Realizing the Predictive Enterprise through Intelligent Process Predictions based on Big Data Analytics: A Case Study and Architecture Proposal

Julian Krumeich[1], Jens Schimmelpfennig[2], Dirk Werth[1], Peter Loos[1]

[1]Insitute for Information Systems (IWi),
German Research Center for Artificial Intelligence (DFKI),
Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany
{Julian.Krumeich, Dirk.Werth, Peter.Loos}@dfki.de

[2]Software AG,
Altenkesseler Str. 17, 66115 Saarbrücken, Germany
Jens.Schimmelpfennig@softwareag.com

Abstract: Today's globalized economy forces companies more than ever to constantly adapt their business process executions to present business situations. Companies that are able to analyze the current state of their processes and moreover forecast its most optimal progress as well as proactively control them based on reliable predictions will be a decisive step ahead competitors. The paper at hands examines, based on a case study stemming from the steel manufacturing industry, which production-related data is currently collectable using state of the art sensor technologies forming a potential foundation for a detailed situation awareness and derivation of accurate forecasts. An analysis of this data however shows that its full potential cannot be utilized without dedicated approaches of big data analytics. By proposing an architecture for implementing predictive enterprise systems, the article intends to form a working and discussion basis for further research and implementation efforts in big data analytics.

1 Introduction

1.1 Vision of the Predictive Enterprise

In order to meet the steadily rising market requirements in today's globalized world, companies are more than ever forced to adapt their business processes to present process and business situations. The rising digitalization of the real world, driven by the era of Internet of Things, allows for unprecedented insights into current business process situations [Kr14a]. In the future, only those companies which are able to analyze their business operations based on this rapidly growing mass of data, predict the best proceeding process flow, and proactively control their processes with this knowledge will survive in increasingly competitive markets. Such a company sketches the vision of a "Predictive Enterprise" as the next stage in the evolution of real-time enterprises within the age of data as a decisive competitive asset [Kr14a, Lu06].

1.2 Motivation and problem definition

Today's corporate practice is beyond the realization of this vision. The potential of massive amounts of data in companies which can already be collected by state of the art sensor technologies, has only received scant attention in business process management (BPM) [Un12]. Frequently, this is due to the lack of technical capabilities to analyze big data in a timely manner and to derive the right decisions. However, business analysts agree that companies are going to face existential difficulties if they are not able to discover or even anticipate problems and defects in business process executions in time [PG11]. Thus, BPM software has to provide means to analyze process executions, forecast their further progress and consequently anticipate potential problems early enough to proactively enact appropriate countermeasures. Yet, current approaches and techniques in predictive analytics rarely take into account context situations, e.g. through the utilization of technologies like Complex Event Processing (CEP), as a basis to derive forecasts. Most predictions remain purely stochastic involving the risk of computing highly-likely forecasts of process progressions that turn out to be entirely unlikely in the present, perhaps exceptional, situation. This calls for a dedicated consideration of given situations respectively events to compute reliable forecasts. Even though available data sets will grow continuously due to enhanced sensing capabilities and provide unprecedented insights into process situations, this entails the need for dedicated big data analytics able to exploit the full potential out of such data.

1.3 Research contribution and applied research method

By means of a case study (see sec. 2 and for preliminary work [Kr14b, Kr14c]), the paper at hand analyzes which process and context data a sample company can currently collect by state of the art sensor technology forming a potential basis for a predictive planning and control of manufacturing processes. In doing so, the paper basically follows case study research, which has been applied in Information Systems research for almost two decades [DP03]. In particular, the paper employs the methodology proposed by Benbasat et al. [BGM87]: as the unit of analysis the steel bar production line at Saarstahl AG, a major German steel producer, was chosen. The two research questions applied to the case study are "What type of data is currently available in industry processes using state-of-the-art sensor technology to realize event-based predictions?" and "Why is it a 'Big Data' challenge to analyze this data appropriately?" The results of the selected case study are considered to be generalizable to other manufacturing enterprises. Since the chosen research design is a single-case study, it is particular appropriate to revelatory cases, which is typical for big data analytics cases as a relatively new phenomenon in information systems research. The case study data was collected by the central department of information and communication technology of the company. As data selection methods, interview techniques were applied and physical artifact—sensor networks—investigated.

