

SmartAlert

Machine learning-based patient-ventilator asynchrony detection system in intensive care units

Pažout, Jaroslav; Němý, Milan; Mikeš, Jakub; Jirman, Jan; Kubr, Jan; Niebauerová, Eliška; Macík, Miroslav; Babuška, Robert; Duška, František; More Authors

DOI

[10.1016/j.cmpb.2025.108927](https://doi.org/10.1016/j.cmpb.2025.108927)

Publication date

2025

Document Version

Final published version

Published in

Computer Methods and Programs in Biomedicine

Citation (APA)

Pažout, J., Němý, M., Mikeš, J., Jirman, J., Kubr, J., Niebauerová, E., Macík, M., Babuška, R., Duška, F., & More Authors (2025). SmartAlert: Machine learning-based patient-ventilator asynchrony detection system in intensive care units. *Computer Methods and Programs in Biomedicine*, 269, Article 108927. <https://doi.org/10.1016/j.cmpb.2025.108927>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Contents lists available at ScienceDirect

Computer Methods and Programs in Biomedicine

journal homepage: www.sciencedirect.com/journal/computer-methods-and-programs-in-biomedicine

SmartAlert: Machine learning-based patient-ventilator asynchrony detection system in intensive care units

Jaroslav Pažout^a, Milan Němý^{b,c,*}, Jakub Mikeš^a, Jan Jirman^c, Jan Kubr^d,
 Eliška Niebauerová^a, Miroslav Macík^d, Michal Pech^c, Michal Štajnrt^a, Jakub Vaněk^c,
 Petr Waldauf^a, Václav Zvoníček^a, Lenka Vysloužilová^c, Robert Babuška^{e,f,#},
 František Duška^{a,#}, on behalf of VentConnect Study group

^a Department of Anesthesiology and Intensive Care, 3rd Faculty of Medicine, Charles University and Kralovské Vinohrady University Hospital in Prague, Šrobárova 50, Prague 100 34, Czech Republic

^b Division of Clinical Geriatrics, Center for Alzheimer Research, Department of Neurobiology, Care Sciences and Society, Karolinska Institute, Blickagången 16, Huddinge 14183, Sweden

^c Department of Cognitive Systems and Neurosciences, Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague, Jugoslávských partyzánů 1580/3, Prague 160 00, Czech Republic

^d Department of Computer Graphics and Interaction, Faculty of Electrical Engineering, Czech Technical University in Prague, Karlovo náměstí 13, Prague 121 35, Czech Republic

^e Department of Artificial Intelligence, Czech Institute of Informatics, Robotics, and Cybernetics, Czech Technical University in Prague, Jugoslávských partyzánů 1580/3, Prague 160 00, Czech Republic

^f Cognitive Robotics, Faculty of 3mE, Delft University of Technology, Mekelweg 2, Delft, CD 2628, The Netherlands

ARTICLE INFO

Keywords:

Patient-ventilator asynchrony
 Mechanical ventilation
 Real-time monitoring
 Deep neural networks

ABSTRACT

Background and Objective: Patient-ventilator asynchronies (PVA) are associated with ventilator-induced lung injury and increased mortality. Current detection methods rely on static thresholds, extensive preprocessing, or proprietary ventilator data. This study aimed to develop and validate a fully online, real-time system that detects and classifies PVAs directly from ventilator screen data while alerting clinicians based on severity.

Methods: The SmartAlert system was developed using ventilator screen recordings from ICU patients. It extracts pressure and flow waveforms from video recordings, converts them into time-series data, and employs deep neural networks to classify asynchronies and assign alarm levels from no urgency to most urgent. A dataset of 381,280 double-breath units was independently annotated by two expert intensivists. Two deep learning models were trained: one for alarm prediction and another for asynchrony classification (ineffective triggering, double cycling, high inspiratory effort, no asynchrony). Performance was evaluated using accuracy, sensitivity, specificity, and AUC-ROC, compared to expert consensus.

Results: SmartAlert demonstrated strong performance for alarm level prediction (overall accuracy: 83.8 %, weighted AUC-ROC: 0.943 [95 % CI: 0.941–0.945]) and PVA classification (weighted accuracy: 89.3 %, weighted AUC-ROC: 0.951 [95 % CI: 0.950–0.953]). It showed high specificity for urgent alarms (99.9 % for level 3) and PVA types (98.5 % for ineffective triggering, 96.9 % for double cycling, 94.8 % for high inspiratory effort).

Conclusions: We developed and internally validated SmartAlert, an automated system that detects PVAs, classifies severity, and alerts clinicians in real time. Its potential to reduce alarm fatigue, optimize ventilator settings, and improve patient outcomes remains to be tested in clinical trials.

* Corresponding author at: Milan Němý, Department of Cognitive Systems and Neurosciences, Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague, Prague, Czech Republic.

E-mail address: milan.nemy@cvut.cz (M. Němý).

Shared senior authors

<https://doi.org/10.1016/j.cmpb.2025.108927>

Received 27 February 2025; Received in revised form 3 June 2025; Accepted 19 June 2025

Available online 21 June 2025

0169-2607/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Positive pressure mechanical ventilation (MV) not only provides essential organ support for critically ill patients but it also carries risk of ventilator-induced lung injury (VILI). The mechanisms of VILI range from shear stress-induced inflammation to classical barotrauma. Pre-existing lung injuries further increase susceptibility to VILI, emphasizing the importance of optimal ventilator settings to minimize harm [1,2]. Evidence supports maintaining tidal volumes, plateau pressures, driving pressures, and mechanical power within safe limits, as recommended by guidelines [3–6]. However, dynamic patient physiology necessitates frequent ventilator adjustments by experienced operators, which is not always feasible in a busy ICU environment. Furthermore, alarm fatigue, where constant monitoring alarms desensitize providers, can delay responses to critical events and endanger patient care [7]. Moreover, clinically significant events that do not trigger alarms may go unnoticed. There is an urgent need for systems that strike an optimal balance between frequency and the detection of critical events [8]. This need became even more apparent during the COVID-19 pandemic, where barrier nursing and the use of personal protective equipment further limited clinicians' ability to monitor and adjust ventilator settings, thereby increasing the risk of VILI.

One critical issue in ventilator management is patient-ventilator asynchrony (PVA), in which the ventilator's mechanical support does not align with the patient's respiratory efforts [9–11]. PVAs including ineffective triggering (IE), double cycling (DC), and high inspiratory effort (hIE), are linked to various adverse effects, such as patient discomfort, reduced sleep quality, increased needs for sedation and/or muscle relaxant, and aggravation of VILI. These complications can result in prolonged mechanical ventilation, and even increased mortality, highlighting the urgent need for timely PVA detection and management [11–14]. Despite these severe consequences, the true incidence of PVA may be significantly underestimated due to inadequate monitoring or clinicians' limited experience in identifying subtle asynchronies [15].

Numerous studies have attempted to address PVA detection using traditional rule-based algorithms that rely on predefined thresholds or feature engineering. Early work by Mulqueeny et al. [16] and Blanch et al. [17] demonstrated automated detection of PVAs, such as IE and DC. However, these methods were limited by their reliance on fixed thresholds that do not adapt to changes in patient physiology (e.g., variations in respiratory effort or mechanical lung properties) and thus lacked the flexibility needed for real-time application in the ever-changing conditions of critically ill patients.

