The Medical Cyber-Physical Systems Activity at EIT: A Look under the Hood

Daniel Sonntag*, Sonja Zillner†, Samarjit Chakraborty‡, Andras Lorincz§, Esko Strommer¶, and Luciano Serafini‖
*German Research Center for Artificial Intelligence (DFKI)
Stuhlsatzenhausweg 3, Saarbruecken, Germany
Email: sonntag@dfki.de
†Siemens AG, Corporate Technology, Otto-Hahn-Ring 6, Munich, Germany
Email: sonja.zillner@siemens.com
‡Real-Time Computer Systems, Technical University of Munich, Arcisstrae 21, Munich, Germany
§Neural Information Processing Group, Eotvos Lorand University (ELTE), Budapest, Hungary
¶VTT Technical Research Centre of Finland, FI-02044 VTT, Finland
‖Data and Knowledge Management at Fondazione Bruno Kessler (FBK), Via Sommarive 18, Povo, I-38100 Trento, Italy

***Abstract*—In this paper, we describe how we combine active and passive user input modes in clinical environments for knowledge discovery and knowledge acquisition towards decision support in clinical environments. Active input modes include digital pens, smartphones, and automatic handwriting recognition for a direct digitalisation of patient data. Passive input modes include sensors of the clinical environment and/or mobile smartphones. This combination for knowledge acquisition and decision support (while using machine learning techniques) has not yet been explored in clinical environments and is of specific interest because it combines previously unconnected information sources for individualised treatments. The innovative aspect is a holistic view on individual patients based on ontologies, terminologies, and textual patient records whereby individual active and passive real-time patient data can be taken into account for improving clinical decision support.**

***Keywords*-computer-based medical systems, cyber-physical systems, knowledge acquisition, clinical decision support.**

I. Introduction

EIT ICT Labs is the Knowledge and Innovation Community of EIT with a focus on future Information and Communication Society.[1] In this framework, we have a specific innovation objective within cyber-physical systems, namely to combine active and passive user input modes in clinical environments for knowledge acquisition and decision support.

Active input modes include digital pens, smartphones, and automatic handwriting recognition for a direct digitalisation of patient data. Passive input modes include sensors of the clinical environment and mobile smartphones. This combination for knowledge acquisition and decision support (while using machine learning techniques) has not yet been explored in clinical environments and is of specific interest because it combines previously unconnected information sources. The innovative aspect is a holistic view on individual patients whereby individual active and passive patient data obtained from sensors can be taken into account for clinical decision support.

We worked towards this goal by defining four milestones [9], [10], [5], [6]: (1) knowledge acquisition by intelligent user interfaces (active modes); (2) networked embedded systems development and sensors development in presence, posture, and activity monitoring of humans by non-intrusive sensors, thereby including a functional sensor architecture (passive modes); (3) data mining: segmentation of behavioural patterns, finding analogies between behaviours, and predicting using these analogies; and (4) the knowledge integration (including a semantic model for clinical information) according to the data intelligence models in combination with clinical patient data towards clinical decision support.

The overall rationale is to address the knowledge acquisition bottleneck: future clinical care relies on digitised patient information; this information can be collected manually (active mode) by a doctor or a patient, or automatically (passive mode) by using suitable environmental sensors, e.g., of a portable device the patient is carrying. Suitable data mining models should enable us to combine those independent information sources for a combined decision support model. For example, a patient's recovery from a back injury can be monitored by the movement data while climbing steps and combined with the digitised patient disease record at the hospital's side.

