

Towards the Second-Order Adaptation in the Next Generation Remote Patient Management Systems

Aleksandra Tesanovic^a and Goran Manev^{a,b}

^aPhilips Research Laboratories
High Tech Campus 37, 5656 AE Eindhoven,
the Netherlands
aleksandra.tesanovic@philips.com
goran.manev@philips.com

Mykola Pechenizkiy^b and Ekaterina Vasilyeva^b

^bEindhoven University of Technology
Department of Computer Science
P.O. Box 513, 5600 MB Eindhoven,
the Netherlands
{m.pechenizkiy, e.vasilyeva}@tue.nl

Abstract

Remote Patient Management (RPM) systems are expected to be increasingly important for chronic disease management as they facilitate monitoring vital signs of patients at their home, alerting the care givers in case of worsening. They also provide patients with educational content. RPM systems collect a lot of (different types of) data about patients, providing an opportunity for personalizing information services. In our recent work we highlighted the importance of using available information for personalization and presented a possible next generation RPM system that enables personalization of educational content and its delivery to patients. We introduced a generic methodology for personalization and emphasized the role of knowledge discovery (KDD). In this paper we focus on the necessity of the second-order adaptation mechanisms in the RPM systems to address the challenge of continuous on-line (re)learning of actionable patterns from the patient data.

1 Introduction

Remote patient management (RPM) systems offer a potential for reducing hospitalization costs and worsening of symptoms for patients with chronic diseases, e.g., coronary artery disease, heart failure, and diabetes. An RPM system both monitors vital signs and provides a feedback to the patient in terms of appropriate education and coaching. Although the large volumes of data collected by RPM systems provide an opportunity for tailoring and personalizing information services, there is a limited understanding of the necessary architecture, methodology, and tailoring criteria to facilitate personalization of the content.

In our recent work [7] we presented an architecture of the next generation RPM systems that facilitates personal-

ization of educational content and its delivery to patients. We also introduced a generic approach for personalization of RPM and thereafter focused on (off-line) knowledge discovery from patients' data from a clinical trial.

In this paper we go one step further highlighting the importance and the necessity of continuous on-line (second order) adaptation in RPM systems, i.e., developing a mechanism which would adjust or adapt the behavior of the adaptive system.

The rest of the paper is organized as follows. We briefly review the state of the art in RPM systems and present our view on the next generation of RPM systems in Section 2. In Section 3 we discuss the potential of second-order adaptation mechanisms in RPM systems and consider a few hypothetical examples of gradual and abrupt concept drift, re-occurring contexts and context-aware learning. We conclude with brief discussion and outline the directions of the future work in Section 4.

2 Remote Patient Management Systems

Existing commercial RPM systems normally provide an end-to-end infrastructure that connects patients at home with health professionals at their institution. The patients at home are equipped with a number of sensors measuring vital signs to obtain objective measurements about their physical condition. The vital sign measurements (e.g. weight, blood pressure, glucose) are transferred to the monitoring and management server. Subjective measurements (e.g., symptoms and quality of life scores) are collected from the patients via questionnaires. Objective and subjective measurements (referred to as RPM data) are presented to the medical professional who, based on the indicated deviations from the normal values, adjusts the patient's treatment plan, including medications and lifestyle goals (e.g., nutrition and physical activity).

The majority of commercial RPM systems only have the link between the patient and professional that enables uploading patient data to the professional for review and treatment changes; these systems are typically referred to as remote patient monitoring systems as they provide only monitoring, but not the management part.

The current commercial systems typically send the same generic non-personalized content to all the patients, regardless of their current health condition, knowledge level, or a mental state.

Research on personalization is ongoing in e-Learning and there is a number of successful implementations of adaptive hypermedia systems like AHA!, Interbook, etc. [1]. However, existing architectures are not adopted in eHealth applications such as RPM systems. Furthermore, in the mentioned systems, the adaptation and personalization is pre-authored and thus remains highly static and often subjective based on some domain expertise translated to the machine readable form.

Developing personalized RPM systems is possible only if we can learn key (potentially changing and dynamic) characteristics of the patients and track them continuously. Personalization can be organized using individual and group (or stereotype) user modeling. In a stereotype approach, the users are classified into several groups. In eHealth applications users can be classified according to their main disease, background in medicine (patients, nurses, and physicians), general education background (no degree, college degree, doctorate, etc), and their tasks (consultation, education, and emergency cases). Individual patient (user) models, besides the user's medical profile, could include also individual characteristics such as cognitive and psychological individual peculiarities, the interaction parameters – the last visited pages, used links, number of the particular pages visits, resource usages etc. Table 1 gives an overview of possible features of various data classes that can play a role in the patient model of an RPM system. A feature can be *static*, e.g. gender, residence, language, or *relatively static*, e.g., age, cognitive impairment (which a patient can develop during the usage of RPM system) and *dynamic*, e.g., values of weight measurements or system usage. The example given is for heart failure, but can be generalized to any of chronic diseases given a specific set of relevant symptoms and vital signs for that chronic disease.

