Critical care

White paper

Introduction to the Philips Critical Care Outcome Prediction Models – Mechanical Ventilation Predictions

The purpose of this white paper is to:

- Provide a brief overview of the methods used for developing the Philips CCOPM-Ventilation models.
- Provide a comprehensive review of the performance of these models for both probability and duration. For ventilation duration, performance statistics are provided in comparison to APACHE IVa and IVb.
- Provide a basis for obtaining feedback from eICU partners for model design/evaluation.

Executive summary

Background: Mechanical ventilation is a critical part of intensive care unit (ICU) patient care. Philips has been providing benchmark reports for customers using APACHE IVa ventilation prediction models. These models allow customers to obtain the total ventilation duration for ICU patients and use that value for benchmark reporting across hospitals and health systems.
In this work we aim to enhance ventilation prediction by providing separate predictions for invasive and non-invasive ventilation.
We also aim to improve the overall model prediction accuracy compared to APACHE models.

In this report, we propose the Philips Critical Care Outcome Prediction Model for Ventilation (Philips CCOPM-Ventilation), which includes:

1. Ventilation-Probability model: This model predicts probabilities of a patient's need to receive: a) Any ventilation b) Any invasive ventilation c) Any non-invasive ventilation
2. Ventilation-Duration models: Two models are developed to predict: a) The duration (in days) of invasive ventilation and b) The duration (in days) of non-invasive ventilation

The objectives of this work were to:

- Develop a set of new ventilation models to provide greater insight into both invasive and non-invasive ventilation needs.
- Improve the model prediction accuracy compared to APACHE IVa and IVb in predicting the total ventilation duration.
- Train models based on more recent data (2010-2019).
- Reduce the documentation burden of obtaining mechanical ventilation predictions.

Methods:

The Philips CCOPM-Ventilation models were trained using eICU data from 2010-2019. Two manuscripts detailing the development of these models are currently under development. This white paper describes the performance of these models on the entire eCareManager database across different years (2010-2022), ICU types, diagnostics groups, and admission sources. We also compare our model to APACHE IVa and APACHE IVb models for the prediction of the total ventilation duration.

Results:

Ventilation-Probability Model: The Philips CCOPM-Ventilation Probability model includes predicting the probability of any ventilation, invasive ventilation and non-invasive ventilation. The model performance was evaluated based on independent measurement of the Area Under the Curve (AUC):

- Model predicting “any ventilation” showed AUC = 0.916.
- Model predicting “any invasive ventilation” showed AUC = 0.934.
- Model predicting “any non-invasive ventilation” showed AUC = 0.827.
- Model performance was consistent across different years (2010-2019), different ICU types, diagnostic groups, and admission sources.

Ventilation-Duration Models: These models include prediction of the duration of invasive ventilation, non-invasive ventilation, and the total ventilation. The total ventilation duration is estimated by adding the duration predicted by the invasive ventilation model and the non-invasive ventilation model. The Philips CCOPM-Ventilation Duration models were evaluated according to mean absolute error (MAE) between the predictions and the true values (in days) with the following results:

- Ventilation model predicting the duration of invasive ventilation has MAE = 1.968 days.
- Ventilation model predicting the duration of non-invasive ventilation has MAE = 0.521 days.
- Ventilation model predicting the total duration of ventilation has MAE = 2.093 days.
- Invasive and non-invasive duration models have consistent mean absolute error across different years, ICU-types, diagnostic groups, and admission sources.

We showed that the total ventilation duration predicted by our model (MAE 2.09 days) consistently outperformed APACHE-IVa and APACHE-IVb models with mean absolute errors (MAE) of 2.62 days and 2.54 days respectively.

Introduction

Ventilation is a critical part of intensive care unit (ICU) patient care. Mechanical ventilation is a lifesaving intervention for patients in respiratory distress. However, its use may create risks for patients, including infections and ventilator-induced lung injury (VILI) which, in some cases, can even lead to patient mortality[1]. Extensive research has been performed to minimize the detrimental effects of mechanical ventilation on patients while maintaining its efficacy[2]. On the other hand, delays in intubation also carry significant risks. Therefore, using ventilation predictive models could potentially offer a novel way of supporting clinical teams to enhance patient ventilation management[3].

Philips has been providing the prediction of ventilation duration using APACHE models for quarterly reporting since 2009. The model predictions are all based on first-day patient characteristics and represent the expected outcomes according to the APACHE national equations methodology[4].

Changes in medical practice and overall population health warrant review and recalibration of predictive models over time, as models may become less accurate. For example, the APACHE-IVa model was developed in a cohort of patients from 2006-2008, and the APACHE-IVb model was later trained on a more recent cohort of patients in 2014-2015.

In this white paper, we introduce new ventilation models that offer insight into ventilation management. They were trained using Philips eICU program customer data including recent years (2010-2019) to predict:

1. A set of probabilities for the need of any ventilation, any invasive ventilation, and any non-invasive ventilation for ICU patients.
2. The duration of invasive ventilation and non-invasive ventilation for ICU patients. We studied the performance of these models as well as the 'total ventilation duration' model performance relative to APACHE IVa/IVb.

