Designing a 3-Stage Patient Deterioration Warning System for Emergency Departments

Thomas Schmidt, Uffe Kock Wiil
The Maersk Mc-Kinney Moller Institute
University of Southern Denmark
Campusvej 55, 5230 Odense M, Denmark
{schmidt,ukwiil}@mmmi.sdu.dk

ABSTRACT
Detecting patient deterioration in Emergency Departments is a substantial challenge due to the variety in admission reasons and range in patient demographics. Recent research has found that 20% of patients triaged to have non-critical symptoms, deteriorate within the first 24 hours after admission. This can lead to serious adverse events in a clinical setting where patient monitoring relies solely on manual observations of monitors at infrequent intervals. In this position paper, we present a novel 3-Stage Patient Deterioration Warning System as a model to mitigate the risk of undetected deterioration and alarm fatigue. This staged approach enables patients to be monitored at levels of increasing descriptiveness, ranging from population through group to individual models of normality. We argue for the validity of the model through related work, clinical observations, and patterns of patient data collected at an emergency department bedside ward. Finally, we present concrete plans for future implementation work.

Categories and Subject Descriptors
I.5.2 [**Pattern Recognition**]: Design Methodology

General Terms
Algorithms, Performance, Design.

Keywords
Health Informatics patient monitoring, emergency departments, predictive modelling, Hidden Markov Models.

1. INTRODUCTION
Roughly 20% of patients arriving at an emergency department (ED) with non-critical vital values, deteriorate within the first 24 hours [11]. In line with the assumption that 3 to 6% of all deaths in hospitals are unexpected [10] and [27], this has spawned several attempts to prevent these situations by increasing formalization and automation of patient observations.

In this regard much scholarly effort has addressed improving the predictive accuracy of Early Warning Systems (EWS) for detecting patient deterioration. Broadly speaking, these attempts can be classified as either improving the system for identifying deteriorating patients or seeking to automate the deterioration detection system to alleviate the cognitive and physical workload on clinicians. Permutations of the first aspect have been researched with regards to improving EWS's in general wards [15], and in ED related contexts similar to our settings [6]. A recurring theme in this research is the inability to definitively determine exact vital sign thresholds and correlations to mark the initial stage of deterioration. Thus most systems have poor quality of supporting evidence [5] and [2]. Most papers do however note an effect of implementing a formalized EWS. This is in line with the second line of attempts, where the process of identifying patients at risk depends on collaboration and communication between multiple actors. This interplay has been coined as the "chain of survival" [21]. This chain consists of: 1) high-quality recording of vital signs; 2) ability to recognize familiar patterns; 3) reporting of abnormality; and 4) a precise and prompt response. With this baseline, this position paper focuses primarily on introducing a model that adheres to the clinical reality, and to provide a system design that enables higher accuracy, while accommodating the cognitive models of clinicians.

Our work includes a field study of an Emergency Department, with the purpose of identifying collaborative and organizational causes for undetected patient deterioration. This field study is based on 15 participatory observation sessions of complete 8 hour shifts with different clinical groups. We conclude that future patient deterioration detection systems must convey information about patient state and trajectory ubiquitously throughout the department, and not just at bedside or in designated offices to overcome both temporal and spatial challenges. Additionally, as each clinical group has different areas of expertise and means of interaction, the deterioration warning system should differentiate patient state representation according to individual clinical groups. The results of our field study are not yet published.

From the field study, we also identified that clinical observations and automated patient monitoring is challenged by the fact that certain groups of patients and individuals fall outside the population based model of normality currently in use at the ED. Several attempts to improve the predictive accuracy of EWS's have already been made (see Section 3). However, our findings indicate that models should accommodate for the individual traits of each patient. Yet, the clinical reality often proceeds at a pace that prohibits this from being attainable in a real time environment. Thus we present a design that counters these challenges through a staged model which allows for a gradual progression of accuracy as the system familiarizes itself with each patient. We have named this system the 3-Stage Patient Deterioration Warning System (3-Stage PDWS).

The paper first describes the settings and structure of an ED. We then exemplify related work in Section 3 to validate our approach. We present our data collection approach and the study methodology which sets the stage for the design of the 3-Stage PDWS. As this is the main contribution of the paper, we conclude with a discussion of the challenges we face and a description of planned future work to handle these.

