

# Incorporating Observed Physiological Data to Personalize Pediatric Vital Sign Alarm Thresholds

Sarah Poole and Nigam Shah

Stanford Center for Biomedical Informatics Research, Stanford University, Stanford, CA, USA.

Biomedical Informatics Insights  
Volume 11: 1–8  
© The Author(s) 2019  
Article reuse guidelines:  
sagepub.com/journals-permissions  
DOI: 10.1177/1178222618818478



**ABSTRACT:** Bedside monitors are intended as a safety net in patient care, but their management in the inpatient setting is a significant patient safety concern. The low precision of vital sign alarm systems leads to clinical staff becoming desensitized to the sound of the alarm, a phenomenon known as alarm fatigue. Alarm fatigue has been shown to increase response time to alarms or result in alarms being ignored altogether and has negative consequences for patient safety. We present methods to establish personalized thresholds for heart rate and respiratory rate alarms. These thresholds are first chosen based on patient characteristics available at the time of admission and are then adapted to incorporate vital signs observed in the first 2 hours of monitoring. We demonstrate that the adapted thresholds are similar to those chosen by clinicians for individual patients and would result in fewer alarms than the currently used age-based thresholds. Personalized vital sign alarm thresholds can help to alleviate the problem of alarm fatigue in an inpatient setting while ensuring that all critical vital signs are detected.

**KEYWORDS:** Alarm fatigue, Vital signs, Inpatient monitoring, Clinical informatics

**RECEIVED:** September 23, 2018. **ACCEPTED:** October 16, 2018.

**TYPE:** Original Research

**FUNDING:** The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported in part by NLM grant R01 LM011369 and by a National Institutes of Health National Center for Advancing Translational Science Clinical and Translational Science Award (UL1TR001085).

**DECLARATION OF CONFLICTING INTERESTS:** The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

**CORRESPONDING AUTHOR:** Nigam Shah, Stanford Center for Biomedical Informatics Research, Stanford University, MSOB, 1265 Welch Road, Stanford, CA 94305, USA. Email: nigam@stanford.edu

## Introduction

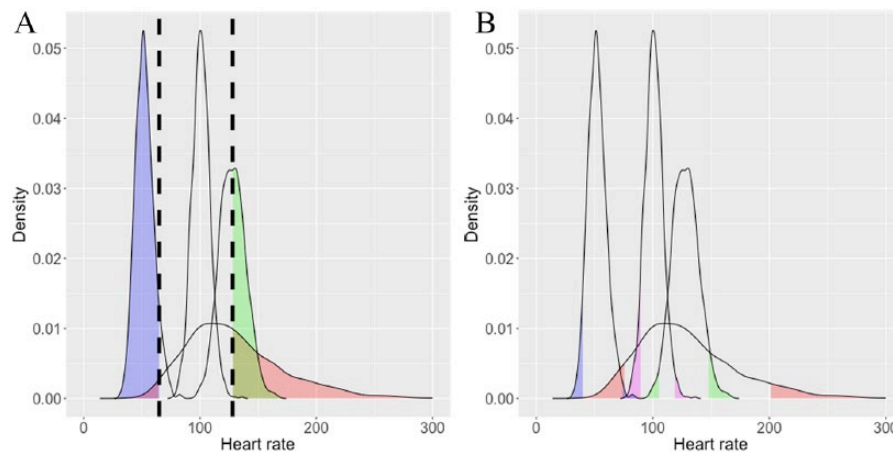
In an inpatient setting, bedside monitors are used to track the vital signs of critically ill patients. The monitors are programmed to sound an audible alarm when these vital signs move outside of a predefined range, alerting clinical staff that the patient may be deteriorating. These alarm systems are typically set up to have very high recall, ensuring that no concerning vital signs are missed. Operating at a high recall often results in a system with low precision, with a large number of unnecessary alarms sounding. Exposure to many unnecessary alarms leads to clinical staff becoming desensitized to the sound of the alarm, a phenomenon known as alarm fatigue. Alarm fatigue has been shown to increase response time to alarms both in the short term, such as during a nurse's shift,<sup>1</sup> and in the long term, such as over the course of a nurse's career,<sup>2</sup> which can result in negative outcomes for patients. The US Food and Drug Administration (FDA) reported over 500 alarm-related patient deaths over a 5-year period,<sup>3</sup> and alarm configuration policies and practices have regularly made the list of top patient safety concerns.<sup>4–6</sup> Surveys of nursing staff across the United States showed that an increasing proportion of nurses believe that alarm fatigue is disruptive to care, reduces trust in alarms, and can lead caregivers to inappropriately deactivate alarms.<sup>7–9</sup> Implementing measures to reduce alarm fatigue in the inpatient environment is therefore of critical importance.<sup>3,10</sup>

