Individual QI projects from single institutions

OPEN

The Vitals Risk Index—Retrospective Performance Analysis of an Automated and Objective Pediatric Early Warning System

Tyler J. Gorham, PhD, MPH*; Steve Rust, PhD*; Laura Rust, MD, MPH†‡; Stacy Kuehn, RN, BSN†; Jing Yang, MS*; James Shuhan Lin, MD‡; Jeffrey Hoffman, MD†‡; Yungui Huang, PhD, MBA*; Simon Lin, MD, MBA*‡; Richard McClead, MD, MHA†; Richard Brilli, MD, MCCM†‡; Ryan Bode, MD†‡; Tensing Maa, MD†‡

Abstract

Introduction: Pediatric in-hospital cardiac arrests and emergent transfers to the pediatric intensive care unit (ICU) represent a serious patient safety concern with associated increased morbidity and mortality. Some institutions have turned to the electronic health record and predictive analytics in search of earlier and more accurate detection of patients at risk for decompensation. **Methods:** Objective electronic health record data from 2011 to 2017 was utilized to develop an automated early warning system score aimed at identifying hospitalized children at risk of clinical deterioration. Five vital sign measurements and supplemental oxygen requirement data were used to build the Vitals Risk Index (VRI) model, using multivariate logistic regression. We compared the VRI to the hospital's existing early warning system, an adaptation of Monaghan's Pediatric Early Warning Score system (PEWS). The patient population included hospitalized children 18 years of age and younger while being cared for outside of the ICU. This dataset included 158 case hospitalizations (102 emergent transfers to the ICU and 56 "code blue" events) and 135,597 control hospitalizations. **Results:** When identifying deteriorating patients 2 hours before an event, there was no significant difference between Pediatric Early Warning Score and VRI's areas under the receiver operating characteristic curve at false-positive rates ≤ 10% ($pAUC_{10}$ of 0.065 and 0.064, respectively; $P = 0.74$), a threshold chosen to compare the 2 approaches under clinically tolerable false-positive rates. **Conclusions:** The VRI represents an objective, simple, and automated predictive analytics tool for identifying hospitalized pediatric patients at risk of deteriorating outside of the ICU setting. *(Pediatr Qual Saf 2020;2:e271; doi: 10.1097/pq9.0000000000000271; Published online March 20, 2020.)*

From the *Abigail Wexner Research Institute, Nationwide Children's Hospital, Columbus, Ohio; †Nationwide Children's Hospital, Columbus, Ohio; and ‡The Ohio State University College of Medicine, Columbus, Ohio.

*Corresponding author. Address: Tyler J. Gorham, PhD, MPH, Abigail Wexner Research Institute, Nationwide Children's Hospital, 431 S. 18th St., Columbus, OH 43205
Email: tyler.gorham@nationwidechildrens.org

PH: (614) 355-5663.

INTRODUCTION

Survival to hospital discharge following pediatric in-hospital cardiac arrests ranges from 16% to 48%.[1–5] Hundreds of in-hospital cardiac arrests occur suboptimally outside of the pediatric intensive care unit (ICU) and are considered preventable harm by the Children's Hospital Association.[6,7] Nationally, many efforts have attempted to address this distinct population with a demonstrable reduction in in-hospital cardiac arrests outside of the ICU, most notably the implementation of rapid response teams.[3,8–11]

However, unplanned transfers to the ICU persist and carry an associated but avoidable increased mortality.[12–14] Institutions looking to prevent these events have begun to diversify their deterioration detection efforts. One such focus is creating an optimal Pediatric Early Warning Score (PEWS) system,[15–21] which is typically composed of rules applied to objective and subjective patient data and an associated actionable mitigation plan. To date, PEWS systems have been unable to demonstrate a significant decrease in hospital mortality but have favorably affected the rate of other clinical deterioration events.[22,23]

Therefore, institutions have sought to broaden their deterioration detection efforts by utilizing scoring systems and increasing situational awareness while continuing to search for the optimal predictive model.[14]

