


RESEARCH ARTICLE

Open Access



Implementing a system for the real-time risk assessment of patients considered for intensive care

Simarjot S. Dahella¹, James S. Briggs^{1*} , Paul Coombes², Nazli Farajidavar³, Paul Meredith⁴, Timothy Bonnici⁵, Julie L. Darbyshire⁶ and Peter J. Watkinson⁶

Abstract

Background: Delay in identifying deterioration in hospitalised patients is associated with delayed admission to an intensive care unit (ICU) and poor outcomes. For the HAVEN project (HICF ref.: HICF-R9–524), we have developed a mathematical model that identifies deterioration in hospitalised patients in real time and facilitates the intervention of an ICU outreach team. This paper describes the system that has been designed to implement the model. We have used innovative technologies such as Portable Format for Analytics (PFA) and Open Services Gateway initiative (OSGi) to define the predictive statistical model and implement the system respectively for greater configurability, reliability, and availability.

Results: The HAVEN system has been deployed as part of a research project in the Oxford University Hospitals NHS Foundation Trust. The system has so far processed > 164,000 vital signs observations and > 68,000 laboratory results for > 12,500 patients and the algorithm generated score is being evaluated to review patients who are under consideration for transfer to ICU. No clinical decisions are being made based on output from the system. The HAVEN score has been computed using a PFA model for all these patients. The intent is that this score will be displayed on a graphical user interface for clinician review and response.

Conclusions: The system uses a configurable PFA model to compute the HAVEN score which makes the system easily upgradable in terms of enhancing systems' predictive capability. Further system enhancements are planned to handle new data sources and additional management screens.

Keywords: Early warning score, Vital signs, Pathology data, Portable format for analytics, PFA, Human-computer interaction, HCI, Information visualisation, Clinical decision making, OSGi

Background

Hospital inpatients whose condition deteriorates are often transferred from a general ward to an Intensive Care Unit (ICU) in order for them to receive a higher level of care. Such unplanned ICU admissions from within hospital typically make up over half of the total ICU admissions [1]. Unplanned ICU admission is associated with a poor patient

outcome, but timely admission may improve that outcome [2]. It is imperative therefore to identify patients who would need ICU admission in time to avert an unplanned admission. Unplanned ICU admission can be effectively averted by using a system that has a good predictive model and stimulates situational awareness of the relevant clinicians by visualisation. Visualisation can facilitate good clinical decision-making processes [3].

To facilitate early recognition of deteriorating patients and to avoid unplanned ICU admission, we have designed and developed the "Hospital alerting via electronic

* Correspondence: jim.briggs@port.ac.uk

¹Centre for Healthcare Modelling and Informatics, University of Portsmouth, Buckingham Building, Lion Terrace, Portsmouth PO1 3HE, UK

Full list of author information is available at the end of the article



© The Author(s). 2020 **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

noticeboard (HAVEN)” system. HAVEN is a decision support system that can be used hospital-wide. It uses a configurable predictive model and a user-defined graphical user interface to display details of patients at risk of unplanned admission to the ICU. This paper describes the architecture and technologies that we used to build the system.

The National Institute for Health and Care Excellence (NICE) has recommended the use of multiple-parameter or aggregate weighted scoring systems for monitoring acutely ill patients since 2007 [4]. Much research has been done in creating models for computing the early warning scores [5–7] and comparing the different models that compute the scores [8, 9]. It has been reported by NHS Digital, however, that recorded critical care periods in English hospitals have increased by 25% from 2010/11 to 2015/16 [1]. In 2015/16, more than half of these ICU admissions were ‘unplanned’ [10, 11]. Delays in admitting ward patients to ICU are common, with high short-term mortality because of the difficulty in triaging deteriorating patients in wards [12]. This is a global problem [13].

