Combining Sonification and Visualization for the Analysis of Process Execution Data

Tobias Hildebrandt
Faculty of Computer Science
University of Vienna, Austria
tobias.hildebrandt@univie.ac.at

Felix Amerbauer
Faculty of Computer Science
University of Vienna, Austria
felix@amerbauer.at

Stefanie Rinderle-Ma
Faculty of Computer Science
University of Vienna, Austria
stefanie.rinderle-ma@univie.ac.at

***Abstract*—Business process execution data is analyzed for different reasons such as process discovery, performance analysis, or anomaly detection. However, visualizations might suffer from a number of limitations. Sonification (the presentation of data using sound) has been proven to successfully enhance visualization in many domains. Although there exist approaches that apply sonification for real-time monitoring of process executions, so far this technique has not been applied to analyze process execution data ex post. We therefore propose a multi-modal system, combining visualization and sonification, for this purpose. The concepts are evaluated by a prototypical ProM plugin as well as based on a use case.**

***Index Terms*—sonification; visualization; multi-modal displays; process data analysis; business process management**

I. Introduction

Based on a business process description or model, during runtime, a few up to thousands of process instances are created, initiated, and started. During the execution of these process instances, execution data is generated and logged. Typically, business process execution data is comprised of events that represent the execution of single process tasks such as *assemble unit* enriched by time stamps, instance identifiers, and further information. Business process execution data is a valuable source for answering process-related questions. Most prominently, process mining refers to a bundle of techniques operating on business process execution data to, for example, discover process models, check the conformance of execution data with a given process model, and conduct process performance analysis. One of the grand challenges in process mining is to provide understandable representations of the analysis results. The reason is that discovered process models and process execution data might hold a high complexity, hampering the understanding and interpretation by non-experts. [1]

The Process Mining Manifesto [1] advocates visual analytics, visualization, and interactive process mining for utilizing the *"the amazing capabilities of humans to see patterns in unstructured data"*. Visualizations along different process perspectives have been developed and some interaction techniques are offered by process mining tools such as ProM[1] [2]. However, visualization of process execution data faces certain challenges, such as the amount of data dimensions that can be conveyed visually, or the ability to convey patterns or details on individual events in huge process execution logs.

One way to overcome limitations of a modality is the combination with other modalities, resulting in a multi-modal representation [3]. Sonification, the presentation of data using non-speech audio, is especially suitable to represent time-based data, such as process execution data. For presenting business process execution data during *runtime*, i.e., for process monitoring, sonification has been successfully applied [4]. However, it has neither been applied nor investigated for supporting the ex post analysis of business process execution data, apart from a few first theoretical considerations by us (e.g. [5]). In addition, approaches for the multi-modal representation of business process data combining visualization and sonification are missing. This leads to the following research questions:

(1) How can multi-modal visualization and sonification be designed to best support users in the ex post analysis of business process execution data?

(2) How well does multi-modal sonification support users in their different analysis tasks?

Although we specifically focus on process execution data, we expect that the results are at least partially transferable to the analysis of other data of similar structure (i.e. event-based data that is mostly qualitative). The main artifact is the concept for the multi-modal representation and analysis of business process execution data described in Section III ($\mapsto$ Question 1). The concept is comprised of visualization, sonification, and an interaction design based on a graphical user interface for the analysis of process execution data. The concept is demonstrated and evaluated based on a prototypical implementation as plugin for the ProM framework with a use case (cf. Section IV) ($\mapsto$ Question 2).

II. Related work

Our main contribution lies in a concept for enhancing process analysis with techniques from sonification, and the visualization aspects serve mainly to demonstrate our approach. Therefore, we do not discuss the related work concerning process data visualization in detail, and instead focus more

[1] http://www.promtools.org

on the body of work concerning sonification and multi-modal approaches.

