12th International Workshop on Business Process Intelligence (BPI 2016)

Introduction to the 12th International Workshop on Business Process Intelligence (BPI 2016)

Boudewijn F. van Dongen[1(✉)], Jochen De Weerdt[2], Andrea Burattin[3], and Jan Claes[4]

[1] Eindhoven University of Technology, Eindhoven, The Netherlands
b.f.v.dongen@tue.nl
[2] KU Leuven, Leuven, Belgium
jochen.deweerdt@kuleuven.be
[3] University of Innsbruck, Innsbruck, Austria
andrea.burattin@uibk.ac.at
[4] Ghent University, Ghent, Belgium
Jan.Claes@UGent.be

1 Aims and Scope

Business Process Intelligence (BPI) is a growing area both in industry and academia. BPI refers to the application of data- and process-mining techniques to the field of Business Process Management. In practice, BPI is embodied in tools for managing process execution by offering several features such as analysis, prediction, monitoring, control, and optimization.

The main goal of this workshop is to promote the use and development of new techniques to support the analysis of business processes based on runtime data about the past executions of such processes. We aim at bringing together practitioners and researchers from different communities, e.g. Business Process Management, Information Systems, Database Systems, Business Administration, Software Engineering, Artificial Intelligence, and Data Mining, who share an interest in the analysis and optimization of business processes and process-aware information systems. The workshop aims at discussing the current state of research and sharing practical experiences, exchanging ideas and setting up future research directions that better respond to real needs. In a nutshell, it serves as a forum for shaping the BPI area.

The 12th edition of this workshop attracted 12 international submissions. Each paper was reviewed by at least three members of the Program Committee. From these submissions, the top six were accepted as full papers for presentation at the workshop. The papers presented at the workshop provide a mix of novel research ideas, evaluations of existing process mining techniques, as well as new tool support. *Ackermann, Schönig* and *Jablonski* present an approach capable of converting DPIL multi-perspective declarative constraints into a logic language, called Alloy, and to simulate them. *De Smedt, Di Ciccio, Mendling* and *Vanthienen* focus on model checking of mixed-paradigm process models (Declare and Petri Net) in the context of

FusionMinerFul. *Rehse* and *Fettke* introduce an approach to learn a hierarchy of reference models from broad but general to detailed but narrow. *Vogelgesang, Rinderle-Ma* and *Appelrath* devise a framework for multidimensional process mining, which provides more interactivity to the analysis. *De Koninck* and *De Weerdt* investigate a novel multi-objective trace clustering approach. Finally, *Ferreira* and *Santos* discuss how often-used abstraction of event logs (i.e., the calculation of direct-follows relation) can be distributed on multiple cores or GPUs.

For the first time this year, all authors were allotted 20 minutes to present their work, followed by 20 minutes of discussion. The discussions this year were very fruitful and many points for future collaboration within the community were identified.

As has become tradition, this year's BPI workshop was accompanied by the BPI Challenge. Due to the sponsorship of GradientECM and their tool Minit, the winners of the BPI challenge were present at the workshop and presented their analysis results in detail.

For the first time, the process discovery contest was co-organized at BPI. This contest was sponsored by Celonis and the winner of this contest was also given time to present their results.

As with previous editions of the workshop, we hope that the reader will find this selection of papers useful to keep track of the latest advances in the BPI area, and we are looking forward to keep bringing new advances in future editions of the BPI workshop. As organizers, we look back at a very succesful workshop and we are looking forward to next year's edition.

Acknowledgement. The organizers thank the PC members for their valuable comments and active participation!

2 Program Committee

Wil van der Aalst	Eindhoven University of Technology, The Netherlands
Andrea Burattin	University of Innsbruck, Austria
Artur Caetano	IST, University of Lisbon, Portugal
Jan Claes	Ghent University, Belgium
Jochen De Weerdt	KU Leuven, Belgium
Boudewijn van Dongen	Eindhoven University of Technology, The Netherlands
Diogo R. Ferreira	IST, University of Lisbon, Portugal
Walid Gaaloul	Computer Science Department Télécom SudParis, France
Gianluigi Greco	University of Calabria, Italy
Daniela Grigori	Laboratoire LAMSADE, University Paris-Dauphine, France
Antonella Guzzo	University of Calabria, Italy
Anna Kalenkova	Higher School of Economics, Russia
Michael Leyer	University of Rostock, Germany
Ana Karla de Medeiros	Achmea, The Netherlands
Jan Mendling	Wirtschaftsuniversität Wien, Austria
Jorge Munoz-Gama	Pontificia Universidad Católica de Chile, Chile

