Comparing regression and neural network techniques for personalized predictive analytics to promote lung protective ventilation in Intensive Care Units

Rachael Hagan [a],[*], Charles J. Gillan [a], Ivor Spence [a], Danny McAuley [b], Murali Shyamsundar [b]

[a] School of Electrical and Electronic Engineering and Computer Science, Queen's University Belfast, Queen's Road, Queen's Island, Belfast, Northern Ireland, BT9 3DT, United Kingdom
[b] The Centre for Experimental Medicine, School of Medicine, Dentistry and Biological Sciences, Queen's University Belfast, 97 Lisburn Road, Belfast, Northern Ireland, BT9 7BL, United Kingdom

ARTICLE INFO

ABSTRACT

Mechanical ventilation is a lifesaving tool and provides organ support for patients with respiratory failure. However, injurious ventilation due to inappropriate delivery of high tidal volume can initiate or potentiate lung injury. This could lead to acute respiratory distress syndrome, longer duration of mechanical ventilation, ventilator associated conditions and finally increased mortality.

In this study, we explore the viability and compare machine learning methods to generate personalized predictive alerts indicating violation of the safe tidal volume per ideal body weight (IBW) threshold that is accepted as the upper limit for lung protective ventilation (LPV), prior to application to patients. We process streams of patient respiratory data recorded per minute from ventilators in an intensive care unit and apply several state-of-the-art time series prediction methods to forecast the behavior of the tidal volume metric per patient, 1 hour ahead.

Our results show that boosted regression delivers better predictive accuracy than other methods that we investigated and requires relatively short execution times. Long short-term memory neural networks can deliver similar levels of accuracy but only after much longer periods of data acquisition, further extended by several hours computing time to train the algorithm. Utilizing Artificial Intelligence, we have developed a personalized clinical decision support tool that can predict tidal volume behavior within 10% accuracy and compare alerts recorded from a real world system to highlight that our models would have predicted violations 1 hour ahead and can therefore conclude that the algorithms can provide clinical decision support.

1. Introduction

Mechanical ventilation to support patients with respiratory failure is one of the primary interventions in Intensive Care Units (ICU) globally. The projected US national estimates for mechanical ventilation suggest 790,257 hospitalizations of adult patients involving mechanical ventilation in 2005 with an estimated national cost of $27 billion accounting for 12% of all hospital costs [1]. In England, Wales and Northern Ireland during 2012, 44.5% of adult patients admitted to ICUs were mechanically ventilated and this equated to 116,000 cases [2].

Despite its undisputed role as a lifesaving organ support tool, injurious mechanical ventilation has been shown to both initiate and potentiate lung injury [3–5]. Injurious ventilation leads to lung injury secondary to excessive lung stress and strain due to both volume and pressure related factors. The harm from high tidal volume has been clearly demonstrated in the pivotal lung protective ventilation (LPV) trial from ARDSNet investigators that has established the role of LPV by using low tidal volume and appropriate use of positive end expiratory pressure in patients with acute respiratory distress syndrome (ARDS) [6]. There is accumulating supportive evidence for the use of LPV to prevent development of lung injury in all patients. A meta-analysis comparing LPV with conventional ventilation demonstrated a reduced incidence of lung injury as well as lower mortality in non-ARDS patients [7]. Similarly, a 28% reduction is seen in the occurrence of ARDS with an associated increase in ICU free days, hospital free days and mortality benefit [8]. Current data suggest clear harm of tidal volume > 10ml/kg

---

body weight with various systematic reviews suggesting a lower tidal volume to be associated with better clinical outcomes.

Development of significant deterioration secondary to worsening ARDS, or the development of ARDS, is characterised by a reduction in lung compliance [9]. A clinical decision support tool (CDS) which is efficient at determining breaches in thresholds of tidal volume for a set pressure will also be able to detect improvement in lung compliance. Seamless integration of a CDS that promotes compliance with LPV leading to early detection of physiological improvement to facilitate early implementation and support of ventilator weaning is crucial during periods of unprecedented pressure on critical care services such as the COVID-19 pandemic. LPV is essential to prevent further lung injury in patients with severe COVID-19 related respiratory failure while aggressive weaning is essential to reducing the duration of mechanical ventilation, currently a median of 17 days, to allow better utilisation of ventilatory resources [10,11].

Despite robust evidence, LPV is still poorly implemented with a third of patients receiving injurious ventilation [12]. A recent multicentre observational study has confirmed ongoing poor adherence to LPV, at 50% which further reduces to 15% if ARDS is unrecognised [13]. While previous studies of CDS include displaying safe thresholds for LPV at time of ventilator set up [14], change of ventilator parameters [15], or even default set ups [16], these solutions do not consider the mode of ventilation, provide decision support to detect a developing condition such as ARDS and ignore the potential changes in physiology between the ventilator interactions. The lack of current methods in improving practice provides further support to the use of automated systems to both diagnose and to use a physician independent alert system to change practice [17].

Intensive care units routinely collect vast volumes of physiological data on their patients. Prior research has shown that these streams have very valuable information buried in them. This trend started in the 1950s at the University of Southern California when physicians realized that the critically ill may have substantially better chances of survival when minute to minute monitoring of vital signs are available [18]. Research in the intervening years has delivered metrics and protocols which are now routinely used in the ICU [19]. In pioneering work, McGregor demonstrated the viability of monitoring physiological parameters to detect sleep apnea in neo-natal ICU [20,21].

Artificial Intelligence (AI) has shown promise in various fields and has potential in the field of ICU, which is a data-rich environment [22, 23]. Multiple studies in the areas of ECG analysis, delirium detection, sedation and identification of septic patients have highlighted the potential superiority of AI over routine clinical decision making [24]. In the field of mechanical ventilation, a treatment policy developed using AI techniques was shown to predict extubation readiness [23]. Similar benefit using clinical decision support tool has demonstrated in the management of patients with sepsis and detection of renal impairment [25,26]. Further, artificial intelligence has been proven to aid in the prediction of mortality and outcomes in ICUs [27–29].

It is the essential that alerts are clinically relevant and studies have suggested artefact related alert rates of 30% [31]. Multiple false alerts could lead to alarm fatigue and clinical inattention [16,32], increased response time [33] and there is evidence that clinicians over ride rate is up to 96% [30]. In a critical care study of a CDS to improve ventilation practices, the positive predictive value was only 59% [34]. This demonstrates the need to improve the quality of the alert generated, along with other measures such as pausing alerts for a specific individual and situation to reduce alert fatigue.

Around 40% of patients in intensive care are supported on invasive mechanical ventilation at any given hour. The ventilators can have many settings that need to be monitored closely and it is important to wean patients off ventilation as soon as possible to avoid dependency or infections. Researchers have utilised numerous machine learning techniques to aid in extubation and ventilator support [23], detect deteriorating patients [35], and distinguish patients at risk [36] and with

diseases such as ARDS and ALI [37].

While there has been a significant amount of work carried out analysing medical data and improving patient outcomes, there has, in so far as we know, been no work carried out of the prediction or monitoring of the tidal volume metric for mechanical ventilation, ensuring lung protection. The main novelty in our work lies in examining the viability of several machine learning methods to construct a personalized predictive alert system for violation of tidal volume thresholds during periods of mechanical ventilation. Our work uses patient data collected in an ICU over several years but is preclinical in the sense that there is no subsequent clinical intervention.

2. Methodology

In a previous paper we introduced the VILIAlert system [38], a quality improvement project, presenting an analysis of the performance of the database systems which underpin the collection of streams of patient respiratory data. The data was collected in the Regional Intensive Care Unit (RICU), Royal Victoria Hospital, Belfast over a three year period. RICU is the regional medical surgical ICU and the regional trauma centre. Hence the data and trends observed will be generalisable and representative of most ICUs. VILIAlert monitors patients in real-time by continuously computing a set of metrics from the received streams of ventilation data. Mathematical kernels process the data streams to allow patients to be monitored against the thresholds for lung protective ventilation (LPV). When a threshold is violated consistently (which we defined initially as a period of 60 min), an alarm is immediately raised and sent by SMS message to clinical staff. The aim of the VILIAlert system is to give the clinician an opportunity to intervene early and mitigate the potential damage of over ventilation.

In this paper we turn our attention to the challenge of predicting violations of the LPV thresholds based on the time series of patient readings from the ventilator. By adding this to the VILIAlert system we can send an alert to a clinician before potential damage from over ventilation starts to occur. We operated the VILIAlert system for nearly three years, recording in excess of four million per minute tidal volume readings for almost one thousand patients. We define the LPV violation threshold to be tidal volume per ideal body weight (IBW) greater than 8 ml/kg IBW. We employ a pipeline of well-known methods, including ensemble methods built on decision trees and the long short-term memory (LSTM) form of neural networks, the details of which are comprehensively described already in the literature and for which software in the Python language is available. These newer methods based on supervised learning have proven to be superior to the older ARIMA models [40].

