Artificial Intelligence in medicine AIME '11

Mor Peleg and Carlo Combi

The program for the Thirteenth European Conference on Artificial Intelligence in MedicinE (AIME'11), was prepared by the Program Committee Chair Mor Peleg (University of Haifa, Haifa, Israel), the Local Chair Nada Lavrač (Jožef Stefan Institute, Ljubljana, Slovenia), and the AIME Chairman of the Board Carlo Combi (University of Verona, Verona, Italy). The conference was held in Bled, Slovenia on July 2nd – 6th, 2011. The conference included also several specialized workshops and tutorial and was considered quite successful by most participants. Continuing a tradition started at AIME 2005, a doctoral consortium was held and included a tutorial by Carlo Combi and Riccardo Bellazzi on how to structure a research paper and the discussion of the contents of the students' doctoral theses by a scientific panel. As was done in AIME'09, two half-day tutorials and five full-day workshops were organized before and after the AIME 2011 main conference. The workshops included "Third Knowledge Representation for Health-Care (KR4HC) Workshop: Data, Processes and Guidelines", "Sixteenth IDAMAP (Intelligent Data Analysis in Biomedicine and Pharmacology) Special Topic: Intelligent Data Analysis for Quality and Safety in Healthcare, "Third Louhi workshop - Text and Data Mining of Health Documents", "Probabilistic Problem Solving in Biomedicine", and "Learning from Medical Data Streams", where the first two workshops are regularly given in conjunction with the AIME conference. The tutorials included "Introduction to Clinical Data Mining Methods" and "Personalized Healthcare Information Access".

The AIME'11 goals were to present and consolidate the international state of the art of AI in biomedical research from the perspectives of theory, methodology, and application: authors were solicited to submit original contributions regarding the development of theory, techniques, and applications of AI in biomedicine, including the exploitation of AI approaches to molecular medicine and biomedical informatics and to healthcare organizational aspects. AIME'11 received 113 abstract submissions, 92 thereof were eventually submitted as complete papers. Submissions came from 29 different countries, including 8 outside Europe. 15 of the papers (10 long and 5 short) were submitted to the special session on AI Applications, and went through a separate review process. At the end of the peer-review process, for the main track, 18 long papers (with an acceptance rate of about 26%) and 20 short papers were accepted.

The authors of the papers which were assigned the top quality scores, as judged by the scores of 3 referees who had reviewed the papers for the purposes of the conference, and which were also considered as suitable for the AIIM Journal, were invited to submit an extended version for re-evaluation for inclusion in a special AIIM issue. After the usual review and revision process, the resulting six papers are presented in the current special issue, dedicated to the AIME '11 conference.

It is intriguing to compare the trend in topics of interest to the AIME community, as reflected by these six selected papers, to the insights pointed out by Reinhold Haux in [1], which reflects about medical informatics as a discipline that evolved over the past 50 years and identifies future research fields within medical informatics. Haux argues for the principle that medical informatics research should contribute to progress in the sciences and contribute to high-quality, efficient health care. He categorizes the contribution to health-care into three categories: (a) medical informatics contributions to good medicine and good healthcare of an individual, (b) medical informatics contributions to good management of medical and healthcare knowledge, and (c) medical informatics contributions to well-organized healthcare. Considering the emerging trends of globalization, the focus on the individual (patient), and the situating of health as an integral part of life, he identifies 13 important future research topics of medical informatics that fall into the three categories (6 subcategories of a, 4 of b, and 3 of c). Below we discuss how the 6 papers of this Special Issue fall into these 3 categories and 13 topics.

Two papers fall into category 1: medical informatics contributions to good medicine and good health care for the individual. The paper by Ferrer et al. [2], Automated Generation of Patient-tailored Electronic Care Pathways by Translating Computer-interpretable Guidelines into Hierarchical Task Networks, falls into Topic #2: knowledge-based decision-support for diagnosis and therapy, and beyond (e.g., context-aware, individualized decision proposals using formally represented knowledge). This paper presents a solution to providing computer-aided support for the planning, visualization, and execution of patient-specific treatment plans in a specific healthcare process, taking into account complex temporal constraints and institutional resource allocation constraints. The authors present a methodology whereby a mapping between the Asbru guideline modeling language and a temporal Hierarchical Task Network planning language (HPDL) is defined and implemented. The translation method focuses on the representation of temporal knowledge. The method is based on the identification and translation of workflow and temporal patterns found in a computer-interpretable guideline (CIG) into HPDL. A later step augments the HPDL specification (translated from the Asbru CIG) with resource allocation constraints. This allows automatic generation, via a planner, of time-annotated and resource-based patient-specific care pathways. While this paper focuses and demonstrates the knowledge-engineering methods involved in generation of HPDL planning domain models directed to obtain patient and resource-based care plans, it also discusses how business-process modeling (BPMN) tools could be used within this framework to facilitating the care pathway's visualization and execution.

