Self-Tracking Reloaded: Applying Process Mining to Personalized Health Care from Labeled Sensor Data

Timo Sztyler[1], Josep Carmona[2], Johanna Völker[1], Heiner Stuckenschmidt[1]

[1]University of Mannheim, Germany
{timo,johanna,heiner}@informatik.uni-mannheim.de
[2]Universitat Politècnica de Catalunya, Spain
jcarmona@cs.upc.edu

Abstract. Currently, there is a trend to promote personalized health care in order to prevent diseases or to have a healthier life. Using current devices such as smart-phones and smart-watches, an individual can easily record detailed data from her daily life. Yet, this data has been mainly used for *self-tracking* in order to enable personalized health care. In this paper, we provide ideas on how process mining can be used as a fine-grained evolution of traditional self-tracking. We have applied the ideas of the paper on recorded data from a set of individuals, and present conclusions and challenges.

1 Introduction

Physical inactivity is a major risk factor for certain types of diseases. Indeed, physical activity does not only prevent or relieve diseases, but also improves public health and well being [6]. In this context, personalized health solutions and lifestyle monitoring can help to ensure that individuals are doing the right activity at the right time. However, the regular use of such methods is critical to achieve the desired result. Barriers for the adoption must be low, and using both software and devices should be as comfortable as possible.

Wearable devices such as smart-phones, smart-watches, and wristbands which do not affect people during their daily routine allow to setup a body sensor network. The provided sensor technology allows to monitor people all day long. In contrast, most of the available software requires substantial user input to specify, e.g., the current activity or even vital parameters like the heart rate or blood pressure.

The goal of our work is the development of an environment that monitors and analyzes the personal lifestyle of users and the provision of insightful visualizations. In this paper we focus on deriving and analyzing personal process models through process mining [32] techniques as a central part of the system. The general goal will only be achievable if the recognition of a person's daily activities (such as different types of sports and desk work) can be automated. In this paper we assume this step is already addressed, i.e, with state of the art activity recognition techniques [18].

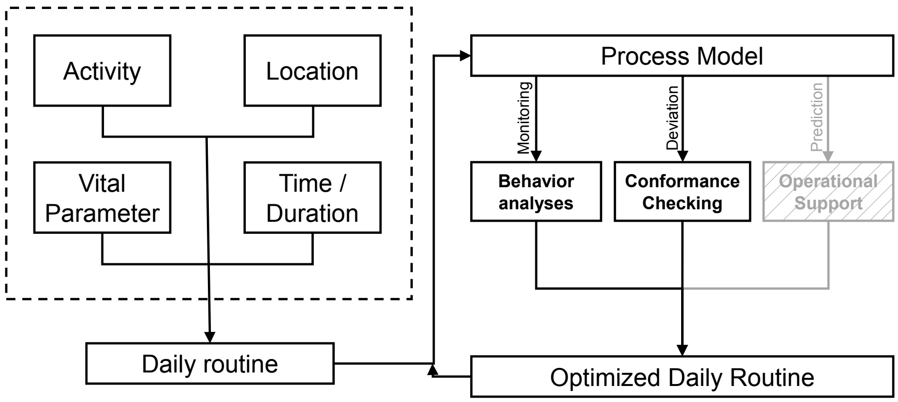

Fig. 1. Optimization of the daily routine to achieve a healthier live. The framework illustrates the interaction of the individual and considered components of the activity recognition (left) and the process mining (right) domain. The latter comprises *Behavior Analyses*, *Conformance Checking* and *Operational Support*. In this paper, we focus on the first two aspects and consider *Operational Support* as future work.

Figure 1 illustrates the main components of our framework in detail. The left part covers the activity recognition system that recognizes the performed physical activity based on data of on-body sensors. For that purpose, it collects context-related information, i.e, the current geographical location, local time, and vital parameter of the patient. Hence, this data represents a sequence of activities which can be considered as event log. The event log denotes the daily routine of an individual and can be transformed into a personal process model. Consequently, the process model enables to examine the daily routine regarding specific patterns, discrepancies, or to predict the next activity using common process mining techniques. This allows to reveal anomalous behavior and non-conformance regarding doctor's prescription. Besides, also operational support can be provided. As a result, the daily routine can be optimized by recommendations and feedback or a carer can be informed. This paper is concerned with this latter part of the problem: the mining of suitable process models from activity and location labels that have been extracted from an event log.

In the following, we consider different process formalisms when illustrating the techniques of this paper. The reason for this decision is twofold: on the one hand, we aim to present the process mining field in general terms. Thus, using the best notation for the problem which is available. On the other hand, the current situation of the process mining field enforces this decision, by not having an unified process notation that is superior in every dimension. For instance, it is well-known that *fuzzy models* are a good visualization aid, but since they lack formal semantics cannot be used for the analysis of the underlying process, for which Petri nets are better suited.

The paper is structured as follows: In Section 2, the related work concerning health care, activity recognition, and process mining is summarized. Section 3 introduces background knowledge regarding process mining that is considered in the following sections. Section 4 describes the possibilities of discovering personal processes and extracting meaningful patterns and rules. Based on this, Section 5 outlines how to analyze and compare these processes to detect deviations and optimize the behavior of the related person. Section 6 describes the experiments of the introduced ideas concerning several different data sets. Finally, Section 7 covers the future work of this paper.

2 Related Work

In this paper, we aim to explore spatio-temporal data with process mining techniques to extract knowledge that facilitates personalized health care. Patients are often required to follow a well defined exercise routine or have to be monitored as part of their treatment. Therefore, detecting wrong behavior or abnormal activities may help to prevent undesirable consequences [16, 36]. Accurate information on people's behavior and their daily routine allows to support them [6] or provide feedback to a caregiver.

The event log which describes the daily routine, and can be transformed to a process model [34], results from common activity recognition techniques. During the past decade, research on human activity recognition rely on wearable and external sensors [18] to determine, e.g., if a person is preparing food or going to work. The wearable senors are attached to the patient and are used to determine the physical activity by sensing the body movement [4]. In addition, external sensors are attached, e.g., to doors and items to recognize with which objects the patient interacts. Commonly, this is the case in a smart-home environment [30]. The result is a sequence of activities including the duration, location [22], and vital parameter [19].

We focus on wearable devices, i.e, smart-phones and smart-watches because they provide variety of sensors and are carried all day long by many people [5]. Besides, the accelerometer enables continuous sensing over a complete day due to a low power consumption.