Another important aspect of outcome prediction in healthcare systems is the predictive model itself. Traditionally, a predictive model is static and embedded in the system [14]. Changing or updating the model in the system is usually a complicated task. This may mean that models are not updated as often as they should be. Hence, clinicians end up using systems that are not optimal.

We used innovative software technologies to enable rapid system updates. This is essential when the predictive model is frequently improved in its means of identifying accurately the sickest patients who will derive most benefit from increased care. The technologies used and details of the system build are described in the following sections.

Implementation architecture

We explored technologies for developing the system in a modular way that can be easily deployed and dynamically updated. Additionally, we evaluated the techniques for defining and executing the predictive model. The model needed to be configurable so that it can be updated dynamically when system capabilities are enhanced.

HAVEN has been implemented as a web-based application. It is therefore accessible in any part of the hospital using any machine with an internet connection without installing any additional software beyond a standard web browser.

The technology stack we decided to use for the application implementation consists of the Java programming language [15], MySQL database [16], Open Services Gateway initiative (OSGi) Framework [17], PFA model [18], and ReactJS [19].

The OSGi specification describes a modular system and a service platform for the Java programming language that implements a complete and dynamic **component model**. Components are deployed as bundles in an OSGi container. Each bundle of the system is complete in itself in terms of its lifecycle management. The OSGi bundles hide their implementations from other bundles while communicating through services. So, OSGi provides a collaborative environment for micro-services to work together to fulfil the overall goal of the system. The dynamic component model makes the dynamic updating of bundles possible. Additionally, system components (bundles) can be remotely installed, started, stopped, and uninstalled without requiring a full system reboot [20]. The fact that no downtime is required for maintenance and enhancement of the system makes it a viable choice over traditional monolithic software development approaches.

MySQL 5.6 server has been used for the aggregated system database. It is an **open-source relational database management system** (RDBMS). MySQL is easy to use, secure, scalable, and yet extremely powerful.

The system is designed to make the statistical model dynamically configurable. To achieve this, we have used the Portable Format for Analytics (PFA). PFA is a new standard for statistical models. It provides a way for the analytics team to describe and implement predictive models. The model can be provided to the deployment team for application in the production environment. A scoring model is an executable module that performs a purely mathematical task. It has a well-defined input and output. The HAVEN model is configured in the aggregated database which is then loaded by the system to compute the HAVEN score, which is generated from physiological data, laboratory test results, and other key variables collated for each patient.

Apart from micro-level deployable components which are deployed as bundles in an OSGi container, the system is conceptually and virtually divided into different layers. Different layers of the system are explained in the following sections, and components within each layer are shown in the system architecture diagram.

Data aggregation

The patient-specific clinical and administrative data are sourced in real-time from different “live” hospital systems, and are integrated by the ‘HAVEN Integration Engine’ (HIE). Currently, the HAVEN system supports data from three clinical systems:

1. The Patient Administration System (PAS) for notification of Admissions, Discharges and Transfers (ADT);

2. The Vital Signs collection system (known as SEND - System for Electronic Notification and Documentation [21]); and
3. The hospital’s pathology laboratory system.

While the system has been configured to use data available within its trial environment, it has been designed to accept commonly available data from a variety of system sources. All coding of data within the clinical systems conforms to current, common terminology. Messages arrive in a standard HL7 version 2 format (used commonly across the NHS and worldwide). The architecture of the HAVEN system (Fig. 1) shows the HIE, deployed in an OSGi container, receiving data from different sources.

The HIE takes those inputs and processes the data for the ‘HAVEN Scoring Engine’ (HSE) to compute the HAVEN score for each patient. Each HL7 message contains a partial package of data that includes patient demographics, clinical physiology, and laboratory results. The HIE receives the incoming HL7 message, validates it against expected HL7 message structure, pre-processes the data items contained within it, and stores them in the integrated database. Abiding by the ‘Separation of Concerns’ design principle, each separate set of linked functionalities is implemented and bundled in a separate component which is then deployed in an OSGi container. There are separate components within the OSGi container that process each type of message, for example ADT messages, vital signs data, and laboratory results.

