Editorial: “Business Process Intelligence: Connecting Data and Processes”

WIL VAN DER AALST, Technische Universiteit Eindhoven

J. LEON ZHAO, City University of Hong Kong

HARRY JIANNAN WANG, University of Delaware

This introduction to the special issue on Business Process Intelligence (BPI) discusses the relation between data and processes. The recent attention for Big Data illustrates that organizations are aware of the potential of the torrents of data generated by today's information systems. However, at the same time, organizations are struggling to extract value from this overload of data. Clearly, there is a need for data scientists able to transform event data into actionable information. To do this, it is crucial to take a process perspective. The ultimate goal of BPI is not to improve information systems or the recording of data; instead the focus should be in improving the process. For example, we may want to aim at reducing costs, minimizing response times, and ensuring compliance. This requires a “confrontation” between process models and event data. Recent advances in process mining allow us to automatically learn process models showing the bottlenecks from “raw” event data. Moreover, given a normative model, we can use conformance checking to quantify and understand deviations. Automatically learned models may also be used for prediction and recommendation. BPI is rapidly developing as a field linking data science to business process management. This article aims to provide an overview thereby paving the way for the other contributions in this special issue.

Categories and Subject Descriptors: H.2.8 [Database Applications]: Data Mining and; H.4.1 [Office Automation]: Workflow Management

General Terms: Design, Management, Performance

Additional Key Words and Phrases: Process mining, business process intelligence, process modeling, performance analysis, compliance checking

ACM Reference Format:

Wil van der Aalst, J. Leon Zhao, and Harry Jiannan Wang. 2015. Editorial: “Business process intelligence: Connecting data and processes”. *ACM Trans. Manag. Inform. Syst.* 5, 4, Article 18e (April 2015), 7 pages.
DOI: <http://dx.doi.org/10.1145/2685352>

1. TOWARDS PROCESS SCIENCE

Our capabilities to store and process data have been increasing exponentially since the 1960s [Hilbert and Lopez 2011; Manyika 2011]. Today, data are collected about anything, at any time, and at any place. Moreover, the emerging “internet of things” and the rapid growth of mobile devices provide new forms of data much closer to our everyday reality (social, physical, and economical). Figure 1 illustrates the different types of data being collected in today's world cumulating into the Internet of Events (IoE). As a result, organizations are seeking ways to exploit the information readily available

Authors' addresses: W. M. P. van der Aalst, Department of Mathematics & Computer Science, Technische Universiteit Eindhoven; email: W.M.P.v.d.Aalst@tue.nl; J. L. Zhao, Department of Information Systems, City University of Hong Kong; email: jlzhao@cityu.edu.hk; H. J. Wang, Department of Accounting and Management Information Systems, University of Delaware; email: hjwang@udel.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2015 ACM 2158-656X/2015/04-ART18e \$15.00

DOI: <http://dx.doi.org/10.1145/2685352>

Fig 1. The “Internet of Events” is composed of the “Internet of Content,” the “Internet of People,” the “Internet of Things,” and the “Internet of Places” [Van der Aalst 2014a].

[Van der Aalst 2014a]. The “Big Data” hype illustrates this trend. Moreover, a new discipline has emerged: Data Science (DS). DS aims to collect, analyze, and interpret data from a variety of sources (social interaction, business processes, cyberphysical systems) [Van der Aalst 2014a; Chiang et al. 2012]. To turn data into actionable information, a comprehensive understanding of the context of the data and the ability to mine and visualize large amounts of data are essential. Therefore, the role of the data scientist is:

- to ensure that the right data is recorded and stored (provenance) and to be able to extract relevant data in a complex IT landscape,
- to use a wide variety of analytical techniques to extract value (models, insights, predictions, recommendations, and visualizations) from data, and
- to present the results to end users, assist in the interpretation of results, and be able to put it in context.

Just like computer science emerged from mathematics when computers became abundantly available in the 1980s, we can now see that today’s data tsunami is creating the need for data scientists. Although data is enabling new and exciting forms of analysis, one should not forget that in the end we would like to improve operational business processes based on data and not just collect data. Therefore, we relate DS to Business Process Management (BPM) [Van der Aalst 2013a].

