

# Risk-Based Proactive Process Adaptation

Andreas Metzger<sup>(✉)</sup> and Philipp Bohn

paluno – The Ruhr Institute for Software Technology,  
University of Duisburg-Essen, Essen, Germany  
{andreas.metzger, philipp.bohn}@paluno.uni-due.de

**Abstract.** Proactive process adaptation facilitates preventing or mitigating upcoming problems during process execution, such as process delays. Key for proactive process adaptation is that adaptation decisions are based on accurate predictions of problems. Previous research focused on improving aggregate accuracy, such as precision or recall. However, aggregate accuracy provides little information about the error of an individual prediction. In contrast, so called reliability estimates provide such additional information. Previous work has shown that considering reliability estimates can improve decision making during proactive process adaptation and can lead to cost savings. So far, only constant cost functions have been considered. In practice, however, costs may differ depending on the magnitude of the problem; e.g., a longer process delay may result in higher penalties. To capture different cost functions, we exploit numeric predictions computed from ensembles of regression models. We combine reliability estimates and predicted costs to quantify the risk of a problem, i.e., its probability and its severity. Proactive adaptations are triggered if risks are above a pre-defined threshold. A comparative evaluation indicates that cost savings of up to 31%, with 14.8% savings on average, may be achieved by the risk-based approach.

**Keywords:** Predictive monitoring · Proactive adaptation · Risk · Business process · Machine learning

## 1 Introduction

Proactive process adaptation allows preventing the occurrence of problems or mitigating the impact of upcoming problems during process execution [29]. Thereby, proactive process adaptation addresses shortcomings of reactive adaptation, such as loss of money (e.g., due to contractual penalties) or time-consuming roll-back and compensation activities [1, 20].

Proactive adaptation relies on predictive process monitoring to forecast potential problems [1]. Predictive process monitoring predicts how an ongoing process instance will unfold up to its completion [19, 22]. If a potential problem is predicted, this problem is analyzed and adaptation decisions are triggered to prevent or mitigate the predicted problem. As an example, a delay in the expected delivery time for a freight transport process may incur contractual

penalties [6, 14]. If during the execution of such freight transport process a delay is predicted, alternative and thus faster transport services (such as air delivery instead of road delivery) can be scheduled before the delay actually occurs, thereby avoiding contractual penalties.

**Problem Statement.** A key requirement for proactive adaptation is that the adaptation decisions are based on accurate predictions. Informally, prediction accuracy characterizes the ability of a prediction technique to forecast as many true violations as possible, while – at the same time – generating as few false alarms as possible [29]. Prediction accuracy is important for two main reasons [23]. First, accurate predictions mean more true violations and thus triggering more required adaptations. Each missed required adaptation means one less opportunity for proactively preventing or mitigating a problem. Second, accurate predictions mean less false alarms, and thus triggering less unnecessary adaptations. Unnecessary adaptations incur additional costs for executing the adaptations, while not addressing actual problems.

Previous research on predictive process monitoring and proactive adaptation (see Sect. 5) focused on aggregate accuracy, such as precision or recall. However, aggregate accuracy does not provide direct information about the error of an *individual* prediction. In contrast, so called *reliability estimates* provide such additional information [2]. As an example, an aggregate accuracy of 75% means that, for all predictions, there is the same 75% chance that a prediction is correct. In contrast, the reliability estimate of one prediction may be 60% while for another prediction it may be 90%. Reliability estimates thus facilitate distinguishing between more and less reliable predictions on a case by case basis. In our previous work, we have introduced a predictive monitoring approach that considers such reliability estimates [21]. Experimental results indicate that considering reliability estimates during proactive process adaptation may lead to better decisions in 83% of the cases, entailing cost savings of 14% on average.

Yet, previous work is based on simplistic cost models that only consider constant cost functions. In practice, however, costs may differ depending on the magnitude of the problem (e.g., see [17, 24, 30]). As an example, a longer delay in the freight transport process may result in higher penalties.

**Contributions.** We introduce a *risk-based proactive process adaptation* approach that can capture different cost functions. We understand risk as the combination of the severity and the probability of a potential problem (e.g., see ISO 31000 [27]). Risk severity is computed by feeding the predicted magnitude of the problem into the respective cost function. The magnitude of the problem is predicted using ensembles of neural-network regression models. Risk probability is given by the reliability estimate of the prediction. Similar to our previous work, reliability estimates are computed from neural-network ensembles. During run time, an adaptation is triggered if a risk is detected that is greater than a pre-defined risk threshold.

We experimentally analyze the effect of considering risks during proactive adaptation in terms of cost savings. To this end, we perform a comparative

evaluation of the risk-based approach with our previous approach, which was based on binary predictions computed from ensembles of classification models.

