service Model (MSM) and it was linked to the models developed, in order to include unambiguous semantics to the service components. All the constructed models were bound to SNOMED-CT and public ontologies (such as Dublin Core and schema.org). Discovery and analysis of CDS services based on machine interpretable models were performed by reasoning over the built ontologies. The evaluation of the proposed approach relied on the implementation of a set of CDS artifacts as linked services. All the CDS artifacts (e.g. Atrial Fibrillation Treatment, Stroke Prevention Medication Recommendation, etc.) were expressed as openEHR archetypes and Guideline Definition Language rules, including terminology bindings. Among the lessons learnt, the authors argued that once the behavior of a CDS service is clearly understood, the easiest tasks are the specification of the model and the definition of non-functional and functional semantics, while the tasks that consume most of the

effort concern the definition of the clinical models and their binding with SNOMED-CT and other domain ontologies.

The third paper is authored by McEvoy et al. [6] and studied the variation in high-priority drug-drug interaction (DDI) alerts across 14 Electronic Health Record (EHR) systems operating in different healthcare institutions in the US. The assessment concentrated on current DDI alerting practice with respect to alert implementation (presence or absence of an alert) and display (alert appearance in an interruptive or passive process). The aim of the study was to prevent patient harm that may occur when providers prescribe drugs without being aware of their potential interactions, or when safer alternatives exist. Among the various strategies provided for mitigating the risk of DDIs, a highly prevalent strategy concerns the implementation of interruptive point-of-care DDI alerts within CPOE systems. Nevertheless, poor implementation within user workflows and out-of-date or poorly tiered DDI knowledge bases (KBs) result in healthcare professionals ignoring or overriding such alerts, while best practices for reducing the noise from undesired alerts remain elusive. The study explored whether there is a standard of care regarding high-priority DDI alert implementation that spans institutions and EHRs, what impact EHR vendors have on DDI alert implementation and display, and what impact healthcare organizations have on DDI alert implementation and display. The DDIs considered in the study had their basis on 15 drug pairs that should “always be alerted on”, as a standard for implementation across EHRs. The list of high-priority DDIs was approved by an expert panel as “contraindicated for concurrent use”. Seventeen medical informaticians completed the evaluation of their CPOE/EHR system, while two freely available EHRs were also evaluated. The findings of the study illustrated that no system alerted on all the DDI pairs tested. Across all systems, 58% of the DDI pairs produced interruptive alerts, while an additional 12% produced passive alerts. However, there was a great variation in alert display across systems. In one system, all alerts were interruptive, while in another all alerts were passive. Only one system used hard stops (alerts that could not be overridden), which

were applicable in seven of the evaluated DDIs. EHR vendors and DDI knowledge bases differed across systems, but nearly all systems had different severity levels of DDI alerts. The study concluded with two important recommendations: a) healthcare institutions shall carefully review their DDI alerting systems, and b) there is a need for the creation of a standardized DDI reference knowledge base, officially approved by a national/international committee including all relevant stakeholders (pharmacists, physicians, and informaticians), given that creating a standardized DDI alerting methodology requires knowledge from different disciplines.

With the development of chronic diseases and because of aging populations, multimorbidity has become a major public health issue, with an increase of complex patient cases to manage. In the fourth paper, Zamborlini et al. [7] tackle this issue from a knowledge modeling perspective. Clinical practice guidelines are usually elaborated for one health problem, typically the management of one single disease. The multitude of possible comorbidities cannot be anticipated in a way that would allow the production of evidence-based, best-practice recommendations for all clinical situations. In the case of multimorbid patients, recommendations issued from single-disease guidelines need to be reconciled. In order to manage the possible interactions between guideline-based recommendations, the authors proposed a generic modeling framework to represent the guideline contents in a formal way. This model, called TMR4I (for Transition-based Medical Recommendations for detecting Interactions), extends previous works of the authors. It basically relies on components that can be shared within the same guideline representation, or between the representation of several guidelines: actions, transitions, and recommendations. A ‘transition’ from one state to another using some ‘action’ (e.g. administration of a drug) might be generated or not by a ‘recommendation’. Logical statements are provided to infer the possible interactions that might occur between two recommendations within the same guidelines or from two different guidelines. The interaction categories are repetition, contradiction, and alternative. A prototype has

