

Applying Process Mining to Smart Spaces: Perspectives and Research Challenges

Francesco Leotta¹(✉), Massimo Mecella¹, and Jan Mendling²

¹ Sapienza Università di Roma, Rome, Italy
{leotta,mecella}@diag.uniroma1.it

² Wirtschaftsuniversität Wien, Vienna, Austria
jan.mendling@wu.ac.at

Abstract. A software system managing a smart space takes, among its inputs, models of human behavior; such models are usually difficult to obtain and to validate. The employment of techniques from business process modeling and mining may represent a solution to both the problems, but a set of challenges need to be faced in order to cope with major differences between human activities and business processes. In this work we provide insights about these challenges, and propose further research activities to tackle them.

Keywords: Smart spaces · Process mining · Human habits

1 Introduction

The main goal of a smart space is to employ raw data coming from *sensors* deployed into the environment to make decisions about the environment itself [2]; these decisions are then applied through *actuators*, either virtual or physical, or proposed to the final user in the form of suggestions or alerts.

The reasoning task is performed according to a set of models representing environmental dynamics and, noteworthy, user habits and desires [10]. These models can be either hand-made by experts or automatically extracted (through learning and mining techniques) from previously acquired *sensor logs*. In the latter case, the amount of labeling work in charge of the final user strongly influences how much applicable a specific technique is in a real setting. From this point of view, automatic learning techniques can be classified as “supervised”, “semi-supervised” and “unsupervised”. The latter ones are particularly interesting in a smart space scenario as they (potentially strongly) reduce the burden on the final users to manually label the dataset fed as input. Semi-supervised learning is often associated with “active learning” approaches, whereas unsupervised techniques are the main tools employed for data mining.

If we focus our attention on models of human habits, many attempts have been made to represent human behavior by means of different kinds of state-based modeling techniques. Supervised solutions (probabilistic models such as Hidden Markov Models or deterministic models such as decision trees) are either

difficult to read or they require a prohibitively large amount of data to be labeled. Recently, approaches that model human habits as workflows have been proposed, thus opening to the possibility of applying process mining [15] techniques to obtain, in a semi or unsupervised manner, human-readable models that can be easily validated by end-users. Still, many new challenges have to be addressed.

As noted in [4], despite the growing availability and maturity of process mining techniques, their applicability in some contexts still faces the problem of bridging the gap between events given as input and actions composing the models obtained as output. This problem particularly applies to the challenge of mining human habits, as here sensor logs contain very fine grained events whereas mined model should contain actions at a higher level of abstraction.

Moreover, human habits are flexible in their nature. In the vast majority of cases, humans do not follow a precise workflow while performing their daily activities; they instead follow a series of best practices that are difficult to be described using a precise sequence of tasks.

Finally, process mining techniques need a log segmented in traces, which is not a safe assumption in many application scenarios, including smart spaces.

In this paper we identify a set of research challenges behind the above mentioned problems through the analysis of a case study tailored on an established dataset employed in the smart space community [5].

2 Case Study and Research Challenges

We can abstract a smart space as producing, at runtime, a sensor log containing raw measurements from available sensors. Given a set S of sensors, a sensor log is a sequence of measurements of the kind $\langle ts, s, v \rangle$ where ts is the timestamp of the measurement, $s \in S$ is the source sensor and v the measured value, which can be either nominal (categorical) or numeric (quantitative). Measurements can be produced by a sensor on a periodic base (e.g., temperature) or whenever a particular event is detected (e.g., a door opening).

Many solutions, especially those based on pattern analysis, expect as input an *event log* instead. Given a set $E = \{e_1, \dots, e_n\}$ of event types, an event log is a sequence of pairs $\langle e, t \rangle$ where $e \in E$ and t is an integer, the occurrence time of the event e . Translating a sensor log into an event log may cause a loss of information, especially if discretization of sensor measurements is required.