Model development

The following steps were followed to develop the Philips CCOPM-Ventilation prediction models for benchmarking:

- Baseline patient features were defined and extracted to accurately reflect the actual patient status at the time of ICU admission.
- Model features were selected by considering clinical domain knowledge, complexity of data collection and employing a data-driven approach. High priority was given to objective measurements not requiring manual data entry. Features such as urinary output, active treatments, and chronic conditions that are relatively cumbersome to obtain were not required.
- Machine-learning modeling techniques were then used to harness the predictive values of the patient baseline characteristics in the eICU patient cohort to predict the need and duration of ventilation.
- The models were designed to be updated frequently, so they reflect the current performance of ICUs in the eICU installed base.

The patient cohort for model development

For training the new models, we used the eICU Research Institute (eRI) database that houses all the historical data collected from participating customers. We included all patient unit stays discharged from the hospital between 2010 and 2019 with approximately 3.8 M patients from 350 hospitals. We defined independent unit stays to be more reflective of the real patient unit stays by applying the following rules:

- Excluded patient unit stays not classified as ICU stays
- Excluded patient unit stays with irrational admission/discharge time stamps
- Combined back-to-back or overlapping stays
- Considered cases where patients are briefly moved out of ICU for surgical intervention and immediately readmitted to the ICU; these two adjacent ICU stays were handled as one continuous ICU stay
- Updated conflicting admission/discharge information using all available data

Based on the reconciled ICU patient unit stays, we excluded stays with a length of stay <4 hours or >365 days.

The baseline time window

We established the first 24 hours of ICU admission as the time window to represent a patient's baseline risk. For each feature, we used the data up to six hours prior to admission when no data was available in the first 24 hours of ICU admission.

Model features

The list of features included demographics, vital signs, critical laboratory measurements, and essential user-documented data of admission diagnosis and Glasgow Coma Score (GCS). We extracted information from all sources from the eRI, including structured and unstructured data, reconciled the conflicts and errors, and generated summarized statistics of individual inputs as model features. The same set of features were used for both probability and duration ventilation models. Vital signs and laboratory measurements were summarized (using mean and variance) over the first 24h of ICU stay if available, and over the 6h prior to ICU admission if not.

Model outcomes

The Philips ventilation probability model predicts a set of probabilities for the ventilation status for each patient ICU stay. These probabilities include the probability of the use of 1) any mechanical ventilation, 2) any invasive ventilation and 3) any non-invasive ventilation.

Similarly, for each patient-stay under ventilation, the model predicts the duration (in days) of 1) invasive ventilation and 2) non-invasive ventilation.
The model also predicts the total ventilation duration by simply summing up the predicted values of the duration of invasive and non-invasive ventilation.

Modeling techniques

For the probability model, we trained a multiclass gradient boosting model[5] to predict whether a patient fell into one of four categories during their ICU stay: none (receiving neither invasive nor non-invasive ventilation); invasive-only (receiving only invasive ventilation but not non-invasive ventilation); non-invasive-only (receiving only non-invasive ventilation but not invasive ventilation); and both (receiving both invasive and non-invasive ventilation). Gradient-boosting models consist of an ensemble of weak prediction models that provide flexibility in modeling to capture non-linear relationships and interactions between individual input features and between input features and the ventilation outcomes.

For the duration model, we developed two gradient-boosting models independently trained for predicting invasive and non-invasive ventilation duration. Features used to train these models are the same features as for the ventilation probability model. The model outcomes are the ventilation duration in hours which are later converted to days.

Model performance

In this report, the overall performance of both ventilation probability and duration models were investigated on all Health Systems in eCareManager database for patient stays during 2010-2022 which included more than 500 hospitals and 6M patient stays. These models were developed based on the data available in eRI (de-identified data from a subset of customer). Two manuscripts are in preparation describing the development process and performance of these ventilation models on the eRI data.

For the probability model, we assess the performances using the area under the receiver operating characteristic curve (AUC), with each binary discriminator having a separate AUC value. In addition to showing overall AUCs across all data, we further analyze AUCs across different years, different diagnostic groups, different ICU types and different admission sources.

In addition, we assess the performance of the 1) invasive ventilation duration model, and 2) non-invasive ventilation duration model, using mean absolute error (MAE). We similarly stratify the measurement of MAEs for different years, different diagnostic groups, different ICU types, and different admission sources. We additionally compare model performance for the prediction of Total Ventilation Duration made by our model to the APACHE IVa and APACHE IVb models.

Results and discussion:

Ventilation probability model

The validation of Philips CCOPM ventilation models was performed on the entire eCareManager archived databases used for quarterly benchmarking which, to date, includes more than 6 million ICU patient stays across more than 500 hospitals during the period of 2010-2022. The ventilation probability model was evaluated through measuring the area under the receiver operator curve (AUC) for the following individual binary predictive tasks: 1) the probability of any ventilation 2) the probability of any invasive ventilation, and 3) the probability of any non-invasive ventilation. Table 1 provides these AUCs across the database.

Table 1. Prediction performance measurements (AUC) for the following tasks: 1) any ventilation, 2) any invasive ventilation, 3) any non-invasive ventilation.

Ventilation type	Any ventilation	Any invasive ventilation	Any non-invasive ventilation
AUC	0.916	0.934	0.827

We further investigated AUC values across different years (Figure 1), different ICU types (Figure 2), admission diagnostics groups (Figure 3) and patient's admission sources (Figure 4). According to the results in Figures 1-4, the Philips ventilation probability model performance remained relatively consistent across different years and ICU types.