2. THE SETTINGS

The ED capacities of the Danish healthcare sector have recently been restructured by merging multiple hospital entry points to a single point of entry. This meant closing emergency departments at minor hospitals and fusing the capacities of larger hospitals. All observations in this paper are based on a field study conducted at an ED belonging to a large university hospital in Denmark. This ED is organized into a receiving ward, known as the Emergency Treatment Center (ETC) and a bedside ward; Center for Accelerated Patient admissions (CAP). The ETC handles both medical and surgical illnesses which can be identified and treated in a day. If the treatment period requires hospitalization in less than 48 hours, the patient will be admitted to the CAP. All patients arriving at the ED with anything but minor injuries will be triaged upon arrival. The variety in patients in EDs poses challenges to define a single warning system to suit all patients [26]. At the ED of this study, the clinicians currently rely on the ADAPT triage model [24], which defines thresholds for each severity score and provides guidelines for how often registered observations are to be scheduled during the stay.

The ED relies on several information technology systems, but in this paper we solely consider the two systems directly involved with patient treatment and collaborative clinical work. The first of these is the Electronic Health Record (EHR) system which is used across all hospitals in the same region, and which through integration with lab systems aggregates multiple streams of information used in the diagnosis and treatment process. The second system is a patient logistic system used by the ED for organizing and visualizing the patient flow through the department, pending tasks and clinical responsibility delegation.

Vital signs monitoring in the ED utilize Philips IntelliVue MP30/50 monitors in a networked setup that enable clinicians to remotely monitor patients from ward offices. How much and at what frequency, a patient is monitored depends on clinical judgment based on the patient's triage level.

A distinctive trait of EDs is that the clinicians plan treatment of patients based on their presented history and symptoms. It is not based on an already established diagnosis as is often the case in other types of hospital departments. In a context where patient throughput is high, and a large part of clinical observations are tacit and thus seldom transferred consistently between shifts. The need for a shared representation which captures a patient's state, trajectory and clinically linked observations is a reality that to the best of our knowledge is not dealt with properly today.

During our field study we observed on multiple occasions nurses muting patient alarms without actually assessing the patient's state. This oversight of alarms was frequently based on assumptions about the patient, or the equipment's reliability. This is in line with similar causing factors for alarm fatigue such as a high number of false positives, usability issues, and faith in own knowledge [19]. The monitoring system issues alarms in stepwise degrees, and even though the most severe alarms still lead to increased levels of observation, inexperienced nurses may be affected by the overall tendency to dismiss non-critical alarms, and thus miss true adverse events in the long run.

As few EDs to our knowledge have the necessary staffing, and budget to integrate the latest generation of automated patient monitoring, we believe that there is a need to identify ways of improving deterioration detection by utilizing existing equipment. This pragmatic approach should be of interest to EDs worldwide.

3. RELATED WORK

In our review of existing related work, we have focused on studies that concentrated on integration into a clinical reality: work that attempts to integrate prospective data, real-time analysis, and an assessment of clinical feasibility. From these criteria, the research contributed by the Oxford Biomedical Research Center [23], [16] stands out. In this body of work, we find efforts to investigate the applicability of latent variable models that merge multiple streams of patient vital values into a model that is built upon machine learning techniques, with the intention of providing an intuitive visualization of patient state and trajectory.

The plausibility of building individual models of normality has been investigated [29]. Although the specific angle in this study is unfit for a large scale real-time system, the research still conceptually shows the possibility of detecting patient deterioration from dynamically created models.

In a study based on observational vital sign data, models of normality were built for a specific post-operative patient population based on three different metrics calculated from the vital sign distributions [17]. In the same study, the authors also found that the majority of observed vital sign types varied substantially from submission to admission.

Priming a clinical warning system by performing risk stratification based on EHR information to determine which patients were in need of continuous monitoring offers several advantages [8]. This vision has been elaborated upon by focusing on the challenges of doing time series analysis on streams of vital signs [14].

Although several contributions to this field have been made. most of the work has been done in parallel, and not in cooperation, with the targeted clinical context. Thus we are motivated to conduct the planning and execution of this project with the intent of providing a solution that strives to fit into the entirety of the problem domain.

4. METHODOLOGY AND DATA GATHERING

This paper is part of larger action-oriented research project which involves a field study, workshops and prototype-driven controlled experiments. As such we follow an action oriented research approach [4]. Consequently, we have participated in ED training courses, managed workshops, and helped plan new standard working procedures. In October 2013, we launched an ongoing automated gathering of vital sign data from patients admitted to the CAP. The registration of vital signs is approved by the Danish Data Protection Agency. Data is stored in a restricted access database in compliance with Danish legislation on privacy concerns.