A typical example of employed sensor is represented by PIR – Presence Infrared – sensors. Usually these sensors are installed on the ceiling following a grid layout; they trigger upon the detection of an object entering their field of view and automatically reset after a fixed amount of time. The CASAS project [5] provides, for example, datasets from different installations and different housing patterns (e.g., number of persons, presence of pets). All these installations contain such a grid of sensors, thus, in the following, we will refer to this particular kind of sensors to provide an insight of the research challenges raised by the application of process mining to the smart space scenario.

2.1 Bridging the Gap Between Sensor Events and Actions

As stated in [4], there exists a clear gap between the granularity of sensor logs and that of traces traditionally employed for process mining. In other words, there is no one-to-one correspondence between sensor measurements and actions performed by a person. Techniques from the machine learning field can be used to aggregate sensor measurements in order to recognize actions. These techniques, originally developed for computer vision applications [13], have been recently moved to a more generic smart space setting [1]. Noteworthy, in smart space literature, the terms “action” and “activity” have a completely different meaning [16]: an activity is a coordinated sequence of actions whose goal is performing a specific daily routine. This distinction influences our approach as well, as we suggest to split the log transformation step from the process mining one.

Sensor measurements and human actions can be directly related, indirectly related or completely unrelated to human actions. In the simplest case, especially when they come from device-attached sensors (e.g., a switch sensor connected to the oven), it is easy to associate sensor measurements to human actions. In other cases, sensors that are, in principle, unrelated to human actions, are instead indirectly influenced by human behavior; as an example, a temperature sensor is usually only influenced by environmental dynamics, but sometimes it may report variations due to the human behavior (e.g., a user is having a shower). In other cases, sensor measurements are completely unrelated to human actions and should be filtered out. Additionally, in some cases sensor measurements are not an effect of human behavior, but act instead as control variables (reflecting some kind of environmental phenomenon) influencing the behavior of the user.

In real usage scenarios, sensor and event logs are affected by a certain degree of uncertainty. Sensors have indeed their own technical limitations as they are prone to breakdowns, disconnections from the system and environmental noise (e.g., electromagnetic noise). As a consequence, measured values can be out of date, incomplete, imprecise, and contradictory to each other. Techniques for cleaning sensor data do exist, but uncertainty of sensor data may still lead to wrong conclusions about the current context.

Referring to our case study, PIR sensors are indirectly influenced by human behavior. In order to (probabilistically) infer human actions from them, it is necessary to first define which devices are installed in correspondence of each one. This association is of the type n -to- n ; on the one extreme, a PIR sensor may cover multiple actions (e.g., interacting with both the fridge and the oven), thus limiting the expressiveness of models if only PIR sensors are available; on the other extreme, a single action (e.g., cooking) may cover several PIR sensors.

Another way of looking at this scenario is to think in terms of human action patterns with respect to PIR sensor triggering, e.g., the following ones:

- Static: only one PIR sensor triggers while a human is performing an action. As an example, when the user is chopping vegetables, the only interested sensor will be the one in correspondence of the kitchen table.
- Interleaved: multiple sensors trigger while the user is performing an action. In the CASAS dataset for example, while the user is sleeping, two PIR sensors

interleave, namely the ones over the two sides of the bed. The same pattern applies to higher level actions (e.g., cooking, cleaning the house).

- Movement: a set of PIR sensors quickly trigger in sequence because the user is moving from one place of the house to another to perform different actions.

These considerations make it necessary to preprocess a sensor log containing PIR sensor measurements by filtering certain sensor measurements and aggregating other ones. Methods to do that should take into account *(i)* the position of the sensor, *(ii)* the set of devices which potentially relates to the sensor, and *(iii)* the amount of time that elapses from a sensor measurement to another one. As an example, while a user is moving we will see a quick sequence of sensor triggering one close to the other, all these sensors will probably belong to movement areas of the house (e.g., a corridor); such a sequence should be filtered out. As another example, the interleaving between the two sensors over the bed should be easily aggregated as they have a single device in common (i.e., the bed).