Dynamic features plays an important role in the patient model of RPM system. This calls for the second-order adaptation mechanisms in the RPM systems to address the challenge of continuous on-line (re)learning of actionable patterns from the patient data.

Table 1. Typical features included in a patient model template

Data class	Feature	Changes	
		Static	Dynamic
Demographic	Gender	x	
	Age	x	
	Country	x	
	Language	x	
Living status	Single/Family	x	
Baseline data	Weight		x
	Height	x	
	Body Mass Index		x
	Edema		x
	Biomarker values		x
Medical history	Cause of disease	x	
	Co-morbidities		x
	Implantables	x	x
Symptoms	Ankle swelling		x
	Breathlessness		x
	Depression		x
	Anxiety		x
Vital signs (Frequencies of values out of band)	weight		x
	heart rate		x
	blood pressure		x
	diastolic blood pressure		x
System usage (Frequency of measurements)	weight		x
	blood pressure		x
	heart rate		x
Learning styles	Verbaliser/Imager	x	
	FD/FI	x	
Cognitive function	Reduced eyesight	x	
	Dementia	x	

Legend: *Frequency of vital sign measurements* - how often the patient has been using a sensor for a measurements (1 – every day, 0- not at all), *FD/FI* – field dependent/independent.

3 Second-order adaptation

In our recent work [7] we suggested a general architecture of an personalized RPM system in which we followed general principles of personalization in e-Learning systems with KDD process as one of the key integrated components. Figure 1 depicts a part of the architecture that could provide a possible foundation for the next generation adaptive RPM systems.

The key components of the system that facilitate personalization and adaptation include: (1) patient (user) model, (2) domain model, (3) adaptation rules, (4) adaptation engine, and (5) KDD process. Further, there are authoring and management tools allowing medical experts and professionals to monitor, control and manage patient models, domain models and adaptation rules.

In this section we consider briefly the role of knowledge discovery for patient modeling and provide motivating examples for handling gradual and sudden changes in




Figure 1. A high level view of the next generation RPM system [7]

the modeled concepts, and for learning contextual features describing reoccurring contexts.

3.1 Knowledge discovery for patient modeling in RPM systems

The *KDD process* is essential for discovering relevant actionable patterns that are the basis for creation of the patient model and the adaptation rules. This KDD process is (initially) done “off-line”, using stable historical data available from an existing RPM database or from completed clinical trials relevant for the disease in question. Via this KDD process we obtain relevant patterns that are used to build a *patient model template*. The same patterns are utilized to build the *adaptation rules* and *domain model* of the available content material that is stored in corresponding databases. The KDD is highly iterative and interactive and involves considerable effort from domain and KDD experts. Moreover, this is by no means one-time activity. With accumulation of new evidence and possible contextual changes, models and rules might need an update or extension.

Even-pattern and time-series analysis are particularly helpful in getting a better understanding of what features and relations between them may potentially describe patient current state and its short-term and long-term dynamics.

In general different types of approaches can be used for discovery of useful patterns, including association analysis, subgroup discovery, etc. In this study we search for discriminating patterns by defining corresponding classification tasks. For example, we searched for rules that would predict next symptom status and change in next symptom values, using the last symptom status, gender, age, and frequency of daily system usage in the period between two

symptoms (typically one month) as predictive features.

Table 2 illustrates example patterns found (with the help of popular J48 and JRIP classification techniques) for two most prominent symptoms, breathlessness and swelling of ankles. We obtained off-line the exact number for the level of the usage of the system (i.e., *freqOfWeightUsage*), but have replaced that with a parameter (Low, Medium, High) because it changes over time. We discuss extensively in the next section how this parameter changes in the context of necessity for second-order adaptation. From Table 2 we can observe that women in general are at higher risk to remain breathless (P1) and remain with swelling ankles if they do not use the system regularly (P4-P5). In general, patients are in risk if they are under-utilizing the system (P2-P3), while male patient population above 75 is at risk for worsening of their condition (P6). We refer an interested reader to [7] for a more detailed discussion of the KDD process and a particular case study.

