Enabling Process Innovation via Deviance Mining and Predictive Monitoring

Marlon Dumas and Fabrizio Maria Maggi

Abstract

A long-standing challenge in the field of business process management is how to deal with processes that exhibit high levels of variability, such as customer lead management, product design or healthcare processes. One thing that is understood about these processes is that they require process designs and support environments that leave considerable freedom so that process workers can readily deviate from pre-established paths. At the same time, consistent management of these processes requires workers and process owners to understand the implications of their actions and decisions on the performance of the process. We present two emerging techniques—deviance mining and predictive monitoring—that leverage information hidden in business process execution logs in order to provide guidance to stakeholders so that they can steer the process towards consistent and compliant outcomes and higher process performance. Deviance mining deals with the analysis of process execution logs offline in order to identify typical deviant executions and to characterize deviance that leads to better or to worse performance. Predictive monitoring meanwhile aims at predicting—at runtime—the impact of actions and decisions of process participants on the probable outcomes of ongoing process executions. Together, these two techniques enable evidence-based management of business processes, where process workers and analysts continuously receive guidance to achieve more consistent and compliant process outcomes and a higher performance.

M. Dumas (✉) • F.M. Maggi
Institute of Computer Science, University of Tartu, J. Liivi 2, Tartu 50409, Estonia
e-mail: marlon.dumas@ut.ee; f.m.maggi@ut.ee

J. vom Brocke, T. Schmiedel (eds.), *BPM – Driving Innovation in a Digital World*,
Management for Professionals, DOI 10.1007/978-3-319-14430-6_10

1 Introduction

Traditional approaches to business process management are geared towards regular and predictable processes, where there is in essence one primary and well-understood way of performing a process, with relatively few and well-scoped variations. It is widely accepted that these approaches do not fit the requirements of more flexible processes, such as customer lead management processes, product design processes, patient treatment and related healthcare processes. Moreover, when said traditional approaches are pushed down to the level of fine-grained automation, bottom-up process innovation is stiffed, as process workers are in essence expected to follow a scripted process that tells them what to do, what data to gather and when and how to make decisions.

A number of approaches for flexible process management have emerged in recent years. One family of such approaches encompassing adaptive case management (Swenson, 2010) is based on *flexibility by underspecification* (Schonenberg, Mans, Russell, Mulyar, & van der Aalst, 2008). The idea is to underspecify the process at design-time so that process workers have the freedom to perform tasks in various ways, at almost any point in time, to leave tasks incomplete, and to collect different subsets of data at different points in the process, with as few restrictions as possible. However, while these approaches provide flexibility to process workers and allow managers to scope this flexibility, they do not per se support managers and workers in deciding what to do and when.

In this paper, we outline two emerging techniques that complement flexible process management approaches by identifying patterns of activities associated with positive or negative deviance (deviance mining) and by continuously estimating the probability that ongoing process executions may lead to undesirable outcomes (predictive monitoring). Together, these techniques turn a business process support system into a recommender system that not only enables process innovation, but also channels it towards more consistent and compliant outcomes and higher performance.

In the following section, we outline the architecture of a monitoring system integrating deviance mining and predictive monitoring. We then present each of these techniques in turn and close the paper with a discussion on future challenges on the way to the adoption of these techniques in practice.

2 Business Process Monitoring Architecture

Figure 1 sketches a high-level architecture of a business process system that supports deviance mining and predictive monitoring. The figure highlights that both techniques take as input a log of completed business process execution traces and a set of business constraints. In this context, a business constraint is any condition that can be evaluated to be true or false over every completed case of the process. A business constraint may be a service level constraint such as "every simple insurance claim should be resolved at most 2 weeks after all required

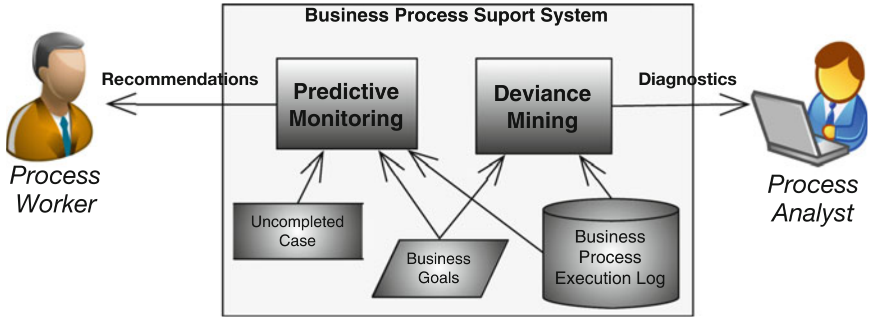

Fig. 1 Business process support with deviance mining and predictive monitoring

documents have been submitted" or it may be a compliance rule such as "every invoice above a given amount should be approved before being paid".