II. Background

In the future, clinical environments will develop into medical cyber-physical systems of their own. Patients will get direct treatment according to a direct data acquisition and interpretation workflow. In the end, the doctor's decision support will be provided according to the data the cyber-physical system (CPS) collects from the individual patients. This approach will be scalable, extend patient monitoring to

[1]The European Institute of Innovation and Technology (EIT) is a body of the European Union based in Budapest, Hungary. It was established by the Regulation (EC) No 294/2008 of the European Parliament and of the Council of 11 March 2008.

data collections at home by using portable sensors providing information about a patient's recovery status, and influence healthcare of the future. Potential users include doctors and patients at hospital and patients at home or workplace. Towards this goal, the EIT activity Medical CPS Systems 2013-2014 develops technical components for a first medical CPS reference architecture, a template solution for an architecture for a particular domain including hardware, software, and knowledge management; the respective technical elements and relations provide templates for concrete architectures in a particular medical application domain and in the family of cyber-physical systems for medical applications: functional active and passive data acquisition methods are developed, as well as functional sensor interpretation methods. In our first running testbed system, sensor data aggregate clinical databases automatically whereupon scalable patient data can be recorded on a daily basis in a clinical database. We hope that the reference architecture will provide new insights into the interplay of active and passive sensor modes in the medical domain which we will additionally exploit in the technology transfer into the European healthcare sector. Our experts predict a change of the medical healthcare sector in Europe (according to the US developments) with three main steps: (1) digitisation of patient data, (2) usage of environmental (medical) sensor technology, (3) distributed data access and real-time clinical decision support. This activity addresses all three steps in very specific but related tasks. The envisioned goal is that the European healthcare market will profit directly by the provided reference architecture, and indirectly by a new co-operation of European institutions in a future European healthcare sector. Rich capture of clinical data, including symptoms and signs rather than just a single diagnosis, are among the objectives of parallel European efforts (e.g., http://www.transformproject.eu/).

III. Technical Components

We developed a first medical CPS reference architecture and testbed system, which combines patient data from different sources that are acquired both manually and automatically, as well as data intelligence methods for clinical decision support. The main technical contributions include the real-time dialogue server and technical components for aggregating digitised patient information, combining manual data acquisition and sensory data; and the integration of data acquisition technology into the clinical test-bed environment. Figure 1 shows the distribution of the (real-time) information sources in a combined architecture.

A. Real-Time Dialogue Server

In recent works [8], [9], [10] we have already studied the usability of novel interaction design strategies in the medical field, such as speech interaction on a mobile device for medical image annotation or classical tools, i.e., writing on mammography forms with a digital pen. Essentially, all

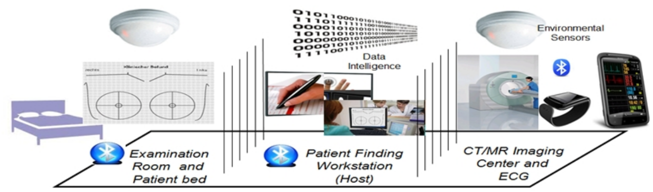

Figure 1. Architecture of the Medical CPS

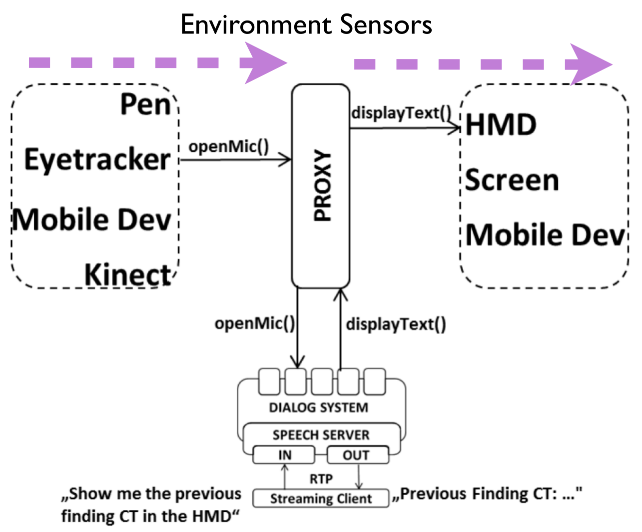

Figure 2. Input/Output Real-Time Dialogue Server

approaches share the same underlying goal to make expert knowledge acquisition fast and easy and retrieve important case information in real-time.

Our main contribution in this activity is the integration of distinct input channels consisting at the present stage of the following active modalities: i) speech-based interface, ii) see-through head-mounted display (HMD) interface, iii) pen-based interface, iv) gesture-based interface (Kinect), v) touch-based interfaces (IOS- and Android- based handheld devices (such as iPad, iPhone and Samsung Galaxy Note). We have made first experiments with the Android integration.