Previous attempts to alleviate alarm fatigue have focused on decreasing the number of bedside alarms. Methods for alarm reduction include a system to escalate alarms to pages sent directly to clinical staff,<sup>11</sup> ensuring that alarms are not missed even if the relevant member of the care team is not physically in the unit. Unnecessary alarms can be decreased by ensuring that bedside monitors are reserved for the most critically ill

patients. To this end, one institution moved to ensure that monitoring of all low-risk patients was discontinued in a timely manner,<sup>12</sup> and another institution implemented a nurse-managed monitor discontinuation process in an inpatient unit.<sup>13</sup> A subset of premature ventricular contraction alarms were disabled in one institution after studies showed that these alarms contribute a large amount of the noise load experienced in an inpatient unit, despite being largely non-actionable.<sup>14</sup> The sound of an alarm is often configured to reflect the acuity of the event, so efforts have been taken to reconfigure the acuity of each alarm type, ensuring that only the most serious events will trigger a high-acuity alarm sound. Wearables have even been used to help nurses quickly triage clinical alarms.<sup>15</sup> Approaches to alarm fatigue within the pediatric setting have been proposed by Karnik and Bonafide.<sup>16</sup> These have included limiting the patients who are monitored to those at the highest level of risk, careful application and regular changing of electrode sensors to avoid artifact alarms, and carefully choosing the thresholds for alarms to minimize unnecessary alarms. Existing research addressing alarm fatigue includes only limited evaluation of the safety and efficacy of the proposed measures. The lack of robust evaluations can be attributed to the absence of large sets of gold-standard labels that indicate which alarms are crucial for patient safety and which alarms are unnecessary and should be suppressed.<sup>17–19</sup> As a result, evaluation has typically been concentrated on the expected number of alarms that will be suppressed, with no consideration given to the appropriateness of this suppression.

This study focuses on addressing alarm fatigue in a pediatric population, as the change in vital signs with age results in higher inter-patient variation compared to the adult population.<sup>17,20</sup> Research has shown that the default heart rate alarm thresholds in a pediatric hospital often fall near the 50th





**Figure 1.** Modeled heart rates for patients from a single age group (8-12 years). Each distribution represents a separate patient. Distributions are obtained by sampling 1000 points from a lognormal distribution parameterized by the observed mean and standard deviation of heart rate over the patient's first 24 hours of monitoring. (A) The dashed vertical lines represent the 5th and 95th percentiles of heart rate observed for patients in this age group, the alarm thresholds as suggested by our previous study. The shaded regions are those that would trigger alarms. We see that the left-most distribution is almost entirely outside these alarm thresholds, so most of heart rate measurements from this patient would trigger a low heart rate alarm. We also see that the heart rate distribution from one of the patients is so narrow that very few heart rate measurements would trigger an alarm. The variation in vital signs among patients of the same group illustrates the weakness of alarm thresholds that are optimized for groups of patients and demonstrates the need for personalized alarm thresholds. (B) The heart rate values that would trigger an alarm if the observed 1st and 99th percentiles of heart rate are used as alarm thresholds. The shaded regions are those that would trigger alarms. Many fewer alarms would sound under this model, and alarms would be triggered for different thresholds for each patient.

percentile of heart rate values observed in the population,<sup>18,21</sup> causing alarms to be triggered for over half of the observed heart rates for some patients.

Our previous work used nurse-charted vital sign data to choose population-specific vital sign alarm thresholds.<sup>18</sup> The thresholds calculated in this study were implemented in a pilot study in a 20-bed cardiac step-down unit, decreasing the number of alarms per monitored-bed-day, with no negative impact on patient safety observed.<sup>22</sup> The age-group-based thresholds calculated were a step toward decreasing alarm fatigue; however, these thresholds had to be chosen conservatively to ensure that they were appropriate for the entire age group. Figure 1 shows several patients from the 8-to-12-year age group for whom the alarm thresholds are not appropriate. If we can instead model the expected distribution of vital sign values for each patient, we can choose alarm thresholds specific to each patient. These personalized alarm thresholds can be set at the 1st and 99th percentiles of the patient's vital sign values, ensuring that only the most extreme events will trigger an alarm.

This study aims to address vital sign alarm fatigue by finding personalized alarm thresholds for heart rate and respiratory rate alarms. Models are trained to find vital sign alarm thresholds on a patient-by-patient basis, rather than using patient groups. The personalized thresholds are initially built using data available at admission and are adjusted after observation of the patient's own vital signs over a 2-hour period. A method to evaluate the resulting patient-specific alarm thresholds is also introduced. We identify alarms that had a setting chosen by a physician specifically for the patient of interest and use

these alarm settings as the best available standard against which to evaluate our proposed thresholds.

## Methods

### Data

Two main sources of data were used for this study. The Philips Research Data Export (RDE) system at Stanford's Lucille Packard Children's Hospital (LPCH) has been recording vital sign waveforms for every patient who has had their vital signs monitored, both in intensive care units and on floor units, for the past several years. An extract from this system, containing 3.5 years worth of data (December 5, 2012 to April 20, 2016) has been made available for research purposes. The extract contains once per minute average heart rate and respiratory rate, as well as records of any vital sign alarms that were triggered. These data have been combined with an extract from the electronic medical record (EMR), obtained through StaRR, the Stanford Medicine Research Data Repository.<sup>23</sup> StaRR contains patient demographics and clinical data including ICD-9 (International Classification of Diseases, Ninth Revision) codes and medication records. These datasets were linked using patient medical record numbers or using data showing which patient was in a specific bed location at the time data is available. This study was approved by the Stanford Institutional Review Board.