Widespread adoption of electronic health records (EHR) has created large repositories of patient data, making the development and validation of predictive analytics models against large datasets feasible. Progress in this realm has led to a prediction model of ICU transfer within 24 hours of admission, which outperformed a modified PEWS[24] and completely automated EHR-based warning systems.[25]

Many of these scoring systems are still dependent on subjective assessments, such as mental status or capillary refill time, and rely on real-time documentation. These limitations can increase the documentation burden and reduce the efficacy of such systems when data entry is delayed.[26] Initiatives have begun to investigate using objective variables available in real-time within the EHR to identify patients at high-risk for clinical deterioration.[24]

Our objective was to develop and validate an automated and objective pediatric early warning score utilizing predictive analytics modeling techniques that would perform as well as, or better than, our institution's current standard, a modified version of Monaghan's PEWS (see Supplemental Digital Content at http://links.lww.com/PQ9/A170 for Table 1).[20,27] Such a model would be devoid of the variability currently present with subjective PEWS components (behavioral assessment, capillary refill time, respiratory distress, persistent vomiting). This model, the Vitals Risk Index (VRI), is composed of entirely objective component inputs measured during routine inpatient clinical care with automated risk score calculation. Predicted outcomes of interest were code blue activations outside of the ICU[6] and emergent transfers to the ICU.[14]

METHODS
Study Population and Exclusion Criteria
The study population included children hospitalized from July 1, 2011, to December 31, 2017, at Nationwide Children's Hospital, a freestanding, quaternary care academic children's hospital. As this work was not human subject research, but rather a quality improvement study, review and approval by the Nationwide Children's Hospital Institutional Review Board was not required per policy. Exclusion criteria for both cases and controls included patients who are 19 years of age or older, neonatal ICU hospitalizations, length of stay <12 hours, and absence of vital sign measurements recorded in the EHR (Epic Systems Corporation, Verona, Wisc.). The study population was also limited to patients with at least 1 PEWS assessment recorded in the EHR.

Case hospitalizations were those with either a code blue event outside of the ICU[6] or an emergent transfer to the ICU (see Supplemental Digital Content at http://links.lww.com/PQ9/A171, for full case definitions).[14] Code blue events required emergency assisted ventilation, chest compressions, or electric shock. We excluded code blue events triggered by a seizure (an internal practice) as they are often considered not predictable. Emergent transfers required 1 or more of the following interventions 1 hour before or after transfer: intubation, initiation of vasoactive medications, 60 mL/kg of fluid boluses, or cardiopulmonary resuscitation. We defined control hospitalizations as those without a clinical deterioration event type, as defined above. During the study period, there were 102 emergent transfers and 56 code blue events outside the ICU for a total of 158 case hospitalizations and 135,597 control hospitalizations.

We built the study dataset from data collected on patients outside the ICU within our inpatient EHR system. Unless otherwise ordered, vital signs and PEWS evaluations were nominally recorded every 4 hours. If a patient received a PEWS score of 3 or 4, vital signs and PEWS assessment frequency increased to every 2 hours, and a PEWS score of 5 or 6 resulted in hourly vital signs and PEWS evaluations. Vital sign measurements and supplemental oxygen flowsheet row data were considered valid for up to 24 hours or until a new measurement was recorded; if the most-recent value recorded in the EHR was over 24 hours old, the field was coded as missing. Missing data points were replaced with age-specific nominally normal values that were ultimately coded to zeros in the final scoring algorithm.

Similarly, if no supplementary oxygen data were available, we assumed the patient was not receiving supplementary oxygen. Handling missing data in this manner assumes that patients without a recent assessment are not at an increased risk of deterioration. Missing data imputation procedures were not considered because the implementation of such procedures would not be possible in real-time with the EHR system. To prevent bias toward hospitalizations with multiple events and to avoid unknown effects of event-driven interventions on the detection of subsequent events, only the first event of each hospitalization was employed.

Local Monaghan's PEWS Implementation
NCH operationalized Monaghan's PEWS (adapted from Monaghan[20] and Tucker et al[27]) across all floor units under consideration by the beginning of the study period. PEWS is recorded in the EHR and was included in the extracted data. Because a PEWS score of 5 or greater (PEWS-5) is used by the study institution as a trigger for evaluation, mitigation, and escalated response, PEWS-5 is the baseline against which we compare VRI performance. Comparisons against PEWS-4 are also provided, as this alternative allows for higher sensitivity when compared with PEWS-5, but lower specificity.

