

Understanding Users' Behavior with Software Operation Data Mining

Stella Pachidi ^a, Marco Spruit ^a, Inge Van de Weerd ^b

^a Department of Information and Computing Sciences, University of Utrecht
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands

^b Department of Information and Logistics, VU University Amsterdam,
De Boelelaan 1105 1081HV Amsterdam, Netherlands

E-mail addresses: s.pachidi@vu.nl, m.r.spruit@uu.nl, i.vande.weerd@vu.nl

Corresponding author: Stella Pachidi

Abstract. Software usage concerns knowledge about how end-users use the software in the field, and how the software itself responds to their actions. In this paper, we present the Usage Mining Method to guide the analysis of data collected during software operation, in order to extract knowledge about how a software product is used by the end-users. Our method suggests three analysis tasks which employ data mining techniques for extracting usage knowledge from software operation data: users profiling, clickstream analysis and classification analysis. The Usage Mining method was evaluated through a prototype that was executed in the case of Exact Online, the main online financial management application in the Netherlands. The evaluation confirmed the supportive role of the Usage Mining method in software product management and development processes, as well as the applicability of the suggested data mining algorithms to carry out the usage analysis tasks.

Keywords: software usage; user behavior; software operation knowledge; software analytics; log data; data mining

Understanding Users' Behavior with Software Operation Data Mining

Abstract. Software usage concerns knowledge about how end-users use the software in the field, and how the software itself responds to their actions. In this paper, we present the Usage Mining Method to guide the analysis of data collected during software operation, in order to extract knowledge about how a software product is used by the end-users. Our method suggests three analysis tasks which employ data mining techniques for extracting usage knowledge from software operation data: users profiling, clickstream analysis and classification analysis. The Usage Mining method was evaluated through a prototype that was executed in the case of Exact Online, the main online financial management application in the Netherlands. The evaluation confirmed the supportive role of the Usage Mining method in software product management and development processes, as well as the applicability of the suggested data mining algorithms to carry out the usage analysis tasks.

Keywords: software usage; user behavior; software operation knowledge; software analytics; log data; data mining

1. Introduction

Software usage concerns the utilization of a software product by the end-users. Software usage data may be collected while the end-users are using the software in the field (El-Ramly & Stroulia, 2004). Simmons (2006) points out the possibility to extract system requirements from usage, rendering the beneficial role of user experience in product innovation and differentiation. Software usage knowledge includes the awareness of how end-users use the software in the field, and how the software itself responds to their actions (van der Schuur, Jansen & Brinkkemper, 2010).

By tracking software usage, we can monitor which applications are most often used, which features are underutilized, and which functionalities could be expanded (Junco, 2013). This information could for example be used to highlight changes in the requirements engineering process. We may also gain insights on how users browse themselves through the user interface in order to perform an operation, with the goal to improve software usability or to reengineer processes. Furthermore, by observing the usage behavior of different customer profiles, the software vendor can implement more directed marketing or customized licensing (Germanakos et al., 2008; van der Schuur et al., 2010). Improved customer satisfaction, and consequently customer retention and increase in sales, are some of the business advantages that could be gained through an automated usage analysis, based on real execution data.

Software usage knowledge may be extracted from *software operation data*, i.e. data that are collected during software operation in the field (van der Schuur et al., 2010). A noticeable amount of research has already been performed in the process of recording software operation data (Bowring, Orso & Harrold, 2002; Nusayr & Cook, 2009). In practice, most vendors tend to handle the acquired data manually, or use general statistics and simple visualization techniques (Kristjansson & Van der Schuur, 2009). However, such analysis cannot yield interesting patterns that are hidden in large datasets (Kantardzic, 2002).

1
2 On the other hand, a lot of development has been seen in the web usage mining field
3 (Cooley, Mobasher & Srivastava, 1997). Although many lessons can be learned from
4 there, the approach for analyzing web usage by website visitors has significant
5 differences, compared to analyzing how software products are used by the users. The
6 techniques that are used in web usage mining (and other related domains) need to be
7 revised for their application in mining usage on software operation data.
8
9

10 While usage knowledge is highly important for making good software products, the rise
11 of cloud computing and Software-as-a-Service (SaaS) applications (Park & Ryoo, 2013)
12 creates an opportunity to mine the easily acquired data. Even though there are algorithms
13 for doing such data analysis, they are hardly ever used for analyzing software usage.
14 Following a meta-algorithmic approach, we will try to answer the research question:
15

16 *How should we inspect software operation data, in order to gain knowledge about*
17 *how the software is used by the end-users?*
18

19 This research suggests how data mining techniques can be integrated to analyze software
20 operation data in a uniform and automated way. Hence, it contributes to the domain of
21 software usage analysis as well as to the software operation knowledge and its use in
22 software product management, development and maintenance processes (Van der Schuur
23 et al. 2010). From a practical perspective, the method that we suggest for usage mining
24 constitutes a reference process model that can be followed by software vendors, to
25 analyze how their customers use their products.
26
27

28 The remainder of this paper has the following structure: In section 2 we review the
29 research that has been performed on the area of extracting usage knowledge from the
30 system utilization. We shortly present our research design in section 3. In section 4 we
31 present the method that has been constructed to extract usage knowledge. In section 5 we
32 describe the usage knowledge subjects that we suggest to extract, and the variables that
33 should be inspected in software operation data, in order to derive conclusions about how
34 software operates in the field. Section 6 describes the data mining techniques that are
35 suggested for mining software usage knowledge. In section 7 we present the prototype
36 that was constructed as an instantiation of the usage mining method. We evaluate the
37 two artifacts in a case study in section 8. Finally, in section 9 we discuss the insights
38 from this research and provide some general conclusions.
39
40
41
42
43
44
45

46 **2. Related Work**

47

48 As far as specific research on software usage analysis is concerned, extraction of in-the-
49 field usage knowledge remains an area that needs a lot of enrichment. Data analysis
50 techniques have been previously applied to this field: for software reengineering
51 purposes (El-Ramly, Stroulia & Samir, 2009; Lefngwell & Widrig, 2003), for program
52 comprehension (Zaidman et al., 2005), for re-documentation of use cases (Smit, Stroulia
53 & Wong, 2008), or for user interface learning agents (Ruvini & Dony, 2001). However,
54 these approaches are not directly related to analyzing how the end-users are utilizing the
55 software in the field. Also, they do not provide any holistic approach to the various usage
56 knowledge types (e.g. user profiles or most frequent navigation paths). Some of them are
57 very old, so they do not use state of the art data mining techniques.
58
59

60 Several techniques have been developed for deriving models based on analysis of log
61
62
63
64
65

1
2 files (Petruch, Tamm & Stantchev, 2012). For example, analyzing the audit trails
3 through sequence analysis techniques can prove to be quite useful for evaluators who are
4 curious to compare the designers' expectations of use with the actual usage patterns
5 followed by the users (Judd & Kennedy, 2004). This approach is similar to the field of
6 Process Mining (Van der Aalst & Weijters, 2004), which involves analysis of event logs
7 with the goal to monitor and/or redesign operational business processes that take place in
8 an information system (Maruster & van Beest, 2009). Process Mining has also been
9 applied on web services workflows mining (Dustdar & Gombotz, 2007).

10
11
12 Practical examples that include usage analysis of logged events can be found in literature
13 (Lin & Tsai, 2011). Transaction logs analysis techniques are used in the usage analysis
14 of a digital library (Jones, Cunningham & McNab, 1998). Shen, Fitzhenry and Dietterich
15 (2009) use a subgraph mining algorithm, in order to partially automate the user's
16 workflows or to create to-do lists, in a desktop assistant application. Sartipi and Safyallah
17 (2009) developed a data mining algorithm for sequential pattern discovery on traces that
18 are generated from the execution of task scenarios.

19
20
21 A closely related area to software usage analysis is web usage mining, one of the subfields
22 in web mining. Web usage mining is the process of automatically discovering and
23 analyzing behavioral patterns and users profiles in clickstream and other associated data,
24 which are generated or collected when users interact with web resources found on one or
25 more websites (Liu, 2006). The most common pattern discovery and analysis tasks
26 include: session and visitor analysis (Liu, 2006), visitor segmentation and profiling (Xie
27 & Phoha, 2001), association analysis (Meo, et al., 2006), navigation analysis or path
28 analysis (Cooley, Mobasher & Srivastava, 1999), and prediction based on web user
29 transactions (Liu, 2006).