The collected vital signs will be coupled with national Danish health registries to cluster all patients using categorical data such as past illnesses from ICD-10 codes, initial triage level, gender, admission package, number of prescribed medications, age, and seven day outcome in a retrospective analysis. The dataset will be segmented into event and non-event subgroups based on the occurrences of heart failures, ICU transfers, and in-hospital death. This retrospective dataset also forms the foundation for the training of the patient state models which we introduce in later sections.

4.1 Vital sign data collection

Vital values are harvested from the Philips IntelliVue patient monitors through a HL7 export interface. From this we receive HL7 Unsolicited Observation Reporting messages with patient vital signs from each bed in 60 second intervals. These messages are parsed and stored in a VitalSigns database. The HL7 messages carries information about arterial blood oxygen saturation (SpO2) and Pulse Rate (PR) measured through pulse oximetry, Respiration Rate (RR) and Heart Rate (HR) measured using 3-lead electrocardiography, mean, systolic, and diastolic blood pressure measured using a Noninvasive Blood Pressure (NBP) cuff. The actual types of vital signs registered for each individual patient depend on the level of criticality and overall mobility of the patient. Clinicians often adjust the frequency of NBP measurements to match the state of the patient, and consequently we register blood pressure measurements in intervals from five to sixty minutes. As pulse oximetry is the least obtrusive vital sign to monitor, SpO2 and PR are by far the most frequent observations in our dataset.

When a patient is received on the CAP ward, we asked the nurses to admit the patient to the Philips IntelliVue system by entering personal identification information such as name and social security number. This information was stored by our system in a Patient database table and coupled with the VitalSigns table to track the progression of patient state.

Table 1 summarizes the collection of vital values in the period from October 2013 to May 2014.

Table 1 Overview of vital sign registrations

Number of patients registered	4.071
Total number of registered vital signs	2.173.686
Heart Rate registrations	1.191.285
Respiratory Rate registrations	1.163.977
Pulse Rate registrations	1.887.869
SpO2 registrations	1.895.054
Blood Pressure registrations	64.756
Mean age male patients	62.7 ± 20.6 years
Mean age female patients	66.6 ± 21.9 years

The number of vital signs registered for each patient varies from a single measurement up to several thousand.

5. DESIGNING THE 3-STAGE PDWS

5.1 Guiding design principles

Our design principles are influenced by related work, clinical observations, and our own assumptions regarding what seems feasible fitting into the settings of the particular ED from the field study.

The ongoing data collection has been preliminarily evaluated to probe for support of our assumption that a granular model of normality would be an appropriate approach for the system. As an example, we have included Figure 1 and Figure 2 which exemplify the conceptual model of our system.

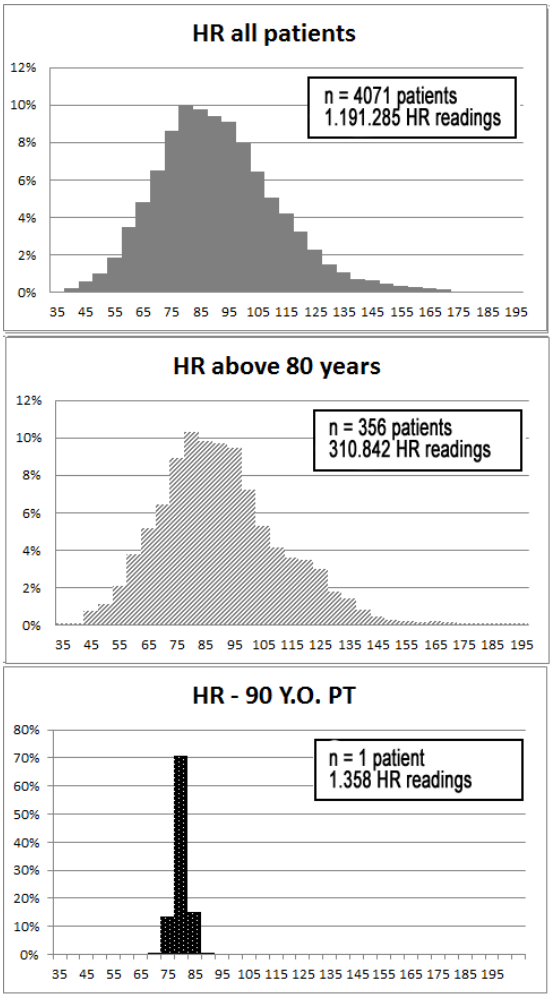

Figure 1. HR distribution for three samples

In Figure 1, all HR data points registered with age metadata were extracted and represented as Probability Mass Functions (PMF) for each group; one for all patients in the entire dataset, one for all patients above 80 years of age, and one for a randomly selected patient from this geriatric patient group. All three distributions were assessed using a two-sample Kolmogorov-Smirnov test in R. All of the samples above where shown to be significantly different. Even so, age is just one of many possible classification features. This is backed by Figure 2 that depicts the distributions of 6.000 randomly sampled heart rates of patients assigned to four different specialties; J for pulmonary patients, N for neurological patients, Q for infectious disease patients, and S for gastroenterological patients.