The tidal volume data set for each patient is a set of N discrete observations V_j recorded at per minute intervals $t_j, j \in (1, ..., N)$. N ranges from several hours to many days. We divide each calendar day into 96 periods of 15 min duration and average the recorded tidal volumes within each period, in line with previous work in the field [41]. This smooths out random fluctuations in the data due to the phenomena of a patient taking random deep breaths or moving in the bed. We denote these averaged readings by $\overline{V}_i$. Following Friedmann [42], we may state the problem as follows. Output variable $\overline{V}_i$, dependent on a vector of n input variables $\mathbf{x} = x_1, x_2, x_3, ..., x_n$ through some function, $F(\mathbf{x})$, the form of which is unknown. This function represents the behavior of the respiratory system of the patient. However, we have a set of m observations, each of which associates one input vector $\mathbf{x}_i$ with one output,

$$F(\mathbf{x}_i) \rightarrow \overline{V}_i \ \forall i \ \in \ 1, 2, 3, 4, ..., m \tag{1}$$

This set of observations forms our training set from which we seek to find an approximation, $F^*(\mathbf{x})$ to the true function $F(\mathbf{x})$.
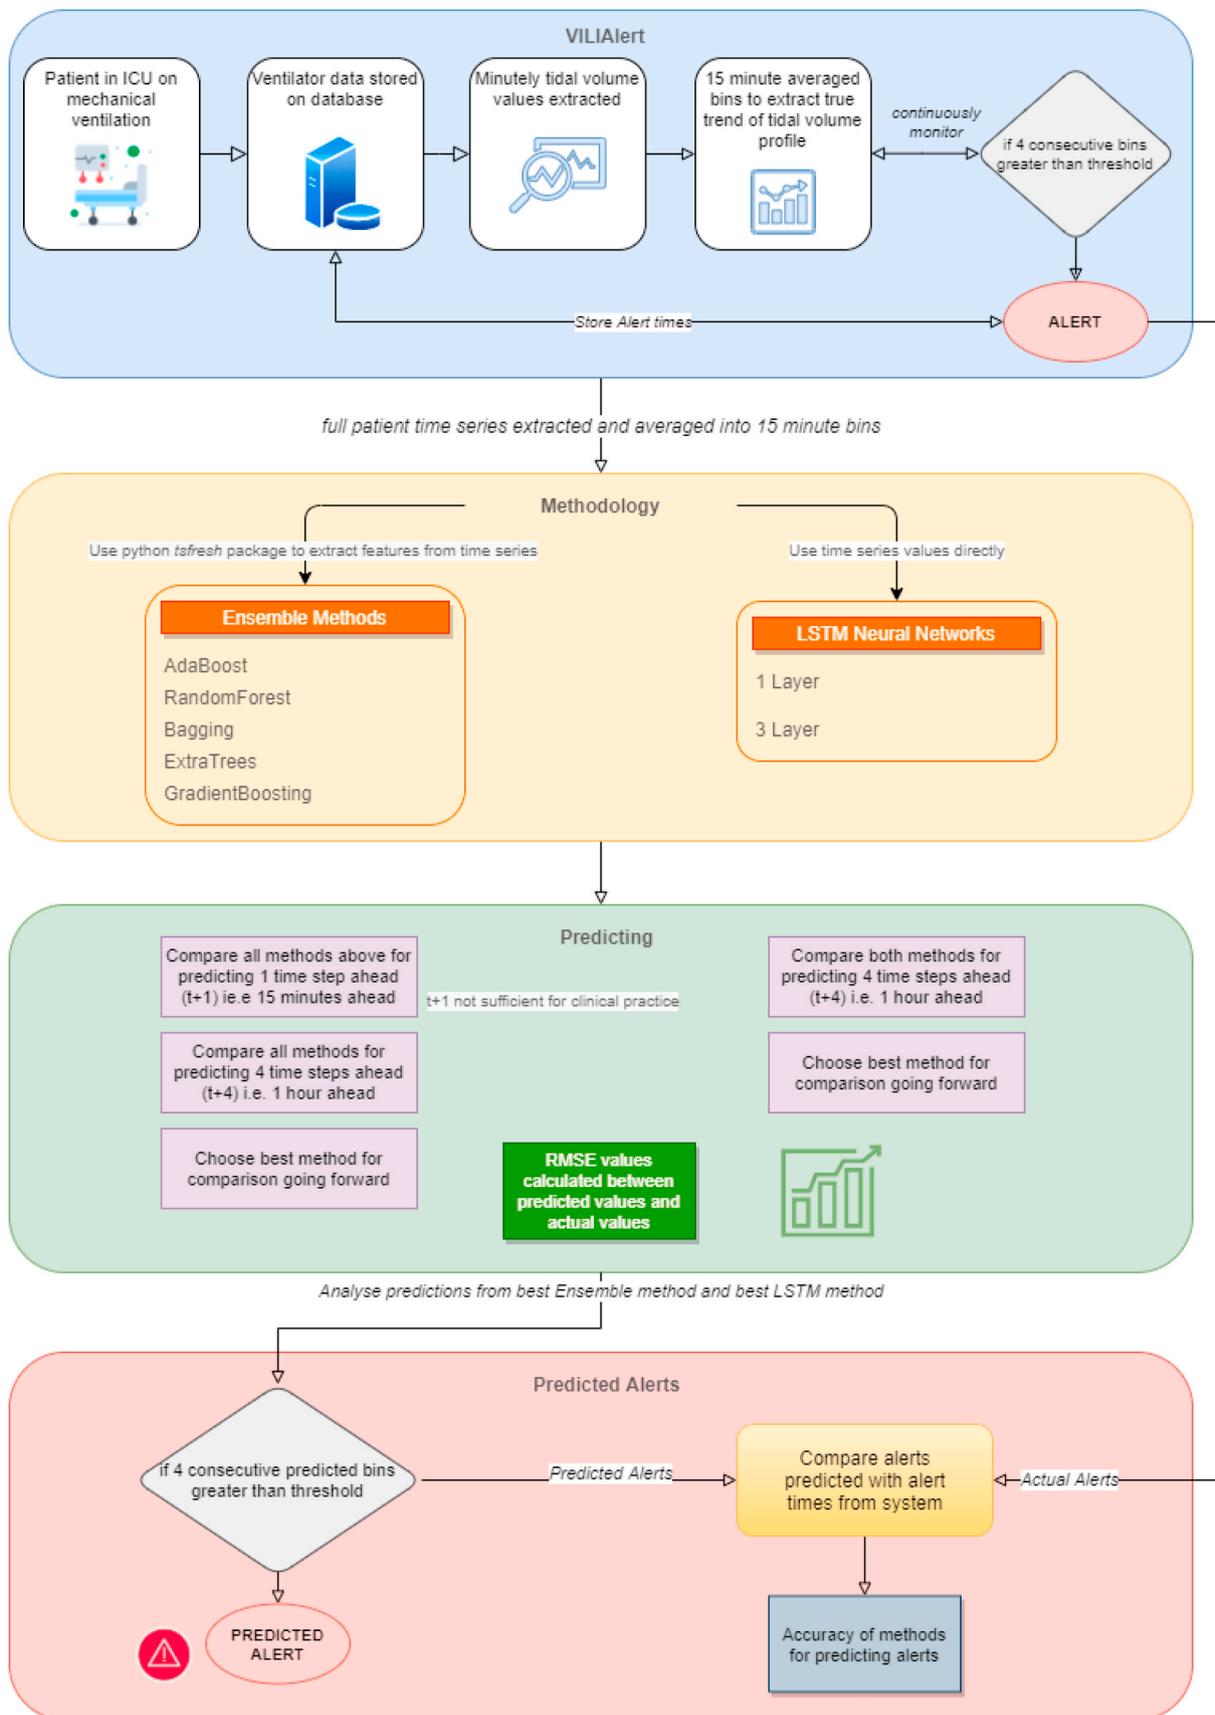
Fig. 1. Flow Diagram highlighting the process and methodology used as described in 2 for the prediction of tidal volume.
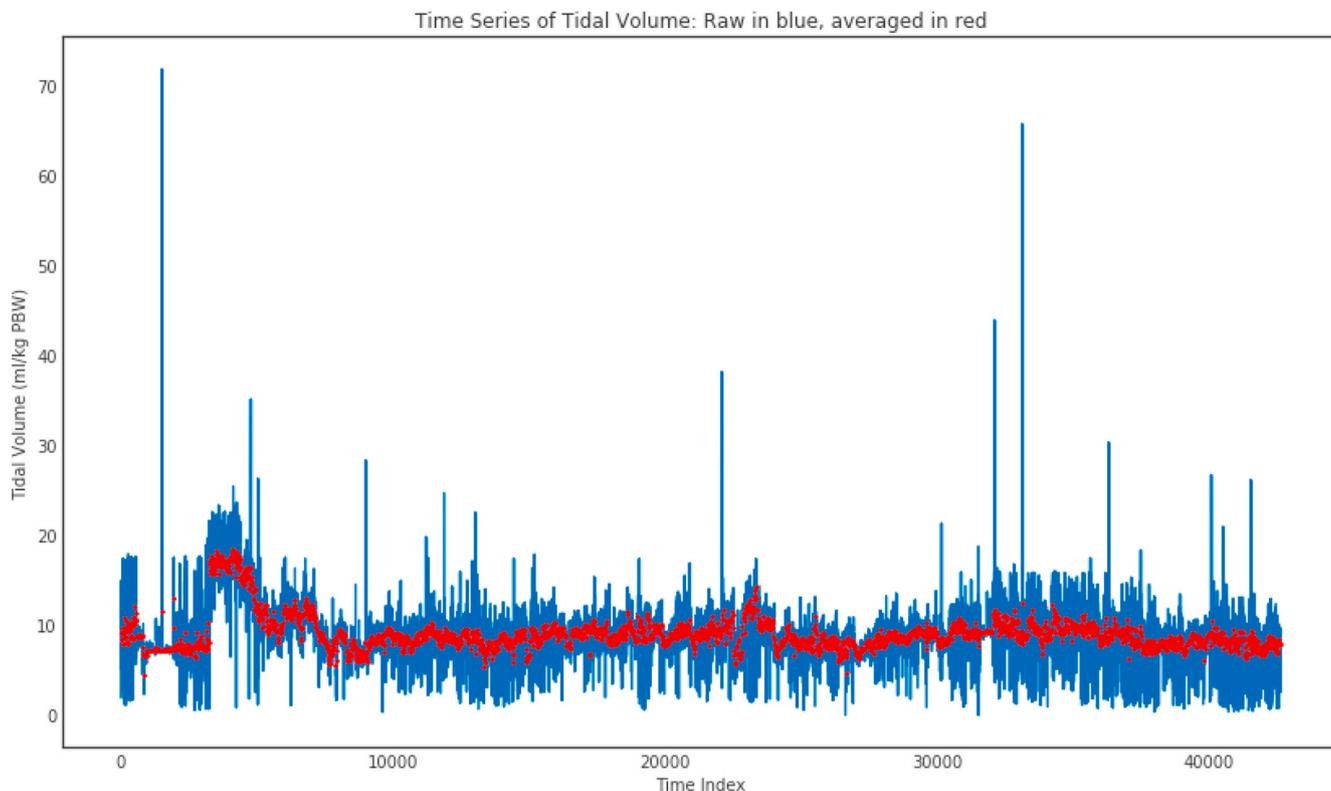
Fig. 2. Tidal volume per kg of predicted body weight for patient 12 raw data in blue. The red points represent the smoothed data from 15 min bins.