Commonly, probabilistic approaches such as *Markov Logic Networks* or *Hidden Markov Models* are used to determine the performed activity or to predict an unobserved state, e.g., the next activity [17]. In this context, researcher also focus on pattern detection, i.e., analyzing a specific sequence of activities [17, 27] to verify given references. In contrast, process mining enables to infer and extract routines that occur during the daily routine of a patient from a hidden structure. Further, the techniques allow to perform a more analytical discussion regarding the performed healthcare process [23]. This means that the mentioned approaches do not exclude but can complement each other.

Several researchers of the processes mining area already addressed similar problems and developed techniques that are suitable for sequences of events and spatio-temporal data. Aztiria et al. showed that learning a habit is very similar to mine a process [3] and Agrawal et al. introduced algorithms that enable to mine sequential patterns that allow to identify common behavior [2, 24]. However, these approaches focus only on the performed activities where we want also to consider the location and time of day. The combination of these dimensions may lead to valuable knowledge.

In this context, trajectory pattern mining allows to consider chronologically ordered geographical locations and the duration of movements between them. This facilitates to examine movement behavior but also the relation between time, location, and activity. Thus, the techniques discover highly frequented places as well as underlying patterns which might be related to other persons due to semantic relations [21]. However, current methods do not address sparsity and noise which is an important concerning our scenario. Hence, behavior that occurs rarely may be an strong evidence for a specific diseases.

Commonly, there are discrepancies between the daily routine of a patient and the desired behavior. For that purpose, Rozinat et al. [28] developed a fitness measure to expose the distinctions between a predefined model and the real behavior. Due to the limitations of this measure, Leoni et al. enhanced this approach by considering further dimensions. In detail, they describe costs and quantities for additional event data that allows to quantify conformance and analyze differences between model and reality [12].

However, a general problem is the handling of unstructured or flexible processes, as it is the case for the daily routine of an individual. In this context, Leotta identified that the human habits are flexible in their nature and addressed this problem by considering declarative models [20]. As a result, they developed a technique that enables to perform mining on declarative models of human habits. The work of this paper can be seen as an extension of Leotta's work where other formalisms like fuzzy maps are considered and the posterior analysis of the derived process is taken into account.

Finally, related to incorporating the modeling of context information like location in the process, the work in [38] represents a promising direction. In this work Petri nets are enriched with location constraints, and the semantic is extended to cope with this new dimension. For tool support, location-aware Petri nets are mapped to colored Petri nets so that the analysis can be done in CPN Tools [15]. Hence, it can be integrated with a Geographical Information System at run-time. Unfortunately, no discovery technique for this location-aware Petri nets exists so far. A general framework to incorporate also other types of context in process models is presented in [29].

3 Preliminaries: Process Mining Techniques

In this section we provide the necessary background to understand the techniques which we consider in the following sections. We will focus on two main process mining disciplines: *process discovery* and *conformance checking*, which represent the core of process mining [32].

A log L is a finite set of traces over an alphabet A representing the footprints of the real process executions of a system S that is only (partially) visible through these runs. Process discovery techniques aim at extracting from a log L a process model M (e.g., a Petri net) with the goal to elicit the process underlying in S. By relating the behaviors of L, M and S, particular concepts can be defined [9]. A log is *incomplete* if $S \backslash L \neq \emptyset$. A model M *fits* log L if $L \subseteq B(M)$, where $B(M)$ denotes the behavior underlying M. A model is *precise* in describing a log L if $B(M) \backslash L$ is small. A model M represents a *generalization* of log L with respect to system S if some behavior in $S \backslash L$ exists in $B(M)$. Finally, a model M is *simple* when it has the minimal complexity in representing $B(M)$, i.e., the well-known *Occam's razor principle*.

Process discovery is challenging because the derived model has to be fitting, precise, general, and simple. Conformance checking techniques are meant to verify these criteria to asses the quality of a model in representing the information contained in a log. In this paper we focus on the *cost-based fitness analysis* [1] which allows to score deviations between log and model. An *optimal alignment* between a log trace and a model is a pair of traces denoting what is the best way for the log trace to be reproduced by the model. An alignment can be seen as a two-row matrix where the top row corresponds to "moves in the log" and the bottom row corresponds to "moves in the model". If a move in the model cannot be mimicked by a move in the log, or vice versa (denoted by the symbol $\gg$ in the corresponding matrix cell), then a fitness problem between the model and the log is revealed. In contrast, when log and model can execute the same activity, it denotes a fitting step. Considering an alignment, if only fitting steps appear then the trace can be reproduced by the model, otherwise a fitting problem is encountered. An example of alignment can be found below:

$$\begin{array}{|c|c|c|c|c|} a & \gg & b & d & e \\ \hline a & c & b & \gg & e \end{array}$$

The first, third, and fifth column are fitting steps while the other denote fitting problems, also called *misalignments*. If unitary costs are assigned to misalignments, while fitting steps have cost zero, the previous example will have cost 2. In general, arbitrary costs can be assigned to the different types of misalignments. Considering the example, the misalignments $(\gg, c)$ and $(d, \gg)$ might have been the costs 1 and 2, respectively, whereas the rest of fitting steps have costs of zero. This will play a crucial role in the context of this paper. Techniques for computing alignments of imperative or declarative models with respect to logs exist in the literature [1, 11].

4 The Discovery of Personal Processes

In this section we provide intuitive descriptions of what type of representations can be obtained through process discovery (4.1) and how these representations can be enhanced to incorporate the information in the context of personal process behavior (4.2).

4.1 Imperative and Declarative Representations of Personal Processes

Since a picture is worth a thousand words, the deployment of graphical representations of event data may lead to a precise awareness of the activities carried out by an individual. We believe that graphs are a strong visualization aid to understand aggregated behavior. Thus, consider this direction as the first use case for understanding personal activity data. This deviates from the typical information that is provided by current tools for self-tracking individuals. In general, such tools focus only on showing numeric correlations between the tracked variables (e.g., eating vs. sport) or the evolution of single variables (weight over the week).

An interesting information a user can get periodically (every day or week) is the personal process model that describes the main activities and their dependencies. As we have introduced in a previous section, two options exist for modeling process behavior: imperative and declarative models.