Visualization of process data often refers to the visualization of process models [6]. In contrast, only few approaches were proposed for visualizing process event data. In [7], the visualization of instance traffic refers to an aggregation of the number of executed process instances along process paths. In process mining, dotted charts are the main visualization for the occurrence of process event data over time [2]. To our best knowledge, there are no approaches that apply sonification for the purpose of analyzing historic business process data, apart from a few first theoretical considerations into this direction by us (e.g. [5]). We have researched the usage of sonification for the real-time monitoring of process executions (e.g. [8]). However, due to several reasons, many results of process monitoring are not directly transferable to process analysis: process monitoring is often performed as a background task in parallel to other work, and executed for a long period of time (often over a complete work day), thus aesthetic aspects are crucial. Process analysis on the other hand is typically a task that is performed for a confined period of time, but with the users' full attention. Therefore, especially sonification techniques that can condense large amounts of data in a short time are adequate.

There is however a growing body of research concerning the usage of sonification for analyzing historic data outside the business process domain. Unfortunately, the majority of research seems to deal with quantitative, continuous data, and less with qualitative data, like discrete events. In [9], sonification is applied to detect anomalies in twitter data. However, unlike the approach presented in this paper, the authors only sonify condensed meta data, and not individual events. The same approach is used by the tool "E-Rhythms Data Sonifier"[2], in which the frequency of different types of event occurrences can be mapped onto volume or pitch. One of the few approaches that sonify discrete events was presented in [10]. The authors present a sonification of nominal event data using melodic Earcons (short abstract sounds or melodies), which has served as an inspiration for the sound design of our approach. [11] presents a tool that sonifies HTTP requests from web server logfiles. Similar to our concept, it sonifies individual events that contain primarily nominal data (such as return codes).

In general, of the few approaches that exist in the area of sonification for the analysis of qualitative and nominal data, the majority are *pure* sonifications, and not multi-modal approaches like the one presented in this paper. Therefore, specifically questions of integrating visualization and sonification could not be answered by previous works. Furthermore, even though we have conducted no real user study yet, our approach offers a preliminary evaluation based on a use case, while most of the mentioned approaches have not been evaluated yet.

[2]http://www.jackjamieson.net/blog/e-rhythms-data-sonifier/

III. Concept for Multi-Modal Process Analysis

The usability and thus user acceptance of systems for multi-modal process data analysis heavily relies on the design decisions that are made, especially the design of the interactive process of creating sonification mappings. This is because most potential users will not be experienced with sonification as in western culture, the visual sense is very predominating and sonification is in many areas still a research topic.

As sonification for exploratory data analysis is a relatively young discipline, there is still a lack of established guidelines for best practices concerning sound- and mapping design in most areas. As log files typically contain quantitative as well as qualitative data, suitable mappings for all data types should be available. There are three basic types of mappings that need to be supported. Timestamp refers to the concept that the time intervals between individual events should be represented. Nominal variables are not represented by numbers and thus contain no inherent ordering, such as different activity names or organizational roles. Continuous variables on the other hand are represented by numerical values, such as costs.

A. Sonification

Systems that apply sonification to convey process execution data should consider several aspects, as suggested in our previous work [12]. Although the referenced literature study focuses on real-time monitoring, these suggestions can serve as a first starting point for a concept. In general, users should be able to customize the mapping from data to sound, and adjust it according to their preferences. Earcons have been proven successful for sonifying discrete events as they exist in process event logs. Especially such Earcons that are based on real world instruments and that adhere to concepts from motif design and melodic contours seem suitable to convey process events. Auditory Icons (recordings of real world sounds that can be mapped to an event, like the paper basket of an operating system being emptied) would also be possible, but Earcons are usually more versatile and more generic. In the context of process execution data, it might be difficult to find auditory analogies to all concepts that should be conveyed, e.g. task names. Thus, our proposed concept bases on parameterized motif Earcons, similar to the ones described in [10] or [4]. Parameterized Earcons are Earcons that enable the mapping of quantitative variables onto acoustic properties, such as pitch.

B. Visualization

The visualization in our multi-modal concept has the main purpose to show the effects of sonification mappings that the users create, and demonstrate the combination of the two modalities. Therefore the visualization should follow the sonification, meaning that changes performed in the sonification mapping should be reflected in the visualization. The basic principles of the visualization are based on the dotted chart visualization available in ProM, which can be seen as a baseline for visualizing business process event data [2]. It intends to offer the user an overview concerning the distribution of traces and events over time at a glance. Color

and shape are used to encode further information, such as the type of activity or the trace, and at the same time represent the sonification mapping that has been applied. Table I presents the different types of process data and how they can be mapped onto visual and auditory properties. While mapping of the timestamp is implicit, and thus cannot be adjusted by the user, mappings concerning nominal and continuous variables can be customized.