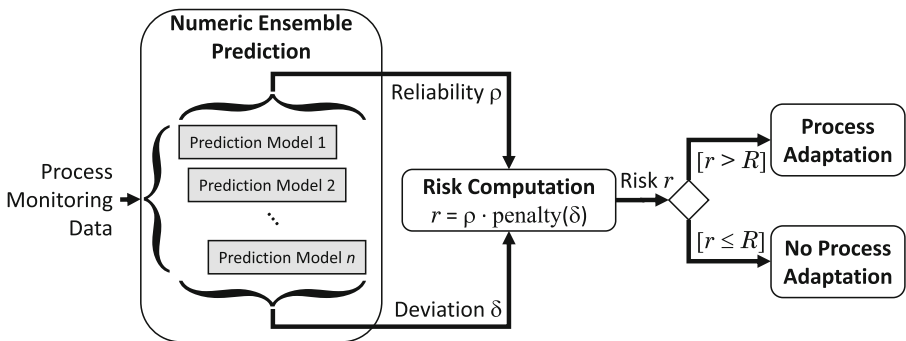
The remainder of the paper is structured as follows. Section 2 describes the risk-based adaptation approach. Section 3 explains the experimental design, while Sect. 4 presents the experimental results. Section 5 discusses related work, and Sect. 6 concludes with an outlook on future work.

## 2 Risk-Based Adaptation

This section provides a conceptual overview of our approach for risk-based proactive process adaptation. It explains how we build and combine the prediction models and implement the approach using machine learning technology.

### 2.1 Conceptual Overview

Figure 1 depicts how our approach computes risk  $r$  during process execution, and how this risk is considered for proactive process adaptation.



**Fig. 1.** Overview of risk-based process adaptation (Ensemble size  $n$ )

As mentioned above, we quantify risk as a combination of two factors: the probability of the occurrence of a risk event and the severity of that risk event (e.g., such as in ISO 31000 [27]). In our approach, a risk event is the violation of a service level objective; e.g., a delay in a transport process.

Our approach uses an ensemble of regression models to compute the two aforementioned risk factors. Ensemble prediction is a meta-prediction technique that combines the predictions of  $n$  prediction models trained to perform the same task [26]. The main aim of ensemble prediction is to increase predictive performance and, in particular, aggregate prediction accuracy. Additionally, ensemble prediction facilitates computing reliability estimates (e.g., see [2, 21]).

In our approach, each prediction model  $i \in \{1, \dots, n\}$  gives a prediction  $a_i$  pertaining to the service level objective of interest.

We use these  $n$  predictions to compute the two risk factors as follows.

**Risk probability.** The reliability estimate,  $\rho$ , of the prediction gives the risk probability. The intuition here is that a higher reliability of a predicted violation indicates a higher probability for that risk event to actually occur.

The reliability estimate  $\rho$  is computed by counting how many models in the ensemble agree on their prediction. Assuming an expected service level objective  $A$  for a given process instance,  $\rho$  is computed from the predictions  $a_i$  as follows:

$$\rho = \max_{i=1,\dots,n} \left( \frac{|i : a_i \models A|}{n}, \frac{|i : a_i \not\models A|}{n} \right),$$

with  $a_i \models A$  meaning that the *predicted* service level objective fulfills the *expected* service level objective.

**Risk severity.** For risk severity, first the average predicted deviation,  $\delta$ , from the expected service level objective  $A$  is computed as:

$$\delta = \frac{1}{n} \cdot \sum_{i=1,\dots,n} (a_i - A)$$

A  $\delta > 0$  indicates a violation. Without loss of generality, we assume that a smaller service level objective value is better.

Using a penalty function,  $\text{penalty}(\delta)$ , this gives the predicted penalty (see Sect. 3.2 for a definition of this function).

**Risk.** Together, these two risk factors give risk as  $r = \rho \cdot \text{penalty}(\delta)$ . During run time, numeric predictions are generated using process monitoring data collected for the specific process instance. A running process instance is proactively adapted if  $r > R$ , with  $R$  being a pre-defined risk threshold.

## 2.2 Implementation

We use artificial neural networks (ANNs [15]) as prediction models, which have shown good success in our earlier work [21, 22]. In particular, we use multilayer perceptrons as a specific form of ANNs. We use the implementation of multilayer perceptrons (with their standard parameters) of the WEKA open source machine learning toolkit<sup>1</sup>. As attributes for the prediction models, we take the expected and actual times for all services of the process until the point of prediction.

To automatically train the ensembles of ANNs and to compute reliability estimates as well as predicted deviations, we developed a Java tool that exploits the libraries of the WEKA machine learning toolkit.

We use *bagging* (bootstrap aggregating) as a concrete ensemble technique. Bagging uses a single type of prediction technique (ANNs in our case), but uses different training data sets to generate different prediction models. Bagging generates  $n$  new training data sets from the whole training set by sampling from the whole training data set uniformly and with replacement. For each of the  $n$  new training data sets an individual prediction model is trained. Bagging is generally recommended and used for ANNs [9].

<sup>1</sup> <http://www.cs.waikato.ac.nz/ml/weka/>.