Recently, machine learning (ML) has shown promise in advancing PVA detection. Zhang et al. [18] employed ventilator waveforms in image format and applied convolutional neural networks (CNNs) to classify various PVA subtypes. Although this image-based approach achieved high accuracy, it required computationally intensive pre-processing and was limited by a small dataset. Pan et al. [19] took a different approach by using 1D CNNs to analyze ventilator waveforms, offering a more efficient method for waveform-based PVA detection. In a subsequent study, Pan et al. [20] further improved the methodology by utilizing a recurrent neural network architecture, specifically a Long Short-Term Memory (LSTM) network, to capture the temporal patterns in the waveforms, thereby enhancing classification performance. While these methods demonstrated significant progress, they often relied on separate models for each PVA subtype, which limits scalability, and were not designed for fully online, real-time clinical applications. Moreover, most existing approaches depend on proprietary ventilator data outputs, reducing their versatility for use across diverse ventilator models.

During the COVID-19 pandemic, we developed an interface called VentConnect (see www.ventconnect.cz for details), which enables the encrypted online transfer of ventilator screen data to any computer device. This platform helps intensivists remotely monitor ventilator screens and identify potentially dangerous patterns, such as PVAs. As an

extension of the VentConnect project, we aimed to further develop this technology into a system called SmartAlert, a fully online, real-time solution capable of (1) automatically detecting and classifying patterns like PVAs from ventilator screen data and (2) notifying clinicians via a custom-designed application with alerts tailored to the severity of detected patterns. SmartAlert is designed to ensure continuous monitoring without overwhelming clinicians, thereby addressing one of the key challenges in effective PVA management.

In this paper, we describe the technical details of developing the SmartAlert system and present the results of a study evaluating its performance against a gold standard, defined as the agreement among two fully qualified intensivists specially trained to recognize PVAs.

2. Methods

2.1. Study workflow and system overview

This study aimed to develop a system capable of capturing ventilator screens in their raw visual form and, through a series of signal transformations and advanced machine learning techniques, detecting and classifying patient-ventilator asynchronies (PVA) as they appear on the screen. The system was designed to output an alarm level that indicates the severity of the situation, thereby providing actionable information to clinicians so that ventilator settings can be adjusted for more optimal lung ventilation. The workflow involved waveform digitization via the VentConnect system, annotation of the pressure and flow waveforms, and training deep learning classifiers for both PVA detection and alarm level prediction.

2.2. Patient selection

All adult patients admitted to the ICU at University Hospital Královské Vinohrady, Czech Republic, between January 2021 and August 2022 were included in the study. The cohort comprised a mixed general ICU population, including patients with major trauma, major elective surgery, and COVID-19. However, patients with burns, solid organ transplants, or those undergoing cardiac surgery were excluded. All included patients required mechanical ventilation.

2.3. Research ethics and data protection

The VentConnect system was originally developed during COVID-19 pandemic to facilitate the clinical care of patients receiving mechanical ventilation. It enabled intensivists to remotely monitor ventilator settings in real time for numerous patients. Anonymized ventilatory waveform data were stored within the system and later used for algorithm training. No other individual patient data were recorded, processed, or stored for this study. The Královské Vinohrady Research Ethics Board reviewed the study, approved retrospective use of waveform data, and waived the need for informed consent.

2.4. VentConnect system overview

This study utilized our in-house VentConnect system to capture and analyze lung ventilator waveforms. VentConnect provides access to video output streams from mechanical ventilators and vital signs monitors, enabling both real-time data streaming and retrospective analysis of stored snapshots.

Fig. 1 illustrates the system's core architecture, which comprises hardware capture units (VentUnits), a central data server, storage infrastructure, and analytical components.

The system captures video from ventilator displays and automatically digitizes waveform data. It incorporates pattern recognition algorithms for detecting patient-ventilator asynchrony (PVA) and performs threshold-based monitoring of critical ventilation parameters. When potentially suboptimal ventilation states are identified, the system

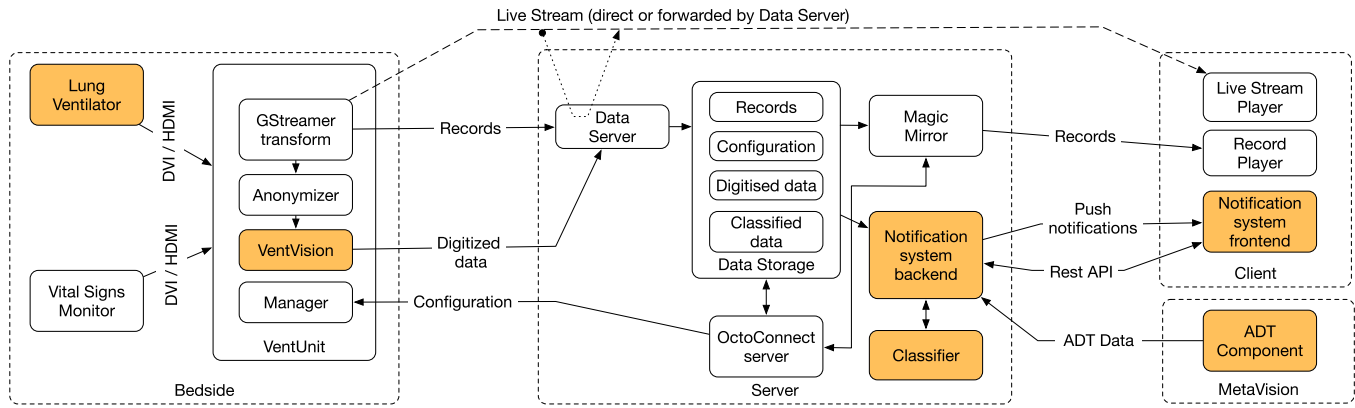


Fig. 1. Schema of VentConnect System Architecture with Notification system extension. Components important for notification system are highlighted in orange.

generates alert notifications for healthcare professionals. This functionality forms the foundation for the analyses described in this study.

For real-time use, only the digitized time-series data are stored and subsequently processed. In addition, a 10-second video recording is captured every 15 min and stored for review purposes and for retrospective analyses in studies like this one. In this configuration, the system processes waveform data with a total latency of approximately 1–1.5 s from screen capture to alarm generation. Waveform extraction is executed locally on each VentUnit with minimal delay (50–100 ms per full screen), and classification is performed centrally on the VentConnect server, with parallelized processing to ensure responsiveness even when monitoring multiple ventilators simultaneously. A comprehensive description of the VentConnect architecture, including detailed component specifications and data flow, is provided in the Supplementary materials.

2.5. Definition of patient-ventilator asynchronies

In this study, we adopted the adult classification of patient-ventilator asynchronies (PVA) to annotate and classify the asynchronies most commonly observed in clinical practice. Among the various PVAs described in the literature, we focused on three severe types: ineffective triggering (IE), double cycling (DC), and high inspiratory effort (hIE). These PVAs were selected due to their prevalence and their potential to disrupt synchrony between the patient's respiratory efforts and the ventilator's mechanical support, posing significant risks to patient comfort and safety if left unaddressed.

Table 1 provides the definitions for these PVAs, adapted for the purposes of our study. These definitions formed the basis for annotating the training dataset and developing the classification algorithm described in subsequent sections.

2.6. Ventilator screen-to-signal conversion

The pressure (airway pressure) and flow (volumetric flow rate) waveforms analyzed in this study were recorded from two types of lung ventilators, the Hamilton G5 and the Dräger Evita Infinity V500, via the VentConnect system. To facilitate subsequent analyses, we developed a robust data processing pipeline to convert these video-recorded waveforms into accurate multivariate time-series data. All processing was performed using MATLAB R2023a.

The pipeline begins by precisely cropping each video to isolate regions displaying the pressure and flow waveforms. To mitigate compression artifacts inherent in lossy image formats, color quantization was applied using k-means clustering.