Imperative process models tend to be well-suited for simple personal behavior, i.e., behavior that only denotes a reduced number of variants. This is the case considering elderly people where the number of performed activities is reduced and also the behavior is limited. However, even if the underlying process is less structured, this model still enables to discover frequent paths of activities. In this context, *fuzzy models* [14] or *heuristic nets* [35] may be a good alternatives. Figure 2(a) illustrates an example of a fuzzy model showing the main behavior of a group of individuals during the working days. In this process model, nodes (representing the occurrence of activities) and arcs (denoting the activity ordering) are drawn in a way that frequent behavior is highlighted: the darker the background of a node (the thicker the arc), the more frequent was the related activity (arc) performed. Thus, it can be observed that particular patterns (sub-traces) like `MealPreparation` $\rightarrow$ `EatingDrinking` $\rightarrow$ `HouseWork` $\rightarrow$ `DeskWork` are dominant in this model. In previous work, we have already used these models to infer interesting conclusions on the behavior of individuals. Thus, the distinctions between working days and weekend behavior, across different type of users [31].

In contrast, declarative process models are adequate regarding flexible or unstructured behavior. Intuitively, declarative process models are denoted by a set of temporal constraints that relate pairs of activities [26]. Those constraints can be partitioned into existence, relation, negation, choice and branching templates, establishing the boundaries between observed and unobserved behavior. For the case of personal processes, declarative constraints seem to be very adequate representations, as it has been already acknowledged in recent work [20]. Figure 2(b) illustrates a declarative process model that results from the same log as Figure 2(a). Considering both models, it is remarkable that the declarative model simplifies the information in a way which emphasizes meaningful rules. Thus, the declarative model covers three types of information. First, any pair of activities in the group {`EatingDrinking`, `Movement`, `Transportation`} are in choice relation, i.e., meaning that at least one of them should be present in any trace of the log. Second, the activity `DeskWork` has an existence constraint of 2 or more. Hence, in case of a workday this activity is repeated at least twice. Finally, the relation

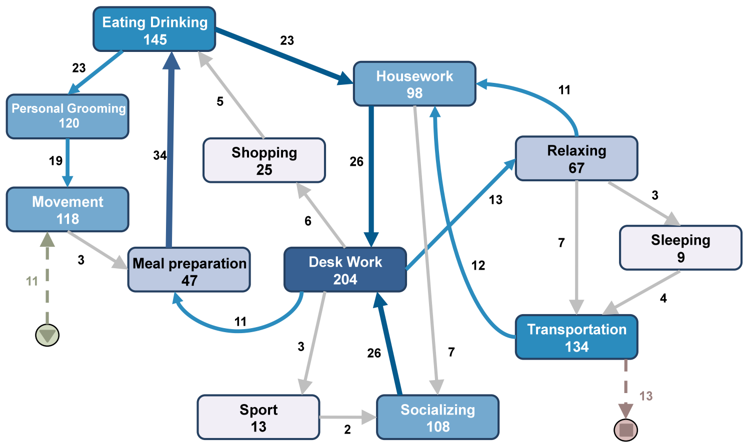

(a) Imperative Model (Fuzzy model)

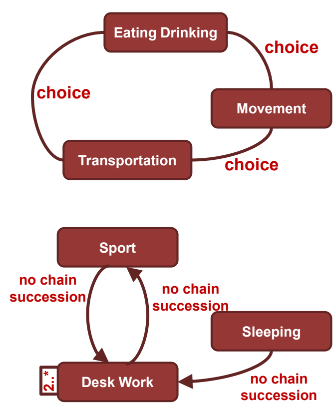

(b) Declarative Model (Declare)

Fig. 2. Main personal activities of a set of users during the week.

constraint *not chain succession* establishes nonexistence of immediate succession between activities, e.g., no trace exists where `DeskWork` directly follows `Sport`.

4.2 Model Enhancement using context-related Information

The model of a personal process can also incorporate a geographical description of the process, i.e., the locations where the activities were performed, the frequency, and the relations between them. We focus on the chronological order and the relation between the location and the duration of the activity. Thus, it has to be considered that the same activity can be performed in different locations and varying duration. This means that it is not possible to adjust easily the trajectory patterns [37]. Instead, it would be necessary to enhance the expressiveness of the trajectory patterns so that it is possible to describe relations between the spatio-temporal data and activities. As a result, the enhanced models could help to optimize the daily routine concerning a healthier life by addressing, e.g., the type of movement between locations or providing beneficial locations for certain activities.

We combine the presented process model and a geographical map to arrange the performed activities with the related context information. As an example, we

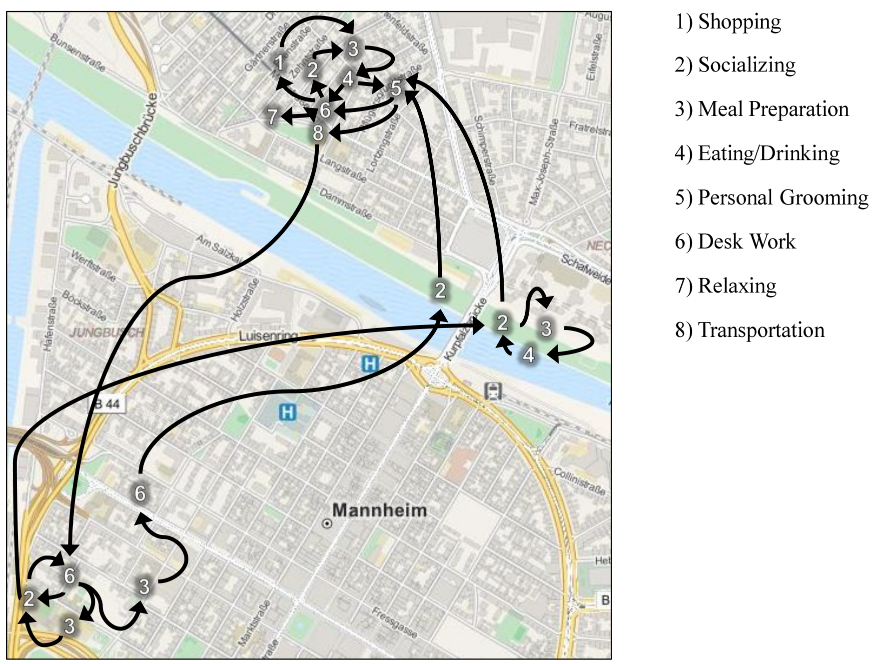

Fig. 3. Main personal activities for an individual including geographical position data: numbers correspond to different activities, and arcs denote control-flow relations extracted from the activity data.