We introduce a normalization factor for deviations in our implementation that ensures that  $\delta$  lies in the interval  $[0, 1]$ . This is not a limitation of the approach, but serves to define alternative cost functions in a comparable way. The normalization factor can be computed using the largest observed actual deviation from the training data set, i.e., the data set which is used to train the prediction models. Together with normalizing the cost functions to  $[0, 1]$  this gives normalized risks values  $r \in [0, 1]$ .

### 3 Comparative Evaluation

This section explains our experimental design and execution, in particular focusing on the cost model with the different considered cost functions.

We aim to experimentally analyze the effect that considering risk has on the overall costs of process execution. We perform a series of experiments, using a real-world process model and data set from the transport and logistics industry. We compare our risk-based approach with the baseline approach introduced previously [21]. This baseline approach uses binary predictions (i.e., “violation”/“non-violation” predictions) computed from ensembles of classification models. The baseline approach only considers constant cost functions.

#### 3.1 Cost Model

We consider two cost factors of proactive process adaptation [21, 23]. On the one hand, one may face penalties in case an adaptation is missed or not effective, as problems remain. On the other hand, an adaptation of the running processes may require effort and resources, and thus incur costs. Figure 2 shows a cost model (in the form of a decision tree) that incorporates these cost factors.




Fig. 2. Costs of proactive process adaptation

In this model, the actual costs of executing a single process instance depend on three main factors: (1) the *predicted risk* and whether it triggers an adaptation; (2) the fulfillment of the service level objective after a triggered *adaptation*; (3) the fulfillment of the service level objective if *no adaptation* is triggered.

### 3.2 Costs Functions

Penalties and adaptation costs may differ depending on the magnitude of deviation ( $\delta$ ) from expected service level objectives and the extent of adaptation required. As an example, penalties faced in the transport process may be higher if the actual delays in a transport process are longer. Also, using an air transport service for an alternative transport leg may be more expensive compared with using a road transport service.

In particular, this means the cost functions for penalties and adaptation costs may take different shapes. Different types of cost functions have been identified in the literature (e.g., see [18,24,30]). These cost functions share two main characteristics [18]: (1) Cost functions are monotonically increasing; e.g., the penalty for a longer delay is never smaller than the penalty for a shorter delay; (2) Cost functions have a point discontinuity. Before that and including that point, the costs are generally 0, beyond this point costs incur. For our experiments we consider the point discontinuity to be at  $\delta = 0$ .

To keep the complexity of our experiments manageable, we have chosen three cost functions that represent typical shapes of costs as described in the literature and which may be faced in the transport and logistics domain (the domain of our data set; see Sect. 3.4). The variants of these shapes, as we use them, are shown in Fig. 3. For the step-wise costs, we consider 5 steps, i.e.,  $s = 5$ , for our experiments (the higher  $s$  the closer the function will be to the linear function).

To ensure a fair comparison among the resulting costs when using the different cost functions, we choose the parameters for the cost models such that their average costs (across  $\delta \in [0, 1]$ ) are the same. This means,  $c_{\text{const}} = 1/2 \cdot c_{\text{lin}}$  and  $c_{\text{step}} = s/(s + 1) \cdot c_{\text{lin}}$ .




Fig. 3. Different shapes of cost functions

### 3.3 Experimental Variables

In our experiments, we consider *cost* as a dependent variable. For each process instance, we compute its individual costs according to the cost model defined in Sect. 3.1. The two cost drivers considered are penalties in case of violations and the costs for proactive process adaptation. The total costs are the sum of the individual costs of all process instances.

We consider the following independent variables.

- *Penalty cost function*: We use each of the three cost functions introduced in Sect. 3.2 for determining penalties. Penalty functions serve two purposes: (1) we use the penalty function to compute the *predicted* penalty and thus severity of the risk (see Sect. 2.1); (2) we use the penalty function to compute the *actual* penalty according to the cost model in Sect. 3.1.
- *Adaptation cost function*: We consider different shapes of adaptation costs by using each of the three cost functions from Sect. 3.2. Together with the three cost functions used for penalties, this leads to nine combinations of cost functions considered during our experiments.
- *Adaptation effectiveness*  $\alpha \in (0, 1]$ : If an adaptation helps achieve the expected service level objectives, we consider such adaptation effective. We use  $\alpha$  to represent the fact that not all adaptations might be effective. More concretely,  $\alpha$  represents the probability that an adaptation is effective; e.g.,  $\alpha = 1$  means that all adaptations are effective.
- *Risk threshold*  $R \in [0, 1]$ : An adaptation is triggered if risk  $r > R$ . We vary  $R$  to reflect different attitudes towards process risks.

Note that for a concrete problem situation in practice, the concrete values for all of the aforementioned independent variables – with the exception of the risk threshold – are given. The penalty cost function is defined by the respective service level agreement (SLA). The adaptation cost function and the adaptation effectiveness are characteristics of the process execution environment.