Higher variabilities were observed across diagnostics groups. We also observed a change in the performance of the model during 2020-2022 likely due to the COVID-19 pandemic. The ventilation needs of patients with COVID-19 would have differed from the patient population the model was trained on (2010-2019).

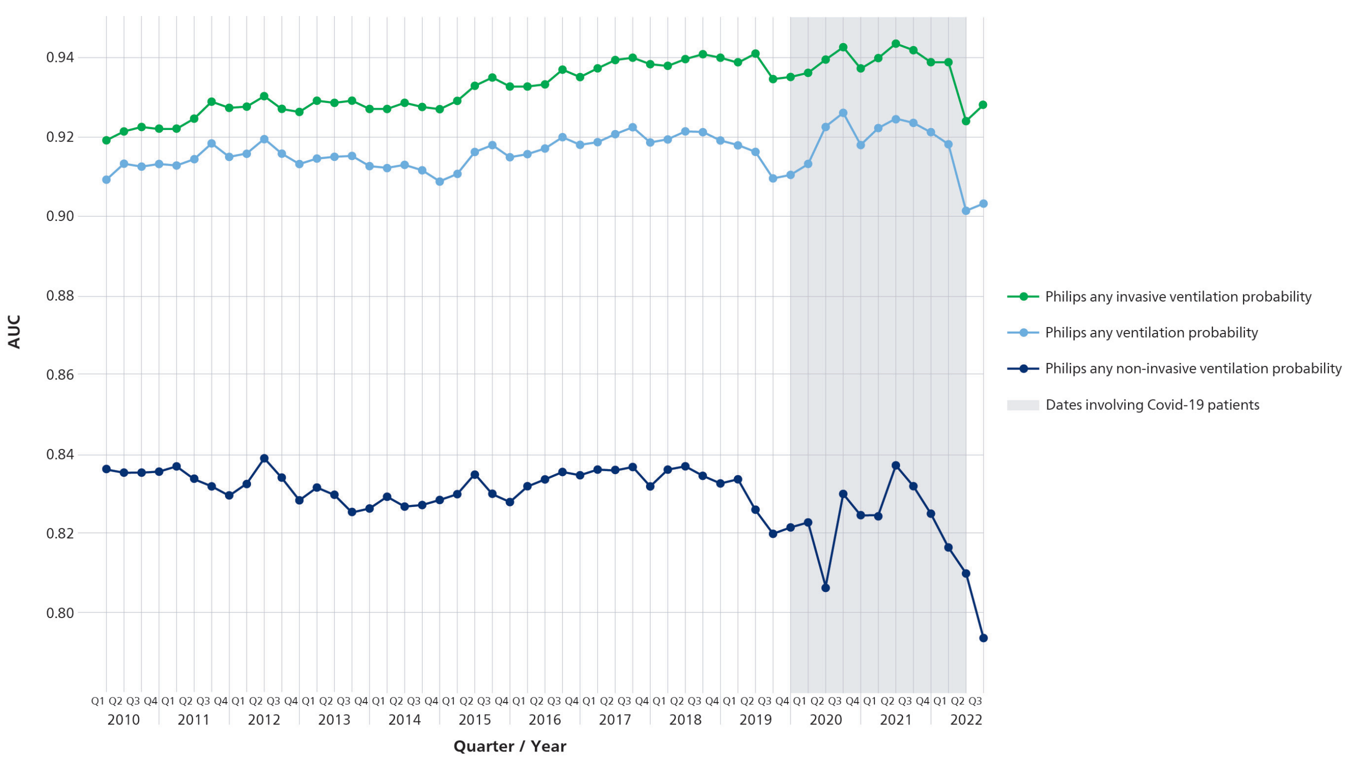

Figure 1. Area under the receiver operator (AUC) for the Philips probability model across different years from 2010-2022. Blurred results for dates involving Covid-19 patients between 2020 Q1 and 2022 Q2.

