Factors Affecting Radiologist's PACS Usage

Daniel Forsberg^{1,2} · Beverly Rosipko² · Jeffrey L. Sunshine²

Published online: 20 April 2016

© Society for Imaging Informatics in Medicine 2016

Abstract The purpose of this study was to determine if any of the factors radiologist, examination category, time of week, and week effect PACS usage, with PACS usage defined as the sequential order of computer commands issued by a radiologist in a PACS during interpretation and dictation. We initially hypothesized that only radiologist and examination category would have significant effects on PACS usage. Command logs covering 8 weeks of PACS usage were analyzed. For each command trace (describing performed activities of an attending radiologist interpreting a single examination), the PACS usage variables number of commands, number of command classes, bigram repetitiveness, and time to read were extracted. Generalized linear models were used to determine the significance of the factors on the PACS usage variables. The statistical results confirmed the initial hypothesis that radiologist and examination category affect PACS usage and that the factors week and time of week to a large extent have no significant effect. As such, this work provides direction for continued efforts to analyze system data to better understand PACS utilization, which in turn can provide input to enable optimal utilization and configuration of corresponding systems. These continued efforts were, in this work, exemplified by a more detailed analysis using PACS usage profiles, which revealed insights directly applicable to improve PACS utilization through modified system configuration.

Keywords PACS management · Efficiency · Workflow

Introduction

In a modern radiology department, a picture archiving and communication system (PACS) is among the most vital IT systems required for radiologists to provide patients care. As such, optimal use of PACS is of great importance to ensure an efficient workflow, and consequently numerous research studies have been dedicated to this matter albeit predominately focused on external elements related to PACS. For example, studies have evaluated the effect of different input devices on radiologists' efficiency [1–3], impact of employed display devices [4], or general ergonomic and environmental aspects of the reading context [5–8]. A very limited set of research studies [9–12] have focused directly on PACS usage itself.

A more detailed analysis of PACS usage holds several potential benefits. First of all, it can provide detailed knowledge about how utilization of a PACS varies among users and for different situations, as exemplified in a study for a web-based electronic medical record (EMR) system [13]. Second, the derived knowledge can be used to configure the PACS to optimize the workspace for different users or situations [14]. Third, such analyses can provide input for improving PACS usage by detecting opportunities for available functionality not employed by the radiologists, or inefficiencies in current PACS usage patterns among specific image viewing or analysis tasks [15]. Fourth, it can provide feedback to the software developers of the PACS and guide them, both in terms of developing new functionality and in maintenance of the existing code base, as shown in a study for an EMR [16]. Fifth, a combination of PACS usage data with more traditional timestamp data available from a PACS or RIS provides opportunities for more fine-grained workflow analysis [17].

✉ Daniel Forsberg
daniel.forsberg@sectra.com

¹ Sectra, Teknikringen 20, SE-583 30 Linköping, Sweden

² Department of Radiology, Case Western Reserve University and University Hospitals Case Medical Center, 11100 Euclid Avenue, Cleveland, OH 44106, USA

The lack of research related to direct PACS usage and the potential benefits thereof present a need as well as a motivation for a more detailed analysis of PACS utilization.

The purpose of this paper is to investigate whether radiologist (User), examination category (ExamCategory), time of week (TimeOfWeek), and week (Week) can be considered to have an effect on PACS usage, where PACS usage is defined as the sequential order of computer commands issued by a radiologist's mouse or keyboard selection of functions in the PACS during interpretation and dictation. For ease of statistical analysis, PACS usage is quantified by four high-level metrics: number of commands, number of command classes, bigram repetitiveness, and time to read. Based upon earlier results [12], we hypothesize that the factors radiologist and examination category are the only two factors that have a significant effect on PACS usage as described by all four metrics and that the interaction effect between these two factors would also significantly affect PACS usage.

Material and Methods

This study was considered as exempt from review by our institutional review board as part of a larger quality project in our department. The following subsections provide information about the included material (command logs recording PACS usage from radiologists) and the statistical methods employed to analyze the data.

Command Logs and Pre-processing

The PACS installed at our hospital provides the possibility to log all commands issued by each PACS user. Command logs, spanning May 4 to July 5, 2015 and recording user activity from residents, fellows, and attending radiologists, served as starting point for our analysis.