Table 1:
Comparison of regressor methods for prediction for patients tidal volume metric one time step ahead. The elapsed time for each computation is reported in the format hh:mm:ss.

Patient	No. data points	AdaBoost		RandomForest		Bagging		ExtraTrees		GradientBoosting	
		RMSE	Time	RMSE	Time	RMSE	Time	RMSE	Time	RMSE	Time
1	517	0.69	00 : 02 : 31	0.68	00 : 08 : 54	0.70	00 : 08 : 43	0.68	00 : 03 : 10	0.84	00 : 01 : 00
2	150	1.05	00 : 00 : 10	1.03	00 : 00 : 16	1.03	00 : 00 : 14	1.05	00 : 00 : 09	1.23	00 : 00 : 04
3	1358	0.38	00 : 20 : 17	0.38	02 : 12 : 15	0.38	02 : 03 : 59	0.36	00 : 22 : 36	0.39	00 : 06 : 19
4	40	1.05	00 : 00 : 00	0.99	00 : 00 : 00	1.02	00 : 00 : 00	1.00	00 : 00 : 00	0.95	00 : 00 : 00
5	162	0.34	00 : 00 : 09	0.32	00 : 00 : 16	0.32	00 : 00 : 16	0.34	00 : 00 : 10	0.41	00 : 00 : 04
6	178	1.38	00 : 00 : 38	1.28	00 : 00 : 26	1.35	00 : 00 : 22	1.39	00 : 00 : 14	1.32	00 : 00 : 06
7	1153	0.62	00 : 15 : 09	0.63	00 : 38 : 59	0.62	00 : 38 : 04	0.62	00 : 13 : 09	0.61	00 : 05 : 38
8	1245	0.83	00 : 17 : 48	0.84	00 : 42 : 29	0.84	00 : 40 : 11	0.84	00 : 14 : 50	0.91	00 : 06 : 44
9	1501	0.32	00 : 24 : 37	0.32	00 : 59 : 46	0.31	00 : 57 : 41	0.31	00 : 20 : 15	0.38	00 : 08 : 32
10	167	1.66	00 : 00 : 10	1.13	00 : 02 : 22	1.14	00 : 00 : 26	1.22	00 : 00 : 12	1.13	00 : 00 : 05
11	133	0.60	00 : 00 : 06	0.65	00 : 00 : 09	0.62	00 : 00 : 11	0.57	00 : 00 : 08	0.67	00 : 00 : 04
12	2682	0.78	01 : 11 : 55	0.75	03 : 14 : 42	0.75	03 : 15 : 09	0.73	01 : 05 : 27	0.98	00 : 26 : 27
13	2530	0.94	01 : 03 : 41	0.94	02 : 49 : 55	0.93	02 : 49 : 01	0.90	00 : 58 : 58	1.02	00 : 26 : 32
14	2107	0.18	00 : 46 : 23	0.13	01 : 26 : 02	0.14	01 : 23 : 10	0.19	00 : 37 : 10	0.50	00 : 18 : 42
15	757	0.84	00 : 06 : 32	0.85	00 : 11 : 42	0.85	00 : 11 : 42	0.86	00 : 05 : 13	0.87	00 : 02 : 30
16	1103	0.72	00 : 13 : 38	0.68	00 : 30 : 43	0.68	00 : 28 : 53	0.67	00 : 11 : 28	0.77	00 : 05 : 04
17	795	0.39	00 : 06 : 48	0.40	00 : 12 : 53	0.40	00 : 12 : 51	0.40	00 : 05 : 43	0.42	00 : 02 : 41
18	853	0.69	00 : 07 : 36	0.68	00 : 18 : 59	0.67	00 : 18 : 36	0.66	00 : 07 : 17	0.74	00 : 03 : 08
19	349	0.47	00 : 01 : 36	0.46	00 : 01 : 51	0.47	00 : 01 : 59	0.45	00 : 01 : 02	0.45	00 : 00 : 29
20	545	0.64	00 : 03 : 09	0.61	00 : 06 : 03	0.62	00 : 06 : 03	0.54	00 : 02 : 55	0.75	00 : 01 : 19
21	205	1.84	00 : 00 : 46	1.85	00 : 00 : 35	1.82	00 : 00 : 32	2.01	00 : 00 : 18	2.02	00 : 00 : 08
22	4202	0.95	03 : 17 : 11	0.94	09 : 43 : 39	0.94	09 : 47 : 20	0.93	02 : 59 : 15	0.98	01 : 20 : 19
Mean±sd		0.79 ± 0.42		0.75 ± 0.38		0.81 ± 0.41		0.76 ± 0.41		0.83 ± 0.38	
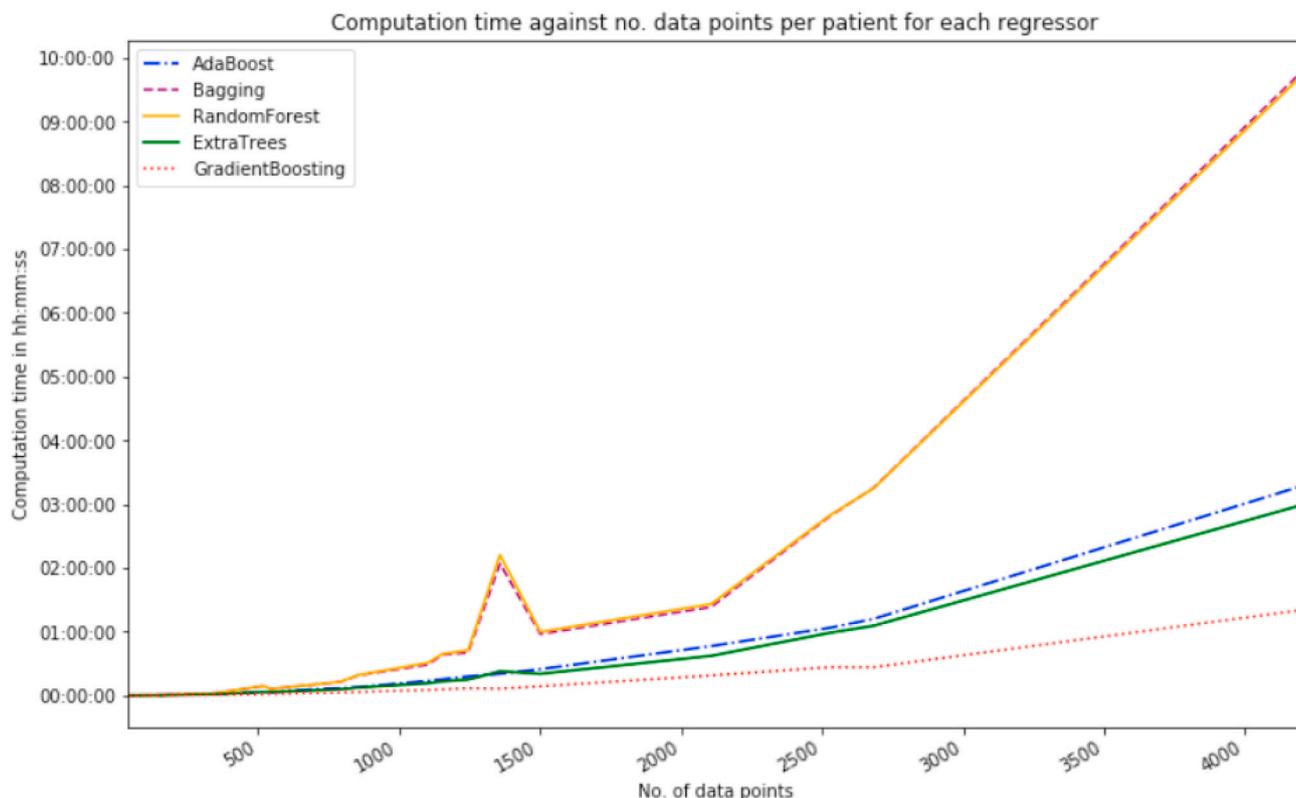
Fig. 3. Computation time taken against the number of data points per patient for each of the 5 regressors predicting 1 time step ahead.