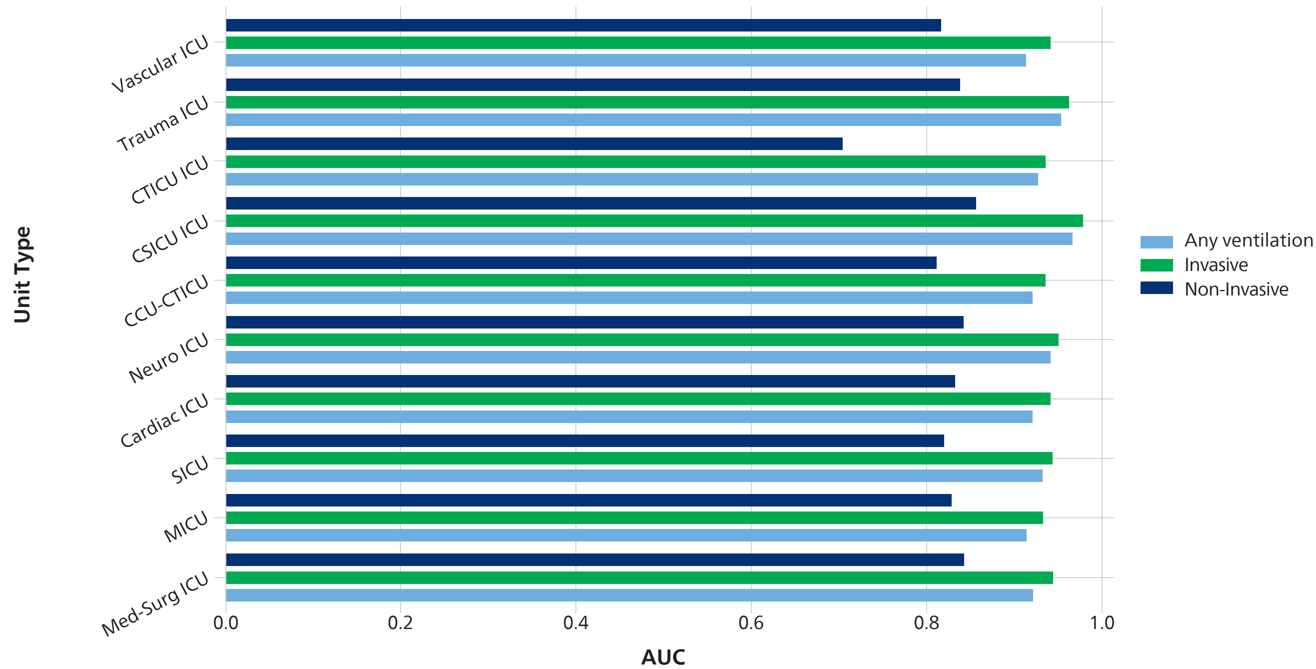

Figure 2. Area under the receiver operator (AUC) for the Philips probability model across ICU unit types.

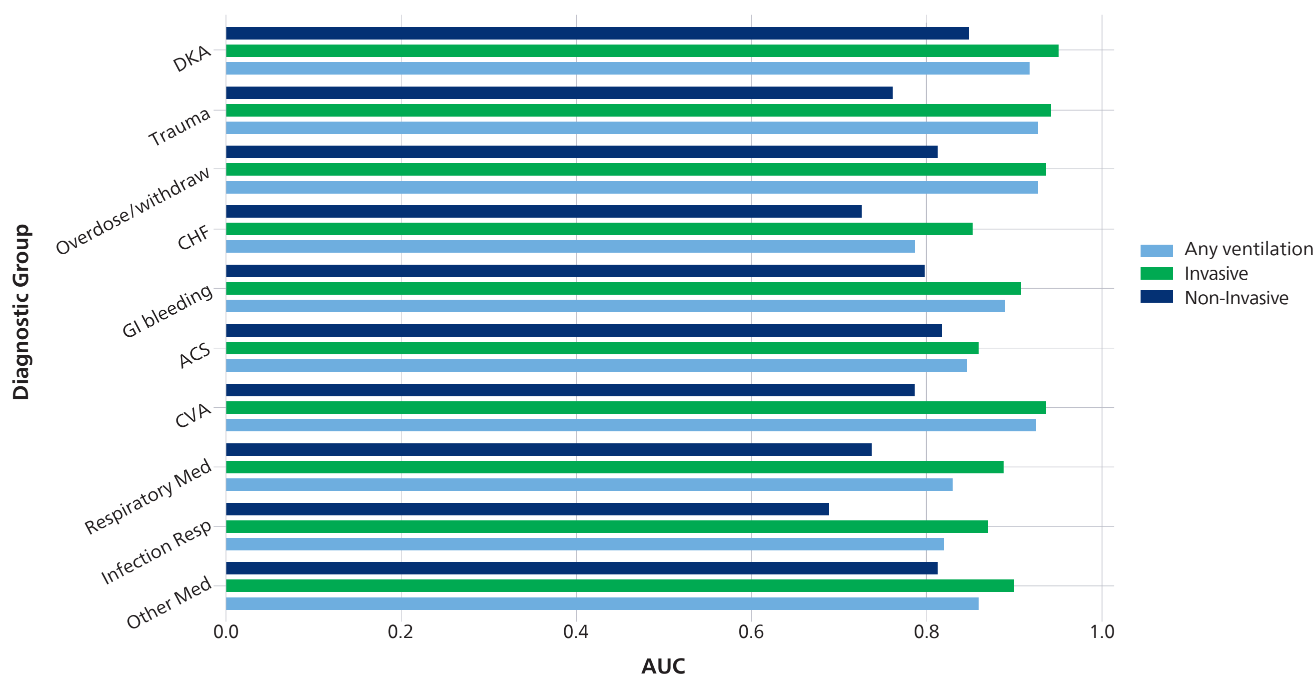

Figure 3. Area under the receiver operator (AUC) for Philips probability model across diagnostics groups.

