Data-driven Deep Learning for Proactive Terminal Process Management

Andreas Metzger[1] [0000−0002−4808−8297], Johannes Franke[2], and Thomas Jansen[2]

[1] paluno – The Ruhr Institute for Software Technology
University of Duisburg-Essen, Essen, Germany
andreas.metzger@paluno.uni-due.de
[2] duisport – Duisburger Hafen AG, Duisburg, Germany
johannes.franke@duisport.de, thomas.jansen@duisport.de

Abstract. Big data offers tremendous opportunities for transport process innovation. One key enabling big data technology is predictive data analytics. Predictive data analytics supports business process management by facilitating the proactive adaptation of process instances to mitigate or prevent problems. We present an industry case employing big data for process management innovation at duisport, the world's largest inland container port. In particular, we show how data-driven deep learning facilitates proactive port terminal process management. We demonstrate the feasibility of our deep learning approach by implementing it as part of a terminal productivity cockpit prototype. The terminal productivity cockpit provides decision support to terminal operators for proactive process adaptation. We confirm the desirability of our approach via interviews. We assess the viability of our approach by estimating the improvements in a key business KPI, as well as experimentally measuring the cost savings when compared to terminal operations without using proactive adaptation. We also present our main technical lessons learned regarding the use of big data for predictive analytics.

Keywords: Business process monitoring, proactive adaptation, prediction, accuracy, earliness, reliability, decision support, terminal operations

1 Introduction

Big data offers tremendous opportunities for transport process innovation and will have a profound economic and societal impact on mobility and logistics. As an example, with annual growth rates of 3.2% of passenger transport and 4.5% of freight transport in the EU [6], transforming the current mobility and logistics processes to become significantly more efficient, will have major impact. Improvements in operational efficiency empowered by big data are expected to save as much as EUR 440 billion globally in terms of fuel and time within the mobility and logistics sector, as well as reducing 380 megatons of CO2 emissions [27]. The mobility and logistics sector is ideally placed to benefit from big data technologies, as it already manages massive flows of goods and people whilst generating vast amounts of data [4].

One key enabling big data technology in transport is predictive data analytics [23]. Predictive analytics is a significant next step from descriptive analytics [13]. Where descriptive analytics aims to answer the question "what happened and why?", predictive analytics aims to answer the question “what will happen and when?”. Predictive analytics is considered a key technology and technical priority within the European big data ecosystem; e.g., see the Strategic Research and Innovation Agenda of the European Big Data Value Association [30].

Predictive analytics – in the form of predictive process monitoring [10,15,24] – supports business process management by facilitating proactive process adaptation. Proactive process adaptation can help prevent the occurrence of problems and it can mitigate the impact of upcoming problems during process execution by dynamically re-planning a running process instance [29,18,26,36,16,19]. As an example, a delay in the expected delivery time for a freight transport process may incur contractual penalties [12]. If during the execution of such freight transport process a delay is predicted, faster transport services (such as air delivery instead of road delivery) can be proactively scheduled to prevent the delay. Proactive process adaptation thereby helps transport operators to be proactive and avoid contractual penalties or time-consuming roll-back and compensation activities.

We present an industry case employing big data for process management innovation at duisport, the world’s largest inland container port. In particular we focus on how data-driven deep learning facilitates proactive port terminal process management. Deep learning employs artificial neural networks with many neurons and layers [11]. Applying such deep neural networks became feasible with recent breakthroughs in learning algorithms and the advent of powerful hardware.

Section 2 describes the situation faced in terminal process management. Section 3 elaborates on the actions taken in order to exploit data-driven deep learning for terminal process management. Section 4 presents results with respect to the impact on terminal operations. Section 5 provides our lessons learned.

2 Situation faced

2.1 Context and challenges for terminal process management

The case we present is located at duisport, an inland container port that handles 4.1 million containers per year. Duisport is situated in the middle of a large city (with close to 1/2 million inhabitants) and at the center of Germany’s largest metropolitan region, the Rhine-Ruhr metropolitan region (with close to 10 million inhabitants). This means that a multitude of roads, tracks and water ways serve as entry and exit points for containers to and from the terminals and ports. In addition, the transport infrastructure (roads and tracks) need to be shared within the metropolitan region.