The ‘Core’ as shown in Fig. 1 contains the common code that is used by multiple other bundles. It encapsulates the database model and data access functionality via services.

Scoring

The ‘Haven Scoring Engine’ computes the HAVEN Score for a patient based on the currently configured HAVEN predictive model (specified as a PFA document) that is stored in the integrated system database. Development and evaluation of the HAVEN risk score will be the subject of a future publication from the project team.

PFA is an emerging standard for statistical models and data transformation engines [18]. A PFA document is a JSON format document with additional constraints. We also explored the use of Predictive Model Markup Language (PMML) - this is an XML-based language-neutral way to encode models, and is a well-established format in the analytics space. However, PMML has limited support for computation as it has only a standard defined set of supported models. Specific needs outside the standard require tremendous effort, or may not be possible. PFA provides the flexibility of arbitrary function composition. A very simple example of a PFA model is shown in Fig. 2 below. It just adds a number to the input and returns the result.

One could implement complex models (e.g. a clustering model, a tree model, a change detection model, a segmented model and many more) using PFA, with the

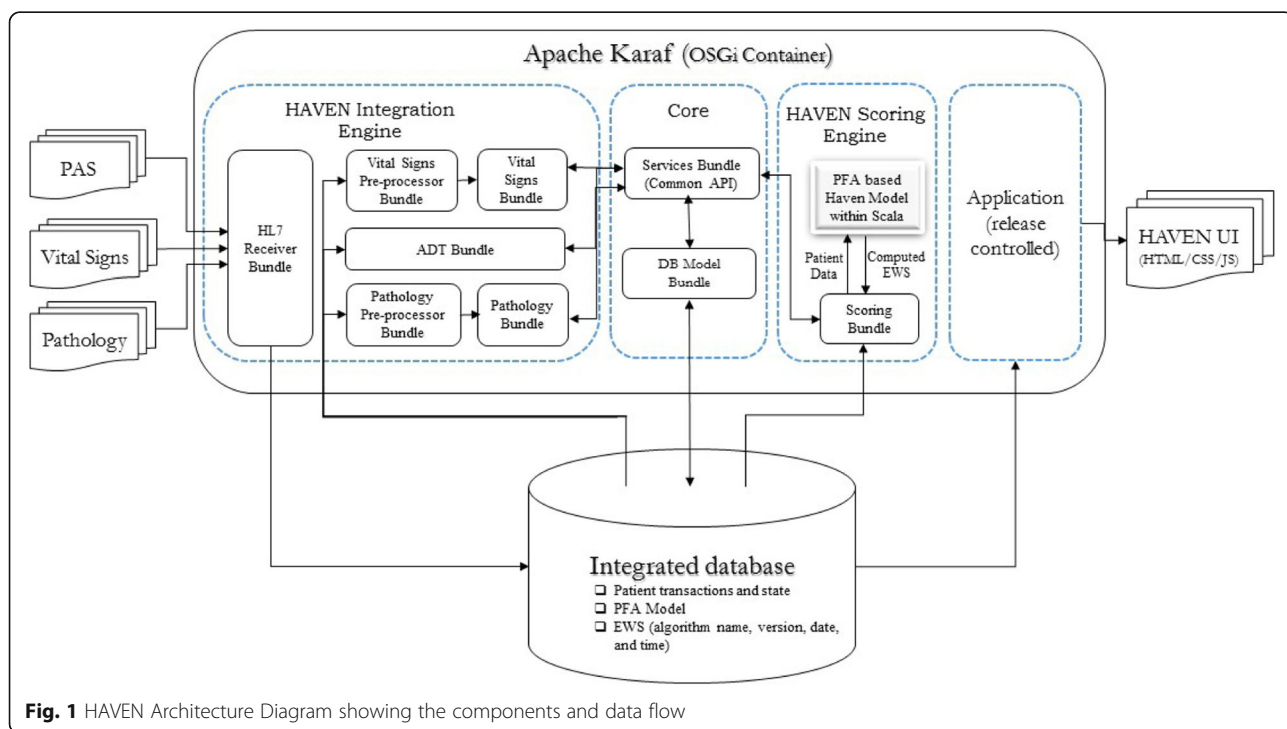


Fig. 1 HAVEN Architecture Diagram showing the components and data flow

```

{
  "input": "double",
  "output": "double",
  "action": {
    "+": [
      "input",
      200
    ]
  },
  "name": "SimpleModel",
  "version": 1
}

```

Fig. 2 Example of PFA

possibility of extending the computational capabilities with customised functions.

Another benefit of using PFA is that the entire scoring flow is represented in a standardized manner in a single executable document. All the additional pre- and post-processing can be encapsulated by the analytics team in that document. No separate code fragment or script is required by the engineering team for model usage. This makes the operationalisation of the models “turn-key”. These features of PFA make it ideal for defining and executing the model to compute the HAVEN score.