By analyzing the collected data, the paper points out that current state of the art sensor technologies already provide data that has to be considered through the lens of big data (see sec. 2.3). To make sufficient business value out of it, especially in terms of process predictions, dedicated big data analytics are required. For this, the paper derives seven requirements to be considered for a system implementation to realize process predictions based on big data analytics (see sec. 3.1). Moreover, the article proposes a general system architecture (see sec. 3.2) intending to form a working and discussion basis for further research and implementation efforts in big data analytics (see sec. 5).

2 Case study: Big data in the steel bar production at Saarstahl AG

2.1 Brief description and case study settings

This case study analyzes a production part within the steel bar production at Saarstahl AG, a major German steel producer, and provides types and sizes of sensor data that is currently collectable throughout the production processes. In the considered production area, half a million tons of steel are produced annually. In order to meet customers' specific quality requirements for various existing end products, Saarstahl AG conducts comprehensive quality checks within its production line. These include diameter controls with laser, surface testing by magnetic particle testing, checks for internal steel errors by ultrasound, and a variety of temperature and vibration measurements. All of these techniques provide continuous sensor data at the lowest system level (L1). In addition, other sensor systems in production (ambient and positioning sensors) are installed to monitor the control of steel bars via a material flow tracking system (L2-system level). Based on this basic data and the available customer orders, the production planning and control system calculates a rough schedule (L3 to L4 system level).

2.2 Current problem statement

Steelmaking processes are significantly influenced by internal and external events that lead to deviations in production quality. Some examples are fluctuating material properties of pig iron or production-induced deviations due to physical processes taking place. This entails a need for frequent adaptations of production instances and subsequently extensive re-planning of underlying production plans. For instance, if a standard further processing of steel provides the steps peeling, heat treatment and finally sawing of steel bars, then the crude steel cannot pass these steps if the sensor network used in the process detects a bending, as an example of a production deviation.

In this case, this intermediate either needs to be post-processed in a dressing and straightening machine to comply with quality requirements, or the processing of the intermediate can continue for another customer with lower quality criteria. Therefore the production planning for this process instance is obsolete and needs to be re-initiated in order to meet quality promises and deadlines to all customers, e.g., steel bars need to be factored into the normal process flow after conducting such an ad hoc processing step. Since the production planning systems compute an almost full capacity for the following days in batch mode, the simple insertion of this process instance into the running production flow may contradict to the planned execution of the production plan.

Therefore, in case a production deviation is detected during the production process, an intelligent recalculation of the entire production plan would be necessary in real-time to ensure an optimal solution from a global planning perspective. If such production influencing events were even anticipatable, corresponding countermeasures could be initiated before its occurrence.

Currently, the data provided by the sensor network, which is integrated into the production processes, exceeds the volume that is analyzable so far. The information and control systems used and the analysis techniques available on the market do not allow the producer to monitor and control the entire production process in real-time. Moreover, no future states and events, such as looming production deviations, can be predicted on time. Hence, so far the control of production processes is rather done in a reactive way; yet, the sketched vision of a predictive enterprise has not been realized.

2.3 Data characterization, challenge of Big Data and requirement derivation

In the following, sample data obtained from the applied sensor networks at Saarstahl AG are described according to the Big Data characteristics proposed by Gartner [BL12]. If the entire sensor networks within the production process is considered—which would be necessary for a comprehensive production planning and control on an L3-/L4-systems' level—, the Big Data challenge will be many times higher.

An example from the sensor network illustrates the immense volume of data in monitoring the production process. In the rolling mills 31 and 32 there are two optical surface test sensors that can continuously provide real-time data for the detection of surface defects during the rolling process. Basically, this allows to take into account the varying customer demands for a particular surface quality. The unit can already prototypically detect errors and differentiate the types of errors. This optical testing generates several hundred terabytes of data annually (volume). Currently, only a sporadic reactive analysis of these data is possible. Also, other context data from the sensor network and the systems settled on L-2 or L-3 level can currently not be linked due to the volume of data to be analyzed. While these systems could in principle detect production deviations in batch mode, this takes too long to allow timely reactions.