30 31 32 33 34 35 **3. Research Design**

36
37 The users' shift to cloud computing applications (Park & Ryoo, 2013) creates the
38 opportunity for software vendors to automatically collect vast amounts of usage data.
39 Although several algorithms have been developed to analyze the behavior of website
40 visitors, they are hardly ever used in the software products domain. This research aims to
41 follow a meta-algorithmic approach, by incorporating the state-of-the-art data mining
42 techniques in a method. Our goal is to show how the appropriate technique can be used for
43 analyzing each aspect of software users' behavior.

44
45 In figure 1 we display a diagram of our research design. We follow the *design science*
46 *research (DSR)* approach (Hevner et al., 2004), as we develop a method and a prototype
47 for software usage mining. We follow the General Design Cycle (Vaishnavi & Kuechler,
48 2007), which includes the phases: problem awareness, suggestion of a tentative design,
49 development of the artifact, evaluation, and conclusion.

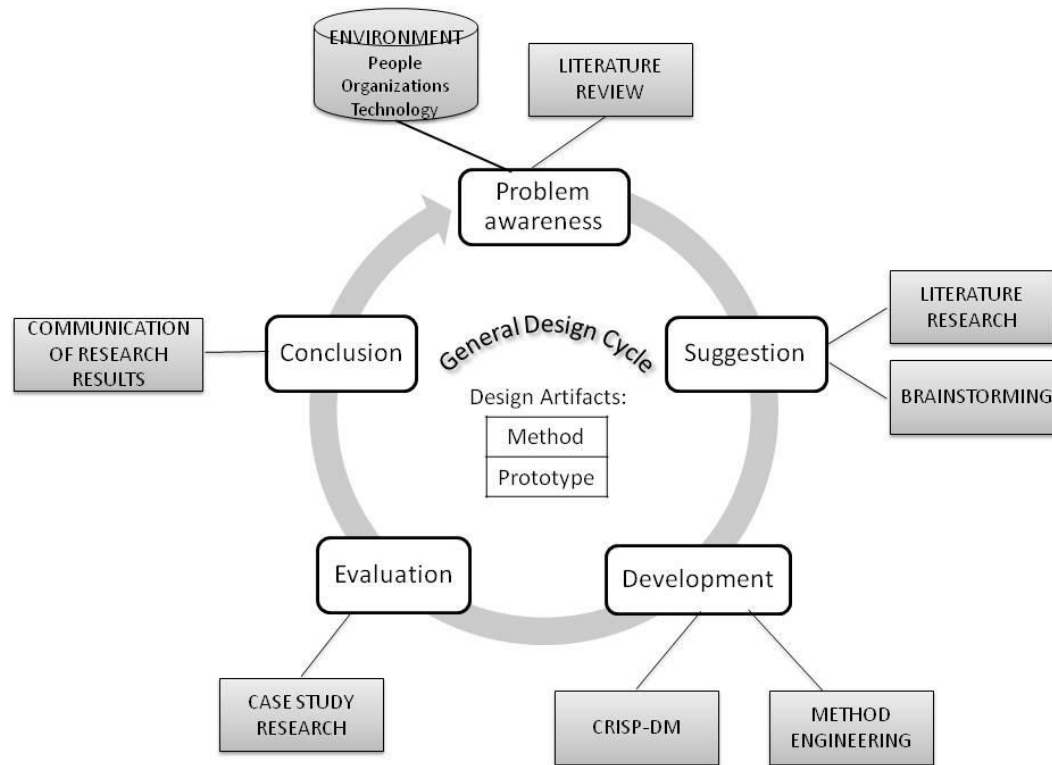


Fig. 1. Diagram of our research design

We construct a method for usage mining using the Method Engineering approach provided by van de Weerd and Brinkkemper (2008). To evaluate the effectiveness and applicability of the method, we perform a case study in an international software company, for analyzing the usage of an online financial application by trial customers.

In order to structure our data mining research, but also to assemble our Usage Mining Method, we follow the CRISP-DM Reference Model (Chapman et al., 2000), which includes six phases of data mining activities: business understanding, data understanding, modeling, evaluation and deployment. In section 4 we show how these activities were incorporated in the method.

For the evaluation of this research we use Case Study Research (Runeson & Höst, 2009) and follow a positivist approach. The case study takes place in the context of Exact Online, an online financial management application, and consists of four phases: (1) design and preparation, (2) conducting, (3) analyzing and (4) reporting. In section 8 we describe these phases in detail.

4. Usage Mining Method

In this section we present the first design artifact that we constructed in this research. The Usage Mining Method suggests an ordered set of activities that should be followed to extract relevant usage knowledge from software operation data.

In order to provide guidance in analyzing software product users' usage behavior, we propose the Usage Mining Method (figure 2). The method has been constructed with the Method Engineering approach, provided by van de Weerd and Brinkkemper (2008). The method is designed for mining usage, user and corporate data of software-as-a-service

1
2 applications, which are collected in a central point on the software vendor's side.
3

4 Figure 2 includes an overview of the method's activities and sub-activities. The activities
5 that are connected through an arrow are sequential, i.e. they need to be carried out in a
6 pre-defined order, for the reason that the outcomes of the former activities are essential
7 for the execution of the latter. Activities drawn in parallel may be executed concurrently.
8

9 The method starts with the activity of *Data Understanding*, which includes the selection
10 of variables that need to be logged, the logging procedure, the description of the data
11 quantity and format, and the evaluation of data quality (in terms of completeness,
12 inconsistencies, and duplicate data). The subsequent activity is called *Data Preparation*,
13 and includes all steps related to data pre-processing (data selection, data transformation,
14 data cleaning, data construction, data integration), in order to produce the final dataset
15 that can be used for the analysis. Then we perform *Exploratory Analysis*, in which we
16 perform statistical analysis on the data set to produce general measures and figures that
17 describe the users' behavior and are used as input in the data mining tasks. The three
18 main analysis tasks correspond to the activities: *Classification Analysis*, in which we
19 build a classification tree and then evaluate the model with cross-validation; *Users*
20 *Profiling*, in which we build several clustering models by running different clustering
21 algorithms, and then we validate the results and select the clustering with optimal scores
22 on the validation measures; and *Clickstream Analysis*, in which we build Markov chains
23 and mine sequential patterns, to select interesting usage paths. The final activity of the
24 method is the *Evaluation*, which consists in evaluating the results from the activities of
25 classification analysis, users profiling and clickstream analysis, in terms of business
26 success criteria.
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