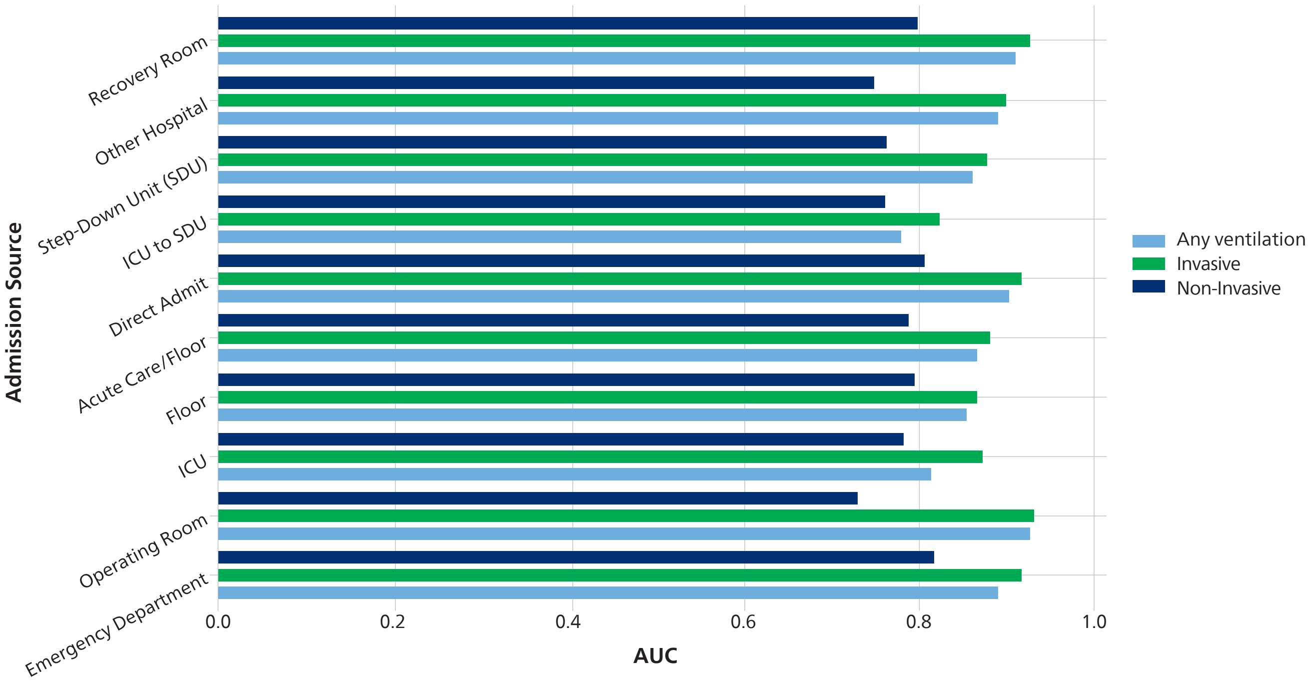

Figure 4. Area under the receiver operator (AUC) for the Philips probability model across various admission sources.

Ventilation duration model

Additionally, we evaluated the performance of the Philips ventilation duration model on the eCareManager database for the period of 2010 to 2022. The mean absolute error (MAE) is computed to evaluate the model performance, where the error is the difference between the true duration values and the predictions.

In Table 2, we provide MAE for both the invasive ventilation duration and the non-invasive ventilation duration across the entire eCareManager database.

In Table 3, we compare the performance of the Philips model on predicting the total ventilation duration with predictions made by APACHE-IVa and APACHE-IVb. The data used in Table 3 is from 2014 as that's when APACHE IVb was made available.

Table 2. Mean absolute error (MAE) in days between the true ventilation duration and Philips ventilation duration model for both invasive ventilation duration and non-invasive ventilation duration.

Ventilation type	Philips invasive ventilation duration	Philips non-invasive ventilation duration
Mean absolute error (days)	1.968	0.521

Table 3. Mean absolute error (MAE) in days between the true total ventilation duration and predictions from the following models: Philips ventilation duration model, APACHE-IVa and APACHE-IVb.

Total duration of ventilation	Philips total duration of ventilation	APACHE IVa total duration of ventilation	APACHE IVb total duration of ventilation
Mean absolute error (days)	2.093	2.623	2.538

Figure 5 shows that the mean absolute error (MAE) values for Philips CCOPM-Ventilation models are consistent across different years. Moreover, MAE for Philips model is consistently lower than both APACHE models across all years between 2010-2022. It also shows that accuracy of APACHE IVb is slightly improved compared to APACHE IVa. Both Philips and APACHE model performances were affected by the Covid patient population during 2020-2022, but Philips models showed more resilience compared to APACHE models (see Figure 5).