While this is just one example of very large resulting data of individual sensors in a particular section of the factory, another example illustrates the high data diversity (variety), which is continuously collected by different sensors and sensor networks at various points throughout the production process. This places high demands on an analysis by big data principles. For instance, the further processing of steel bars as of now already provides half a million of sensor data records, which reflect one production area to a particular context. In the next couple of months the sensor performance will be advanced, such that over 1.5 million sensor data will be available on L1 and L2 level. According to the principles of CEP; however, only the identification of relevant events in this torrent of both homogeneous and heterogeneous data as well as their correlation allows to derive patterns and deviations. This is possible only by using highly structured, technical knowledge. At this point, the basic claim to a scalable solution becomes clear, since sensor networks should be flexibly expandable, but must also allow analyses and forecasts within a required time frame. The company plans to increase sensing in this subsection to an output of more than five million records, which underlines the need for scalability. Thus, in terms of the analysis of these large and diverse data, the responding time is crucial, since speed is a decisive competitive factor in the analysis (velocity / analytics). Classic reporting or batch processing would be significantly too slow, so that so-called high velocity technologies must be performed in near-real-time analyses. For the purposes of the outlined vision of predictive enterprises, it is also crucial to conduct accurate forecasts of the process sequences. Each day, an average of one terabyte of video data is recorded in a single subsection of the plant. However, a pure video analysis method is not sufficient for predictive analytics methodologies. In the existing system, it has been shown that only some production deviations could be detected by this classical approach. In addition, there is no feedback for process optimization. Therefore process data need to be included in the model formation and forecasting. Here, as outlined, over one million data sets are incurred in the coming months. For analyzing the dependencies among process and video data, data from a long period of time must be used for model training. In this case, the data volume may rapidly exceed 50 terabytes. For a real-time adaptive prediction, on average one-tenth of the data should be used. At present, however, such a number of data can hardly be processed in real-time. A direct compression of the data is impossible because of its variety to be considered.

As of now, due to its big data characteristics, the production process is away from the envisioned optimum in terms of planning and control. Technically, Saarstahl AG can integrate further sensor technology into its production processes to achieve a better monitoring. However to increase the demanded business value, current analytical methods take too long to complete an analysis.

3 Towards a system architecture proposal for realizing intelligent process predictions based on Big Data analytics

3.1 Requirement analysis

In past, datasets sensed from manufacturing processes were rather small covering only insufficiently operational context situations; hence, leading to imprecise forecasts. While in principle, it was possible to expand the data base, this information gathering has proven to be too complex, too expensive and not accomplishable in a timely manner. Conventionally, predictions had been computed using stochastic means; consequently, percentage values indicate which outcome or result a process will have. Such stochastic means however neither take into account nor reflect the current context situation in which a manufacturing process takes place resulting in inaccurate forecasts (cf. [Kr14b] for a comparison of stochastic and event-based predictions). Today, manufacturing companies still perform insufficient analyzes and forecasts based on their collected sensor data, even though it will positively influence their economic and ecological performance [Un12]. This is in contrast to industries such as insurance or banking that have fully implemented predictive analytics in their business models [MCD13]. Yet, recent technological progress in the fields of Internet of Things and Cyber-Physical Systems enable to equip production processes relatively cost-neutral by sensors. This allows to measure internal and external production process parameters in a previously unprecedented level of detail. Thereby the technical foundation has been created to establish and continuously enrich a data base allowing for highly-accurate predictions.

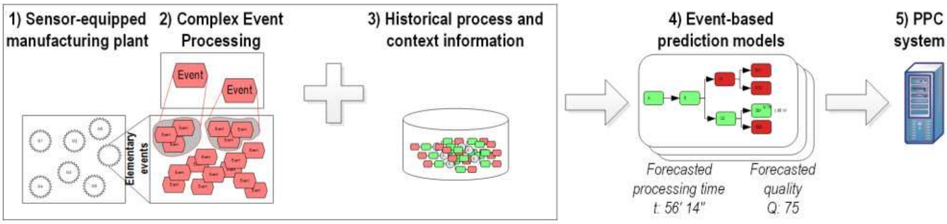

Figure 1: Basic concept for an event-based process forecast and control

As the use case analysis showed, the overall challenge of computing accurate predictions has shifted from being able to capture current context situations—as the baseline for deriving forecasts—to being able to manage and adequately consider this mass of data. In addition, the use case illustrated that the mass of data is progressively growing possessing the necessity of having a scalable platform for integrating sensor technology (cf. Fig. 1, 1). This leads to the first requirement of the system architecture.

Requirement 1. Provide scalable means for extending sensors throughout production processes and to store the mass of collectable data in descriptive process and context models.