Given the location of duisport within a dense metropolitan region, the increase in container volumes (due to the growth of freight transport) cannot be captured by a growth in space. It requires an improvement of terminal productivity.

The duisport case we report focuses on improving the productivity of a specific terminal: logport III. The logport III terminal covers an area of 15 hectares, offers nine rail connections, runs seven transhipment tracks and operates two gantry cranes. The terminal is interconnected with other duisport port areas and to more than 80 destinations in Europe and Asia. This includes daily rail and barge shuttles to the seaports of Antwerp and Rotterdam, as well as more than 30 trains per week between duisport and China.

We developed the duisport case in the context of the EU-funded lighthouse project TransformingTransport [3]. TransformingTransport is part of the European Big Data Value Public-Private Partnership[3]. The project started in January 2017 and brings together knowledge, solutions and impact potential of major European ICT and big data technology providers with the competence and experience of key European industry players and public bodies in the mobility and logistics domain. TransformingTransport developed 13 pilot cases that demonstrate how various transport sectors will benefit from big data solutions and the increased availability of data.

2.2 Big data availability as opportunity for process management innovation

The main driver to explore big data technologies for process management innovation at duisport was the increasing availability of data due to the instrumentation and digitization of terminal equipment. To illustrate, the two gantry cranes of the terminal that move the containers between trains and towing vehicles (i.e., trucks) produce data about 100 variables in 5 second intervals. These variables include information such as the crane's current state, its position, its current speed, its energy consumption, whether it transports a container or not, as well as observed faults.

At the time of writing, eight different data sets and over 30 million data entries from nine devices were available at the duisport terminal logport III. Figure 1 illustrates some of the available data. On the left hand side, the figure shows the terminal and its equipment. On the right hand side, the figure shows a visualization of the integrated and aggregated data in the form of a heat map, which shows the density of containers over the last 96 hours (ranging from low density = "green" to high density = "red").

2.3 Requirements towards predictive process management solutions

With respect to the usefulness of predictions as input for proactive process adaptation, we had to address two important requirements.

Requirement 1 – "Prediction accuracy". Informally, prediction accuracy characterizes the ability of a prediction technique to forecast as many true violations as possible, while generating as few false alarms as possible [32]. Prediction accuracy is important due to several reasons. Accurate predictions deliver more true violations and thus trigger more required adaptations. Each missed required adaptation means one less opportunity for preventing or mitigating a problem. Also, accurate predictions mean less false alarms, which in turn means triggering less unnecessary adaptations [22]. Unnecessary adaptations incur additional costs for executing the adaptations, while not addressing actual problems. A too high rate of false alarms will mean that a terminal operator will not trust the predictions and thus will not use them for decision making.

Requirement 2 – "Prediction earliness". Predictions should be produced early during process execution, as this leaves more time for adaptations. An adaptation typically has a non-negligible latency, i.e., it may take some time until an adaptation becomes effective [25,14]. As an example, dispatching additional personnel to mitigate delays in container transports may take several hours. Also, the later a process is adapted, the