Modeling Approach
Five patient vital sign measurements and a proxy for supplemental oxygen requirements were used to develop the

VRI algorithm. Vital signs included heart rate, respiratory rate, temperature, percutaneous oxygen saturation, and systolic blood pressure. We captured supplemental oxygen as either flow rate (liters per minute) or fraction of inspired oxygen delivered when the flow rate was not recorded.

We trained a multivariate logistic regression model to predict the occurrence of an inpatient deterioration event within the subsequent 24-hour period. All vital signs and supplemental oxygen variables were discretized into bins (eg, "very low," "low," "'normal," "high," and "very high"), by age group when appropriate, following the work of Duncan et al (2006)[15] and Fleming et al (2011).[28] Note that this previous work was only referenced to threshold-continuous values by age group for algorithm development and is unrelated to Monaghan's PEWS scoring and thresholds. High and very high supplemental oxygen thresholds were determined by observing change points in the receiver operating characteristic (ROC) curves for these flowsheet rows, predicting emergent transfers and cardiopulmonary failure events.

For each case hospitalization, we limited data used in model training to the 6 hours leading up to the deterioration event, to allow the model to be more specifically trained to data very near an event. For case events occurring <6 hours after admission, we included all data before the event. For controls, we used data from the middle 24 hours of hospitalization if the hospitalization was >24 hours; otherwise, all data were included. This middle 24-hour period is intended to represent a period of relative clinical stability among controls. In the final model, for future EHR implementation, statistically insignificant predictors (P-value $\geq$ 0.05) with negative coefficients were manually dropped. The remaining coefficients were reweighted so that the smallest and largest possible VRI values would be 0 and 100, respectively.

For performance evaluation and comparisons, we included case data for the 24 hours before the first deterioration event and control data for the middle 24 hours of hospitalization. Comparisons were performed using all data up until the time of the event as well as removing data recorded within 1, 2, and 3 hours of the event. This strategy helps address the challenge of identifying at-risk patients with sufficient lead time to either prevent the event from occurring or limit its severity. To allow for the discussion of the VRI as an early warning tool, the primary model results chosen for discussion hereafter are those representing the removal of case data within 2 hours of the deterioration event. VRI performance results were generated by performing 10-fold cross-validation (dividing the data into 10 folds, then using 9 folds to produce predicted values for the remaining fold). The 10-fold cross-validation process was repeated 10 times to reduce the variability in the reported performance results. For graphical and statistical comparisons, PEWS performance is compared with an averaged VRI ROC curve, constructed by taking the mean predicted response at the patient level across the 10 repetitions of cross-validation. We did not use a train-test split in performance assessment due to the small number of available cases.

Confidence intervals for PEWS and VRI ROC curves and sensitivities were constructed using non-parametric stratified bootstrapping,[29] via the "pROC" package in R.[30] All figures were created using the "ggplot2" package in R.[31]

RESULTS

Final Model
The final model coefficients are presented in Table 1. "High" and "low" respiratory rate categories as well as "low" and "very low" temperature ranges were found not to be predictive of deterioration and were collapsed into the "normal" ranges. The VRI is the sum of the applicable model coefficients for the 6 model components.

ROC curves for Monaghan's PEWS and the VRI, based on at least 2 hours of lead time, are presented in Figure 1. There was no significant difference in the area under the ROC curve (AUC) of the VRI 0.76 (95% CI, 0.72–0.80) compared with PEWS 0.73 (0.69–0.78) (P = 0.16; Fig. 1). Additionally, there was no significant difference in the areas under the ROC curve between VRI and PEWS at false-positive rates $\leq$ 10% ($pAUC_{10}$), a threshold chosen to compare the 2 approaches under clinically tolerable false positive rates ($pAUC_{10}$ of 0.065 and 0.064, respectively; P = 0.74). The $pAUC_{10}$ metric can take on values from 0 to 0.1 with a random algorithm taking on the value 0.005.