BPM is the discipline that combines knowledge from information technology and knowledge from management sciences and applies this to operational business processes [Van der Aalst 2013a; Dumas et al. 2013; Reichert and Weber 2012; Weske 2007; Stohr and Zhao 2001]. BPM received considerable attention in recent years because of its potential for significantly increasing productivity and saving costs. Moreover, today there is an abundance of BPM systems. These systems are realized using generic software products that are driven by explicit process designs to enact and manage operational business processes. However, BPM is definitely not limited to BPM systems. BPM has a broader scope: from process automation and process analysis to operations management and the organization of work [Van der Aalst 2013a; Zhao et al. 2000]. On the one hand, BPM aims to improve operational business processes, possibly without the use of new technologies. For example, by modeling a business process and analyzing it using simulation, management may get ideas on how to reduce costs while improving service levels. On the other hand, BPM is often associated with software to manage,

Fig. 2. Overview of the different skills relevant for today's data scientist [Van der Aalst 2014b] and the relation to BPM. The highlighted area demarcates the focus of this special issue.

control, and support operational processes. Traditional BPM approaches do not exploit event data in a systematic way: the focus of BPM practitioners is often on modeling processes and measuring key performance indicators. However, it is obvious that the BPM discipline will become more data driven.

Figure 2 shows the focus of this special issue. Next to the typical DS skills, the connection to BPM is emphasized in the figure. Next to data scientists, there is an urgent need for process scientists, who are analysts that combine BPM skills with DS skills to improve processes by exploiting event data.

2. BUSINESS PROCESS INTELLIGENCE

Business Process Intelligence (BPI) is on the interface between DS and BPM. BPI is closely related to topics such as Business Intelligence (BI) and Process Mining (PM).

Boris Evelson of Forrester Research defines BI as “a set of methodologies, processes, architectures, and technologies that transform raw data into meaningful and useful information used to enable more effective strategic, tactical, and operational insights and decision making.” Although this definition does not exclude a process focus, BI methodologies and tools are typically not process aware.

PM aims to discover, monitor and improve real processes by extracting knowledge from event logs readily available in today's information systems as well as from corporate policy manuals [Van der Aalst 2011; Li et al. 2010; Wang et al. 2009]. The starting point for process mining is an *event log*. Each event in such a log refers to an activity (i.e., a well-defined step in some process) and is related to a particular case (i.e., a process instance). The events belonging to a case are ordered and can be seen as one “run” of the process. Event logs may store additional information about events. In fact, whenever possible, PM techniques use extra information such as the resource (i.e., person or device) executing or initiating the activity, the timestamp of the event, or data elements recorded with the event (e.g., the size of an order).

Event logs can be used to conduct three basic types of PM [Van der Aalst 2011]. The first type of PM is *discovery*. A discovery technique takes an event log and produces a model without using any a priori information. Process discovery is the most prominent PM technique. For many organizations, it is surprising to see that existing techniques are indeed able to discover real processes merely based on example behaviors stored in event logs. The second type of PM is *conformance*. Here, an existing process model is compared with an event log of the same process. Conformance checking can be used to check if reality, as recorded in the log, conforms to the model and vice versa. The third type of PM is *enhancement*. Here, the idea is to extend or improve an existing process model, thereby using information about the actual process recorded in some event log. Whereas conformance checking measures the alignment between model and reality, this third type of PM aims at changing or extending the a priori model. For instance, by using timestamps in the event log, one can extend the model to show bottlenecks, service levels, and throughput times.

Orthogonal to three types of PM (process discovery, conformance checking, and enhancement), one can look at the *offline* setting (i.e., historic data on completed cases) or the *online* setting (i.e., also consider event data on still running cases that can still be affected by corrective actions). For example, discovered and subsequently enhanced process models can be used for *operational support* and thus enable *predictive analytics* (e.g., predicting the remaining flow time of a case, recommending a suitable activity or resource).

Another application of the three basic types of PM is comparative process mining, which uses process cubes [Van der Aalst 2013b] to compare different subprocesses or groups of cases. The focus is no longer on a single process and a homogenous set of cases. Comparative PM can be used to identify differences between two branches of the same organization or two groups of customers.

Figure 3 positions process mining as a bridge between process-centric approaches and more data-centric approaches like data mining and machine learning. It is difficult to distinguish between PM and BPI. BPI could be defined as BI with particular attention for processes. However, this is exactly what process mining is about. Therefore, we use the terms “process mining” and “BPI” interchangeably.