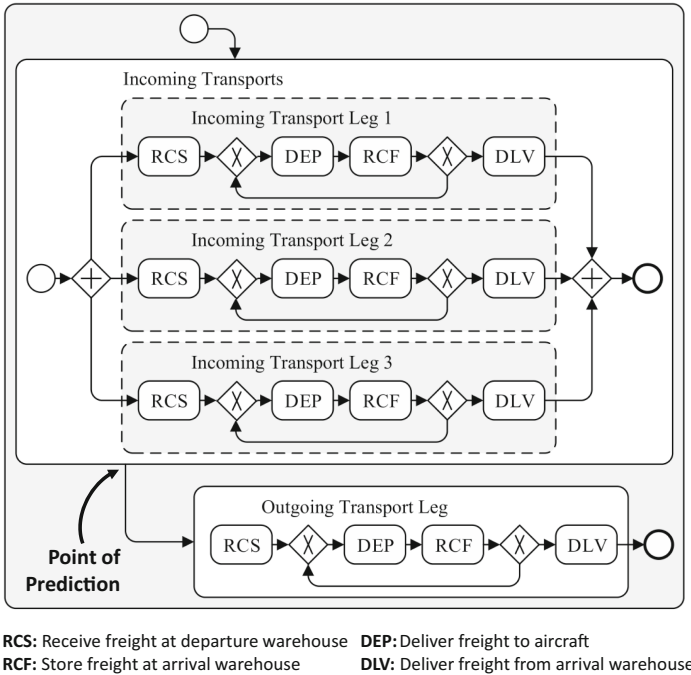
### 3.4 Industry Data Set and Experiment Execution

The data set we use in our experiments stems from operational data of an international freight forwarding company. The data set covers five months of business operations and includes event logs of 3,942 business process instances, comprising a total of 56,082 service executions<sup>2</sup>.

The processes and event data comply with IATA's Cargo 2000 standard<sup>3</sup>. Figure 4 shows the BPMN model of the business processes covered by the data set. Up to three smaller shipments from suppliers are consolidated and in turn shipped together to customers to benefit from better freight rates or increased cargo security. The process involves the execution of transport and logistics services, which are labeled using the acronyms of the Cargo 2000 standard.

<sup>2</sup> The industry data set is available from <http://www.s-cube-network.eu/c2k>. The predictions used in our experiments are available from <https://uni-duisburg-essen.sciebo.de/index.php/s/oYnNH2PAudkWDfg>.

<sup>3</sup> Cargo 2000 (now Cargo iQ: <http://cargoiq.org/>) is an initiative of IATA.



**Fig. 4.** Cargo 2000 transport process and services

In our experiments, we predict, during process execution, whether a transport process instance violates its stipulated delivery deadline. Predictions may be performed at any point in time during process execution, but the point of prediction has an impact on prediction accuracy; e.g., earlier points usually imply lower prediction accuracy [22]. For our experiments, we perform the predictions immediately after the synchronization point of the incoming transport processes as indicated in Fig. 4. Our earlier work indicated reasonably good prediction accuracy (>70%) for this point in process execution, while still leaving time to execute actions required to respond to violations or mitigate their effects [22].

## 4 Experimental Results


Here, we present and discuss the results of our experimental evaluation.

### 4.1 Results

Figure 5 gives a first impression of the effect of risk-based proactive adaptation on costs. The figure shows the relative cost savings of our risk-based approach compared with the baseline approach of our previous work for all nine combinations of cost functions. We have chosen  $\alpha = 0.9$ , which is a relatively high



probability of effective process adaptations. Our previous approach has already shown high cost savings for such high  $\alpha$ , and thus poses a more challenging baseline for further savings.



**Fig. 5.** Relative cost savings [%] when considering risk-based adaptation for  $\alpha = 0.9$

As can be seen from Fig. 5, the risk-based approach performs worse than the baseline if we face constant penalties. This is not surprising, as in such case the severity of the risk is constant and thus does not have an effect of risk-based decision making. However, the risk-based approach shows clear cost savings if penalties are non-constant. For the chosen  $\alpha = 0.9$ , cost savings can be as high as 26%. Cost savings are achieved for all combinations of the non-constant cost functions for risk thresholds that are greater than 0.3.

Table 1 shows the maximum cost savings for different values for  $\alpha$ . These results show an interesting trend. The smaller the chance of effective process

adaptation, the higher the cost savings of the risk-based approach when compared to the baseline. We attribute this to the fact that the risk-based approach is more conservative and precise when deciding on whether to proactively adapt, and thus would rather avoid an adaptation. This in particular leads to benefits in situations where adaptations might not be effective.

