BMJ Open A protocol for developing early warning score models from vital signs data in hospitals using ensembles of decision trees

Michael Xu,¹ Benjamin Tam,² Lehana Thabane,³ Alison Fox-Robichaud²

To cite: Xu M, Tam B, Thabane L, *et al.* A protocol for developing early warning score models from vital signs data in hospitals using ensembles of decision trees. *BMJ Open* 2015;**5**:e008699. doi:10.1136/bmjopen-2015-008699

► Prepublication history for this paper is available online. To view these files please visit the journal online (<http://dx.doi.org/10.1136/bmjopen-2015-008699>).

Received 5 June 2015
Revised 12 August 2015
Accepted 19 August 2015

CrossMark

¹Bachelor of Health Sciences Program, McMaster University, Hamilton, Ontario, Canada

²Department of Medicine, McMaster University, Hamilton, Ontario, Canada

³Department of Epidemiology and Biostatistics, McMaster University, Hamilton, Ontario, Canada

Correspondence to

Dr Alison Fox-Robichaud;
afoxrob@mcmaster.ca

ABSTRACT

Introduction: Multiple early warning scores (EWS) have been developed and implemented to reduce cardiac arrests on hospital wards. Case-control observational studies that generate an area under the receiver operator curve (AUROC) are the usual validation method, but investigators have also generated EWS with algorithms with no prior clinical knowledge. We present a protocol for the validation and comparison of our local Hamilton Early Warning Score (HEWS) with that generated using decision tree (DT) methods.

Methods and analysis: A database of electronically recorded vital signs from 4 medical and 4 surgical wards will be used to generate DT EWS (DT-HEWS). A third EWS will be generated using ensemble-based methods. Missing data will be multiple imputed. For a relative risk reduction of 50% in our composite outcome (cardiac or respiratory arrest, unanticipated intensive care unit (ICU) admission or hospital death) with a power of 80%, we calculated a sample size of 17 151 patient days based on our cardiac arrest rates in 2012. The performance of the National EWS, DT-HEWS and the ensemble EWS will be compared using AUROC.

Ethics and dissemination: Ethics approval was received from the Hamilton Integrated Research Ethics Board (#13-724-C). The vital signs and associated outcomes are stored in a database on our secure hospital server. Preliminary dissemination of this protocol was presented in abstract form at an international critical care meeting. Final results of this analysis will be used to improve on the existing HEWS and will be shared through publication and presentation at critical care meetings.

INTRODUCTION

Deterioration of patients' condition in hospitals is frequently preceded by abnormal vital or other physiological signs.¹ The failure of clinicians and staff responsible for the care of the patient to recognise and intervene in the deterioration of a patient can result in

Strengths and limitations of this study

- Novel approach with the use of 'big data'.
- Validation of a new warning score in comparison to previous published scores.
- The need to impute data for missing vital signs.
- The relatively low event rate for the composite end point, particularly in the setting of a mature rapid response system.

increased risk of death or cardiopulmonary arrest. The failure to recognise deterioration of a patient can also result in avoidable and unwanted admissions to intensive care units (ICUs). Hospitals are constrained by their resources with regard to how they can manage patient care; with few ICU beds available, it is preferable that avoidable admissions to the unit are intervened and treated appropriately prior to severe deterioration of condition.

Early warning scores (EWS) vary in the design and inclusion of which physiological parameters are assessed. At the most simplistic level, they can be thought of as models that assess the risk of mortality following a given set of vitals. EWS usage has been on the rise and has also been widely implemented in different forms with Subbe's Modified EWS (MEWS),² VitalPac EWS,³ the NHS' National EWS (NEWS)⁴ and, most recently, the Bedside⁵ Paediatric EWS. This was accomplished through the assignment of a score to the patient's physiological parameters to evaluate how ill a patient is. The rationale for such a score is earlier evaluation of the patient prior to deterioration. Categorisation of the deviation of a patient's physiological parameters may help to guide care and intervention.

The Hamilton Early Warning Score (HEWS; figure 1) uses a combination of systolic blood pressure, heart rate, respiratory

Figure 1 HEWS limits and vitals used to assess patient condition. BP, blood pressure; CAM, Confusion Assessment Method; HEWS, Hamilton Early Warning Score; HR, heart rate.