Table 2

Comparison of the RMSE of five regressor methods for prediction of patient 1 up to four timesteps ahead.

	RMSE		
Regressor	T+2	T+3	T+4
AdaBoost	0.69	0.71	0.73
RandomForest	0.66	0.68	0.71
Bagging	0.66	0.67	0.70
ExtraTrees	0.64	0.69	0.71
GradientBoosting	0.82	0.83	0.83
Mean±sd	0.69 ± 0.07	0.72 ± 0.07	0.74 ± 0.05

2.1. Regression models

We studied five different regression tree approximations, including bagging and boosting approaches which combine simpler decision trees in various ways to obtain the most accurate predictions. In the bootstrap aggregation approach, also known as bagging, we explore the Bagging, ExtraTrees and RandomForest methods. The alternative boosting approach is covered with the AdaBoost and GradientBoosting methods. For each of the methods we utilised the tsfresh software toolkit [43], to extract the features used as input into the models.

2.2. LSTM neural networks

As an alternative to using regression methods, we investigated the use of long short-term memory neural networks (LSTM), using the Keras and Tensorflow libraries [44]. The LSTM form of the recurrent neural network architecture is quite complex relative to the original Elman form [45] of the RNN and this enables LSTMs to store information over longer periods, an ideal attribute for modelling time series. Instead of working with the feature vectors derived from the time series, in this approach we work directly with the time bins and the observed values y_i.

Gorr [46] argued that machine learning methods do not require pre-processing of the observed data to achieve stationarity of both the mean and the variance. This is the approach that we therefore follow however we note that there is a contrary view expressed by some authors [47].

The key parameters which distinguish one neural network model from another in our work are: the number of layers and the shape of the input and the output layers.

We investigated two models which used different numbers of intermediate layers. The first and last layers are dense layers identical in each model and exist to accommodate the input and the output. The intermediate layers distinguishing the models are as follows: ModelNeuralA has one LSTM layer and ModelNeuralB has three LSTM layers, with a 20% dropout layer between the second and third layers to avoid overfitting.

Each LSTM layer in our model has 50 nodes; this was chosen as 2.5 times the number of input neurons a figure that we believe to be representative of the patients recent breathing history. We trained both models using 70% of the available data and then made predictions for the remaining 30% of the data. We trained the models for 500 epochs and for both models we performed computations using the direct forecasting method [48] where we used 20 input points for forecast 4 steps ahead. Our network used the activation function *relu*, the rectified linear unit and the optimizer used was *adam*. The rectified linear activation function is a piecewise linear function that will output the input directly if is positive, otherwise, it will output zero. It is commonly used because it is easy to compute relative to other activation functions [49,50]. For ModelNeuralA the above choices created a network with $10,604$ trainable parameters and for ModelNeuralB there was a total of $51,004$ trainable parameters.

All analysis for this work was carried out using the Python programming language and for each method calculated the root mean square of the absolute error (RMSE) between the observed values at a time point and the values predicted for the time point an interval ahead.
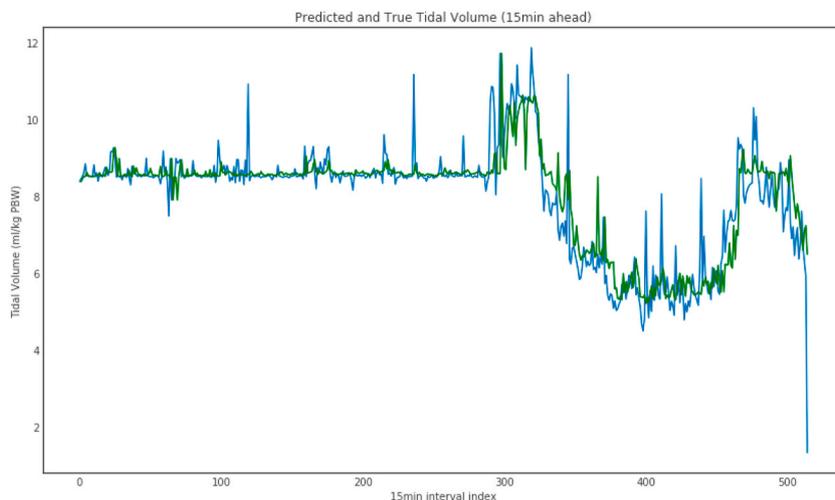
(a) Tidal volume per kg of predicted body weight for patient 1 in blue. The line in green is the prediction one time bin ahead (15 minutes) using the AdaBoost regression technique.
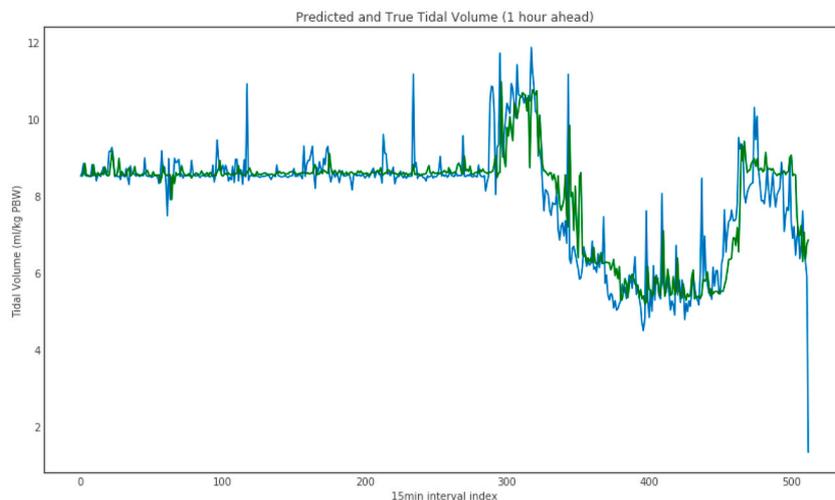
(b) Tidal volume per kg of predicted body weight for patient 1 in blue. The line in green is the prediction four time bins ahead (60 minutes) using the AdaBoost regression technique.
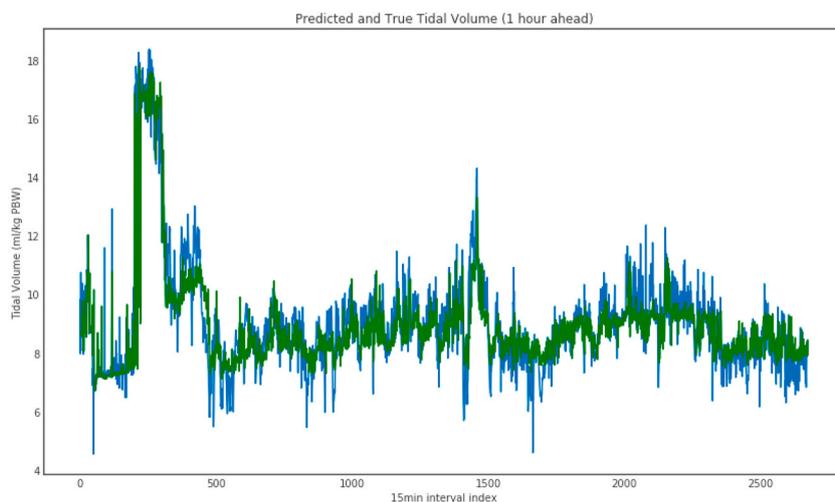
(c) Tidal volume per kg of predicted body weight for patient 12 in blue. The line in green is the prediction four time bins ahead (60 minutes) using the AdaBoost regression technique.

Fig. 4. Comparison of tidal volume per kg of predicted body weight for patient 1 predicting one and four time steps ahead shown in Fig. 4a and b. And 4c showing the prediction of patient 12 from cohort 2 four time steps ahead. Raw data in blue, predicted values in green.
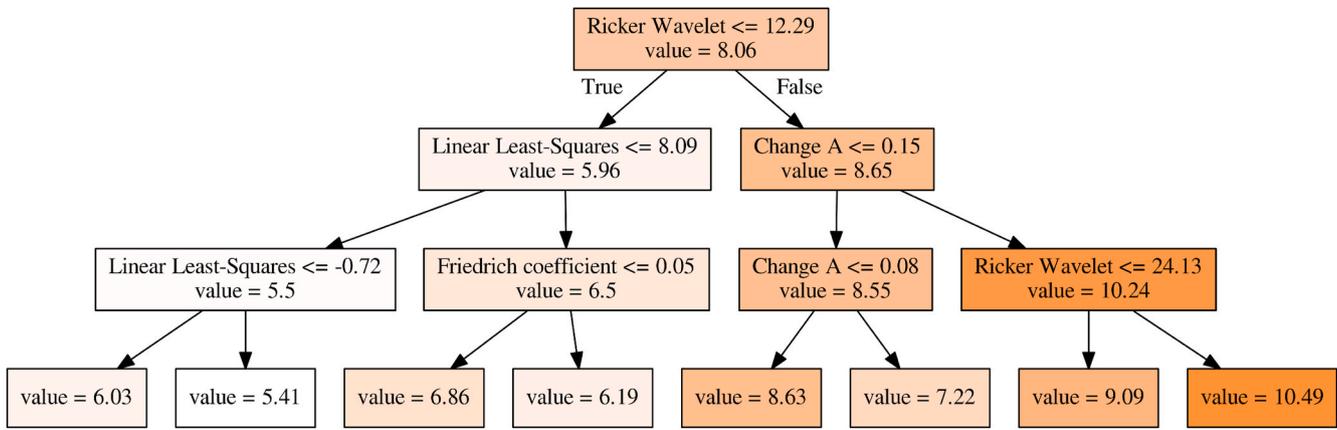
Fig. 5. One of the ten regression trees generated by the AdaBoost kernel for patient 1. Refer to Table 6 in appendix for feature explanations.