Figure 2 visualises the interplay of the active input components and passive environmental sensors in the context of the multi-device infrastructure. In the center of the infrastructure we have a proxy that is responsible to route and forward method invocations to any target recipient, i.e., an I/O device. The invocation of the actual method on the target is passed on through the network by means of the proxy.

The proxy architecture includes the new version of a mobile sensor architecture (mobile devices, sensor packages, mobile clients) so that real-time data can be captured and interpreted on the server. Individual datasets can be generated and data mined on the server in real-time.

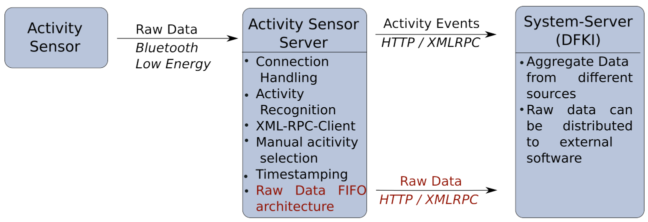

Figure 3. Sensor Data Flow

B. Mobile Sensor Architecture

The sensor system for real-time tasks distributed on multiprocessors [2] consists of three parts: (1) The activity sensor measures acceleration data in three dimensions and transmits them via a Bluetooth Low Energy (BLE)-link to the activity sensor server. Connection and disconnection happens automatically as soon as sensor and activity sensor server are in range. In case of the loss of a connection, reconnection procedures are initiated. (2) The activity sensor server receives the streaming data. It converts the raw acceleration signals from the sensors to physical units. The signal is then processed, transformed to the frequency domain and activity detection performed. From the detected activities, activity detection events are generated. These events contain timestamps together with the detected activities and are forwarded to the system server depicted in figure 2. The activity sensor server runs on a Linux-PC and can be installed on the same physical machine as the system server or on a different computer connected by a standard network (LAN, Internet). Optionally, raw sensor data can be signalised from the activity sensor server to the system server to allow further processing in the sense of real-time data mining and distributed activity recognition. (The server provided by DFKI aggregates the data from different sources and sends it to other (remote) data mining related clients/servers.) The sensor system uses the XML-RPC model to call methods of a remote server; the data flow is shown in figure 3.

C. Human Posture and Activity Level Detection

The motivation for additional environmental sensors lies in tracking the behaviour of patients or care home residents and detecting abnormal living patterns. In addition, the medical CPS perspective includes the monitoring of clinical care of bedridden patients. We take an approach for such an (eHealth) monitoring detector by implementing an intelligent furniture network (here, a bedplate). Human behaviour in the form of postures and activity levels is monitored by using a set of very low cost low-intrusive capacitive proximity sensors. The sensor system relies on wireless sensor network technologies and is extended with data management and real-time monitoring over these XML-RPC model to connect to the dialogue server. Our experimental tests show that compact algorithms based on nearest neighborhood classifiers

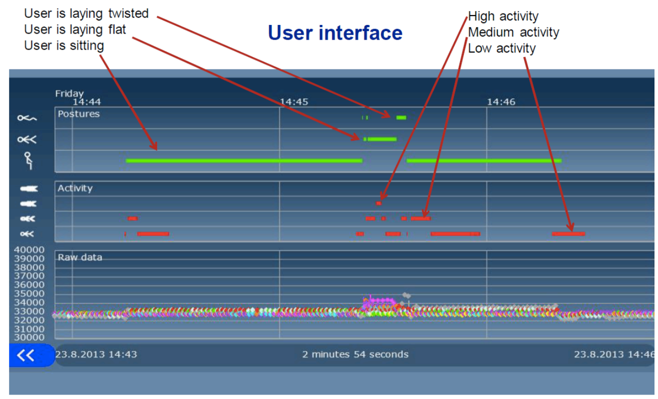

Figure 4. Monitoring Interface for Bedplate Sensors

and filter banks with Infinite Impulse Response (IIR) filters or Haar wavelets can identify the state of the bedplate user in the form of postures and activity levels. [5]