The VRI is a continuous score, ranging from 0 to 100, so a threshold must be set along this scale to calculate metrics like sensitivity and specificity. Testing VRI thresholds that match the false alarm rate of PEWS-4 (5%) or PEWS-5 (1%) allows for fair comparisons of the sensitivities of these 2 approaches, controlling for the "cost" of false-positive alarms. When the threshold for the VRI is set to match the specificity of PEWS-4 (0.95), the VRI and PEWS have very similar sensitivity levels from the time of the event to 3 hours before an event (Fig. 2A). Matching on the very high specificity of PEWS-5 (0.99), the sensitivity of the VRI (0.25, 95% CI, 0.19–0.32) is significantly lower than that of PEWS-5 (0.46, 95% CI, 0.38–0.54) at the time of the event (Fig. 2B). However, the advantage of PEWS-5 over the VRI disappears if a successful alarm is required to occur 2 or more hours before the deterioration event. Both approaches perform similarly when requiring this much lead time at such high specificity.

DISCUSSION
The goal of this study was to develop a fully automated and objective early warning predictive analytics tool—the VRI—that would perform at least as well as PEWS in the accurate and timely identification of hospitalized pediatric patients at risk for clinical deterioration outside of the
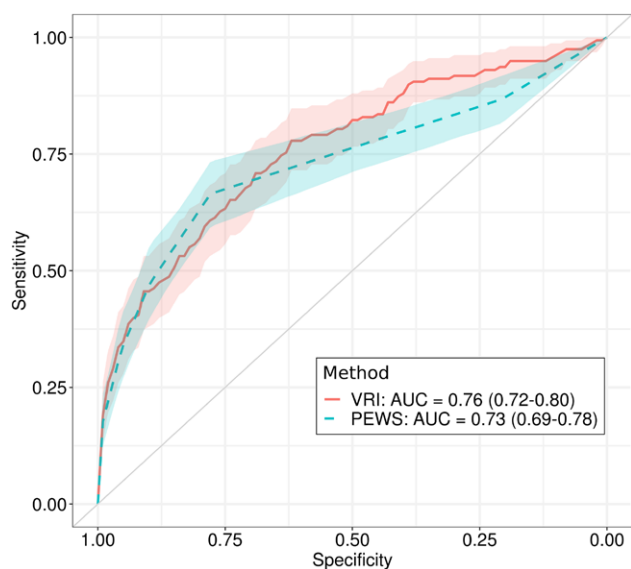
Table 1. Item Sub-scores of the VRI

Item	Age	Item sub-score				
Heart rate (beats/min)	Coefficient	8.79	1.84	0	4.31	18.29
	0–3 months	<90	90–109	110–150	151–180	>180
	3–12 months	<80	80–99	100–150	151–170	>170
	1–4 years	<70	70–89	90–120	121–150	>150
	4–12 years	<60	60–69	70–110	111–130	>130
	>12 years	<50	50–59	60–100	101–120	>120
Respiratory rate (breaths/min)	Coefficient	13.58		0		12.36
	0–3 months	<20		20–80		>80
	3–12 months	<20		20–70		>70
	1–4 years	<15		15–60		>60
	4–12 years	<12		12–40		>40
	>12 years	<8		8–24		>24
Systolic blood pressure (mm Hg)	Coefficient	20.94	4.16	0	0.57	1.31
	0–3 months	<50	50–59	60–80	81–100	>100
	3–12 months	<70	70–79	80–100	101–120	>120
	1–4 years	<75	75–89	90–110	111–125	>125
	4–12 years	<80	80–89	90–120	121–130	>130
	>12 years	<85	85–99	100–130	131–150	>150
Temperature (°F)	Coefficient			0	9.01	14.27
	All ages			<101.3	101.3–104	>104
Oxygen saturation (%)	Coefficient	19.80	5.06	0		
	All ages	<85	85–95	>95		
Supplemental oxygen (flow rate, L/min; FiO_2, %)	Coefficient			0	4.94	13.12
	All ages			Flow rate ≤ 0.45 or (flow rate is NA and FiO_2 ≤ 21)	(0.45 < flow rate ≤ 6.5) or (flow rate is NA and 21 < FiO_2 ≤ 31)	Flow rate > 6.5 or (flow rate is NA and FiO_2 > 31)