Given an execution log and a set of business constraints, deviance mining allows a process analyst to obtain a diagnostic that explains why certain cases deviate from the intended behavior, meaning that they outperform or underperform the given service level objectives, or fail to fulfill the compliance rules. The diagnostic can take different forms as discussed later, but in any case its purpose is to enable the analyst to identify process improvement opportunities.

Predictive monitoring on the other hand produces recommendations for process workers during the execution of a case. These recommendations refer to a specific (uncompleted) case of the process and tell the user what is the impact of a given action on the probability that the case at hand will fail to fulfill the relevant performance objectives or compliance rules. In particular, predictive monitoring can be used to raise alerts when certain actions are likely to lead to violations of business constraints. In this way, rather than prescribing what to do, the business process support system acts as a compliance monitoring and recommender system, raising flags whenever certain actions heighten the probability of undesirable deviations.

3 Deviance Mining

Business process deviance mining is a family of process mining techniques aimed at analyzing business process execution logs in order to explain the reasons why a business process deviates from its normal or expected execution. Such deviations may be of a negative or of a positive nature—cf. theory of positive deviance (Spreitzer & Sonenshein, 2004). Positive deviance corresponds to executions that lead to high process performance, such as achieving positive outcomes with low execution times, low resource usage or low costs. Negative deviance refers to the executions of the process with low process performance or with negative outcomes or compliance violations.

A concrete example of negative deviance mining in a large Australian insurance company has been reported by Suriadi, Wynn, Ouyang, ter Hofstede, and van Dijk (2013). In this case, a team of analysts sought to find the reasons why certain simple claims that should normally be handled within a few days were taking substantially longer to be resolved. In other words, they needed to understand the difference between "simple quick claims" that were handled in less than x days and "simple slow claims" that took longer to be handled. The team decided to use *delta-analysis* in this context. In other words, two sub-logs were extracted: one containing only traces of "simple slow claims" and another containing "simple quick claims". A process model was then discovered from each of the sub-logs separately—using the Disco process discovery tool[1]—and the resulting models were manually compared. It was found that certain paths and cycles were considerably more frequent for slow claims than for quick claims. The authors also found that two activity metrics were helpful in discriminating slow versus quick claims, namely "average number of occurrences of a given activity X (per case)" and "percentage of cases where a given activity X appears at least once". These two metrics basically indicate how often an activity is executed once, multiple times, or skipped. By calculating these metrics for each activity, the team found it was possible to track down the sources of delays to specific activities in the process, which then allowed them to extract specific process improvement recommendations.

A similar idea was applied by Sun, Du, Chen, Khoo, and Yang (2013) in the context of software defect handling processes in a large commercial bank in China. The authors took a log of over 2,600 defect reports of 4 large software development projects and examined the differences between reports that had led to a correct resolution (normal cases) versus those that had led to complaints by users (anomalous cases). The team defined a number of features to distinguish between normal and anomalous complaints, including "number of occurrences of a given activity X in a case" (for each possible activity X) and "number of occurrences of activity B after an activity A". Since there are many such combinations (A,B) and to avoid having too large a number of features, the authors employed a discriminative itemset mining technique to identify the most relevant of such pairs (A,B). Based on the resulting features, the authors constructed a decision tree that classifies cases into normal and anomalous. Finally, from the decision tree they extracted a set of seven rules that explained the majority of the anomalous cases, thus leading to potential improvement ideas.