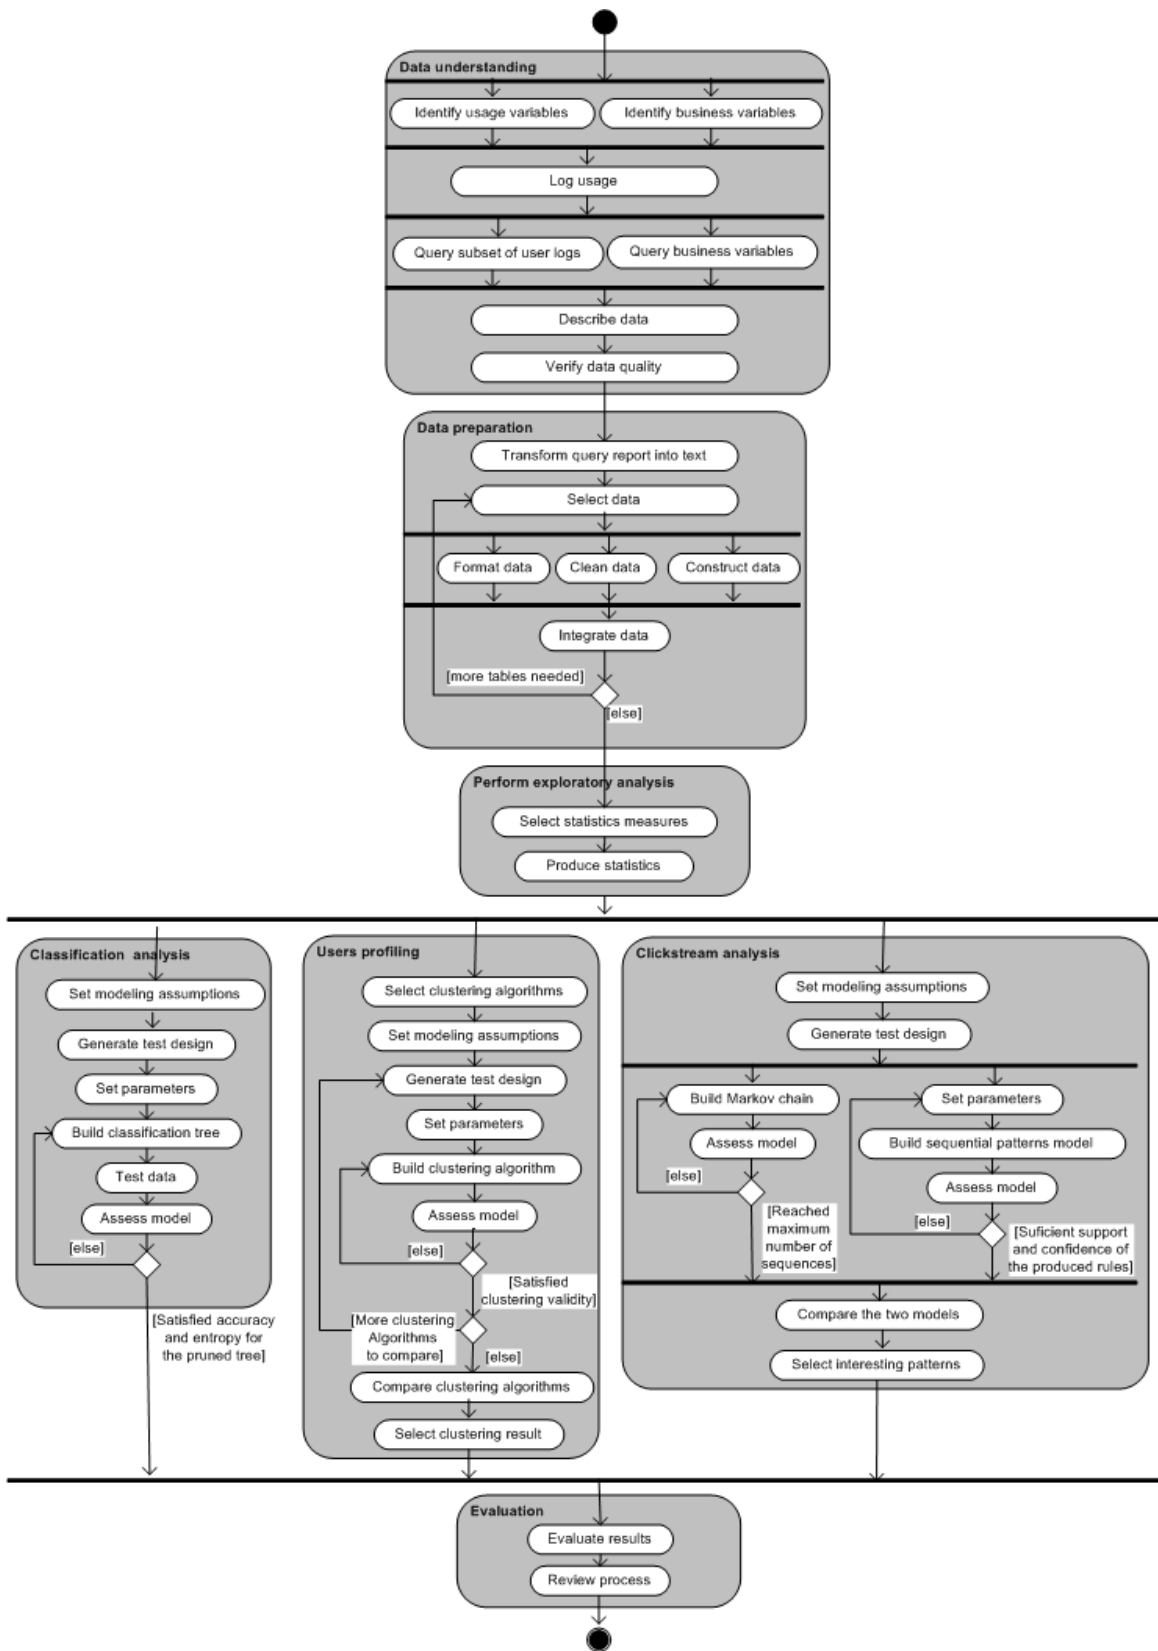


Fig. 2. Process Diagram of the Usage Mining Method

5. Software Usage Knowledge

In this section we suggest what types of knowledge should be extracted from software operation data to gain insights about how the end-users are using a software product. Subsequently, we present the fundamental variables that should be inspected during software operation, in order to gather the data that are necessary to analyze usage.

Based on our findings from our literature research in the domains of usage analysis in software systems (El-Ramly et al., 2009; Simmons, 2006) and web usage mining (Liu, 2006; Srivastava et al., 2000) as well as the insights from our case study, we suggest four categories of usage knowledge that could be derived from software operation data:

- *Statistical summary of sessions and users' behavior*: A statistical summary of the usage related data helps us infer conclusions about the general behavior of all users and their sessions. Examples: most frequently used functionalities, average session duration.

- *Factors that influence the customers' decisions*: Decisions such as renewing or upgrading the license may be predicted if we analyze the factors that influence the customers, based on how they have been using the product (Liu, 2006).

- *Users Profiles*: Extracting users profiles based on their usage behavior (Srivastava et al., 2000) consists in creating segments of users who use different sets of functionalities and visit different sets of pages.

- *The Most Frequent Navigation Paths*: Knowing the most frequent navigation paths helps us analyze the usage and usability of the software product. Furthermore, by analyzing inter-session patterns of users we can make predictions about their future actions (Lee, Podlaseck, Schonberg & Hoch, 2001).

In order to extract the aforementioned knowledge categories from software operation data, we can apply various data analysis techniques, which we study in section 6. However, it is also important to first specify what variables should be inspected in software operation data, in order to draw conclusions about usage. Based on the usage model of Simmons (2006) and the data types used in web mining research (Srivastava et al., 2000), we distinguish three *categories of variables*:

- *Usage data*: This category includes 18 operation details variables, recorded during the software usage, which describe the end-users actions on the software product. We want to log: who is using the product (customer id, user id, IP address); where the application is being hosted (web server, database); what the end user does (application, page, method, function, button that is accessed, action that is performed); when the user performs the operation (data and time, session id); how long it takes to complete an operation (duration, query duration); and other operation details (errors, background tasks, number of records loaded).

- *User data*: This category refers to variables that describe the profile of each user, such as demographic information. The 10 most common variables that are used include: age of the end-user; level of education; region, city, area of residence; the type of license that the user has purchased; the size of data records stored in the system's database for the specific user; user ratings on various objects; information on the user settings such as type of browser used when accessing the application and operating system; date that the user account was created.

- *Corporate data*: This category is applicable to corporate customers, and includes variables that provide information about the customer organization, such as: in which

domain the organization operates (e.g. construction, logistics, government, education, etc.); head-count of the organization; number of administrations that the organization consists of; number of users who use the software product; license information such as type of license and data of purchase.

In this section different knowledge types related to software usage and different variables that may be inspected in software operation data were presented. In the next section we will look into the data mining techniques that may be used for the analysis.

6. Usage Mining Tasks and Data Mining Techniques

In this section, we are going to suggest which data mining techniques could be performed on the software operation data, in order to analyze the software usage. More specifically, in order to produce the various usage knowledge types presented in the previous section, we suggest the following usage analysis tasks:

1. *Classification Analysis*, to understand the factors which influence the decisions that customers take, in the context of the software product utilization.
2. *Users Profiling*, i.e. segmentation of the user records (sessions or transactions) or the users, in order to extract profiles, that describe different navigation behaviors, or different groups of users with similar interests, respectively.
3. *Clickstream Analysis*, to extract different usage scenarios, represented by the most frequently followed navigation paths.

In the following sub-sections we will present the main data mining techniques that can be used for each of the remaining analysis tasks. More emphasis is given on the techniques that we implemented in our prototype. These techniques were selected having as criteria their implementation in R (advising the available packages in R library as of July 2011), their ease of use (by experimenting with the algorithms in the prototype construction and by communicating with data mining experts from the R-help mailing list), and their prior use in related work (based on the literature review in the domains described in section 2). Table 1 enlists a rating of the presented techniques based on these criteria.

