ICU MANAGEMENT & PRACTICE

ICU Volume 13 - Issue 1 - Spring 2013 - Matrix

Racing to Improve Early Warning

The unique challenges in developing early warning for children have led to a patient-specific early warning approach. The methodology, which is based on the altered patterns of physiological derangement, associated with compensation and decompensation in clinical deterioration rather than population normal distributions, is suitable for patients of all ages. A software platform that is used for Formula One telemetry, which is able to analyse continuous data in real-time and produce predictive models for the future, has been adapted for use in critically ill patients. This enables real-time principal component analysis (PCA) and predictive modelling, which are promising solutions for developmental physiological changes and patient specific variations, whilst avoiding false alarms.

Introduction

Formula One motor racing is all about winning. Healthcare is all about reducing distress and saving lives. Both are most successful when underpinned by detailed knowledge of the context and behaviour involved as well as the drive for safety and quality. In addition, they both require prediction of likely outcomes, pre-planning to mitigate anticipated complications and rapid, real-time decision-making to achieve the best results.

There are, of course, many differences between managing a racing car and managing a sick patient. The most obvious is that while racing cars are relatively predictable, patients have homeostasis with known, but often unpredictable, closed loop feedback systems that aim to compensate for physiological derangement until a point where decompensation is inevitable. The physiological patterns of deterioration prior to respiratory or cardiac arrest are well described in adults and children and underpin the global development of early warning systems (Devita et al. 2006; National Institute for Health and Clinical Excellence 2011; NHS 2011). These systems involve an aggregation of physiological derangement to identify the trend of deterioration and have been shown to reduce avoidable death and life-threatening events. The UK National Early Warning Score (NEWS) is an aggregate scoring system to identify adult patients during the compensation phase before they decompensate to cardiac arrest (Royal College of Physicians 2012). Identification of deterioration triggers a response to generate increased monitoring, management of the cause of decline and a call for more expert help.

Children provide unique challenges in early warning that have led to developing patient specific monitoring. The wide range of “normal” physiology which develops as children grow, and the complex “at risk” patients that have their own specific range of normality, for example pulse oximetry of 75% in cyanotic heart disease, has provided the drive to develop early warning that learns the normal and abnormal physiological patterns for that patient. Once successfully developed, this age-agnostic approach will almost certainly be equally applicable to adult and geriatric patients. We describe an approach where the patterns generated by individual susceptibility, acute physiological response, compensation and decompensation thresholds create early warning.

Early Warning Systems in Adults and Children

Early warning scores are substantially better than relying on chance identification of deterioration, achieving a 25% reduction in life-threatening illness and death (Priestley et al. 2004; Tibballs et al. 2009). However, there are two good reasons why we need to explore beyond the current aggregate systems.

Firstly, current early warning systems rely on categorising a patient as normal or increasingly abnormal in comparison to a population normal distribution. Patients have individual normal physiology and resilience and responses to illness are based on their background health or illness, medication and age (Bion 2000; Smith et al. 2008). We need to know what is normal for that specific patient. A well 10-year-old, an immunocompromised 20-year-old with leukaemia and an elderly 80-year-old patient with heart failure will each respond completely differently to abdominal sepsis. Current early warning systems are not tailored to this background-dependent resilience or susceptibility but focus only on the acute, generic physiological changes. In addition, some patients deteriorate and decompensate within the normal range for their age and for them. This is because the effects of deterioration aren’t limited to changes in high and low thresholds, but are related to the pattern of variation within physiological parameters. The patient shows a clear change in pattern of variability prior to cardiac arrest, but would not have triggered high or low alarm threshold until the acute life-threatening event. Some of this altered pattern is measured in heart rate variability: a marker that has benefit in identifying sepsis in adults, neonates and foetal distress (Ahmad et al. 2009; Moorman et al. 2011; Van Laar et al . 2008).

Secondly, current early warning systems rely on relatively infrequent (every one to 12 hours) and varied vital sign measurements, from a choice of up to 36 parameters. Furthermore, the warning can in some cases be associated with missing measurements (Royal College of Physicians 2012; Duncan 2007; Royal College of Nursing 2011). There is no information on the optimal frequency of observations for a specific patient or population and this decision is often left to relatively inexperienced, busy bedside nurses and healthcare assistants. The assumption is that the bedside staff will recognise a deteriorating change in trends and alter the observation frequency, but this opportunity is often missed. There is also variation in which vital signs should be measured.