Table 3
Predicting 4 time steps ahead using AdaBoost Regression for all 22 patients.

Patient	No. Data points	RMSE	Time
1	517	0.74	00 : 01 : 08
2	150	1.29	00 : 00 : 05
3	1358	0.37	00 : 07 : 18
4	40	1.10	00 : 00 : 00
5	162	0.49	00 : 00 : 06
6	178	1.43	00 : 00 : 08
7	1153	0.65	00 : 05 : 20
8	1245	0.91	00 : 06 : 22
9	1501	0.37	00 : 08 : 59
10	167	1.11	00 : 00 : 07
11	133	0.66	00 : 00 : 05
12	2682	1.02	00 : 33 : 30
13	2530	1.04	00 : 28 : 40
14	2107	0.65	00 : 19 : 30
15	757	0.87	00 : 02 : 36
16	1103	0.82	00 : 05 : 08
17	795	0.42	00 : 02 : 43
18	853	0.79	00 : 03 : 05
19	349	0.50	00 : 00 : 30
20	545	0.72	00 : 01 : 14
21	205	2.32	00 : 00 : 11
22	4202	0.99	01 : 21 : 26
Mean±sd		0.88 ± 0.43	

Table 4
Predicting 4 time steps ahead using both LSTM models for all 22 patients.

Patient	No. Data points	ModelNeuralA		ModelNeuralB	
		RMSE	Time	RMSE	Time
1	517	4.97	00 : 00 : 39	2.74	00 : 21 : 57
2	150	2.24	00 : 00 : 12	1.90	00 : 05 : 59
3	1358	0.64	00 : 01 : 38	0.66	01 : 02 : 21
4	40	–	–	–	–
5	162	0.97	00 : 00 : 13	1.24	00 : 06 : 22
6	178	1.83	00 : 00 : 14	2.18	00 : 08 : 18
7	1153	1.00	00 : 01 : 25	1.03	00 : 55 : 41
8	1245	1.26	00 : 01 : 32	1.70	00 : 56.49
9	1501	0.94	00 : 01 : 43	0.79	01 : 06 : 49
10	167	2.14	00 : 00 : 13	2.25	00 : 06 : 30
11	133	1.10	00 : 00 : 11	0.87	00 : 06 : 23
12	2682	1.44	00 : 03 : 15	1.54	01 : 41 : 25
13	2530	2.30	00 : 03 : 06	2.57	01 : 39 : 36
14	2107	1.08	00 : 02 : 33	1.19	01 : 22 : 38
15	757	1.40	00 : 00 : 55	1.34	00 : 37 : 56
16	1103	1.62	00 : 01 : 21	1.52	00 : 51 : 35
17	795	0.74	00 : 01 : 00	1.08	00 : 39 : 35
18	853	0.99	00 : 01 : 02	0.98	00 : 38 : 59
19	349	0.50	00 : 00 : 25	0.47	00 : 17 : 16
20	545	1.69	00 : 00 : 39	1.68	00 : 27 : 44
21	205	4.60	00 : 00 : 16	5.03	00 : 10 : 52
22	4202	1.37	00 : 04 : 41	1.41	02 : 24 : 30
Mean±sd		1.66 ± 1.16		1.61 ± 0.99	

We used RMSE [51], a commonly used tool in regression analysis, to quantify the accuracy of our predictions for each patient. RMSE is sometimes criticized as being sensitive to outliers, but we see this property as valuable in clinical application work since it flags more significant differences between the predicted and the actual patient readings when they occur [52]. Fig. 1 shows the process of our work.

3. Results

From our selection of 22 patients that represent a coverage of the profiles seen we were able to identify two cohorts representing the two most frequently used modes of ventilation. From visual analysis we identified cohort 1 to mimic controlled followed by support mode of ventilation and cohort 2 to be pure support mode. While patients can show any combination of controlled and support modes and can even move from one to the other more than once as demanded by disease progression, clinical need etc., we will stick to just two patterns, controlled followed by support or just support alone. Patients 1 to 11

represent cohort 1 and patients 12 to 22 labelled as cohort 2.

3.1. Smoothing raw time series

As discussed in Section 2 we apply smoothing to our raw data in order to extract the true trends in the patients tidal volume. As described we take averaged 15 min bins as our patient data going forward. Fig. 2 shows how smoothing the data can remove the large anomaly and variation in the data that would potentially throw off our predictive models but still captures the overall trend of the patients data.

3.2. Evaluating an optimal regression forecasting method

Initially we investigated predicting one 15 min time step ahead for each of the patients. For each point in each of the time series, we extracted feature vectors using all preceding points and then used the

Table 5
Prediction of Alerts: Table 5a shows the Total number of alerts with TP and FN reported for AdaBoost and LSTM. Table 5b reports the accuracy using AdaBoost regression for all 22 patients.

(a) Prediction of Alerts. Total being the total number of alerts generated, TP giving the true positives and FN stating the false negatives.

	AdaBoost			LSTM		
	Using all data			Only last 30% of data		
Patient	**Total**	**TP**	**FN**	**Total**	**TP**	**FN**
1	84	81	3	4	3	1
2	25	23	2	10	5	5
3	0	0	0	0	0	0
4	2	1	1	–	–	–
5	3	2	1	1	0	1
6	11	3	8	4	3	1
7	0	0	0	0	0	0
8	167	142	25	86	82	4
9	64	44	20	24	14	10
10	7	3	4	5	1	4
11	13	8	5	4	2	2
12	430	382	48	177	173	4
13	627	627	0	189	184	5
14	79	35	44	48	31	17
15	3	0	3	0	0	0
16	42	25	17	10	2	8
17	0	0	0	0	0	0
18	50	18	32	1	0	1
19	0	0	0	0	0	0
20	32	27	5	25	4	21
21	48	48	0	15	10	5
22	47	17	30	0	0	0

(b) The accuracy of predicted alerts for all patients using AdaBoost Regression.

Patient	Accuracy
Patient 1	0.96
Patient 2	0.92
Patient 3	1.00
Patient 4	0.50
Patient 5	0.67
Patient 6	0.27
Patient 7	1.00
Patient 8	0.85
Patient 9	0.69
Patient 10	0.43
Patient 11	0.62
Patient 12	0.89
Patient 13	1.00
Patient 14	0.44
Patient 15	0.00
Patient 16	0.59
Patient 17	1.00
Patient 18	0.36
Patient 19	1.00
Patient 20	0.84
Patient 21	1.00
Patient 22	0.36

five different types of regression methods to predict the value of the time series one bin ahead i.e. 15 min. We repeated this process for all points in the time series, for each patient and report the RMSE value per method.

Table 1 presents RMSE calculated over all patients for each type of regression method. In addition the table shows the elapsed time taken for the computations, reflecting the number of arithmetic and logical operations needed to implement each regressor kernel. In all models we set the maximum number of trees to be generated at 10 to prevent overfitting. We compared the depth of the trees for each bagging regressor approximation. The mean depth for each of the RandomForest, ExtraTrees and Bagging trees were 32 ± 5, 37 ± 4 and 33 ± 5 respectively. The boosted regression methods, however, create trees of depth four by default. We therefore compared the effect of increasing the number of trees created for these models for patient 1, finding a decreasing trend of RMSE for increasing number of trees, as expected.

Table 1 shows that for patient 1 the GradientBoosting method yields a relatively poorer result than others with an RMSE of 0.84 whereas the other four regressors yield similar values of RMSE lying in the range $[0.68, 0.70]$. This is consistent across our experiments with 17 out of the 22 patients yielding a higher RMSE when predicting with a GradientBoosting regressor. We can further distinguish among the regressors by taking the computational performance into consideration. Across all patients, GradientBoosting yields the quickest compute time of all five regressors, where RandomForest and Bagging have significantly higher compute times. Fig. 3 highlights how the computation time increases as the number of data points increases for each of the five methods. AdaBoost and ExtraTrees regressors give the best trade-off between RMSE and computational time, having an RMSE of 0.69 and 0.68, and taking only 02:31 and 03:10 (mm:ss) for patient 1 respectively. We further report that AdaBoost has the smaller RMSE range $[0.18, 1.84]$ compared to that off ExtraTrees $[0.19, 2.01]$.

We next investigated predicting up to four time steps ahead. One might at first expect that the further into the future one predicts then the larger the RMSE would be. However the change in RMSE is in the second decimal place, for all five methods in Table 2 as we move from predicting two to predicting four time steps ahead, for patient 1. The RandomForest method had an RMSE of 0.66 predicting two steps ahead, and 0.71 for four steps ahead. Similarly for AdaBoost, the change is from 0.69 to 0.73. This difference is further highlighted in Fig. 4a and b, presenting the raw data for patient 1, in blue and the predictions one time step and four time steps ahead in green. Further, Fig. 4c presents patient 12, from cohort 2, predicting 4 time steps ahead. The computational time was independent of the timestep, matching the computational time values reported in Table 1.