Analysis task	Data mining technique	Implementation possibilities in R	Ease of use	Prior use in related work
Classification Analysis	Logistic Regression Models	***	***	*
	Classification Tree Models	***	**	***
	Multilayer Perceptron Models	*	*	*
Users Profiling	Cluster Analysis	***	***	***
	Kohonen Maps	**	**	*
Clickstream Analysis	Sequential Pattern Mining	*	**	***
	Probabilistic Expert Systems	*	*	*
	Markov Chains	**	***	***

Table 1. Data Mining Techniques for Usage Analysis

6.1. Classification Analysis

This task is performed with the further purpose to improve the conversion rates (Lee et al., 2001) related to the context of the software product utilization. We aim to understand the factors that influence the decisions that customers take when they use a software product, such as whether they will convert from the trial to the regular version of the product or whether they will update their license. By extracting these factors, we can predict their next decisions, and gather feedback to improve the product design and customization, or the placement of advertisements in a web application.

This task requires all types of variables that we presented in section 5. User data and corporate data are directly related to each customer who uses the software product. However, as far as the usage data are concerned, *aggregate variables* need to be created, to describe the usage data, so that in the end we have one record per customer. These aggregate variables are produced during the exploratory analysis task (Giudici, 2003), and they may be for example the frequency between logins, the average session length of the user, or the total time spent on utilizing the product.

Three different techniques can be applied in this task:

- *Logistic Regression Models*, through which we can predict the expected outcome value for an object, based on the related scores on the predictor variables (Field, 2009). A *binary logistic regression* model is suggested, in which the response variable describes the outcome of the customer's decision, e.g. whether the customer will convert or not. The set of predictor variables consists of the several usage, user and corporate attributes.
- *Classification Tree Models*, a popular data mining technique for predicting the outcome class of an object which offers easy interpretation, automatic selection of the relevant attributes, and the ability to handle both numeric and categorical attributes (Breiman, Friedman, Olshen & Stone, 1984). The structure of a *classification tree* (also referred to as decision tree) looks like a flowchart, in which each internal node (i.e. non-leaf node) represents a test on an attribute, each branch denotes a test outcome, and each leaf-node corresponds to a class label (Han, 2005). To construct a classification tree (Fayyad & Irani, 1992), we use historical data that have pre-defined classes and other attributes. The tree building algorithm calculates *tests/splits* which predict the available objects' classes with the best accuracy. In the case of analyzing the factors that influence the users' decisions, the classes represent the outcomes of the decision (yes/no or 1/0). We get to have as a sample the total number of customers, with attributes the aggregate usage variables, the user and the corporate variables. From this set of cases, we could construct (and cross validate) a classification tree that best describes our data. We can then study the splitting variables, which indicate the factors that are related to the decision of the user. But also, we can use this tree as future reference to predict the outcome of decision of other users.
- *Multilayer Perceptron Models*, which can be used for the task of credit scoring (Giudici, 2003). A multilayer perceptron (MLP) is an artificial neural network model and is represented by a directed graph, which comprises multiple interconnected layers of nodes (Haykin, 1998). In our case we can choose a perceptron model to classify the customers into two groups (e.g. decide to convert and not convert) according to the values of the usage, user and corporate variables, and then make predictions.

6.2. Users Profiling

In the task of users profiling, we wish to create profiles of the users, based on their

1 navigation behavior, which group the users according to their similar interests. The
2 retrieval of users profiles can increase the marketing intelligence and help attract new
3 customers more effectively, help the vendors offer solutions that are effectively targeted
4 to each user, and in the end retain the old customers in a more adept way. Another
5 interesting advantage that the navigation profiles can give us is the observation of a user's
6 navigation behavior over time, thus we can observe the evolution of individual customers
7 but also the evolution of the type of customers who use the software product over time.
8

9 In order to perform this task, we need to have usage data, i.e. user records which
10 represent the operations that users perform on the software product, and have as
11 attributes the operation details (such as which application was accessed, when, by which
12 user, etc.). In order to prepare the dataset for this task, we need to perform exploratory
13 analysis on these data, and generate aggregate variables, which describe the navigational
14 behavior of each user. The goal is to create a matrix, containing one record per user,
15 with aggregate attributes that describe how often, or for how long, the specific user used
16 each function of the product. When a product has many functions, it is advised to
17 organize them in homogeneous categories, which reflect their logical meaning and
18 functionality (Giudici, 2003).
19
20

21 We suggest two techniques for performing this task:
22

- 23 • *Cluster Analysis: Clustering* is the process of organizing objects into a number of
24 groups, in such a way that the objects within the same group are similar with each
25 other and are dissimilar with objects that are members of other groups (Kaufman &
26 Rousseeuw, 2008). Each object may be described by a *data point* in a
27 multidimensional space, in which the dimensions correspond to the variables that
28 describe the data (Jain, Murty & Flynn, 1999). A cluster is viewed as an aggregation
29 of points that are positioned closely to each other. The main goal in clustering is to
30 minimize the within-groups distances and maximize the between-groups distances
31 (Han, 2005).
32

33 In the context of usage analysis, our data set is a matrix with m rows that correspond
34 to the different users and n columns that correspond to the categories of the
35 product's pages/functions etc. Each cell $[i, j]$ in the matrix represents the frequency
36 of uses or the total time of using the pages of the category j by the user i . Then we
37 have an n -dimensional space, in which each dimension corresponds to a category of
38 pages and each user is represented by a data point. Thus, the objective of clustering
39 is to group together the users who have similar viewing with each other and
40 dissimilar with users from other clusters. A simplistic example drawn from our case
41 study is reflected in figure 3. Each cluster represents a usage profile, which could be
42 represented by the centre of the corresponding cluster; the centre indicates the
43 navigational behavior of the users who share the same profile.
44

45 Several clustering techniques have been developed in the data mining domain. We have
46 tested and suggest hierarchical clustering (Tan, Steinbach & Kumar, 2005) and
47 specifically Ward's method (Hastie, Tibshirani & Friedman, 2009), as well as
48 partitional clustering (Jain et al., 1999) and specifically: the K-Means algorithm
49 (Everitt, Landau & Leese, 2001), the K-Medoids algorithm (Everitt et al., 2001) and
50 fuzzy clustering (Xie & Beni, 1991).
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Fig. 3. Simplistic example of clustering the users in an online accounting solution

- *Kohonen Maps: Self-organizing maps (SOMs) or Kohonen networks* (Kohonen, 2001) are a typical implementation of *unsupervised neural networks* that can be used in the context of descriptive data mining, when we want to cluster a set of observations into homogeneous groups. A Kohonen network has one input layer and one output layer. Our objective is to divide n observations (each of which is a p -dimensional vector that contains numeric and categorical attributes) to a pre-defined number k of groups (clusters). The multivariate n observations correspond to n input neurons of the network, while the k output nodes are described by discrete values $1, \dots, k$ which indicate the different groups. The objective of a Kohonen map is to map each p -dimensional input observation to an output space, which is represented by the spatial grid of output neurons (Kohonen, 2001). The winning output for each input observation is the closest cluster based on a predetermined distance function. However, the assignment is done in such a way that the neighboring relations between the clusters are preserved to some degree.

Going back to the usage analysis, our data set consists of n observations that correspond to the different users, and each observation is a p -dimensional vector with the frequencies that the corresponding user has visited the pages of the p categories. The n user records will be the n input neurons in the Kohonen Network (Giudici, 2003). We need to select the desired number k of usage profiles that we want to extract. We also need to select the number of rows a and the number of columns b in the grid space that will characterize the map, so that $axb = k$. Then the objective of the Kohonen map is to map each of the n user records to one of the k output neurons, which are positioned in the axb map. In order to understand the cluster configurations, we can inspect the center of each cluster, which will be indicative of the navigation behavior of the users in the same cluster.

6.3. Clickstream Analysis

While the end-users are using a SaaS application, the logging procedure keeps track of all their actions, which are centrally stored in a log file on a server (Stieger & Reips, 2010). We can therefore capture the navigational paths that each end-user follows when using the product. The analysis of these navigational paths is often called *clickstream analysis* (Giudici, 2003), and can be used to extract different usage scenarios, analyze the usability of the software, and predict the next actions of the user.