Before analysis, the logs were pre-processed. This pre-processing consisted of splitting user sessions into individual command traces for each opened examination and removal of commands irrelevant to the reading workflow in the PACS, for example, commands for navigation (among worklists or in an image stack), search commands, and login/logout commands. Each command trace would record user, role, information about the opened examination, and the commands issued. Table 1 provides an example of such a recorded command trace.

In the next step, each command trace was linked to its corresponding examination as recorded in the PACS database, in order to ensure that data included was related to examinations performed between May 4 and June 28, 2015 (8 weeks). In addition, this allowed us to update each command trace with information about examination category (in this work coarsely categorized based upon modality and section).

Table 1 Example command trace of recorded PACS usage

User	001
Role	Radiologist
Examination information	7/12/2015, 9:54 AM, ABDOMEN AP VIEW
DateTime	Command
2015-06-12 12:25:18	SetCurrentExam
2015-06-12 12:25:25	Dictate report
2015-06-12 12:25:47	Pan
2015-06-12 12:25:48	Next hanging
2015-06-12 12:25:54	Next hanging
2015-06-12 12:25:58	Previous hanging
2015-06-12 12:26:33	ArtificialEnd

Next, each set of command traces related to the same examination was analyzed and only command traces from examinations where only a single user during a single occasion had dictated and reviewed the examination were retained. Hence, each command trace will correspond to a unique examination. This approach was decided upon to ensure that analyzed PACS usage would not be confounded by factors related to the teaching workflow or related to the interpretation of individual examinations spread across multiple traces (occasions), where some work is bound to be replicated. In addition, command traces recording a time difference of 5 min or more between subsequent commands were excluded to avoid traces recording radiologists who have gone on longer breaks. Finally, only command traces from attending radiologists were retained for subsequent analysis.

Factors and PACS Usage Variables

For each command trace, the factors User, ExamCategory, TimeOfWeek, and Week were recorded, where values for User are self-explanatory. The examination categories were limited to a smaller set of categorical values, {"CT TH", "CR BN", "CR PD", "MR TH", "MR BD", "CT BA", "CT BN", "MR BR", "CT BD", "GI", "MR BN", "CT NR", "MG", "CR TH", "US", "MR NR"}. The category values refer to modality and section, where CR = computed radiography, CT = computed tomography, MR = magnetic resonance, MG = mammography, US = ultrasound, BA = body angiography, BD = body, BN = bone, BR = breast, GI = gastrointestinal imaging, NR = neuroradiology, and TH = thorax. For TimeOfWeek, the timestamp of the first command in the command trace was used to determine time of week. Four categorical values were used, daytime/nighttime weekday/weekend, same as in [12]. The variable Week was simply set to the week number of the recorded data but in groups of two, i.e., 1–2, 3–4, 5–6, and 7–8.

For each recorded command trace, the PACS usage variables number of commands (how many commands used),

number of command classes (how many different types of commands), bigram repetitiveness (how repetitive a command trace is), and time to read (the time between the first and the last command in the recorded command trace as measured in seconds) were computed. All PACS usage variables except time to read excluded the initial “SetCurrentExam” and the final “ArtificialEnd” commands. These two commands were excluded, since they were present in all command traces and as such do not add any value in describing PACS usage.

The bigram repetitiveness assesses how repetitive a command sequence is by computing the number of available bigram command sequences divided by the maximum number of bigrams for the given command sequence length. For example, given a command sequence represented by the symbols “abcdab,” the available bigrams are “ab”, “bc”, “cd”, and “da”, i.e., four bigrams, and the maximum number of bigrams is five. Hence, for this command sequence, the bigram repetitiveness would equal $4/5 = 0.8$. For command sequences containing a single command, the bigram repetitiveness was set to 1.0. Note that a higher value corresponds to a less repetitive sequence and where 1.0 is maximum, i.e., the command sequence contains no repetitions.

Overview of Included Data

The initially extracted command traces contained data from 70 radiologists and 16 examination categories, in total 42,494 individual command traces (corresponding to 42,494 reviewed examinations). However, the data was not evenly distributed across the main effects (apart from week) and many combinations of the interaction effects missed data completely. For example, no radiologist had read all examination categories, most radiologists would only read during daytime weekdays, a few radiologists had not read during all weeks, and most examination categories were only read during daytime. Because of this, we extracted subsets of data free of gaps in order to test different aspects of the formulated hypotheses as detailed next.