Another case study showing the potential of deviance mining, this time in the healthcare domain, is reported by Lakshmanan, Rozsnyai, and Wang (2013). Here, the team applied deviance mining techniques to understand the differences between cases leading to positive clinical outcomes versus those leading to negative outcomes in the process of congestive heart failure treatments at a large US-based healthcare provider. In this case, the team employed a combination of delta-analysis (as in the Australian insurer case study mentioned above) with

[1] http://www.fluxicon.com/

sequence mining techniques. Specifically, the authors used sequence mining to detect typical sequences of activities (e.g. activity B occurring some time after activity A) that were common for positive outcomes but not common for negative ones or vice-versa. The observations made using sequence mining were complemented with additional observations obtained by comparing a process model discovered from cases with positive outcomes with the model obtained for cases with negative outcomes. In this way, the authors extracted a number of pathways and patterns that discriminate between positive and negative cases.

In yet another case study, Bose and van der Aalst (2013) apply a technique for extracting patterns of activities that discriminate between event traces associated to malfunctions (versus normal traces) in components of remotely monitored X-ray machines. The techniques they employ fall under a wider family of techniques known as *discriminative sequence mining techniques* (Lo, Cheng, & Lucia, 2011), which in a nutshell allow one to extract sequential patterns that discriminate between multiple types of sequences (e.g. sequences with positive outcomes versus sequences with negative outcomes).

Finally, Swinnen, Depaire, Jans, and Vanhoof (2012) present a case study in a large European financial institution where analysts sought to understand the reasons for deviations from normative pathways in a procurement process. A dataset of close to 30,000 cases of a procurement process were extracted from the institution's SAP system. Using a process discovery tool, it was found that about 29 % of cases corresponded to deviations from expected (mandated) pathways. Association rule mining was then applied to extract rules to characterize deviant cases. It was found that a total of ten rules could explain almost all deviant cases. Analysis of these rules by business experts revealed characteristic situations where control points in the procurement process were being bypassed, leading to potential weaknesses in the process. The case study highlights that even highly standardized business processes, such as a procurement process, are characterized by frequent deviations. Analysis of such deviations can help to identify and to rectify potential weaknesses. In other words, deviance mining is not only relevant in ad hoc processes, but equally well in standardized scenarios.

The above case studies show that delta-analysis in combination with association rule and sequence mining—particularly discriminative sequence mining—provide a basis for discovering patterns of activities that distinguish negative deviance from normal cases. Setiawan and Sadiq (2013) show that similar techniques can be applied to distinguish positive deviance (i.e. high-performing cases) from normal cases. Specifically, Setiawan and Sadiq propose a method to identify activities or workers that are associated with positive deviance in a business process. The method assumes that the analyst is interested in understanding positive deviance with respect to a given set of process performance measures. The first step of the method is to identify cases that are associated with higher performance according to the given performance measures. Given that there are typically multiple, and sometimes contradicting, process performance measures, the method uses the notion of Pareto frontier to identify cases that strike the best tradeoffs between these multiple measures. In a second step, the method analyzes the performance of process workers to determine which process workers perform better for different types of activities. Finally, in the third step the method analyzes the impact of the performance of different activities on the overall performance of the process. The outcome is a characterization of the best cases of the process with respect to process workers and activity performance observed for these cases, which can be used to identify best practices.

A similar method for positive deviance analysis is outlined by Tregear (2013). Tregear's so-called $^{+}$D method shares common points with that of Setiawan and Sadiq. As in the latter method, the $^{+}$D method starts by determining how success is measured, which entails selecting and defining performance measures. In a second step, data is collected with respect to the chosen performance measures. This is followed (third and fourth steps) by identifying samples of exceptional performance from the data and analyzing the data in an exploratory manner in order to identify what factors might underpin the identified exceptional performance (positive deviance). In a fifth step, statistical tests are used to identify correlations and causal links between the identified factors and positive deviance. This last step leads to the formulation of hypotheses to explain positive deviance. In a sixth step, controlled tests are undertaken in order to validate the hypotheses. Finally, the validated hypotheses are used as a basis to formulate new practices that are implemented and communicated across all relevant process stakeholders.

4 Predictive Monitoring

The execution of business processes is generally subject to internal policies, norms, best practices, regulations, and laws. For example, a doctor may only perform a certain type of surgery a pre-operational screening is carried out beforehand. Meanwhile, in a sales process, an order can be archived only after the customer has confirmed receipt of all ordered items.