The dataset consists of usage data, which describe the users' actions. More specifically, it is a long matrix in which each row corresponds to a page-view (i.e. any kind of user

1 action such as the view of a certain functionality) and each column corresponds to a
2 specific usage variable (timestamp of the action, user id, the page viewed). The data need
3 to be organized in *sessions*, each of which describes the succession of pages viewed by a
4 user during a limited time period (e.g. from the moment he opens until he turns off the
5 application or from the moment he logs in until he logs out). The simplest form of
6 clickstream analysis is to remove the user information after deriving the sessions, and
7 hence try to find similarities between sessions. The more complex form is to organize the
8 data into sessions per user, and look for common patterns, not only between page views,
9 but also between successive logins.

10
11 We are suggesting three techniques for analyzing navigational patterns:

- 12 • *Sequential Pattern Mining*: Sequences of events that describe the behavior and actions
13 of users or systems can be generated often in several domains. Sequential pattern
14 mining is a technique based on frequent pattern mining, but here the rules that we try
15 to mine have the format $A \rightarrow B$, meaning that if episode A occurs, then episode B is
16 likely to occur subsequently.

17 As far as the usage data are concerned, the transactions represent user sessions, i.e.
18 trails of the users' actions. A simple example of a rule that we could extract could be
19 $catalogue \rightarrow product$ which means that when the user views the "catalogue" he will then
20 view the "product", or $\{catalogue \rightarrow product\} \rightarrow addcart$ which means that if the user
21 views the "catalogue" and then the "product" page, he will likely view the "addcart"
22 page next.

- 23 • *Probabilistic Expert Systems*: Probabilistic expert systems are graphical networks that
24 try to process the encoded knowledge from a knowledge base, in order to model the
25 uncertainty and decisions in large complex domains (Cowell, Dawid, Lauritzen &
26 Spiegelhalter, 2007). The strategy of probabilistic expert systems to build up a global
27 statistical model is based on subsequent local factorizations. Here, a rule $A \rightarrow B$ means
28 that the page B will be visited only if the page A has been visited. If the support of this
29 rule is higher than a predefined threshold, it is established as a valid rule.

- 30 • *Markov Chains*: A *Markov chain* is a mathematical system that goes through
31 transitions from one state to another based on probabilities in a chainlike manner
32 (Grinstead & Snell, 2006). We call as *order* of the Markov model the number of prior
33 events that contribute to the prediction of a future event. In the simple case of a first-
34 order Markov chain, what happens at time t depends only on the event that happened at
35 time $t-1$ (Giudici, 2003). Markov models of higher order predict with higher accuracy,
36 but result in lower coverage (recall) and very high computational complexity (Liu,
37 2006).

38 The dataset that is considered for the clickstream analysis consists of rows, which
39 represent the actions of users, roughly called pageviews, with columns the details of
40 the operation case (timestamp, page viewed, action, user id, etc.), organized in sessions,
41 whereas each session constitutes a sequence of pageviews. We could consider each
42 pageview as a possible *state* and create Markov chains to model the usage
43 (navigational) behavior of the end-users. We can then produce navigational paths,
44 which are represented by paths that connect nodes through the most likely transitions.
45 Also, we can obtain useful information such as the likelihood of moving from one page
46 to another; or the most frequent entry/exit point for a session.

47 In this section, we presented how we can use data mining techniques to analyze data
48 related to the end-users behavior, actions and decisions. For the classification analysis
49 three techniques were presented: logistic regression models, classification tree models,
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 multilayer perceptron models. Two techniques were suggested for users profiling: cluster
2 analysis and Kohonen maps. Clickstream analysis may be performed through three
3 techniques: sequential pattern mining, probabilistic expert systems and Markov chains.

4 **7. A Prototype for Usage Mining**

5 The Usage Mining Method presented in section 4 is instantiated in a prototype, which we
6 developed in R (R Development Core Team, 2008) and implements the method's
7 activities. The prototype can be used to analyze the software usage of SaaS products with
8 embedded logging procedures that record the operation data. The prototype has the
9 format of an *R script*, which performs successively the activities of Data Preparation,
10 Exploratory Analysis, Classification Analysis, Users Profiling and Clickstream Analysis.
11

12 For the *classification* task, we used the *R* package *rpart*, which includes programs that
13 build regression and classification models (Therneau & Atkinson, 1997), to grow a
14 customer classification. The case of a new customer may be inserted in the classification
15 tree, and, through several tests, it may be assigned with the class label 1 (the customer
16 will convert) or 0 (the customer will not convert).
17

18 Since the result of *clustering* analysis is quite dependent to the method used and to the
19 selected number of clusters, in the prototype we implement several methods; based on a
20 set of evaluation criteria, the analyst can select the clustering result that is more
21 appropriate each time. More specifically, the prototype implements: hierarchical
22 agglomerative clustering with Ward's method (R function *hclust*); partitional clustering
23 with the K-means (R function *kmeans*) and the K-medoids algorithm (R function *pamk*);
24 fuzzy clustering with the C-Means algorithm (R function *cmeans*) and Kohonen maps
25 (functions *somgrid* and *som* of the R package *kohonen*).
26

27 For the *clickstream analysis* task, the prototype creates the probability transition matrix,
28 in which the element (i, j) denotes the probability of viewing page j after viewing page i .
29 The probability transition matrix can then be used to build first-order Markov chains that
30 model the users' navigational patterns. Other interesting conclusions may be inferred
31 from the matrix, such as the most common entry and exit points of a session.
32

33 The output of the prototype execution is exported in images and documents of CSV
34 format, in a specified output folder, so that they can be inspected at any time. In order to
35 use the prototype with the logs of a specific product, the analyst has to modify some
36 fields in the R script (guided by comments on the code), and then select the lines and
37 execute.
38

39 **8. Case Study**

40 Following the Design Science Research approach, we just presented the two design
41 artifacts that we constructed in this research: the Usage Mining Method and the prototype
42 developed in R. In order to evaluate the two artifacts, we performed a case study in a
43 Dutch software company, to implement the Usage Mining Method and run the prototype
44 in the context of a real software product. This section comprises the design of the case
45 study, as well as the execution and interpretation of the results.
46

8.1. Case Study Design

8.1.1. Case Study Objectives

The case study was performed to evaluate the Usage Mining Method and the Usage Mining Prototype, in response to our main research question: how we should inspect software operation data, in order to gain knowledge about how the software is used by the end-users.

8.1.2. Case Selection

In order to evaluate the Usage Mining Method and Prototype, we performed a case study in Exact¹, a Dutch software company that serves small to medium enterprises with information technology, by delivering business software solutions. Our case study was performed in the context of *Exact Online*, an internet-based accounting solution which constitutes one of the main software products of the company. Exact Online is a *Software-as-a-Service* application, with over 10,000 customers and more than 3,500 users per day.

Exact Online logs performance and usage data through a log-generating code that is incorporated in the system layer of the application. Every time the application is used by an end-user, several variables are stored in log tables in the administration database, on the server's side. The information that is logged contains: usage data (user id, application used, date/time, action taken, etc.), performance data (how long it took to process a query, how long it took to load an application, etc.), quality data (e.g. errors that appeared during usage) and other useful information such as which background tasks were running during usage. All logs are stored in the database of Exact Online for the period of 90 days. Currently, such log information is analyzed by manual inspection, through the use of standard statistics. However, this kind of analysis turns out to be very difficult, error-prone and time consuming. Soon, as users of Exact Online continue to increase significantly, it will even become more difficult. At the same time, manual inspection fails to identify complex patterns between these data.

Exact offers the possibility to anyone interested in trying out the functionalities of Exact Online, to register for a 30-day free trial. After registering for the trial, the customers can use the functionalities of the product and decide whether to buy a subscription for one of the packages offered.

The software operation data of Exact Online could provide us with valuable information on the first experiences of the end-users with the product, such as how they browse the pages in order to discover the functionalities. This kind of information is particularly useful: to understand how the usability of the product could be improved; to understand what factors influence users to buy a license; but also to segment customers according to their usage profiles, and help the marketing department target them more effectively (Okazaki, 2007). We applied the suggested Usage Mining Method and used our prototype to extract such kind of usage knowledge from the operation of the trial version of Exact Online. We discuss the process and results from this analysis in the following paragraphs.

8.1.3. Data Collection Procedure

Data was collected through documentation (reports, software manuals, software

¹ <http://www.exact.com/>

1 architecture specifications, memos, etc.), direct observations (to understand what the
2 developers need to know about the usage of their software), exploratory interviews (with
3 product managers, software engineers, research engineers and a data mining expert) and
4 participation in experts meetings.