1. To study the effect of which week the examination was read, a subset of data based upon certain radiologists and examination categories was extracted to ensure a complete dataset for all main and interaction effects, avoiding effects from the factor TimeOfWeek by limiting the data to only include daytime weekday. The extracted subset contained data from 10 radiologists and 6 examination categories and corresponded to 11,374 individual command traces.
2. To study the effect of when during a week the examination was read, another subset of data was extracted from all 8 weeks without considering the factor Week. This assumption was necessary to ensure a reasonable size of the included data. The extracted subset contained data

from five radiologists and six examination categories and corresponded to 6060 individual command traces.

3. Finally, to study the effects of user and examination category, we assumed that week and time of week would have little effect on the PACS usage variables, and extracted a dataset without considering week and time of week. The extracted subset contained data from 23 radiologists and 6 examination categories and corresponded to 27,070 individual command traces.

During the extraction of the subsets of data, we attempted to maximize the number of included users and examination categories by iteratively removing users with few read examination categories or examination categories read by few users. Note that for all extracted subsets, the values of the variable ExamCategory always corresponded to {“CR BN”, “CR TH”, “CT BD”, “CT NR”, “CT TH”, “US”}.

Statistical Analysis

The data was analyzed using IBM SPSS Statistics Version 23 and generalized linear models. Initial significance level was set to 0.05, but since each of the four PACS usage variables were tested separately on the same data sets, the significance level was corrected using a conservative Bonferroni correction. Thus, the actual significance level for each dataset and PACS usage variable was set to $0.0125 \approx 0.01$.

Before analysis, the PACS usage variable bigram repetitiveness was converted to a nominal variable by distributing the data into six bins. Five of the bins had a width of 0.2 (bin number 1 [0.0, 0.2), bin number 2 [0.2, 0.4), etc.) but where the last bin only included the value 1.0.

Suitable distributions and link functions for the applied generalized linear models were determined based upon selecting the combination with the best goodness of fit value. The PACS usage variables commands, command classes, and time to read were modeled using a gamma distribution and the identity function as link function, and bigram repetitiveness was modeled using a multinomial distribution and a cumulative logit function as link function.

Results

The results from analyzing the effects of the factors on the extracted datasets using generalized linear models are given in Table 2 for excluding the influence of time of week, in Table 3 for neglecting the effect of week, and in Table 4 for neglecting the effect of both time of week and week. We summarize the pattern of effects among the PACS usage variables in each table and analysis as follows.

From data set no. 1 (Table 2) the results show that the factor Week has no significant influence as main effect on three of

Table 2 *P* values from investigating the effects of the factors User, ExamCategory, and Week and their interaction effects

Factor	<i>P</i> values			
	Number of commands	Number of command classes	Bigram repetitiveness	Time to read
User	<0.001*	<0.001*	<0.001*	<0.001*
ExamCategory	<0.001*	<0.001*	<0.001*	<0.001*
Week	0.227	0.323	1.000	<0.001*
User*ExamCategory	<0.001*	<0.001*	<0.001*	<0.001*
User*Week	0.996	0.609	0.991	0.189
ExamCategory*Week	0.679	0.987	0.800	0.131
User*ExamCategory*Week	0.060	0.574	0.651	0.270

Significant effects are marked with an asterisk (*)

the PACS usage variables (number of commands, number of command classes, and bigram repetitiveness) nor as interaction effect on any of the PACS usage variables, and with Week only significant for time to read. In contrast, the main effects of User and ExamCategory and their interaction effects were found to be significant for all four PACS usage variables.

From dataset no. 2 (Table 3), the results show that the factor TimeOfWeek alone does not have a significant effect on the PACS usage variables, neither as a main effect nor as a part of any two-way interaction effect and for most three-way interactions. However, TimeOfWeek as a part of the full three-way interaction effect is significant for time to read. The main effects of User and ExamCategory and their interaction effects are again significant for all four PACS usage variables.

For dataset no. 3 (Table 4) the factors User and ExamCategory were shown to be significant as main effects and interaction effects for all four PACS usage variables.