For this reason, compliance monitoring is an everyday imperative in many organizations. Accordingly, a range of research proposals have addressed the problem of monitoring business processes with respect to business constraints (Birukou et al., 2010; Ly, Rinderle-Ma, Knuplesch, & Dadam, 2011; Maggi, Montali, Westergaard, & van der Aalst, 2011; Maggi, Montali, & van der Aalst, 2012; Weidlich et al., 2011). Given a process model and a set of *business constraints*, these techniques provide a basis to monitor ongoing executions of a process (a.k.a. *cases*) in order to assess whether they comply with the constraints in question. However, these monitoring approaches are *reactive*, in that they allow users to identify a violation only *after it has occurred* rather than supporting them in *preventing* such violations in the first place.

Predictive Monitoring (Maggi, Di Francescomarino, Dumas, & Ghidini, 2013) is an emerging paradigm based on the continuous generation of predictions and recommendations on what activities to perform and what input data values to provide, so that the likelihood of violation of business constraints is minimized.

In this paradigm, a user specifies a *business goal* in the form of business rules.[2] Based on an analysis of execution traces, the idea of predictive monitoring is to continuously provide the user with estimations of the likelihood of achieving each business goal for a given case. Such predictions generally depend both on: (1) the sequence of activities executed in a given case; and (2) the values of data attributes after each activity execution in a case.

As an example, consider a doctor who needs to choose the most appropriate therapy for a patient. Historical data referring to patients with similar characteristics can be used to predict what therapy will be the most effective one and to advise the doctor accordingly. Meanwhile, in the context of a business process for managing loan applications, the applicant can be advised on the combinations of the loan amount and the length of loan that are the most likely to lead to acceptance of the application, given contextual information about the application and the personal data of the applicant (e.g., age, salary, etc.).

In a previous work (Maggi et al., 2013), we have put forward a specific framework for predictive monitoring aimed at generating predictions at runtime based on user-defined business goals. This technique estimates, for each enabled activity in an ongoing case, and for every data input that can be given to this activity, the probability that the execution of the activity with the corresponding data input will lead to the fulfillment of the business goal. To this aim, we apply a combination of simple string matching techniques with decision tree learning.

An approach for the prediction of abnormal terminations of business processes has been presented by Kang, Kim, and Kang (2012). Here, a fault detection algorithm (local outlier factor) is used to estimate the probability of a fault occurring. Alarms are provided for an early notification of probable abnormal terminations, in order to prevent risks rather than to simply react to them. Castellanos, Salazar, Casati, Dayal, and Shan (2005) present a business operations management platform equipped with time series forecasting functionalities. This platform allows for predictions of metric values on running process instances as well as for predictions of aggregated metric values of future instances (e.g., the number of orders that will be placed next Monday).

Other predictive monitoring techniques have been proposed that are targeted at generating predictions and recommendations focused on temporal aspects. For example, van der Aalst, Schonenberg, and Song (2011) propose predictive monitoring techniques for estimating case completion times and deadline violations based on annotated transition systems encoding temporal information extracted from event logs. Meanwhile, Folino, Guarascio, and Pontieri (2012) propose a predictive clustering approach in which context-related execution scenarios are discovered and modeled through state-aware performance predictors. Finally, Rogge-Solti and Weske (2013) introduce a method for predicting the remaining execution time of a process based on stochastic Petri nets.

---

[2] In line with the forward-looking nature of predictive monitoring, we use the term *business goal* rather than *business constraint* to refer to the monitored properties.

Other approaches focus on generating predictions to reduce risks. Conforti, de Leoni, Rosa, and van der Aalst (2013) for example present a technique to support process participants in making risk-informed decisions by traversing decision trees generated from the logs of past process executions. In a similar vein, Pika, Aalst, Fidge, Hofstede, and Wynn (2013) propose an approach for predicting time-related process risks by identifying indicators observable in event logs that highlight the possibility of deadline transgression.

Finally, initial case studies of predictive process monitoring in the field of transportation and logistics are presented by Metzger, Franklin, and Engel (2012) and Feldman, Fournier, Franklin, and Metzger (2013). These case studies show in particular how predictive process monitoring can be used to explain and predict "late show" events in a transportation process. Here, a "late show" refers to a delay between expected and actual time of delivering the goods to a carrier (e.g. airline). In this case study, standard statistical techniques are used to find correlations between "late show" events and external variables such as weather conditions or road traffic. The uncovered correlations are then used to define complex event processing rules that detect situations where "late show" events are likely to occur. A challenge for predictive process monitoring in this setting is that transportation processes are generally not "case-based" because goods from different customers are often aggregated and dis-aggregated at different points in the process. In other words, multiple "cases" of a transportation process will typically merge and split at runtime and thus delays affecting one delivery might end up affecting others.