Algorithm 1 Personal activity-position map

Input: A_H : **location-enhanced event log**
T_G : **maximum geographical distance for the same activity**
T_F : **minimum number of elements in a cluster**
Output: a personal activity-position map

$A_H \leftarrow \{\langle(act_{1,1}, lat_{1,1}, long_{1,1}, t_{1,1}), ..., (act_{1,n}, lat_{1,n}, long_{1,n}, t_{1,n})\rangle ...$
$\quad\quad \langle(act_{m,1}, lat_{m,1}, long_{m,1}, t_{m,1}), ..., (act_{m,k}, lat_{m,k}, long_{m,k}, t_{m,k})\rangle\}$ {*enhanced event log description as a set of enhanced traces*}
$C \leftarrow [\ ]$ {*empty list of clusters*}
for `each` *trace* `in` A_H **do** {*initialize a cluster for each event in each trace*}
  for `each` $(act_i, lat_i, long_i, t_i)$ `in` *trace* **do**
    $c_{new} \leftarrow$ newCluster(a_i, lat_i, $long_i$)
    add c_{new} to C
  end for
end for
$D \leftarrow (\ )$ {*geographical distance matrix of all clusters*}
for `each` $c_i, c_j \in C$: $(\forall x \in D : D(c_i, c_j) \leq x) \wedge D(c_i, c_j) < T_G$ **do** {*merge clusters that are close to each other and represent the same activity*}
  $D(c_i, c_j) \leftarrow \infty$ {*forces to inspect each pair only once*}
  if $label(c_i) = label(c_j)$ **then**
    $C \leftarrow C \setminus \{c_i, c_j\} \cup \{c_i \cup c_j\}$
    recompute *centroid* of the new cluster $\{c_i \cup c_j\}$
    $D \leftarrow$ update geographical distances matrix
  end if
end for
for `each` $c_i \in C$ **do** {*remove clusters that cover an insufficient number of elements*}
  if $|c_i| \leq T_F$ **then**
    $C \leftarrow C \setminus \{c_i\}$
  end if
end for
$L \leftarrow ProjectAndRelabel(A_H, C)$ {*an event log is obtained from A_H with the activities from C*}
$(nodes, edges) \leftarrow FuzzyMap(L)$ {*a fuzzy miner is invoked on L*}

explain how to combine the imperative control-flow process models (see Figure 2(a)) with the geographical position data to derive a *personal activity-position map*. This map illustrates geographically the control-flow with respect to the real geographical position of the activities. Compared with a trajectory-based graph, this map can be considered as a set of connected sub-graphs where each sub-graph represents the activities for a specific location.

The computation of *personal activity-position maps* can be done by aligning the timing information (*start*, *end*) of an event with the corresponding time of the related geographical position. As a result, the locations that correspond to a specific activity can be extracted and analyzed. For instance, in Figure 3, activity 2 (`Socializing`) was performed in four different locations (nodes). Ideally, to have a simpler graph, the number of locations per activity should be small. Therefore, the locations for an activity can be computed by clustering a set of

Algorithm 2 ProjectAndRelabel Method

Input: A_H : **location-enhanced event log**
  C : **set of clusters**
Output: an event log

$L \leftarrow \{\}$ $\{$*empty event log*$\}$
for `each` *trace* `in` A_H **do** $\{$*traverse the traces of the enhanced log*$\}$
  $\sigma \leftarrow$ empty trace
  for `each` $(act_i, lat_i, long_i, t_i)$ `in` *trace* **do**
    $c \leftarrow$ a cluster $x \in C$ originated from act_i and $(long_i, t_i) \in x$
    if $|c| > 0$ **then**
      $\sigma \leftarrow \sigma \cdot (label(c), t_i)$
    end if
  end for
  if $\sigma \neq \epsilon$ **then**
    $L \leftarrow L \cup \{\sigma\}$
  end if
end for
return L

geographic coordinates and considering a fixed radius of k meters. In this context, the centroids as well as the frequency of the performed activities can be used to optimize the clusters. Finally, the nodes which correspond to activities in certain locations are displayed on top of a real map. Arcs from the control-flow are then routed from the corresponding locations in the map. Algorithm 1 describes this procedure in detail.

The algorithm needs as input the introduced enhanced event log A_H as well as the threshold values T_F and T_G. The thresholds specify the maximum geographical distance between the same activity for the same location (T_G), and the minimal cluster size that has to be considered (T_F). Then, the algorithm computes a set of clusters which contain events that share the same label and that are close enough in terms of their geographical position (lines 1–23). Subsequently, the *ProjectAndRelabel* method is applied (see Algorithm 2) where an event log is extracted. In general, this method simply traverses the traces in A_H, computing a normal trace (built of events with activity name and timestamp) that results from: i) projecting only events that are covered by a cluster, and ii) relabeling the events to guarantee that different clusters originated from the same activity will be represented by different activities in the derived event log. Then, in line 25 of Algorithm 1 a fuzzy miner is invoked, which returns the corresponding fuzzy model. Note that in here any other miner could be used, since the input is now a traditional event log. Finally, the model is rendered by taking into consideration the geographical position of labels.

5 The Analysis of Personal Processes

Self-tracking is a meaningful way to verify if certain requirements with respect to reference quantities are accomplished. Concerning a healthier life, many associa-

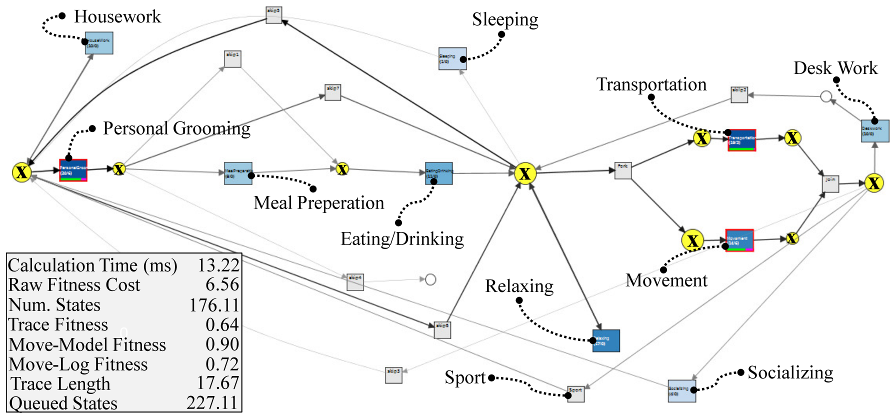

Fig. 4. Example of fitness analysis in ProM[2] of an individual with respect to a reference model (Petri net): places with yellow background (X) represent situations where the individual deviates from the process model. Transitions without a label denote silent events not appearing in the event log.

tions advise to do at least 30 minutes of moderate physical activity per day or eat fish at least twice a week. Those guidelines for a good lifestyle offer a rough description for individuals, mainly concerning about quantities and frequencies. However, some ways of satisfying these guidelines are probably less healthy than others, e.g., it may not be the best decision to eat fish while doing physical activity. Hence, a reference model that describes precisely how certain activities should be carried out in order to satisfy a guideline is required. If reference models are not available, simple rules can be used which should be satisfied by individuals on their daily routine. These rules may describe patterns that should satisfy an individual, e.g., `takingMedicines` should be followed by `EatingDrinking`. This can be formally specified with Linear Temporal Logic (LTL) formulas to be satisfied by the event log of activities [33]. Checking (temporal) rules on the event log may suffice in many situations. However, in this section we go a step further and try to use reference models for the analysis of personal processes, with the aim of providing a fine-grained analysis.