The PFA model can be executed using the Hadrian library [22]. Hadrian is designed to be embedded in applications or used as a scoring engine container. The technology stack used to execute the PFA model in HAVEN system is shown in Fig. 3.

The Haven Scoring Engine is also deployed as a bundle in the OSGi container. Scala, a general-purpose programming language, has been embedded within the Haven Scoring Engine to provide language interoperability with Java and to embed the Hadrian API, which is capable of executing a PFA-based model.

On receiving the latest data for each patient, the system computes their HAVEN score using the PFA-based HAVEN model that has been configured. The system has the capability to compute multiple early warning scores based on differently configured PFA models.

User interface

The HAVEN software is built as a web-based application to enable greater flexibility for deployment across multiple hardware platforms. The Spring-boot [23] framework has been used to implement the server side of the User Interface (UI), handling requests coming from the client’s browser. It uses a Rest API [24] with data in JSON [25] format. The data provided by the Rest API is

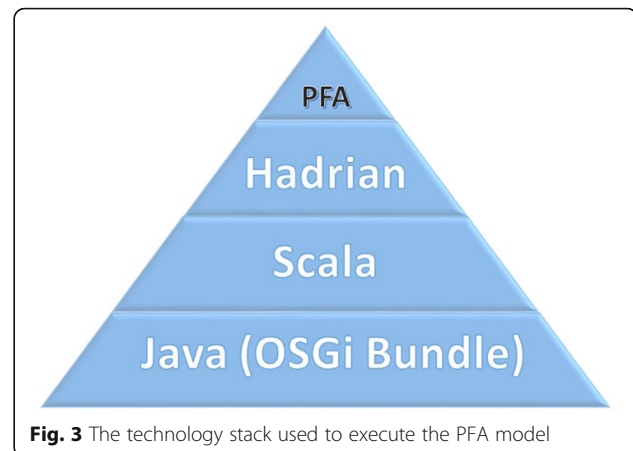


Fig. 3 The technology stack used to execute the PFA model

subsequently used by the client-side implementation of the UI.

The client side of the UI is implemented in ReactJS [19], a JavaScript library. It allows the vital sign observation data to be pre-processed within the web-browser, thus minimising the number of requests made to the server. ReactJS enables the creation of large-scale web-applications that use dynamic data and can change over time without reloading the page, hence providing speed, simplicity, and scalability. It is used for building highly-dynamic and interactive user interfaces. The client-side implementation is then packaged using the JavaScript package manager, NPM [26].

The configurability of the system enables different user groups to easily view patient groups of interest from one system.

The UI facilitates data accessibility by providing user-customised views for carers, nursing staff, and clinicians. It summarises the patients’ recent status and represents the computed HAVEN score for each. It also provides the patients’ demographic overview along with an option to review the physiological data in detail. It focusses on summarising and visualising the data in a manner that suits a user group purpose.