Viara Popova	University of Tartu, Estonia
Manfred Reichert	University of Ulm, Germany
Andreas Rogge-Solti	Vienna University of Economics and Business, Austria
Michael Rosemann	Queensland University of Technology, Australia
Anne Rozinat	Fluxicon, The Netherlands
Domenico Sacca	University of Calabria, Italy
Pnina Soffer	University of Haifa, Israel
Alessandro Sperduti	University of Padua, Italy
Suriadi Suriadi	Massey University, Australia
Seppe Vanden Broucke	KU Leuven, Belgium
Eric Verbeek	Eindhoven University of Technology, The Netherlands
Hans Weigand	Tilburg University, The Netherlands

Summary of the Business Process Intelligence Challenge (BPI Challenge 2016)

Marcus Dees[1,2] and Boudewijn F. van Dongen[1]

[1] Eindhoven University of Technology, Eindhoven, The Netherlands
{b.f.v.dongen,m.dees}@tue.nl
[2] UWV, Amsterdam, Netherlands
marcus.dees@uwv.nl

1 Background

Since 2011, the IEEE Task Force on process mining organizes a yearly Business Process Intelligence Challenge, or BPI Challenge. The goal of this challenge is to bring together practitioners and researchers in the field to show the direct impact of academic work when facing the challenges real-life cases bring. For the BPI challenge, we provide participants with a real-life event log, and we ask them to analyze these data using whatever techniques available, focusing on one or more of the process owner's questions or proving other unique insights into the process captured in the event log.

For 2016, the data was provided by the UWV (Employee Insurance Agency) and the challenge was hosted at the BPI Workshop in Rio de Janeiro, on September 18, 2016. The challenge was sponsored by Gradient ACM, and their tool Minit v.2.

2 Case

The UWV (Employee Insurance Agency) is a Dutch autonomous administrative authority (ZBO) and is commissioned by the Ministry of Social Affairs and Employment (SZW) to implement employee insurances and provide labour market and data services in the Netherlands.

The Dutch employee insurances are provided for via laws such as the Unemployment Insurance Act (WW). The data in this year's data collection pertains to customer contacts over a period of 8 months and UWV is looking for insights into their customers' journeys. The data is focused on customers in the unemployment benefits (WW) process.

Data has been collected from several different sources, namely:

1. Clickdata from the site www.werk.nl of visitors that were not logged in,
2. Clickdata from the customer specific part of the site www.werk.nl,
3. Contact data, showing when customers sent messages to UWV through a digital channel called "werkmap",
4. Call data from the call center, showing when customers contacted the call center by phone, and
5. Complaint data showing when customers complained.

UWV was interested in insights on how their channels are being used, when customers move from one contact channel to the next and why and if there are clear customer profiles to be identified in the behavioural data. Furthermore, recommendations are sought on how to serve customers without the need to change the contact channel. The full dataset is available from https://data.4tu.nl/repository/uuid:360795c8-1dd6-4a5b-a443-185001076eab [1] and further information is available at http://www.win.tue.nl/bpi/doku.php?id=2016:challenge including various documents detailing the information in the data.

3 Sponsors

This year's challenge was sponsored by GRADIENT ECM (http://www.gradientecm.com/). They not only provided free Minit (http://www.minitlabs.com/) licenses for participants, but they also allowed for two selected winners to come to Rio de Janeiro to present their work at the 12th international BPI Workshop held there and to receive the award during the dinner of the International Conference on Business Process Management (BPM2016).

4 The Results

We received several submissions from all over the world, both from academia and from industry. The jury was pleased with the high quality of the contributions in general, but in the end, two submissions were selected as the best one to present their work in Rio:

– Ube van der Ham, with his submission entitled "Marking up the right tree: understanding the customer process at UWV", showing that by manually inspecting the data and using relatively standard data analysis tools many insights can be obtained, and
– Sharam Dadashnia, Tim Niesen, Philip Hake, Peter Fettke, Nijat Mehdiyev and Joerg Evermann, with their submission entitled: "dentification of Distinct Usage Patterns and Prediction of Customer Behavior" showing an innovative technique to predict the next action undertaken by users on the basis of the preceding ten tasks.