TABLE I: Visualization and sonification mappings

Type	Visualization	Sonification	Assignment mandatory
Timestamp	position (x-axis)	playback time	implicit
Nominal	shape, color	instrument, melody	yes
Continuous	not implemented	volume, panning	no

C. Interaction Design

The interaction design should support the user in the process analysis workflow as best as possible. One major aspect in reaching this goal is to simplify the creation and testing of a sonification mapping and help the user to memorize how information aspects available in the event logs are sonified. As the user may not be trained in defining a sonification, an iterative creation of the mapping might be preferable. In such an iterative process, as demonstrated in Fig. 1, users should be able to quickly change mappings in order to try out different options.

Both visualization as well as sonification follow the information-seeking mantra ”Overview first, zoom and filter, then details-on-demand” [13], which is reflected in the user interface. This is incorporated by the ability to zoom in visually, as well as speed up and slow down for the sonification. Furthermore, although there is a plethora of ProM plugins that allow to filter and select e.g. for specific traces or activity types, we believe that switching between plugins interrupts the workflow, which is why we tried to incorporate the most important options directly in the plugin. A typical sequence of steps when trying out different types of sonification mappings could be to assign different melodies to different activity types, listen to the result, filter to hear only one activity type, try out a different mapping and play the filtered sonification, then enable all activity types again and listen to the result. If during this interactive process the user would have to switch plugins, this could interrupt the concentration.

IV. Evaluation and Presentation of Concept

In general, as we do not suggest to use sonification as a replacement neither for visualization, nor for automatic processing and filtering, we do not compare our approach to existing visualization or process mining approaches. Due to space limitations, we decided against a process-discovery use case. The reason for this decision is that process discovery is an area that is already very well targeted by sophisticated and well-working process mining algorithms. Furthermore, we believe that due to the characteristics of sonification, our approach is more targeted towards use cases that are

Fig. 1: Interaction flow

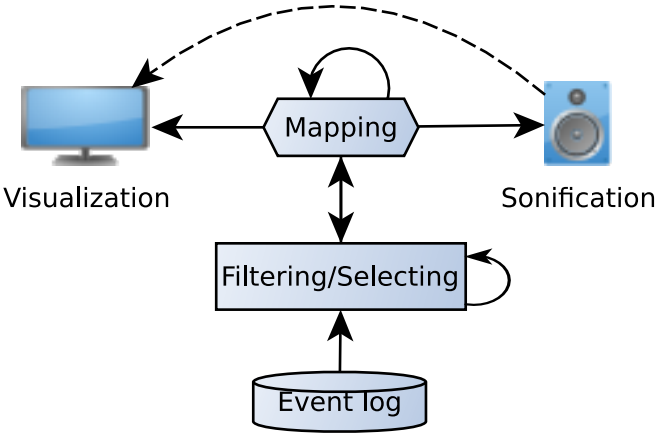

based on an already discovered or defined process model, such as anomaly detection, process improvement, root cause analysis of irregularities or errors, and retrospective process performance monitoring. The main reason for this is that our auditory perception is especially suitable to hear deviations and irregularities, especially in otherwise regular or repeating sound streams. Process discovery on the other hand is typically more focused on discovering *normal* and *typical* behavior instead of outliers. The use case at hand concentrates on anomalies and irregularities, especially such that are not targeted by *classical* process anomaly detection, such as anomalies that are not based on a single trace compared to normal execution, but instead on instance-spanning event data over time.