5 Discussion and Outlook

Process innovation requires business process support systems that depart from traditional normative approaches to business process execution. Rather than imposing a specific and preconceived course of action, business process support systems are increasingly required to provide process workers with sufficient autonomy to enable continuous adaptation and innovation.

In this setting, we position deviance mining and predictive monitoring as two keystones in modern business process support systems. Predictive monitoring and deviance mining are related since they both try to identify deviations with respect to expected behavior. However, while deviance mining tries to do this off-line (by analyzing process logs), predictive monitoring provides feedback on-the-fly to prevent violations. Together, these techniques turn a business process support system into a recommender system that provides guidance to process analysts and process workers, helping them to recognize actions and decisions that typically drive a process towards desired outcomes and higher performance.

While deviance mining and predictive monitoring techniques are still in their infancy, they are already applicable in real-life scenarios as evidenced by the several case studies discussed in this paper. Going forward, we foresee more sophisticated and automated techniques for deviance mining emerging. For example, techniques for extraction of predictive sequence patterns (Xing, Pei, Dong, &

Yu, 2008) could find useful applications in the context of both deviance mining and predictive monitoring. Another family of techniques that could find applications in this space is that of discriminatively trained hidden Markov models (Collins, 2002). These techniques produce models that could be applied to extract probabilities of an ongoing case falling into either a “deviant” or “normal” category.

In parallel to technical developments, we foresee more sophisticated case studies being carried out, where deviance mining and predictive monitoring are applied not only in the context of specific improvement initiatives, but on an ongoing basis as part of continuous process improvement programs. To achieve this goal, analysts will benefit from more methodological guidance and more user-friendly tool support that allows them to readily apply these techniques on potentially large and complex business process execution logs.

Acknowledgments This work is supported by ERDF via the Estonian Centre of Excellence in Computer Science.

References

Birukou, A., D’Andrea, V., Leymann, F., Serafinski, J., Silveira, P., Strauch, S., et al. (2010). An integrated solution for runtime compliance governance in SOA. In *Proceedings of international conference on service-oriented computing (ICSOC)* (Vol. 6470). Berlin: Springer.

Bose, R. P. J. C., & van der Aalst, W. M. P. (2013). Discovering signature patterns from event logs. In *Proceedings of the IEEE symposium on computational intelligence and data mining (CIDM)* (pp. 111–118). IEEE.

Castellanos, M., Salazar, N., Casati, F., Dayal, U., & Shan, M.-C. (2005). Predictive business operations management. In *Proceedings of the workshop on databases in networked information systems (DNIS)* (pp. 1–14). Springer.

Collins, M. (2002). Discriminative training methods for hidden Markov models: Theory and experiments with perceptron algorithms. In *Proceedings of the ACL conference on empirical methods in natural language processing* (pp. 1–8). Philadelphia, PA: Association for Computational Linguistics.

Conforti, R., de Leoni, M., Rosa, M. L., & van der Aalst, W. M. P. (2013). Supporting risk-informed decisions during business process execution. In *Proceedings of international conference on advanced information systems engineering (CAiSE)* (pp. 116–132). Berlin: Springer.

Feldman, Z., Fournier, F., Franklin, R., & Metzger, A. (2013). Proactive event processing in action: A case study on the proactive management of transport processes. In *Proceedings of ACM international conference on distributed event-based systems (DEBS)* (pp. 97–106). ACM.

Folino, F., Guarascio, M., & Pontieri, L. (2012). Discovering context-aware models for predicting business process performances. In *Proceedings of on the move to meaningful internet systems (OTM)* (pp. 287–304). Berlin: Springer.

Kang, B., Kim, D., & Kang, S.-H. (2012, April). Real-time business process monitoring method for prediction of abnormal termination using KNNI-based LOF prediction. *Expert Systems and Applications*, *39*(5), 6061–6068.

Lakshmanan, G. T., Rozsnyai, S., & Wang, F. (2013). Investigating clinical care pathways correlated with outcomes. In *Proceedings of the international conference on business process management* (pp. 323–338). Berlin: Springer.