Representatives of both submissions came to Rio and presented their analysis to the BPI audience. Their presentations were well-received and showed true professional value and direct applicability of academic research in the field of business process intelligence.

Acknowledgement. The organizers thank the jury members for their valuable comments and active participation!

Reference

1. Dees, M., van Dongen, B.: Bpi challenge 2016 (2016). https://doi.org/10.4121/uuid:360795c8-1dd6-4a5b-a443-185001076eab

Summary of the Process Discovery Contest 2016

Josep Carmona[1], Massimiliano de Leoni[2], Benoît Depaire[3], and Toon Jouck[3]

[1] Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
jcarmona@cs.upc.edu
[2] Eindhoven University of Technology, Eindhoven, The Netherlands
m.d.leoni@tue.nl
[3] Hasselt University, Hasselt, Belgium
{benoit.depaire,toon.jouck}@uhasselt.be

1 Background

Process Mining is a relatively young research discipline that aims to discover, monitor and improve processes based on real facts (and not assumptions) by extracting knowledge from event logs readily available in today's (information) systems [1]. The lion's share of attention of Process Mining has been devoted to Process Discovery, namely extracting process models - mainly business process models - from an event log.

In the last decade, several new techniques for process discovery have been put forward. Each technique has been evaluated on separate event data, thus making it difficult to perform a comparative evaluation. However, in light of a continuously growing of strength and interest in Process Mining as a discipline, it becomes crucial to finally foster a comparison of existing discovery techniques. With this need at hand, we organized the first edition of the Process-Discovery contest, which was co-located with the BPM-2016 Conference in Rio de Janeiro (Brazil).

2 Objectives and Context

The Process Discovery Contest aims to compare the efficiency of techniques for what concerns discovering process models that provide a proper balance between *overfitting* and *underfitting*. A process model is overfitting (the event log) if it is too restrictive, disallowing behavior which is part of the underlying process. This typically occurs when the model only allows for the behavior recorded in the event log. Conversely, it is underfitting (the reality) if it is not restrictive enough, allowing behavior which is not part of the underlying process. This typically occurs if it overgeneralizes the example behavior in the event log. Interested readers are referred to [2] (e.g., Sect. 6.4.3).

The starting point was 10 different *reference* process models that were randomly generated and contained the most typical process-model constructs. For each process model, in February 2016, a *training* event log with 1000 compliant traces was generated and made available to the contestants to be used as input for their techniques to discover the underlying process model. Clearly, the ideal situation was that the reference process model was rediscovered. In fact, in many situations, the reference model was different from the models discovered by contestants.

But, how can one measure how far are the discovered models from the reference models? Since we do not want to give preference to any modeling notation, we could not leverage on existing measures of overfitting and underfitting, which are notation dependent. Therefore, we used a classification perspective to evaluate the quality of a discovered model. For each reference model, we generated a *test* event log containing 20 traces, out of which 10 were compliant and 10 were not with respect to the reference model. A model is good in balancing overfitting and underfitting if it is able to correctly classify the traces in the test event log: Given a trace representing real process behavior, the model should classify it as allowed; Given a trace representing a behavior not related to the process, the model should classify it as disallowed. *With a classification view, the winner is the group that can correctly classify the largest number of traces in all the test event logs. All event logs will have the same weight*.

It is also worth mentioning that two *calibration* event logs were shared on 15 April and 15 May 2016. The *calibration* logs had the same structure as the test logs, namely 10 compliant and 10 non-compliant traces. However, we did not disclose which traces were (not) compliant. The contestants could submit their classification attempt and we replied stating how many traces were correctly classified. The feedback loops were intended to support participants with assessing the algorithm effectiveness and, consequently, with adjusting their techniques. In fact, as discussed below, the best performing groups profited from the calibration event log. Further information is available at http://www.win.tue.nl/ieeetfpm/doku.php?id=shared:edition_2016.

3 Report on the Results

The result and the winner group were announced on September, 18th, 2016 during the BPI workshop, co-located with the BPM-2016 conference. The winner group was also given a chance to present the work during the workshop. The contest was successful and attracted 14 submissions from Europe, Australia and Asia. The remainder of this section reports on the three groups that scored the best.