In addition to collecting and storing process data in historical description models—as a knowledge base for process predictions (cf. Fig. 1, 3)—, it is compulsory to investigate current real-time conditions of processes through analyzing their current events and context situations. Due to the wide variety of heterogeneous sensors, which are used in enterprises, this is particularly a challenge of big data analytics in terms of data variety (cf. use case analysis in sec. 2.3). Eventually, these streams of atomic data have to be searched for patterns in real-time, e.g. by utilizing CEP technology (cf. Fig. 1, 2).

Requirement 2. Provide means for detecting and filtering of complex events within tremendous streams of atomic sensor data.

Such CEP allows to correlate current conditions with historical ones as the baseline for deriving event-based predictions. Through this correlation, event-based forecast models can be derived and repeatedly adapted (see Fig. 1, 4). As a technical infrastructure, a platform must be available, which combines a batch-oriented analysis with that of a distributed stream mining analysis. Whereas batch-oriented methods are important for training prediction models; stream-oriented ones are compulsory for the actual real-time analysis of incoming data streams. In particular, for the automatic analysis of image and video sensors—as outlined in the use case—dedicated algorithms needs to be available in order to derive structured information from unstructured data. To realize a real-time correlation of complex events with historical process data lead to the third requirements.

Requirement 3. Provide real-time data storage capabilities to correlate and analyze big data collections and streams in terms of high volume, high variety and high velocity.

Based on such correlations, prediction models comprising forecasts about the future progress of business processes can be computed, which allows a proactive reaction to predicted problems (cf. Fig. 1, 4). Several prediction techniques and algorithms can be applied, all of which have to be capable of processing big data and to cope with real-time requirements, which is another challenge in terms of big data analytics. In this regard, sophisticated CEP techniques are required predicting likelihoods of the occurrence of future atomic events that will eventually trigger complex ones. These probability assumptions will realize a predictive complex event detection.

Requirement 4. Derive and continuously adapt (event-based) prediction models.

Based on event-based prediction models forecasts of substantially higher accuracy can be computed, since predictions are not purely based on stochastic, but instead, the actual current state is decisive for the computation (see Fig. 1, 4). Hence, process progressions can be forecasted and certain deviations from planned and required process execution objectives can be proactively detected leading to corresponding system responses.

Requirement 5. Create alerts as responses to predicted deviations from planned process objectives based on calculated forecasts.

Within computed prediction models not only one single possible future process progress will be forecasted, but several multiple ones whose occurrence are depended both on non-influenceable events as well as on influenceable actions. Thus, recommendations or automatic decisions and actions should be provided to positively influence and control the outcome of business processes in accordance to specific process objectives. This will allow to realize a proactive incident mangement in contrast to a reactive incident handling as a prerequisite of a predictive enterprise.

Requirement 6. Derive recommendation and automatic decisions for mitigation actions.

As a result of intelligent algorithms it should further be determined whether changes enacted within one process instance—as a response to detected deviations, defects and problems—will impact other running instances that in turn will affect recommendations and automatic actions. These automatic or manually triggered actions based on real-time event-based predictions eventually realize an intelligent proactive process planning and control (PPC, cf. Fig. 1, 5) and thus the vision of a predictive enterprise.

Requirement 7. Enact sophisticated proactive process adaptations on the basis of computed recommendations and decisions.

3.2 Architecture proposal

This section outlines a functional description of components required for realizing a system for event-based process prediction based on big data analytics (cf. Fig. 2). On the basis of the previously derived requirements, the architecture comprises three inherent layers that are typically considered within analytical processes [Ak13]: a descriptive (cf. requirements 2, 3), a predictive (cf. requirement 4) and a prescriptive layer (cf. requirements 5, 6). On top of them, an adaptation layer (cf. requirement 7) allows for intelligent actions incorporated into business process management engines as responses to prescriptive decisions. As the foundation, an integration layer (cf. requirement 1) realizes the system's physical interweaving into production process and facilities.

