

Data Modelling for Dynamic Monitoring of Vital Signs: Challenges and Perspectives

Natalija Kozmina^{1(✉)}, Emil Syundyukov¹, and Aleksejs Kozmins²

¹ Faculty of Computing, University of Latvia, Raina blvd. 19, Riga, Latvia
natalija.kozmina@lu.lv, e.syundyukov@gmail.com

² Accenture Latvia, Brivibas gatve 214, Riga, Latvia
aleksejs.kozmins@accenture.com

Abstract. The use-case described in this paper covers data acquisition and real-time analysis of the gathered medical data from wearable sensor system. Accumulated data is essential for monitoring vital signs and tracking the dynamics of the treatment process of disabled patients or patients undergoing the recovery after traumatic knee joint injury (e.g. post-operative rehabilitation). The main goal of employing the wearable sensor system is to conduct rehabilitation process more effectively and increase the rate of successful rehabilitation. The results of data analysis of patient's vital signs and feedback allow a physiotherapist to adjust the rehabilitation scenario on the fly. In this paper, we focus on the methodology for data modelling with a purpose to design a computer-aided rehabilitation system that would support agility of changing information requirements by being flexible and augmentable.

Keywords: Knee joint · Rehabilitation · Real-time monitoring · Data modelling

1 Introduction

One of the most common cases of injuries that occur in people of all ages is an injury of the knee joint. As reported in [1], in the year of 2010 there were roughly 10.4 million patients' visits to doctors' offices because of common knee injuries such as fractures, dislocations, sprains, and ligament tears. In most of the cases, a surgical assistance is required for healing process. Then a rehabilitation routine follows that is aimed to minimize swelling and return the range of movement as well as to strengthen leg muscles by taking into account limitations set by a physiotherapist (e.g. flexion limitation to 90°, length of the rehabilitation session, etc.).

Rehabilitation is one of the most important parts of the treatment process. Laboratory-based equipment could be used to collect data on angular motion, and though this kind of research advances valuable information, the results remain valid only in conditions where no anticipation or reaction to the real world environment is required. According to [2], a body sensor network consisting of wireless sensors attached to a patient provides a promising method to collect clinically relevant information about knee function in everyday situations. In [3], it is recommended that healthcare specialists in the hospital would supervise the rehabilitation process,

however, patients should undergo necessary procedures on their own taking doctor's prescriptions into consideration. Continuous real-time monitoring can be achieved by using wearable devices and wireless sensor networks. Meanwhile, long-term monitoring of the physiological data could lead to significant improvement in the diagnosis and treatment of the diseases as stated in [2].

Real-time data analysis is critical to perform diagnostics and trace the dynamics in changes of the patient's state of health. Typically, data to be ingested is acquired in real time from a set of wireless sensor networks consisting of embedded devices with inertial measurement unit (IMU) and other sensor units. Our developed system aims to fill in the gap and provide an approach for patient post-operative rehabilitation monitoring. The system consists of a combination of embedded devices, sensor nodes for data acquisition, and a mobile application for data storage (NoSQL database) and analysis. The mobile application provides data to a web-based framework for data visualisation, analytics, and Cloud-based storage. In this paper, we concentrate on the aspect of data modelling suitable for performing both real-time and a more complex historical analysis of patients' data and discovered challenges that persist.

The remaining part of the paper is organized as follows. Section 2 gives a summary of related work both in the area of knee joint rehabilitation monitoring and data modelling for such systems, Sect. 3 presents a short overview of our framework, Sect. 4 describes our envisioned approach for data model re-design, and Sect. 5 finalizes the paper with conclusions and future work in different directions.

2 Related Work

To explore the existing wearable sensor systems used for knee joint rehabilitation monitoring and find possible gaps in this area, we accomplished a small-scale literature study. We performed the search via Google Scholar and IEEE Sensors Journal databases, and limited the timeframe to the last ten years to get the most recent papers (e.g. data from proposed systems is compared to new technological product like Microsoft Kinect or XSens Technologies systems). We used the following search string in Google Scholar: *"biofeedback" AND "rehabilitation" AND "wearable" AND "sensor" AND "knee joint" AND (year >= 2006 AND year <= 2016) NOT patents*.