5 As far as the collection of the data needed for the analysis tasks is concerned, our main
6 dataset consisted of usage data logged during the use of the trial version of Exact Online
7 over a period of four months (March-July 2010). This dataset consisted of approximately
8 440,000 rows with 12 attributes, and corresponded to 908 customers. Additional data
9 were collected regarding the accounts, software quality/errors, access to help
10 documentation, and statistics on the size of each customer account. Demographic and
11 corporate information regarding the users was not provided for reasons of protecting the
12 customers' privacy.
13
14
15
16

17 *8.1.4. Analysis Procedure*

18 The analysis of qualitative data was performed incrementally with the data collection
19 procedure. Coding and tabulation of the data helped identify what types of data would
20 need to be queried (from the logs and other data sets). The configuration of the Usage
21 Mining Prototype could then take place, as some pre-processing code had to be added to
22 serve this analysis, but also several parameters had to be arranged in the analysis code.
23
24

25 Before performing the data mining tasks, we first performed an exploratory analysis, in
26 order to get some general statistics on the customers of the trial Exact Online and their
27 usage behavior. We calculated the conversion rate for the customers who were included
28 in our logs, the average view time per page, and so on. The three main analysis tasks
29 followed: classification analysis, users profiling and clickstream analysis. The results are
30 provided in section 8.2.
31
32
33
34

35 *8.1.5. Validity*

36 Attention was paid to all aspects of validity (Runeson & Höst, 2009): To ensure construct
37 validity, we used multiple sources of evidence (triangulation): documentation, direct
38 observations, exploratory interviews. Internal validity was threatened for the
39 classification analysis task and for the users profiling task of the method, because we
40 focused on the usage variables and omitted most demographic and corporate variables,
41 as these were not provided for reasons of privacy. The results will be biased in these two
42 cases. However, the two analysis tasks would be performed in a similar way if those
43 types of variables would also be included, while their usefulness has been previously
44 discussed in literature - e.g. in the work of Srivastava et al. (2000). In order to ensure the
45 external validity, we performed an extensive literature review; had our method reviewed
46 by a peer researcher (data mining expert); and we repeated the analysis with data
47 collected during a different time period (within-case examination). By developing and
48 maintaining a detailed case study protocol and spending sufficient time with the case, we
49 ensured that the procedures used were well documented, and that they could be repeated
50 again (reliability).
51
52
53
54
55
56

57 **8.2. Case Study Results**

58 In the following sub-sections we examine the analysis tasks that were performed on the
59
60
61
62
63
64
65

usage data of the trial version of Exact Online. For each task we describe how the analysis was executed and how the data mining results were validated, while afterwards we provide an interpretation of the results.

8.2.1. Classification Analysis

Analysis Execution

As far as the classification analysis is concerned, we studied the variables that are related to the customers' decision to buy a license, after having used the trial version of Exact Online. In order to validate the classification results, we decided to split our dataset in a set of training data (650 user records) which will be used for growing the classification tree and a set of test data (258 records) the classes of which we will predict with the fitted model(s). We selected which variables would be included in the data set and ran the classification tree building function.

Validation of the data mining results

The resulting classification tree was pruned (figure 4) to avoid overfitting and contained two factors: number of pageviews, and total size of the stored records. Classifying the 258 test observations with the resulting tree was accurate for 85.7% of the cases.

Interpretation

The number of page views denotes the amount of activity of the end-users in the application. Naturally, the users who decided to convert seem to have spent sufficient time using the trial version. On the other hand, it is somewhat surprising that higher size of a customer's administration data, leads to a decision not to convert. Perhaps there were delays in the loading of these records, which demotivated the users.

Through the classification tree analysis, the stakeholders of Exact Online found some interesting insights on how usage was related to the customer's decision to convert or not. However, we would expect the results to have been more insightful if demographic and corporate data had also been included in the variables of the classification tree.




Fig. 4. The pruned classification tree, relating the number of page views and the size of the user's stored transactions to the conversion decision

8.2.2. Users Profiling

Analysis Execution

In the users profiling task we grouped the users into clusters based on their navigational behavior in the trial version. We used all the aforementioned clustering algorithms. A test design was generated to run the clustering algorithm multiple times before selecting the correct number of clusters and stopping criterion. Then the clustering models were built.

Validation of the data mining results

Five clusters were extracted with the k-means, two clusters with the k-medoids, five clusters with the fuzzy c-means, and ten clusters with the Kohonen maps procedure. In order to validate the clustering results, we used the internal validation measures of connectivity, Silhouette width and Dunn index (Brock et al., 2008). The measurements can be viewed in table 2. The k-medoids algorithm gave the optimal scores (lowest connectivity, highest Dunn index and Silhouette width) by giving two clusters of users: the ones with low usage frequency the ones with high usage in some categories.

	Hierarchical clust.	K-means	K-medoids	Fuzzy clust.	Kohonen maps
Number of clusters	5	5	2	5	10
Connectivity	182.19	85.04	49.22	317.83	151.71
Dunn Index	0.003	0.025	0.028	0.001	0.021
Silhouette width	0.426	0.706	0.811	0.166	0.669

Table 2. Validation of the clustering results for the users of trial Exact Online

Interpretation

The centers of the five clusters from the k-means algorithm can be viewed in figure 5, indicating the amount of usage in each category of pages. In all partitions shown in table 2 we noticed the existence of a cluster (cluster 5 in figure 5) that contains most of the observations, indicating the profile of users who had very low usage. This is not unexpected, since we clustered users of the trial version, and the majority of them had used it very few times. Analyzing the customers who have bought the commercial version and use it regularly, would have provably given more distinct cluster profiles.

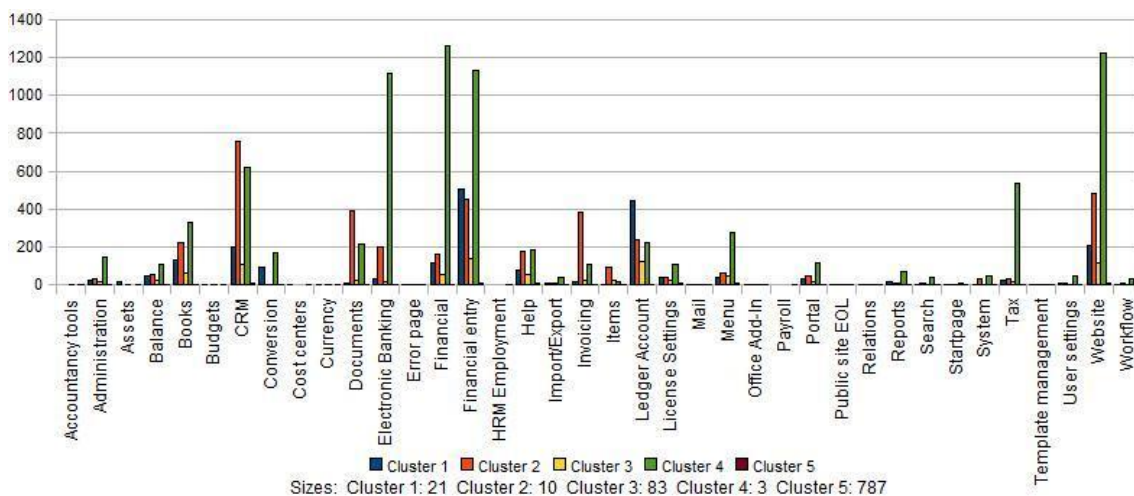


Fig. 5. Cluster centers for the k-means algorithm

8.2.3. Clickstream Analysis

Analysis Execution

We performed Markov chains analysis to find common navigational patterns. A test design was generated, which instructed the creation of 10 different possible Markov chains and one chain with the most probable transitions, as well as the extraction of other interesting clickstream data, such as the most likely entry and exit points and the most likely transitions to/from a set of specific pages. An example of how a Markov chain could be visualized is exposed in figure 6.