ICU setting. We chose the specific and defined outcomes of code blue events outside of the ICU and emergent transfers to the ICU, which represent late or unrecognized inpatient clinical deterioration, because of their associated increased morbidity and mortality.[1,32] To eliminate these preventable harm events, we have focused on earlier detection, mitigation, and potential escalation or transfer to a higher level of care before decompensation. Currently, we have been unable to sustain zero emergent transfers and continue to have patients with insidious


Fig. 1. ROC curves for Monaghan's PEWS and the VRI. The AUC of VRI, when identifying patients at risk of deterioration with 2 hours of lead time, was 0.76 (95% CI, 0.72–0.80) compared with 0.73 (0.69–0.78) for PEWS.

deterioration undetected by Monaghan's PEWS and the healthcare team. Specifically, there have been concerns regarding the predictive ability of PEWS due to its subjective aspects (eg, behavior and mental status) and non-automated nature. Thus, the emphasis is on developing an automated, objective predictive tool that would augment or ultimately replace PEWS.

The VRI did not include any subjective components that exist in Monaghan's PEWS. Despite this advantage, the VRI still requires accurate and timely documentation of vital signs and supplemental oxygen requirements. Fortunately, this behavior is "hard-wired" and embedded in our current workflow. In this study, we intentionally chose a relatively simple modeling approach of logistic regression, because the published coefficients can be easily implemented into any EHR system with built-in clinical decision support. We believe the VRI model is simple, purely objective, automated, real-time, and easily reproducible in other pediatric institutions. Further, we anticipate VRI will complement rather than replace PEWS and add a layer of decision support for clinicians to consider. An automated alert may particularly be beneficial during the high census and high acuity winter months when resources are stretched thin.

Zhai et al[33] reported a machine learning-based algorithm predicting ICU transfer within 24 hours of initial admission. This model used 36 measurements and 155 variables, including vital signs and nursing assessments, with an AUC of 0.91 when identifying patients with at least 2 hours of lead time. Rubin et al[24] developed a predictive model that used objective vital signs as well as automatically calculated pulse pressure, mean arterial pressure, and shock index. Their case encounters
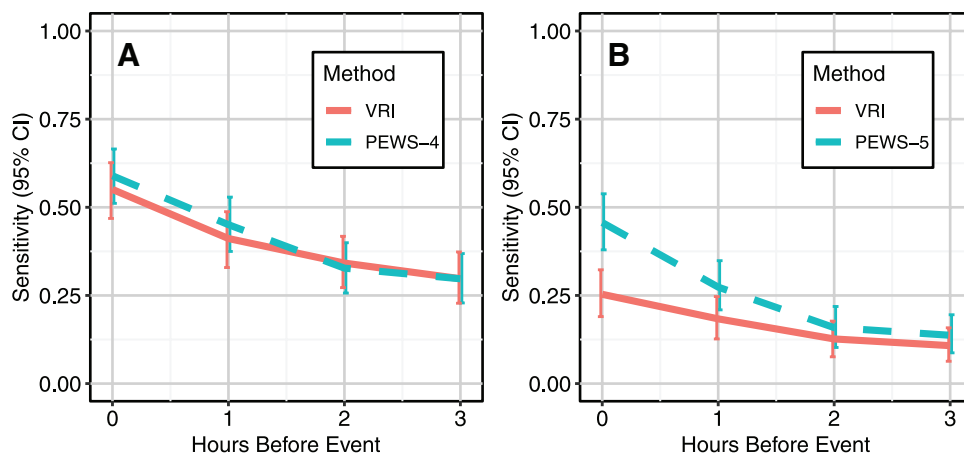
Fig. 2. Sensitivity of the VRI when matching the specificity of (A) PEWS-4 and (B) PEWS-5. A, The VRI has a similar sensitivity to PEWS-4 from 0 to 3 hours before an event (with both approaches operating at a specificity of 0.95). B, PEWS-5 (specificity = 0.99) has a higher sensitivity than VRI at event time, but the 2 approaches perform similarly at 2 and 3 hours before an event.