	3	2	1	0	1	2	3
HR/pulse		<40	41 - 50	51 - 100	101 - 110	111 - 130	>130
Sys BP	<70	71 - 90		91-170		171 - 200	>200
Resp Rate	<8	8 - 13		14 - 20		21 - 30	>30
Temp	<35		35.1 - 36	36.1 - 37.9	38 - 39	≥39.1	
O ₂ Sat	<85	85-92		>92			
O ₂ Therapy				Room Air	<5 L/min Or <50% by mask		>5 l/min Or 50% by mask
Change in CNS from Baseline		CAM +ve		Alert	Voice	Pain	Unresponsive

rate, temperature and Alert-Voice-Pain-Unresponsive (AVPU) scale in combination with Confusion Assessment Method (CAM) delirium. The score was developed on the basis of a review of published scores and consensus from an interprofessional group of health experts in acute care medicine. Like other scores, HEWS was developed on the basis of clinical judgements and a trial and error process to find an optimal threshold.

The limitation to this is that clinical judgement and trial and error methods may miss subtle trends or patterns in a patient's parameters that indicate deterioration. These trends or patterns can be noticed or detected through the use of computer algorithms. Without appropriate involvement from clinical judgement though, a computer model may develop a score that is either too complex to be used or lacks clinical relevance to patient care. In this protocol, we adopt the notion that a model needs to be guided by clinical judgement, but at the same time clinical judgement may not evolve fast enough to detect certain cases that would otherwise be undetected by a conventional EWS. Patient populations change, and the demands of healthcare from a community may change as well, so it is sensible that an EWS should evolve with the patient population.

OBJECTIVES

The primary objective of this study will be to validate the current HEWS through the development of an EWS using decision tree methods. The secondary objective will be to compare the existing HEWS, which evaluates each vital sign independently, with a second decision tree generated score that tracks all vitals and evaluates them in relation to each other. This secondary objective will allow us to compare the predictive performance of the decision tree model with that of the current existing HEWS, and determine, if the decision tree model has

superior predictive ability, which vitals take priority when determining patient deterioration.

DECISION TREES

A decision tree attempts to classify data items by recursively posing a series of questions about parameters and features that describe the items.⁶ A graphic example of this can be seen in figure 2 where a series of yes/no questions are used to sort data into nodes. The advantage to such a model is that it is more interpretable and understandable than other classifiers such as neural networks or support vector machines, as simple questions are asked and answered. Decision trees have been successfully used to shape guidelines regarding decision-making processes.⁷ They also possess flexibility with

Figure 2 Illustration of the heart rate aspect of the Hamilton Early Warning Score divided into a decision tree. Grey indicates a terminal node at which point a score would be given to the vital sign.

regard to the types of data they can handle and, once constructed, can classify new items quickly.

Building trees

Decision trees are a compilation of questions that seek to classify events with rules. A series of good questions will separate the data set into subsets that are nearly homogeneous, which can then be separated again into classes. The goal is to have as little variance as possible between each class, either through reduction of entropy or Gini index, and thus increase in information gain.⁸

Decision trees work from a top-down approach, where questions are continually selected recursively to form smaller subsets. A crucial step in the building of a decision tree is determining where and how to limit the complexity of the learned trees. This is necessary to avoid the decision tree overfitting to its training data.⁹

Boosting, bagging and random forests

A collection of decision trees can improve on the accuracy of a single decision tree by combining the results of the collection. These collections are sometimes among the best performing at classification tasks.¹⁰

Boosting creates multiple decision trees that have different questions regarding the same data set and same features. On generation of a tree that misclassifies an event, a new tree is generated that weighs the relative importance of that event more heavily. This is repeated multiple times until trees are combined and evaluated.

Bagging involves bootstrapping the data to decrease variance in the population by producing multisets of the original data. Using these multisets, trees are generated and through a process of voting their classification rules are generated. The predictive value of the rules may not increase through this method, but a reduction in variance of predictions can occur.¹⁰

METHODS

The data set for both the development and testing of the decision tree score will be retrospectively acquired from a continuous set of electronic vitals and patient notes, of all patients who were admitted to the medical or surgical wards between the dates of 1 January 2014 and 30 September 2014, at two sites in Hamilton, Ontario.