The winner was the group composed by **H.M.W. Verbeek** and **F. Mannhardt**, from Eindhoven University of Technology (The Netherlands), with the so-called *DrFurby Classifier* [3]. For each model, it takes the training log and a test log and classifies every trace in the test log whether it matches the training log (positive trace) or not (negative trace). To reduce the number of misclassifications, the DrFurby Classifier uses a combination of two orthogonal approaches. To reduce the number of false negatives (i.e. compliant traces classified as non-compliant), the DrFurby Classifier only uses process-discovery techniques that generates models that classify all training-log traces as compliant. Also, multiple techniques matching this criterion are combined to reduce the number of false positives (i.e. non-compliant traces classified as compliant). This means that, for each reference model, multiple models are discovered: A trace is classified as compliant if and only if the trace is compliant with all discovered models. In particular, Verbeek and Mannhardt employ two techniques that guarantee perfect fitness: the Inductive Miner with maximal decomposition and the Hybrid ILP

Miner with no decomposition. The choices fall on these two techniques because they provide the best classification on the calibration event logs. Their approach was able to correctly classify 193 out of 200 test-log traces (i.e., 20 traces for each of the 10 processes).

We want to give special mention to two runner-ups: Their approaches could correctly classify 192 traces, namely just one trace less than the winner. The first runner-up is **Raji Ghawi**, from American University of Beirut (Lebanon) [4]. For five models, dr. Ghawi employed the Inductive Miner for five processes and the ILP Miner with maximum decomposition for the other five. Similarly to the winner, the choice whether to opt for Inductive Miner or ILP Miner with maximum decomposition for a specific process was driven by the outcomes obtained on the calibration event logs in April and May. Interesting enough, the winner and one runner-up have obtained very good results by employing decomposition. However, the additional trace correctly classified by the winner group, which made the difference, was due to the employment of two discovery techniques to reduce the false positives. The second runner-up was the group composed by **Moshe Steiner** and **Liat Bodaker** under the supervision of **Arik Senderovich**, from Technion–Israel Institute of Technology (Israel) [5]. The approach is based on the Alpha+ algorithm and is used by all 10 models. To overcome the limitations of Alpha+, the mined models are improved/repaired, based on the log footprints. Last but not least, some models were subsequently improved in an ad-hoc fashion. This submission is worth of interest because it tries to overcome the limitation of Alpha+; but, on the other hand, many adjustments are rather ad-hoc and, hence, they are not generally applicable.

It is worth concluding by mentioning that two groups submitted approaches based on Recurrent Neural Networks (by N. Tax and N. Sidorova, Eindhoven University of Technology, The Netherlands) and on Bayesian Networks (by B. Blaskovic, University of Zagreb, Croatia). Although their approaches do not provide traditional process models, they are perfectly legitimate in consideration of the classification nature of the contest. Looking at these submissions, pure data-mining techniques seem to be outperformed by process-mining techniques. This is yet another case that illustrates the importance of process mining how it differs from data mining: Process mining promotes time- and sequence-related information as first-class citizens.

Given the large success, we plan to repeat the experience at BPM 2017 in Barcelona. We plan to improve the contest in the light of several valuable comments received during the BPM-2016 conference.

Acknowledgement. The organizers want to thank all contestants that made an invaluable effort to participate. Special mention goes to the winner and the runner-up groups to be willing to prepare detailed technical reports, which are cited in this summary.

References

1. van der Aalst, W.M.P., et al.: Process mining manifesto. In: Daniel, F., et al. (eds.) BPM 2011 Workshops, Part I. LNBIP, vol. 99, pp. 169–194. Springer, Berlin (2012)

2. van der Aalst, W.M.P.: Data science in action. In: Process Mining, 2nd edn. Springer, Heidelberg (2016)
3. Verbeek, H., Mannhardt, F.: The DrFurby classifier submission to the process discovery contest @ BPM 2016. Technical report BPM center report BPM-16-08, bpmcenter.org (2016)
4. Ghawi, R.: Submission to the process discovery contest @ BPM2016. Technical report (2016). arXiv:1610.07989
5. Shteiner, M., Bodaker, L., Senderovich, A.: Process discovery contest 2016: Heuristic Alpha+ Miner (HAM). Faculty of Industrial Engineering and Management Report Number IE/IS-2016-02, Technion–Israel Institute of Technology (Israel) (2016)