were defined as any ICU transfer during the hospitalization, rather than just ICU transfers within 24 hours of admission. Their predictive model performed better than modified PEWS at both study institutions tested, with a false positive rate of 26%–27%. Finally, Rothman et al[19] published the development of a pediatric Rothman Index (pRI) using vital signs, nursing assessments, laboratory tests, and cardiac rhythms. AUC for 24-hour mortality was reported over 3 hospitals to be 0.93, 0.93, and 0.95. Due to the relative rarity of pediatric mortality compared with adults, the authors proposed using their model for prediction of "unplanned transfer to the ICU" as an outcome metric. The authors noted a "trend in physiologic deterioration before and after unplanned transfer to the ICU further validates the pRI."[19] However, no sensitivity, specificity, or AUC data were provided to support this conclusion.

Our work confirms that EHR-derived data can be used to successfully develop a predictive model that performs similar to PEWS in detecting pediatric inpatient deterioration outside of the ICU. Compared with the Zhai et al[33] and Rothman et al[19] models, the VRI has added benefits of simplicity and objectivity using vital signs and level of supplemental oxygen while not relying on potentially subjective or delayed nursing assessments. The VRI adds, to prior prediction models, the ability to detect a set of focused, and, in our opinion, meaningful clinical deterioration events that have been demonstrably associated with increased morbidity and mortality. We avoided targeting the set of all ICU transfers because it is very heterogeneous and has not been demonstrably associated with negative healthcare outcomes.

Identification and detection of impending clinical deterioration need to be done in advance of the actual event to mitigate successfully and, if necessary, escalate or transfer to a higher level of care and intensity of resources. Neither VRI nor Monaghan's PEWS perform well (sensitivities <15%) when removing data within the 2 hours before the event under strict specificity levels (eg, ≥0.99). The drop in sensitivity for PEWS-5 from 0 to 2 hours before the event may provide some insight as to why PEWS has not been associated with measurable decreases in mortality and morbidity. The sensitivity of 45% at event time fades to <15% at 2 hours before the event when there is still time to intervene.

There is a balance of "alert fatigue" and the "needle in the haystack" phenomena when trying to predict clinical deterioration. Equally important as predicting and identifying at-risk patients is the evaluation and mitigation response. These responses—including potential physician assessments, rapid response and code teams, increased monitoring, and increased EHR documentation—require personnel, time, and resources. Excessive false positives can lead to alert fatigue and undermine the utility of a predictive analytic tool and the accompanying response. An important question to address moving forward is to determine the targeted or acceptable false-positive rate in predicting clinical deterioration and adverse patient outcomes. To achieve zero-preventable events such as these, we will likely need to tolerate an increase in the false-positive rates associated with pediatric early warning systems and predictive analytic tools.

Limitations

There are multiple potential limitations to our study. This work represented patients from a single, albeit large, free-standing academic children's hospital. While the control population was robust with over 135,000 hospitalizations, there were only 158 case hospitalizations despite including 7 years' worth of data. Also worth mentioning is the heterogeneity in EHR systems utilized by institutions, potentially limiting the generalizability of the described model.

In developing the VRI, while the calculation of the individual components is objective, the categories used to

discretize vital sign measurements were based on expert opinion and could potentially be optimized by a fully data-driven approach. The need to collapse some categories after initial model fitting may be a result of this limitation. Additionally, our approach utilized logistic regression to allow final model coefficients to be implemented in the study institution's EHR, but we acknowledge that deep learning approaches, such as recurrent neural networks, may present an opportunity for improved predictive performance.[34,35]

Future Studies

The next steps include validating model performance prospectively in the patient care environment as well as determining how to integrate VRI into our current workflow, selecting a false positive rate that allows for accurate patient identification but not at the expense of alert fatigue or exhaustion of available resources. Specific patient populations whose baseline vital signs are abnormal for age (eg, oncologic, single ventricle, or ventilator-dependent patients) may require a different algorithm for triggering an alert. Finally, we may incorporate other objective or electronically captured patient characteristics that would improve accuracy and identify at-risk patients, such as past medical history, technology dependence, or prior recent critical event (ie, ICU transfer, rapid response team). Future parameter considerations will balance improved performance with simplicity, need for automation, and objectivity.