Although Figure 2 displays distinct differences, specialty department as such is not a sufficiently accurate classification feature. Although the assignment to a specific specialty tells us something about the current admission reason, it provides less insight into the actual traits of the patient. Later investigations will include selected ICD-10 codes and are expected to yield more conclusive insight.

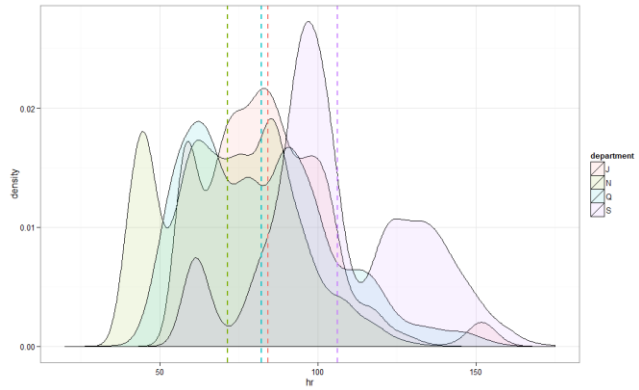

Figure 2. HR distribution for different specialties

The groupings provided by Figure 2 also depict the mean of each group as a vertical dashed line, and show a much more distinct difference than what we have found by looking at grouping by age alone.

Figure 1 and Figure 2 also depict our attempt to probe the validity of establishing a staged normality model rather than attempt to specify a specific model. The actual process of defining and refining such a model requires access to much higher dimensional data-sets. These will be obtained through the coupling of the VitalSigns database to the Danish national health registry. Thus, the specific model features are currently being selected through our cooperation with ED researchers and clinicians which will simultaneously allow us to attempt a very open ended approach, and a clinical a-priori strategy.

Another key assumption which we have yet to sufficiently validate is the model of patient trajectory that is shown in Figure 3.

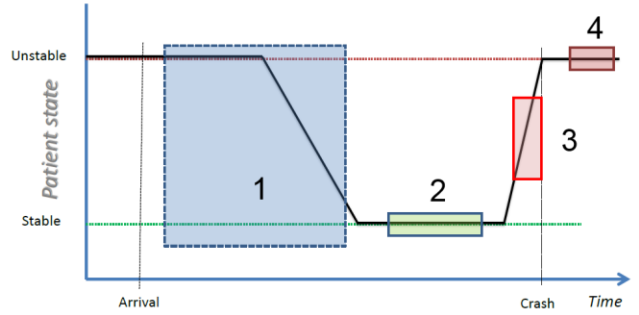

Figure 3. Progression of patients

Figure 3 illustrate the assumption that patients arriving at the ED are classifiable as unstable as they are in a health related imbalance compared to their normal state of being. As clinicians initiate the diagnosis and treatment process during hospitalization, the patient is assumed to gradually stabilize, as illustrated by box 1 in Figure 3. Some patients are expected to stabilize during treatment, which marks the period where the 3-Stage PDWS trains to detect deterioration from personalized stabilization, illustrated as box 2 in Figure 3. This individualized approach is expected to enable earlier realization of the cases where a patient departs from individual stability, illustrated as box 3 in Figure 3. Automated recognition of deterioration is expected to precede human clinical observation of the same deterioration, which is shown as box 4 in Figure 3. Even if the discriminatory ability of the system is poor, it might still provide warning of ongoing deterioration at an earlier stage as the automated system continuously evaluates the patients, whereas clinicians are confined to spot observations due to a busy schedule.

The patient trajectory assumption will be tested by analyzing the vital signs time series data collected from patients to check for any significant progression of vital signs during the admission period of the patients. Similar investigations have been carried out and found that the initial trend of a patient's state is correlated to in-hospital death [12].