[3] `http://www.big-data-value.eu/`

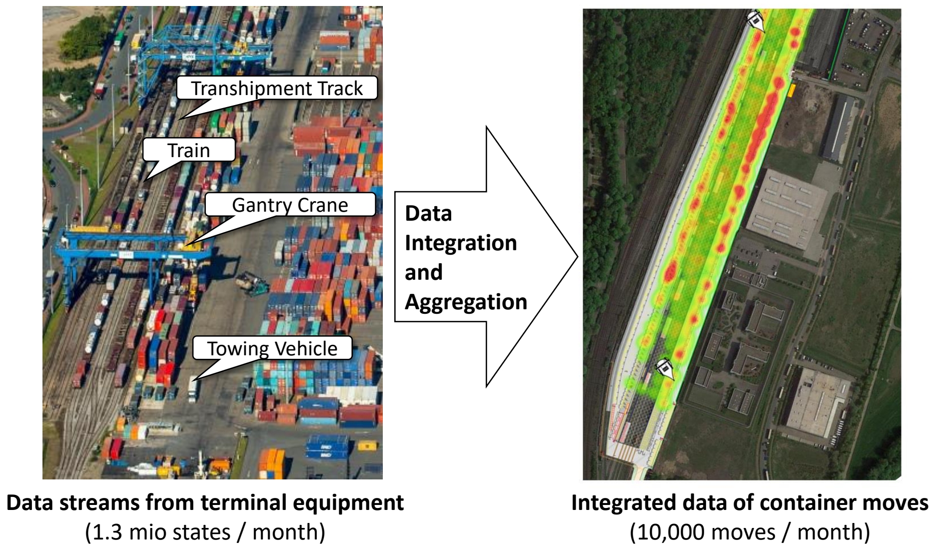

Fig. 1. Illustration of terminal data availability

fewer options may be available for adaptation. As an example, while at the beginning of a transport process one may be able to transport a container by train instead of ship, once the container is on-board the ship, such an adaption is no longer be feasible. Finally, if an adaptation is performed late in the process and turns out to be non-effective, not much time remains for remedial actions or further adaptations.

There is an important tradeoff between these two requirements. Later predictions typically have a higher accuracy (as depicted in Figure 2), because more information about the ongoing process instance is available. This means later predictions have a higher chance to be correct predictions. Therefore, one should favor later predictions as basis for proactive process adaptation. However, later predictions leave less time for process adaptations.

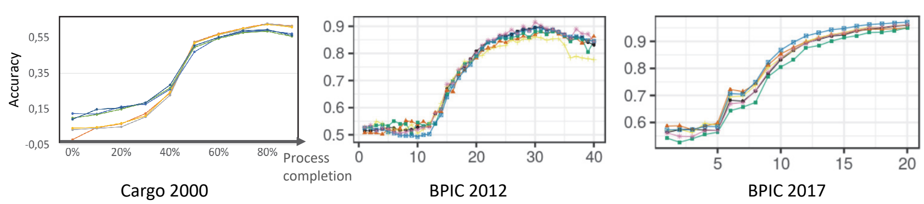

Fig. 2. Prediction earliness vs. prediction accuracy for different data sets (from [20] and [34])

3 Actions taken

3.1 Exploiting advanced analytics for decision support

One of the main actions to leverage the data availability described in Section 2.2 was to employ advanced analytics to provide decision support for terminal operators, thereby helping them better manage terminal processes.

The key concept we prototypically developed in the duisport case is the so-called *terminal productivity cockpit (TPC)*. The TPC exploits advanced data processing and predictive analytics capabilities to facilitate terminal operators in proactive decision making and process adaptation. In particular, the terminal productivity cockpit leverages data-driven deep learning techniques for predictive business process monitoring (see Sections 3.2–3.3). Figure 3 shows a screenshot of the TPC prototype, which visualizes the current and predicted situation in the duisport terminal.

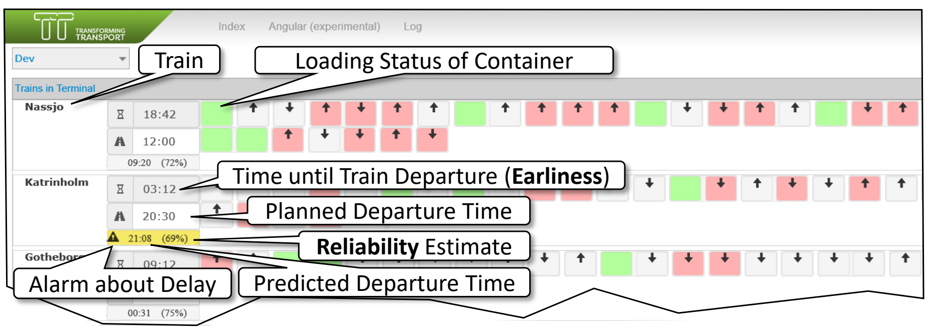

Fig. 3. Screenshot of Terminal Productivity Cockpit – TPC prototype (excerpt)

For each train that is currently in the terminal (in one of the seven transhipment tracks), the TPC shows the following information:

- *Loading status of container:* Each train can carry multiple containers. The TPC shows the status for each of the containers of a train. The arrows indicate the scheduled activities per container, with an upward-facing arrow indicating that a container is to be offloaded, while a downward-facing arrow indicating that a container is to be loaded onto the train. Green means a container has been successfully loaded onto the train, while red indicates a potential problem in container loading.
- *Planned departure time:* For each train, the scheduled departure time is shown. This is essential information, as each train usually has a fixed time slot when it has to depart. On the one hand, such fixed time slots are imposed by the use of the public train infrastructure when leaving the terminal. On the other hand, the train may have to meet fixed departure windows of sea vessels if it connects to a sea port.
- *Time until train departure:* To inform the operators of how much time remains for any potential proactive actions, the TPC shows the time remaining until the planned train departure. This contributes to addressing the *earliness* requirement.