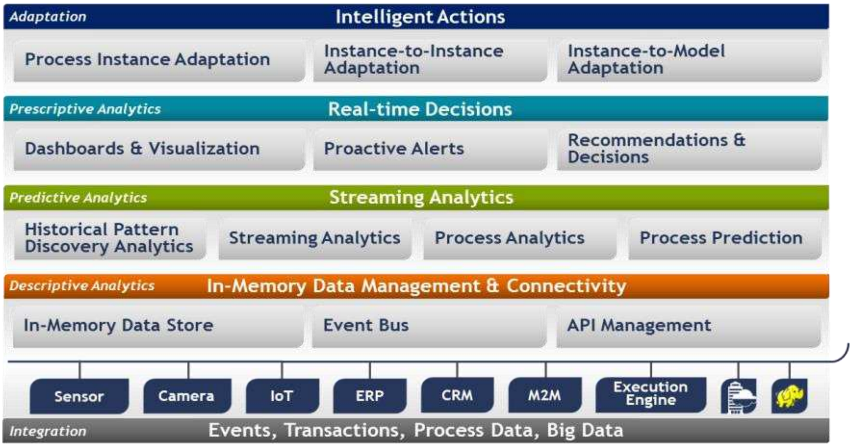

Figure 2: Functional building blocks for a predictive enterprise architecture

Integration layer L0: Event, transactions, process data, Big Data. The foundation of this system is a solid integration platform that connects the system to a company's existing IT infrastructure. In an industry 4.0 context additional adapters for sensor and IoT object integration are required. Due to the wide variety of heterogeneous sensors, which are used throughout enterprises, an initial classification into an ontology-based enterprise data model is required, which also forms the cornerstone to depict context situations in detail. Such a model should consist of at least three subtypes: an enterprise resources model, a business process context model and an enterprise external stimuli model. The use of these models realizes the semantic interoperability for the obtained context data and thus enables in principle an automated analysis.

Descriptive analysis layer L1: In-memory data management & connectivity. Due to the high volume of data and the velocity in which it is generated, an in-memory data management platform is utilized, allowing for distributed in-memory data management with extremely low, predictable latency and fast data access (microseconds) for real-time data handling. An in-memory data store will act as a central point of coordination, aggregating and distribution. Besides data management, events such as alerts or system messages are communicated using an event bus to which components can publish and subscribe. To manage the diverse data sources and connected enterprise systems, an API management component is introduced.

Predictive analysis layer L2: Streaming analytics. Real-time data accessible via the in-memory data management platform can be preprocessed. In particular, for the intelligent evaluation of image and video sensors as outlined in the use case, special algorithms have to be developed in order to derive structured information from unstructured multimedia data. The results are fed back to the in-memory data store. The aggregated data is used for both ex-post and ex-ante analysis. From an ex-post viewpoint, historic data can be analyzed for pattern detection and correlated with respective business process behaviors. Based on these patterns, real-time event detection, such as deviations of production progress from an expected state, can be learned, optimized and applied to monitor real-time data streams. Data analysis components can communicate detected patterns to the event bus on which a complex event processing engine (CEP) is operated to correlate business process relevant events from data analysis results. These can then be used for process prediction.

Prescriptive analysis layer L3: Real-time decisions. In order to enable process owners to make qualified real-time decisions, all relevant data needs to be aggregated and visualized appropriately. Here, dashboarding functionalities similar to current business activity management solutions can be applied. In addition to that, process owners must be notified proactively if a decision is required or when a deviation in the current state of a process instance is detected. Besides pure visualization and notification, a recommendation is generated based on historic process analysis. To understand why a recommendation was made, drill-down functionalities allow to navigate to previous process instance information and enable process users to make qualified decisions.

Adaptation layer L4: Intelligent actions. Based on the data gathered and the resulting process prediction, business processes can either be adapted on an instance base (process instance adaptation) by adjusting the current process execution, or by optimizing the entire process type (instance-to-model). However, adaptations in a process instance can lead to necessary adaptations in other correlated process instances such as supporting or following processes (instance-to-instance). Here, the process owner is also supported. Once the adaptations were decided on, a governance process ensures a consistent transition of changes back into the process execution system(s).

4 Related work

4.1 Process data and context data acquisition and its real-time analysis

In order to detect context situations and progress of ongoing processes, data must be continuously collected. Technically, this can be done by means of physical and virtual sensors, which are often connected in a wireless network. Physical sensors can measure for instance pressure, temperature or vibration and recognize complex structures such as audio signals [SLT07]. Additional data is obtained from IT systems by evaluating for example log files or exchanged messages at the lowest system level [Ch10].

The next step is to analyze this mass of collected data. A common approach is that atomic events, which are detected by sensors, are being first cumulated to more complex ones, which is called CEP. CEP combines methods, techniques and tools to analyze mass of events in real-time and to obtain higher aggregated information from this most elementary data. Such approaches can be found for example in [WDR06], who apply CEP on real-time data streams, and with particular respect to RFID-based data streams in [Wa06]. Often ontologies are used to identify semi-automatically the semantic context of events [TL11]. Operations on data of atomic and complex events require algorithms that are well-adapted to the process context as well as on big data characteristics.