5.2 Model architecture

The conceptual model in Figure 4 depicts each of the three stages in our proposed system:

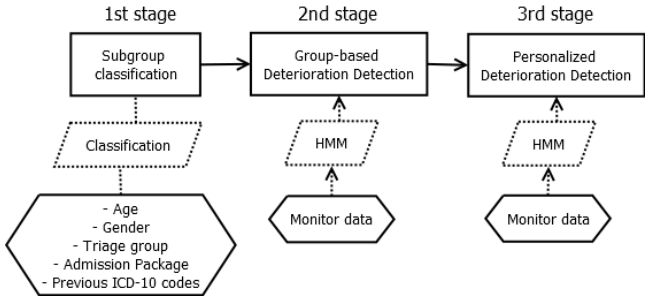

Figure 4. The 3-Stage PDWS model

Each stage in Figure 4 serves a particular purpose:

1. *Subgroup Classification:* each patient will be classified based on arrival parameters which are already being collected by the triage nurse, or easily accessible information from the patient's EHR. Based on these parameters, the patient is classified according to the segments identified in the clustering phase of our model definition process.
2. *Group-based Deterioration Detection:* when a patient is classified as belonging to a given group, all received vital values from the patient will be assessed according to what is identified as normality for this group of patients.
3. *Personalized Deterioration Detection:* given the assumption that some patients have models of normality that differ from any group, the system will evaluate if the given patient seem to be in a stable deviation from the model of normality under which the patient is currently monitored. If so, the patient should be monitored according to an individual model.

Although we are still evaluating machine learning techniques, using Hidden Markov Models (HMM) as a modelling approach for patient state transitions is interesting because the properties and traits of HMMs resemble the clinical reality found in ED's. Namely, that clinicians monitor a set of vital sign observations from which they seek to deduct the actual state of a patient. This is in line with the hidden state nature of HMMs [18]. Although the clinicians operate with a multitude of observation channels, the clinical assessment is essentially still a process of uncertainty and interpretation of the hidden, actual state of the patient.

HMMs have been used to model clinical relevant situations such as real-time daily activity monitoring [25] and hepatitis C disease progression [22]. Although [20] investigated the development of a multi-state Kalman filter algorithm for patient monitoring, and [7]

proposed a method for clustering multivariate time series of both numerical and categorical features in healthcare, our survey of published research indicated that HMMs have not been investigated for modelling patients in an ED context. This is likely related to the Markov property of HMMs, which states that the current state is independent of all prior states. Further investigation is needed to determine if relaxation of the Markov property is needed in order to detect patient deterioration using HMMs.

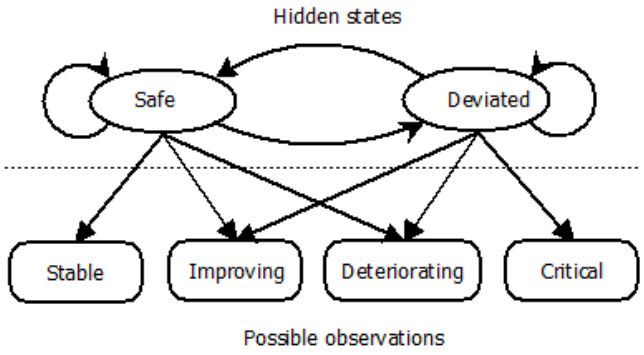

Figure 5 Proposed HMM

Through our observations of patients, and discussions with clinicians, it is evident that it is difficult to quantify the state of a patient. Instead the trajectory of a patient is often mentioned as a noticeable registration by clinicians, which raises the question if dynamic changes in patients can be used to identify patients at risk [12]. Hence, Figure 5 conveys our proposal for the states of the HMM and its transition relationships. The observable states of the HMM illustrate how patients are considered being either stable, or being in a direction of either towards or away from stability (improving/deteriorating), and can be considered critical if sufficiently far away from being stable. The hidden states of the HMM are the unknown actual state of the patient who is either in a safe zone or deviating from normality.

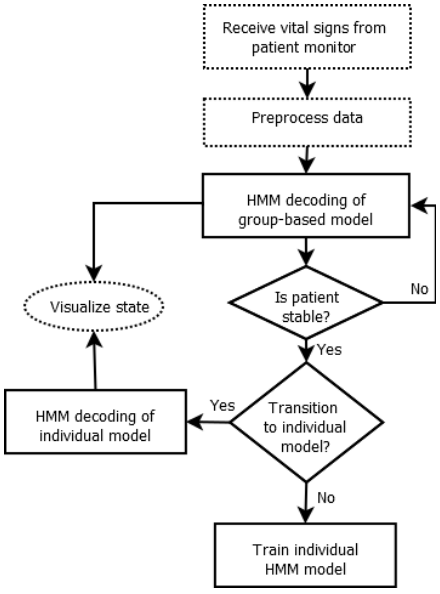

Figure 6. HMMs for groups and individuals.