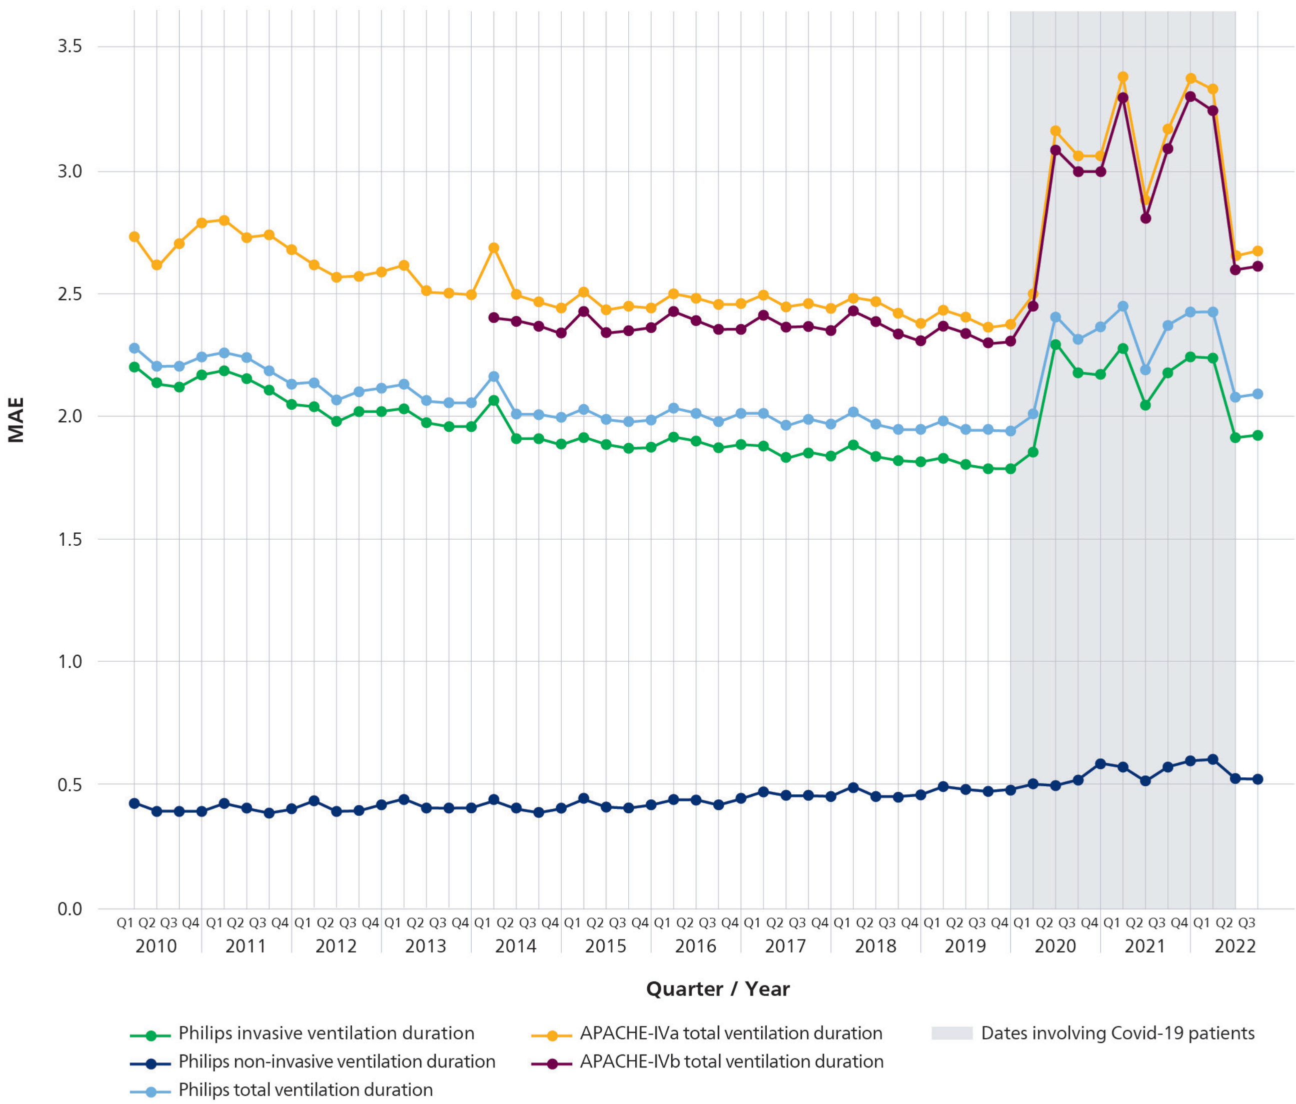

Figure 5. Mean absolute error (MAE) across multiple years from 2010 to 2022 for Philips invasive ventilation duration model (red), Philips non-invasive ventilation duration model (blue), Philips total ventilation duration model (purple), APACHE-IVa total ventilation duration model (gray), APACHE-IVb total ventilation duration model (yellow). Shadedresult for period involving Covid-19 patients between 2020 Q1 and 2022 Q2.

We demonstrate the total duration performance of the Philips model, APACHE IVa and APACHE IVb across different years (Figure 5), diagnostics groups (Figure 6), ICU types (Figure 7) and admission sources (Figure 8). The results indicate that while there are some variabilities in the performance of these models across these categories, the Philips model consistently outperforms both APACHE-IVa and APACHE-IVb across all diagnostics groups, ICU types, and admission sources.