2.2 Improving on Process Mining and Modeling Techniques

A basic question for the application of process mining techniques to human habits is whether human behavior is structured enough to be described using a process model. If the answer to this question is affirmative, the obtained model will probably resemble a “spaghetti” process [15]. Approaches to deal with unstructured processes do exist for both imperative and declarative modeling formalisms.

Declarative modeling formalisms, e.g., DECLARE [12], are usually based on temporal logics; even though different graphical representations are available to describe them, reading a declarative model is not a trivial task, requiring the knowledge of the logic language behind symbols. Additionally, as should be clear from the definitions of sensor and event logs, time is a first-class property of a measurement, but the notion of time in declarative formalisms is qualitative and not quantitative. Nevertheless, attempts to support a quantitative notion of time are available [17], and temporal logics have been applied to smart spaces [11].

Another typical approach to deal with unstructured processes is fuzzy mining [9]; it borrows concepts from the world of maps and cartography and apply them to zoom in and out on a process model highlighting the importance of certain tasks and connection between tasks just like they were points and paths on a map. It is worth to note how the implicit uncertainty coming from the transformation that turns sensor measurements into tasks should be taken into account. As an example, let us consider the derived actions resulting from the filtering and aggregation as described in Sect. 2.1, as task instances and apply a declarative miner [7, 8] to the resulting dataset. Assume now that a constraint between two PIR sensors M001 and M002 has been discovered imposing that the former only triggers if the latter has previously triggered. If we say that M001 triggers because of the user interacting with either the device D001 or the device D002, this results in a branching condition. In that case we should take into account the fact that this is the result of a probabilistic decision. From this point of view, techniques like fuzzy mining appear to be particularly suitable in the smart space scenario as, by definition, they support uncertainty.

2.3 Automatic Log Segmentation

A common prerequisite of process mining techniques is to have a trace log explicitly segmented in traces (process instances). This assumption is usually not met by sensor logs as labeling is generally an expensive task to be performed by humans. This problem can be translated into a precise question: what is a habit? Is that a specific routine (e.g., cleaning the kitchen) or is it more generally what happens in a specific time frame of the day (e.g., between 8:00 am and 9:00 am).

On the one hand, choosing the second solution makes it easier to segment the log as we cannot assume that a user will not perform different routines at the same time and there is no way, from the log solely, to identify such a situation. On the other hand, according to what criteria should we choose the time ranges? The trivial solution of segmenting the day in fixed length stints of time fails if something anomalous happens (e.g., holidays). Possible solutions are:

- Associating specific combinations of sensor measurements to the beginning and to the end of a time range. E.g., the morning habit starts whenever a movement sensor different than the ones mounted over the bed triggers and ends whenever the movement sensor in correspondence of the exit door triggers and no other movement sensor triggers for a certain amount of time.
- Active learning. In order to let the system learn his habits, the user explicitly signals to the system the beginning and the end of a specific habit.

However, time is not the only dimension to take into account for segmentation. Differently from the common process mining scenario, usually sensor logs do not contain any information about which user/s caused a certain sensor to trigger or to provide a specific measurement. Thus, mining habits in a multi-user scenario is significantly harder especially with respect to the transformation task introduced in Sect. 2.1. Our use case, based on PIRs, is particularly challenging from this point of view as, even though multiple users can be identified by the spatial distance between PIRs triggering close in time, when trajectories intersect, tracking techniques [14] must be employed to keep following users.

3 Concluding Remarks

Learning human habits is a hot topic in the pervasive computing community. Such models can be employed to assist humans during activities of daily living or for energy saving strategies. Many works faced this challenge proposing techniques from machine learning and data mining. Nonetheless, the proposed approaches produces unreadable models or require an excessive labeling effort.

As shown in [3], using process modeling techniques (workflows in particular) for human habits represents a promising approach to solve the readability. Other approaches instead use techniques from the data mining area to extract subsequences of symbols that represents different activity templates [6].