One possible approach to perform a more detailed study of differences in PACS usage is to count occurrences of particular command classes in a set of command traces and to normalize each count by the number of included command traces, here referred to as PACS usage profiles. In this work, the PACS usage profiles were focused on a set of commands related to when users modify the hanging of an examination,

for example in the case when the default hanging protocol is deemed inadequate. Relevant commands for this analysis include “MatrixWindow”—opening the workspace where a hanging is modified, “DragImageSeries”—selecting and starting to drag an image series, “HangImageSeries”—hanging a selected image series, “AddOrRemoveHanging”—adding or removing a hanging, “AddOrRemoveMonitor”—adding or removing a monitor in a hanging, “ChangeMonitorLayout”—changing monitor layout in a hanging, “ForceTypes”—forcing either examination or image series types, and “Save”—saving a modified hanging.

Two different sets of PACS usage profiles were created for this more detailed study of differences in PACS usage. For the first set of usage profiles, we visualize the differences in PACS usage for each included examination category from data set no. 3 in Fig. 1. As can be seen for the examination categories CR BN, CR TH, and US, very few manual changes to the hangings are performed. The situation is completely different for the CT-based examination categories and in particular for CT BD and CT TH, where especially a large number of drag and hang operations are performed.

In the second set of PACS usage profiles analyzed, a subset of PACS usage profiles for the examination category CT BD were extracted and grouped based upon radiologist. These

Table 3 *P* values from investigating the effects of the factors User, ExamCategory, and TimeOfWeek and their interaction effects

Factor	<i>P</i> values			
	Number of commands	Number of command Classes	Bigram repetitiveness	Time to read
User	<0.001*	<0.001*	<0.001*	<0.001*
ExamCategory	<0.001*	<0.001*	<0.001*	<0.001*
TimeOfWeek	0.179	0.712	1.000	0.681
User*ExamCategory	<0.001*	<0.001*	<0.001*	<0.001*
User*TimeOfWeek	0.951	0.999	0.665	0.734
ExamCategory*TimeOfWeek	0.387	0.376	0.995	0.342
User*ExamCategory*TimeOfWeek	0.449	0.890	0.454	0.004*

Significant effects are marked with an asterisk (*)

Table 4 *P* values from investigating the effect of the factors User and ExamCategory and their interaction effect

Factor	<i>P</i> values			
	Number of commands	Number of command classes	Bigram repetitiveness	Time to read
User	<0.001*	<0.001*	<0.001*	<0.001*
ExamCategory	<0.001*	<0.001*	<0.001*	<0.001*
User*ExamCategory	<0.001*	<0.001*	<0.001*	<0.001*

Significant effects are marked with an asterisk (*)

usage profiles are visually represented in Fig. 2, where each displayed profile represents a different user. As apparent, there is a large variation between the different radiologists, where some perform on average six drag and hang operations per examination whereas some perform less than one.

Discussion

In this study, we set out to investigate the effects of the factors User, ExamCategory, TimeOfWeek, and Week on PACS usage, as described by the variables number of commands, number of command classes, bigram repetitiveness, and time to read. To enable a reliable statistical evaluation of all the different combinations of the categorical factors, we extracted subsets of data free of gaps in order to test different aspects of the formulated hypotheses in a step-wise manner. We hypothesized that only radiologist and examination category would have any significant effect on PACS usage.

In the first extracted subset of data, the primary goal was study of the effect of the factor Week, i.e., to determine whether there is a difference in PACS usage among separate weeks over 8 weeks. In this case, the results in Table 2 reveal only a single significant effect for the factor Week, on time to read,

whereas the other effects involving the variable Week were found insignificant. One can extrapolate that over a time-span as short as these 8 weeks, the vast majority of command traces for reading radiological images are not likely to change significantly, although some variation in time to read may occur depending on actual case load. As the analyzed data originated from attending radiologists over only 8 weeks, whether user behavior may change over longer time-spans, many months, or years or whether training radiology residents change their behavior within the range of a few weeks remain outside the scope of this work.

In the second subset of data analyzed, the focus was on the effect of the factor TimeOfWeek, i.e., whether there is a difference in PACS usage between work performed during combinations of daytime, nighttime, weekdays, and weekends. Again, the results met the expected hypothesis, that is, that time of week has no significant effect on PACS usage, neither as main effect nor as part of an interaction effect. Recall that to ensure a reasonable size of the included data for number of radiologists and examination categories, we neglected the effect of the factor Week for this analysis. The results from the first dataset showing the factor Week to a large extent with no significant effect on usage demonstrates the validity of this assumption.