The UI was developed in collaboration with the clinical teams. This ensured a user-centred approach in keeping with best user-design practice. Applied Cognitive Task Analysis interviews [27] and Card Sort tasks [28] were conducted with clinical staff. These and Process Mapping [29, 30] of workflows enabled a broader understanding of how variability of individual patient conditions can affect the whole team.

The information from these activities guided decisions on both UI content and layout. The UI development was an iterative process. Each development stage was evaluated using standard human factors measures of usability,

understanding, efficiency, and accuracy. The final version (see Fig. 4) performed well across all areas, resulting in a UI that affords accurate understanding with low levels of cognitive workload.

Discussion

The system has been used as part of the HAVEN research project using patient data at Oxford University Hospitals NHS Foundation Trust. The system has so far processed > 164,000 vital signs observations and > 68,000 laboratory results for > 125,000 hospital inpatients and shown its capability to process large amounts of data effectively. It is being used as part of the HAVEN project’s evaluation phase to review the effectiveness of the system in identifying patients who might benefit from transfer to ICU. No clinical decisions have been made based on the output from the prototype system. All evaluations have been completed using retrospective data. The clinical and the user interface design parts of the project will be the subject of separate forthcoming publications.

We have demonstrated the technological benefits of our system architecture.

PFA is an emerging standard for describing the statistical model in data analysis and the statistical domain. It is usually executed within a Python or standalone Java application. HAVEN has included the PFA based statistical scoring model in its system. It is the first health-

related software system we are aware of that uses the OSGi (Open Services Gateway initiative) Framework in combination with PFA. OSGi is a well-established technology and widely used by many renowned organisations. It replaces the traditional monolithic architecture software development approach with a new micro-services one.

The HAVEN system exploits the benefits of both technologies. By being able to execute the PFA model by embedding the required technologies/libraries like Scala and Hadrian in an OSGi bundle, it is possible to compute a patient’s early warning score based on different scoring configurable models. This makes the system much more configurable and upgradable in terms of computing an early warning score by deploying the latest model. The other benefits that OSGi provides to the HAVEN system are reduced complexity, system scalability, a modular approach, easy deployment and dynamic updates. The system is easily manageable and the components as well as the predictive model can be upgraded without restarting the system.

Conclusions

We have designed and developed a system to support the identification of patients on wards who would need ICU admission in time to avert the unplanned admission. The system is being evaluated in the context of a