In IEEE Sensors Journal we set the search restrictions to: Keywords = knee motion; Journal = sensors. We reviewed to most relevant 100 papers out of 483 found in Google Scholar and all 8 retrieved from IEEE Sensors. A part of the papers were eliminated because the search keywords were used in a different context (e.g. non-computer science field, systems developed for other applications, literature reviews, surveys). Snowballing applied for the paper with the highest number of citations [4] resulted in 15 papers for further analysis.

We were interested in investigating the following research questions. *RQ1*: What type of biofeedback is provided to the patient during rehabilitation session? *RQ2*: What kind of sensors is used for knee joint rehabilitation monitoring? *RQ3*: What type of communication/feedback is returned to healthcare specialists? Regarding *RQ1*, Foody et al. [5] described a biofeedback mechanism and proposed a video game that is capable of providing audible and visual feedback/instructions to a user. Matsubara et al.

[6] analysed the impact of biofeedback given to patients during rehabilitation sessions. To answer *RQ2*, mainly IMU sensors for data acquisition are employed for dynamics monitoring. Five studies included only accelerometers and gyroscopes [7–11]. For example, Seel et al. [11] stated that they avoided using magnetometers in their system, so that it would not rely on a homogeneous magnetic field. In [12], gyroscopes were avoided, and data acquired exclusively from accelerometers and magnetometers was combined. Yamada et al. [7] presented a system that had conductive textile sensors for data acquisition. They proposed a class of wearable and stretchable devices fabricated from thin films of aligned single-walled carbon nanotubes. Considering *RQ3* on communication/feedback, Daponte et al. [12] claimed that clinical staff is able to remotely monitor the movements of the subject via 3D reconstruction. Yurtman et al. [13] stated that the feedback could also be in the form of the notification alerts to inform physicians/therapists only when needed.

Though other research groups have conducted a number of studies, we noticed that little attention is paid to the data modelling as such. It yielded another research question, *RQ4*: Is there a data modelling approach to enhance both real-time and historical data analysis? While looking for relevant papers on data model design for use-cases similar to ours, we performed a search in Google Scholar by the following string: (“*data model*”) AND (“*post-traumatic*” OR “*post-operative*”) AND - “*post-traumatic stress disorder*” AND *rehab**. However, out of 72 sources no relevant work regarding data modelling for rehabilitation procedures was found. The typical approach for storing data is a traditional relational model. In [14] an elaborate description of the relational database to store human motion data is presented. Nevertheless, either of the sources lacks discussion about handling the data model if it evolves over time.

In our paper, we put an emphasis on selection of the approach for data storage guided by a set of restrictions of the system. The restrictions include the need for both real-time and historical data analysis, frequently changing information requirements and, as a result, new data objects and structures, and so on.

3 Framework Architecture for Knee Joint Dynamics Monitoring and Its Further Extensions

The initial version of the framework for knee joint dynamics monitoring is presented in [15] and depicted in Fig. 1. A physiotherapist is able to view real-time and historical analysis data of a patient’s knee joint measurements, and perform an assessment of a patient’s state of health and monitor dynamics of convalescence. The rehabilitation program consists of the number of stages during which a patient is prescribed to fulfil a set of exercises on a regular basis either under supervision of the physiotherapist or at home on his/her own in unattended mode. The developed framework prototype consists of wearable devices and mobile software for health data acquisition. The aim is to ensure that a physiotherapist is able to monitor patient knee joint movement angle in real time during rehabilitation procedures and adjust the rehabilitation program if necessary. A mobile application was developed to make calculations, analyse collected data from sensor nodes, store and visualize it, and provide feedback. Application

functionality includes communication with a patient via real-time feedback. Alerting about the dangerous situations such as exceeding doctor prescriptions on time can help decrease the chance of repeated injuries. Notification system is implemented based on health specialist/patient communication style aiming to create a similar feeling of safety during in home rehabilitation when a patient is under direct supervision of the physiotherapist. A patient is notified automatically when exceeding the flexion limit threshold, violating the rule of exercise execution, reaching the end of rehabilitation exercise session, or the rehabilitation exercise session is over.