In this context, the reference model has to provide the opportunity to describe certain actions in a specific *order* (e.g., `Sport` should be followed by `PersonalGrooming`), should allow explicit *choices* (e.g., after `DeskWork` only `EatingDrinking`, `Socializing`, or `Transportation` are expected actions) and should also consider *concurrency* actions. (e.g., `Transportation` and `Movement` may be overlapping activities).

Reference models can be obtained in several ways. One possibility would be to ask a domain expert to create manually the desired reference model for a given goal. A second option would be to collect event logs from successful individuals. These logs can be combined with the introduced techniques of the previous section

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

to discover a reference model. Finally, a third option would be to translate the textual guidelines into process models, using recent techniques that apply *natural language processing* to elicit process models [13].

The resulting reference model enables to apply *conformance checking* techniques to assess the adequacy of the reference process model in representing the traces of individuals [32]. Since the reference model describes the ideal behavior, it is meaningful to focus on analysis of the fitness of the reference model with respect to the traces of individuals. As already mentioned, a process model fits a given trace if it can reproduce it. An example of such analysis can be seen in Figure 4 where an individual is analyzed with respect to an invented process model meant to represent a healthy behavior.

Obviously, for the case of cost-based fitness analysis (see Section 3) of personal processes an important part is the determination of costs for certain deviations. Hence, the selection influences the derivation of a model traces with respect to the observed behavior. In order to avoid to interfere the search for model traces that must be as much similar as possible to the observed ones, unitary cost will be assigned when computing an alignment. However, once the alignment is computed, the misalignments costs will be reassigned to pinpoint important missalignments, if they exist. Thus, the majority of deviations from reference models may be not penalized. Instead, only certain deviations should be penalized by assigning a-posteriori high costs to particular misalignments. For instance, given the following partial alignment between a log trace (first row) and a reference model (second row):

$\gg$	$\gg$	`EatingDrinking`	$\gg$	`EatingDrinking`	$\gg$
`Sport`	`MealPreparation`	`EatingDrinking`	`Movement`	$\gg$	`Relaxing`

Deviations like ($\gg$, `Sport`) have high costs since the individual reached a situation where it was expected to do sport. Further, deviations that represent missing activities from the user perspective like ($\gg$, `Meal- Preparation`) or ($\gg$, `Movement`) are assigned with lower costs whereas deviations likes ($\gg$, `Relaxing`) have costs near to zero. Symmetrically, the cost of misalignments denoting activities observed in reality but not expected in the model must also be considered, e.g., (`EatingDrinking`, $\gg$). Finally, groupings of misalignments can be used to improve or correct the costs of the whole alignment. For instance, looking at the misalignment of the example ($\gg$, `MealPrepa- ration`), which may have a low penalization since it does not represent a serious issue (denoting situations where the individual did not prepared a meal but the model requires this action), can be penalized only if it goes next to a synchronous step for the activity `EatingDrinking`.

There are techniques for deriving cost-based fitness analysis of imperative or declarative models [1,11]. These techniques can also be extended to consider other perspectives, i.e., costs or quantities for additional event data [12]. A typical advice on dietary guidelines is *to eat as many calories as one burns* [8]. These kind of checks can be incorporated into the reference model by using the data conformance approach from van der Aalst et al. [12]. Therefore, deviations on quantities can also be verified with respect to the reference model.

6 Experiments

In this section, we present our own data set as well as the experiments that address the introduced usage scenarios. For the experiments, we considered our own data set as well as data sets of related work. Table 1 summarizes those and outlines the characteristics of them. In general, they describe the activities of daily living of individuals, e.g., at home and were manual created. Hence, they represent the actual daily routine of several different persons. Concerning the activity analysis, we focus on the distinction between working and weekend days. Further, we do not compare the results across the different data sets but expose that they support our introduced use cases.

Reference	Scenario	Sensors	Name	Events	Description
Sztyler [31]	daily routine	GPS, ACC, ORI	DailyR	1,386	data set that describes the daily routine of seven individuals.
Cook [10]	smart apartment	Movement	hh102	736	daily routine of different people in an apartment for one month.
			hh104	2,842	
			hh110	837	
Ordónez [25]	life at home	MAG, PRE, PIR, ELE	uniS	691	simple daily routine of two persons for several days at home.
			uniD	870	detailed daily routine of two persons for several days at home.

Table 1. Overview of the considered data sets.

6.1 Data sets

Originally, the authors of these data sets created them for different purposes. Therefore, the data sets cover different aspects and provide also a different granularity concerning the considered labels. In the following, we describe these purposes and also present our own data set in more detail. Besides, these different purposes are also the reason that we created our own data set. Hence, the available data sets do not satisfy our entire introduced requirements.

Sztyler. Our data set covers seven subjects (age 23.1±1.81) that recorded their daily routine for several days. The group covered five students, a worker, and a researcher which collected GPS data and labeled their current location, posture, and activity for the whole day. The subjects were not supervised but got an introduction and guidelines, e.g, we explained the meaning of the predefined labels to avoid that they choose different labels for the same situation. The data was collected using a regular smart-phone and smart-watch combined with a self-developed sensor data collector and labeling framework (see Figure 5). Besides, we also recorded the on-body device position and the acceleration and orientations sensor but do not consider this data during the experiments.

Fig. 5. Collector and labeling framework: Wear App (smart-watch, 1) and Hand App (smart-phone, 2). Our app is free available.[4]

The framework consists of two parts, namely *Wear* and *Hand*. The *Wear* application allows to update the parameters (location, posture, and activity) immediately where the *Hand* application manages the settings and the storing of the data. The labels for the mentioned parameters were predefined and could not be changed or extended (see Table 2).