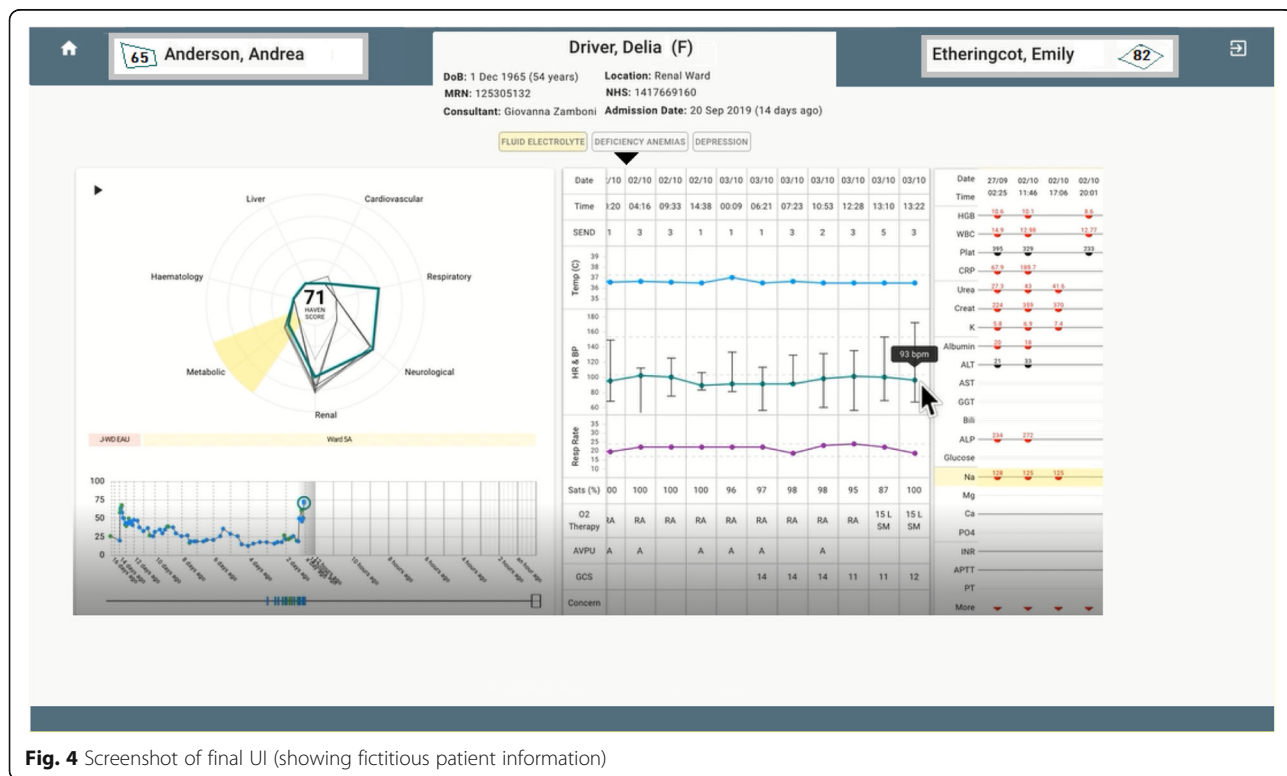


Fig. 4 Screenshot of final UI (showing fictitious patient information)

single hospital and will need considerable evolution before it becomes ready for widespread adoption.

The solution we have designed uses innovative technologies with the possibility of system scalability and enables simple future upgrades of the predictive model. We are using feedback from our human factors colleagues within the wider HAVEN team to make further enhancements and improvements to the system user interface.

Availability and requirements

Project name: HAVEN (Hospital alerting via electronic noticeboard).

Project home page: <https://www.ndcn.ox.ac.uk/research/critical-care-research-group-kadoorie-centre/research-studies/hospital-alerting-via-electronic-noticeboard-haven>

Operating system(s): Platform independent.

Programming language: Java, Scala.

Other requirements: Modern web browser.

License: Proprietary licence.

Any restrictions to use by non-academics: Licence needed.

Acknowledgments

We would like to thank the human factors team (Antonia Burt, Jody Ede, Lauren Morgan, Verity Westgate and Marta Wronikowska) for undertaking the initial exploratory work to understand how this electronic system might best assist clinical teams in their routine work. Also for conducting the UI evaluations and assessment. We would also like to thank the design team (Christopher Fulton and David Wong) who translated the user guidance into visual representations on screen.

Authors' contributions

SSD evaluated the technologies, designed and programmed the software and drafted the manuscript. NF programmed initial test versions of the GUI of the system and helped to draft the manuscript. PC explored and evaluated the technologies to be used, designed and programmed the software and also helped to draft the manuscript. TB contributed in initial system design, formalised the requirements for expected functionality of the system and helped to draft the manuscript. PM contributed in initial system design and helped to draft the manuscript. JLD had overall responsibilities for project management and helped to draft the manuscript. JSB led the technical team and helped to draft the manuscript. PJW is the Chief Investigator for the HAVEN project and helped to draft the manuscript. All authors read and approved the final manuscript.