been developed to demonstrate the automatic detection of recommendation interactions. This detection is not done at the execution time in the context of a patient case, but during the guideline modeling phase. Authors used semantic web technologies to implement their prototype. The TMR4I model has been represented in OWL (Web Ontology Language) and the definitions of recommendation interactions as rules. Additional external knowledge about DDIs was used and came from a linked data resource (DrugBank). An evaluation of the approach was performed for two realistic use cases originating from prior works on the same topic by different authors, in order to allow for comparisons. The first case relies on the modelling of two guidelines (for the management of Duodenal Ulcer and Transient Ischemic Attack) and the second case on the representation of three guidelines for Osteoarthritis, Hypertension, and Diabetes. Authors discussed the characteristics of their approach in comparison with others'. Important features of this work on multiple guidelines management are its genericity and its ability to handle many guidelines concurrently.

Among the 1,145 reviewed papers for the Decision Support section of the 2017 edition of the IMIA Yearbook, several contributions brought to light some interesting results and developments and deserve to be cited in this synopsis. For example, with respect to CPOE alerting, Wipfli et al. [8] conducted a crossover randomized trial to illustrate how regrouping alerts in the CPOE layout influences physicians' prescription behavior. Heringa et al. [9] focused on decreasing the alert rates by clustering relevant drug interaction alerts, while Lilih et al. [10] proposed a method for improving the effectiveness of drug safety alerts to increase adherence to guidelines for gastrointestinal prophylaxis. In addition, Amato et al. [11] analyzed a large number of CPOE-related patient safety reports that occurred in the medication ordering phase from six sites participating to a project examining CPOE safety. In the domain of drug safety, Botsis et al. [12] presented a decision support environment for medical product safety surveillance, which has been developed by the Food and Drug

administration (FDA). Similar to the work of Zamborlini et al. [7], Wilk et al. [13] presented a mitigation framework for the concurrent application of multiple clinical practice guidelines. In the scope of guideline-based CDSSs, Kilsdonk et al. [14] conducted a comprehensive systematic review and gap analysis regarding the factors that influence the successful implementation of guideline-based CDSSs. Regarding CDSSs incorporating genomic data, Evans et al. [15] presented a new framework and a prototype solution for genomic-based CDSSs. An interesting methodological contribution has been made by the study of Halpern et al. [16], in which electronic medical record phenotyping relied on the anchor and learn framework. Tsai et al. [17] conducted a randomized controlled trial to explore whether false positive alerts in naïve clinical decision support system lead to false adoption by physicians, while Tham et al. [18] presented the development of and the lessons learned from a CDSS for pediatric head trauma in the scope of a multicenter trial with the participation of 13 Emergency Departments.

To supplement this panorama, we cannot but direct the reader to the informal review published in the IMIA Yearbook in 2016 by Middleton et al. [19] that summarizes the evolution of the state-of-the-art in clinical decision support in the past 25 years, and that provides a vision of what CDSSs might look like 25 years ahead. Notwithstanding progress on many dimensions, this review of the 2016 literature and the selection of four best papers among many illustrates that research in the field is still very active concerning diverse aspects, such as the types of CDSSs, e.g. guideline-based, machine-learning-based, knowledge-fusion-based, etc. and addresses challenging areas such as linking multiple concurrent guidelines for comorbid patients, the resolution of interoperability issues, as well as the evaluation of CDSSs. Nevertheless, this process also showed that CDSSs are not yet fully part of the digitalized healthcare ecosystem. Yet, many challenges remain to be faced with regard to the evidence of their output, the dissemination of their technologies, as well as their adoption for better and safer healthcare delivery.

Acknowledgement

We would like to thank Martina Hutter and Adrien Ugon for their support, and the reviewers for their participation to the selection process of the Decision Support section of the IMIA Yearbook.