The second paper that falls into Category 1, is the paper by Minne et al. [3], Assessing and Combining Repeated Prognosis of Physicians and Temporal Models in the Intensive Care. This paper falls into topic #3: *patient-centered data analysis and mining (with representations of patient data based on appropriate semantic concepts*). After having devised a method to develop prognostic models incorporating patterns of sequential organ failure to predict the eventual hospital mortality at each day of intensive care unit (ICU) stay, in this paper the authors investigate using a real world setting how these models perform compared to physicians, who are exposed to additional information than the models. They developed prognostic models for day 2 to day 7 of ICU stay by data-driven discovery of patterns of sequential qualitative organ failure (SOFA) scores and embedding the patterns as binary variables in three types of logistic regression models. Type A models include the severity of illness score at admission (SAPS-II) and the

SOFA patterns. Type B models add to these covariates the mean, maximum and delta (increments) of SOFA scores. Type C models include, in addition, the mean, maximum and delta in expert opinion (i.e. the physicians' prediction of mortality). Physicians had a statistically significantly better discriminative ability compared to the models without subjective information (AUC range over days: 0.78-0.79 vs. 0.71-0.74) and comparable accuracy (Brier score range: 0.15-0.18 vs. 0.16-0.18). However when both sources of predictions are combined in type C models, a significantly superior discrimination as well as accuracy than the objective and subjective models alone (AUC range: 0.80-0.83; Brier score range: 0.13-0.16) are reached. The models and the physicians draw on complementary information that can be best harnessed by combining both prediction sources. Extensive external validation and impact studies are imperative to further investigate the ability of the combined model.

The four other papers selected for this Special Issue fall into Category 2: *medical informatics contributions to good management of medical and healthcare knowledge*. Two of the four papers fall into Topic #7: *systematization of medical/health knowledge (with formal representation, automated knowledge collection, beyond languages)*. The paper by Golbreich et al. [4], The FMA (Foundational Model of Anatomy) in OWL 2 and its Use, represents a novel translation of the entire FMA into OWL 2. The translation uses lexical patterns to extend, clarify, and simplify the FMA semantics. The original FMA is very large and its original specification contains many errors due to its size, the fact that it was developed over a period of ten years by several developers, and because the detection of errors is not easy as it involves several inference steps. Due to the size and large number of inconsistencies in the original FMA, it is not possible to classify the entire FMA in OWL. Yet, the authors have attempted to classify a portion of the FMA in OWL and, from this exercise, they have traced down many of the errors to conflicting values of data properties, issued from different origins. The authors found several classes of inconsistencies that can help in the correction and maintenance of the FMA. In addition, they have applied the FMA in OWL to create a lightweight terminology for human anatomy (FMA-TERM) and have included it in a cross-lingual portal of European terminologies/ontologies EHTOP devoted to resources indexing and search. Recently, EHTOP is being used daily by librarians to index French health resources in the CISMeF catalog. From the lessons learned from their experience, the authors have formulated a set of guidelines for transforming an ontology into OWL 2.

The paper by Groznik et al. [5], Elicitation of Neurological Knowledge with ABML, falls into the same topic. This paper describes the process of knowledge elicitation for a neurological decision support system (DSS). A DSS has been developed to help neurologists differentiate between three types of tremors: Parkinsonian, essential, and mixed tremor (comorbidity). The system is intended to act as a second opinion for neurologists, and most importantly to help them reduce the number of patients in the "gray area" that require a very costly further examination (DaTSCAN). To alleviate the difficult problem of knowledge elicitation from data and domain experts, the authors used a recently developed technique called ABML (Argument Based Machine Learning). ABML guides the expert to explain critical special cases that cannot be handled automatically by machine learning. This very efficiently reduces the expert's workload, and combines it with automatically learned knowledge. 122 patients were enrolled into the study. The classification accuracy of the final model was 91%. Both the initial and the final models were also evaluated for their comprehensibility by the neurologists, and the final model was deemed quite appropriate to be able to support its decisions with good explanations. The paper demonstrates ABML's advantage in combining machine learning and expert knowledge. The accuracy of the system is very high with respect to the current state-of-the-art, and the system's knowledge base is very consistent from a medical point of view. This opens up the possibility to use the system also as a teaching tool.

The two last papers in this Special issue fall into Topic 8 of Category 2: *analysis of medical and health knowledge (including knowledge generation, semantic integrity, assessing and certifying quality of knowledge)*. The paper by Rubrichi et al. [6], A System for the Extraction and Representation of Summary of Product Characteristics Content, concerns drug usage safety. Information about medications is critical in supporting decision-making during the prescription process and thus in improving the safety and quality of care. The Summary of Product Characteristics (SPC) represents the basis of information for health professionals on how to use medicines. However, this information is locked in free-text and, as such, cannot be actively accessed and elaborated by computerized applications. In this paper, the authors propose a methodology for the automatic recognition of drug-related entities (active ingredient, interaction effects, etc.) in SPCs, and their further location in a previously developed domain ontology. The approach exploits a combination of machine learning and rule-based methods. It consists of two stages. Initially it learns to classify this information in a structured prediction framework, relying on conditional random fields. The classifier is trained and evaluated using a corpus of about a hundred SPCs. They have been hand-annotated with different semantic labels that have been derived from the domain ontology. At a second stage the extracted entities are added in the domain ontology as new instances, using a set of rules manually-constructed from the corpus. Evaluations show that the extraction module exhibits high overall performance, with an average F1-measure of between 88-90%.