The envisioned flow of the second and third stage is illustrated in Figure 6, which shows the parallel deterioration detection, training of the individual model and concurrent visualization of patient state.

The decision threshold between transitioning from group-based normality to individual normality will be investigated using distance measures between the two HMMs [28] and [13].

5.3 Model validation

The 3-Stage PDWS model has currently only been conceptually validated by its composition from the body of existing published research, preliminary data analysis, and clinical observations from the field study, and in collaboration with nurses and physicians at the ED.

The exact number of clusters found through the investigation of patient characteristics, is still unknown. However, we intend to validate the clusters by relying on the judgment of experienced ED physicians who will review the similarity measures of each cluster.

The final model and its implementation needs validation in two dimensions: a retrospective evaluation of its accuracy in identifying patients at risk of deterioration, and in its ability to convey the patient state in a way that makes sense to different clinical professions.

Each stage of the PDWS will be assessed individually in the retrospective validation. The accuracy of the first stage will be reviewed through the ability of our subgroup classification to produce similar labels for unobserved patients as found by the physicians. The 2nd and 3rd stages will be evaluated based on the ability to accurately classify patients as deteriorating. In this regard we are interested in model accuracy and the ability to predict deterioration onset earlier than currently possible by the existing alarming thresholds.

The clinical utility of our model will be assessed by comparing the misclassification rate of our system with the generic thresholds used at the specific ED in this study.

6. DISCUSSION

Automating patient deterioration detection can be approached from multiple entry points. The first challenge is gathering the vital signs in an unobtrusive way that does not enforce a potentially unjustified sense of illness on the patient, and which does not hinder the workflow of clinicians, and treatment trajectories of patients. Our approach is pragmatic in the sense that we seek to design and build a solution that utilizes the existing equipment at the ED. We have found that the clinicians are prone to not attach the most cumbersome sensors to patients who are scheduled for frequent tests outside the ward. Mobile monitoring technology would help overcome this obstacle, but is outside the scope of our current research approach.

Our approach to subgroup classification resembles that of [30], who investigated the feasibility of using decision trees and probabilistic algorithms for classification of patients into severity levels similar to the clinical triage classification. However, our intent is not to replace existing severity indexes but instead to improve the accuracy of vital sign monitoring by deploying increasingly specific thresholds.

Although the data we are currently collecting only consist of a few dimensions, we are challenged by commonly found problems such as variation in what vital signs are measured and occasional holes in the time series. This issue has been dealt with by replacing the

missing values with either the last registered measurement or with the mean of the vital sign over the entire historical dataset [14]. An alternative approach is to utilize Gaussian Processes which have proved useful in predicting the distribution of missing physiological data [3]. Additionally, we face a sizeable task in ensuring the validity of the vital signs by having to check that the data series can be linked to a given patient in a reliable manner. To ensure this we hope to couple our VitalSigns database with the EDs internal logistic system. This provides accurate information about which patients resided in each bed at a given point in time.

Modifying the harvesting and registration of vital sign data using existing equipment by asking clinicians to revise their standard working practices, has unsurprisingly proven difficult. The staff group as a whole recognizes the importance of registering vital values. But the clinical reality is such that if a system does not yield immediate and tangible benefits, the perception of added utility is generally low, causing the clinicians to abstain from integrating new admission procedures. In our case this is admitting patients to the Philips IntelliVue system, but we find that the concept of clinical utility and the interplay between healthcare organization and health information technology is a topic worthy of further exploration. This interplay relates to the notion of "meaningful use" [1], and the SUMMIT method for modelling the meaningful use of an IT tool as a function of its overall utility [9]. It seems feasible to deploy this framework in the planned controlled experiments with nurses and physicians to structure the evaluation of how the 3-Stage PDWS represents and visualizes knowledge and information about patients. Other research has pointed out a clinical skepticism towards black-box expert systems. Consequently, the aim of our system is to support decision making rather than replace it.

Finally, while our initial approach will rely on datasets tightly coupled to the Danish healthcare system, the core model assumptions are applicable globally. Although the specialty department features proposed in Figure 2 may be of little meaning to other healthcare systems, we expect that the addition of ICD-10 features will provide both interesting insight for clinical researchers and practitioners. The challenges of patient variation is known to all EDs and thus we believe that the 3-Stage PDWS can be of use wherever it is possible to classify arriving patients.

Another aspect is the availability of vital sign data. Several patient monitors already support exporting vital sign data, and as such our solution is independent of particular equipment. As monitoring platforms are becoming increasingly unobtrusive, a wider spectrum of patients can be included in continuous monitoring. This expansion into a broader part of the patient population further justifies building more specific deterioration detection models.