*Diagnosis information.* Baseline vital signs are likely to vary based on the patient's diagnosis.<sup>24,25</sup> To incorporate this into our model, we use two surrogates for diagnosis: (1) the

diagnosis-related group (DRG) and (2) the unit that the patient is being treated in.

The EMR extract includes DRGs,<sup>26</sup> which are designed to group patients according to the medical services they receive, but can also be used to provide a rough grouping by clinical complaint. A total of 45 DRGs were present as admit diagnoses in the cohort of interest. DRGs that contained less than 10 observations were combined into an “OTHER” group, leaving a total of 22 distinct DRGs. One-hot encoding was used to convert this categorical variable into 22 variables with Boolean values, with the constraint that for a given sample only one of the values can be set to 1. The inpatient units at LPCH are arranged such that patients with particular care needs are grouped together. For example, one unit typically houses patients with cardiac issues, while patients with pulmonary-related problems are cared for in another unit. The department in which the patient is located was used as a feature in our model, as it provides a rough grouping according to diagnosis.

*Imputing patient weight.* Including patient weight in the model is an important consideration, as body size is known to impact heart rate. Some patients did not have any weight data available or only had weight recorded distantly in time. We used standard pediatric growth charts<sup>27</sup> to impute patient weights. If a weight measurement was available for a patient, growth charts were used to find which percentile the patient fell into for their age at the time that weight was recorded. The growth charts were then used to determine the weight that the patient would be at the time of vital sign recording, assuming that they remained in the same percentile. If patients had multiple weights recorded, the mean percentile was used. In total, 466 patients had no weight data recorded, so were assumed to be at the 50th percentile of weight for their age. The percentile found for each patient was also recorded and used as an input to the models.

### Modeling personalized alarm thresholds

There are two outcomes of interest for each patient, corresponding to the high and low alarm thresholds. The proposed ideal values for these are the 1st and 99th percentiles of the patient’s observed vital signs over the first day of hospitalization. The first 24 hours of monitoring for each patient was isolated and processed. All data within this 24-hour period were considered, regardless of whether data were continuously recorded or included periods of missing vital sign data. Four values were extracted for each patient: the mean, standard deviation, 1st percentile, and 99th percentile of the vital sign data available in this 24-hour period.

*Lognormal assumption.* As stated above, the proposed ideal values for the low and high alarm thresholds are the 1st and 99th percentiles of the patient’s vital signs over the first day of

hospitalization. Rather than using these values directly as the outcomes of interest, we model each patient’s vital signs over the first day of hospitalization using a lognormal distribution. Figure 2 shows that the 1st and 99th percentiles of the patient’s heart rate and respiratory rate are able to be accurately recovered using this lognormal assumption. The lognormal distribution is parameterized by the mean and the standard deviation, as shown in equation (1), where  $VS$  refers to the vital sign of interest, which in our study is heart rate or respiratory rate.

$$\log(VS) \sim N(\mu_{VS}, \sigma_{VS}^2) \quad (1)$$

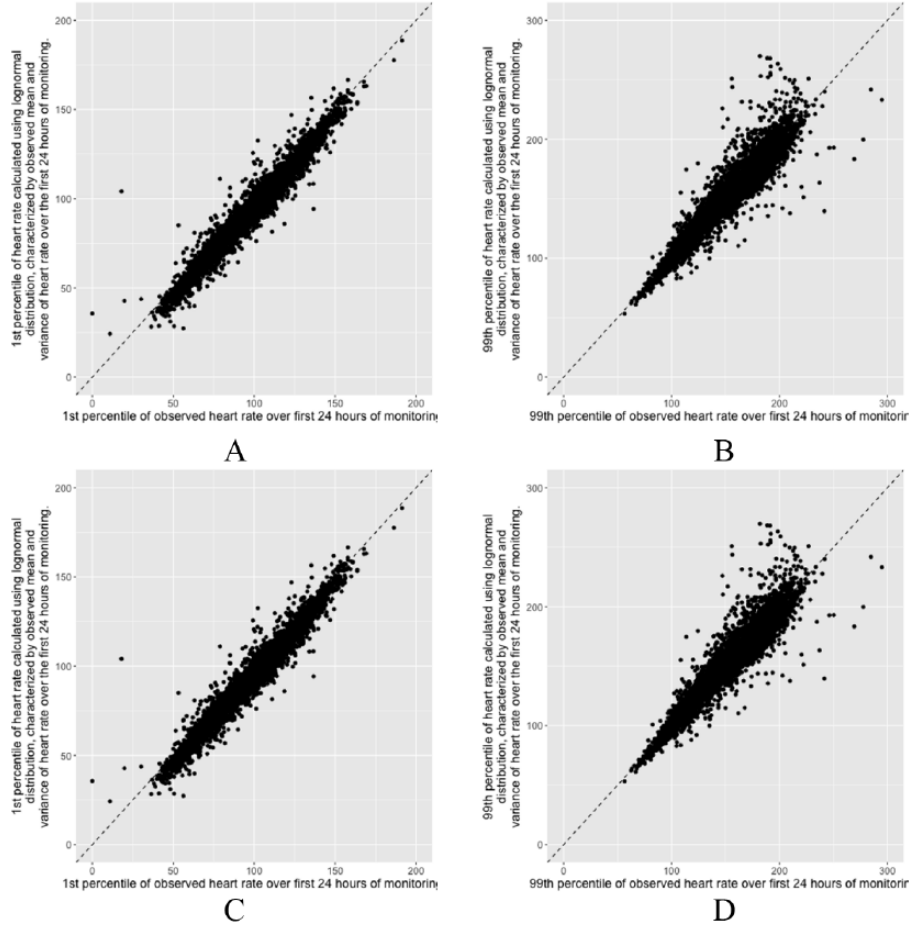
We build two models for each vital sign, one to recover the mean value and the second to recover the standard deviation. The outputs of these models are used as the parameters of patient-specific lognormal distributions, from which the expected 1st and 99th percentiles of the patient’s vital signs are found, as shown in equations (2) and (3), where the constant defines how many standard deviations the percentile of interest is from the mean. These percentiles are the proposed alarm thresholds.