To improve the sensor features and their classification performance, and the applicability of the capacitive human behaviour tracking platform developed earlier [5], the following requirements have been met in order to integrate the bedplate into the medical cyber-physical system platform: (1) Reducing noise and cross-talk between the sensor elements by applying a distributed sensor node architecture (two sensor nodes instead of one). This leads to an improved reliability of the presence, posture, and activity detection due to reduced sensor signal noise and cross-talk; (2) Allowing for a more flexible customisation and lower manufacturing costs by a customised sensor node design (instead of a commercial AD7147-EVAL board); (3) Facilitating comprehensive sensor plate installation by a self-calibration feature of the sensor nodes and improving the monitoring user interface (figure 4) by more self-explanatory features of history data for lay users; (4) Real-time raw date capture to be sent to the distributed data mining suite for more fine-grained human posture classification according to medical guideline tasks.

D. Data Mining Suite

Data mining can be efficient for a range of medical CPS events form simple to sophisticated ones depending on a number of factors: (a) the information it receives (can be very broad range), (b) the expert knowledge provided in the form of medical ontologies (can be very little), (c) the variety of the cases, (d) the size of the database available, and (e) the expected predictions. Our medical CPS scenario will—eventually—require sophisticated spatio-temporal data mining tools. The use of sensory information (active and passive) and events such as being at places (rooms, corridors, toilette), moving between places, and medical events listed by the experts define a special spatio-temporal data mining environment.

Predictions, however, can be made on various levels. This is a relevant aspect for CPS in general because communication bottlenecks and the resulting delays can corrupt synchronisation within and between levels. The requirements for a successful data mining application are as follows: data mining needs expert information about analogies, e.g., the set of sensors that involve the notification of, e.g., the decision support when the sensor runs *in the alarm region.*

Methods: Our data mining uses entropy estimation to estimate predictability. We started from entropy estimation of places (or spatial processes according to medical guidelines). Transition probabilities and the probabilities of event series have been estimated and utilised through the prediction by partial matching method, an adaptive statistical data compression technique based on context modeling and prediction [3], [7]. PPM models use a set of previous symbols in the uncompressed symbol stream to predict the next symbol in the stream. We extended the set of information mined by using the infogain algorithm together with greedy selection since it is known to be close to optimal for our conditions [1]. This selection can be conditioned on the events to be predicted, thus restricting the domain of predictability, but increasing the time interval and precision of the prediction of the events of interest via a goal-oriented selection of additional sensors. We also compiled a dictionary of spatio-temporal events with regard to different (ontological) goals. Regular activities in the clinical environment, e.g., natural patterns of sleep, make us believe that predictions are viable. But the collection time could be in the order of months or years according to the semantic model of clinical information.

E. Semantic Model for Clinical Information

As of today, clinical data are usually stored in a large number of heterogeneous data repositories. For instance, the university hospital of Erlangen in Germany stores the patient data (clinical, administrative, financial and device data) in more than 20 different information systems such as HIS, LIS, RIS, PACS, etc. In addition, we face the situation that the patient health data is very heterogeneous. It consists of unstructured data, such as medical images or medical reports as well as structured data, such as lab report, claims, etc. In order to enable a technical access to clinical data, we need to provide means that a) allow us to store the comprehensive data set in a semantic consistent manner as well as means that b) allow us to capture the meta information describing the relevant content from unstructured clinical data to enable their subsequent automatic processing for a seamless integration within clinical applications. To address these requirements as well as to overcome the given constraints regarding patient data sharing, we developed a generic clinical data access strategy within this medical CPS activity that allows us to seamlessly and flexibly align patient data from various data sources, including sensor data. Our approach relies on two building blocks: First, an integrated clinical information model that establishes the foundation to integrate and structure clinical data by providing concepts covering meta-information and interpretations of clinical patient data; Second, an natural language processing (NLP) information extraction pipeline that provides the basis algorithms needed to extract the relevant information from unstructured clinical textual documents. In the project Theseus MEDICO, we already addressed the challenge of extracting meta-information for medical images from textual ontology description [11], within this activity we mainly focus on the extraction of information from clinical textual documents for their combination with sensory data towards real-time clinical decision support. Through the semantic representation of clinical and administrative data, we established the basis for the flexible aggregation and integration of clinical data with sensor data. Moreover, we established the basis for clinical knowledge integration and data intelligence application, such as disease-symptom models. Through the use of semantic technologies, the clinical patient data stored within hospitals can be easily aggregated with available external data sources, such as dynamic user input and dynamic sensor data. (Currently, we are in the process of aligning the semantically described clinical data with the patient sensor data.)