Funding

This publication presents independent research supported by the Health Innovation Challenge Fund (HICF-R9-524; WT-103703/Z/14/Z), a parallel funding partnership between the Department of Health & Social Care and Wellcome Trust. The views expressed in this publication are those of the author(s) and not necessarily those of the Department of Health & Social Care or Wellcome Trust. The funders had no input into the design of the study, the collection, analysis, and interpretation of data nor in writing the manuscript. PW is supported by the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre (BRC).

Availability of data and materials

Not applicable.

Ethics approval and consent to participate

REC (Research Ethics Committee) approval 16/SC/0264 and CAG (Confidentiality Advisory Group) approval 16/SC/0066.

Consent for publication

Not applicable.

Competing interests

PW co-developed the System for Electronic Notification and Documentation (SEND), for which Sensyne Health has purchased a sole licence. The company has a research agreement with the University of Oxford and royalty agreements with Oxford University Hospitals NHS Trust and the University of Oxford. PW was the Chief Medical Officer for Sensyne Health and holds shares in the company. His department has received funding from Sensyne Health. PW and TB receive royalties on sales of SEND from Sensyne Health. The other authors declare that they have no competing interests.

Author details

¹Centre for Healthcare Modelling and Informatics, University of Portsmouth, Buckingham Building, Lion Terrace, Portsmouth PO1 3HE, UK. ²IM&T, John Radcliffe Hospital, Headley Way, Oxford OX3 9DU, UK. ³Department of Engineering, Institute of Biomedical Engineering, University of Oxford, Old Road Campus Research Building, Oxford OX3 7DQ, UK. ⁴Research & Innovation, Portsmouth Hospitals NHS Trust, Queen Alexandra Hospital, Portsmouth PO6 3LY, UK. ⁵Critical Care Department, University College London NHS Foundation Trust, 235 Euston Road, London NW1 2BU, UK. ⁶Nuffield Department of Clinical Neurosciences, University of Oxford, John Radcliffe Hospital, Headley Way, Oxford OX3 9DU, UK.

Received: 5 March 2019 Accepted: 2 July 2020

Published online: 16 July 2020

References

- NHS Digital. Hospital Adult Critical Care Activity. 2017.
- Liu V, Kipnis P, Rizk NW, Escobar GJ. Adverse outcomes associated with delayed intensive care unit transfers in an integrated healthcare system. *J Hosp Med.* 2012;7(3):224–30.
- Craft M, Dobrenz B, Dornbush E, Hunter M, Morris J, Stone M, et al. An assessment of visualization tools for patient monitoring and medical decision making. In: 2015 Systems and Information Engineering Design Symposium, SIEDS 2015; 2015. p. 212–7.
- National Patient Safety Agency. Recognising and responding appropriately to early signs of deterioration in hospitalised patients. 2007.
- Churpek MM, Yuen TC, Park SY, Gibbons R, Edelson DP. Using electronic health record data to develop and validate a prediction model for adverse outcomes in the wards*. *Crit Care Med.* 2014;42(4):841–8.
- Tarassenko L, Clifton DA, Pinsky MR, Hravnak MT, Woods JR, Watkinson PJ. Centile-based early warning scores derived from statistical distributions of vital signs. *Resuscitation.* 2011;82(8):1013–8.
- Smith GB, Prytherch DR, Meredith P, Schmidt PE, Featherstone PI. The ability of the National Early Warning Score (NEWS) to discriminate patients at risk of early cardiac arrest, unanticipated intensive care unit admission, and death. *Resuscitation.* 2013;84(4):465–70.
- Smith GB, Prytherch DR, Schmidt PE, Featherstone PI. Review and performance evaluation of aggregate weighted “track and trigger” systems. *Resuscitation.* 2008;77(2):170–9.
- Yu S, Leung S, Heo M, Soto GJ, Shah RT, Gunda S, et al. Comparison of risk prediction scoring systems for ward patients: a retrospective nested case-control study. *Crit Care.* 2014;18(3):R132.
- Ludikhuijsen J, Smorenburg SM, de Rooij SE, de Jonge E. Identification of deteriorating patients on general wards; measurement of vital parameters and potential effectiveness of the Modified Early Warning Score. *J Crit Care.* 2012;27(4):424.e7–424.e13.
- Bapoje SR, Gaudiani JL, Narayanan V, Albert RK. Unplanned transfers to a medical intensive care unit: causes and relationship to preventable errors in care. *J Hosp Med.* 2011;6(2):68–72.
- Harris S, Singer M, Rowan K, Sanderson C. Delay to admission to critical care and mortality among deteriorating ward patients in UK hospitals: a multicentre, prospective, observational cohort study. *Lancet.* 2015;385:S40.
- Jäderling G, Calzavacca P, Bell M, Martling CR, Jones D, Bellomo R, et al. The deteriorating ward patient: a Swedish-Australian comparison. *Intensive Care Med.* 2011;37(6):1000–5.
- Knaus WA, Draper EA, Wagner DP, Zimmerman JE. APACHE II: a severity of disease classification system. *Crit Care Med.* 1985;13(10):818–29 Available from: <http://europepmc.org/abstract/MED/3928249>.