Next, the X (time) and Y (pressure or flow) axes of the waveforms were digitized. Tick marks on the axes were identified, and adjacent

Table 1

Patient-ventilator asynchronies selected for annotating respiratory cycles.

Descriptions of these three severe types of PVAs were adopted from adult classification [21] to annotate and classify ventilator waveforms.

PVA Type	Definition
Ineffective Triggering (IE)	A patient's inspiratory effort fails to trigger a ventilator-delivered breath. This is characterized by a positive deflection in flow and a negative deflection in airway pressure that is not followed by an inspiratory phase, increasing the patient's work of breathing and leading to discomfort.
Double Cycling (DC)	Two consecutive ventilator-delivered breaths triggered by a single patient inspiratory effort, appearing as two breaths in rapid succession on the pressure and flow curves. This phenomenon is often caused by inadequate ventilator settings, resulting in unnecessary mechanical insufflations.
High Inspiratory Effort (hIE)	An excessive decrease in pleural pressure during inspiration, observed as either a slower increase in inspiratory pressure or a negative pressure deflection in the initial phase of inspiration. This reflects a mismatch between the ventilator's flow delivery and the patient's effort, leading to discomfort and suboptimal ventilation. It is also known as flow dyssynchrony.

numeric labels were located and processed to extract their values. Each numeric label, represented as a graphical element, was compared against a library of pre-extracted templates using the Dice coefficient to accurately decode its corresponding numerical value. The positions of the tick marks and their numerical equivalents were then used to compute scales for the axes, enabling the mapping of each pixel within the waveform region to its true temporal and amplitude values. This mapping provided a precise transformation of the visual waveform data into quantifiable time-series data.

The waveform signal itself was extracted by thresholding of the relevant color components in the segmented images. Minor discontinuities in the extracted signal, typically spanning one to two pixels and attributable to encoding artifacts, were rectified using linear interpolation to ensure continuity.

This process yielded high-fidelity, multivariate pressure and flow time-series data, which formed the basis for subsequent analyses, including the detection of inspiratory events, segmentation into double-breath units, and their classification.

For the purpose of this study, waveform extraction was performed offline as described above. However, in our real-time SmartAlert system, waveform extraction is performed directly on each VentConnect unit using code written in the C programming language. The process takes approximately 1 % of the time needed to fill a full screen. To accommodate display layout variability, the system uses automated layout detection. This detection relies on distinctive features for each

ventilation mode and layout, such as the color scheme, shapes, and positions of waveform panels and control buttons. Templates were prepared for the most common ventilation modes and settings used in our ICU, and personnel are instructed to use only supported layouts to ensure consistent extraction. Additionally, a built-in validity check based on visual landmarks ensures that only screenshots with complete and correctly aligned waveforms are processed.

2.7. Validation of screen-to-signal conversion against data logger records

To validate the accuracy of our screen-to-signal conversion pipeline, we compared the derived time-series data with reference signals recorded using the Hamilton Datalogger software. This software communicates with Hamilton ventilators via an RS232 interface and provides real-time recordings of waveforms in a tabular format. These data logger signals were used as a proxy for ground truth in the validation process.

The data logger recorded signals at a sampling frequency of 31.25 Hz. To enable a direct comparison, both the data logger signals and our screen-derived time-series were upsampled to a common frequency of 100 Hz using linear interpolation. The accuracy of the derived signals was assessed using the following metrics:

1. Root Mean Squared Error (RMSE): This metric quantifies the absolute differences between corresponding points in the two signals using the formula:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - y_i)^2},$$

where x_i and y_i represent the derived and ground truth values at each time point, respectively, and N is the total number of points.

2. Normalized Root Mean Squared Error (NRMSE): This metric normalizes the RMSE relative to the dynamic range of the ground truth signal, allowing for comparisons across signals of varying magnitudes:

$$NRMSE = \frac{RMSE}{\max(y) - \min(y)},$$

where $\max(y)$ and $\min(y)$ are the maximum and minimum values of the ground truth signal.

3. Pearson Correlation Coefficient: This coefficient was calculated to evaluate the linear relationship between the two signals, quantifying how well the patterns in both signals correspond regardless of their magnitudes.

To ensure robust validation, we analyzed signals from 466 consecutive full-screen recordings in SPONT ventilation mode, yielding a comprehensive dataset that spans multiple ventilation cycles and varying conditions.

2.8. Detection of inspiratory events

For automatic classification, we first identified inspiratory events using proprietary in-house software. The multivariate pressure/flow signal was uniformly resampled to 100 Hz and subsequently low-pass filtered using a 25th-order Finite Impulse Response (FIR) filter, with a bandpass frequency of 10 Hz and a stopband frequency of 20 Hz. An approximate derivative of both the pressure and flow signals was then computed utilizing a first-order difference method.

The onset of the inspiratory event was defined as the beginning of a rapid rising edge, observable concurrently in both the pressure and flow waveforms. Candidate time points for inspiratory events were flagged when the derivative of the flow signal exceeded 1.5 times its standard deviation (calculated over a single screen length) and the derivative of the pressure signal was positive.

To refine these detections, the algorithm retrospectively examined the preceding 25 data points of the flow derivative for each candidate event. The definitive initiation point of the inspiratory event was determined as the last instance within this window where the flow derivative fell below its standard deviation.

Using the detected inspiratory events, the pressure/flow time series was segmented into sections spanning three consecutive inspiratory events (i.e., two complete breath cycles). These segments, referred to as double-breath units, formed the input for further processing and classification. This segmentation approach was chosen to provide sufficient temporal context for detecting patterns like double cycling, which by definition span across adjacent breaths. It can also benefit the detection of other PVA types by allowing the model to evaluate short waveform sequences, consistent with how clinicians interpret ventilator data.

Fig. 2 presents a schematic overview of the inspiratory event detection algorithm.

To validate the temporal accuracy of the detected inspiration onsets, we compared the timing of our detection approach to the reference onset times using the same RS-232-based dataset described in Section 2.7. Temporal alignment was assessed by calculating the difference between the algorithm-detected onset and the reference.

2.9. Clustering-aided annotation

To enable automatic classification of the waveforms, the segmented double-breath units were annotated by medical experts (J. P., J. M., P. W., M. S.), using a custom-developed, web-based application. The experts evaluated each double-breath unit to identify the type of patient-ventilator asynchrony (PVA) present and assign an appropriate alarm state. The annotated PVAs included IE, DC, hIE, as well as the absence of asynchrony (NO).

Alarm states were defined to indicate the severity of the condition (Table 2), supplementing the ventilator's built-in alert system, with levels ranging from 0 (no immediate attention needed) to 3 (urgent attention required).

To streamline the annotation process for the large volume of signals, the double-breath units were initially grouped into clusters based on their similarity. This was achieved using partitional clustering with Dynamic Time Warping as the distance measure and Partition Around Medoids for centroid determination [22]. The signals were first pre-processed through z-normalization, and the optimal number of clusters per each ventilation mode was determined by exploring a range from 2 to 1000 clusters using silhouette analysis.

Annotations were performed independently by three medical experts. In cases of disagreement, the experts discussed the contentious instances until a consensus was reached. If a consensus could not be achieved, the final decision was made by the most experienced expert.

For clarity, Table 3 summarizes the key characteristics of the ventilation modes included in this study. These labels correspond to the modes reported by the ventilator interface and were used consistently throughout the clustering and annotation process.