Concerning the activity labels, we focused on food intake, sport, different type of movements, but also (house) work so that we can compare the daily routine of several individuals to detect common activity patterns but also to analyze the different behaviors. The set of activity labels was minimized and structured to decrease the time which the individual needs to choose a suitable label. There are 12 activities and 33 sub-activities where an activity could be `EatingDrinking` and a corresponding sub-activity `Breakfast`[5]. It was possible to select several activity labels at the same time to record the current situation with a high accuracy (e.g., `Movement/gotoWork`, `Transportation/Train`, and `Sleeping`). Thus, the individual could describe the current situation from several points of view. To keep the set of activity labels as small as possible, we provided some generic labels such as `DeskWork`. This label should be used if the individual works in an office (worker), attends a lecture or class room (student), or visits a school (pupil).

Summarizing, we recorded 74 cases which cover 1,386 events. A case is represented by one individual in one particular day and has an average duration of 12.1

[4] https://play.google.com/store/apps/details?id=de.unima.ar.collector&hl=en
[5] So far, we do not consider the sub-activities in the presented use cases.

Parameter	**Labels**
Device Position	Chest, Hand, Head, Hip, Forearm, Shin, Thigh, Upper Arm, Waist
Environment	Building, Home, Office, Street, Transportation
Posture	Climbing, Jumping, Lay, Running, Sitting, Standing, Walking
Activity	Desk Work, Eating/Drinking – (Breakfast, Brunch, Coffee Break, Dinner, Lunch, Snack), Housework – (Cleaning, Tidying Up), Meal Preparation, Movement – (Go for a Walk, Go Home, Go to Work), Personal Grooming, Relaxing – (Playing, Listen to Music, Watching TV), Shopping, Socializing – (Bar/Disco, Cinema, at Home), Sport – (Basketball, Bicycling, Dancing, Gym, Gymnastics, Ice Hockey, Jogging, Soccer), Transportation – (Bicycle, Bus, Char, Motorcycle, Scooter, Skateboard, Train, Tram)

Table 2. Labeling parameters that have to be updated immediately when they had changed. The subjects had to select at least one of these activity labels to specify their current action. The selection of a sub-activity was optional.

hours. Tables 3 and 4 illustrate the recorded data. The high standard deviation of the numbers of postures results from the different movement behavior.

Labels	**Records** (avg±sd)
Activities	20 ± 7
Postures	80 ± 62
Environment	16 ± 4
Dev. Position	8 ± 6

Table 3. Annotated labels per day and individual.

Raw Data	**Records** (avg)
Acceleration	$2.7 * 10^6$
Orientation	$2.3 * 10^6$
geo. Location	70

Table 4. Number of recorded values per day and individual.

Cook and Ordónez. Their data sets were recorded for different purpose. Cook [10] et al. created the data sets to evaluate a lightweight smart homes design to avoid customization and training. Originally, they considered primarily only movement sensors that record the movement pattern of one or several persons in a apartment. Afterwards, they labeled the record sensor data with the corresponding activity. In contrast, Ordónez et al. [25] investigated the possibility to derive activities of daily living from binary sensor streams in a home setting considering machine learning techniques.

In contrast to our data set, Cook and Ordónez only represent the home environment of the daily routine. However, they considered a broader set of activity labels which results in a more precise description of the behavior.

6.2 Results

In the following, we outline the results of our experiments based on the introduced data sets. The created personal process models from these data sets are available[6]. Based on the derived models, we just inspect them without additional tools. We distinguish between *Workdays* and the *Weekends* and focus on common activity patterns across several days and persons. In this context, we examined the differences between personal processes that consider more general activities (e.g., `grooming`) and such that breakdown the activities (e.g., `washing`, `showering`). As a result, we detected that the personal processes of several people that describe only the behavior at home are more similar than those that illustrates the whole day.

Table 5 illustrates the characteristics of the derived personal process models. The *Density* value represents the degree of connectedness, i.e., number of existing edges in proportion to number of possible edges. A lower value indicates that the personal process has fewer direct transitions between activities, i.e., it is simpler. Considering the models of the data sets *uniD*, *hh102*, *hh104*, and *hh110* it points out that they have the lowest density values but cover the largest set of activities (nodes). This shows that *zooming* into the daily routine of an individual does not lead to a complex structure but uncover common patterns and sequences of specific activities (e.g., `MealPrep.` → `EatingDrinking` → `Cleaning`). Besides, the density of the second model makes clear that the clustering of similar activities leads to a higher density (e.g., `grooming` vs. `toilet`, `wash`, and `shower`).

	Weekday				**Weekend**			
Data set	**Nodes**	**Edges**	**Density**	**Duration**	**Nodes**	**Edges**	**Density**	**Duration**
DailyR[1]	12	19	0.144	18.62	12	20	0.152	22.70
uniS[2]	10	15	0.167	25.23	10	15	0.167	24.33
uniD[3]	14	22	0.121	25.23	13	17	0.109	24.33
hh102[4]	16	24	0.100	28.73	16	27	0.110	19.70
hh104[5]	17	25	0.092	14.06	17	26	0.096	12.88
hh110[6]	15	20	0.095	30.44	14	17	0.093	18.68

Table 5. Characteristics of the derived imperative personal process models.

[6] http://sensor.informatik.uni-mannheim.de/dataset.zip

	Weekday			Weekend		
Data set	**Nodes**	**Edges**	**Density**	**Nodes**	**Edges**	**Density**
DailyR[1]	6/0	7/0	0.233/-	10/0	11/0	0.122/-
uniS[2]	12/4	81/3	0.614/0.250	12/5	100/4	0.758/0.200
uniD[3]	16/5	144/6	0.600/0.300	15/5	158/9	0.752/0.450
hh102[4]	18/4	160/6	0.523/0.500	18/5	197/12	0.644/0.600
hh104[5]	19/3	128/4	0.374/0.667	19/5	157/12	0.459/0.600
hh110[6]	17/4	158/5	0.581/0.417	14/10	156/18	0.857/0.200

Table 6. Characteristics of the derived declarative personal process models.

Further, we identified common patterns that occur in personal processes of different persons (see Patterns 1-4)[7]. For instance, for most people it is very common to go to the bathroom after the turn out. However, there are also patterns that depend on work or weekend days as it is the case for Pattern 3. The activity `Outdoors` has different meanings, i.e., during the week it represents working whereas in context of the weekend it is associated with free time activities. In this context, we detected that `Relax` is the usual activity which is performed after `Outdoor` for workdays. Considering the weekend, the behavior differs, i.e., also `MealPreperation` is a common activity.

$$(\texttt{Medication} \rightarrow)\ \texttt{MealPrep.} \rightarrow \texttt{EatingDrinking} \rightarrow \texttt{Cleaning} \tag{1}$$

$$\texttt{Sleep} \rightarrow \texttt{Toilet/Bath} \tag{2}$$

$$\texttt{Outdoors} \rightarrow \texttt{Relax} \tag{3}$$

$$\texttt{PersonalHygiene/Washing} \rightarrow (\texttt{Medication} \rightarrow)\ \texttt{Sleep} \tag{4}$$

We also noticed that the spend time on specific activities differs across different people but also different days and daytimes. The data sets which distinguish between breakfast, lunch, and dinner, showed that typically the used time for preparing the breakfast is significant lower than for lunch. Moreover, for activities such as `sleeping`, `grooming` (showering, toileting), and `relaxing` (sparetime/tv), we observed that the spend time increased during the weekend.