The proposed architecture tries to make data accessible at the data acquisition stage and provide new chances for direct return of investments in terms of direct interaction (towards interactive decision support), knowledge discovery from textual documents, and intelligent information presentation. Increased data availability should make individualised treatments possible; the problem we, however, face is that these data are not semantically integrated. As a result, most of the available data, such as sensor data, are simply not used in their full strength in clinical decisions. What we need is an integrated and standardised representation of clinical patient data reflecting the health status since this is the basis for various clinical applications like outcome analysis or other decision support systems. A standardised representation requires the use of established ontologies, vocabularies or coding systems like, e.g., the ICD10, LOINC1 or SNOMED CT2. In addition, an information model is needed where the coded data (data with references to standardised vocabularies) is stored and structured.

We identified the following requirements for a model attempting to represent clinical information and sensor data using existing ontologies according to the first clinical data acquisition tests of the medical CPS: *Integration*: data from various sources and of different format are integrated and linked; *Standards*: data should be expressed using established coding systems and terminologies: *Interpretation*: the semantics of clinical data must be defined in a consistent way; *Coverage*: it should be possible to represent all clinical data using the model in combination with other ontologies and all clinically relevant high-level concepts.

Starting from these requirements, we propose a semantic Model for Clinical Information (MCI) based on the Ontology of General Medical Science (OGMS). MCI has the purpose of integrating and structuring clinical data by providing concepts that cover meta-information and interpretations of clinical patient data. In this way, all basic concepts, which are needed to describe clinical information objects on the meta-level like diagnosis, findings, reports, health care provider, procedures, IDs, etc. are covered by MCI. Patient data that include real-time application data from active user input or passive sensors are represented using MCI in combination with large domain ontologies and coding systems (figure 5).[2]

Link to Terminologies:
The representation model of heterogeneous clinical data using terminologies might be instantiated through links from instance data (individuals) to terminologies (mainly classes). These links can be established by either using the *rdf:type* or an annotation property. If we have a class instance in the terminology exactly matching our need such as the relation to ICD codes, then *rdf:type* is a good solution for the link and reasoning mechanisms can be easily applied. If a statement is more complex and there is no single corresponding concept in the terminology, we need to post-coordinate concepts from the terminology for further location specifications (right, dorsal, ...) for example. For this purpose, we defined annotation properties such as, e.g., *mci:has qualifier*. Here it would not make sense to use the *rdf:type property*. The current version of MCI has links to ICD-10, OPS10, ATC11, LOINC, the FMA, and RadLex. By using federated SPARQL queries, it becomes possible to combine data of MCI with knowledge contained in other ontologies (as schematically shown in figure 5). This shows one of the big advantages of a semantic model formalised in RDF: we can query data from different repositories through one query—something which is not straightforward when using classical database techniques and SQL queries.

Storage of Application Data:
For clarity and improved query performance, we separate the triples using different datasets for MCI, the instance data, and the referenced ontologies. Additionally, the separation allows us to have different reasoning levels for the different datasets. MCI is held with OWL-reasoning while the patient data and the referenced ontologies without any reasoning. Further, we use named graphs in order to group patient data triples for the context of clinical encounters. The separation of triples belonging to different clinical encounters