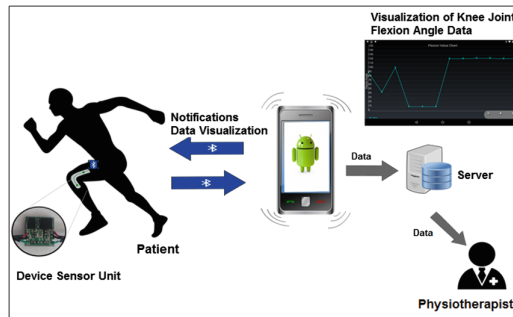


Fig. 1. Structure of proposed framework architecture for knee joint dynamics monitoring.

Another task is the quality check of the real-time data streams and employed equipment; the data received from each sensor should be verified and validated. If the produced data is incorrect (e.g. duplicated/bad/absent data), then the malfunction of a particular sensor is detected, and this sensor needs to be serviced or replaced.

Currently, the framework provides a real-time monitoring of flexion angle during two types of exercises: squats and patient's knee flexion. Wearable sensor system consists of a circuit board with MSP430 microcontroller for data sampling with 50 Hz rate and 4 sensor nodes that include sensors (3-axial accelerometers and magnetometers). For knee joint flexion/extension angle calculation a network consisting of four 3-axial accelerometers mounted on the knee is used. Precision of the developed prototype is $\pm 0.79^\circ$ in comparison to industrial goniometer (precision: ± 0.1).

It is planned to broaden the spectrum of recovery exercises (e.g. short-arc exercise, straight leg raise as in [10]), extend the number of sensor units required for taking measurements (e.g. pulse, blood pressure), gather more patient-related data (e.g. feedback from the patient on how he/she feels during/after exercise sessions), perform statistical analysis of grouped patients' data (e.g. by cohorts of patients determined by their characteristics and habits, state of health features, past diseases). All of this leads to larger volumes of data, more sophisticated analytical tasks that may be set as the framework evolves, and implies a need to re-design a data model, altogether serving as a motivation for this study.

4 An Approach to Data Model Re-design and Its Challenges

The data has to be accumulated in the special data storage and pre-processed during ingestion process in a form suitable to perform further real-time and historical analysis of patients' data. In this section, we discuss issues and requirements for the data storage suitable for our use-case that covers the task of organizing data on patients and rehabilitation programs.

4.1 Structural Flexibility and Automatisatisation Capabilities

Collected data should be arranged so that the maximum structural flexibility would be feasible. The exact and final structure of data is not yet fully known, and it is not considered to be a reachable goal. A tentative conceptual model of the framework for knee joint dynamics monitoring is given in Fig. 2; '*' indicates the extensibility of the entity or reference to another entity, while '^' indicates a sub-process.

Our goal is to design a data storage that would ensure the restructuring of data as the framework evolves. For instance, in the conceptual model (Fig. 2), such entities as *Exercise*, *PhysicalFinding*, *BaselinePattern*, *Habit*, *StateOfHealth*, etc. may be extended with a number of attributes, which are not specified yet. Attributes such as *ExerciseType*, *SmokingFreq*, *DailyActivityType*, etc. should be represented as separate entities not included in the model for readability reasons. Furthermore, a similar principle applies to the data representation of the rehabilitation program in general, which is defined as a complex hierarchical structure. For example, *GeneralEvaluation*,