– *Predicted departure time:* The TPC shows the predicted departure time for the train, which takes into account the current status and data from terminal equipment.
– *Alarm about delay:* To facilitate a quick identification of problems, the TPC visibly highlights alarms, i.e., predictions which indicate a potential delay. Thereby, the attention of the operators can focus on important information, which helps address potential cognitive overload [7].
– *Reliability estimate:* In addition to showing an alarm in the case of a delay, the TPC also shows a reliability estimate. The reliability estimate gives the probability (in %) of the predicted delay being accurate, i.e., whether the alarm indeed is a true alarm. This is quite similar to today's weather forecasts. For instance, in addition to predicting that it will rain, a forecast typically also gives the probability that it will rain. Reliability estimates facilitate distinguishing between more and less reliable predictions on a case by case basis [19]. Reliability estimates can help decide whether to trust an individual prediction (and thus alarm) and consequently whether to perform a proactive adaptation of the given process instance [18,35,9].

The reliability estimates together with the earliness indicators of the TPC provide additional information to the terminal operator for decision making. For example, in the situation visualized in Figure 3, the terminal operator is informed that the train to Katrinholm scheduled for 20:30 may be delayed (until 21:08) with a probability of 69% and that there would be 3:12 hours remaining for any proactive action. Given the relatively low probability that the prediction is correct and the little time remaining for taking proactive actions (e.g., rescheduling the terminal workforce may take around 3 hours), the terminal operator may decide not to act in this specific case.

3.2 Ensemble deep learning for predictive process monitoring

We compute the aforementioned predictions and reliability estimates by using ensembles of deep learning models. Ensemble prediction is a meta-prediction technique where the predictions of m prediction models are combined into a single prediction [28].

In the literature, ensemble prediction is primarily used to increase aggregate prediction accuracy. In our case, using ensembles of deep learning models provided an 8.4% higher accuracy when compared to a single deep learning model (as used in [20]). However, increased accuracy is not the main reason why we use ensembles in our approach. We use ensembles in order to compute good estimates of the prediction reliability [19].

Fig. 4 gives an overview of our approach. Each of the individual deep learning models of the ensemble delivers a prediction of the train departure time $T_{i,j}$ for each so-called *checkpoint* j. Using these individual predictions, the three main pieces of information shown in the TPC are computed employing the strategies defined in [19,18,21]: (1) the predicted train departure time T_j, (2) the alarm about a potential delay A_j, and (3) the reliability estimate ρ_j for the alarm.

For computing the predicted departure time T_j, we follow the recommendations in [1] and compute the mean value of the individual predictions $T_{i,j}$, i.e.,

$$T_j = \frac{1}{m} \cdot \sum T_{i,j}$$

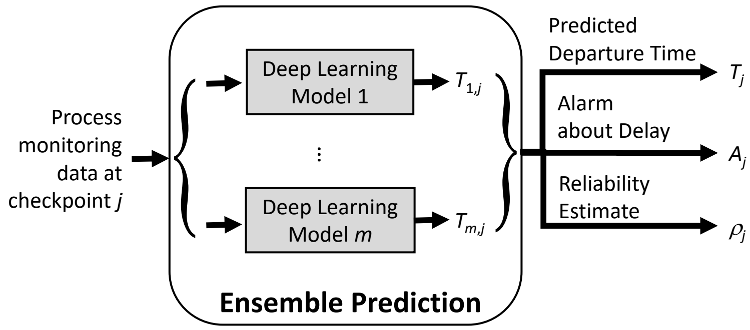

Fig. 4. RNN ensemble for predictive transport process monitoring

For computing the alarm A_j, we first determine for each of the individual predictions $T_{i,j}$ whether they indicate a delay or not by comparing the prediction with the scheduled departure time. This means $A_{i,j} = \mathsf{true}$ indicates a predicted delay. Then, A_j is computed as a majority vote over $A_{i,j}$, i.e.,

$$A_j = \{\mathsf{true\ if}\ |i : A_{i,j} = \mathsf{true}| \geq \frac{m}{2}; \mathsf{false}\ \text{otherwise}\}$$