2.10. Deep neural network architecture and training

We developed two deep neural networks (DNNs) to classify lung ventilator waveforms based on segmented double-breath units. Both models were trained on the same annotated dataset but were designed for distinct classification tasks: one for predicting alarm levels (ranging from 0 to 3) and the other for identifying the type of asynchrony (IE, DC, hIE, NO). Due to the different output structures and evaluation criteria,

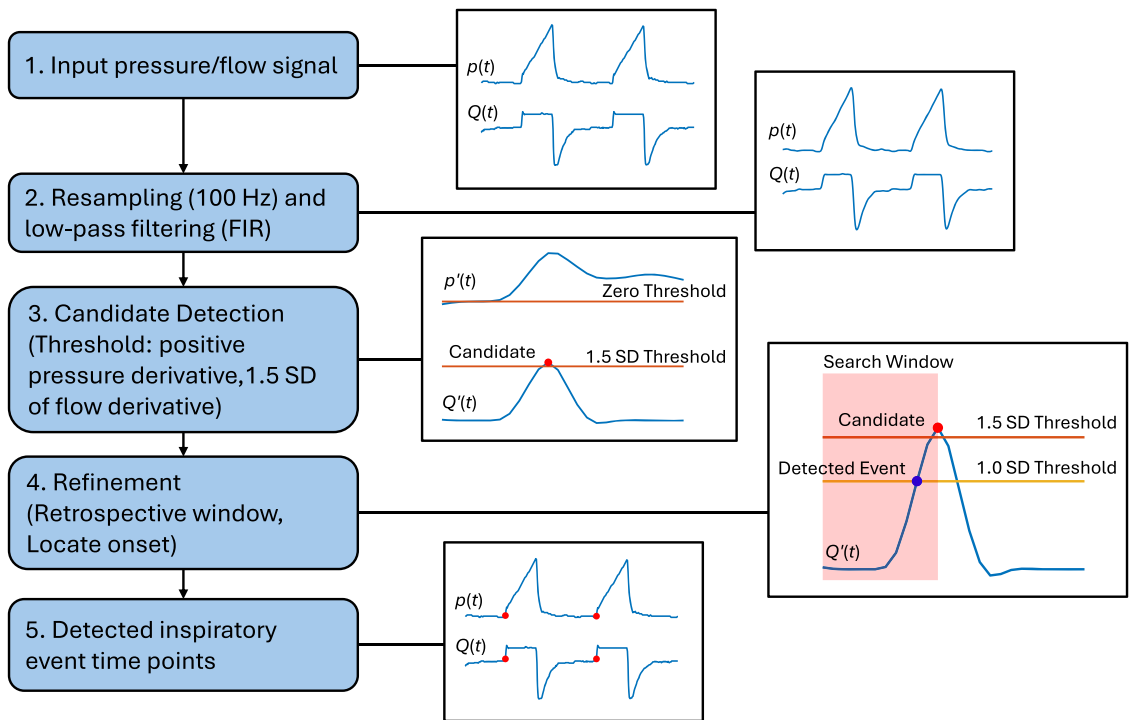


Fig. 2. Workflow of the detection algorithm for inspiratory event identification. SD, standard deviation; FIR, finite impulse response filter; $p(t)$, pressure waveform (time signal); $Q(t)$, flow waveform (time signal); $p'(t)$, approximation of the first time derivative of the pressure signal; $Q'(t)$, approximation of the first time derivative of the flow signal.

Table 2
Description of alarm levels in the SmartAlert system. PVA, patient-ventilator asynchrony; hIE, high inspiratory effort.

Alarm Level	Interpretation	Typical Examples
0	No immediate attention needed	Regular breathing pattern with no signs of PVA
1	Mild abnormality, monitor if persistent	Occasional ineffective effort or minor breath stacking not yet clinically impactful
2	Clinically relevant issue, attention warranted	Recurrent double cycling, frequent high inspiratory effort (hIE) episodes
3	Urgent attention required	Sustained PVA patterns (e.g., persistent hIE), combination of PVAs

and to avoid the complexity of loss balancing, we did not adopt a multitask learning framework. While alarm levels and PVA types are often related in practice, they are not perfectly aligned, and training the models separately allowed for better optimization and flexibility. In deployment, both models are applied in parallel to the same input data.

2.10.1. Input representation

The models processed two input streams: (1) variable-length time-series data representing the double-breath units, this input consisted of two channels (pressure and flow) and (2) a fixed-length numerical vector containing 11 features. This vector included a one-hot encoded representation of the ventilation mode (five selected modes and an 'OTHER' category) and five additional ventilation parameters: peak inspiratory pressure (Ppeak), expiratory tidal volume (VTE), inspiratory-to-expiratory ratio (I:E), expiratory minute volume (ExpMinVol), and respiratory rate (fTotal). Missing numerical values were encoded as -1. The time-series waveform input was FIR-filtered, and no normalization was applied along the time or amplitude axes during either training or inference.

Table 3
Ventilation mode descriptions used in this study.

Ventilation Mode	Description
SPONT (Spontaneous mode)	A continuous positive airway pressure (CPAP) mode in which the patient breathes spontaneously, triggering pressure-supported breaths with flow-cycled termination.
ASV (Adaptive Support Ventilation)	A closed-loop mode that automatically adapts pressure and respiratory rate based on patient mechanics and effort.
P-SIMV (Pressure-controlled Synchronized Intermittent Mandatory Ventilation)	Ventilator combines pressure-controlled mandatory breaths with spontaneous breathing opportunities.
SIMV (Synchronized Intermittent Mandatory Ventilation)	Ventilator delivers mandatory breaths at a set interval, with synchronization to spontaneous effort, but allows unassisted spontaneous breathing between them.
DuoPAP (Dual Positive Airway Pressure)	A bilevel mode in which the ventilator alternates between two pressure levels, allowing spontaneous breathing at both levels to support gas exchange and reduce patient effort.

2.10.2. Model architecture

As illustrated in Fig. 3, both DNNs share a similar architecture, implemented using the TensorFlow/Keras functional API. The time-series input is first passed through a masking layer to accommodate sequences of different lengths. This is followed by a bidirectional Long Short-Term Memory (BiLSTM) layer with 128 hidden units, which captures temporal dependencies in the data. A dropout layer (with a rate of 0.5) is applied to reduce overfitting. The output from the BiLSTM is then concatenated with the fixed-length numerical input vector, forming a combined representation that feeds into subsequent layers.