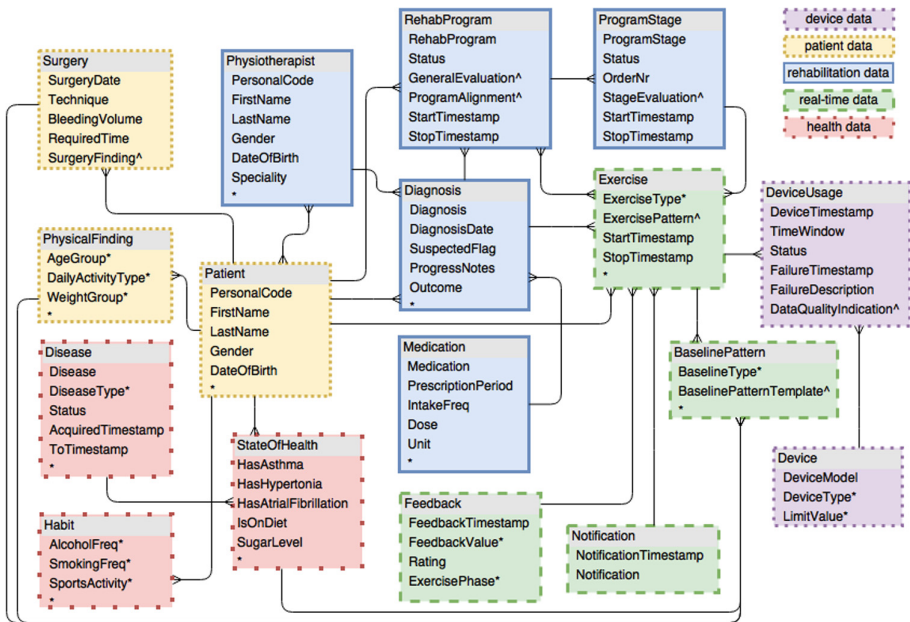


Fig. 2. Conceptual model of the framework for knee joint dynamics monitoring.

ProgramAlignment, *BaselinePatternTemplate*, *DataQualityIndication* refer to a set of entities that describe separate sub-processes of the rehabilitation program (e.g. pattern matching between patient's exercise pattern and a baseline acquired by machine learning algorithms) and are not reflected in the model.

Data structures should be used to select an appropriate rehabilitation program automatically. In fact, a mapping has to be built which takes data on *Patient*, *Diagnosis*, *Surgery*, *Disease*, *Medication*, *Habit*, *PhysicalFinding*, and *StateOfHealth* that serves as an input to suggest an appropriate rehabilitation program for each patient. Taking into consideration that the final version of the data model does not exist, we formulate a set of overall requirements that have to be met in the course of the physical implementation of the data model. It should ensure maximum simplicity of maintaining: (i) changes in the structure of all entities of the data model, (ii) changes of inter-component relations, and (iii) history of all changes made.

4.2 Complexity of the Data Model Re-design

A relational data model would ensure the implementation of the above-mentioned requirements, but still it is not adaptable to support the frequent changes: it requires significant time and effort of software developers for re-designing data structures and adjusting analytics. As adaptation to changes should be carried out unpredictably often, this approach will not give proper effectiveness and is unlikely to be justified. A "check-list" for the data model could be as follows: it should combine transactional, historical tracking/data warehousing and analytical processing features, being effort-/time-saving in development and further support.

According to the results of recent studies [16, 17], other typical data models (e.g. the ones used in data warehousing, *DW*) are not efficient enough for providing capabilities of the real-time data acquisition, whereas Data Vault (*DV*) model suggested by Dan Linsted [18, 19] fits better for this purpose.

State-of-the art studies [17, 20, 21] show that Data Vault (*DV*) model is designed for solving the problems of flexibility and performance, enabling maintenance of a permanent system of records ("all data, all the time"). The *DV* incorporates concepts from massively parallel architectures, Big Data, real-time and unstructured data, and enables agility in data model re-composition in the context of rapidly changing information requirements; these features determine our data model. As formulated in [18]: "Data Vault resolved major (flexibility and performance related) *DW* problems by elevating staging area into a persistent system of records". In our use-case, we do not deal with *DW*; nonetheless, the *DV* approach classified as operational is the best choice for our system. As stated in [22], it is an extension to the *DV* that is immediately accessed by operational systems, and is used when the real-time support as well as reading or writing directly to the data storage is required.

DV includes *hubs*, *links*, and *satellites*. We illustrate *DV* data structures with an example in Fig. 3a. Hubs correspond to objects of the conceptual model represented as unique lists of business keys that are used to track and identify key information (e.g. *HUB_PATIENT* and *HUB_REHAB_PROGRAM*). Links define relations between objects (e.g. *LNK_PATIENT_REHAB_PROGRAM*). Satellites contain descriptive attributes of the objects (e.g. *SAT_PATIENT*, *SAT_PHYSICAL_FINDING*).