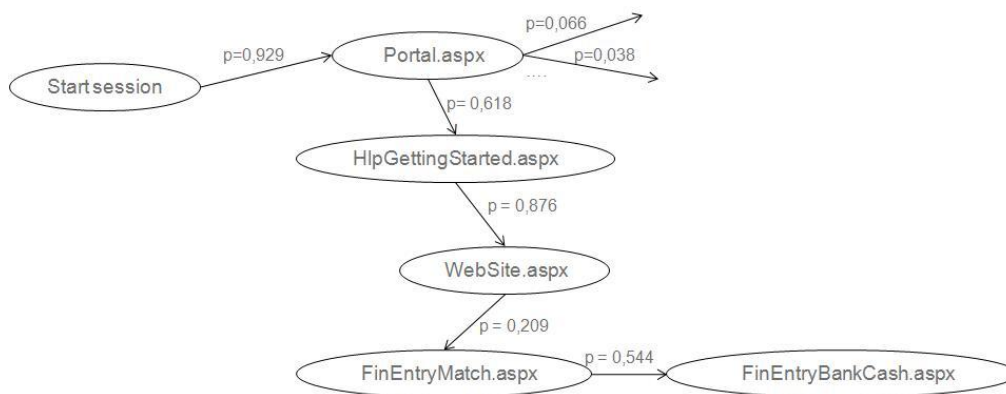


Fig. 6. Visualization of a Markov Chain

Interpretation

From the Markov chains, the developers of Exact Online could gain insights about the navigational patterns that the users follow in the trial version and thus analyze the first experiences of the users, how they browsed themselves through the application as well as where they mostly decided to exit the application. Usability could also be tested by identifying how many clicks are usually performed by the users to execute some functionality and also by analyzing how easily the users get to learn how to use the application.

This section has presented how we applied our suggested method and prototype for Usage Mining in the case of Exact Online. More specifically, we analyzed the usage of the trial version of Exact Online for the period of four months. We customized our prototype code in order to process the logs received from Exact Online and provided the stakeholders with the codes and documentation to perform the analysis in the future.

9. Discussion

In this paper we have investigated *how we can inspect software operation data, in order to gain knowledge about how the software is used by the end-users*. We reviewed related literature on software usage analysis. We constructed and presented a method that could be used to analyze how the end-users are using a software product. We explicated this knowledge by distinguishing four different categories (statistical summaries of sessions and users' behavior, factors that influence the customers' decisions, users profiles, and the most frequent navigation paths) and we presented the variables that need

1 to be inspected through software operation data to later extract them. For this, we
2 suggested three analysis tasks (classification analysis, users profiling, and clickstream
3 analysis) and we proposed data mining techniques that can be used to perform each
4 analysis task. The method and data mining techniques were evaluated through a
5 prototype, which was developed in R, and was used in the case of the trial users of Exact
6 Online. The classification analysis task yielded factors that influence the trial customers
7 in their decision to convert, with 85.7% accuracy. In the users profiling task, five different
8 clustering algorithms were tested to create profiles of the customers based on their
9 navigation behavior. In the clickstream analysis, we built Markov Chains to create
10 possible patterns that users follow when they browse the trial version of Exact Online, as
11 well as to study the probability of using a specific feature, and to find the most common
12 entry and exit points of the application.
13
14

15 The outcomes of this research enrich the domain of software usage analysis (El-Ramly &
16 Stroulia, 2004; Simmons, 2006) and generally IS/IT use (Sun & Teng, 2012). Although
17 data analysis techniques had been previously developed e.g. for software reengineering
18 purposes (El-Ramly, Stroulia & Samir, 2009; Smit, Stroulia & Wong, 2008) they were
19 hardly ever used to analyze the behavior of end-users while utilizing the software in the
20 field (Kristjansson & Van der Schuur, 2009). Also, they did not provide any holistic
21 approach to the various usage knowledge types (user profiles, most frequent navigation
22 paths, etc.). On the other hand, the similar domain of web mining had met a lot of
23 development in the web usage mining field (Cooley, Mobasher & Srivastava, 1997).
24 Although a lot of lessons could be learned from this field, analysis of in-the-field usage
25 for software products would have some significant differences. In this research, we
26 revised data mining techniques, and we explored how we could employ them to extract
27 usage knowledge from software operation data.
28
29
30

31 The usage mining method and techniques that were suggested in our research can be used
32 by the software vendors to inspect how end-users actually use the product, which
33 functionalities they prefer ,and what kind of paths they follow to perform a related task.
34 They can also test how the end-users deal with the new features of an updated release and
35 discover usability issues. The usage profiles can help reorganize the product
36 functionalities in different packages as well as attract and retain customers more
37 effectively. Finding what features the users prefer may also help determine future
38 extensions.
39
40

41 While implementing this research, we were constrained by the availability of data mining
42 functions in R for the implementation of the suggested techniques, by R's static memory
43 allocation settings combined with the system resources that were available, and the
44 limitation of data that we had available from Exact Online, due to privacy issues. The
45 method and prototype were designed to analyze usage of Software-as-a-Service products,
46 but they could easily be extended and adjusted for other types of software products.
47
48

49 Further research could be performed on the usage mining topic: First, the method and
50 prototype should be further extended with analyzing operation data from other types of
51 products. We would also like to test the techniques that were suggested but not
52 implemented in the current prototype. Furthermore, the prototype could be evaluated by
53 users in terms of usability. In addition, the usage analysis could be extended to
54 incorporating also historical software operation information or other data (e.g., release
55 schedules or bug tracker data). Finally, we would like to see how software vendors are
56 utilizing the extracted usage knowledge.
57
58
59
60
61
62
63
64
65

References

- Bowring, J., Orso, A., & Harrold, M. J. (2002). Monitoring deployed software using software tomography. *SIGSOFT Software Engineering Notes*, 28, 2-9.
- Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984), *Classification and Regression Trees*, Wadsworth.
- Brock, G., Pihur, V., Datta, S., & Datta, S. (2008). clvalid: An r package for cluster validation. *Journal of Statistical Software*, 25(4), 1-22.
- Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., & Wirth, R. (2000). *CRISP-DM 1.0 Step-by-step data mining guide*, Technical report, The CRISP-DM Consortium.
- Cooley, R., Mobasher, B., & Srivastava, J. (1997). Web mining: Information and pattern discovery on the world wide web. *Proceedings of the 9th IEEE International Conference on Tools with Artificial Intelligence* (5-58). Los Alamitos, CA: IEEE Computer Society.
- Cooley, R., Mobasher, B., & Srivastava, J. (1999), Data preparation for mining world wide web browsing patterns. *Knowledge and Information Systems*, 1(1), 5-32.
- Cowell, R. G., Dawid, A. P., Lauritzen, S. L., & Spiegelhalter, D. J. (2007). *Probabilistic Networks and Expert Systems: Exact Computational Methods for Bayesian Networks*. (1st ed.). Springer Publishing Company.
- Dustdar, S., & Gombotz, R. (2007). Discovering web service workflows using web services interaction mining. *International Journal of Business Process Integration and Management*, 1, 256-266.
- El-Ramly, M., & Stroulia, E. (2004). Mining software usage data. *International Workshop on Mining Software Repositories in 26th International Conference on Software Engineering* (pp. 64-68).
- El-Ramly, M., Stroulia, E., & Samir, H. (2009). Legacy systems interaction reengineering. In A. Sefah, J. Vanderdonck & M. C. Desmarais (Eds.), *Human-Centered Software Engineering* (pp. 316-333). London: Springer.
- Everitt, B., Landau, S., & Leese, M. (2001). *Cluster analysis*. An Arnold Publication.
- Fayyad, U. M., & Irani, K. B. (1992). On the handling of continuous-valued attributes in decision tree generation. *Machine Learning*, 8, 87-102.
- Field, A. (2009). *Discovering Statistics Using SPSS*. (3rd ed.). SAGE Publications.
- Germanakos, P., Tsianos, N., Lekkas, Z., Mourlas, C., & Samaras, G. (2008). Capturing essential intrinsic user behaviour values for the design of comprehensive web-based personalized environments. *Computers in Human Behavior*, 24(4), 1434-1451.
- Giudici, P. (2003). *Applied Data Mining: Statistical Methods for Business and Industry*. Wiley.
- Grinstead, C. M. and Snell, L. J. (2006). *Introduction to Probability*. American Mathematical Society.