One of the sites was in the process of implementing an EWS, and the other has an established EWS with a rapid response team. The sample size calculation was based on our analysis for code blue rates in 2012. We found the code rate to be 1.57/1000 patient days at the first site and 2.41/1000 patient days at the second site. To determine a relative risk reduction of 50% with a power of 80%, the sample size needed is 17 151 patient days, assuming 200 beds are filled on a daily basis. This approximates to 3 months of consecutive patient enrolment; our timeline was extended to ensure appropriate power for comparisons. The data set will be further subdivided into two sections; the first 6 months of data will

be used to train the decision tree and the latter 3 months will be isolated as a testing set. The decision trees will be generated using the sci-kit package in Python, documentation regarding specific usage can be found at <http://scikit-learn.org/stable/documentation.html>

The outcome predicted by the decision tree model will be a composite outcome containing unanticipated ICU transfer, code blue and unanticipated patient death. The predictor vitals to be measured and extracted from the electronic charting system are heart rate, respiratory rate, systolic blood pressure, AVPU, confusion according to CAM, and temperature. These vitals were chosen as they are the most commonly tracked vitals when nurses assess patients, as well as the most common vitals included in other EWS.^{3 11}

The difficulty of a computer model lies in being able to translate it back into a robust and simple tool that clinicians can both understand and want to use to support their judgement, while at the same time maintaining a high degree of accuracy.¹² Selection of the ideal form of analysis is therefore crucial; too simple a model and the accuracy of the model suffers, too complex and it will be too complicated to implement in a clinical environment. In addition, robustness and accuracy must also be tested through the external validation of the model.¹³ This can be achieved either through more patient data being collected or the process of bootstrapping, the former providing more data and the latter generating simulated data sets from the initial set of observations. In the context of this protocol, boosting as an ensemble method will be used as the approach to increasing the value of decision trees as it combines clinical judgement through the use of preselected features, and is more easily interpreted in the form of one final decision tree rather than a voting system.^{13 14}

Planned statistical analysis

Both HEWS and the decision tree scores will be evaluated to determine their ability to discriminate patients who are at risk of the above outcomes within a 72 h period following observation of an abnormal vital sign. The ability for both to do this will be evaluated using an area under the receiver operator curve (AUROC). AUROC values for the generated decision tree will be compared with that of the AUROC for HEWS. An efficiency curve will then be plotted comparing the percentage of observations that experienced the composite outcome with the percentage of observations that exceeded or were at a given score. External validation will be determined through the application of the model to the testing data, and internal validation through comparison to the original training data. A secondary analysis will be conducted examining the trend of a patient's HEWS, and whether this may also be predictive of a patient's outcomes in addition to HEWS values at a given point in time.

Missing data will be dealt with using multiple imputation when possible, specifically using the Multiple Imputation using Chained Equations (MICE) method.¹⁵

DISCUSSION

Currently, most EWS, including HEWS, were developed using a trial and error approach through round-table discussions such as described by the National Early Warning Score Development and Implementation Group responsible for the development of NEWS and MEWS.^{11 16 17} Decision trees have been used by Badriyah *et al*¹⁸ to validate NEWS, though a key difference between the proposed method and the one conducted by Badriyah *et al*¹⁸ will be the generation of a decision tree that encompasses all vitals rather than a separate tree per vital sign. The use of decision trees was a choice made on the basis of the relative robustness of its classification ability as well as the clarity and ease of translation between a model and a rule set that can be interpreted by clinical staff. Other more complex models, such as support vector machining, may be more accurate but generated rule sets that are difficult to translate and interpret. The second decision tree to be generated using ensemble-based methods and which accounts for all vitals in one tree can help to determine if priority or precedence needs to be given to certain vitals over others, at the moment all vitals in all EWS are weighed equally. One limitation of the current study will be the missing values for vitals at the site where EWS implementation was ongoing, as vitals were poorly charted prior to score introduction.

We anticipate having the HEWS score to be very similar in performance and structure to the first decision tree which evaluates all vitals independently, given that both use the same predictors. Given that prior studies comparing the performance of a single decision tree to ensemble-based decision trees have favoured the predictive ability of the latter, we believe that the ensemble-based trees will provide a more accurate predictive ability.¹⁹ The potential clinical use for either method used to generate decision tree EWS would be providing a relatively low cost and quick method of developing an EWS or for the evaluation of a currently in place EWS.