**Table 1.** Maximum relative cost savings [%] for given  $\alpha$

Cost model		Max. avg. savings [%] for given $\alpha$			
Penalty	Adapt. cost	$\alpha = 0.3$	$\alpha = 0.5$	$\alpha = 0.7$	$\alpha = 0.9$
const	const	2.6	0.0	-1.0	-1.0
const	step-wise	5.7	3.4	0.9	-1.0
const	lin	7.2	4.9	2.5	0.1
step-wise	const	17.0	14.0	12.0	8.6
step-wise	step-wise	22.0	20.0	18.0	16.0
step-wise	lin	24.0	22.0	20.0	18.0
lin	const	22.0	20.0	18.0	16.0
lin	step-wise	27.0	26.0	25.0	23.0
lin	lin	30.0	28.0	27.0	26.0

Table 2 shows how different risk threshold values impact on cost savings. The table shows the cost savings for given values of  $R$ , averaged over  $\alpha = \{0.1, 0.2, 0.3, \dots, 1\}$ . As can be seen from the results, higher thresholds ensure that cost savings (highlighted in gray) will be achieved in more situations. Yet, these cost savings may be smaller than the cost savings that may be achieved for thresholds in the medium range.

**Table 2.** Relative cost savings [%] averaged over  $\alpha = \{0.1, 0.2, 0.3, \dots, 1\}$

Cost model		Avg. savings [%] for given $R$				
Penalty	Adapt. cost	$R = 0.1$	$R = 0.3$	$R = 0.5$	$R = 0.7$	$R = 0.9$
const	const	-27,0	-27,0	-15,0	-2,2	-0,6
const	step-wise	-20,0	-21,0	-16,0	-2,7	1,9
const	lin	-19,0	-20,0	-17,0	-3,0	3,1
step-wise	const	-3,4	11,0	13,0	10,0	1,7
step-wise	step-wise	-12,0	12,0	18,0	17,0	6,4
step-wise	lin	-14,0	12,0	20,0	20,0	8,6
lin	const	8,6	18,0	19,0	15,0	2,8
lin	step-wise	2,6	21,0	25,0	23,0	8,4
lin	lin	0,6	21,0	27,0	26,0	11,0

Overall, in our experiments the risk-based approach led to cost savings of 14.8% on average. Considering non-constant cost models only, the risk-based approach led to cost savings of 23.4% on average. The maximum savings we measured in our experiments were 31%.

## 4.2 Discussion

Below we discuss the experimental results with respect to potential threats to validity and limitations in practice.

**Internal Validity.** To minimize the risk of bias in our results, we performed a 10-fold cross-validation for training and testing the prediction models.

The success of ensemble prediction depends on the accuracy of the individual models, but also on the so called *diversity* among these models [4]. To ensure diversity of the ensemble, we used bagging to generate the individual models (see Sect. 2.2). As bootstrap size (which is the size of the newly generated training data set), we used 80%. Our previous experiments indicated that different bootstrap sizes did not impact the general shape of the experimental results [21].

We used an ensemble of size 100 in our experiments. The size of the ensemble did not lead to different principal findings in our experiments. Yet, by using such a large ensemble, we gain more fine-grained reliability estimates than by using a smaller ensemble. In our case, the ensemble of size 100 delivers reliability estimates with a granularity of 0.01. Training such a large ensemble, however, takes more time than training a smaller ensemble. In our experimental setting, training the ensemble took around one day on a standard desktop PC.

**External Validity.** Our experimental results are based on a relatively large, industry data set. We have specifically chosen different risk thresholds ( $R$ ), different probabilities of effective process adaptations ( $\alpha$ ), as well as different shapes of penalties and adaptation costs to cover different possible situations that may be faced in practice. The process model covers many relevant workflow patterns [31]: sequence; exclusive choice and simple merge; cycles; parallel split and synchronization. Still, our data set is from a single application domain which thus may limit generalizability.

**Construct Validity.** We took great care to ensure we measure the right things. In particular, we used normalized costs as a common reference to perform the comparative evaluation between the risk-based approach and the baseline approach. Yet, so far, we have not considered aggregate or frame SLAs. In these kinds of SLAs, the presence of multiple service level objective violations incurs penalties; e.g., if more than 5% of the process instances are delayed (e.g., see [6, 14]). To address these kinds of SLAs, we plan to explore approaches for predicting aggregate process outcomes (e.g., see [25]).

## 5 Related Work

We discuss related work from three angles: reliability, cost and risk.

**Reliability-based Prediction and Adaptation.** Research on predictive process monitoring (such as [3, 5, 7, 12, 13, 22, 33]) and proactive adaptation (such as [1, 20, 23]) focused on aggregate prediction accuracy. Only recently, reliability estimates have been considered in the context of predictive process monitoring.

Maggi et al. [19] use decision tree learning for prediction and for computing reliability estimates. They observe that filtering predictions using reliability esti-

mates may improve aggregate accuracy. However, they do not factor in reliability for decision making during process adaptation.

Francescomarino et al. [11] use decision trees and random forests. Only if the reliability of a prediction is above a certain threshold, the prediction is considered. In their experiments, they measure “failure rate” to assess the performance of their predictions. “Failure rate” is defined as the percentage of process instances for which no reliable prediction could be given. Yet, they do not further analyze the effectiveness of their approach in case a process adaptation is made.