A well-known approach for big data analytics is MapReduce, which decomposes a problem into independent sub-problems that are distributed and solved by so-called mappers. However, traditional forms of MapReduce follow a batch approach, which is inapplicable for data streams. In this regard, the field of research of stream mining and stream analytics has been formed recently [Br11].

4.2 Process forecast calculation and simulation

Predictive analytics refers to the forecast of prospectively occurring events. The necessary mathematical modelling can be done by using historical data. A characteristic example of a simple forecast calculation is the moving average. Such simple models only work if they are mostly independent of external influencing events, which is rarely the case. Especially in business processes several dependencies and influencing factors exist. Thus, modern statistical methods are required to recognize dependencies and patterns in large amounts of data, such as decision trees, univariate and multivariate statistics and data warehouse algorithms [LFC11].

Approaches combining predictive analytics with BPM are coined as Business Process Intelligence (BPI) (cf. [LFC11] for a discussion on existing definitions). Since it is feasibly to increase the context awareness in the era of Internet of Things, it is promising to use data collected by sensors to increase the accuracy of business process predictions. Reference [Fü12] present a first conceptual framework for combining CEP with predictive analytics. Also [JMM12] provide first conceptual works; however, they state that "[e]vent-based or event-driven business intelligence" approaches are only rudimentary researched and find limited integration in software tools. Even if [SMJ13] already combines BAM with BPI, the question on how findings from this connection can be used for adapting and optimizing business processes is still open.

4.3 Business process adaptation and optimization

In traditional BPM approaches, processes are typically adapted and improved on type level at design time after passing a business process lifecycle. This ex post handling—as it is regularly applied in business process controlling and mining—however causes considerable time delay. Since an aggregation of process execution data has to be done in a first place, before processes can eventually be optimized at type level, problematic process executions have already been completed and can no longer be optimized. Thus, scientists from the process mining community explicitly call for a stronger "operational support" [va12]. This requires the existence of immediate feedback channels to the business process execution. Without the possibility of such feedback loops, only ex-post considerations are possible. At this level of run time adaptation and optimization of business processes, first approaches already exist [JMM12].

In the scope of commercial BPM systems, there are some early adopters characterized by the term Intelligent Business Process Management Suites [Si12]. For example, IBM's Business Process Manager provides means to analyze operations on instance level in real-time, to control them and to respond with further process steps in an ad-hoc manner [IBM13]; yet, a fully automation cannot be attested. On the market of real-time CEP, further vendors are present; e.g. Bosch Software Innovations have implemented CEP functionality into BPM solutions [Si12]. However, to consider them as intelligent BPM systems, a more powerful automated analysis of the variety of events must be implemented as a basis to realize event-based predictions and process control [Kr14a].

5 Summary and outlook

The paper at hands examined, based on a case study stemming from the steel manufacturing industry, which production-related data is currently collectable using state of the art sensor technologies forming a potential foundation for a detailed situation awareness and derivation of accurate process forecasts. An analysis of this data however showed that its full potential cannot be utilized without dedicated approaches of big data analytics. By proposing an architecture for implementing predictive enterprise systems, the article provides a working and discussion basis for further research and implementation efforts in big data analytics.

Acknowledgement

This research was funded in part by the German Federal Ministry of Education and Research under grant number 01IS12050 (project IDENTIFY).

References

[Ak13] Akerkar, R: Advanced data analytics in business. In (Akerkar et al., Eds.): Big data computing. CRC Press, Boca Raton, FL, 2013; pp. 373-400.

[BGM87] Benbasat, I.; Goldstein, D. K.; Mead, M.: The case research strategy in studies of information systems. In: MIS Quarterly, vol. 11, no. 3, 1987; pp. 369-386.

[Br11] Brito et al.: Scalable and low-latency data processing with stream MapReduce. In: Proc. IEEE Third Int. Conf. Cloud Computing Technology and Science (CLOUDCOM '11). IEEE Computer Society, 2011; pp. 48–58.

[BL12] Beyer, M. A.; Laney, D.: The importance of "big data": A definition. In: https://www.gartner.com/doc/2057415, 2012; accessed 2014-07-10.