Our fictive evaluation use case bases on the order process of a web shop. The execution time of each activity varies within a given range. The generated log file contains the web shop activity between 27/11/2015, 6 PM and 28/11/2015, 6 AM. Within this 12 hour period, 11,918 events were produced. As the path of each trace depends to a certain degree on a random variable, several anomalies are contained in the generated execution log *natively*, such as traces with an usual number of loops. We furthermore explicitly applied two fictive anomalies to the order process:

1) The activity decreases throughout the night, except for a short, 5 minute long burst of activity.
2) Between 12 PM and 2 AM, 100% of availability checks lead to the *Cancel System* activity.

In this scenario, we assume the owner of the web shop to be the fictive user of our prototype, who performs process data analysis every morning based on the logs from the previous night. He wants to answer the following questions:

- Have there been problems or irregularities with the technical execution of the process, or with the availability of the server?
- Have there been phases of especially high or low activity?
- Have there been users who where browsing and selecting more articles compared to the average?

- Have there been other noteworthy/irregular occurrences in individual traces?

All sound files referenced in this chapter by the speaker icon can, together with the log file that has been used and video recordings of the plugin, be downloaded from: **http://cs.univie.ac.at/wst/research/projects/project/infproj/1063/**. In order to analyze the event log of last night, the web shop owner first loads the generated log file into the prototype (see Fig. 2).

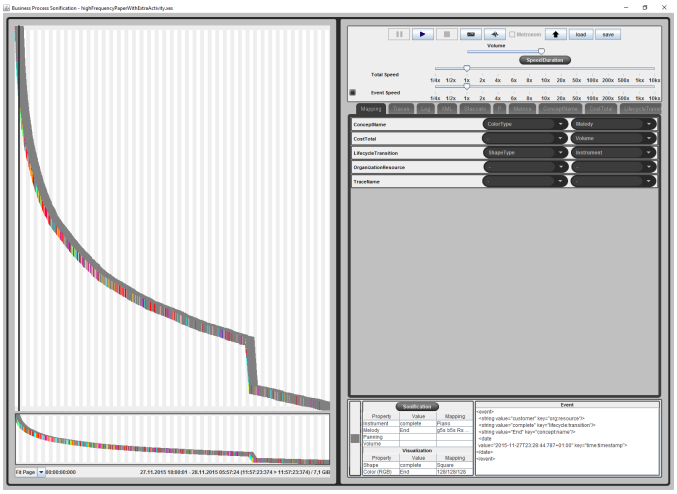

Fig. 2: ProSon GUI with unfiltered event log loaded.

On the right hand side, the menu with different options for playback, filtering and mapping as well as detailed event information can be found. There is such a high number of events, that one cannot derive much information from the visualization without filtering or zooming in. Playing back the events at the same speed as they have been logged is in this case not a viable option, as the playback would have a duration of 12 hours (as does the log file). Thus, when one wants to obtain a overview over the data in a short amount of time, the playback speed has to be increased.

A. Development of Traffic

The first thing our fictive user wants to do, is to analyze the development of user activity of the previous night, and if there have been activity peaks. Thus, in the beginning he is not interested in all activity types, but just in such of the type *start page*, which is why he filters for this activity type using the right hand side menu. As even after filtering there exists a high number of events, the user enters a melody that consists only of one note (a *C* in 3rd octave) to be played whenever the event appears. A longer melody consisting of 3 or 4 short notes, perhaps with different durations and short breaks between them would be more recognizable. However, as such a melody would take more time to be played, this would restrict the possible increase in playback speed. A single note on the other hand can be played in a short period of time. As *start page* is the only activity enabled, it is not necessary to be able to distinguish different melodies (i.e. activities). As no variable is mapped onto the concept *instrument*, all notes are played using the default instrument ”piano” (which of course can be adapted in the menu). The user increases the playback speed, and the event playback speed accordingly. This way, not only the pauses between the events are decreased, but the playback speed of the events is increased as well, meaning that the duration of the notes is decreased to ”fit” all events into the shortened time frame.

tracescomplete.mp3 shows a compressed playback that has a duration of just 20 seconds, which constitutes a compression factor of 1:2160. As in these 20 seconds, 1650 events of the type *start page* are contained (82.5 per second), individual piano notes can hardly be made out anymore. Instead, as the short piano notes are played whenever the respective event occurs, several events of simultaneous traces often overlay each other. The resulting sonification has similarities with that of a Geiger counter, both in terms of the resulting output and the principles behind the sonification mapping. Just as a Geiger counter measures the level of radiation, this resulting sonification measures the level of process activity.