CONCLUSIONS

In this study, we developed a novel pediatric early warning systems—the VRI—based solely on objective vital sign measurements and supplemental oxygen demand. The VRI was shown to be as sensitive as Monaghan's PEWS as implemented at Nationwide Children's Hospital, when predicting patient deterioration outside of the ICU 2 to 3 hours before an event. In settings where an early warning system has not been implemented, the VRI may serve as an important clinical decision-support tool utilizing clinical workflows that are likely already incorporated into the EHR.

DISCLOSURE

ACKNOWLEDGMENTS

REFERENCES

1. Reis AG, Nadkarni V, Perondi MB, et al. A prospective investigation into the epidemiology of in-hospital pediatric cardiopulmonary resuscitation using the international Utstein reporting style. *Pediatrics*. 2002;109:200–209.
2. Suominen P, Olkkola KT, Voipio V, et al. Utstein style reporting of in-hospital paediatric cardiopulmonary resuscitation. *Resuscitation*. 2000;45:17–25.
3. Nadkarni VM, Larkin GL, Peberdy MA, et al.; National Registry of Cardiopulmonary Resuscitation Investigators. First documented rhythm and clinical outcome from in-hospital cardiac arrest among children and adults. *JAMA*. 2006;295:50–57.
4. Knudson JD, Neish SR, Cabrera AG, et al. Prevalence and outcomes of pediatric in-hospital cardiopulmonary resuscitation in the United States: an analysis of the Kids' Inpatient Database*. *Crit Care Med*. 2012;40:2940–2944.
5. Girotra S, Spertus JA, Li Y, et al.; American Heart Association Get With the Guidelines–Resuscitation Investigators. Survival trends in pediatric in-hospital cardiac arrests: an analysis from get with the guidelines-resuscitation. *Circ Cardiovasc Qual Outcomes*. 2013;6:42–49.
6. Hayes LW, Dobyns EL, DiGiovine B, et al. A multicenter collaborative approach to reducing pediatric codes outside the ICU. *Pediatrics*. 2012;129:e785–e791.
7. Berg RA, Sutton RM, Holubkov R, et al.; Eunice Kennedy Shriver National Institute of Child Health and Human Development Collaborative Pediatric Critical Care Research Network and for the American Heart Association's Get With the Guidelines-Resuscitation (formerly the National Registry of Cardiopulmonary Resuscitation) Investigators. Ratio of PICU versus ward cardiopulmonary resuscitation events is increasing. *Crit Care Med*. 2013;41:2292–2297.
8. Shin AY, Longhurst C, Sharek PJ. Reducing mortality related to adverse events in children. *Pediatr Clin North Am*. 2012;59:1293–1306.
9. Brilli RJ, Gibson R, Luria JW, et al. Implementation of a medical emergency team in a large pediatric teaching hospital prevents respiratory and cardiopulmonary arrests outside the intensive care unit. *Pediatr Crit Care Med*. 2007;8:236–46; quiz 247.
10. Sharek PJ, Parast LM, Leong K, et al. Effect of a rapid response team on hospital-wide mortality and code rates outside the ICU in a children's hospital. *JAMA*. 2007;298:2267–2274.
11. Tibballs J, Kinney S. Reduction of hospital mortality and of preventable cardiac arrest and death on introduction of a pediatric medical emergency team. *Pediatr Crit Care Med*. 2009;10:306–312.
12. Odetola FO, Rosenberg AL, Davis MM, et al. Do outcomes vary according to the source of admission to the pediatric intensive care unit? *Pediatr Crit Care Med*. 2008;9:20–25.
13. Humphreys S, Totapally BR. Rapid response team calls and unplanned transfers to the pediatric intensive Care Unit in a Pediatric Hospital. *Am J Crit Care*. 2016;25:e9–13.
14. Brady PW, Muething S, Kotagal U, et al. Improving situation awareness to reduce unrecognized clinical deterioration and serious safety events. *Pediatrics*. 2013;131:e298–e308.
15. Duncan H, Hutchison J, Parshuram CS. The pediatric early warning system score: a severity of illness score to predict urgent medical need in hospitalized children. *J Crit Care*. 2006;21:271–278.
16. Bell D, Mac A, Ochoa Y, et al. The texas children's hospital pediatric advanced warning score as a predictor of clinical deterioration in hospitalized infants and children: a modification of the PEWS tool. *J Pediatr Nurs*. 2013;28:e2–e9.
17. Bonafide CP, Roberts KE, Weirich CM, et al. Beyond statistical prediction: qualitative evaluation of the mechanisms by which pediatric early warning scores impact patient safety. *J Hosp Med*. 2013;8:248–253.
18. Parshuram CS, Hutchison J, Middaugh K. Development and initial validation of the bedside paediatric early warning system score. *Crit Care*. 2009;13:R135.
19. Rothman MJ, Tepas JJ 3rd, Nowalk AJ, et al. Development and validation of a continuously age-adjusted measure of patient condition for hospitalized children using the electronic medical record. *J Biomed Inform*. 2017;66:180–193.
20. Monaghan A. Detecting and managing deterioration in children. *Paediatr Nurs*. 2005;17:32–35.