In the third analysis using generalized linear models, only the effects of the factors User and ExamCategory were considered. The results clearly show the significance of both factors as main effects along with their interaction effect on all PACS usage variables, confirming the original hypothesis.

To exemplify how found differences can be analyzed in more detail and how obtained knowledge can guide efforts to improve PACS usage, we studied two sets of PACS usage profiles describing how much manual work goes into changing the default hanging protocols. The results here clearly showed a difference between CR and US examinations on the one hand and CT examinations on the other hand, where the CT examinations require notably more rehanging. This highlights improvement opportunities for the default hanging protocols applied to CT, especially for body and thoracic examination categories. Further, the second set of analyzed command usage profiles indicated that there also exist large variations between radiologists for the same exam category. Extended analysis is required

Fig. 1 PACS usage profiles for interacting with hanging protocols over included examination categories from dataset no. 3

Fig. 2 PACS usage profiles for interacting with hanging protocols over included users from data set no. 3 but only for examination category “CT BD.” Each line profile represents a different radiologist

to assess whether these differences are due to preference or dependent on, for example, differences in the actual case load.

Based upon previous work related to analysis of PACS usage as recorded by command logs [12], a number of precautions were taken in order to ensure a high quality of the analyzed data. These include, for example, the filtering of command traces to only include those where it remains highly likely that only a single radiologist has interpreted and dictated the corresponding examinations. Our approach can be considered as rather conservative, since a large amount of command traces were discarded and, possibly, this could cause an unknown bias to be included in the extracted data. On the other hand, this ensured that analyzed command traces are likely to be free of confounding factors related to variable teaching workflows along, or interrupted reading sessions. In this work, we selected to avoid the latter and to accept that an unknown bias could possibly be present.

The choice of factors and PACS usage variables also limits the presented study. Given an appropriately scaled data set, it would be interesting to include other factors such as specialty, years of experience, or years of working with a specific PACS system. In terms of the actual examinations, we only had information on the examination category, but no information available on complexity of each case or on how many prior examinations were included in the reading of each case, factors that most definitely will affect PACS usage. However, it can be assumed that the impact of factors like case complexity and number of priors were reasonably uniformly distributed over the included examinations and should as such have limited impact on our presented results. The factor TimeOfWeek was in this study limited to four categorical values, and given the results, it can be questioned whether the chosen categories and time limits were the most relevant or whether other contextual time factors would show more relevance, for example,

weeks with holidays, major society meetings, or beginning of training year. Another contextual time factor, instead of time of week, would be to focus on the level of workload.

Further, our approach reduced the full complexity of PACS usage to a few quantitative metrics; hence, wider information from the command traces has been excluded (for example order of commands used and underlying usage patterns). Although discovering underlying usage patterns from PACS usage logs has been successfully implemented earlier using process mining [12], we did not include it here since process mining currently lacks standards for comparison of different usage patterns when described by separate process models. However, we do intend to revisit these data to compare differences in derived process models, once standardized methods and metrics for this purpose have been defined, with some initial suggestions already available [18–20].

From a statistical perspective, our choice to reduce the PACS usage to a few quantitative metrics allowed use of generalized linear models to assess significance of multiple factors simultaneously. The use of generalized linear models was preferred given the non-normality of the data, for which using n-way ANOVA/MANOVA was not a suitable option for analysis. As noted earlier, the selected distribution and linking function for each PACS usage variable were chosen empirically based upon the goodness of fit as computed from the estimated models.

Our results may be generalizable though the limitations of a short time-span and data only for attending radiologists as discussed earlier remain. Further, as the data only stems from a single PACS system and a specific PACS vendor, there can be differences intrinsic to site-specific configuration, software design, and available PACS functionality. All are relevant factors to address in future research to better understand the factors most broadly affecting PACS usage.

Conclusions

In conclusion, we have performed a detailed analysis of PACS usage to determine the effects of radiologist, examination category, time of week, and week on PACS usage as recorded by command traces and quantified by number of commands, number of command classes, bigram repetitiveness, and time to read. Using generalized linear models, we confirmed our initial hypotheses that the factors radiologist and examination category have significant effects on PACS usage, both as main and as interaction effects. In addition, both time of week and week had limited significance as main effects and as part of interaction effects on most PACS usage variables.