$$VS_{1\%} = e^{-2.3263\sigma + \mu} \quad (2)$$

$$VS_{99\%} = e^{2.3264\sigma + \mu} \quad (3)$$

*Model structure.* Modeling the change in expected vital signs with age is an important component of building personalized alarm thresholds. We use loess models to capture the non-linear variation in the mean and the variance of the vital signs with age. The output of the loess models is then used as input to Bayesian additive regression tree (BART) models, along with additional demographic (age, weight, gender, ethnicity, and race) and diagnostic features (DRG and hospital department). Five-fold cross validation is used, training the loess models on data from 80% of patients and using these models to obtain proposed thresholds for the remaining 20% of patients. By combining the testing data from each fold, we are able to obtain proposed thresholds for every patient. The parameters of the BART model were chosen using a second five-fold cross validation within each outer fold. The output from these models are used as the parameters of lognormal distributions to find the expected 1st and 99th percentiles of vital signs for each patient. These percentile values are proposed as personalized alarm thresholds and will be referred to as *personalized thresholds*.

The BART models calculate the posterior distribution around the variable of interest, allowing us to obtain not only the expected value of the variable but also an estimate of the error in this value. The BART model for the mean vital sign value therefore gives us both  $\hat{\mu}_{VS}$  and  $Var(\hat{\mu}_{VS})$ , and the model for the variance of the vital sign gives us  $\hat{\sigma}_{VS}^2$  and  $Var(\hat{\sigma}_{VS}^2)$ .



**Figure 2.** Comparison of actual 1st and 99th percentiles of vital sign data observed over the first 24 hours of monitoring with the expected percentiles calculated using a lognormal model characterized by the observed mean and variance for the vital sign of interest. (A) 1st percentile of heart rate,  $R^2 = 0.929$ . (B) 99th percentile of heart rate,  $R^2 = 0.857$ . (C) 1st percentile of respiratory rate,  $R^2 = 0.933$ . (D) 99th percentile of respiratory rate,  $R^2 = 0.749$ .

*Adapting thresholds.* The BART models allow us to estimate the mean and standard deviation of the patient's vital signs, given the characteristics available at the time of admission. Once the patient has had their vital signs monitored for 2 hours, we have additional data to inform the expected distribution of vital signs. This 2-hour period length was chosen to give enough time for the patient's condition to begin to stabilize after being admitted, while still ensuring that the personalized thresholds are available as quickly as possible. We can use the output from the BART models to define the prior on each parameter of interest and update this distribution using the data observed, resulting in an expected distribution that has been informed by the patient's own data. Using the 1st and 99th percentiles of this distribution as vital sign alarm thresholds ensures highly personalized thresholds.

Our prior distribution is defined in equation (1). To ensure that the posterior distribution is of the same form, we model the mean using a normal distribution and model the variance using an inverse-gamma distribution.

$$\mu_{VS} \sim N(\hat{\mu}_{VS}, k\hat{\sigma}_{VS}^2) \quad (4)$$

$$\sigma_{VS}^2 \sim \text{Inv-Gamma}(\alpha, \beta) \quad (5)$$

The inverse-gamma distribution is characterized by  $\alpha$  and  $\beta$ , but from the BART model, we have the mean and variance of this distribution. We can use  $\hat{\sigma}_{VS}^2$  and  $\text{Var}(\hat{\sigma}_{VS}^2)$  to obtain  $\hat{\alpha}$  and  $\hat{\beta}$  as shown in equations (6) and (7).

$$\hat{\alpha} = \frac{(\hat{\sigma}_{VS}^2)^2}{\text{Var}(\hat{\sigma}_{VS}^2)} + 2 \quad (6)$$

$$\hat{\beta} = \frac{(\hat{\sigma}_{VS}^2)^3}{\text{Var}(\hat{\sigma}_{VS}^2)} + \hat{\sigma}_{VS}^2 \quad (7)$$

A Bayesian update can now be performed to incorporate the  $n_{data}$  points of observed data and obtain a posterior distribution for the vital sign of interest. To do this, we need one additional parameter:  $n_0$ , the number of data points used to calculate the prior distribution. As our distribution came from the BART model and not directly from the data, the appropriate value for  $n_0$  is not clear. Experimentation showed that the posterior distributions were relatively constant when varying  $n_0$  between 1 and 0.5 times the number of unique patients in the training dataset. The parameters defining the posterior

distribution can then be calculated using the following equations. A tilde is used to indicate the updated parameters of the distribution.