Lo, D., Cheng, H., & Lucia. (2011). Mining closed discriminative dyadic sequential patterns. In *Proceedings of the international conference on extending database technology (EDBT)* (pp. 21–32). Springer.

Ly, L. T., Rinderle-Ma, S., Knuplesch, D., & Dadam, P. (2011). Monitoring business process compliance using compliance rule graphs. In *Proceedings of on the move to meaningful Internet systems (OTM)* (pp. 82–99). Berlin: Springer.

Maggi, F., Di Francescomarino, C., Dumas, M., & Ghidini, C. (2013). Predictive monitoring of business processes. In *Proceedings of the international conference on advanced information systems engineering (CAiSE)*. Springer.

Maggi, F., Montali, M., Westergaard, M., & van der Aalst, W. (2011). Monitoring business constraints with linear temporal logic: An approach based on colored automata. In *Proceedings of the international conference on business process management (BPM)* (pp. 132–147). Heidelberg: Springer.

Maggi, F. M., Montali, M., & van der Aalst, W. M. P. (2012). An operational decision support framework for monitoring business constraints. In *Proceedings of the international conference on fundamental approaches to software engineering (FASE)* (pp. 146–162). Berlin: Springer.

Metzger, A., Franklin, R., & Engel, Y. (2012). Predictive monitoring of heterogeneous service-oriented business networks: The transport and logistics case. In *Proceedings of the SRII global conference* (pp. 313–322).

Pika, A., Aalst, W., Fidge, C., Hofstede, A., & Wynn, M. (2013). Predicting deadline transgressions using event logs. In *Proceedings of the BPM'2012 workshops* (pp. 211–216). Berlin: Springer.

Rogge-Solti, A., & Weske, M. (2013). Prediction of remaining service execution time using stochastic petri nets with arbitrary firing delays. In *Proceedings of international conference on service-oriented computing (ICSOC)* (pp. 389–403). Berlin: Springer.

Schonenberg, H., Mans, R., Russell, N., Mulyar, N., & van der Aalst, W. M. P. (2008). Process flexibility: A survey of contemporary approaches. In *Proceedings of the CIAO! and EOMAS 2009 workshops* (pp. 16–30). Berlin: Springer.

Setiawan, M. A., & Sadiq, S. W. (2013). A methodology for improving business process performance through positive deviance. *International Journal of Information System Modeling and Design, 4*(2), 1–22.

Spreitzer, G. M., & Sonenshein, S. (2004). Toward the construct definition of positive deviance. *American Behavioral Scientist, 47*(6), 828–847.

Sun, C., Du, J., Chen, N., Khoo, S.-C., & Yang, Y. (2013). Mining explicit rules for software process evaluation. In *Proceedings of the international conference on software and system process (ICSSP)* (pp. 118–125). ACM.

Suriadi, S., Wynn, M. T., Ouyang, C., ter Hofstede, A. H. M., & van Dijk, N. J. (2013). Understanding process behaviours in a large insurance company in Australia: A case study. In *Proceedings of the international conference on advanced information systems engineering (CAiSE)* (pp. 449–464). Springer.

Swenson, K. D. (2010). *Mastering the unpredictable: How adaptive case management will revolutionize the way that knowledge workers get things done*. Tampa, FL: Meghan-Kiffer.

Swinnen, J., Depaire, B., Jans, M. J., & Vanhoof, K. (2012). A process deviation analysis – A case study. In *Proceedings of the BPM'2011 workshops* (pp. 87–98). Springer.

Tregear, R. (2013, January). Insignificant and exceptional. *BPTrends.*

van der Aalst, W. M. P., Schonenberg, M. H., & Song, M. (2011). Time prediction based on process mining. *Information Systems, 36*(2), 450–475.

Weidlich, M., Ziekow, H., Mendling, J., Günter, O., Weske, M., & Desai, N. (2011). Event-based monitoring of process execution violations. In *Proceedings of the international conference on advanced information systems engineering (CAiSE)* (pp. 182–198). Springer.

Xing, Z., Pei, J., Dong, G., & Yu, P. S. (2008). Mining sequence classifiers for early prediction. In *Proceedings of the SIAM international conference on data mining (SDM)* (pp. 644–655). SIAM.