In 2008, a paediatric early warning system (PEWS) was introduced in all wards at Birmingham Children’s Hospital (Duncan et al. 2006; Parshuram et al. 2011). This is a paper-based aggregate score embedded as colour-coded, age-dependent thresholds on the four standardised observation charts. It is associated with simulation-based training on taking routine observations and appropriate decision-making relating to deterioration. The type and frequency of observations are guided by a comprehensive evidence- and expert-based observation, monitoring and escalation policy. All life-threatening events are tracked, in keeping with international recommendations (Devita et al. 2006), and they are classified into timely and untimely for intensive care referral and admission, and whether or not they are predictable and/or potentially preventable for acute life threatening events. It is this detailed forensic review of all episodes of critical deterioration that has provided the insight into how best to approach early warning.

Since the introduction of early warning systems, in-hospital cardiac arrests have reduced and more patients are receiving optimal pre-intensive care. These are direct indicators of more timely treatment of acute illness. But it could be better: measurements and observations could be more frequent, processing the data could be automated, data entry mistakes could be avoided and warnings or alarms could be tailored more specifically to individual patients.

Birmingham Children's Hospital cares for children from birth to 16 years old, with weights ranging from 450g to 120kg. Our patients are frequently complex with cyanotic heart disease, chronic lung disease, neuro developmental disorders, multi- organ involvement and they epitomise individuality. Four age-appropriate observation charts are needed to accommodate ab/normal physiological parameters for the age-groups: birth to one year, one to five years, five to 12 years and older than 12 years. Infants, in particular, and older children can deteriorate very quickly; in between infrequent observations. These situations can erroneously be interpreted as unpredictable; however, parents are often adamant that a change had occurred in the child's condition prior to an acute lifethreatening event that routine or even enhanced monitoring did not detect. These challenges have led to our exploration of Real-time continuous, Adaptive patientspecific, Predictive Indicators of Deterioration (RAPID).

How Does Formula One Help Solve These Problems?

Based on the problems identified so far, we determined that what is required is a system with the following requirements:

• Real-time analysis - to identify changes in patterns of physiological compensation and decompensation, and predict or form a model for the future;

• Adaptive - to have real-time analytical ability to learn normal for that specific patient;

• Continuous - until an optimal observation frequency can be determined; and

• Scalable and not reliant on expensive individual monitors - to measure as many at risk patients as possible.

A solution is not yet available for medical monitoring, but is routinely used in motor racing telemetry. You need a fast car, great driver and good strategy to win races in the complex and highly competitive world of Formula One. It is human endeavour at its most extreme, characterised by relentless development with new innovations appearing throughout the year in everchanging forms. Each must be anticipated, its influence evaluated, and then put into action quickly.

The cars are changing continually to make them faster, stronger and safer, and they go into intense competition every two weeks between March and November. Hence, it is unsurprising that the world of Formula One is underpinned by data. Quickly making sense of what you see and hear is often the difference between winning and losing. In this respect, healthcare is little different. Recognising problems quickly is the first step towards effective treatment, but each patient is different and early signs can be subtle and complex. Nonetheless, they are usually there to be found in the data. Recognising deterioration early provides a real opportunity for reducing distress and saving lives.

McLaren Electronics Systems provides telemetry for all Formula One teams so they can measure, visualise and respond to changes during development of the cars as well as during the time-critical race situation. The Formula One realtime data system comprises SQL-Race, an application processing interface (API) that manages a large population of individual sources of time-series and associated data; vTAG server, a data logging and processing platform upon which real-time models run; and ATLAS, the data analysis and viewing software used by teams and engine makers throughout Formula One.

Each car is fitted with over a hundred sensors. Live health and performance data is sent back via telemetry to the garage and over the Internet to the team's factory, often on the other side of the world. Over 750 million numbers from each car are processed in real-time during a two-hour race. Over the race weekend the data is used to make the cars better and faster. The data tells the engineers how much life remains in the engine, how quickly the tyres are degrading and how much fuel is being used (as well as how much is left in the tank). The data tells them whether setup changes are effective or not, and the system has the ability to run thousands of models simultaneously to predict the consequences of different treatment strategies.

In healthcare, it is not feasible to have the equivalent of a Formula One team's engineers focusing on just two patients, but it is possible to use the data platform to analyse patient-specific data in realtime, and to predict the future. If such analysis of changing physiology and variation could be visible to bedside or remote clinicians, then a much higher incidence of subtle signs of compensation and decompensation could trigger more sophisticated alerts, and we could see the predicted consequences of treatment and observation strategies.