What can be detected from the sonification (and from the visualization) is that there is a short burst of activities towards the end of the log file. What also can be heard, but not so easily seen in the visualization without zooming and scrolling, is that the density of activities decreases more or less steadily towards the end of the log file (apart from the mentioned anomaly). The user can now click on individual events that are of interest to obtain details on them, e.g. their timestamp, in order to conduct root cause analysis.

B. Development of Orders

In the next step, the user would also like to know if the development of orders is similar to that of new traces, as not every user is placing an order. Therefore, he uses the menu to filter out all activities except such of the type *order*.

order.mp3 shows a 10 second long sonification.

What can be deducted from this sonification is that the placement of orders in general follows the development of traces. However, the web shop owner can see already in the small bottom left overview visualization (see Fig. 3) that there is a relatively large period of time towards the end where no orders have been placed. He could then zoom into that period and, by activating the other activity types again, find out which traces are concerned to find out what could have caused this interruption of orders.

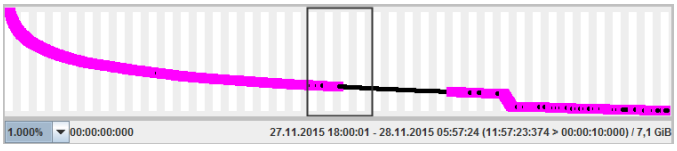

Fig. 3: Overview visualization of log filtered for event type *order*.

C. Browsing and Choosing

Another point the web shop owner is interested in, is if there have been users that where browsing and selecting more articles than the average, which can indicate that they did not find what they were looking for. The web shop owner filters only for the two activity types he is interested in, in this case *browse* and *choose*, and assigns two single notes, a lower one for *browse* (C in 4th octave) and a higher one for *choose* (D in 6th octave). This follows a logical analogy, as the activity *browse* always precedes the activity *choose*, therefore the increase in pitch follows the process path.

Next, to be able to mentally assign an activity to a trace, the user maps trace to the concept *instrument*. This means, that all events of a particular trace are played using the same instrument. Thus, if e.g. trace 1 is mapped to the instrument "piano", a *browse* event in trace 1 is sonified by playing the note *C* in 4th octave on a piano. For each trace, a different instrument is mapped, until after 127 traces the available instruments repeat themselves.

browsechoose.mp3 presents a sonification condensed to 20 seconds. As could be anticipated, individual events cannot be made out anymore at this speed, and it is very difficult to remember if an instrument has already been played or to mentally group events that belong to the same trace. A possible solution would be to load the log file in the *untangled* mode, which means that while loading the log files are sorted in such a way that a new trace only starts after the previously started trace has completed, with a pause of a few seconds in between. For a log file with a relatively low number of traces it is in such a way possible to detect anomalies or deviances, such as traces with higher numbers of browsing and choosing than usual, by listening to one trace after another.

browsechoosesequence.mp3 presents the same mapping loaded in the untangled fashion. The resulting sonification sounds a bit more ordered and less chaotic, as each trace is played before the next one, but the basic challenge remains the same: at this playback speed it is impossible to make out individual traces or even events. Slowing down the playback speed on the other hand would make a sonification of all traces in a short time unfeasible.

D. Other Irregularities

Finally, the user wants to detect other types of irregularities and anomalies that he was not actively looking for. Thus, he filters out any activity types he is not interested in at all, and maps the others to be visualized and sonified. Fig.4 shows the activities our fictive user is interested in, and how they are mapped to short notes. Again, the pitch increases with advancing process state. Increasing melodies symbolize positive events, decreasing pitches negative ones.

To hear when one trace ended and the next started, *trace* is mapped onto *instrument*.

complete.mp3 presents the sonification with the playback speed increased to 500x.