In Table 3 we present possible adaptation rules based on previously discovered patterns P1-P6. These rules mostly identify patients at risk for worsening of their condition, notify the medical professional about risk, and send adequate content to the patient so that worsening can be prevented. E.g., the first rule based on pattern P1 would send content material to the patient to help her master her breathing, while at the same time notify the medical professional that this woman is at risk to remain breathless. Similarly the second rule based on patterns P2-P3 would identify patient at risk and send appropriate educational and instructional material to the patient and notification about risk to the professional. In this case a patient needs to be motivated to use the system, and properly instructed how to do so.

Level of system usage (denoted *freqOfWeightUsage*

Table 3. Examples of adaptation rules

P#	Possible Rule	Desired effect	
		Patient	Medical professional
P1	If Sex=F and BreathlessSymptom=B then Send videos with breathing exercises	Regain control over the breathing	Notification for patient at risk
P2, P3	If BreathlessSymptom=A and Age=(37.5-81.5] and (freqOfWeightUsage <0.4 or freqOfPulseUsage < 0.4) then Send Motivational content	Motivation, instruction for using the system, education on breathlessness	Notification of patient at risk
P4	If SwellingSymptom = 'B' and Sex = 'F' and freqOfWeightUsage < 0.6 then Send Motivational video	Motivation, instruction for using the system, education on swelling ankles	Alert for additional action
P5	If StartSymptom = 'S' and Sex = 'F' and freqOfWeightUsage < 0.6 and Age < 74.5 then Send motivational content		Notification of patient at risk
P6	If SewllingSymptom = 'S' and Sex = 'M' and Age > 74.5 then Send educational content	Motivation, education on importance of managing condition	Notification of patient at risk

Table 2. Examples of discovered patterns

Patterns	Symptom
P1 (StartSymptom = 'B') & (Sex = F) => NextSymptom = B	Breathlessness
P2 (StartSymptom = 'A') & (Age = '(37.5-81.5]') & (freqOfWeightUsage <0.4) => NextSymptom = 'B'	
P3 (StartSymptom = 'A') & (Age = '(37.5-81.5]') & (freqOfPulseUsage <0.4) => NextSymptom = 'B'	
P4 (StartSymptom = 'B') & (Sex = 'F') & (freqOfWeightUsage < 0.6) => NextSymptom = 'B'	Swelling of ankles
P5 (StartSymptom = 'S') & (Sex = 'F') & (freqOfWeightUsage < 0.6) & (Age < 74.5) => NextSymptom = 'W'	
P6 (StartSymptom = 'S') & (Sex = 'M') & (Age >=74.5) => NextSymptom = 'B'	
Legend: Start/Next Symptom: G = good (no problem), S = small problems, A = average, B = bad (many problems), W = worse, I = improved	

in Table 3) is one of the user model parameters that clearly has significant impact on adaptation. We consider the possibly changing nature of the *freqOfWeightUsage* concept to motivate the necessity for second-order adaptation in terms of handling a concept drift.

3.2 Coping with sudden and gradual changes

The user parameters discovered in the patterns, such as level of system usage is likely to change over time due to (i) change in patient motivation to use the system because of the educational and instructional material that has been sent, (ii) change in patients' lifestyle habits impact of these to the

usage of the system), (iii) change in seasonal patterns as patients tend to measure themselves differently during summer compared to winter or during working days compared to public holidays, or (iv) any other hidden context. The system needs to be able to detect and adapt to the changes quickly and without additional input from the patient, e.g., by collecting additional evidence for motivation or lifestyle habit changes¹.

Consider the following example presented in Figure 2. It shows how the progression of the disease can develop over time. The patients could be placed on the system with only one co-morbidity, and due to their age and progression of the chronic disease (heart failure) could develop a number of other conditions. The bottom figure shows what could be the effect of additional conditions to the system usage. Namely, with the increase in number of conditions the patients' overall health could significantly decade (while possibly keeping the main parameters of heart failure still in the normal ranges), directly influencing the patient's motivation and ability to measure him/herself.

Another example is the change of the patient's cognitive abilities. The decline in cognitive abilities could occur for patient over time and it could possibly effects the usage of the system. For example, the patient could become more forgetful, develop initially very mild, mild, or sever case of dementia. This could directly impact his ability to use the system - (s)he would start forgetting to weight or measure blood pressure, initially sporadically, and then more and more often.

It is rather intuitive that besides system usage, such parameters as patient's weight and weight change can trigger

¹ Even if one would like to measure motivation by asking the patient, that would be least preferred option as already patients are faced with many symptom questions they need to answer on daily basis.