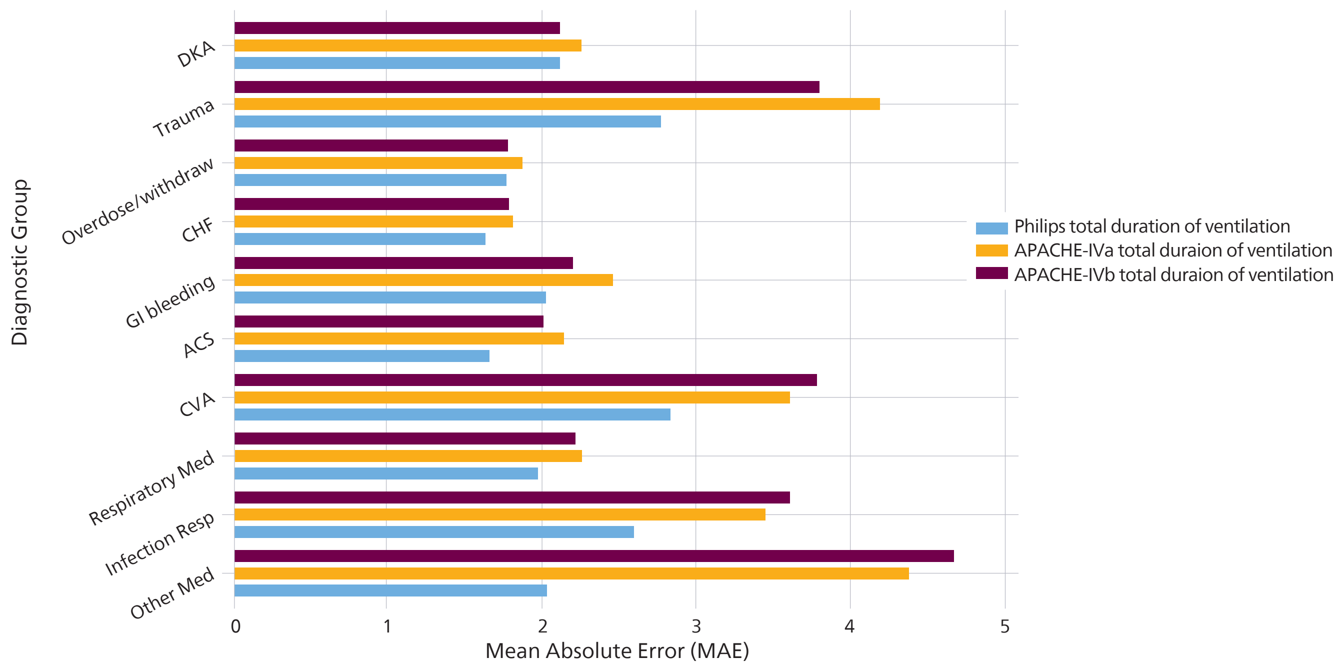

Figure 6. Mean absolute error (MAE) in days for the prediction of total ventilation duration across diagnostic groups for 1) Philips total duration model, 2) APACHE-IVa total duration model and 3) APACHE-IVb total duration model.

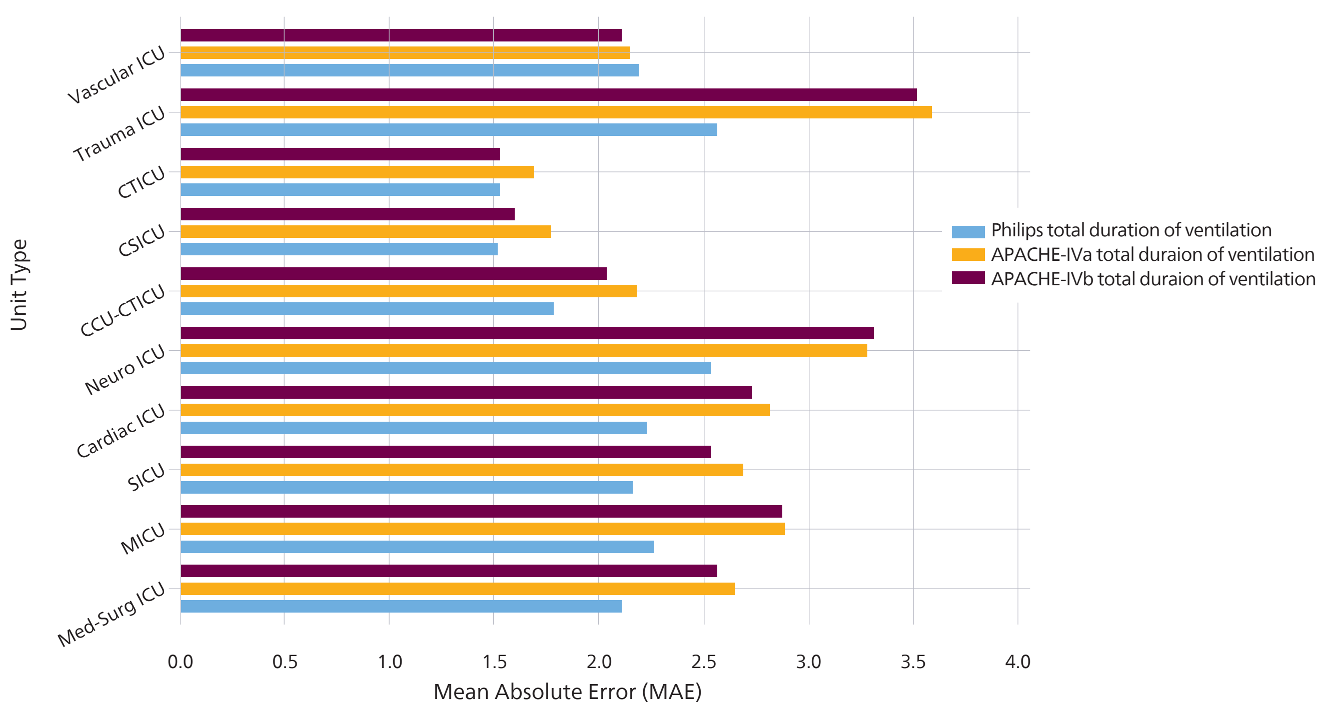

Figure 7. Mean absolute error (MAE) in days for the prediction of total ventilation duration across unit types for 1) Philips total duration model, 2) APACHE-IVa total duration model and 3) APACHE-IVb total duration model.