Fig. 5 shows one of the ten regressor trees created by the AdaBoost method for the prediction 4 time steps ahead for patient 1. The tree splits the data at each node based on the condition given, derived from the features shown and arrives at a prediction by asking a series of questions to the data. The features utilised in Fig. 5 have been extracted using tsfresh as the most significant features for the prediction of tidal volume for the given patient. It is interesting here to note what some of the features can mean in our problem domain. The Ricker Wavelet is used to process seismic data propagated through viscoelastic homogeneous media. Further the Friedrich coefficient, derived from the Langevin model, aims at describing the random movement of a particle in a fluid, taking into account the viscosity and temperature. This would indicate that possibly the amount of fluid in the lungs and thus a viscous medium would have a high impact in how the patients tidal volume changes over time, something that we plan to explore going forward.

Based on the results in Tables 1 and 2, and further on the basis of easier interpretation of the decision trees made, we selected the AdaBoost regressor to examine the effectiveness of predicting 1 h ahead for all patients. Table 3 highlights these results, showing an RMSE range $[0.37, 2.32]$.

3.3. Prediction using neural networks

We then proceeded to analyse our two LSTM models as described in Section 2. It is important to note that because our LSTM models use 70% training data and uses 20 input points to predict 4, predictions can not be made on patient 4 as the patient only had 40 data points, a drawback of this method.

Table 4 shows the results for predicting four time steps ahead for all patients using LSTM ModelNeuralA and ModelNeuralB. We notice that the computation times are a lot quicker compared to the regressor models for ModelNeuralA. While patient 22 takes up to 1 h, 21 min and 26 s to predict 4 time steps ahead using AdaBoost, here the compute time is only 4 min and 41 s. However, we do notice a significantly higher RMSE range of $[0.5, 4.97]$ indicating that our LSTM one layer model does not perform as well as AdaBoost for the prediction of tidal volume values 1 h ahead. Further, our LSTM three layer model has both higher
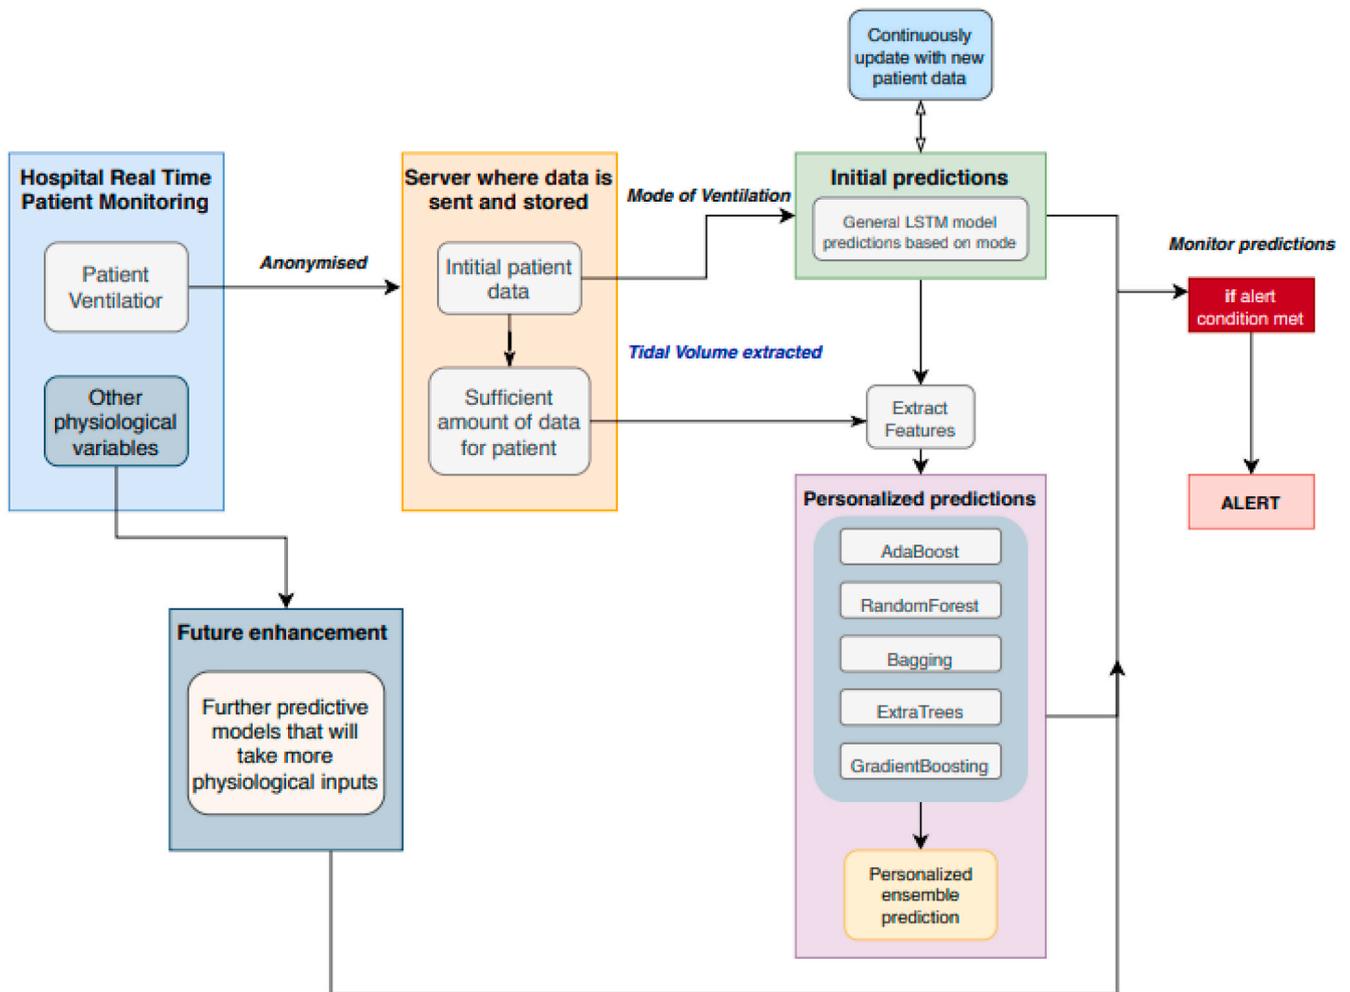
Fig. 6. Bigger picture of how the software could expand to real world application.

computation times and higher RMSE values and we therefore decide to proceed our work with LSTM ModelNeuralA.

3.4. Prediction of alerts

The VILIAlert system created an SMS alert to the clinicians whenever the LPV threshold was violated for four time bins consecutively. All of the generated alerts were stored in the database, and used to compare our predictions with. We take the time point for each of the generated alerts and cross reference with our predictions; if the four previous time point predictions are greater than the LPV threshold defined, then our system would have predicted the alert 1 h ahead. Therefore we compared the recorded SMS alerts with the predictions shown in Fig. 5.

Due to the drawback of our LSTM model using 70% training data, we can only test our predictions on the last 30% of data remaining for all patients, highlighted in the differences in results shown in Table 5a. With true positives (TP) being alerts that were generated from the VILIAlert system that our model would have predicted 1 hour ahead and false negatives (FN) being alerts that would not have predicted, we can evaluate the accuracy of our predictive models. We report the accuracy using equation (2) in Table 5b.

$$Accuracy = \frac{TP}{TP + FN} \qquad (2)$$

We can see from Table 5a that our AdaBoost model performs accurately for the prediction of alerts generated by the VILIAlert system. For the 84 alerts generated for patient 1, 81 of these would have been

predicted an hour ahead of time and therefore could have been prevented and ensured safer ventilation of the patient. Further, our results show that the different modes of ventilation and thus cohorts of patients does not have an effect on the predictive accuracy. In turn, this would indicate that any patient with any tidal volume profile can be predicted within an accuracy of 10% using our AdaBoost model.

4. Discussion

Ventilation is a valuable tool for treatment of patients in the ICU but has to be managed so that it does not in itself lead to lung injury. Early recognition of the potential for such damage is vital to assist the clinician. In this paper we have studied the viability of methods for the prediction of tidal volume, methods based on machine learning techniques to provide early warning of over ventilation. We further utilised smoothing techniques and have demonstrated a smart alert system that has a predictive accuracy within 10% of true values.

It is important to ensure the quality of alerts in clinical decision support tools in order to reduce alarm fatigue. As discussed alarm fatigue can cause increased response time and alerts can even be overridden. Dependent on the tidal volumes, the results for patients with values oscillating around the 8 ml/kg IBW threshold, the accuracy is low, e.g. in patient 6, the accuracy is only 0.27, thus while we are predicting these values with an RMSE 1.43 they may not always be flagged as alerts. Going forward we would deem it appropriate to have a threshold range ±0.5 ml/kg IBW in order to improve true alert detection. Further, we would propose using a traffic light alert system in real

time; green suggesting no breaches predicted, amber indicating that within the next 1 h period the 4 predictions are within 8 ml/kg ± 0.5 and further a red alert if all 4 predicted values for the next hour are above the 8 ml/kg threshold.