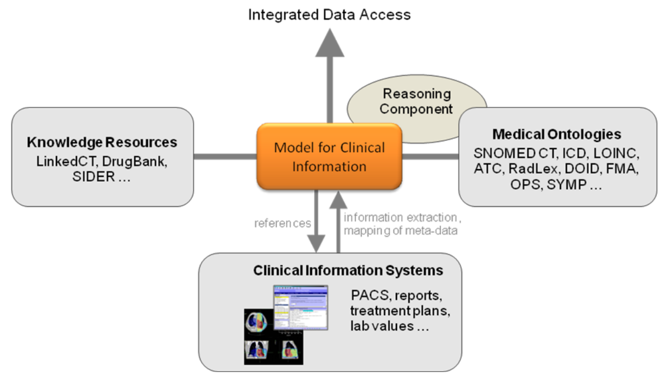

Figure 5. Coding System

is necessary because some clinical departments might have implemented different roles within the context of different encounters (e.g., admission/discharge role). Similarly, the *mci:age* at admission makes sense only within the context of some clinical encounter. The MCI covers clinical patient data as well as administrative data such as provider information and meta data for clinical encounters. Sensor data about the behaviour of the patient can now be integrated into the model. The resulting semantic model allows us to monitor for example the sleeping status which can be associated with health information; or the duration a patient is bedridden can be associated with administrative data.

F. Modelling of Clinical Guidelines

MoKi (http://moki.fbk.eu) is a tool that supports the creation of articulated enterprise models through structured wiki pages. Moki enables heterogeneous teams of (medical) experts with different knowledge engineering skills to actively collaborate by inserting knowledge; transforming knowledge; and revising knowledge at different degrees of formality. Active collaboration is guaranteed by an automatic translation between formal and informal specifications of the different (medical) experts, and by facilitating an integrated construction of the different parts of the integrated model.

Main features which we use are in the context of modeling clinical guidelines in the context of the management of exceptional flows in medical processes [4]: (1) support for the construction of integrated domain and process models; (2) easy editing of a wiki page by means of forms accessible to the medical experts; (3) automatic import and export in OWL and BPMN; easy import of lists of elements organised according to predefined semantic structures (taxonomy or partonomy); (4) term extraction functionalities; (5) graphical browsing/editing of the domain and process models; and (6) fully-integrated model evaluation functionalities (model checklist and quality indicators).

In Medical CPS, we modelled the ACR guidelines: http://www.acr.org/Quality-Safety/Standards-Guidelines/Practice-Guidelines-by-Modality/Breast-Imaging.

[2]SNOMED CT—Systematized Nomenclature of Medicine Clinical Terms; ICD—International Classification of Diseases; LOINC—Logical Observation Identifier Names and Codes: ATC—Anatomical Therapeutic Chemical Classification System; RadLex—Radiological Lexicon; DOID—Disease Ontology; FMA—Foundational Model of Anatomy; OPS—Operationen- und Prozedurenschluessel (German coding system for procedures); SYMP—Symptom Ontology; LinkedCT—Linked Clinical Trials; DrugBank—Open Drug Data; SIDER—Side Effect Resource.

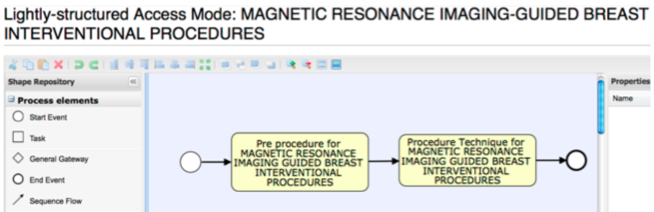

Figure 6. The MRI intervention process as composed of two subprocesses executed in sequence. Each of the subprocesses is specified in an individual wiki page.

From the Medical CPS perspective, textual input pages with additional clinical guideline annotations (conditions/actions) with formal semantics (semi-formal fragments, see figure 6) and wiki templates pages for clinicians of the mammography routine turned out to be most valuable.