References

1. Jenders R. Advances in clinical decision support: Highlights of practice and the literature 2015-2016. *Yearb Med Inform* 2017;125-32
2. Lamy JB, Séroussi B, Griffon N, Kerdelhué G, Jaulent MC, Bouaud J. Toward a formalization of the process to select IMIA Yearbook best papers. *Methods Inf Med* 2015;54(2):135-44.
3. Bouaud J, Koutkias V. Computerized Clinical Decision Support: Contributions from 2015. *Yearb Med Inform* 2016 Nov 10;(1):170-7.
4. Lin FP, Pokorny A, Teng C, Dear R, Epstein RJ. Computational prediction of multidisciplinary team decision-making for adjuvant breast cancer drug therapies: a machine learning approach. *BMC Cancer* 2016 Dec 1;16(1):929.
5. Marco-Ruiz L, Pedrinaci C, Maldonado JA, Panziera L, Chen R, Bellika JG. Publication, discovery and interoperability of Clinical Decision Support Systems: A Linked Data approach. *J Biomed Inform* 2016 Aug;62:243-64.
6. McEvoy D, McEvoy DS, Sittig DF, Hickman TT, Aaron S, Ai A, Amato M, et al. Variation in high-priority drug-drug interaction alerts across institutions and electronic health records. *J Am Med Inform Assoc* 2017 Mar 1;24(2):331-8.
7. Zamborlini V, Hoekstra R, Da Silveira M, Pruski C, ten Teije A, van Harmelen F. Inferring recommendation interactions in clinical guidelines. *Semantic Web* 2016;7(4):421-46.
8. Wipfli R, Ehrler F, Bediang G, Bétrancourt M, Lovis C. How regrouping alerts in computerized physician order entry layout influences physicians' prescription behavior: Results of a crossover randomized trial. *JMIR Hum Factors* 2016 Jun 2;3(1):e15.
9. Heringa M, Siderius H, Floor-Schreuderling A, de Smet PA, Bouvy ML. Lower alert rates by clustering of related drug interaction alerts. *J Am Med Inform Assoc* 2017 Jan;24(1):54-9.
10. Lilih S, Pereboom M, van der Hoeven RT, Mantel-Teeuwisse AK, Becker ML. Improving the effectiveness of drug safety alerts to increase adherence to the guideline for gastrointestinal prophylaxis. *Int J Med Inform* 2017 Jan;97:139-44.
11. Amato MG, Salazar A, Hickman TT, Quist AJ, Volk LA, Wright A, et al. Computerized prescriber order entry-related patient safety reports: analysis of 2522 medication errors. *J Am Med Inform Assoc* 2017 Mar 1;24(2):316-22.
12. Botsis T, Jankosky C, Arya D, Kreimeyer K, Foster M, Pandey A, et al. Decision support environment for medical product safety surveillance. *J Biomed Inform* 2016 Dec;64:354-62.
13. Wilk S, Michalowski M, Michalowski W, Rosu D,

- Carrier M, Kezadri-Hamiaz M. Comprehensive mitigation framework for concurrent application of multiple clinical practice guidelines. *J Biomed Inform* 2017 Feb;66:52-71.
14. Kilsdonk E, Peute LW, Jaspers MW. Factors influencing implementation success of guideline-based clinical decision support systems: A systematic review and gaps analysis. *Int J Med Inform* 2017 Feb;98:56-64.
 15. Evans JP, Wilhelmsen KC, Berg J, Schmitt CP, Krishnamurthy A, Fecho K, et al. A new framework and prototype solution for clinical decision support and research in genomics and other data-intensive fields of medicine. *EGEMS (Wash DC)* 2016 Apr 19;4(1):1198.
 16. Halpern Y, Horng S, Choi Y, Sontag D. Electronic medical record phenotyping using the anchor and learn framework. *J Am Med Inform Assoc* 2016 Jul;23(4):731-40.
 17. Tsai CY, Wang SH, Hsu MH, Li YC. Do false positive alerts in naïve clinical decision support system lead to false adoption by physicians? A randomized controlled trial. *Comput Methods Programs Biomed* 2016 Aug;132:83-91.
 18. Tham E, Swietlik M, Deakyne S, Hoffman JM, Grundmeier RW, Paterno MD, et al. Clinical decision support for a multicenter trial of pediatric head trauma: Development, implementation, and lessons learned. *Appl Clin Inform* 2016 Jun 15;7(2):534-42.
 19. Middleton B, Sittig DF, Wright A. Clinical Decision Support: a 25 Year Retrospective and a 25 Year Vision. *Yearb Med Inform* 2016 Aug 2;Suppl 1:S103-16.