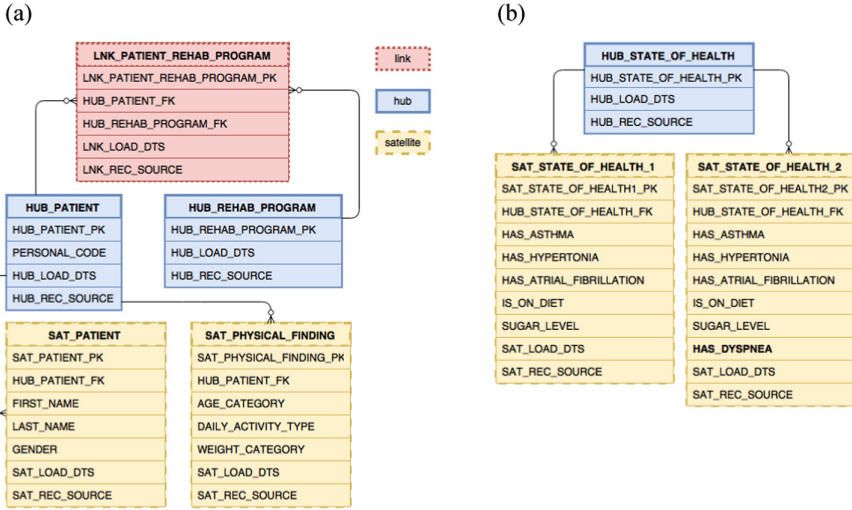


Fig. 3. (a): Representing entities *Patient* and *RehabProgram* in hub/link/satellite data structure, (b): Two versions of *StateOfHealth* satellite table (before/after adding *HasDyspnea* attribute).

Each structure element of the DV has its load date timestamp (<structure>_LOAD_DTS) and record source (<structure>_REC_SOURCE).

In our use-case, the determining factors are the flexibility of restructuring and historicity of the data. Change management should be available not only at the development stage, but also to framework users (e.g. support/non-technical staff) in the process of its operation. Hence, it leads to introducing a functional layer to describe the transformations both in data structures and mappings by means of the graphical user interface (GUI). A subject-oriented data definition language (DDL) should be provided to operate freely with conceptual objects. The concept-based statements should be translated to the 2nd layer DDL statements applied to DV objects. We are evaluating the existing solutions on the subject of their applicability to our domain of study.

Let's consider an example of representing *StateOfHealth* entity of the conceptual model (Fig. 2) as a DV object. The corresponding satellite table **SAT_STATE_OF_HEALTH_1** (Fig. 3b) contains a set of attributes that got extended with an attribute **HAS_DYSYPNEA** represented in a new satellite table **SAT_STATE_OF_HEALTH_2**. Starting from the moment when the attributes are added to the existing satellite table, their values in the old records of the table will be empty. Thus, the application of a multi-instanced approach to represent satellites is more rational when creating a new satellite with a new set of attributes linked to the same hub table (**HUB_STATE_OF_HEALTH**). Usage of satellite tables is determined by the time of creation of a new satellite, and also by a particular treatment program. Thus, the whole history of changes in entity structure is being preserved.

The core of any rehabilitation program is a set of exercises during each stage. In the course of time the set may be altered (e.g. some exercises get eliminated/added). This way, hub tables **HUB_EXERCISE** and **HUB_REHAB_PROGRAM** should be

re-linked maintaining the history of all changes. For example, if two new exercises are added to the link table, then for the rehabilitation program with the code, say, *RPI* the old exercise code *E1* remains as a historical record and two new exercise codes (exercise code is *E2* or *E3*) are linked as current ones. Data on physiotherapists that restructures a rehabilitation program is also stored in the link table.

4.3 Analytical Reports to Assess the Effectiveness of Rehabilitation Programs and Technical Implementation Features

Analytical reports available in the framework are supposed to be of various types with filtering capabilities on a set of parameters such as an arbitrary combination of attributes, time restrictions, etc., also including ad-hoc queries (e.g. rehabilitation progress of patients with a certain diagnosis in a given time period of time).

We propose to employ DV as a base data model instead of the DWH, since in our case there is no data delivered from source systems by an appropriate ETL. All stored information should be used for analytics with minimal possible structural additions. The difficulty of this approach is dictated by the specificity of the DV model, for which as far as we know, off-the-shelf applications of this kind haven't been developed yet. Thus, another task would be to create an additional – analytical – layer over DV data for reporting. Similarly to the subject-oriented DDL, it is necessary to implement an effective language for defining analytical queries, which should be translated to 2nd layer DML operating with DV objects. The susceptibility and adaptability of the reporting tool to frequent changes of both DV objects and inter-component relations must be a prerequisite. Either a DV structure or a traditional relational model to answer both operational and analytical queries would be suitable for constructing an analytical layer to store pre-aggregated data. The Starry Vault [23] approach for finding a multidimensional structure in the DV is an advanced option for this purpose.