Twitter Follow Michael Xu at @mikekxu

Contributors MX designed the database for vital signs collection, wrote the statistical analysis plan, developed the methodology, drafted and revised the paper, and developed the idea behind the framework. BT filed for ethics and funding, revised the paper and contributed to the development of the methodology of the framework. LT provided statistical oversight and guidance, and revised the paper. AF-R revised and drafted the draft paper, as well as filing for ethics and funding.

Funding This project was funded by residency safety grants from Hamilton Health Sciences and the Department of Medicine, McMaster University.

Competing interests None declared.

Ethics approval Hamilton Integrated Research Ethics Board (#13-724-C).

Provenance and peer review Not commissioned; externally peer reviewed.

Open Access This is an Open Access article distributed in accordance with the Creative Commons Attribution Non Commercial (CC BY-NC 4.0) license, which permits others to distribute, remix, adapt, build upon this work non-commercially, and license their derivative works on different terms, provided the original work is properly cited and the use is non-commercial. See: <http://creativecommons.org/licenses/by-nc/4.0/>

REFERENCES

1. Goldhill DR, McNarry AF, Mandersloot G, *et al*. A physiologically-based early warning score for ward patients: the association between score and outcome. *Anaesthesia* 2005;60:547–53.
2. Subbe CP, Kruger M, Rutherford P, *et al*. Validation of a modified Early Warning Score in medical admissions. *QJM* 2001;94:521–6.
3. Prytherch DR, Smith GB, Schmidt PE, *et al*. ViEWS—towards a national early warning score for detecting adult inpatient deterioration. *Resuscitation* 2010;81:932–7.
4. McGinley A, Pearse RM. A national early warning score for acutely ill patients. *BMJ* 2012;345:e5310.
5. Parshuram CS, Hutchison J, Middaugh K. Development and initial validation of the Bedside Paediatric Early Warning System score. *Crit Care* 2009;13:1–10.
6. Breiman L. Random forests. *Mach Learn* 2001;45:5–32.
7. Quinlan JR. Generating production rules from decision trees. *Proceedings of the 10th International Joint Conference on Artificial Intelligence*. 1987;1:304–7.
8. Quinlan JR. Induction of decision trees. *Mach Learn* 1986;1:81–106.
9. Esposito F, Malerba D, Semeraro G, *et al*. A comparative analysis of methods for pruning decision trees. *IEEE Trans Pattern Anal Mach Intell* 1997;19:476–91.
10. Caruana R, Niculescu-Mizil A. An empirical comparison of supervised learning algorithms. *Proceedings of the 23rd International Conference on Machine Learning*. New York, NY, USA: ACM, 2006:161–8. (cited 9 Nov 2014). <http://doi.acm.org/10.1145/1143844.1143865>
11. Smith GB, Prytherch DR, Meredith P, *et al*. 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;84:465–70.
12. Kannampallil TG, Schauer GF, Cohen T, *et al*. Considering complexity in healthcare systems. *J Biomed Inform* 2011;44:943–7.
13. Collins GS, de Groot JA, Dutton S, *et al*. External validation of multivariable prediction models: a systematic review of methodological conduct and reporting. *BMC Med Res Methodol* 2014;14:40.
14. Schapire RE, Singer Y. Improved boosting algorithms using confidence-rated predictions. *Mach Learn* 1999;37:297–336.
15. White IR, Royston P, Wood AM. Multiple imputation using chained equations: issues and guidance for practice. *Stat Med* 2011;30:377–99.
16. Kyriacos U, Jelsma J, James M, *et al*. Monitoring vital signs: development of a modified early warning scoring (Mews) system for general wards in a developing country. *PLoS ONE* 2014;9:e87073.
17. Kyriacos U, Jelsma J. Early warning scoring systems versus standard observations charts for wards in South Africa: a cluster randomized controlled trial. *Trials* 2015;16:103.
18. Badriyah T, Briggs JS, Meredith P, *et al*. Decision-tree early warning score (DTEWS) validates the design of the National Early Warning Score (NEWS). *Resuscitation* 2014;85:418–23.
19. Banfield RE, Hall LO, Bowyer KW, *et al*. A comparison of decision tree ensemble creation techniques. *IEEE Trans Pattern Anal Mach Intell* 2007;29:173–80.