In our previous work [21], we employed ensembles of classification models to compute reliability estimates. We have analyzed the effect of using these reliability estimates for process adaptation and have measured an increase of non-violation rates, i.e., the rates of successful process executions.

**Cost-based Adaptation.** Different ways of factoring in costs during predictive monitoring and proactive adaptation have been presented in the literature. On the one hand, costs may be considered by the prediction technique itself. A prominent class of approaches is cost-sensitive learning [10]. Cost-sensitive learning incorporates asymmetric costs into the learning of prediction models to minimize costs due to prediction errors [34,35]. However, existing cost-sensitive learning techniques do not consider reliability estimates.

On the other hand, costs may be considered when deciding on proactive process adaptations. Cost-based adaptation attempts to minimize the overall costs of process execution and adaptation. Leitner et al. [17] consider the costs of adaptation when deciding on the adaptation of service-oriented workflows. They formalize an optimization problem taking into account costs of violations and costs of applying adaptations. Their experimental results indicate that cost reductions of up to 56% may be achieved. Aschoff and Zisman [1] consider response time and cost values during the proactive adaptation of service compositions. Their experimental results indicate that the cheapest executions were selected in 85% of the cases. However, both aforementioned cost-aware proactive adaptation approaches do not consider prediction reliability.

In our previous work [21], we analyzed the effect of using reliability estimates on costs. In 82.9% of the situations, considering reliability estimates had a positive effect on costs, leading to cost savings of up to 54%, with 14% savings on average. However, we only used a simple cost model with constant costs and did not consider different shapes of costs.

**Risk-aware Process Management.** Risk-aware process management aims to (1) minimize risks in business processes by design, and (2) to monitor the emergence of risks and apply risk mitigation actions at run time [32]. While research focused mainly on (1), a few approaches have been presented for (2).

Conforti et al. [8] augment process models with so called risk sensors, which collect information from running process instances and exploit historical process data. Each sensor is associated to a risk condition that combines the probability of a problem with a risk threshold. If the probability is greater than the threshold, process managers are notified. Pika et al. [25] follow an approach similar to risk sensors. They propose defining so called process risk indicators, which are

patterns observable in event logs that may indicate risks. While the aforementioned authors focused on risk detection and prediction, Kim et al. [16] propose an integrated risk management approach that facilitates proactively mitigating risks at run time. Risk mitigation strategies are expressed as event-condition-action rules. All three aforementioned approaches, however, only quantify the probability of a risk event, but do not quantify the severity of the risk event. Also, the approaches have not been evaluated with respect to how such risk information may improve overall process adaptation; e.g., in terms of costs.

Rogge-Solti and Weske [28] focus on the risk of missing a process deadline. They consider costs incurred by a deadline violation in addition to the probability of that deadline violation. They use stochastic Petri nets to predict the risk of missing a given process deadline. For each of these deadlines specific costs may be assigned. However, the specific costs for each deadline are always constant and thus independent of the magnitude of the deadline violation. In contrast, our approach takes into account cost functions that depend on the magnitude of deviation from expected service level objectives.

## 6 Conclusion

We have introduced a risk-based approach for proactive process adaptation, which exploits ensembles of regression models to compute the probability and the severity of service level objective violations. Our comparative evaluation provided empirical evidence that risk-based proactive process adaptation may lead to additional cost savings when compared with proactive adaptation based on probability only. Additional cost savings were 14.8% on average (23.4% if considering non-constant cost models), with maximum savings of 31%.

Building on these promising results, we plan to gather further empirical evidence by replicating our experiments using other process models and data sets, such as from port logistics and e-commerce. Further, to handle aggregate and frame SLAs, we will extend our approach to consider penalties caused by multiple service level objective violations.

**Acknowledgments.** We cordially thank Christina Bellinghoven, Felix Föcker, and Adrian Neubauer for helpful pointers during earlier drafts of the paper. Research leading to these results received funding from the EU's Horizon 2020 research and innovation programme under grant agreement no. 731932 (TransformingTransport) and from the EFRE co-financed operational program NRW.Ziel2 under grant agreement 005-1010-0012 (LoFIP – Cockpits for Operational Management of Transport Processes).

## References

1. Aschoff, R., Zisman, A.: QoS-driven proactive adaptation of service composition. In: Kappel, G., Maamar, Z., Motahari-Nezhad, H.R. (eds.) ICSOC 2011. LNCS, vol. 7084, pp. 421–435. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25535-9\\_28](https://doi.org/10.1007/978-3-642-25535-9_28)
2. Bosnic, Z., Kononenko, I.: Automatic selection of reliability estimates for individual regression predictions. *Knowl. Eng. Rev.* **25**(1), 27–47 (2010)