The need to implement the subject-oriented DV DDL and analytical layer raises the issue of technical implementation of the data model. From the perspective of the language, it is hardly possible to find an alternative to SQL; nevertheless, the question is what kind of database management system (DBMS) to employ. On one hand, the developed database should function as a conventional OLTP system supplemented by subject-oriented DDL. On the other hand, high efficiency is required both in carrying out structural transformations and in performing analytical queries. For this purpose, the following new kind of DBMSs are suitable: (i) DBMS that ensure effective tools for working with columnstore indexes, or (ii) hybrid DBMSs [24] that provide capabilities of Hybrid transaction/analytical processing (HTAP), and are equally productive for OLTP and analytical tasks that include high efficiency in-memory processing of complex analytical queries. The need for turning to the DBMSs of the new breed is dictated by the fact that for both DDL operations and analytics at a low level of data access high performance of indexes is required for an arbitrary combination of DV table objects. Namely, this is a typical feature of the above-mentioned types of DBMS.

5 Conclusions and Future Work

A prototype of the currently developed system consists of wearable device and mobile software for health data acquisition. A mobile application was developed to make calculations, analyse collected data from sensor nodes, store and visualize it. The framework includes functionality for communication with a patient via notifications.

We plan to provide functionality for risk factor detection for patients undergoing the post-operative rehabilitation procedure. The system should be able to detect abnormal values in certain individuals/cohorts of patients to help physiotherapists discover the potential problems and prevent the risk of gaining a repeated trauma in the course of time. The existing prototype will be examined to study positioning of the sensor nodes for more precise harvesting of vital signs. Adding supplementary sensors such as gyroscopes and more advanced signal processing including improved data filtering and sensor fusion algorithms can also affect system performance.

In this paper, we summarize that the overall requirements for the data model that would ensure the flexibility and performance are the following – it should give maximum simplicity to the maintaining of: (i) changes in the structure of all entities, (ii) changes of inter-component relations, and (iii) history of all changes made. DV could be adapted to frequent changes in information requirements – new elements are easy to add to the model without re-design of the existing structures and data load routines.

Other subjects for future research are: (i) representation and usage of baseline patterns for each exercise in the database, (ii) remote (unattended) control over rehabilitation exercise execution, and (iii) tracing the wearable device workability in real-time. A baseline is a person-oriented pattern typically derived for a cohort of patients and is based on patients' attribute data. In a data storage, it is represented as a pattern applied to the real-time streaming data coming from embedded devices with IMU during the exercise execution. The purpose of baseline matching is to trace the correctness of patient's exercise execution in real-time and to provide notifications to help a patient perform the exercise in accordance with the rehabilitation program.

Another direction of the study is detection of bad data coming from the wearable devices. The key points of investigation are: (i) pre-configuration for tracing bad data by deriving a pattern to be applied to the data flows coming from wearable devices; (ii) bad data pattern recognition in real-time during concurrently performed exercise sessions. We envision the process of bad data detection as stream analytics with flexible reconfiguration.

References

1. American Academy of Orthopaedics: Common Knee Injuries. <http://orthoinfo.aaos.org/topic.cfm?topic=a00325>. Accessed 10 June 2017
2. Bergmann, J.H.M., et al.: An attachable clothing sensor system for measuring knee joint angles. *Sensors* **13**(10), 4090–4097 (2013)
3. Pasa, L., Visna, P.: Suture of meniscus. *Scr. Med. (Brno)* **78**(3), 135–150 (2005)