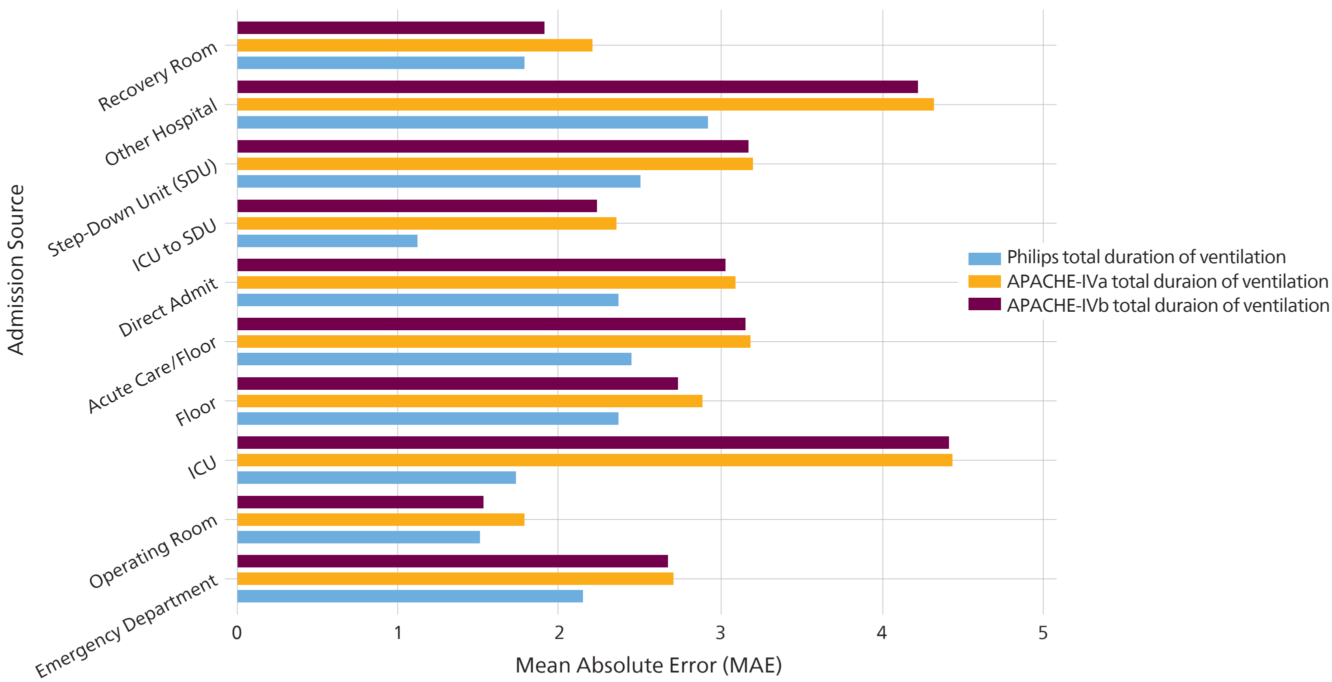

Figure 8. Mean absolute error (MAE) in days for the prediction of total ventilation duration across admission sources for 1) Philips total duration model, 2) APACHE-IVa total duration model and 3) APACHE-IVb total duration model.

Conclusion:

In this analysis, we introduced the Philips CCOPM-Ventilation probability duration models to provide better insight and more accurate outcome predictions to improve ventilation benchmarking in the ICU. We showed strong predictive performance of both ventilation probability model and ventilation invasive/non-invasive duration models across different years, ICU types, diagnostic groups, and admission sources. We also demonstrated that the Philips CCOPM-Ventilation models, predicting total ventilation duration, consistently outperformed both APACHE IVa and IVb ventilation models.

References

[1] Slutsky, A. S., & Ranieri, V. M. (2013). Ventilator-induced lung injury. New England Journal of Medicine, 369(22), 2126-2136.
[2] Nabian, M., & Narusawa, U. (2018). Patient-specific optimization of mechanical ventilation for patients with acute respiratory distress syndrome using quasi-static pulmonary PV data. Informatics in Medicine Unlocked, 12, 44-55.
[3] Sayed, M., Riaño, D., & Villar, J. (2021). Predicting Duration of Mechanical Ventilation in Acute Respiratory Distress Syndrome Using Supervised Machine Learning. Journal of Clinical Medicine, 10(17), 3824.
[4] Zimmerman, J. E., Kramer, A. A., McNair, D. S., & Malila, F. M. (2006). Acute Physiology and Chronic Health Evaluation (APACHE) IV: hospital mortality assessment for today's critically ill patients. Critical care medicine, 34(5), 1297-1310.
[5] Natekin, A., & Knoll, A. (2013). Gradient boosting machines, a tutorial. Frontiers in neurorobotics, 7, 21.

www.philips.com

00000051-00-00 * FEB 2023