Our data was collected as part of an observational study and as historical data, has allowed us to investigate a significant number of data points, with over 4 million per minute tidal volume readings recorded. We compare two different machine learning methods for the prediction of this metric 1 h ahead, ensuring enough time for clinical intervention to prevent a threshold breach. We have found that decision trees are an adequate solution in as much as they deliver relevant predictions of threshold breaches within a few hours of starting ventilation and require minimal computational resources. Furthermore, we identified that the magnitude of the Rickler wavelet is a critical determinant in the analysis of the tidal volume waveform. This wavelet analysis arises in the study of seismic wave propagation through viscoelastic homogeneous media, under the approximation that Newtonian viscosity is valid. The viscoelastic characteristic of lung parenchyma and additional fluid component such as haematocrit of blood in the pulmonary circulation have been studied [53], however the changes associated with the development of extraalveolar oedema are yet to be studied. Development of an automated detection tool based on changes in viscoelastic properties could enable rapid detection of development of cardiogenic and non-cardiogenic oedema in the lungs. Earlier and automated detection will guide fluid balance strategies that is associated with clinical outcomes [54], as well as earlier institution of investigations lung ultrasound to assess extravascular lung water [55]. We aim to investigate this further in future studies where we select patients with specific lung pathophysiologies. More generally, we are building on the foundational work in this paper in a new project which seeks to optimize mechanical ventilation to deliver lung protective ventilation, predict the development of ventilatory associated conditions and guide weaning. While our work does present limitations as being a single centre, retrospective study, we are confident that our techniques and models work independent of the patient profile or mode of ventilation utilised. Therefore, our models are generalisable despite not being externally validated which we plan on doing in future work.

Human physiology is a complex dynamic system and thus we will incorporate additional physiological data streams to predict deterioration more accurately and gain a more in depth understanding of patient states. Building such systems involves consideration of many variability points and within that several configuration settings. In our work each of the tree based models are ensembles in their own right. Other systems have also been studied recently proving the benefit of machine learning in healthcare [56]. In the financial domain, Krauss and co-workers [57] found that applying a higher level of ensemble proved to be a powerful model and we intend to investigate similar ensembles applied to physiology in future work, as highlighted in Fig. 6.

Author contributions

MS and CJG designed the VILIAlert system and CJG implemented the software to gather the data. RH created and implemented all of the analytics software which generated the results reported in this paper. MS and DMcA provided clinical input on the algorithms and analysis of the results while CJG and IS provided input on the software implementation and computational validation of the results. RH led the writing of the manuscript with feedback from all co-authors.

5. Data availability

The relevant data and code will be made available on request and will be released for replication of result purpose.

Declaration of competing interest

The authors declare no competing interests.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.compbiomed.2020.104030.

References

[1] H. Wunsch, W.T. Linde-Zwirble, D.C. Angus, M.E. Hartman, E.B. Milbrandt, J. M. Kahn, The epidemiology of mechanical ventilation use in the United States, Crit. Care Med. 38 (10) (2010) 1947–1953, https://doi.org/10.1097/CCM.0b013e3181ef4460.

[2] Icnarc, Number of mechanically ventilated patients during 2012 (January). Available from: https://www.icnarc.org/DataServices/.../194bd654-e14a-e411-a65b-d48564544b14, 2014.

[3] O. Gajic, J. Lee, C.H. Doerr, J.C. Berrios, J.L. Myers, R.D. Hubmayr, Ventilator-induced cell wounding and repair in the intact lung, Am. J. Respir. Crit. Care Med. 167 (8) (2003) 1057–1063, https://doi.org/10.1164/rccm.200208-889OC.

[4] J.C. Parker, L.A. Hernandez, K.J. Peevy, Mechanisms of ventilator-induced lung injury, Crit. Care Med. 21 (1) (1993) 131–143.

[5] D. Dreyfuss, G.U.Y. Basset, P. Soler, G. Saumon, Intermittent positive-pressure hyperventilation with high inflation pressures produces pulmonary microvascular injury in rats, Am. Rev. Respir. Dis. 132 (4) (1985) 880–884.

[6] R.G. Brower, M.A. Matthay, A. Morris, D. Schoenfeld, B.T. Thompson, A. Wheeler, Ventilation with lower tidal volumes as compared with traditional tidal volumes for acute lung injury and the acute respiratory distress syndrome, N. Engl. J. Med. 342 (18) (2000) 1301–1308, https://doi.org/10.1056/NEJM200005043421801.

[7] Serpa Neto, S.O. Cardoso, J.A. Manetta, V.G. Pereira, D.C. Espósito, O. Pasqualucci Mde, M.C. Damasceno, M.J. Schultz, Association between use of lung-protective ventilation with lower tidal volumes and clinical outcomes among patients without acute respiratory distress syndrome: a meta-analysis, J. Am. Med. Assoc. 308 (16) (2012 Oct 24) 1651–1659, https://doi.org/10.1001/jama.2012.13730.

[8] A.S. Neto, F.D. Simonis, C.S.V. Barbas, M. Biehl, R.M. Determann, J. Elmer, M. J. Schultz, Lung-protective ventilation with low tidal volumes and the occurrence of pulmonary complications in patients without acute respiratory distress syndrome: a systematic review and individual patient data analysis, Crit. Care Med. 43 (10) (2015) 2155–2163, https://doi.org/10.1097/CCM.0000000000001189.

[9] J.J. Cordingley, B.F. Keogh, The pulmonary physician in critical care 8: ventilatory management of ALI/ARDS, Thorax 57 (8) (2002) 729–734, https://doi.org/10.1136/thorax.57.8.729.

[10] Nhs, Clinical guide for the management of cancer patients during the coronavirus pandemic, Retrieved from, https://www.england.nhs.uk/coronavirus/wp-content/uploads/sites/52/2020/03/specialty-guide-itu-and-coronavirus-v1-16-march-2020.pdf, 2020.

[11] N. Chen, M. Zhou, X. Dong, J. Qu, F. Gong, Y. Han, L. Zhang, Epidemiological and clinical characteristics of 99 cases of 2019 novel coronavirus pneumonia in Wuhan, China: a descriptive study, Lancet 395 (10223) (2020) 507–513, https://doi.org/10.1016/S0140-6736(20)30211-7.

[12] G. Bellani, J.G. Laffey, T. Pham, E. Fan, L. Brochard, A. Esteban, A. Pesenti, Epidemiology, patterns of care, and mortality for patients with acute respiratory distress syndrome in intensive care units in 50 countries, JAMA - Journal of the American Medical Association 315 (8) (2016) 788–800, https://doi.org/10.1001/jama.2016.0291.

[13] M. Schwede, R.Y. Lee, H. Zhuo, et al., Clinician recognition of the acute respiratory distress syndrome: risk factors for under-recognition and trends over time, Crit. Care Med. 48 (6) (2020) 830–837, https://doi.org/10.1097/CCM.0000000000004328.

[14] S. Eslami, N.F. de Keizer, A. Abu-Hanna, E. de Jonge, M.J. Schultz, Effect of a clinical decision support system on adherence to a lower tidal volume mechanical ventilation strategy, J. Crit. Care 24 (4) (2009) 523–529, https://doi.org/10.1016/j.jcrc.2008.11.006.

[15] S. Bagga, D.E. Paluzzi, C.Y. Chen, J.M. Riggio, M. Nagaraja, P.E. Marik, M. Baram, Better ventilator settings using a computerized clinical tool, Respir. Care 59 (8) (2014) 1172–1177, https://doi.org/10.4187/respcare.02223.

[16] C.P. Bourdeaux, M.J.C. Thomas, T.H. Gould, G. Malhotra, A. Jarvstad, T. Jones, I. D. Gilchrist, Increasing compliance with low tidal volume ventilation in the ICU with two nudge-based interventions: evaluation through intervention time-series analyses, BMJ Open 6 (5) (2016) 6–11, https://doi.org/10.1136/bmjopen-2015-010129.

[17] R. Kalhan, M. Mikkelsen, P. Dedhiya, J. Christie, C. Gaughan, P.N. Lanken, B. D. Fuchs, Underuse of lung protective ventilation: analysis of potential factors to

explain physician behavior, Crit. Care Med. 34 (2) (2006) 300–306, https://doi.org/10.1097/01.CCM.0000198328.83571.4A.

[18] G Ristagno and M H Weil, The History of Critical Care Medicine: the Past Present and Future, Intensive and Critical Care Medicine, A. Gullo, et al. (eds.), pub: Springer-Verlag pp.3–17.

[19] N. Webster, Monitoring the critically ill patient, J. R. Coll. Surg. Edinb. (1999) 386–393, available on the web at, http://www.rcsed.ac.uk/RCSEDBackIssues/journal/vol44_6/4460032.htm.

[20] A. Thommandram, J.E. Pugh, J.M. Eklund, C. McGregor, A.G. James, Classifying neonatal spells using real-time temporal analysis of physiological data streams: algorithm development, in: IEEE EMBS Spec. Top. Conf. Point-of-Care Healthc. Technol. Synerg. Towar. Better Glob. Heal. PHT, 2013, pp. 240–243, 2013.

[21] C.M.C. Gregor, N. Percival, K.P. Smith, D. Sow, Real-time analysis for intensive care ©, Eng. Med. Biol. Mag. IEEE 29 (2) (2010) 110–118.