- 1 Han, J. (2005). *Data Mining: Concepts and Techniques*. San Francisco, CA: Morgan
2 Kaufmann Publishers.
- 3 Hastie, T., Tibshirani, R., & Friedman, J. (2009). *Hierarchical clustering*. New York:
4 Springer.
- 5 Haykin, S. (1998). *Neural Networks: A Comprehensive Foundation*. (2nd ed.), Upper
6 Saddle River, NJ: Prentice Hall.
- 7 Hevner, A. R., March, S. T., Jinsoo, P., & Ram, S. (2004). Design Science in
8 Information Systems Research. *MIS Quarterly*, 28(1), 75-105.
- 9 Jain, A. K., Murty, M. N., & Flynn, P. J. (1999). Data clustering: a review. *ACM*
10 *Computing Surveys*, 31(3), 264-323.
- 11 Jones, S., Cunningham, S., & McNab, R. (1998). An analysis of usage of a digital
12 library. In *Research and Advanced Technology for Digital Libraries, Vol. 1513 of*
13 *Lecture Notes in Computer Science* (pp. 518-518). Berlin/Heidelberg: Springer.
- 14 Judd, T., Kennedy, G. (2004). Basic sequence analysis techniques for use with audit trail
15 data. *Journal of Educational Multimedia and Hypermedia*, 17(3), 285-306.
- 16 Junco, R. (2013). Comparing actual and self-reported measures of Facebook use.
17 *Computers in Human Behavior*, 29(3), 626-631.
- 18 Kantardzic, M. (2002), *Data Mining: Concepts, Models, Methods and Algorithms*,
19 John Wiley & Sons, Inc., New York, NY, USA.
- 20 Kaufman, L., & Rousseeuw, P. J. (2008). *Finding groups in data: an introduction to*
21 *cluster analysis*. Hoboken, NJ: John Wiley & Sons.
- 22 Kohonen, T. (2001), *Self-Organizing Maps*. (3rd ed.). Springer.
- 23 Kristjansson, B., & Van der Schuur, H. (2009). *A Survey of Tools for Software*
24 *Operation Knowledge Acquisition*, Technical Report UU-CS-2009-028, Department of
25 Information and Computing Sciences, Utrecht University.
- 26 Lee, J., Podlaseck, M., Schonberg, E., & Hoch, R. (2001). Visualization and analysis of
27 clickstream data of online stores for understanding web merchandising. *Data Mining*
28 *and Knowledge Discovery*, 5(1), 59-84.
- 29 Lefngwell, D. and Widrig, D. (2003), *Managing Software Requirements: A Use Case*
30 *Approach*, 2 edn, Pearson Education.
- 31 Lin, C., & Tsai, C. (2011). Applying social bookmarking to collective information
32 searching (CIS): An analysis of behavioral pattern and peer interaction for co-
33 exploring quality online resources. *Computers in Human Behavior* 23(3), 1249-1257.
- 34 Liu, B. (2006). *Web Data Mining: Exploring Hyperlinks, Contents, and Usage Data*
35 *(Data-Centric Systems and Applications)*. Secaucus, NJ: Springer-Verlag New York.
- 36 Maruster, L., & van Beest, N. (2009). Redesigning business processes: a methodology
37 based on simulation and process mining techniques. *Knowledge and Information*
38 *Systems* 21(3), 267- 297.
- 39 Meo, R., Lanzi, P., Matera, M., & Esposito, R. (2006). Integrating web conceptual
40 modeling and web usage mining. In B. Mobasher, O. Nasraoui, B. Liu and B. Masand,
41 (Eds.), *Advances in Web Mining and Web Usage Analysis* (pp. 135-148). Springer Berlin
42 /Heidelberg.
- 43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 Nusayr, A., & Cook, J. (2009). AOP for the domain of runtime monitoring: breaking
2 out of the code-based model. *Proceedings of the 4th workshop on Domain-specific*
3 *aspect languages* (pp. 7-10). New York, NY: ACM.
- 4 Okazaki, S. (2007). Lessons learned from i-mode: What makes consumers click wireless
5 banner ads?. *Computers in human behavior*, 23(3), 1692-1719.
- 6
- 7 Park, S.C., & Ryoo, S.Y. (2013). An empirical investigation of end-users' switching
8 toward cloud computing: A two factor theory perspective. *Computers in Human*
9 *Behavior* 29(2013), 160-170.
- 10
- 11 Petruch, K., Tamm, G., & Stantchev, V. (2012). Deriving In-Depth Knowledge from
12 IT-Performance Data Simulations. *International Journal of Knowledge Society*
13 *Research (IJKSR)*, 3(2), 13-29. doi:10.4018/jksr.2012040102
- 14
- 15 R Development Core Team (2008), *R: A Language and Environment for Statistical*
16 *Computing*, Vienna, Austria: R Foundation for Statistical Computing.
- 17
- 18 Runeson, P., & Höst M. (2009). Guidelines for conducting and reporting case study
19 research in software engineering, *Empirical Software Engineering*, 14(2), 131-164.
- 20
- 21 Ruvini, J.D., & Dony, C. (2001). *Learning users' habits to automate repetitive tasks*
22 (pp. 271-297). San Francisco, CA: Morgan Kaufmann Publishers Inc.
- 23
- 24 Sartipi, K., & Safyallah, H. (2009). Dynamic knowledge extraction from software
25 systems using sequential pattern mining. *International Journal of Software Engineering*
26 *and Knowledge Engineering*, 20(06), 761-782.
- 27
- 28 Shen, J., Fitzhenry, E., & Dietterich, T. G. (2009). Discovering frequent work procedures
29 from resource connections. *Proceedings of the 13th International Conference on*
30 *Intelligent User Interfaces* (pp. 277-286), New York: ACM.
- 31
- 32 Simmons, E. (2006). The usage model: Describing product usage during design and
33 development. *IEEE Software*, 23(3), 34-41.
- 34
- 35 Smit, M., Stroulia, E., & Wong, K. (2008). Use case redocumentation from gui event
36 traces. *Proceedings of the 12th European Conference on Software Maintenance and*
37 *Reengineering* (pp. 263-268). Washington, DC: IEEE Computer Society.
- 38
- 39 Srivastava, J., Cooley, R., Deshpande, M., & Tan, P.-N. (2000). Web usage mining:
40 discovery and applications of usage patterns from web data. *ACM SIGKDD*
41 *Explorations Newsletter*, 1(2), 12-23.
- 42
- 43 Stieger, S., & Reips, U. D. (2010). What are participants doing while filling in an online
44 questionnaire: A paradata collection tool and an empirical study. *Computers in Human*
45 *Behavior*, 26(6), 1488-1495.
- 46
- 47 Sun, J., & Teng, J. T. (2012). Information Systems Use: Construct conceptualization and
48 scale development. *Computers in Human Behavior*.
- 49
- 50 Tan, P.N., Steinbach, M., & Kumar, V. (2005). *Introduction to Data Mining*. (1st ed.).
51 Boston, MA: Addison-Wesley Longman Publishing.
- 52
- 53 Therneau, T. M., & Atkinson, E. J. (1997). An introduction to recursive partitioning
54 using the RPART routines. Technical Report 61, Section of Biostatistics, Mayo Clinic,
55 Rochester. URL <http://www.mayo.edu/hsr/techrpt/61.pdf>.
- 56
- 57 Vaishnavi, V. K., & Kuechler, Jr., W. (2007). *Design Science Research Methods and*
58 *Patterns: Innovating Information and Communication Technology*. (1st ed.). Boston,
59
60
61
62
63
64
65

MA: Auerbach Publications.

1
2 Van der Aalst, W. M. P., & Weijters, A. J. M. M. (2004). Process mining: a research
3 agenda. *Computers in Industry*, 53(3), 231-244.

4
5 Van der Schuur, H., Jansen, S., & Brinkkemper, S. (2010). A Reference Framework for
6 Utilization of Software Operation Knowledge. *36th EUROMICRO Conference on*
7 *Software Engineering and Advanced Applications* (pp. 245-254). IEEE Computer
8 Society.

9
10 Van de Weerd, I., & Brinkkemper, S. (2008). *Meta-modeling for situational analysis*
11 *and design methods* (pp. 38-58). Hershey: Idea Group Publishing.

12
13 Xie, X. L., & Beni, G. (1991). A validity measure for fuzzy clustering. *Pattern Analysis*
14 *and Machine Intelligence, IEEE Transactions on*, 13(8), 841-847.

15
16 Xie, Y., & Phoha, V. V. (2001). Web user clustering from access log using belief
17 function. *Proceedings of the 1st international conference on knowledge capture* (pp. 202-
18 208). New York: ACM.

19
20 Zaidman, A., Calders, T., Demeyer, S., & Paredaens, J. (2005). Applying web mining
21 techniques to execution traces to support the program comprehension process.
22 *Proceedings of the Ninth European Conference on Software Maintenance and*
23 *Reengineering* (pp. 134-142). Washington, DC: IEEE Computer Society.

24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65