15. Java Software | Oracle United Kingdom. Available from: <https://www.oracle.com/uk/java/>. Accessed 10 Feb 2020.
16. MySQL. Available from: <https://www.mysql.com/>. Accessed 10 Feb 2020.
17. OSGi™ Alliance – The Dynamic Module System for Java. Available from: <https://www.osgi.org/>. Accessed 10 Feb 2020.
18. PFA. Portable Format for Analytics. Available from: <http://dmg.org/pfa/>. Accessed 10 Feb 2020.
19. React – A JavaScript library for building user interfaces. Available from: <https://reactjs.org/>. Accessed 10 Feb 2020.
20. Benefits of Using OSGi – OSGi™ Alliance. Available from: <https://www.osgi.org/developer/benefits-of-using-osgi/>. Accessed 10 Feb 2020.
21. Wong D, Bonnici T, Knight J, Morgan L, Coombes P, Watkinson P. SEND: a system for electronic notification and documentation of vital sign observations. *BMC Med Inform Decis Mak.* 2015;15:68.
22. GitHub - modelop/hadrian: Implementations of the Portable Format for Analytics (PFA). Available from: <https://github.com/modelop/hadrian>. Accessed 10 Feb 2020.
23. Spring Boot. Available from: <https://spring.io/projects/spring-boot>. Accessed 10 Feb 2020.
24. What is REST – Learn to create timeless REST APIs. Available from: <https://restfulapi.net/>. Accessed 10 Feb 2020.
25. JSON. Available from: <https://www.json.org/json-en.html>. Accessed 10 Feb 2020.
26. npm | build amazing things. Available from: <https://www.npmjs.com/>. Accessed 10 Feb 2020.
27. Militello LG, Hutton RJB. Applied cognitive task analysis (ACTA): a practitioner's toolkit for understanding cognitive task demands. *Ergonomics.* 1998;41(11):1618–41.
28. Card Sorting | The Encyclopedia of Human-Computer Interaction, 2nd Ed. Available from: <https://www.interaction-design.org/literature/book/the-encyclopedia-of-human-computer-interaction-2nd-ed/card-sorting>. Accessed 10 Feb 2020.
29. Improvement Hub » Improvement Leaders' Guide – Process Mapping, Analysis and Redesign – General Improvement Skills. Available from: <https://www.england.nhs.uk/improvement-hub/publication/improvement-leaders-guide-process-mapping-analysis-and-redesign-general-improvement-skills/>. Accessed 10 Feb 2020.
30. Colligan L, Anderson JE, Potts HW, Berman J. Does the process map influence the outcome of quality improvement work? A comparison of a sequential flow diagram and a hierarchical task analysis diagram. *BMC Health Serv Res.* 2010;10(1):7. [cited 2020 Feb 10]. <https://doi.org/10.1186/1472-6963-10-7>.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