2.10.3. Alarm prediction model

For predicting alarm levels, the combined representation is passed

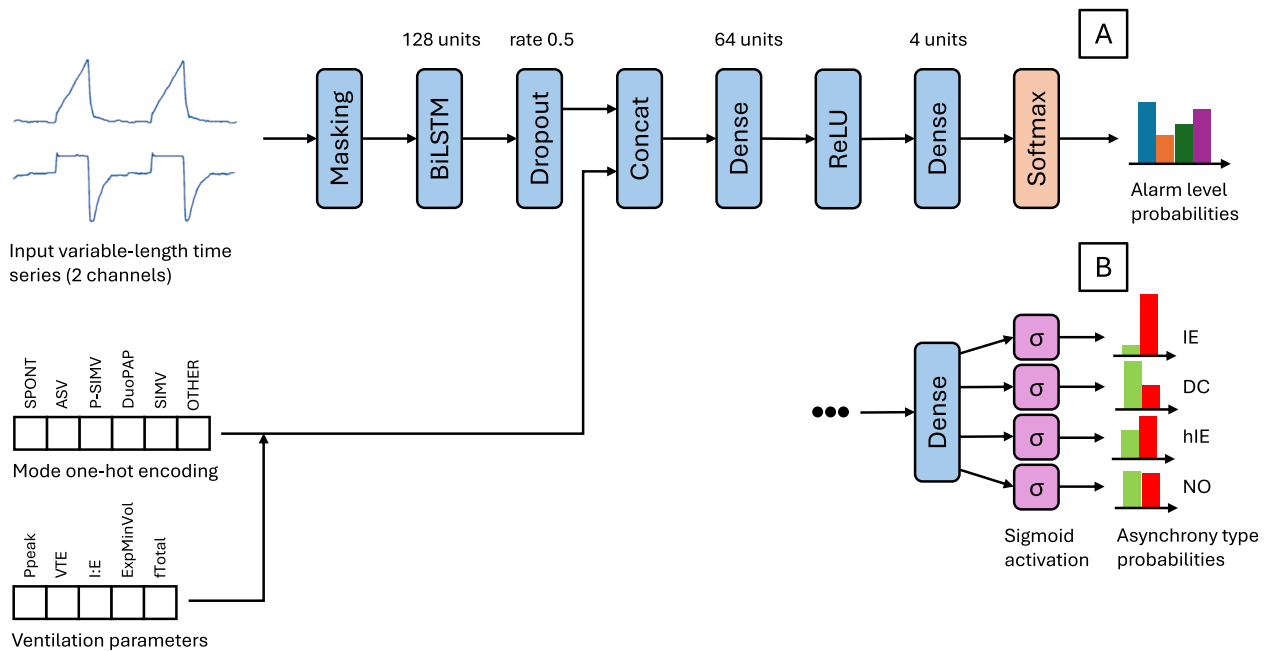


Fig. 3. Architecture of the deep neural networks for classifying lung ventilator waveforms. The models take two inputs: (1) variable-length time-series data of pressure and flow signals (two channels) and (2) a numerical vector comprising one-hot encoded ventilation modes and additional ventilation parameters. The time-series input is processed through a masking layer, a bidirectional LSTM (BiLSTM) layer with 128 units, and a dropout layer (rate = 0.5). The BiLSTM output is concatenated with the numerical input vector, followed by a dense layer with 64 units and ReLU activation. Two separate outputs are generated: (A) alarm level probabilities (softmax activation, 4 classes) predicting the urgency of attention required, and (B) probabilities of asynchrony types (sigmoid activation, multilabel classification for ineffective triggering [IE], double triggering or cycling [DC], high inspiratory effort [hIE], and no asynchrony [NO]).

through a dense layer with 64 units and ReLU activation. The final output layer employs softmax activation for multiclass classification, corresponding to the four alarm levels (0 to 3). This model was trained using the Adam optimizer with a learning rate of 0.0001 and a sparse categorical cross-entropy loss function.

2.10.4. Asynchrony prediction model

This network predicted the presence of one or more asynchrony types (IE, DC, hIE, NO), as annotated by medical experts. Similar to the alarm prediction model, the combined representation from the BiLSTM output and the fixed-length numerical vector was processed through a dense layer with 64 units and ReLU activation. However, the output layer utilized sigmoid activation to accommodate multilabel classification. The model was trained using the Adam optimizer (with a learning rate of 0.0001) with a binary cross-entropy loss function.

The dataset was shuffled and divided into training (70 %), validation (15 %), and testing (15 %) sets. Both models were trained for up to 500 epochs with a batch size of 64. Early stopping with a patience of 50 epochs and model checkpointing based on validation loss were employed to prevent overfitting and to select the best-performing model based on validation metrics.

The models were implemented using TensorFlow version 2.4.1 running in a Singularity container. This container utilized NVIDIA's TensorFlow framework stack, which includes CUDA 11.2 and cuDNN 8.1 libraries.

2.11. Statistical analysis

After training the networks, we evaluated the performance of both models using the testing dataset. The statistical evaluation was performed on separate double-breath units without consideration of their temporal sequence.

For the alarm prediction model, we computed overall accuracy, sensitivity, and specificity across the entire dataset. In addition, these metrics were calculated separately for each alarm level (0, 1, 2, and 3) to

assess the model's ability to differentiate between various degrees of clinical urgency. Both weighted and macro-averaged sensitivity, specificity, and F1-score were reported. The weighted averages account for the frequency of each alarm level in the dataset, while the macro averages provide an unweighted assessment across all levels.

For the asynchrony prediction model, accuracy, sensitivity, and specificity were determined for each label (IE, DC, hIE, and NO) to evaluate the model's performance in identifying each type of asynchrony independently. Similarly, weighted and macro-averaged versions of these metrics were calculated offering both a frequency-weighted evaluation and an equal-weighted overview across all labels.

To further quantify model performance, we calculated one-vs-all AUC-ROC scores for each label or alarm level. In addition, macro- and weighted-average AUC-ROC scores were computed to provide both a balanced view across all classes and a frequency-weighted evaluation. Confidence intervals (CIs) for the AUC-ROC were derived using a bootstrapping procedure with 1000 resamples and the 95 % CIs were reported for both the macro- and weighted-average AUC-ROC scores, as well as for the one-vs-all scores of each individual label or alarm level.

High inspiratory effort (hIE) has an inherently subjective component, and its detectability may vary depending on the ventilation mode. To account for this potential variability, we performed an additional mode-stratified analysis. Specifically, we evaluated the distribution of hIE events and the performance of the classifier for each ventilation mode.

Statistical analysis was performed using Python 3.10. The analyses leveraged the NumPy library (v1.23.5) for numerical operations, and scikit-learn (v1.3.0) for computing classification evaluation metrics such as sensitivity, specificity, accuracy, F1-score, and AUC-ROC.

3. Results

3.1. Patient characteristics

A total of 1863 patients from a mixed general ICU population were included in the study. This cohort encompassed patients with major

trauma, major elective surgery, and COVID-19 cases, while excluding those with burns, solid organ transplants, or cardiac surgery. The mean age of patients was 58.4 ± 17.8 years, with the first and third quartiles at 44.8 and 73.0 years, respectively. The majority of patients were men ($n = 1174$, 63.0 %). The most frequent primary diagnoses were respiratory failure ($n = 448$, 24.0 %), trauma ($n = 318$, 17.1 %), COVID-19 ($n = 260$, 14.0 %), postoperative complications ($n = 245$, 13.2 %), and stroke ($n = 199$, 10.7 %). The overall ICU mortality rate was 11.9 %. Detailed cohort characteristics are presented in Table 4.

3.2. Validation of screen-to signal conversion and inspiration onset detection

The accuracy of the derived signals from the screen-to-signal conversion pipeline was evaluated by comparing them to the ground truth signals recorded using the Hamilton Datalogger. Three metrics were used in this evaluation: Root Mean Squared Error (RMSE), Normalized Root Mean Squared Error (NRMSE) relative to the signal range, and the Pearson correlation coefficient.

For the airway pressure signal, the RMSE was 0.698 cmH₂O, the NRMSE was 2.56 %, and the Pearson correlation coefficient was 0.984 ($p < 0.001$). For the flow signal, the RMSE was 3.633 l/min, the NRMSE was 1.47 %, and the Pearson correlation coefficient was 0.993 ($p < 0.001$).

Next, we assessed the temporal accuracy of the detected inspiration onset. The mean temporal difference between the algorithm-detected and reference onsets was -0.035 s (SD 0.027 s), with a median of -0.030 s (IQR: -0.050 s to -0.020 s).

3.3. Signal preprocessing and transformation

The dataset for training the classification algorithm was derived from an extensive collection of 635,806 video recordings of ventilator screens, each lasting 10 s. These recordings were processed to extract and parse double-breath units for model training.

After applying optical character recognition to the video recordings, 464,097 recordings were identified as originating from Hamilton ventilators, 97,058 from Dräger ventilators, and 74,651 were deemed unrecognizable or faulty. Among the Hamilton recordings, 223,373 were captured in active (non-standby) mode while 240,724 in standby mode. For the Dräger recordings, 54,608 were in non-standby mode and

42,450 in standby mode. Due to the significantly smaller number of recordings from the Dräger devices, only the active recordings from Hamilton devices were retained for the subsequent analysis.