$$\mu_{data} = \frac{1}{n_{data}} \sum_{i=1}^{n_{data}} VS_i \quad (8)$$

$$\sigma_{data}^2 = \frac{1}{n_{new}} \sum_{i=1}^{n_{new}} (VS_i - \mu_{data})^2 \quad (9)$$

$$\tilde{\alpha} = \hat{\alpha} \frac{n_{new}}{2} \quad (10)$$

$$\tilde{\beta} = \hat{\beta} + \frac{\hat{\sigma}_{data}^2 n_{data}}{2} + \frac{n_{data} n_0 (\mu_{data} - \hat{\mu})^2}{2(n_{data} + n_0)} \quad (11)$$

$$\tilde{\sigma}^2 = \frac{\tilde{\beta}}{(\tilde{\alpha} - 1)} \quad (12)$$

$$\tilde{\mu} = \frac{n_{data} \mu_{data}}{n_{data} + n_0} + \frac{n_0 \hat{\mu}}{n_{data} + n_0} \quad (13)$$

We can then use equations (2) and (3) to obtain the new suggested alarm thresholds from these updated estimates of the distribution mean and variance. The alarm thresholds obtained from this adaptive step are referred to as *adapted thresholds*.

### Evaluations

The performance of the models can be evaluated by directly comparing the estimated 1st and 99th percentiles of the vital signs to the observed values for each patient. As a second evaluation, a record of the alarms that sounded can be used to estimate the proportion of alarms that would be suppressed if the suggested thresholds were implemented. However, both of these metrics fail to evaluate the clinical appropriateness of the proposed thresholds.

*Physician-selected alarm thresholds.* To estimate the appropriateness of the proposed alarm thresholds, we use the record of alarms that previously sounded in LPCH to find alarm thresholds that do not match the default values, indicating that clinical staff manually chose this threshold for the patient. Alarms with a setting that matched the policy at any point of the study were excluded. The most commonly used setting for each age group was identified, and alarms with this setting were also excluded.

*Total number of alarms.* A record of the alarms that sounded is available, including the reading that caused the alarm to be triggered. We can examine these historic alarms to determine whether each alarm would have been triggered if the proposed thresholds were in use. As we only have a record of alarms that were triggered, we are unable to calculate the number of additional alarms that would have been triggered if the new thresholds were used. To estimate the additional alarm load associated with the proposed thresholds, we compare the proposed thresholds to the age-based thresholds that were implemented as part of our previous study.<sup>18,22</sup> Where the new thresholds are wider than the existing thresholds, no additional alarms would have sounded. If the new thresholds are narrower than the existing thresholds, implementation could lead to additional alarms.

### Results

As shown in Figure 2, using the mean and standard deviation of a patient's heart rate over a 24-hour period as parameters in a lognormal distribution gives an accurate estimate of the 1st and 99th percentiles of the heart rate over this period.

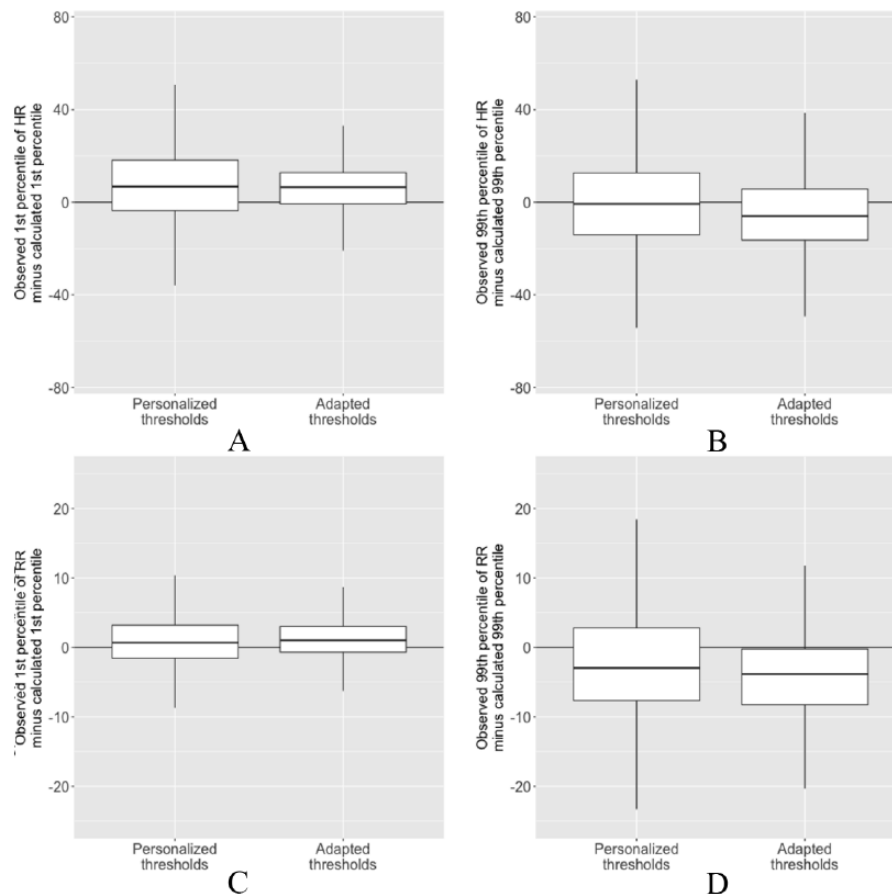
Figure 3 shows that the models are able to recover the 1st and 99th percentiles of the observed data. Incorporating the patient's own data generally improves the estimate of the percentiles. This figure also includes a comparison of the age-based thresholds proposed in our previous work with the observed percentiles.<sup>18</sup>

Figure 4 compares the thresholds from each source to the physician-selected alarm thresholds that were recorded in the RDE dataset. The low errors between the physician-selected thresholds and those produced by the models suggest that the use of 1st and 99th percentiles as threshold values is an appropriate one.