Concerning declarative models, similar conclusions can be reached as it is illustrated in Table 6 where two experiments are reported: the models obtained with and without simplification. For simplification, we have filtered the process models obtained by using simple heuristics (e.g., removing negative constraints, or fake start/end nodes). For the DailyR benchmark, the obtained models are already simplified but only filtered constraints are derived, which implies to empty the model after manual simplification.

The results show that the personal process models lead to a better understanding of the personal activity data. Further, the resulting graphs, patterns,

[7] Sequence mining techniques may in principle extract similar patterns. One difference is the inability for these techniques to present process view of the extracted patterns.

and features allow to verify certain requirements, e.g., regarding health care or a good lifestyle. As a result, the detected procedures and duration of certain activities can be used to determine the fitness of the derived model.

7 Future Work

So far, we only considered manually created event logs describing personal behavior. However, the automatic creation of them from personal data may enable full automation of the techniques of this paper. This entails a lot of unsolved problems such as the correct recognition of the activities as well as how granular they need to be. Hence, it may be easy to recognize that a person interacts with something in a living room, but it is more difficult to distinguish between watching TV or reading a book. In this context, the personal behavior recognized may differ depending on indoor or outdoor activity recognition available technology. Further, semi-supervised or unsupervised approaches may not allow to consider a predefined set of labels which may result in problems regarding the interpretation and evaluation.

When process mining is applied on personal data, different challenges and directions can be considered that will be explored in the future. First, the aggregation of collected data on different levels of abstraction (e.g., activities like `Reading`, `WatchingTV`, or `Gaming` into `Entertainment`) may enable the simplification of the derived process models. Another challenge is dealing with uncertain data. In particular, the data generated by classification-based methods for activity recognition will most probably be uncertain, since these methods are never a hundred percent accurate. However, provenance information such as explicit uncertain values will be available in most cases, and might serve as an additional input to process mining methods. Hence, process mining methods may need to be adapted in such new context.

With respect to future directions, we focus on two main aspects. On the one hand, the process models derived may be used for something more than just visualization or analysis, i.e., to support the activity of individuals on their daily routine. Notice that historical data of an individual is a rich source of information which may be crucial to influence the daily routine in order to reach a particular goal. In this context, process models can be enhanced and used at each decision point to assess the influence of the next step in satisfying the targeted goal. For instance, following the guideline of the previous section that advice to eat as many calories as one burns, activities can be annotated with respect to calorie levels (e.g., `EatingDrinking` produces an amount of calories while `Movement` takes an amount of calories). Then, historical activity data can be aggregated with this information to learn for all decision points the impact of the decision regarding the likelihood of satisfying the targeted goal, e.g., the balanced consumption of calories. Thus, when an individual is about to start a new activity, recommendations can be provided on the basis on the model's aggregated data corresponding to the current state. This deviates from current prediction and recommendation practices that do not consider the current state of the model explicitly.

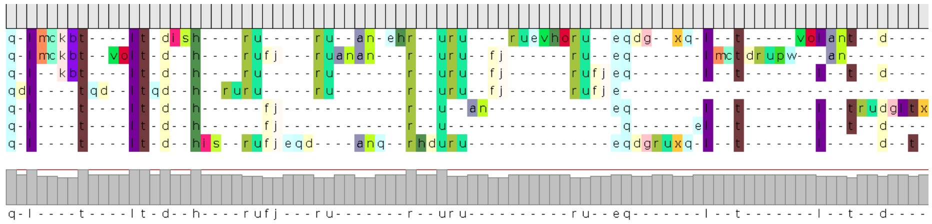

Fig. 6. Example of discovered trace cluster: letters in the bottom denote activities with high consensus. The Y-axis represents seven different traces where the X-axis illustrates the different events per traces.

Finally, another research line will be to preprocess the log with the goal of extracting patterns, and then transform the log accordingly, either by introducing hierarchy, or by ignoring outlier activities not following the learned patterns. For this purpose, *Trace alignment* techniques from van der Aalst et al. [7] can be applied. As an example, in Figure 6 seven traces have been aligned together from the one of the log of the previous section.

8 Conclusions

This paper discusses challenges and opportunities for process mining in the area of personalized health care. It represents the first step towards providing a fine-grained analysis and monitoring of personal processes, which may have very important applications in some domains (e.g., elderly care).

Acknowledgments. This work as been partially supported by funds from the Ministry for Economy and Competitiveness (MINECO) of Spain and the European Union (FEDER funds) under grant COMMAS (ref. TIN2013-46181-C2-1-R).

References

1. A. Adriansyah, N. Sidorova, and B.F. van Dongen. Cost-based Fitness in Conformance Checking. In *Application of Concurrency to System Design conference (ACSD 2011)*, Kanazawa, Japan, June 2011.
2. Rakesh Agrawal and Ramakrishnan Srikant. Mining sequential patterns. In *Data Engineering, 1995. Proceedings of the Eleventh International Conference on*, pages 3–14. IEEE, 1995.
3. A. Aztiria, A. Izaguirre, R. Basagoiti, J. C. Augusto, and D. J Cook. Automatic modeling of frequent user behaviours in intelligent environments. In *Intelligent Environments (IE), 2010 Sixth International Conference on*, pages 7–12. IEEE, 2010.
4. Ling Bao and Stephen S Intille. Activity recognition from user-annotated acceleration data. In *Pervasive computing*, pages 1–17. Springer, 2004.