Correspondence to:

Dr Vassilis Koutkias
Institute of Applied Biosciences
Centre for Research & Technology Hellas
6th Km. Charilaou — Thessaloniki
P.O. Box 60361
GR — 57001 Thessaloniki, Greece
Tel. +30 2311 25 76 15
E-mail: vkoutkias@certh.gr

Appendix: Content Summaries of Selected Best Papers for the 2017 IMIA Yearbook, Section Decision Support

Lin FP, Pokorny A, Teng C, Dear R, Epstein RJ
Computational prediction of multidisciplinary team decision-making for adjuvant breast cancer drug therapies: a machine learning approach
BMC Cancer 2016 Dec 1;16(1):929

Decisions about cancer management rely on the combined expertise of different cancer

specialists and are currently made during multidisciplinary team (MDT) meetings. This collective expertise is however non transferable between centers as a computational tool. The objective of this work is to design a predictive model of MDT decisions using machine learning techniques. The study focused on decision-making for adjuvant breast cancer therapies, i.e. the therapeutic decision after surgery, restricted to drug treatment. A cohort of 1,065 retrospective MDT decisions made in a single oncology department of an Australian hospital over an 8-year period (2007-2015) was collected. MDT decisions were considered with respect to the three modalities of systemic treatments: chemotherapy, endocrine therapy, and targeted therapy. For each modality, the outcomes of the MDT decision were either recommended, non recommended, or discussable. Each decision context was described by the clinical and pathological characteristics of the patient and the tumor at the decision time. Eighteen methods, based on 10 supervised machine learning classifiers, were trained using stratified ten-fold cross-validation for the prediction of MDT decision outcomes. Additionally, predictions were also computed for widely recognized cancer guidelines (ESMO and NCCN). For the employed dataset, results evidenced the best classifiers as those which accurately predicted the three MDT decision outcomes, using ten-fold cross validation. Considering guideline-based predictions, there was no significant difference with MDT decisions of endocrine therapy and targeted therapy. However, for chemotherapy decisions, the difference between guideline-based prediction and MDT decisions was significant and machine learning methods performed better. The authors suggested that these discrepancies for adjuvant chemotherapy might be explained by hidden, non clinicopathologic criteria, like patient preferences and resource availability, taken into account by MDT clinicians, which are captured by learning models, but not considered in guidelines.

Marco-Ruiz L, Pedrinaci C, Maldonado JA, Panziera L, Chen R, Bellika JG
Publication, discovery and interoperability of Clinical Decision Support Systems: A Linked Data approach
J Biomed Inform 2016 Aug;62:243-64

This study introduced a comprehensive paradigm for publication, discovery, and interoperability of CDSSs by employing the linked-data approach. Having their basis on the Service-oriented Architecture (SOA), linked services represent the evolution of the Semantic Web Services paradigm to process linked data. The authors provided extensive background information regarding the technical aspects which are part of their framework, as well as the technological challenges in comparison with existing CDS standards. In particular, the proposed approach entails semantics at four levels, i.e. functional, data, execution, and non-functional semantics. The definition of linked services for CDS involves the description of the service with a Web service modeling ontology and the development of ontologies to attach non-functional, functional, and clinical data semantics to service descriptions. Thus, the creation of the proposed semantic framework of interoperable CDS services relies on the models that the authors defined, comprising machine interpretable ontologies in compliance with linked-data principles. The constructed models were bound to SNOMED-CT and publicly available ontologies. These ontologies facilitate the discovery and analysis of CDS services through automated reasoning. The study illustrated the proposed approach by implementing a set of CDS artifacts as linked services. For expressing these CDS artifacts, openEHR archetypes and Guideline Definition Language rules were employed along with the appropriate terminology bindings. The authors envision the use of the proposed approach inside medium-large health networks that aim to decouple CDS functionalities from the EHR, while CDS services would be offered to any Clinical Information System in the network that requires its functionality based on a shared local knowledge base.