[22] C.A. Lovejoy, V. Buch, M. Maruthappu, Artificial intelligence in the intensive care unit, Crit. Care 23 (1) (2019) 23–25, https://doi.org/10.1186/s13054-018-2301-9.

[23] N. Prasad, L.F. Cheng, C. Chivers, M. Draugelis, B.E. Engelhardt, A reinforcement learning approach to weaning of mechanical ventilation in intensive care units, in: Uncertainty in Artificial Intelligence - Proceedings of the 33rd Conference, 2017. UAI 2017.

[24] S. Nemati, A. Holder, F. Razmi, M.D. Stanley, G.D. Clifford, T.G. Buchman, An interpretable machine learning model for accurate prediction of sepsis in the ICU, Crit. Care Med. 46 (4) (2018) 547–553, https://doi.org/10.1097/CCM.0000000000002936.

[25] M. Komorowski, L.A. Celi, O. Badawi, A.C. Gordon, A.A. Faisal, The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care, Nat. Med. 24 (11) (2018) 1716–1720.

[26] N. Tomašev, X. Glorot, J.W. Rae, M. Zielinski, H. Askham, A. Mottram, S. Mohamed, A clinically applicable approach to continuous prediction of future acute kidney injury, Nature 572 (7767) (2019) 116–119, https://doi.org/10.1038/s41586-019-1390-1.A.

[27] A. Jalali, D. Bender, M. Rehman, V. Nadkarni, C. Nataraj, Advanced analytics for outcome prediction in intensive care units, in: Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS, 2016, pp. 2520–2524, https://doi.org/10.1109/EMBC.2016.7591243.

[28] Z. Che, S. Purushotham, R. Khemani, Y. Liu, Interpretable deep models for ICU outcome prediction, in: AMIA … Annual Symposium Proceedings. AMIA Symposium, 2016, pp. 371–380, 2016.

[29] G. Boverman, S. Genc, Prediction of mortality from respiratory distress among long-term mechanically ventilated patients, in: 2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC, 2014, pp. 3464–3467, https://doi.org/10.1109/EMBC.2014.6944368, 2014.

[30] J.S. Ancker, A. Edwards, S. Nosal, D. Hauser, E. Mauer, R. Kaushal, Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system, BMC Med. Inf. Decis. Making 17 (1) (2017) 1–9, https://doi.org/10.1186/s12911-017-0430-8.

[31] F. Schmid, M.S. Goepfert, D. Kuhnt, V. Eichhorn, S. Diedrichs, H. Reichenspurner, D.A. Reuter, The wolf is crying in the operating room: patient monitor and anesthesia workstation alarming patterns during cardiac surgery, Anesth. Analg. 112 (1) (2011) 78–83, https://doi.org/10.1213/ANE.0b013e3181fcc504.

[32] A.K. Sowan, T.M. Gomez, A.F. Tarriela, C.C. Reed, B.M. Paper, Changes in default alarm settings and standard in-service are insufficient to improve alarm fatigue in an intensive care unit: a pilot project, JMIR Human Factors 3 (1) (2016) e1, https://doi.org/10.2196/humanfactors.5098.

[33] F. Schmid, M.S. Goepfert, D.A. Reuter, Patient monitoring alarms in the ICU and in the operating room, Crit. Care 17 (2) (2013), https://doi.org/10.1186/cc12525.

[34] V. Herasevich, M. Tsapenko, M. Kojicic, A. Ahmed, R. Kashyap, C. Venkata, K. Shahjehan, S.J. Thakur, B.W. Pickering, J. Zhang, R.D. Hubmayr, Limiting ventilator-induced lung injury through individual electronic medical record surveillance, Crit. Care Med. 39 (1) (2011) 34–39.

[35] C.G. Rusin, S.I. Acosta, L.S. Shekerdemian, E.L. Vu, A.C. Bavare, R.B. Myers, D. J. Penny, Prediction of imminent, severe deterioration of children with parallel circulations using real-time processing of physiologic data, J. Thorac. Cardiovasc. Surg. 152 (1) (2016) 171–177, https://doi.org/10.1016/j.jtcvs.2016.03.083.

[36] J. Parreco, A. Hidalgo, J.J. Parks, R. Kozol, R. Rattan, Using artificial intelligence to predict prolonged mechanical ventilation and tracheostomy placement, J. Surg. Res. 228 (2018) 179–187, https://doi.org/10.1016/j.jss.2018.03.028.

[37] C.M. Ennett, K.P. Lee, L.J. Eshelman, B. Gross, L. Nielsen, J.J. Frassica, M. Saeed, Predicting respiratory instability in the ICU, in: Proceedings of the 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS'08 - "Personalized Healthcare through Technology, 2008, pp. 2848–2851, https://doi.org/10.1109/iembs.2008.4649796.

[38] C.J. Gillan, A. Novakovic, A.H. Marshall, M. Shyamsundar, D.S. Nikolopoulos, Expediting assessments of database performance for streams of respiratory parameters, Comput. Biol. Med. 100 (2018) 186–195, https://doi.org/10.1016/j.compbiomed.2018.05.028.

[40] Michael J. Kane, Natalie Price, Matthew Scotch, Peter Rabinowitz, Comparison of ARIMA and Random Forest time series models for prediction of avian influenza H5N1 outbreaks, BMC Bioinf. 15 (1) (2014) 276.

[41] V. Herasevich, M. Tsapenko, M. Kojicic, et al., Limiting ventilator-induced lung injury through individual electronic medical record surveillance, Crit. Care Med. 39 (1) (2011) 34–39, https://doi.org/10.1097/CCM.0b013e3181fa4184.

[42] J.H. Friedman, Greedy function approximation, the IMS 1999 ritz lecture available on the web at:. http://statweb.stanford.edu/~jhf/ftp/trebst.pdf.

[43] M. Christ, N. Braun, J. Neuffer, A.W. Kempa-Liehr, "Time series FeatuRe extraction on basis of scalable hypothesis tests (tsfresh – a Python package), Neurocomputing 307 (2018) 72–77.

[44] I. Vasilev, Python Deep Learning : Exploring Deep Learning Techniques and Neural Network Architectures with PyTorch, Keras, and TensorFlow, second ed., pub: Packt Publishing (Birmingham), 2019. -13: 978-178-9348460.

[45] J.L. Elman, Finding structure in time, Cognit. Sci. 14 (2) (1988) 179–211, https://doi.org/10.1207/s15516709cog1402_1.

[46] W.L. Gorr, Research prospective on neural network forecasting, Int. J. Forecast. 10 (1) (1994) 1–4, https://doi.org/10.1016/0169-2070(94)90044-2.

[47] G.P. Zhang, M. Qi, Neural network forecasting for seasonal and trend time series, Eur. J. Oper. Res. 160 (2) (2005) 501–514, https://doi.org/10.1016/j.ejor.2003.08.037.

[48] Spyros Makridakis, Evangelos Spiliotis, Vassilios Assimakopoulos Statistical and Machine Learning forecasting methods: Concerns and ways forward, https://doi.org/10.1371/journal.pone.0194889. Published. (Accessed 27 March 2018).

[49] Yann LeCun, Yoshua Bengio, Geoffrey Hinton, Deep learning, Nature 521 (7553) (2015) 436–444, https://doi.org/10.1038/nature14539. Bibcode:2015 Natur.521..436L, PMID 26017442.

[50] R. Arora, A. Basu, P. Mianjy, A. Mukherjee, Understanding deep neural networks wiht rectified linear units, in: Proceedings of the 6th International Conference on Learning Representations (ICLR), Vancouver Canada, 2018 available on the web at: https://arxiv.org/abs/1611.01491.

[51] Rob J. Hyndman, Anne B. Koehler, Another look at measures of forecast accuracy, Int. J. Forecast. 22 (4) (2006) 679–688, https://doi.org/10.1016/j.ijforecast.2006.03.001. CiteSeerX 10.1.1.154.9771.

[52] Robert Pontius, Olufunmilayo Thontteh, Hao Chen, Components of information for multiple resolution comparison between maps that share a real variable, Environ. Ecol. Stat. 15 (2) (2008) 111–142, https://doi.org/10.1007/s10651-007-0043-y.

[53] F. Peták, G.H. Fodor, B. Babik, W. Habre, Airway mechanics and lung tissue viscoelasticity: effects of altered blood hematocrit in the pulmonary circulation, J. Appl. Physiol. 121 (1) (2016) 261–267.

[54] H.P. Wiedemann, A.P. Wheeler, G.R. Bernard, Comparison of two fluid-management strategies in acute lung injury, J. Vasc. Surg. 44 (4) (2006) 909.

[55] M. Shyamsundar, B. Attwood, L. Keating, A.P. Walden, Clinical review: the role of ultrasound in estimating extra-vascular lung water, Crit. Care 17 (5) (2013) 237.

[56] S.L. Hyland, M. Faltys, M. Hüser, X. Lyu, T. Gumbsch, C. Esteban, T.M. Merz, Machine learning for early prediction of circulatory failure in the intensive care unit, Nat. Med. 26 (March) (2019), https://doi.org/10.1038/s41591-020-0789-4.

[57] Christopher Krauss, Xuan Anh Do, Nicolas Huck, Deep neural networks, gradient-boosted trees, random forests: statistical arbitrage on the S and P 500, FAU discussion papers in economics, No. 03/2016, friedrich-alexander-universität ErlangenNürnberg, institute for economics, erlangen available on the web at. http://hdl.handle.net/10419/130166, 2016.