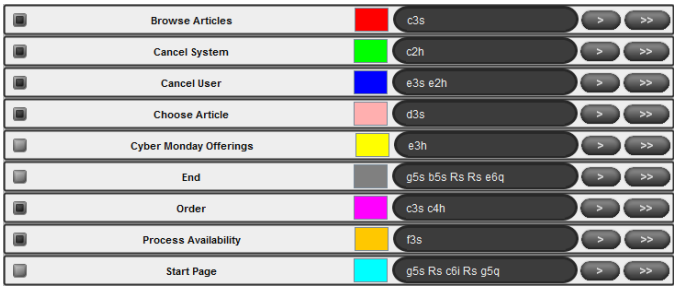

Fig. 4: Mapping of melodies to activities.

completereversed.mp3 presents a reversed mapping, where the activity type is mapped onto *instrument* instead of onto *melody*, *trace* is not mapped. In this example, one can hear instantly that during the period without orders, many users have canceled the process.

V. Conclusion and Discussion

Though there exist a large variety of algorithms for process mining and automated processing of event data, e.g. in the form of ProM plugins, visual analytics is an important addition to process analysis. To further build on our cognitive abilities of pattern recognition, we propose to enhance visualization with techniques from sonification, forming a concept for multi-modal process analysis. To demonstrate our multi-model approach, we developed a prototype as a ProM plugin, which combines visualization and sonification. At hand of a fictive use case, it has been evaluated in order to answer our research questions.

A. Research Questions

Based on the use case evaluation, the two use research questions can be answered:

1) How can multi-modal visualization and sonification be designed to best support users in the ex post analysis of business process execution data?

Based on the evaluation, but also considering existing literature, Earcons, especially such that are based on concepts of Motif design, seem to be suitable to sonify individual events. In general, as demonstrated in the use case evaluation, multi-modal approaches in this domain ideally couple both modalities tightly, in a way that changes in the sonification mapping are directly applied to the visual part. Such systems further should allow the user to customize the visual and acoustic mappings as easily and directly as possible, in order to allow him/her for extensive trial and error. Such customizations should support the already mentioned visual information-seeking mantra ("Overview first, zoom and filter, then details-on-demand")[13], which seems to fit both to visual and to auditory information seeking.

Furthermore, it seems recommendable to avoid the need for switching between plugins or even tools as much as possible, in order to minimize concentration losses on part of the user. Therefore, extensive mechanisms of data filtering and sorting

should be included directly in a multi-modal system. As not all functionalities that are required for mining and analyzing process execution data can and should be reinvented each time, it is recommendable to develop systems as part of a more comprehensive framework such as ProM, in order to avoid the need of switching between applications.

As an example, a typical process analysis workflow for high-frequency data sets might look like the following:

- 1) Get a first overview, and see/hear interesting segments and anomalies.
- 2) Filter and sort either concerning specific activity types, attribute values or time periods in which anomalies have already been detected in the previous step. Such filtering is partially supported by our proposed plugin, but other plugins can also be used.
- 3) Reload filtered logfile, detect more or more specific irregularities, perform root cause analysis by filtering/selecting/zooming, or, if necessary:
- 4) Resort to other ProM plugins or even other systems.
- 5) Repeat steps until all questioned are answered.

2) How well does multi-modal sonification support users in their different analysis tasks?

To summarize, it seems that sonification for the purpose of process data analysis is suitable for the same tasks as visualization: obtaining overviews over large event logs, discovering general trends and large scale anomalies, and comparing selected individual events and traces for anomaly detection and root cause analysis.

With both, visualization and sonification, it seems difficult to detect small anomalies or deviations that concern only single traces in large amounts of event data, such as detecting users who were browsing and choosing more than usual.

Of course, as we did not compare our approach directly to *sole* visualization, no assumptions concerning efficiency improvement over visual-only approaches can be made. However, as both modalities are each better suited to convey different features of data, and different users differ concerning how trained and evolved their perceptions concerning the different modalities are, it can safely be assumed that a combination of both modalities in most cases is preferable to a sole visualization.