From the non-standby recordings, a total of 219,929 Hamilton recordings were successfully parsed into complete sets of segmented waveforms along with their associated statistics. For further processing, only the largest subsets, SPONT, ASV, P SIMV, DuoPAP, and SIMV, were retained, while the remaining 22,496 recordings representing a mixture of ventilation modes with low counts, were excluded. As shown in Table 5, the SPONT mode accounted for the largest subset with 90,808 recordings, followed by ASV with 43,205 recordings and P SIMV with 34,258 recordings.

The segmented waveforms were then transformed into double-breath units, which serve as the fundamental input for the classification algorithm. Table 5 provides details on the segmentation and clustering outcomes for Hamilton ventilators. In total, Hamilton ventilators contributed 381,280 double-breath units across all ventilation modes, organized into 2490 clusters. Out of these, 340,871 double-breath units were successfully annotated for subsequent training of the classification model.

The annotation process involved identifying and excluding clusters and recordings that contained artifacts or excessive noise, ensuring that only high-quality data were used for model training.

3.4. Prevalence of patient-ventilator asynchronies and alarm level in annotated samples

We used a total of 289,739 annotated double-breath units for training (85 % of the sample) and validation of the deep learning classifiers, while 51,132 double-breath units (15 % of the sample) were reserved for final testing (Table 6). The dataset included annotations for both PVA types and alarm levels, with the distribution of each category remaining consistent across the training and test sets. Based on the full dataset, the “no asynchrony” (NO) label was the most common, accounting for 55.80 % of all double-breath units, followed by hIE at 34.93 %, DC at 13.71 %, and IE at 5.70 %. Regarding alarm levels, the majority of units were classified as non-urgent (level 0, 60.92 %), with progressively fewer instances for levels 1 (29.55 %), 2 (9.24 %), and 3 (0.29 %).

3.5. Performance of deep learning models to recognise patient-ventilator asynchronies

The performance of the deep learning classifier for predicting alarm levels is summarized in Table 7. The overall accuracy (Exact Match), which represents the proportion of instances where all predicted alarm levels exactly matched the true levels, was 0.838. The per-class

Table 4
Baseline characteristics and outcomes of patients with mechanical ventilation. Data represent the number of patients (percentage) for categorical variables. Continuous variables are reported as mean \pm standard deviation.

Characteristic	
<i>n</i>	1863
Age (yr)	58.4 ± 17.8
Weight (kg)	86.9 ± 21.7
Height (cm)	174.6 ± 9.5
BMI	28.5 ± 6.9
Sex (M/F)	1174/689 (63.0 %/37.0 %)
Length of Stay (days)	8.1 ± 10.0
ICU Mortality	221 (11.9 %)
Major diagnosis	
Respiratory Failure	448 (24.0 %)
Trauma	318 (17.1 %)
COVID-19	260 (14.0 %)
Postoperative	245 (13.2 %)
Stroke	199 (10.7 %)
Traumatic Brain Injury	117 (6.3 %)
Sepsis	95 (5.1 %)
Resuscitation	60 (3.2 %)
Poisoning	19 (1.0 %)
Epilepsy	17 (0.9 %)
Other	85 (4.6 %)

Table 5
Counts of segmented, clustered, and annotated double-breath units. The “Double-Breath Count in Clusters” column represents the total number of segmented units, while the “Number of Clusters” indicates how these units were grouped for annotation. The “Annotated Double-Breath Count” reflects the number of high-quality units retained after excluding recordings with artifacts or noise. Abbreviations used include: SPONT (Spontaneous mode), ASV (Adaptive Support Ventilation), P SIMV (Pressure Synchronized Intermittent Mandatory Ventilation), DuoPAP (Dual Positive Airway Pressure), SIMV (Synchronized Intermittent Mandatory Ventilation).

Ventilation mode	Recording Count	Double-Breath Count in Clusters	Number of Clusters	Annotated Double-Breath Count
SPONT	90,808	90,729	800	69,463
ASV	43,205	165,079	600	161,255
P SIMV	34,258	81,858	530	72,714
DuoPAP	14,767	36,385	400	31,095
SIMV	13,945	7229	160	6344
Total	196,983	381,280	2490	340,871

Table 6

Distribution of PVA types and alarm levels in the full and test datasets. The counts represent annotated double-breath units, and the percentage of PVAs does not necessarily add up to 100 % because more than one PVA can be present in the same double-breath unit. Here IE stands for ineffective triggering, DC for double cycling, hIE for high inspiratory effort, and NO for no asynchrony.

Category	Full dataset		Test dataset	
	Count	Percentage, %	Count	Percentage, %
PVA type				
IE	19,421	5.70	2894	5.66
DC	46,742	13.71	7046	13.78
hIE	119,056	34.93	17,799	34.81
NO	190,216	55.80	28,578	55.89
Alarm level				
0 (no urgency)	207,649	60.92	31,156	60.93
1	100,732	29.55	15,091	29.51
2	31,502	9.24	4700	9.19
3 (most urgent)	988	0.29	185	0.36

accuracies for alarm levels 0, 1, 2, and 3 were 0.884, 0.848, 0.947, and 0.997, respectively. Sensitivity showed a decreasing trend with increasing alarm severity (0.915 for level 0, 0.733 for level 1, 0.681 for level 2, and 0.372 for level 3). Conversely, specificity increased with alarm severity (0.836 for level 0, 0.897 for level 1, 0.974 for level 2, and 0.999 for level 3). The classifier achieved AUC-ROC scores exceeding 0.950 for all alarm levels except level 1, which had an AUC-ROC of 0.915. Aggregated results showed weighted accuracy of 0.880, reflecting the model's performance adjusted for the frequency of each alarm level in the dataset, while the macro accuracy, which treats all alarm levels equally, was higher at 0.919, underscoring the classifier's balanced performance across all levels. The weighted AUC-ROC was 0.943, and the macro AUC-ROC reached 0.956, further confirming the model's robust classification capabilities across all alarm levels.

The performance of the deep learning classifier for predicting PVA types is detailed in Table 8. Among the four PVA types, the classifier achieved the highest accuracy for ineffective triggering (IE) at 0.960,

followed by double cycling (DC) at 0.918, no asynchrony (NO) at 0.888, and high inspiratory effort (hIE) at 0.878. Sensitivity varied across PVA types, with the highest sensitivity observed for NO (0.913), followed by hIE (0.745), DC (0.596), and IE (0.539). Specificity was consistently high across all PVA types, ranging from 0.840 for NO to 0.985 for IE. The classifier achieved AUC-ROC scores exceeding 0.940 for all PVA types, with the highest AUC-ROC observed for NO (0.954 [95 % CI: 0.953–0.956]) and the lowest for DC (0.940 [95 % CI: 0.937–0.942]). Aggregated metrics further highlight the model's performance with a weighted accuracy of 0.893, reflecting the overall performance adjusted for the frequency of each PVA type and a macro accuracy of 0.940, emphasizing the model's balanced performance across the types.

An additional analysis showed that the occurrence of hIE and the performance of the classifier in detecting hIE varied across ventilation modes. As shown in Supplementary Table 1 and Supplementary Table 2, hIE was most frequent in SPONT (57.1 %) and least in DuoPAP (22.4 %). Sensitivity of the prediction ranged from 0.574 (DuoPAP) to 0.876 (SIMV), while AUC-ROC values remained high across all modes, from 0.870 (SPONT) to 0.968 (SIMV).

Performance metrics for the training and validation sets of both models are provided in the Supplementary materials (Supplementary Table 3, 4, 5, 6) for comparison.