21. Ross CE, Harrysson IJ, Goel VV, et al. Integration of single-center data-driven vital sign parameters into a modified pediatric early warning system. *Pediatr Crit Care Med*. 2017;18:469–476.

22. Parshuram CS, Dryden-Palmer K, Farrell C, et al.; Canadian Critical Care Trials Group and the EPOCH Investigators. Effect of a pediatric early warning system on all-cause mortality in hospitalized pediatric patients: the EPOCH Randomized Clinical Trial. *JAMA*. 2018;319:1002–1012.

23. Endsley MR. Toward a theory of situation awareness in dynamic systems. *Human Factors*. 1995;37(1):32–64.

24. Rubin J, Potes C, Xu-Wilson M, et al. An ensemble boosting model for predicting transfer to the pediatric intensive care unit. *Int J Med Inform*. 2018;112:15–20.

25. Umscheid CA, Betesh J, VanZandbergen C, et al. Development, implementation, and impact of an automated early warning and response system for sepsis. *J Hosp Med*. 2015;10:26–31.

26. de Groot JF, Damen N, de Loos E, et al. Implementing paediatric early warning scores systems in the Netherlands: future implications. *BMC Pediatr*. 2018;18:128.

27. Tucker KM, Brewer TL, Baker RB, et al. Prospective evaluation of a pediatric inpatient early warning scoring system. *J Spec Pediatr Nurs*. 2009;14:79–85.

28. Fleming S, Thompson M, Stevens R, et al. Normal ranges of heart rate and respiratory rate in children from birth to 18 years of age: a systematic review of observational studies. *Lancet*. 2011;377:1011–1018.

29. Carpenter J, Bithell J. Bootstrap confidence intervals: when, which, what? A practical guide for medical statisticians. *Stat Med*. 2000;19:1141–1164.

30. Robin X, Turck N, Hainard A, et al. pROC: an open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics*. 2011;12:77.

31. Wickham H. *ggplot2: Elegant Graphics for Data Analysis*. New York: Springer-Verlag; 2016.

32. Bonafide CP, Roberts KE, Priestley MA, et al. Development of a pragmatic measure for evaluating and optimizing rapid response systems. *Pediatrics*. 2012;129:e874–e881.

33. Zhai H, Brady P, Li Q, et al. Developing and evaluating a machine learning based algorithm to predict the need of pediatric intensive care unit transfer for newly hospitalized children. *Resuscitation*. 2014;85:1065–1071.

34. Choi E, Schuetz A, Stewart WF, et al. Using recurrent neural network models for early detection of heart failure onset. *J Am Med Inform Assoc*. 2017;24:361–370.

35. Rajkomar A, Oren E, Chen K, et al. Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine*. 2018;1(1):18.