The reliability estimate ρ_j for A_j is computed as the fraction of predictions $A_{i,j}$ that predicted the delay, i.e.,

$$\rho_j = \frac{1}{m} \cdot |i : A_{i,j} = \mathsf{true}|$$

We use bagging (bootstrap aggregating [5]) as a concrete ensemble technique. Bagging generates m 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 m new training data sets an individual deep learning model is trained. We use bagging with a sample size of 60% to increase the diversity of the ensemble. Bagging contributes to the scalability of our approach, as training the individual models can happen in parallel.

3.3 RNN-LSTMs as Deep Learning Models

We use RNN-LSTMs (Recurrent Neural Networks – Long Short-term Memory) as the individual deep learning models in the ensemble. RNN-LSTMs offer the following advantages over other prediction models:

- *High Accuracy.* RNN-LSTMs have shown significant improvements in prediction accuracy when compared to other prediction models [2,33]. As an example, we experimentally measured accuracy improvements of up to 42% when compared to Multi-Layer Perceptrons [20].
- *Arbitrary Length Sequences.* RNNs can handle arbitrary length sequences of input data [11]. Thus, a single RNN can be employed to make predictions for business processes that have an arbitrary length in terms of process activities. In contrast,

other prediction models (such as random forests or multi-layer perceptrons) may require the special encoding of the input data [15,34,24]. However, these encodings entail information loss and thus may limit prediction performance.
- *Scalability.* RNNs facilitate the scalability of our approach. Assume we have c checkpoints in the business process. A single RNN model can make predictions at any of these c checkpoints [8,33]. If we want to avoid information loss, other prediction models (such as random forests or multi-layer perceptrons) require the training of c prediction models, one for each of the c checkpoints. Performance measurements using a benchmark data set indicate a training time of ca. 8 minutes per checkpoint for multi-layer perceptrons on a standard PC, while the training time for an RNN was 25 minutes[4]. RNNs provide better scalability if the process has many potential checkpoints (in our example already if $c > 3$).

We use RNNs with LSTM cells as they better capture long-term dependencies in the data [33,17]. We use a shared multi-tasks layer architecture as presented by Tax et al. as this provided higher prediction accuracy [33]. In addition to a shared layer, we use three separate layers to predict (1) the next process activity, (2) the time stamp, and (3) the binary process outcome (delay / no delay). Our implementation is available online[5].

4 Results achieved

4.1 Feedback from terminal operators

Based on demonstrations and structured interview sessions with the logport III terminal operator, a qualitative assessment of the TPC with respect to its usefulness and usability was collected. The general feedback was very positive. However one key point was raised during the first rounds of interviews. Given the amount and diversity of data available for the TPC, the terminal operator felt overwhelmed by the amount of information displayed in the TPC. Thus, the terminal operator suggested only providing information that could indicate a problem and its root cause. As a result, the current version of the TPC shows only the information deemed relevant and – as depicted in Section 3.1 – visibly highlights alarms about potential problems in terminal operations.

While the terminal operator was interacting with the TPC, an important side effect was observed. The terminal operator became aware of the broad range of existing data about the terminal and thereby the possibilities that data may provide in finding answers for hitherto unanswerable questions.

4.2 Potential improvements in terminal operations

To quantify the usefulness of the TPC, we analyzed the potential improvements in terminal operations with respect to terminal productivity and costs.

[4] Further performance speedups are possible via special-purpose hardware and RNN implementations. RNN training time reduced to 8 minutes on GPUs (using CuDNN), and further to 2 minutes on TPUs (Tensor Processing Units).

[5] `https://github.com/Chemsorly/BusinessProcessOutcomePrediction`

Productivity. For what concerns the productivity of terminal operations, we set out to measure the improvement of a specific business KPI: “Number of trains leaving the terminal on-time”. This is one of the critical success factors, because – as mentioned above – trains have designated time slots. If a train misses its time slot, re-scheduling is necessary and penalties for late deliveries can occur. Using historic data about terminal operations, we estimated that the use of the TPC may increase the rate of number of trains leaving the terminal on time by up to 4.7%.