DS, consisting of data mining and machine learning, is concerned with four main questions [Van der Aalst 2014a]:

- Reporting: *What happened?*
- Diagnosis: *Why did it happen?*
- Prediction: *What will happen?*
- Recommendation: *What is the best that can happen?*

The first two questions aim at understanding the past. The last two questions use knowledge learned from past experiences to say something about the future. Obviously, all four types of questions can be posed in the context of process improvement. Typical BPI questions are:

- What is the process that people in an organization really follow?
- Where are the bottlenecks in my business process and how to remove them?
- Where do organizations deviate from the expected or idealized process?
- What are the “highways” in my process?
- What factors are influencing the service level provide to customers?
- Can we predict problems (delay, deviation, risk, etc.) for running cases?
- Can we recommend countermeasures?
- How to redesign the process given a set of goals?
- Are noncompliant cases more costly?

Fig. 3. Positioning PM as the linking pin between BPM and data-driven approaches like data mining.

PM/BPI software aims to answer these questions. Mainstream BI software like IBM Cognos Business Intelligence, Oracle Business Intelligence, SAP BusinessObjects, WebFOCUS (Information Builders), Microsoft SQL Server, MicroStrategy, NovaView (Panorama Software), QlikView (QlikTech), SAS Enterprise Business Intelligence, TIBCO Spotfire Analytics, Jaspersoft, and Pentaho BI Suite are not process oriented and, therefore, less suitable for answering the above questions. PM tools like ProM (open source), Disco (Fluxicon), Perceptive Process Mining, ARIS Process Performance Manager (Software AG), QPR ProcessAnalyzer, Celonis Discovery, Interstage Process Discovery (Fujitsu), Discovery Analyst (StereoLOGIC), and XMAalyzer (XMPPro) aim at answering at least a subset of the aforementioned questions. Existing tools provide widely applicable and very robust types of analysis. However, PM/BPI is a relatively young discipline with many open questions.

3. IN THIS SPECIAL ISSUE ON BPI

There is a growing interest in BPI. This is not surprising, as it is clear that management information systems need to become more process-centric while utilizing the event data at hand. This special issue presents state-of-the-art results in BPI. Out of 23 submissions, 8 articles were accepted after an initial screening and multiple rounds of reviews and revisions. A brief overview of these articles is given next.

Two articles are related to the process discovery aspect of BPI. In "On the Discovery of Declarative Control Flows for Artful Processes," Di Ciccio and Mecella focus on the study of artful processes, where decisions are based on the experience, intuition, and knowledge of the process actors. Compared with operational business processes, artful processes are much less structured, very flexible, and even completely unknown a priori. The authors propose an algorithm to automatically discover artful processes, which

is critical for better understanding knowledge-intensive artful processes and providing improvements in many scenarios. Wang et al. present a unique application of PM to knowledge management for online Q&A communities in “An Analytical Framework for Understanding Knowledge Sharing Processes in Online Q&A Communities.” They innovatively assign dialog act tags to posts in discussion threads and treat each thread as a process instance of dialog act tags. Then, different process patterns are discovered and analyzed to identify patterns that are more likely to lead to helpful knowledge sharing.

Conformance checking is another important part of BPI. In “Compliance Checking of Organizational Interactions,” Jiang et al. look into the problem of regulatory process compliance checking. They propose a normative structure named *Norm Nets* (NNs) to formally model narrative process regulations and develop mechanisms to automatically check process compliance by mapping NNs to colored Petri nets. Their approach enables the systematic analysis of process compliance with dynamic interrelated regulations and business partner interactions.

The next two articles investigate the intelligent task-agent assignment problem in BPI. In “Process Analytics Approach for R&D Project Selection,” Silva, Ma, and Yang, propose a research analytics framework to facilitate the reviewer assignment task in the scientific proposal selection process. Text mining and clustering techniques are used to cluster proposals and social network analysis and bibliometric analysis are applied to generate reviewers’ network with no conflict of interests and their research qualification scores. Based on that, a reviewer assignment function is developed to optimize the research relevance between proposals and reviewers. In “Mining Agents’ Goals in Agent-Oriented Business Processes,” Yan et al. propose a framework to mining process agents’ goals from process event logs. Belief-Desire-Intention (BDI) logic is used to formally represent agents’ goals and a goal-mining algorithm is developed. This framework can be used to detect and analyze inconsistencies between the design and execution of business processes in terms of agents’ goal and overall process goal, which facilitate process improvements and redesign.