4. Discussion

4.1. Main results

In our study, we present the development and validation of the SmartAlert system, a machine learning-based tool designed to enhance the detection of patient-ventilator asynchronies (PVAs) in ICU settings. The system uses state-of-the-art signal processing techniques and deep learning models to analyze data from ventilator screens, capturing and classifying PVAs. SmartAlert operates solely on the graphical information displayed on ventilator screens, making it adaptable to a wide range of lung ventilator models with only minimal adjustments to the

Table 7

Performance of the deep learning classifier to predict alarm levels. Metrics include accuracy, sensitivity, specificity, AUC-ROC with 95 % confidence intervals (in square brackets), F1 score, and mean absolute error (MAE). These metrics are evaluated both per alarm level (0, 1, 2, and 3) and as aggregated results. Aggregated metrics are reported as weighted averages (weighted by the frequency of each alarm level) and as macro averages (unweighted mean across classes). Overall Accuracy (Exact Match) represents the proportion of instances where all predicted alarm levels exactly match the true levels.

Alarm level	Accuracy	Sensitivity	Specificity	AUC-ROC	F1 score	MAE
Per-Class Metrics						
0	0.884	0.915	0.836	0.952 [0.951–0.954]	0.906	0.092
1	0.848	0.733	0.897	0.915 [0.913–0.918]	0.741	0.267
2	0.947	0.681	0.974	0.967 [0.965–0.969]	0.703	0.366
3	0.997	0.372	0.999	0.990 [0.987–0.993]	0.449	1.014
Aggregated Metrics						
Overall Accuracy (Exact Match)	0.838					0.171
Overall (Weighted)	0.880	0.838	0.867	0.943 [0.941–0.945]	0.837	0.171
Overall (Macro)	0.919	0.675	0.926	0.956 [0.955–0.957]	0.700	0.435

Table 8

Performance of the deep learning classifier to predict PVA types. Metrics include accuracy, sensitivity, specificity, AUC-ROC with 95 % confidence intervals (in square brackets), and F1 score, evaluated per PVA type (IE, DC, hIE, NO) and as aggregated results. Aggregated metrics are reported as weighted averages (weighted by the frequency of each PVA type) and macro averages (unweighted mean across types). IE stands for ineffective triggering, DC for double cycling, hIE for high inspiratory effort, and NO for no asynchrony.

PVA type	Accuracy	Sensitivity	Specificity	AUC-ROC	F1 score
Per-Class metrics					
IE	0.960	0.539	0.985	0.941 [0.937–0.945]	0.605
DC	0.918	0.596	0.969	0.940 [0.937–0.942]	0.667
hIE	0.878	0.745	0.948	0.952 [0.950–0.953]	0.809
NO	0.888	0.913	0.840	0.954 [0.953–0.956]	0.915
Aggregated Metrics					
Overall (Weighted)	0.893	0.808	0.894	0.951 [0.950–0.953]	0.839
Overall (Macro)	0.940	0.638	0.957	0.962 [0.961–0.963]	0.699

measuring system. This universality ensures that the system can be easily deployed across different clinical setups.

The SmartAlert system accurately identified various types of PVAs, closely matching expert clinician assessments and correctly triggering alarms tailored to the needs of intensivists. Addressing a multi-class classification problem, one that requires distinguishing between multiple categories simultaneously, is inherently more complex than binary decision-making. This added complexity often results in lower performance metrics, as documented in machine learning research [23]. Despite these challenges, SmartAlert achieved an overall accuracy of 83.8 % for alarm level predictions, with a weighted sensitivity, of 83.8 %, specificity of 86.7 %, and a weighted AUC-ROC of 0.943 [95 % CI: 0.941–0.945]. Its performance in detecting PVA matched methods using hysteretic lung mechanics (≥ 89.5 % sensitivity, ≥ 96.8 % specificity) and is comparable to Zhang et al.'s image-based approach (~ 90 % accuracy) [18,24]. Per-class analysis revealed particular strengths; for instance, the detection of alarm level 0 achieved an F1 score of 0.906 and an AUC-ROC of 0.952 [95 % CI: 0.951–0.954]. Although alarm level 3 was less prevalent, its high specificity (99.9 %) minimized false alarms, and the low mean absolute error (MAE) for lower alarm levels reflected the model's precision in ranking severity. Even when predictions deviated, the classifier consistently assigned high alarm levels to critical cases, effectively maintaining the distinction between critical and non-critical alarms.

Similarly, the classifier for PVA types performed robustly, with a weighted accuracy of 89.3 % and AUC-ROC of 0.951 [95 % CI: 0.950–0.953]. It demonstrated strong detection of ineffective effort (IE), double cycling (DC), and high inspiratory effort (hIE), achieving specificities exceeding 94 %. For cases labeled as no asynchrony (NO), the classifier achieved the highest sensitivity (91.3 %) and an F1 score of 0.915. These results confirm the model's ability to capture clinically relevant patterns while maintaining a low false-positive rate, an essential criterion for deployment in the ICU.

The detection of high inspiratory effort, in particular, deserves further consideration, as this asynchrony type includes an inherently subjective component and may manifest differently depending on the ventilation mode. To address this, we performed a mode-stratified analysis, which revealed that while the frequency of hIE varied across modes, the model maintained high predictive power (AUC-ROC) throughout. This finding supports the robustness of the classifier despite physiological and detection-related variability.

Interpreting the classifier's performance requires consideration of several factors. Although the training and testing datasets were curated, they still reflect real-world ICU variability, which inevitably introduces some noise. Training on such data enhances generalizability of the model, as system trained on perfectly cleaned datasets may struggle with actual ICU inputs. Given the system operates in real time, flawless classification of every double-breath unit is not the ultimate goal. Instead, SmartAlert aggregates classification results over time using sequence-based decision mechanisms to improve reliability. Our group is currently developing this aggregation decision system, and future studies will include its performance analysis to further refine its integration into clinical workflows.

4.2. Advantages and potential clinical application

The SmartAlert system offers several distinct advantages over existing approaches. One key advantage is its straightforward technical solution, which exports real-time ventilator screen data directly to an internet browser, clinical information software, or a mobile application. This design builds on clinicians' natural reliance on pattern recognition to diagnose PVAs, and by incorporating a reliable machine learning (ML)-based recognition system. SmartAlert has the potential to significantly improve the detection and management of PVAs in busy ICU environments. While its effectiveness in actual clinical settings still awaits confirmation through randomized controlled trials (RCTs), the

system's innovative design already addresses critical limitations of current solutions.

A notable strength of SmartAlert is its universality. Unlike studies that depend on proprietary data outputs from ventilator machines [18, 20], SmartAlert utilizes video output. This approach makes the system highly adaptable, allowing it to accommodate various screen layouts with minimal customization and ensuring compatibility across a wide range of ventilators.

We have also validated the quality of the extracted waveform signals by comparing them to those obtained via the RS-232 output of the ventilator using a manufacturer-provided data logger. This comparison showed a high level of agreement in shape, timing, and amplitude, which supports the technical validity of our screen-based extraction approach and its comparability to traditional signal-based methods.

Furthermore, SmartAlert system leverages a deep learning approach to analyze minimally pre-processed pressure and flow waveforms. This contrasts sharply with traditional rule-based detection algorithms, which rely on computing engineered features and applying fixed thresholds [14,16,17,25]. By eliminating the need for handcrafted features, SmartAlert learns to identify relevant features and their nonlinear combinations directly from raw data. This self-learning capability enhances both the robustness and adaptability of the system compared to standard machine learning techniques that depend on predefined features [26].