Table 1 shows the proportion of patients for whom the new proposed thresholds are wider than the age-based thresholds proposed in our previous work.<sup>18</sup> Wider thresholds will result in fewer alarms. Figure 5 shows the number of alarms that would sound if each of the different sets of thresholds had been used. This analysis uses the list of alarms that were actually triggered over a 3.5-year period, so it is not a complete count of all alarms that would have sounded if the proposed thresholds had been implemented.

### Discussion

The aim of this study was to find personalized vital sign alarm thresholds for hospitalized pediatric patients without requiring input from clinical staff. One of the challenges to achieving this aim was determining an appropriate target for these thresholds. Due to the lack of a large set of labeled alarms to use in training a model, we needed to use the patient's observed vital signs to choose targets for the threshold values. We chose to use the 1st and 99th percentiles of the observed vital sign values as the target alarm thresholds. Models were



**Figure 3.** Comparison of the 1st and 99th percentiles of heart rate and respiratory rate, as observed over the first 24 hours of monitoring, with the 1st and 99th percentiles calculated using our method. We propose these calculated percentiles as the personalized and adapted thresholds. (A) Comparison of the observed 1st percentile of heart rate with the calculated 1st percentile. (B) Comparison of the observed 99th percentile of heart rate with the calculated 99th percentile. (C) Comparison of the observed 1st percentile of respiratory rate with the calculated 1st percentile. (D) Comparison of the observed 99th percentile of respiratory rate with the calculated 99th percentile.

trained to choose personalized thresholds using the patient characteristics available at the time the patient was admitted to the hospital. These thresholds were then adapted using the patient's observed vital signs over the first 2 hours of monitoring. As shown in Figure 3, these models were able to accurately estimate the 1st and 99th percentiles of each patient's vital signs.

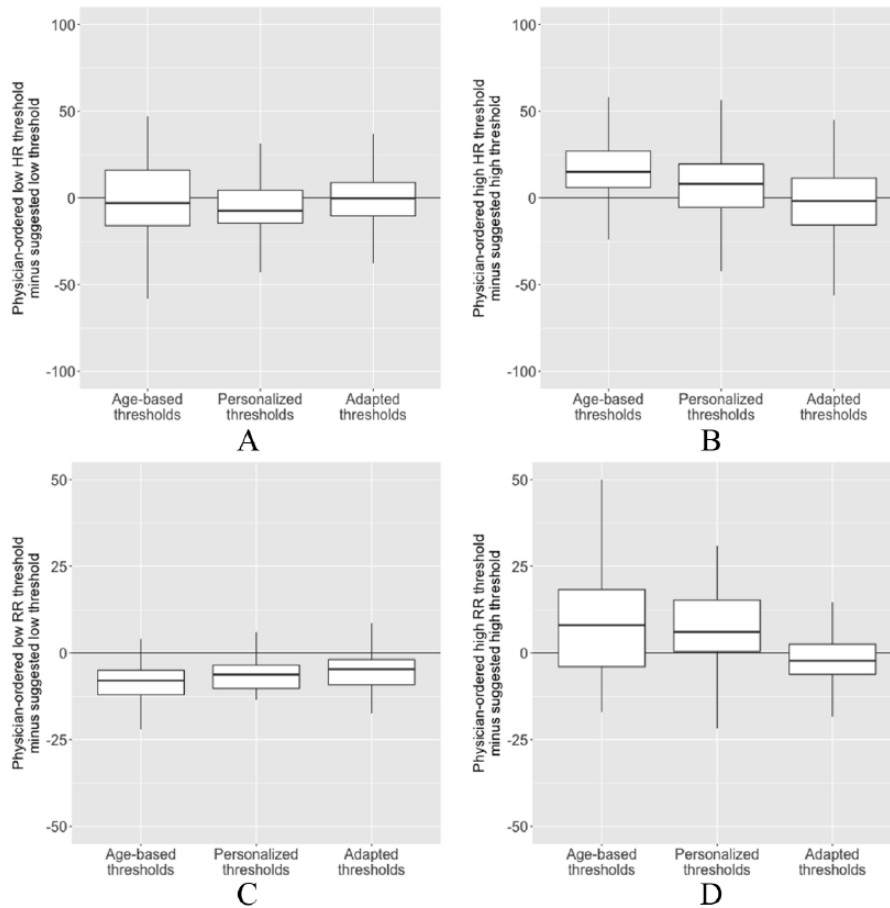
Being able to predict the observed percentiles of each patient's vital signs is only useful for this study if these percentiles are appropriate for use as alarm thresholds. By identifying alarms where the threshold did not match any default values, we were able to observe cases where an alarm threshold was chosen specifically for the patient of interest. We compared these physician-selected alarm thresholds to those produced by our models. As shown in Figure 4, our models propose thresholds that are similar to the physician-selected thresholds. Incorporating the patient's own data generally gives us a better estimate of the physician-selected thresholds. The agreement between our proposed thresholds and these physician-selected thresholds suggests that the choice of 1st and 99th percentiles of the observed vital signs as alarm threshold targets was an appropriate one. It also suggests that implementing our pro-

posed thresholds would decrease the number of patients requiring physician-selected thresholds.