[Ch10] Choi, J.: RFID context-aware systems. In (Turcu et al., Eds.): Sustainable radio frequency identification solutions. InTech, Rijeka, 2010; pp. 307–330.

[DP03] Dubé, L.; Paré, G.: Rigor in information systems positivist case research: current practice, trends and recommendations. In: MIS Quarterly, vol. 27, no. 4, 2003; pp. 597-636.

[Fü12] Fülop et al.: Predictive complex event processing: A conceptual framework for combining complex event processing and predictive analytics. In: Proc. Fifth Balkan Conf. Informatics (BCI '12), ACM, 2012; pp. 26–31.

[IBM13] IBM (Eds.): IBM Business Process Manager Standard. In: http://www-03.ibm.com/software/products/en/business-process-manager-standard, 2013, accessed 2014-07-10.

[JMM12] Janiesch, C.; Matzner, M.; Müller, O.: Beyond process monitoring: A proof-of-concept of event-driven business activity management. In: Business Process Management Journal, vol. 18, no. 4, 2012; pp. 625–643.

[Kr14a] Krumeich, J.; Weis, B.; Werth, D.; Loos, P.: Event-driven business process management: Where are we now? – A comprehensive synthesis and analysis of literature. In: Business Process Management Journal, vol. 20, no. 4, 2014; pp. 615-633.

[Kr14b] Krumeich, J.; Jacobi, S.; Werth, D.; Loos, P.: Towards planning and control of business processes based on event-based predictions. In: Proc. 17th Int. Conf. Business Information Systems (BIS 2014), LNBIP vol. 176, Springer, 2014; pp. 38-49.

[Kr14c] Krumeich, J.; Jacobi, S.; Werth, D.; Loos, P.: Big Data Analytics for Predictive Manufacturing Control – A Case Study from Process Industry. In: Proc. 3rd IEEE Int. Congress on Big Data; in press.

[LFC11] Linden, M.; Felden, C.; Chamoni, P.: Dimensions of business process intelligence. In: Business Process Management Workshops (BPM 2010), LNBIP vol. 66, Springer, 2011; pp. 208–213.

[Lu06] Lundberg, A.: Leverage complex event processing to improve operational performance. In: Business Intelligence Journal, vol. 11, no. 1, 2006; pp. 55–65.

[MCD13] Minelli, M.; Chambers, M.; Dhiraj, A.: Big data, big analytics: Emerging business intelligence and analytic trends for today's business. John Wiley & Sons, Hoboken, NJ, 2013.

[PG11] Pettey, C.; Goasduff, L.: Gartner says between now and year-end 2014, overlooked but easily detectable business process defects will topple 10 global 2000 companies. In: http://www.gartner.com/newsroom/id/1530114, 2011, accessed 2014-07-10.

[SMJ13] Schwegmann, B.; Matzner, M.; Janiesch, C.: A method and tool for predictive event-driven process analytics. In: Proc. 11th Int. Conf. Wirtschaftsinformatik (WI2013), University of Leipzig; pp. 721–735.

[SLT07] Stäger, M.; Lukowicz, P.; Tröster, G.: Power and accuracy trade-offs in sound-based context recognition systems. In: Pervasive and Mobile Computing, vol. 3, no. 3, 2007; pp. 300–327.

[TL11] Taylor, K.; Leidinger, L.: Ontology-driven complex event processing in heterogeneous sensor networks. In: 8th Extended Semantic Web Conf. (ESWC 2011), LNCS vol. 6644, Springer, 2011; pp. 285–299.

[Un12] Unni, K.: Steel manufacturing could use more sensing and analysis. In: http://www.sensorsmag.com/process-industries/steel-manufacturing-could-use-more-sensing-and-analysis-10249, 2012, accessed 2014-07-10.

[va12] van der Aalst et al.: Process mining manifesto. In: Business Process Management Workshops (BPM 2011), LNBIP vol. 99, Springer, 2012; pp. 169–194.

[Wa06] Wang, F.; Liu, S.; Liu, P.; Bai, Y.: Bridging physical and virtual worlds: Complex event processing for RFID data streams. In: 10th Int. Conf. Extending Database Technology (EDBT 2006), LNCS vol. 3896, Springer, 2006; pp. 588–607.

[WDR06] Wu, E.; Diao, Y.; Rizvi, S.: High-performance complex event processing over streams. In: Proc. Int. Conf. Management of Data (SIGMOD '06), ACM, 2006; pp. 407–418.