Costs. For what concerns costs, we performed controlled experiments using the public Cargo2000 transport data set[6]. The cost models we employed considered various penalty costs in the case of actual delays, as well as various adaptation costs for adapting the running process instance. Details of these experiments are reported in [19,18,21]. Here, we summarize the key outcomes.

We first used a fixed point for predictions (the 50% mark of process execution), and thus did not consider the requirement of prediction earliness. We computed reliabilities via ensembles of classification models (delay/non-delay predictions). When using these reliability estimates to decide on proactive process adaptation, we measured cost savings of 14% on average [19]. When also including the magnitude of a predicted violation (computed from ensembles of regression models), we measured additional cost savings of 14.8% on average [18].

To consider prediction earliness and thus find a trade-off between earliness and accuracy, we used the reliability estimates to dynamically determine the earliest prediction with sufficiently high reliability and used this prediction as basis for proactive adaptation [21]. This meant that the actual checkpoint chosen for a proactive adaptation decision varied among the different process instances, in the same way the reliability estimates varied among the predictions and process instances. Experimental results suggest that dynamically determining the checkpoint offers cost savings of 9.2% on average when compared to using a fixed, static checkpoint. Dynamically determining the checkpoints thus effectively addresses the tradeoff between prediction accuracy and prediction earliness and thus meets the requirements as identified in Section 2.3.

5 Lessons learned

To complement the results from above, we present our main recommendations based on the technical lessons learned regarding the use of big data for predictive analytics:

- *Deep learning works well without extensive hyper-parametrisation.* If enough good quality data is available (like in our case), we experienced that deep learning techniques provide high prediction accuracy without the need for extensive hyper-parameter tuning. In addition, the deep learning models we used (RNNs) did not require special encoding of the input data. Thus, consider using deep learning to make the engineering of data-driven predictive process monitoring solutions more productive!
- *Data quality is a key concern for the usefulness of data analytics.* Data quality is an important concern in data analytics (“garbage in – garbage out”), but also a

[6] Available from `https://archive.ics.uci.edu/ml/datasets`

very resource- and time-intensive activity. With respect to data quality we had to face missing data (e.g., because it was not available in digital form or because of network outages), cope with low data accuracy (due to imprecise measurements), and handle data timeliness (due to delays in data collection). Thus, plan sufficient time and effort for data quality and refinement of data collection!

– *Data processing and integration can consume significant time and resources.* We estimate that data processing, integration and quality assurance consumed around 80% of the resources spent in the duisport pilot case. The reasons were manifold. Oftentimes, we did not have control over the data from third parties (such as equipment manufactures), or data collection and semantics drifted over the course of development. Other examples were telemetry data using different coordinate systems (such as GPS vs XYZ) and timestamps being based on non-synchronized clocks. Thus, plan sufficient time for data processing and integration!
– *Operators benefit from information about data reliability.* Getting additional information about how reliable an individual prediction helps operators decide whether to act on a prediction or not. It supports operators in finding the earliest prediction with sufficient accuracy, thereby allowing more time for proactive actions. In addition, we observed that operators benefit from understanding how reliable the actual data is; e.g., in the form of descriptive analytics outcomes or when visualized in the terminal productivity cockpit. Thus, consider augmenting descriptive and predictive analytics results with reliability estimates, confidence intervals, error ranges, etc. in order to provide additional support to process operators for decision making!

6 Conclusions and Perspectives

The duisport case we presented in this paper shows how data-driven deep learning can deliver profound transport process innovation. We have shown the *feasibility* of our deep learning approach by implementing it as part of a terminal productivity cockpit prototype. The terminal productivity cockpit provides decision support to terminal operators for proactive process adaptation. The *viability* of our approach is supported by an estimated improvement in a key business KPI, as well as experimentally measured cost savings when compared to terminal operations without using proactive adaptation. The *desirability* of our approach is confirmed by positive feedback received from the terminal operator during interviews.

The continuing significant growth of transport data volumes and the rates at which such data is generated will be an important driver for the next level of business process innovation in transport: Data-driven Artificial Intelligence [31]. From an industrial point of view, artificial intelligence means algorithm-based and data-driven computer systems that enable machines and people with digital capabilities such as perception, reasoning, and learning, as well as autonomous decision making and actuation. Building on today's promising results in using artificial intelligence, we can expect artificial intelligence to deliver the next level of productivity improvements in transport.