3. Breuker, D., Delfmann, P., Matzner, M., Becker, J.: Designing and evaluating an interpretable predictive modeling technique for business processes. In: Fournier, F., Mendling, J. (eds.) BPM 2014. LNBP, vol. 202, pp. 541–553. Springer, Cham (2015). doi:[10.1007/978-3-319-15895-2\\_46](https://doi.org/10.1007/978-3-319-15895-2_46)
4. Brown, G., Wyatt, J.L., Tiño, P.: Managing diversity in regression ensembles. *J. Mach. Learn. Res.* **6**, 1621–1650 (2005)
5. Cabanillas, C., Di Ciccio, C., Mendling, J., Baumgrass, A.: Predictive task monitoring for business processes. In: Sadiq, S., Soffer, P., Völzer, H. (eds.) BPM 2014. LNCS, vol. 8659, pp. 424–432. Springer, Cham (2014). doi:[10.1007/978-3-319-10172-9\\_31](https://doi.org/10.1007/978-3-319-10172-9_31)
6. Marquezan, C.C., Metzger, A., Franklin, R., Pohl, K.: Runtime management of multi-level SLAs for transport and logistics services. In: Franch, X., Ghose, A.K., Lewis, G.A., Bhiri, S. (eds.) ICSSOC 2014. LNCS, vol. 8831, pp. 560–574. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-45391-9\\_49](https://doi.org/10.1007/978-3-662-45391-9_49). (Industry paper)
7. Castellanos, M., Salazar, N., Casati, F., Dayal, U., Shan, M.C.: Predictive business operations management. *Int. J. Comput. Sci. Eng.* **2**(5/6), 292–301 (2006)
8. Conforti, R., Rosa, M.L., Fortino, G., ter Hofstede, A.H.M., Recker, J., Adams, M.: Real-time risk monitoring in business processes: a sensor-based approach. *J. Syst. Softw.* **86**(11), 2939–2965 (2013)
9. Dietterich, T.G.: Ensemble methods in machine learning. In: Kittler, J., Roli, F. (eds.) MCS 2000. LNCS, vol. 1857, pp. 1–15. Springer, Heidelberg (2000). doi:[10.1007/3-540-45014-9\\_1](https://doi.org/10.1007/3-540-45014-9_1)
10. Elkan, C.: The foundations of cost-sensitive learning. In: Nebel, B. (ed.) 7th Intl Joint Conference on Artificial Intelligence (IJCAI 2001), Seattle, Washington, pp. 973–978. Morgan Kaufmann (2001)
11. Di Francescomarino, C., Dumas, M., Federici, M., Ghidini, C., Maggi, F.M., Rizzi, W.: Predictive business process monitoring framework with hyperparameter optimization. In: Nurcan, S., Soffer, P., Bajec, M., Eder, J. (eds.) CAISE 2016. LNCS, vol. 9694, pp. 361–376. Springer, Cham (2016). doi:[10.1007/978-3-319-39696-5\\_22](https://doi.org/10.1007/978-3-319-39696-5_22)
12. Ghosh, R., Ghose, A., Hegde, A., Mukherjee, T., Mos, A.: QoS-driven management of business process variants in cloud based execution environments. In: Sheng, Q.Z., Stroulia, E., Tata, S., Bhiri, S. (eds.) ICSSOC 2016. LNCS, vol. 9936, pp. 55–69. Springer, Cham (2016). doi:[10.1007/978-3-319-46295-0\\_4](https://doi.org/10.1007/978-3-319-46295-0_4)
13. Grigori, D., Casati, F., Castellanos, M., Dayal, U., Sayal, M., Shan, M.C.: Business process intelligence. *Comput. Ind.* **53**(3), 321–343 (2004)
14. Gutiérrez, A.M., Cassales Marquezan, C., Resinas, M., Metzger, A., Ruiz-Cortés, A., Pohl, K.: Extending WS-Agreement to Support Automated Conformity Check on Transport and Logistics Service Agreements. In: Basu, S., Pautasso, C., Zhang, L., Fu, X. (eds.) ICSSOC 2013. LNCS, vol. 8274, pp. 567–574. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-45005-1\\_47](https://doi.org/10.1007/978-3-642-45005-1_47)
15. Haykin, S.: *Neural Networks and Learning Machines: A Comprehensive Foundation*, 3rd edn. Prentice Hall, Englewood Cliffs (2008)
16. Kim, J., Lee, J., Lee, J., Choi, I.: An integrated process-related risk management approach to proactive threat and opportunity handling: a framework and rule language. *Knowl. Process Manag.* **24**(1), 23–37 (2017)
17. Leitner, P., Hummer, W., Dustdar, S.: Cost-based optimization of service compositions. *IEEE Trans. Serv. Comput.* **6**(2), 239–251 (2013)
18. Leitner, P., Michlmayr, A., Rosenberg, F., Dustdar, S.: Monitoring, prediction and prevention of SLA violations in composite services. In: International Conference on Web Services (ICWS 2010), Miami, Florida, pp. 369–376. IEEE Computer Society (2010)