IV. Conclusion

We described how we combined active and passive user input modes in clinical environments for knowledge discovery and knowledge acquisition towards real-time decision support in clinical environments. The main contributions are, first, technical components for aggregating digitised patient information, combining manual data acquisition and sensory data together with a dialogue server, and second, the integration of data acquisition technologies into a clinical testbed environment based on semantic models for clinical information. Data acquisition and integration may lead to knowledge acquisition, but do not constitute knowledge acquisition. The expected outcomes within EIT include enabling technologies for clinical decision support; and a medical CPS reference architecture that can contribute to the definition of a European Healthcare Infrastructure in Horizon 2020, see http://ec.europa.eu/programmes/horizon2020.

In the next project phase (2015), we will focus on CPS related programming questions, i.e., physiological close-loop control with human in the loop extension: The use of automatic control in clinical scenarios raises the stakes for the application of control theory in medical applications. Medical device systems for patients with complicated conditions may involve application of several treatments simultaneously, which affect several body systems in complicated and often insufficiently understood ways. These treatments also can interfere with each other. Effects of each treatment can differ widely from patient to patient. Critical variables are often not directly observable, adding to the uncertainty. Control-theoretic methods designed to operate under high parametric uncertainty and in real-time, such as supervisory adaptive control where the dialogue server proactively request users for real-time expert input, may be helpful in this context.

Acknowledgements

This work is partly supported by the EIT ICT labs, http://www.eitictlabs.eu, in the Medical CPS activity, see http://www.dfki.de/MedicalCPS/. The authors would like to thank the industrial and academic project partners and software engineers.

References

[1] M.-F. Balcan and N. J. Harvey. Learning submodular functions. In P. Flach, T. Bie, and N. Cristianini, editors, *Machine Learning and Knowledge Discovery in Databases*, volume 7524 of *Lecture Notes in Computer Science*, pages 846–849. Springer Berlin Heidelberg, 2012.

[2] J.-J. Chen and S. Chakraborty. Resource augmentation for uniprocessor and multiprocessor partitioned scheduling of sporadic real-time tasks. *Real-Time Systems*, 49(4):475–516, 2013.

[3] J. G. Cleary and I. Witten. Data compression using adaptive coding and partial string matching. *IEEE Transactions on Communications*, 32(4):396–402, Apr 1984.

[4] C. Ghidini, C. D. Francescomarino, M. Rospocher, P. Tonella, and L. Serafini. Semantics-based aspect-oriented management of exceptional flows in business processes. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, 42(1):25–37, 2012.

[5] T. Heikkila, E. Strommer, S. Kivikunnas, M. Jarviluoma, M. Korkalainen, V. Kyllonen, E. Sarjanoja, and I. Peltomaa. Low intrusive ehealth monitoring: human posture and activity level detection with an intelligent furniture network. *Wireless Communications, IEEE*, 20(4):57–63, August 2013.

[6] A. Lörincz, Z. Palotai, and G. Szirtes. Sparse and silent coding in neural circuits. *Neurocomputing*, 79:115–124, 2012.

[7] T. Schürmann and P. Grassberger. Entropy estimation of symbol sequences. *CHAOS*, 6(3):414–427, 1996.

[8] D. Sonntag and C. Schulz. A multimodal multi-device discourse and dialogue infrastructure for collaborative decision-making in medicine. In J. Mariani, S. Rosset, M. Garnier-Rizet, and L. Devillers, editors, *Natural Interaction with Robots, Knowbots and Smartphones*, pages 37–47. Springer New York, 2014.

[9] D. Sonntag, M. Weber, M. Hammon, and A. Cavallaro. Integrating digital pens in breast imaging for instant knowledge acquisition. In *Proceedings of the Innovative Applications of Artificial Intelligence Conference (IAAI)*, 2013.

[10] D. Sonntag, S. Zillner, C. Schulz, M. Weber, and T. Toyama. Towards medical cyber-physical systems: Multimodal augmented reality for doctors and knowledge discovery about patients. In A. Marcus, editor, *Design, User Experience, and Usability. User Experience in Novel Technological Environments*, volume 8014 of *Lecture Notes in Computer Science*, pages 401–410. Springer Berlin Heidelberg, 2013.

[11] S. Zillner and D. Sonntag. Aligning medical ontologies by axiomatic models, corpus linguistic syntactic rules and context information. In *CBMS*, pages 1–6. IEEE, 2011.