The data perspective of BPI is discussed in the following two articles. In “Classification Models for RFID-Based Real-Time Detection of Process Events in the Supply Chain: An Empirical Study,” Keller, Thiesse, and Fleisch aim to address the noisy process data problem in RFID-enabled supply chains. A number of data mining techniques have been applied to filter and aggregate RFID raw data and a customized decision-rule-based classifier is developed. The method proposed in this article can greatly improve the process data quality and thus pave the way for better real-time process monitoring and optimization. Business process integration has been an important issue for B2B collaboration. In “A Dataflow Perspective for Business Process Integration,” Guo, Sun, and Vogel propose a dataflow formalism to facilitate interorganizational process integration, which provides a set of constructs to model interorganizational data exchange and a method to calculate public dataset to facilitate collaboration.

Lastly, Partington et al. present a comprehensive case study of applying process mining to the administrative and clinical data from four healthcare providers in “Process Mining for Clinical Processes: A Comparative Analysis of Four Australian Hospitals.” This case study provides many useful insights into conducting practical process mining projects.

The guest editors would like to thank all reviewers who provided timely and constructive reviews for this special issue; without their dedication and expertise, this special issue would not be possible. We also want to thank Hsinchun Chen for his kind guidance and strong support in the process of developing this special issue. Last but not least, we extend our thanks to Cathy Larson for her help in every step of the submission and review process.

REFERENCES

- W. M. P. van der Aalst. 2011. *Process Mining: Discovery, Conformance and Enhancement of Business Processes*. Springer-Verlag, Berlin.
- W. M. P. van der Aalst. 2013a. Business process management: A comprehensive survey. *ISRN Software Engineering* 2013, Article 507984, 37 pages. doi:10.1155/2013/507984
- W. M. P. van der Aalst. 2013b. Process cubes: Slicing, dicing, rolling up and drilling down event data for process mining. In *Proceedings of the Asia Pacific Conference on Business Process Management (AP-BPM'13)*, M. Song, M. Wynn, and J. Liu (eds.), volume 159 of Lecture Notes in Business Information Processing, vol. 159. Springer-Verlag, Berlin, 1–22.
- W. M. P. van der Aalst. 2014a. Data scientist: The engineer of the future. In *Proceedings of the I-ESA Conference*, K. Mertins, F. Benaben, R. Poler, and J. Bourrieres (eds.), Enterprise Interoperability, vol. 7. Springer-Verlag, Berlin, 13–28.
- W. M. P. van der Aalst. 2014b. Process Mining: Data Science in Action. Coursera course. Retrieved from <https://www.coursera.org/course/procmin>.
- R. H. L. Chiang, P. Goes, and E. A. Stohr. 2012. Business intelligence and analytics education, and program development: A unique opportunity for the information systems discipline. *ACM Transactions on Management Information Systems* 3, 3, 12:1–12:13.
- M. Dumas, M. La Rosa, J. Mendling, and H. Reijers. 2013. *Fundamentals of Business Process Management*. Springer-Verlag, Berlin.
- M. Hilbert and P. Lopez. 2011. The World's Technological Capacity to Store, Communicate, and Compute Information. *Science* 332, 6025, 60–65.
- J. Li, H. J. Wang, Z. Zhang, and J. L. Zhao. 2010. A policy-based process mining framework: mining business policy texts for discovering process models. *Information Systems and e-Business Management* 8, 2 (March), 169–188.
- J. Manyika, M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, and A. Byers. 2012. *Big Data: The Next Frontier for Innovation, Competition, and Productivity*. McKinsey Global Institute.
- M. Reichert and B. Weber. 2012. *Enabling Flexibility in Process-Aware Information Systems: Challenges, Methods, Technologies*. Springer-Verlag, Berlin.
- E. A. Stohr and J. L. Zhao. 2001. Workflow automation: Overview and research issues. *Information Systems Frontiers* 3, 3, 281–296.
- H. J. Wang, J. L. Zhao, and L. J. Zhang. 2009. Policy-driven process mapping (PDPM): Discovering process models from business policies. *Decision Support Systems* 48, 1, 267–281.
- M. Weske 2007. *Business Process Management: Concepts, Languages, Architectures*. Springer-Verlag, Berlin.
- J. L. Zhao, A. Kumar, and E. A. Stohr. 2000. Workflow-centric information distribution through email. *Journal of Management Information Systems* 17, 3, 45–72.

Received September 2014; accepted October 2014