4. Kuhn, H.H., Kimbrell, W.C.J.: Electrically conductive textile materials and method for making same. Patent US4803096, USA, 7 February 1989
5. Foody, J., et al.: A prototype sourceless kinematic-feedback based video game for movement based exercise. In: Proceedings of IEEE International Conference on Engineering in Medicine and Biology Society, vol. 1, p. 5366 (2006)
6. Matsubara, M., et al.: The effectiveness of auditory biofeedback on a tracking task for ankle joint movements in rehabilitation. In: Proceedings of ISON 2013, Fraunhofer IIS, Germany, pp. 81–86 (2013)
7. Yamada, T., et al.: A stretchable carbon nanotube strain sensor for human-motion detection. *Nat. Nanotechnol.* **6**, 296–301 (2011)
8. Guo, Y., et al.: A low-cost body inertial-sensing network for practical gait discrimination of hemiplegia patients. *Telemed. E-Health* **18**(10), 748–754 (2012)
9. Raya, C., et al.: KneeMeasurer. A wearable interface for joint angle measurements. In: Proceedings of DRT4all, pp. 89–96 (2007)
10. Chen, K., et al.: Wearable sensor-based rehabilitation exercise assessment for knee osteoarthritis. *Sensors* **15**(2), 4193–4211 (2015)
11. Seel, T., Raisch, J., Schauer, T.: IMU-based joint angle measurement for gait analysis. *Sensors* **14**(4), 6891–6909 (2014)
12. Daponte, P., et al.: Design and validation of a motion-tracking system for ROM measurements in home rehabilitation. *Measurement* **55**, 82–96 (2014)
13. Yurtman, A., Barshan, B.: Automated evaluation of physical therapy exercises using multi-template dynamic time warping on wearable sensor signals. *Comput. Methods Programs Biomed.* **117**(2), 189–207 (2014)
14. Riaz, Q.: Human Motion Analysis Using Very Few Inertial Measurement Units. Bonn, Germany (2015). <http://hss.ulb.uni-bonn.de/2016/4243/4243.pdf>. Accessed 10 June 2017
15. Hermanis, A., et al.: Demo: wearable sensor system for human biomechanics monitoring. In: Proceedings of EWSN 2016, pp. 247–248. Junction Publishing, Canada (2016)
16. van der Lans, R.F.: Developing a bi-modal logical data warehouse architecture using data virtualization (2016). Denodo whitepapers <http://www.denodo.com/en/document/whitepaper/developing-bimodal-logical-data-warehouse-architecture-using-data-virtualization>. Accessed 10 June 2017
17. Mesaglio, M., Mingay, S.: Bimodal IT: how to be digitally agile without making a mess (2014). Gartner <https://www.gartner.com/doc/2798217/bimodal-it-digitally-agile-making>. Accessed 10 June 2017
18. Jovanovic, V., Bojicic, I.: Conceptual data vault model. In: Proceedings of SAIS 2012, vol. 22 (2012)
19. Linstedt, D.: Data vault basics. Dan Linstedt <https://danlinstedt.com/solutions-2/data-vault-basics/>. Accessed 10 June 2017
20. Inmon, W.H., Linstedt, D.: Data Architecture: A Primer for the Data Scientist - Big Data, Data Warehouse and Data Vault. Morgan Kaufmann, 378p. (2015)
21. Krneta, D., Jovanovic, V., Marjanovic, Z.: A direct approach to physical data vault design. *Comput. Sci. Inf. Syst.* **11**(2), 569–599 (2014)
22. Linstedt, D., Olschimke, M.: Building a Scalable Data Warehouse with Data Vault 2.0. Morgan Kaufmann, 684p. (2015)
23. Golfarelli, M., Graziani, S., Rizzi, S.: Starry vault: automating multidimensional modeling from data vaults. In: Pokorný, J., Ivanović, M., Thalheim, B., Šaloun, P. (eds.) ADBIS 2016. LNCS, vol. 9809, pp. 137–151. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44039-2_10
24. Pezzini, M., et al.: Hybrid transaction/analytical processing will foster opportunities for dramatic business innovation (2014). Gartner <https://www.gartner.com/doc/2657815/hybrid-transactionanalytical-processing-foster-opportunities>. Accessed 10 June 2017

Advances in Conceptual Modeling

ER 2017 Workshops AHA, MoBiD, MREBA, OntoCom, and
QMMQ, Valencia, Spain, November 6–9, 2017,

Proceedings

De Cesare, S.; Frank, U. (Eds.)

2017, XVII, 249 p. 77 illus., Softcover

ISBN: 978-3-319-70624-5