Acknowledgments. Research leading to these results received funding from the EU's Horizon 2020 R&I programme under grant agreement no. 731932 (TransformingTransport) and 732630 (BDVe).

References

1. Breiman, L.: Bagging predictors. Machine Learning 24(2), 123–140 (1996)
2. Camargo, M., Dumas, M., Gonzalez-Rojas, O.: Learning accurate LSTM models of business processes. In: Hildebrandt, T., van Dongen, B., Rglinger, M., Mendling, J. (eds.) BPM 2019, Vienna, Austria, September 1-6, 2019. LNCS, Springer (2019)
3. Castiñera, R., Metzger, A.: The TransformingTransport project - Mobility meets big data. In: 7th Transport Research Arena (TRA 2018), Vienna, Austria, April 16-19, 2018. Elsevier Transportation Research Procedia (2018)
4. DHL: Big data in logistics: A DHL perspective on how to move beyond the hype (2014)
5. Dietterich, T.G.: Ensemble Methods in Machine Learning, pp. 1–15. Springer Berlin Heidelberg, Berlin, Heidelberg (2000)
6. Directorate-General for Mobility and Transport (European Commission): EU transport in figures: Statistical pocketbook 2018 (2018)
7. Endsley, M.R.: Designing for Situation Awareness: An Approach to User-Centered Design, Second Edition. CRC Press, Inc., Boca Raton, FL, USA, 2nd edn. (2011)
8. Evermann, J., Rehse, J., Fettke, P.: Predicting process behaviour using deep learning. Decision Support Systems 100 (2017)
9. Fahrenkrog-Petersen, S.A., Tax, N., Teinemaa, I., Dumas, M., de Leoni, M., Maggi, F.M., Weidlich, M.: Fire now, fire later: Alarm-based systems for prescriptive process monitoring. CoRR abs/1905.09568 (2019), `http://arxiv.org/abs/1905.09568`
10. Francescomarino, C.D., Ghidini, C., Maggi, F.M., Milani, F.: Predictive process monitoring methods: Which one suits me best? In: Weske, M., Montali, M., Weber, I., vom Brocke, J. (eds.) BPM 2018, Sydney, Australia, September 9-14, 2018. LNCS, vol. 11080, pp. 462–479. Springer (2018)
11. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016)
12. Gutierrez, A., Cassales Marquezan, C., Resinas, M., Metzger, A., Ruiz-Cortés, A., Pohl, K.: Extending WS-Agreement to support automated conformity check on transport & logistics service agreements. In: Basu, S., et al. (eds.) ICSOC 2013, Berlin, Germany, December 2-5, 2013. LNCS, vol. 8274, pp. 567–574. Springer (2013)
13. Khatri, V., Samuel, B.M.: Analytics for managerial work. Commun. ACM 62(4), 100 (2019)
14. Leitner, P., Ferner, J., Hummer, W., Dustdar, S.: Data-driven and automated prediction of service level agreement violations in service compositions. Distributed and Parallel Databases 31(3), 447–470 (2013)
15. Márquez-Chamorro, A.E., Resinas, M., Ruiz-Cortés, A.: Predictive monitoring of business processes: A survey. IEEE Trans. Services Computing 11(6), 962–977 (2018)
16. Mehdiyev, N., Emrich, A., Stahmer, B.P., Fettke, P., Loos, P.: iprodict - intelligent process prediction based on big data analytics. In: Brambilla, M., Hildebrandt, T. (eds.) BPM 2017 Industry Track, Barcelona, Spain, September 10-15, 2017. CEUR Workshop Proceedings, vol. 1985, pp. 13–24. CEUR-WS.org (2017)
17. Mehdiyev, N., Evermann, J., Fettke, P.: A multi-stage deep learning approach for business process event prediction. In: Conf. on Business Informatics (CBI 2017), Thessaloniki, Greece, July 24-27, 2017 (2017)
18. Metzger, A., Bohn, P.: Risk-based proactive process adaptation. In: Maximilien, E.M., Vallecillo, A., Wang, J., Oriol, M. (eds.) ICSOC 2017, Malaga, Spain, November 13-16, 2017. LNCS, vol. 10601, pp. 351–366. Springer (2017)