A set of gold-standard alarm labels would allow more rigorous evaluation of the proposed alarm thresholds. In the absence of such labels, and given that there is currently no standard method for evaluating alarm thresholds before implementation, we believe that our attempt to assess the proposed alarm thresholds using physician-selected thresholds was appropriate and sufficient.

As shown in Table 1, the personalized and adapted alarm thresholds are generally wider than the age-group-based thresholds. Therefore, implementing our proposed thresholds will result in fewer alarms being triggered. The exception is low heart rate alarms, where the majority of patients are given narrower thresholds. The narrower threshold for low heart rate alarms is unsurprising, as during the initial study it was noted that very few low heart rate values triggered alarms, and the low alarm thresholds were raised slightly.

This study involves updating the alarm thresholds to incorporate the patient's own data at a single point in time. Although the same methodology could be used to repeatedly or even continuously update the alarm thresholds to better fit the



**Figure 4.** Comparison of physician-selected alarm thresholds with the personalized and adapted thresholds we derive. Comparison with the age-based thresholds from our previous study are also shown as an indication of current practice. (A) Comparison of physician-selected thresholds for low heart rate with the proposed personalized and adapted thresholds. (B) Comparison of physician-selected thresholds for high heart rate with the proposed personalized and adapted thresholds. (C) Comparison of physician-selected thresholds for low respiratory rate with the proposed personalized and adapted thresholds. (D) Comparison of physician-selected thresholds for high respiratory rate with the proposed personalized and adapted thresholds.

**Table 1.** Proportion of patients who have wider alarm thresholds using the proposed thresholds compared to the age-group-based thresholds. Wider thresholds will result in fewer alarms.

THRESHOLD SOURCE	HEART RATE (%)		RESPIRATORY RATE (%)	
	LOW	HIGH	LOW	HIGH
Personalized thresholds	92	82	95	89
Adapted thresholds	81	64	93	81


patient’s data, we chose to limit to one update. This decision was made because of the difficulty in distinguishing between a slow trend in a vital sign due to the patient’s state improving and a trend due to deterioration. Continually adapting the thresholds with observed data may result in thresholds that move as the patient deteriorates, never sounding an alarm. A dataset that included precise labels of patient state would allow models to be built to distinguish between trends due to clinical improvement and those due to clinical deterioration, address-

ing this problem and ensuring that thresholds could be regularly updated without fear of missing patient deterioration.

The age-based alarm thresholds from our previous study were easily implemented, as the Philips monitors allow profiles to be pre-programmed with alarm thresholds and activated when a new patient is admitted. Implementation of the thresholds proposed in this study will be more difficult, as the methods rely on data from both the bedside monitor and the EMR. Currently, the Epic EMR is able to interface with the bedside monitors, pulling in data in real time to display in the patient’s record. To implement the proposed thresholds, data would have to travel in the opposite direction to set the alarm thresholds within the bedside monitor, a functionality that is not currently supported. To avoid creating such a link, the alarm functionality could be moved to within the EMR itself, combining patient characteristics with the data passed by the bedside monitor to determine when to sound an alarm.

### Conclusions

In conclusion, this study describes methods to choose personalized heart rate and respiratory rate alarm thresholds for



**Figure 5.** Number of alarms that would have sounded using the different sets of proposed thresholds. Note that only alarms that were actually historically triggered are part of this analysis. The dashed lines show the total number of alarms that were triggered. The alarms making up this analysis span a 3.5-year period.

pediatric inpatients. These thresholds are initially chosen based on patient characteristics available at admission and are adapted to incorporate vital signs observed over the first 2 hours of vital sign monitoring. The resulting adapted thresholds are similar to physician-selected thresholds and result in fewer alarms than the currently used thresholds. The adoption of personalized thresholds for interpretation of vital signs has the potential to reduce the number of unnecessary alarms sounding, alleviating alarm fatigue among clinical staff and improving patient outcomes.

## Acknowledgements

The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

## Author Contributions

SP and NS conceived of the study design, SP pre-processed the data and completed the analysis, SP and NS contributed to the manuscript and visualizations.

## REFERENCES

- Bonafide CP, Lin R, Zander M, et al. Association between exposure to nonactionable physiologic monitor alarms and response time in a children's hospital. *J Hosp Med.* 2015;10:345–351.