Process mining could offer the best of two worlds but three research challenges need to be faced, namely *(i)* the gap between sensor logs and tasks employed in process mining, *(ii)* designing process modeling and mining techniques fitting the variability of human habits, and *(iii)* the problem of segmentation of logs. In this paper we briefly discussed them, by providing examples based on the PIR-based CASAS dataset, in order to stimulate further research in the field.

Acknowledgments. The work of F. Leotta and M. Mecella has been partly supported by the EU project VOICE, the Italian clusters SM&ST (Social Museum & Smart Tourism) and Smart Living Technologies, the project RoMA, and the Intel Galileo University Donations.

References

1. Aggarwal, J.K., Ryoo, M.S.: Human activity analysis: a review. *ACM Comput. Surv.* **43**(3), 16 (2011)
2. Aztiria, A., Izaguirre, A., Augusto, J.C.: Learning patterns in ambient intelligence environments: a survey. *Artif. Intell. Rev.* **34**(1), 35–51 (2010)
3. Aztiria, A., Izaguirre, A., Basagoiti, R., Augusto, J.C., Cook, D.J.: Automatic Modeling of Frequent User Behaviours in Intelligent Environments. In: *Intelligent Environments*, pp. 7–12 (2010)
4. Baier, T., Mendling, J.: Bridging abstraction layers in process mining by automated matching of events and activities. In: Daniel, F., Wang, J., Weber, B. (eds.) *BPM 2013*. LNCS, vol. 8094, pp. 17–32. Springer, Heidelberg (2013)
5. CASAS Project Website. <http://wsucasas.wordpress.com>
6. Cook, D., Krishnan, N., Rashidi, P.: Activity discovery and activity recognition: a new partnership. *IEEE Trans. Cybern.* **43**(3), 820–828 (2013)
7. Di Ciccio, C., Maggi, F.M., Mendling, J.: Discovering target-branched declare constraints. In: Sadiq, S., Soffer, P., Völzer, H. (eds.) *BPM 2014*. LNCS, vol. 8659, pp. 34–50. Springer, Heidelberg (2014)
8. Di Ciccio, C., Mecella, M.: A two-step fast algorithm for the automated discovery of declarative workflows. *Symposium on Computational Intelligence and Data Mining*, pp. 135–142 (2013)
9. Günther, C.W., van der Aalst, W.M.P.: Fuzzy mining – adaptive process simplification based on multi-perspective metrics. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) *BPM 2007*. LNCS, vol. 4714, pp. 328–343. Springer, Heidelberg (2007)
10. Leotta, F.: *Instrumenting and Mining Smart Spaces*. Ph.D. thesis (2014)
11. Magherini, T., Fantechi, A., Nugent, C.D., Vicario, E.: Using temporal logic and model checking in automated recognition of human activities for ambient-assisted living. *IEEE Trans. Hum. Mach. Syst.* **43**(6), 509–521 (2013)
12. Pesic, M., Schonenberg, H., van der Aalst, W.M.P.: Declare: full support for loosely-structured processes. In: *IEEE Conference on Enterprise Distributed Object Computing (EDOC)*, pp. 287–300 (2007)
13. Poppe, R.: A survey on vision-based human action recognition. *Image Vis. Comput.* **28**(6), 976–990 (2010)
14. Smeulders, A., Chu, D., Cucchiara, R., Calderara, S., Dehghan, A., Shah, M.: Visual tracking: an experimental survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(7), 1442–1468 (2014)

15. Van Der Aalst, W.: Process mining: discovery, conformance and enhancement of business processes, Springer Science & Business Media (2011)
16. Ye, J., Dobson, S., McKeever, S.: Situation identification techniques in pervasive computing: a review. *Pervasive Mob. Comput.* **8**(1), 36–66 (2012)
17. Westergaard, M., Maggi, F.M.: Looking into the future. In: Meersman, R., Panetto, H., Dillon, T., Rinderle-Ma, S., Dadam, P., Zhou, X., Pearson, S., Ferscha, A., Bergamaschi, S., Cruz, I.F. (eds.) *OTM 2012, Part I. LNCS*, vol. 7565, pp. 250–267. Springer, Heidelberg (2012)