7. FUTURE WORK AND CONCLUSION

The overall goal is to improve the detection of deteriorating patients by identifying the onset of adverse events earlier and to embed this detection ubiquitously into clinical practices by assuming a holistic approach to the integration of patient monitoring. If the system proves successful, we expect to see a reduction in the rate of seven-day mortality of patients and in increased clinical utility.

The intent of our research is to target the solution domain as a whole and not explicitly focusing on particular parts, e.g. providing a revolutionary real-time analysis model, conceiving new machine learning techniques, or developing new monitoring platforms. We expect to draw out more systemic findings which can support more depth-oriented research approaches.

The system is currently under development, as we have undertaken initial analysis of the vital signs, and how to utilize these with machine learning techniques that are sensible to clinicians. Coupling of the collected vital signs with the national Danish health registries is planned for Q3 2014, and we expect the prototype to be ready for initial clinical controlled experiments by early 2015. Fine tuning of the predictive capabilities of the 3-Stage PDWS is planned for Q2 in 2015.

8. REFERENCES

[1] Blumenthal, D. and Tavenner, M. 2010. The "Meaningful Use" Regulation for Electronic Health Records. *New England Journal of Medicine*. 363, 6 (2010), 501–504.

[2] Brabrand, M., Folkestad, L., Clausen, N.G., Knudsen, T. and Hallas, J. 2010. Risk scoring systems for adults admitted to the emergency department: a systematic review. *Scandinavian Journal of Yrauma, Tesuscitation and Emergency Medicine*. 18, (Jan. 2010), 8.

[3] Clifton, L., Clifton, D.A., Pimentel, M.A.F., Watkinson, P.J. and Tarassenko, L. 2012. Gaussian process regression in vital-sign early warning systems. *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. x (Aug. 2012), 6161–6164.

[4] Easterbrook, S., Singer, J., Storey, M.-A. and Damian, D. 2008. Selecting empirical methods for software engineering research. *Guide to Advanced Empirical Software Engineering*. F. Shull, J. Singer, and D.I.K. Sjøberg, eds. Springer London. 285–311.

[5] Gao, H., McDonnell, A., Harrison, D. a, Moore, T., Adam, S., Daly, K., Esmonde, L., Goldhill, D.R., Parry, G.J., Rashidian, A., Subbe, C.P. and Harvey, S. 2007. Systematic review and evaluation of physiological track and trigger warning systems for identifying at-risk patients on the ward. *Intensive Care Medicine*. 33, 4 (Apr. 2007), 667–79.

[6] Geier, F., Popp, S., Greve, Y., Achterberg, A., Glöckner, E., Ziegler, R., Heppner, H.J., Mang, H. and Christ, M. 2013. Severity illness scoring systems for early identification and prediction of in-hospital mortality in patients with suspected sepsis presenting to the emergency department. *Wiener klinische Wochenschrift*. 125, 17-18 (Sep. 2013), 508–515.

[7] Ghassempour, S., Girosi, F. and Maeder, A. 2014. Clustering multivariate time series using Hidden Markov Models. *InternationalJjournal of Environmental Research and Public Health*. 11, 3 (Mar. 2014), 2741–63.

[8] Hackmann, G., Chen, M., Chipara, O., Lu, C., Chen, Y., Kollef, M. and Bailey, T.C. 2011. Toward a two-tier clinical warning system for hospitalized patients. *AMIA Annual Symposium proceedings*. American Medical Informatics Association 2011, (Jan. 2011), 511–9.

[9] Haynes, S.R., Winkler, T.E. and Ritter, F.E. 2013. Modeling Meaningful Use as Utility in Emergency Medical

Services. *2013 IEEE International Conference on Healthcare Informatics*. (Sep. 2013), 455–460.

[10] Hayward, R.A. and Hofer, T.P. 2001. Estimating hospital deaths due to medical errors: preventability is in the eye of the reviewer. *The Journal of the American Medical Association*. 286, 4 (2001), 415–420.

[11] Henriksen, D.P., Brabrand, M. and Lassen, A.T. 2014. Prognosis and risk factors for deterioration in patients admitted to a medical emergency department. *PloS one*. 9, 4 (Jan. 2014), e94649.

[12] Kellett, J., Wang, F., Woodworth, S. and Huang, W. 2013. Changes and their prognostic implications in the abbreviated VitalPAC™ Early Warning Score (ViEWS) after admission to hospital of 18,827 surgical patients. *Resuscitation*. 84, 4 (Apr. 2013), 471–6.