B. Limitations and Future Work

In general, most of the anomalies and trends that can be detected by our multi-modal approach can also be found using process mining or other types of algorithmic data processing, at least if one knows what one is looking for. However, the same argument can be brought forward concerning *pure* visualization, and yet tools such as "Dotted Chart analysis" are amongst the most used ProM plugins [14]. We are not suggesting multi-modal approaches as a replacement for the state of the art, neither for process mining and other algorithmic processing, nor for visual analytics. Instead, we merely propose to complement and enhance current approaches with sonification, which we believe will help users in detecting additional features in their process logs. Sonification has, as does visualization, certain strengths and weaknesses (see [4]) that need to be considered. It can be expected that most users will need a training period to familiarize themselves with our prototype, but also with the concept of sonification in general, before being productive. Based on findings from research it can be expected, that this training period may be shorter for people with musical and/or acoustic training. Furthermore, we expect the productivity to be greater, when the user is not only accustomed to our tool, but also how his/her processes *sound*. Thus, if a user analyzes the events of the same process regularly, he will be even more efficient in detecting deviations and irregularities after a while. This is because in the beginning, he has no comparison to *normal* behavior.

We plan to enhance our approach to visualize and sonify also the development of KPIs and other trace-spanning measurements over time. This might require other types of sonification (such as continuous sound streams instead of singular sound events) and mapping techniques. More formal long term evaluations with professional data analysts are planned. Future evaluations will also analyze the potential for process discovery.

REFERENCES

[1] W. v. d. Aalst and e. al, "Process Mining Manifesto," in *Business Process Management Workshops*, ser. LNBIP, 2011.

[2] M. Song and W. M. van der Aalst, "Supporting process mining by showing events at a glance," *Proceedings of the 17th Annual Workshop on Information Technologies and Systems (WITS)*, pp. 139–145, 2007.

[3] M. Rauterberg and E. Styger, "Positive Effects of Sound Feedback During the Operation of a Plant Simulator," ser. EWHCI '94. London: Springer, 1994, pp. 35–44.

[4] T. Hildebrandt, "Towards Enhancing Business Process Monitoring with Sonification," in *Intl. Workshop on Theory and Applications of Process Visualization (TAProViz 2013)*, ser. Lecture Notes in Business Information Processing. Berlin, Germany: Springer, 2013.

[5] T. Hildebrandt, S. Kriglstein, and S. Rinderle-Ma, "On Applying Sonification Methods to Convey Business Process Data," in *CaISE 2012 Forum*. CEUR, 2012.

[6] W. Grossmann and S. Rinderle-Ma, *Fundamentals of Business Intelligence*, ser. Data-Centric Systems and Applications. Springer, 2015. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-46531-8

[7] S. Kriglstein, G. Wallner, and S. Rinderle-Ma, "A Visualization Approach for Difference Analysis of Process Models and Instance Traffic," in *Business Process Management - 11th Int'l Conf.*, 2013, pp. 219–226.

[8] T. Hildebrandt, J. Mangler, and S. Rinderle-Ma, "Something Doesn't Sound Right: Sonification for Monitoring Business Processes in Manufacturing," in *IEEE CBI*, 2014, pp. 174–182.

[9] M. Ballora, R. J. Cole, H. Kruesi, H. Greene, G. Monahan, and D. L. Hall, "Use of sonification in the detection of anomalous events," vol. 8407, Jun. 2012, p. 84070S.

[10] C. Cullen and E. Coyle, "TrioSon: a Graphical User Interface for Pattern Sonification," ser. ICAD 2005, Jan. 2005.

[11] M. Ballora, B. Panulla, M. Gourley, and D. L. Hall, "Preliminary Steps in Sonifying Web Log Data." Washington: International Community for Auditory Display, 2010.

[12] T. Hildebrandt and S. Rinderle-Ma, "Toward a Sonification Concept for Business Process Monitoring," in *Int'l Conf. on Auditory Display*, 2013, pp. 323–330.

[13] B. Shneiderman, "The eyes have it: A task by data type taxonomy for information visualizations," ser. VL '96. IEEE Computer Society, 1996, pp. 336–.

[14] J. Claes and G. Poels, *BPM 2012 International Workshops. Revised Papers*. Springer Berlin Heidelberg, 2013, ch. Process Mining and the ProM Framework: An Exploratory Survey, pp. 187–198.