Figure 2. Changes in co-morbidities and expected system usage over time.

different The weight of the patient could increase over the time. Normally the rule for alerting the care giver about the risk of heart failure is very simple: if weight increase is more than 2 kilos over 1-2 days, raise an alert. However, there are patients who decompensate with slow increase of weight over period of 10-15 days. Hence, it could happen that the patient gets hospitalized without the alert being raised due to slow increase. The system should be able to learn the slow increase in weight for these patients and adopt the alerting rules accordingly. Furthermore different external events like national and religious holidays can affect normal eating habits. Such events can be recognized from the data and model adjusted accordingly, or proactive handling of possible change could be implemented,

As discussed, in RPM systems (and eHealth domain in general) the concept of interest (user parameters) depends on a changing context that is not necessarily given explicitly in the form of predictive features. Hence, stationary data distribution assumed by majority of traditional data mining techniques is no longer the case. Rather, here we are faced with *concept drift* [11], i.e., unforeseen changes over time in the phenomenon of interest. The phenomenon here would be the usage of system behavioral pattern relevant to current potentially hidden context. The concept we are trying to learn (value of the level of system usage) for the true patient model parameter depends on the observed behavior.

Changes in the hidden context may not only be a cause of a change of the target concept, but may also cause a change of the underlying data distribution. Even if the target concept remains the same, and it is only the data distribution that changes, this may often lead to the necessity of revising the current model, as the model's error may no longer be acceptable with the new data distribution (e.g., more females added to the system may change the behavior). The need to the change of current model due to the change of data distribution is called *virtual concept drift* [10].

Three approaches to handling concept drift can be distinguished: (1) instance selection; (2) instance weighting; and (3) ensemble learning [8]. In instance selection, the goal is to select instances relevant to the current concept. The most common concept drift handling technique is a sliding window and consists in generalizing from a window that moves over recently arrived instances and uses the learnt concepts for prediction only in the immediate future [11].

Instance weighting uses the ability of learning algorithms such as Support Vector Machines (SVMs) to process weighted instances [2]. Instances can be weighted according to different features such as "age" or competence with regard to the current concept. Klinkenberg [2] demonstrates in his experiments that instance weighting techniques handle concept drift worse than analogous instance selection techniques.

Ensemble learning is among the most popular and effective approaches to handle concept drift, in which a set of concept descriptions built over different time intervals is maintained, predictions of which are combined using a form of voting, or the most relevant description is selected [4]. Street and Kim [6] and Wang *et al.* [9] suggest that simply dividing the data into sequential blocks of fixed size and building an ensemble on them may be effective for handling concept drift. Stanley [5] and Kolter and Maloof [3] build ensembles of incremental learners in an online setting, starting to learn new base classifiers after fixed intervals, while continuing to update the existing ones.

3.3 Reoccurring contexts and context-aware adaptation

Figures 3 and 4 show an example of changes in the level of system usage due to seasonal patterns. Seasonal "index" (based on simple average) is used for seasonal construction. Two types of seasonality are shown: period of the year (fall, winter, spring, summer), and week days (Monday to Sunday). Expected seasonal behavior of male patients from Figure 3 would be to use the system at most in fall while female patients in winter and fall. In summer both male and female patients use the system less, and in spring men use the system (in average) more than women (less than average). Figure 4 shows a tendency of both men and women to use the system less during weekends, and to use the system at most on Friday. Additionally, local behaviors, such as, holidays can change the seasonality concept as during the holiday (regardless of the day of the week) patients are under-utilizing the system.

Thus, there exists seasonality behavior in the level of the usage of the system for different patients, which changes the concept. This behavior can be detected and handled by methods which learn from data under the assumption of concept drift. Once the seasonality is discovered, features




Figure 3. Seasonal index (per periods of year) of the use of the system.




Figure 4. Seasonal index (per days of the week) of the use of the system.

representing seasonal context, such as holiday, season, and week day, become a part of the observation set that the system is monitoring. Thereby, in the summer or during the holidays the adaptation rules should be modified to include, for example, additional content that would motivate the patients to use the system. Examples of additional rules are given in Table 4. The difference between rule P8 and rule P10 is due to the expected decline in usage of the system during summer, compared to winter. Hence, when the level of usage of the system is Medium, an alert is sent only in the winter. Note that behavior described with these additional rules can further change over time.