Further, we performed an in-depth analysis of PACS usage related to rehandling examinations by means of PACS usage profiles. This clearly demonstrated how detailed PACS usage analysis can be relevant for pinpointing system inefficiencies and to provide data supporting continuous improvement efforts.

This work provides a starting point for continued efforts in analyzing system usage data at a very detailed level to ensure a better understanding on how healthcare computer systems are used in practice, which in turn provides input to ensure that corresponding systems are optimally configured and utilized.

Compliance with Ethical Standards

Funding D. Forsberg is supported by a grant (2014–01422) from the Swedish Innovation Agency (VINNOVA).

Disclosures D. Forsberg is an employee of Sectra, a PACS vendor. B. Rosipko has nothing to disclose. J. Sunshine has nothing to disclose.

IRB Statement This study was considered as exempt from review by our institutional review board as part of a larger quality project in our department.

References

- Weiss DL, Siddiqui KM, Scopelliti J: Radiologist assessment of PACS user interface devices. *J Am Coll Radiol* 3:265–273, 2006
- Atkins MS, Fernquist J, Kirkpatrick AE, Forster BB: Evaluating interaction techniques for stack mode viewing. *J Digit Imaging* 22: 369–382, 2009
- Lidén M, Andersson T, Geijer H: Alternative user interface devices for improved navigation of CT datasets. *J Digit Imaging* 24:126–134, 2011
- Goo JM, Choi JY, Im JG, et al: Effect of Monitor Luminance and Ambient Light on Observer Performance in Soft-Copy Reading of Digital Chest Radiographs. *Radiology* 232:762–766, 2004
- Nagy P, Siegel EL, Hanson T, Kreiner L, Johnson K, Reiner B: PACS reading room design. *Semin Roentgenol* 38:244–255, 2003
- Siddiqui KM, Chia S, Knight N, Siegel EL: Design and ergonomic considerations for the filmless environment. *J Am Coll Radiol* 3: 456–467, 2006
- Goyal N, Jain N, Rachapalli V: Ergonomics in radiology. *Clin Radiol* 64:119–126, 2009
- Hugine A, Guerlain S, Hedge A: User evaluation of an innovative digital reading room. *J Digit Imaging* 25:337–346, 2012
- Horii SC, Feingold ER, Kundel HL, et al: PACS Workstation functions: usage differences between radiologists and MICU physicians. *Med Imaging Int Soc Opt Photon* 2711:266–271, 1996
- Erickson BJ, Ryan WJ, Gehring DG, Beebe C: Clinician usage patterns of a desktop radiology information display application. *J Digit Imaging* 11:137–141, 1998
- Erickson BJ, Ryan WJ, Gehring DG: Functional requirements of a desktop clinical image display application. *J Digit Imaging* 14:149–152, 2001
- Forsberg D, Rosipko B, Sunshine JL: Analyzing PACS Usage Patterns by Means of Process Mining: Steps Toward a More Detailed Workflow Analysis in Radiology. *J Digit Imaging* 29: 47–58, 2016
- Li X, Xue Y, Malin B: Towards understanding the usage pattern of web-based electronic medical record systems. *World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. 2011 I.E. Int Symp on a 1–7, 2011
- Jameson A: Adaptive interfaces and agents. In: *Human-Computer Interaction: Design Issues, Solutions, and Applications*. 2009
- Reiner B: Automating radiologist workflow, part 2: hands-free navigation. *J Am Coll Radiol* 5:1137–1141, 2008
- Zheng K, Padman R, Johnson MP, Diamond HS: An interface-driven analysis of user interactions with an electronic health records system. *J Am Med Inform Assoc* 16:228–237, 2009
- Gregg Jr, WB, Randolph M, Brown DH, Lyles T, Smith SD, D'Agostino H: Using PACS audit data for process improvement. *J Digit Imaging* 23:674–680, 2010
- van der Aalst WM: Business alignment: using process mining as a tool for Delta analysis and conformance testing. *Requir Eng* 10: 198–211, 2005
- Dijkman R, Dumas M, van Dongen B, Käärrik R, Mendling J: Similarity of business process models: Metrics and evaluation. *Inf Syst* 36:498–516, 2011
- van der Aalst WM: Process cubes: Slicing, dicing, rolling up and drilling down event data for process mining. In *Asia Pacific Business Process Management*. 1–22, 2013