Saving Young Lives

In 2011, Birmingham Children's Hospital and McLaren Electronic Systems installed a real-time data system to gather and process live physiological data from all beds in paediatric intensive care and from the trolley in one of the specialist child transport ambulances. Through the "Young Lives" project, supported by the Health Foundation SHINE programme and applied mathematics academics from Aston University, we developed a system that would stream data from all of the bedside monitors and quickly tease out patterns in the data, with a purpose of alerting doctors and nurses to changing conditions.

The reason for starting in paediatric intensive care was twofold: it is where the sickest children are treated with 1:1 nursing and it is where the physiological data was already routinely collected (but previously overwritten after 96 hours). In the first twelve months of running the system, we collected physiological data from more than 1000 different patients. By streaming the data into the Formula One data system, we have been able to provide a richer display and manipulation of data and store it longer for the purposes of clinical review and research.

The bedside instruments provide data from a range of sensors, but initial focus was placed on pulse oximetry (SpO_2) because it is readily measured and is rich in information about respiration and cardiac activity. The real-time data platform can gather and process data from a large population of individual patients. The data processing can be applied to any of the physiological sensors and uses principal component analysis (PCA) to extract the characteristic patterns from the data as it changes with time. This technique is used for analysis and prediction in financial, environmental, military and aeroplane engineering applications. A patient who is stable exhibits patterns that change little over time. Plotting two principal components against each other creates a model "distance". Deterioration is reflected in an increase in the model "distance", a parameter which characterises how well the principal components correlate with the evolving data.

The PCA approach not only teases out characteristic patterns, but also provides the means to predict how the data should look in the future. It does this by extrapolating and then reconstructing the physiological data for a later time. Currently, we predict about two minutes ahead. The importance of the prediction is that it enables quite tight margins to be applied in testing for divergence from normality. This can lead to much earlier reliable detection of change for individual patients.

We are testing the PCA model distance alongside an automated version of a modified paediatric early warning score (mPEWS). Early indications show that changing conditions are apparent in the PCA distance and scatter plots well before the mPEWS or raw data are seen to change. Further clinical interpretation is needed before changes may be characterised in terms of deterioration.

What Does the Future Hold?

Formula One has been using telemetry data to develop and race cars for over 25 years. The engineers and drivers believe and act upon the information they see, using it to understand and continually improve their cars and race-craft. Analytical techniques and fidelity checking between parameters has managed false alarms out of the system. Much of the work done in setting up the race car and developing a winning strategy takes place away from the track using live data sent across the world via standard fibre and wireless networks. It is no longer always necessary for the engineer and car to be in the same location in order to make a difference.

However, exploiting this approach in healthcare involves more than simply transferring technology. The healthcare environment is less structured, people can be more complicated and less predictable, clinical interventions can be frequent and varied and the culture in secondary (and primary) health is not always one that is immediately receptive to change. The next stages of development at Birmingham Children's Hospital will be to:

(1) Establish more rigorous clinical interpretation of the changing patterns;

(2) Ensure that false detection of deterioration cannot happen;
(3) Move the system beyond the walls of intensive care and into the high dependency and general wards through-out the hospital; and
(4) Create new patient pathways and resourcing models that make use of the better clinical cues.

A lot has been achieved, but there is much more to do (Nangalia et al. 2010; Bion 2008). Embedding knowledge into the system of what constitutes normality, how characteristic changes in patterns relate to treatment and outcomes, and how alarm thresholds could and should be set, will all come with detailed clinical scrutiny of the data. Properly engineered, our approach will present physiological data clearly, immediately and in context, so that every patient, regardless of age, might be seen by the right people, in the right place and at the right time. Ideally, the changing conditions of the population of individual patients would help inform the most appropriate allocation of nurses throughout the hospital and direct doctors and other clinicians to the sickest patients. There is no reason, however, why an approach like RAPID should be confined to the hospital. Once developed, the applications that detect deterioration could operate remotely or be embedded in local devices, such as smart phones or tablets. Patients with acute and chronic conditions could be monitored at home with the reassurance that expert help could be informed quickly should a condition suddenly worsen.

Acknowledgements

The real-time Principal Component Analysis application that is used to extract and display changing patterns was developed by Dr. Rajeswari Matam of Birmingham Children's Hospital. She also liaised directly with the software and support engineers from McLaren Electronic Systems throughout the development and commissioning of the system. The development of the analysis approach, and ensuring that it was both practical and relevant to this new application, was supported strongly by Prof. David Lowe of Aston University

Published on : Thu, 14 Mar 2013