[13] Lyngsø, R.B., Pedersen, C.N. and Nielsen, H. 1999. Metrics and similarity measures for hidden Markov models. *Proc. Int. Conf. Intell. Syst. Mol. Biol* (Jan. 1999), 178–86.

[14] Mao, Y., Chen, Y. and Hackmann, G. 2011. Medical Data Mining for Early Deterioration Warning in General Hospital Wards. *Proceedings of the 2011 IEEE 11th International Conference on Data Mining Workshops*. IEEE Computer Society. 1042–1049.

[15] Mcgaughey, J., Alderdice, F., Fowler, R., Kapila, A., Mayhew, A. and Moutray, M. 2007. Outreach and Early Warning Systems (EWS) for the prevention of Intensive Care admission and death of critically ill adult patients on general hospital wards (Review). *Cochrane Database System Rev*. 3, (2007).

[16] Orphanidou, C., Clifton, D., Khan, S., Smith, M., Feldmar, J. and Tarassenko, L. 2009. Telemetry-based vital sign monitoring for ambulatory hospital patients. *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE* (Jan. 2009), 4650–3.

[17] Pimentel, M.A.F., Clifton, D.A., Clifton, L., Watkinson, P.J. and Tarassenko, L. 2013. Modelling physiological deterioration in post-operative patient vital-sign data. *Medical & Biological Engineering & Computing*. 51, 8 (Aug. 2013), 869–77.

[18] Rabiner, L. 1989. A Tutorial on Hidden Markov Models and Selected Applications in Speed Recognition. *Proceedings of the IEEE*. 77, 2 (1989), 257–286.

[19] Sijs, H. Van Der, Aarts, J., Vulto, A. and Berg, M. 2006. Overriding of drug safety alerts in computerized physician order entry. *Journal of the American Medical Informatics Association*. 13, 2 (2006), 138–147.

[20] Sittig, D.F. and Factor, M. 1990. Physiologic trend detection and artifact rejection: a parallel implementation of a multi-state Kalman filtering algorithm. *Computer Methods and Programs in Biomedicine*. 31, 1 (Jan. 1990), 1–10.

[21] Subbe, C.P. and Welch, J.R. 2013. Failure to rescue: using rapid response systems to improve care of the deteriorating patient in hospital. *Clinical Risk*. 19, 1 (Jan. 2013), 6–11.

[22] Sweeting, M.J., Farewell, V.T. and De Angelis, D. 2010. Multi-state Markov models for disease progression in the presence of informative examination times: an application to hepatitis C. *Statistics in Medicine*. 29, 11 (May 2010), 1161–74.

[23] Tarassenko, L., Hann, a and Young, D. 2006. Integrated monitoring and analysis for early warning of patient deterioration. *British Journal of Anaesthesia*. 97, 1 (Jul. 2006), 64–8.

[24] Triage-Manual: 2009. *http://www.hillerodhospital.dk/NR/rdonlyres/D20F6C68-ABB6-402D-B463-C7293185C372/0/Triagemaster.pdf*.

[25] Wei, H., He, J. and Tan, J. 2011. Layered hidden Markov models for real-time daily activity monitoring using body sensor networks. *Knowledge and Information Systems*. 29, 2 (Jun. 2011), 479–494.

[26] Windle, J. and Williams, J. 2009. Early warning scores: are they needed in emergency care? *Emergency nurse : the journal of the RCN Accident and Emergency Nursing Association*. 17, 2 (May 2009), 22–6.

[27] Zegers, M., de Bruijne, M.C., Wagner, C., Hoonhout, L.H.F., Waaijman, R., Smits, M., Hout, F. a G., Zwaan, L., Christiaans-Dingelhoff, I., Timmermans, D.R.M., Groenewegen, P.P. and van der Wal, G. 2009. Adverse events and potentially preventable deaths in Dutch hospitals: results of a retrospective patient record review study. *Quality & Safety in Health Care*. 18, 4 (Aug. 2009), 297–302.

[28] Zeng, J., Duan, J. and Wu, C. 2010. A new distance measure for hidden Markov models. *Expert Systems with Applications*. 37, 2 (Mar. 2010), 1550–1555.

[29] Zhang, Y., Silvers, C.T. and Randolph, A.G. 2007. Real-time evaluation of patient monitoring algorithms for critical care at the bedside. *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE* (Jan. 2007), 2783–6.

[30] Zmiri, D., Shahar, Y. and Taieb-Maimon, M. 2012. Classification of patients by severity grades during triage in the emergency department using data mining methods. *Journal of Evaluation in Clinical Practice*. 18, 2 (Apr. 2012), 378–88.