Table 4. Examples of additional rules

#P	Additional rules
P9	Season = Winter and Sex = Female and freqOfWeightUsage = Medium then send additional motivational content, alert medical professional
P10	Season = Winter and Sex = Female and freqOfWeightUsage = Low then send additional motivational material, alert medical professional
P11	Season = Summer and Sex = Female and freqOfWeightUsage = Medium then send motivational material, do not alert
P12	Summer = Winter and Holiday = False and freqOfWeightUsage = Low then send alert to the medical professionals
P13	Summer = Winter and Holiday = True and freqOfWeightUsage = Low then do nothing

The direct potential relation between co-morbidities and usage of the system considered in the earlier example implies that the rules triggering an alert due to non-usage of the system might need additional conditions that would take into account number of co-morbidities. The context of co-morbidities could be known in the system, e.g., the patient goes to the clinical assessment and notifies caring nurse, but it could also be hidden in the sense that the patient does not notify the caring professional about the new disease (which could be as simple as breaking the leg or falling, or as complex as diabetes or renal failure).

The possibility that usage of the system declines with the similar rate of decline of patient mental abilities should be reflected in the system. The system should (i) be able to detect the pattern of slow decline of the system usage, (ii) when that is detected, send appropriate cognitive tests to re-confirm the cognitive decline, and (iii) potentially modify the alerting rules such that they incorporate the cognition ability and thereby have the threshold of alerting based on

the usage of the system different. Moreover, the delivery of care should be modified to possibly include cognitive ability tests more regularly.

4 Conclusions and further work

Remote Patient Management (RPM) systems are expected to be increasingly used in the near future. The current generation of RPM systems follows the one-size-fits-all approach despite of the wide acceptance of the benefits of personalization and adaptation of information services.

In our work we focused on the data driven approaches to adaptation. We considered motivating examples which illustrate ideas behind patient profiling and tailoring the educational or motivational content. This allowed us to come closer to the following challenge – the second order adaptation in RPM systems. We illustrated with further more detailed consideration of naive seasonal and generally time-changing patterns, the benefits of online learning, concept drift handling mechanism, and discovery and use of contextual features for adapting the set of adaptation rules, and user modeling procedures.

In this paper we considered mostly either hypothetical or rather fragmented examples of patient modeling in the context of the second order adaptation. Our further work includes knowledge discovery from data collected during the real clinical studies to justify the advantages of systems equipped with concept drift handling and context-sensitive learning mechanisms.

References

- [1] P. Brusilovsky and E. Millán. User models for adaptive hypermedia and adaptive educational systems. *The Adaptive Web*, pages 3–53, 2007.
- [2] R. Klinkenberg. Learning drifting concepts: Example selection vs. example weighting. *Intelligent Data Analysis*, 8(3):281–300, 2004.
- [3] J. Z. Kolter and M. A. Maloof. Dynamic weighted majority: A new ensemble method for tracking concept drift. In *ICDM '03: Proceedings of the Third IEEE International Conference on Data Mining*, page 123, Washington, DC, USA, 2003. IEEE Computer Society.
- [4] L. I. Kuncheva. Classifier ensembles for changing environments. In F. Roli, J. Kittler, and T. Windeatt, editors, *Multiple Classifier Systems*, volume 3077 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2004.
- [5] K. Stanley. Learning concept drift with a committee of decision trees, 2001.
- [6] W. N. Street and Y. Kim. A streaming ensemble algorithm (sea) for large-scale classification. In *KDD '01: Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 377–382, New York, NY, USA, 2001. ACM.
- [7] A. Tesanovic, G. Manev, M. Pechenizkiy, and E. Vasilyeva. eHealth personalization in the next generation rpm systems. In *Proceedings of IEEE International Symposium on Computer-based Medical Systems (CBMS 2009)(to appear)*. IEEE, 2009.
- [8] A. Tsymbal, M. Pechenizkiy, P. Cunningham, and S. Puuronen. Dynamic integration of classifiers for handling concept drift. *Information Fusion*, 9(1):56–68, 2008.
- [9] H. Wang, W. Fan, P. S. Yu, and J. Han. Mining concept-drifting data streams using ensemble classifiers. In *KDD '03: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 226–235, New York, NY, USA, 2003. ACM.
- [10] G. Widmer and M. Kubat. Effective learning in dynamic environments by explicit context tracking. In *ECML '93: Proceedings of the European Conference on Machine Learning*, pages 227–243, London, UK, 1993. Springer-Verlag.
- [11] G. Widmer and M. Kubat. Learning in the presence of concept drift and hidden contexts. *Mach. Learn.*, 23(1):69–101, 1996.