- Bonafide CP, Localio AR, Holmes JH, et al. Video analysis of factors associated with response time to physiologic monitor alarms in a children's hospital. *JAMA Pediatr.* 2017;171:524–531.
- The Joint Commission. Sentinel event alert issue 50: medical device alarm safety in hospitals. [https://www.jointcommission.org/sea\\_issue\\_50/](https://www.jointcommission.org/sea_issue_50/). Technical report. Published 2013.
- The Joint Commission. The Joint Commission announces 2014 national patient safety goal. [https://www.jointcommission.org/assets/1/18/JCP0713\\_Announce\\_New\\_NSPG.pdf](https://www.jointcommission.org/assets/1/18/JCP0713_Announce_New_NSPG.pdf). Technical report. Published 2014.
- The Joint Commission. 2015 hospital national patient safety goals. [https://www.jointcommission.org/assets/1/6/2015\\_HAP\\_NPSG\\_ER.pdf](https://www.jointcommission.org/assets/1/6/2015_HAP_NPSG_ER.pdf). Technical report. Published 2015.
- The Joint Commission. 2016 hospital national patient safety goals. [https://www.jointcommission.org/assets/1/6/2016\\_NPSG\\_HAP\\_ER.pdf](https://www.jointcommission.org/assets/1/6/2016_NPSG_HAP_ER.pdf). Technical report. Published 2016.
- Ruppel H, Funk M, Clark JT, et al. Attitudes and practices related to clinical alarms: a follow-up survey. *Am J Crit Care.* 2018;27:114–123.
- Petersen EM, Costanzo CL. Assessment of clinical alarms influencing nurses' perceptions of alarm fatigue. *Dimens Crit Care Nurs.* 2017;36:36–44.
- Casey S, Avalos G, Dowling M. Critical care nurses' knowledge of alarm fatigue and practices towards alarms: a multicentre study. *Intensive Crit Care Nurs.* 2018;48:36–41.
- Varpio L, Kuziemsy C, Macdonald C, King WJ. The helpful or hindering effects of in-hospital patient monitor alarms on nurses: a qualitative analysis. *Comput Inform Nurs.* 2012;30:210–217.
- Cvach MM, Frank RJ, Doyle P, Stevens ZK. Use of pagers with an alarm escalation system to reduce cardiac monitor alarm signals. *J Nurs Care Qual.* 2014;29:9–18.
- Dandoy CE, Davies SM, Flesch L, et al. A team-based approach to reducing cardiac monitor alarms. *Pediatrics.* 2014;134:e1686–e1694.
- Perrin K, Ernst N, Nelson T, Sawyer M, Pfoh E, Cvach M. Effect of a nurse-managed telemetry discontinuation protocol on monitoring duration, alarm frequency, and adverse patient events. *J Nurs Care Qual.* 2017;32:126–133.
- Srinivasa E, Mankoo J, Kerr C. An evidence-based approach to reducing cardiac telemetry alarm fatigue. *Worldviews Evid Based Nurs.* 2017;14:265–273.
- McFarlane DC, Doig AK, Agutter JA, Brewer LM, Syroid ND, Mittu R. Faster clinical response to the onset of adverse events: a wearable metacognitive attention aid for nurse triage of clinical alarms. *PLoS ONE.* 2018;13:e0197157.
- Karnik A, Bonafide CP. A framework for reducing alarm fatigue on pediatric inpatient units. *Hosp Pediatr.* 2015;5:160–163.
- Paine CW, Goel VV, Ely E, et al. Systematic review of physiologic monitor alarm characteristics and pragmatic interventions to reduce alarm frequency. *J Hosp Med.* 2016;11:136–144.
- Goel VV, Poole SF, Longhurst CA, et al. Safety analysis of proposed data-driven physiologic alarm parameters for hospitalized children. *J Hosp Med.* 2016;11:817–823.
- de Waele S, Nielsen L, Frassica J. Estimation of the patient monitor alarm rate for a quantitative analysis of new alarm settings. Paper presented at: 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC); November 6, 2014; Chicago, IL:5727–5730. New York, NY: IEEE.
- Lawless ST. Crying wolf: false alarms in a pediatric intensive care unit. *Crit Care Med.* 1994;22:981–985.
- Bonafide CP, Brady PW, Keren R, Conway PH, Marsolo K, Daymont C. Development of heart and respiratory rate percentile curves for hospitalized children. *Pediatrics.* 2013;131:e1150–e1157.
- Kipps AK, Poole SF, Slaney C, et al. Inpatient-derived vital sign parameters implementation: an initiative to decrease alarm burden. *Pediatrics.* 2017;140:e20162458.
- Lowe HJ, Ferris TA, Hernandez PM, Weber SC. Stride—an integrated standards-based translational research informatics platform. Paper presented at: AMIA Annual Symposium Proceedings; November 14–18, 2009; San Francisco, CA: 391. Bethesda, MD: American Medical Informatics Association.
- Fox K, Borer JS, Camm AJ, et al. Resting heart rate in cardiovascular disease. *J Am Coll Cardiol.* 2007;50:823–830.
- Kenzaka T, Okayama M, Kuroki S, et al. Importance of vital signs to the early diagnosis and severity of sepsis: association between vital signs and sequential organ failure assessment score in patients with sepsis. *Intern Med.* 2012;51:871–876.
- Fetter RB, Shin Y, Freeman JL, Averill RF, Thompson JD. Case mix definition by diagnosis-related groups. *Med Care.* 1980;18:1–53.
- National Center for Health Statistics. *CDC Growth Charts.* Hyattsville, MD: National Center for Health Statistics; 2000.