McEvoy DS, Sittig DF, Hickman TT, Aaron S, Ai A, Amato M, Bauer DW, Fraser GM, Harper J, Kennemer A, Krall MA, Lehmann CU, Malhotra S, Murphy DR, O'Kelley B, Samal L, Schreiber R, Singh H, Thomas EJ, Vartian CV, Westmorland J, McCoy AB, Wright A

Variation in high-priority drug-drug

interaction alerts across institutions and electronic health records

J Am Med Inform Assoc 2017 Mar 1;24(2):331-8

This study assessed the variation in high-priority drug-drug interaction (DDI) alerts across a number of diverse EHR systems operating in different healthcare institutions in the US. The main questions posed in the study concerned whether there is a standard of care regarding high-priority DDI alert implementation that spans institutions and EHRs, what impact EHR vendors have on DDI alert implementation and display, and what impact healthcare organizations have on DDI alert implementation and display. The material for conducting the study relied on 15 drug pairs approved by an expert panel as “contraindicated for concurrent use” and that should “always be alerted on”. They were used as a standard for implementation across EHR systems. For each DDI pair, the following information was recorded: a) the presence of an alert; b) the alert severity level; c) the alert display; d) the passive alert appearance, e) the override capability, and f) the override reason requirement. Seventeen medical informaticians completed the evaluation of their CPOE/EHR system, while two freely available EHRs were also evaluated. The findings of the study include: a) no system alerted on all of the DDI pairs tested; b) across all systems, 58% of the DDI pairs produced interruptive alerts, while an additional 12% produced passive alerts; c) a

great variation in alert display across systems was recorded; d) in one system, all alerts were interruptive, while in another system, all alerts were passive; e) only one system used hard stops, which were applicable in seven of the DDIs evaluated, and f) EHR vendors and DDI definition repositories differed across systems, but nearly all systems had different severity levels of DDI alerts available. In view of the original questions posed, the study concluded with two relevant recommendations: a) healthcare institutions shall carefully review their DDI alerting approaches, and b) there is a need for creating an officially approved, standardized DDI reference resource by a national or international committee comprising all relevant stakeholders.

Zamborlini V, Hoekstra R, Da Silveira M, Pruski C, ten Teije A, van Harmelen F

Inferring recommendation interactions in clinical guidelines

Semantic Web 2016;7(4):421-46

The management of multimorbidities is a growing concern in medical practice. Clinical practice guidelines provide recommendations for single diseases, as does their computerized version in decision support systems. In case of multimorbid patients, guideline-based recommendations issued for each pathology may interact, and possibly be conflicting. A unique modeling framework to represent clinical guidelines is proposed to allow for the reuse

and combination of knowledge from multiple guidelines. The formal model to represent guidelines, named TMR4I (Transition-based Medical Recommendations for detecting Interactions), is based on the descriptions of a) ‘actions’, like drug prescriptions, b) ‘transitions’ between a current state and an expected state through an ‘action’, and c) ‘recommendations’ to perform or not a ‘transition’. Logical descriptions of different types of interactions between recommendations are specified. These interactions may be internal, within the same guideline, or external, between distinct guidelines. Semantic web technologies (mainly ontologies and rules) as well as the possibility to access linked data about DDIs were used to implement a prototype that automatically infers recommendation interactions. This framework has been tested on two realistic cases on the multimorbidity management extracted from prior works published by different authors that address the same topic to allow for comparisons. The first case combines two guidelines for Duodenal Ulcer and Transient Ischemic Attack. The second case mixed three guidelines for Osteoarthritis, Hypertension, and Diabetes, respectively. In this work, the detection of interactions between recommendations is performed during the knowledge modeling phases and not at the execution time, i.e. without any given patient case. The overall framework is generic, as well as not restricted to specific guidelines, or to their number.