5. Louise Barkhuus and Valerie E Polichar. Empowerment through seamfulness: smart phones in everyday life. *Personal and Ubiquitous Computing*, 15(6):629–639, 2011.
6. Steven N Blair and Tim S Church. The fitness, obesity, and health equation: is physical activity the common denominator? *Jama*, 292(10):1232–1234, 2004.
7. RP J. C. Bose and Wil M. P. van der Aalst. Process diagnostics using trace alignment: Opportunities, issues, and challenges. *Inf. Syst.*, 37(2):117–141, 2012.
8. Nancy Brown. American Heart Association. http://www.heart.org. Last Access: 29.04.2015.
9. Joos C. A. M. Buijs, Boudewijn F. van Dongen, and Wil M. P. van der Aalst. Quality dimensions in process discovery: The importance of fitness, precision, generalization and simplicity. *Int. J. Cooperative Inf. Syst.*, 23(1), 2014.
10. Diane J Cook, Aaron S Crandall, Brian L Thomas, and Narayanan C Krishnan. Casas: A smart home in a box. *Computer*, 46(7), 2013.
11. Massimiliano de Leoni, Fabrizio Maria Maggi, and Wil M. P. van der Aalst. An alignment-based framework to check the conformance of declarative process models and to preprocess event-log data. *Inf. Syst.*, 47:258–277, 2015.
12. Massimiliano De Leoni and Wil M. P. van der Aalst. Aligning event logs and process models for multi-perspective conformance checking: An approach based on integer linear programming. In *Business Process Management*, pages 113–129. Springer, 2013.
13. F. Friedrich, J. Mendling, and F. Puhlmann. Process model generation from natural language text. In *Advanced Information Systems Engineering - 23rd International Conference, CAiSE 2011, London, UK, 2011. Proceedings*, pages 482–496, 2011.
14. C.W. Günther and Wil M. P. van der Aalst. Fuzzy mining - adaptive process simplification based on multi-perspective metrics. In *BPM*, pages 328–343, 2007.
15. Kurt Jensen, Lars Michael Kristensen, and Lisa Wells. Coloured Petri nets and CPN tools for modelling and validation of concurrent systems. *STTT*, 9(3-4):213–254, 2007.
16. Yanjun Jia. Diatetic and exercise therapy against diabetes mellitus. In *Intelligent Networks and Intelligent Systems, 2009. ICINIS'09. Second International Conference on*, pages 693–696. IEEE, 2009.
17. Eunju Kim, Sumi Helal, and Diane Cook. Human activity recognition and pattern discovery. *Pervasive Computing, IEEE*, 9(1):48–53, 2010.
18. O.D. Lara and M.A. Labrador. A survey on human activity recognition using wearable sensors. *Communications Surveys & Tutorials, IEEE*, 15(3):1192–1209, 2013.
19. Óscar D Lara, Alfredo J Pérez, Miguel A Labrador, and José D Posada. Centinela: A human activity recognition system based on acceleration and vital sign data. *Pervasive and mobile computing*, 8(5):717–729, 2012.
20. F. Leotta. *Instrumenting and Mining Smart Spaces*. PhD thesis, Universita di Roma – La Sapienza, March 2014.
21. Zhenhui Li. In *Frequent Pattern Mining*, pages 283–306. Springer, 2014.
22. Lin Liao. *Location-based activity recognition*. PhD thesis, University of Washington, 2006.
23. Ronny S Mans, Wil MP van der Aalst, and Rob JB Vanwersch. *Process Mining in Healthcare: Evaluating and Exploiting Operational Healthcare Processes*. Springer, 2015.
24. Fosca Giannotti Mireo Nanni. Efficient mining of temporally annotated sequences. In *Proceedings of the Sixth SIAM International Conference on Data Mining*, volume 124, page 348. SIAM, 2006.

25. Fco Javier Ordónez, Paula de Toledo, and Araceli Sanchis. Activity recognition using hybrid generative/discriminative models on home environments using binary sensors. *Sensors*, 13(5):5460–5477, 2013.
26. M. Pesic. *Constraint-Based Workflow Management Systems: Shifting Control to Users*. PhD thesis, Technische Universiteit Eindhoven, 2008.
27. Daniele Riboni, Claudio Bettini, Gabriele Civitarese, Zaffar Haider Janjua, and Rim Helaoui. Fine-grained recognition of abnormal behaviors for early detection of mild cognitive impairment. In *Pervasive Computing and Communications (PerCom), 2015 IEEE International Conference on*, pages 149–154. IEEE, 2015.
28. Anne Rozinat and Wil M. P. van der Aalst. Conformance checking of processes based on monitoring real behavior. *Information Systems*, 33(1):64–95, 2008.
29. Estefanía Serral, Johannes De Smedt, Monique Snoeck, and Jan Vanthienen. Context-adaptive petri nets: Supporting adaptation for the execution context. *Expert Syst. Appl.*, 42(23):9307–9317, 2015.
30. Hong Sun, Vincenzo De Florio, Ning Gui, and Chris Blondia. Promises and challenges of ambient assisted living systems. In *Information Technology: New Generations, 2009. ITNG'09. Sixth International Conference on*, pages 1201–1207. Ieee, 2009.
31. Timo Sztyler, Johanna Völker, Josep Carmona, Oliver Meier, and Heiner Stuckenschmidt. Discovery of personal processes from labeled sensor data–an application of process mining to personalized health care. In *Proceedings of the International Workshop on Algorithms & Theories for the Analysis of Event Data, ATAED*, pages 31–46, 2015.
32. Wil M. P. van der Aalst. *Process Mining - Discovery, Conformance and Enhancement of Business Processes*. Springer, 2011.
33. Wil M. P. van der Aalst, H. T. de Beer, and Boudewijn F. van Dongen. Process mining and verification of properties: An approach based on temporal logic. In *CoopIS, Cyprus*, pages 130–147, 2005.
34. Wil M. P. van der Aalst, Ton Weijters, and Laura Maruster. Workflow mining: Discovering process models from event logs. *Knowledge and Data Engineering, IEEE Transactions on*, 16(9):1128–1142, 2004.
35. A.J.M.M. Weijters, Wil M. P. van der Aalst, and A.K. Alves de Medeiros. Process mining with the heuristics miner-algorithm. Technical Report WP 166, BETA Working Paper Series, Eindhoven University of Technology, 2006.
36. Jie Yin, Qiang Yang, and Jeffrey Junfeng Pan. Sensor-based abnormal human-activity detection. *Knowledge and Data Engineering, IEEE Transactions on*, 20(8):1082–1090, 2008.
37. Yu Zheng. Trajectory data mining: an overview. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 6(3):29, 2015.
38. Xinwei Zhu, Guobin Zhu, Seppe K. L. M. vanden Broucke, Jan Vanthienen, and Bart Baesens. Towards location-aware process modeling and execution. In *Proceedings of the Workshop on Data and Artifact-Centric BPM (DAB'14), Haifa (Israel), September 7-11*, 2014.