19. Maggi, F.M., Di Francescomarino, C., Dumas, M., Ghidini, C.: Predictive monitoring of business processes. In: Jarke, M., Mylopoulos, J., Quix, C., Rolland, C., Manolopoulos, Y., Mouratidis, H., Horkoff, J. (eds.) CAiSE 2014. LNCS, vol. 8484, pp. 457–472. Springer, Cham (2014). doi:[10.1007/978-3-319-07881-6\\_31](https://doi.org/10.1007/978-3-319-07881-6_31)
20. Metzger, A., Chi, C.H., Engel, Y., Marconi, A.: Research challenges on online service quality prediction for proactive adaptation. In: ICSE 2012 Workshop on European Software Services and Systems Research (S-Cube), Zurich, Switzerland. IEEE (2012)
21. Metzger, A., Föcker, F.: Predictive business process monitoring considering reliability estimates. In: Dubois, E., Pohl, K. (eds.) CAiSE 2017. LNCS, vol. 10253, pp. 445–460. Springer, Cham (2017). doi:[10.1007/978-3-319-59536-8\\_28](https://doi.org/10.1007/978-3-319-59536-8_28)
22. Metzger, A., Leitner, P., Ivanović, D., Schmieders, E., Franklin, R., Carro, M., Dustdar, S., Pohl, K.: Comparing and combining predictive business process monitoring techniques. *IEEE Trans. Syst. Man Cybern. Syst.* **45**(2), 276–290 (2015)
23. Metzger, A., Sammodi, O., Pohl, K.: Accurate proactive adaptation of service-oriented systems. In: Cámara, J., de Lemos, R., Ghezzi, C., Lopes, A. (eds.) *Assurances for Self-Adaptive Systems*. LNCS, vol. 7740, pp. 240–265. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-36249-1\\_9](https://doi.org/10.1007/978-3-642-36249-1_9)
24. Pernici, B., Siadat, S.H., Benbernou, S., Ouziri, M.: A penalty-based approach for QoS dissatisfaction using fuzzy rules. In: Kappel, G., Maamar, Z., Motahari-Nezhad, H.R. (eds.) *ICSOC 2011*. LNCS, vol. 7084, pp. 574–581. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25535-9\\_43](https://doi.org/10.1007/978-3-642-25535-9_43)
25. Pika, A., van der Aalst, W.M.P., Wynn, M.T., Fidge, C.J., ter Hofstede, A.H.M.: Evaluating and predicting overall process risk using event logs. *Inf. Sci.* **352–353**, 98–120 (2016)
26. Polikar, R.: Ensemble based systems in decision making. *IEEE Circ. Syst. Mag.* **6**(3), 21–45 (2006)
27. Purdy, G.: ISO 31000: 2009 - setting a new standard for risk management. *Risk Anal.* **30**(6), 881–886 (2010)
28. Rogge-Solti, A., Weske, M.: Prediction of business process durations using non-markovian stochastic petri nets. *Inf. Syst.* **54**, 1–14 (2015)
29. Salfner, F., Lenk, M., Malek, M.: A survey of online failure prediction methods. *ACM Comput. Surv.* **42**(3), 10:1–10:42 (2010)
30. Schuller, D., Siebenhaar, M., Hans, R., Wenge, O., Steinmetz, R., Schulte, S.: Towards heuristic optimization of complex service-based workflows for stochastic QoS attributes. In: *International Conference on Web Services (ICWS 2014)*, Anchorage, Alaska, pp. 361–368. IEEE Computer Society (2014)
31. Skouradaki, M., Ferme, V., Pautasso, C., Leymann, F., van Hoorn, A.: Micro-benchmarking BPMN 2.0 workflow management systems with workflow patterns. In: Nurcan, S., Soffer, P., Bajec, M., Eder, J. (eds.) CAiSE 2016. LNCS, vol. 9694, pp. 67–82. Springer, Cham (2016). doi:[10.1007/978-3-319-39696-5\\_5](https://doi.org/10.1007/978-3-319-39696-5_5)
32. Suriadi, S., et al.: Current research in risk-aware business process management - overview, comparison, and gap analysis. *Commun. Assoc. Inf. Syst. (CAIS)* **34**, 52 (2014)
33. Verenich, I., Dumas, M., La Rosa, M., Maggi, F.M., Di Francescomarino, C.: Complex symbolic sequence clustering and multiple classifiers for predictive process monitoring. In: Reichert, M., Reijers, H.A. (eds.) *BPM 2015*. LNBIP, vol. 256, pp. 218–229. Springer, Cham (2016). doi:[10.1007/978-3-319-42887-1\\_18](https://doi.org/10.1007/978-3-319-42887-1_18)
34. Zadrozny, B., Elkan, C.: Learning and making decisions when costs and probabilities are both unknown. In: Lee, D., Schkolnick, M., Provost, F.J., Srikant,

- R. (eds.) 7th International Conference on Knowledge Discovery and Data Mining (KDD 2001), San Francisco, California, pp. 204–213. ACM (2001)
35. Zhao, H., Sinha, A.P., Bansal, G.: An extended tuning method for cost-sensitive regression and forecasting. *Decis. Support Syst.* **51**(3), 372–383 (2011)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