19. Metzger, A., Föcker, F.: Predictive business process monitoring considering reliability estimates. In: Dubois, E., Pohl, K. (eds.) CAiSE 2017, Essen, Germany, June 12-16, 2017. LNCS, vol. 10253. Springer (2017)
20. Metzger, A., Neubauer, A.: Considering non-sequential control flows for process prediction with recurrent neural networks. In: 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA 2018), Prague, Czech Republic, August 29-31, 2018. IEEE Computer Society (2018)
21. Metzger, A., Neubauer, A., Bohn, P., Pohl, K.: Proactive process adaptation using deep learning ensembles. In: Giorgini, P., Weber, B. (eds.) CAiSE 2019, Rome, Italy, June 3-7, 2019. LNCS, vol. 11483. Springer (2019)
22. Metzger, A., Sammodi, O., Pohl, K.: Accurate proactive adaptation of service-oriented systems. In: Camara, J., de Lemos, R., Ghezzi, C., Lopes, A. (eds.) Assurances for Self-Adaptive Systems. pp. 240–265. Springer (2012)
23. Metzger, A., Thornton, J., Valverde, F., Lopez, J.F.G., Rublova, D.: Predictive analytics and predictive maintenance innovation via big data: The case of TransformingTransport. In: 13th Intelligent Transport Systems - European Congress (ITS Europe), Brainport-Eindhoven, The Netherlands, June 3-6 (2019)
24. Metzger, A. et al.: Comparing and combining predictive business process monitoring techniques. IEEE Trans. Systems Man Cybernetics: Systems 45(2), 276–290 (2015)
25. Moreno, G.A., Cámara, J., Garlan, D., Schmerl, B.R.: Flexible and efficient decision-making for proactive latency-aware self-adaptation. ACM Trans. Autonomous and Adaptive Systems 13(1), 3:1–3:36 (2018)
26. Nunes, V.T., Santoro, F.M., Werner, C.M.L., Ralha, C.G.: Real-time process adaptation: A context-aware replanning approach. IEEE Trans. Systems, Man, and Cybernetics: Systems 48(1), 99–118 (2018)
27. OECD: Exploring data-driven innovation as a new source of growth – mapping the policy issues raised by 'big data' (2013)
28. Polikar, R.: Ensemble based systems in decision making. IEEE Circuits and Systems Magazine 6(3), 21–45 (2006)
29. Poll, R., Polyvyanyy, A., Rosemann, M., Röglinger, M., Rupprecht, L.: Process forecasting: Towards proactive business process management. In: Weske, M., Montali, M., Weber, I., vom Brocke, J. (eds.) BPM 2018, Sydney, Australia, September 9-14, 2018. LNCS, vol. 11080, pp. 496–512. Springer (2018)
30. S. Zillner, E. Curry, A. Metzger, R. Seidl (Eds.): European big data value strategic research and innovation agenda (SRIA). Version 4.0, October (2017)
31. S. Zillner, J.A. Gomez, A. Garcia, E. Curry (Eds.): Data for Artificial Intelligence for European economic competitiveness and societal progress – BDVA position statement (2018)
32. Salfner, F., Lenk, M., Malek, M.: A survey of online failure prediction methods. ACM Comput. Surv. 42(3), 10:1–10:42 (2010)
33. Tax, N., Verenich, I., Rosa, M.L., Dumas, M.: Predictive business process monitoring with LSTM neural networks. In: Dubois, E., Pohl, K. (eds.) CAiSE 2017, Essen, Germany, June 12-16, 2017. LNCS, vol. 10253. Springer (2017)
34. Teinemaa, I., Dumas, M., Rosa, M.L., Maggi, F.M.: Outcome-oriented predictive process monitoring: Review and benchmark. ACM Transactions on Knowledge Discovery from Data (TKDD) 13 (2019)
35. Teinemaa, I., Tax, N., de Leoni, M., Dumas, M., Maggi, F.M.: Alarm-based prescriptive process monitoring. In: Weske, M., Montali, M., Weber, I., vom Brocke, J. (eds.) BPM Forum 2018, Sydney, Australia, September 9-14, 2018. LNBIP, vol. 329, pp. 91–107. Springer (2018)
36. Weber, B., Sadiq, S.W., Reichert, M.: Beyond rigidity - dynamic process lifecycle support. Computer Science - R&D 23(2), 47–65 (2009)