The last paper, by Dinh et al., A study on Factors Affecting the Effectiveness of Multi-terminology based Document Indexing and Retrieval, belongs to Topic 8 as well. The authors developed a new method for effective retrieval of biomedical documents. The method is based on using multiple terminologies. They use an approximate concept extraction method to identify concepts in each document using a mono terminology. Candidate concepts are weighted to measure their relevance to the document using a particular term weighting model (e.g., probabilistic model); second, they apply the concept extraction algorithm on several terminologies and combine several concept lists using several voting techniques. Each concept identified from each document using multiple terminologies is an implicit vote for the document. The combination of the document’s global context (domain knowledge sources or termino-ontological resources) and the query’s local context (top-ranked documents) may be a source evidence to improve the biomedical information retrieval effectiveness. The subject matters of documents can be detected via concepts extracted from documents using several terminologies (global context). For gathering contextual information of users’ queries, they use a local context query expansion method for expanding the original user query with relevant terms from the top-ranked expanded documents (local context) retrieved from the first retrieval stage. Their experimental results showed that the multi-terminology based biomedical IR approach significantly outperforms a state-of-the-art baseline.

The third category proposed by Haux, medical informatics contributions to well-organized health care, includes topics such as models for integration of health data and knowledge for enabling multiple organizations and researchers to communicate, view and exchange data, and use knowledge for reasoning based on distributed data, as well as individualized health advice and education. This category is not well covered by the six papers selected for this Special Issue, and is only partly addressed by the paper by Ferrer et al. [2], which concerns decision-support for individual patient management combined with organizational considerations of resources. The contribution to the level of multiple organizations is a relatively new topic for the AI in Medicine research community. This topic, especially data and knowledge integration, is researched in other related communities and conferences, such as International Health Informatics Symposium of the ACM Special Interest Group on Health Informatics (SIGHIT), Software Engineering for Healthcare, and Knowledge Representation for Healthcare, which is associated with the AI in Medicine community.

We trust that the readers will find these papers as an intriguing reading material; they represent, although certainly in an incomplete fashion, several of the research directions currently prominent in the Artificial Intelligence in (bio)Medicine community.

Finally we would like to thank the reviewers of the submitted papers, who helped us in creating this special issue:

Ameen Abu-Hanna, Dizza Beimel, Petr Berka, Ivan Bratko, James Hunter, Katharina Kaiser, Elpida Keravnou-Papailiou, Giorgio Leonardi, Niels Peek, Christian Popow, Silvana Quaglini, Alan Rector, Constantine D. Spyropoulos, Annette ten-Teije, Paolo Terenziani, Frans Voorbraak, Blaz Zupan, and Pierre Zweigenbaum.

References

[1] Reinhold Haux. Medical informatics: Past, present, future. Intl J Med Inform 2010;79(9):599-610.

[2] Arturo González Ferrer, Annette Ten Teije, Juan Fdez-Olivares, and Krystyna Milian. Careflow Planning: from Time-annotated Clinical Guidelines to Temporal Hierarchical Task Networks. This Issue.

[3] Lilian Minne, Evert de Jonge and Ameen Abu-Hanna. Repeated Prognosis in the Intensive Care: How well do Physicians and Temporal Models perform? This Issue.

[3] Christine Golbreich, Julien Grosjean and Stefan Darmoni . The FMA in OWL 2. This Issue.

[4] Vida Groznik, Matej Guid, Aleksander Sadikov, Martin Možina, Dejan Georgiev,Veronika Kragelj, Samo Ribarič, Zvezdan Pirtošek and Ivan Bratko. Elicitation of Neurological Knowledge with ABML. This Issue.

[5] Stefania Rubrichi, Silvana Quaglini, Alex Spengler and Patrick Gallinari. Extracting Information from Summary of Product Characteristics for Improving Drugs Prescription Safety. This Issue.

[6] Duy Dinh and Lynda Tamine . Voting Techniques for a Multi-terminology based Biomedical Information Retrieval. This Issue.

Mor Peleg*
Department of Information Systems
University of Haifa,
Haifa, Israel 31905

Carlo Combi
Department of Computer Science
University of Verona,
strada le Grazie 15, I - 37134 Verona - VR - Italy

*Corresponding author. Tel. +972-4-824-9641; fax: +972-4-828-8522
Email addresses: morpeleg@is.haifa.ac.il (M. Peleg); carlo.combi@univr.it (C. Combi)