The SmartAlert system also stands out for its efficiency and practicality in real-time applications. The system uses a single, relatively simple deep-learning architecture to predict multiple types of PVAs simultaneously. In our study, three types of PVAs were accurately predicted in real time with minimal pre-processing, a significant improvement over approaches such as Pan et al. [20], which required computationally intensive transformations from 1D to 2D data and complex architectures. Similarly, Zhang et al. [18] employed LSTM networks but required separate models for each PVA subtype, adding to the complexity of their solution. In contrast, SmartAlert not only integrates predictions for multiple PVA subtypes within a single model but is also capable of simultaneous operation across multiple ventilators, ensuring accurate, real-time predictions in complex ICU environments. By addressing these practical challenges, SmartAlert sets a new standard for real-time, multi-class PVA detection and management in clinical settings.

A key innovation of the SmartAlert system is its novel alarm functionality. By incorporating a severity scale that assesses the patient's condition regarding problematic ventilator settings, the system can assign tailored alarm levels. This feature not only alerts physicians to immediate dangers but also provides detailed information about the detected type of PVA, thereby enhancing clinical decision-making.

Finally, SmartAlert addresses the critical issue of alarm fatigue. Frequent alarms in ICU settings can desensitize staff and compromise patient care. By providing targeted, clinically relevant alerts that focus on significant asynchronies, SmartAlert reduces unnecessary notifications and supports optimized ventilator settings. This targeted alerting approach builds on the foundation of effective asynchrony management strategies demonstrated in previous studies [14,27] and sets a new standard for real-time, multi-class PVA detection and management in clinical practice.

4.3. Limitations

The development and validation of the SmartAlert system come with several important limitations that must be considered when interpreting the study's results. First, the reliance on a controlled dataset from a single type of ventilator, may limit the generalizability of our findings to other devices or clinical settings. Second, the gold standard for validation was based on consensus among two trained intensivists (or three in cases of disagreement), which, despite being robust, introduces potential human error and subjective bias. Additionally, our study focused on

specific patient-ventilator asynchronies, such as ineffective triggering, double cycling, and high inspiratory effort, which may exclude other, less prevalent yet clinically significant asynchronies that occur in real-world settings. The clustering process used to detect similar patterns might also have overlooked some clinically relevant signals that could have escaped detection by human evaluators. Next, variability in waveform display zoom between screen captures may also affect the effective sampling frequency of the extracted signal, although the system adapts to each screen individually by detecting axis scales from tick marks and labels. Another consideration is that approximately 11.7 % of videos were excluded during preprocessing due to unrecognizable screen content. These cases typically involved unsupported layouts or modes, HDMI capture issues (e.g., cable disconnection or corrupted signal), or failures in OCR due to compression artifacts. In deployment, where the system operates with continuous screen capture and the underlying physiological processes evolve relatively slowly, the impact of such exclusions is minimal. Finally, this study did not evaluate the real-time application and responsiveness of the SmartAlert system in a live ICU environment.

4.4. Future directions

There is an urgent need to investigate the impact of the SmartAlert system on clinical, patient-centered outcomes. To address this, we have designed a randomized controlled trial (RCT) to assess how integrating SmartAlert into a mobile app influences the safety and duration of mechanical ventilation (registered at Open Science Framework: <https://osf.io/s25e6/>). This study aims to evaluate how real-time alerts generated by SmartAlert affect clinical decision-making and patient outcomes in ICU settings. By examining the system's ability to reduce patient-ventilator asynchrony and potentially enhance ventilatory safety, the trial seeks to provide concrete evidence of its benefits and operational effectiveness within the complex ICU environment.

5. Conclusion

In conclusion, we have developed a tool capable of detecting patient-ventilator asynchronies in real time with high accuracy, designed to alert clinicians to this potentially dangerous condition that standard ventilator alarms often fail to identify. By using machine learning to detect and notify healthcare professionals of critical asynchronies, the SmartAlert system has the potential to significantly reduce the risk of VILI and improve outcomes for mechanically ventilated patients.

Statement of ethical approval

This study was approved by the Multicentric Ethics Committee at the University Hospital Královské Vinohrady (approval number: EK-VP/49/0/2021). The Committee reviewed the study, approved retrospective waveform data use, and waived the need for informed consent.

Availability of data and material

The datasets generated and/or analyzed during the current study are not publicly available due to the inclusion of potentially sensitive patient information and because they constitute the private intellectual property of the authoring institutions. However, the data may be made available from the corresponding author upon reasonable request, subject to institutional and ethical approvals.

Funding

This research was supported by Ministry of Health of the Czech Republic, grant nr. NU22-06-00625, institutional resources of Czech Technical University in Prague, and a research project RCI (reg. no. CZ.02.1.01/0.0/0.0/16_019/0000765) supported by EU.

Computational resources were provided by the e-INFRA CZ project (ID:90254), supported by the Ministry of Education, Youth and Sports of the Czech Republic.

CRediT authorship contribution statement

Jaroslav Pažout: Writing – original draft, Investigation, Data curation. **Milan Němý:** Writing – original draft, Software, Methodology, Investigation, Formal analysis. **Jakub Mikeš:** Writing – review & editing, Investigation, Data curation. **Jan Jirman:** Software, Resources, Conceptualization. **Jan Kubr:** Software, Resources. **Eliška Niebauerová:** Investigation, Data curation. **Miroslav Macík:** Writing – review & editing, Writing – original draft, Software, Resources, Conceptualization. **Michal Pech:** Software, Resources. **Michal Štajnrt:** Investigation, Data curation. **Jakub Vaněk:** Software, Resources, Conceptualization. **Petr Waldauf:** Writing – review & editing, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Václav Zvoníček:** Writing – review & editing, Methodology, Data curation, Conceptualization. **Lenka Vyslouzilová:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization. **Robert Babuška:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization. **František Duška:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no competing interests.

Acknowledgements

Not applicable.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.cmpb.2025.108927](https://doi.org/10.1016/j.cmpb.2025.108927).

References

- [1] A.S. Slutsky, Ventilator-induced lung injury: from barotrauma to biotrauma, *Respir. Care* (2005) 50.
- [2] D. Dreyfuss, G. Saumon, Ventilator-induced lung injury: lessons from experimental studies, *Am. J. Respir. Crit. Care Med.* (1998) 157.
- [3] Ventilation with lower tidal volumes as compared with traditional tidal volumes for acute lung injury and the acute Respiratory distress syndrome, *New England J. Med.* 342 (2000), <https://doi.org/10.1056/nejm200005043421801>.
- [4] H. Yasuda, M. Sanui, T. Nishimura, et al., Optimal upper limits of plateau pressure for patients with acute Respiratory distress syndrome during the first seven days: a meta-regression analysis, *J. Clin. Med. Res.* 13 (2021), <https://doi.org/10.14740/jocmr4390>.
- [5] L. Gattinoni, E. Carlesso, P. Cadringer, et al., Physical and biological triggers of ventilator-induced lung injury and its prevention, In: *Euro. Resp. J., Supplement* (2003).
- [6] L. Vignaux, F. Vargas, J. Roesler, et al., Patient-ventilator asynchrony during non-invasive ventilation for acute respiratory failure: a multicenter study, *Intensive Care Med.* 35 (2009), <https://doi.org/10.1007/s00134-009-1416-5>.
- [7] S. Sendelbach, M. Funk, Alarm fatigue: a patient safety concern, *AACN. Adv. Crit. Care* 24 (2013), <https://doi.org/10.1097/NCI.0b013e3182a903f9>.
- [8] M. Cvach, Monitor alarm fatigue: an integrative review, *Biomed. Instrum. Technol.* (2012) 46.
- [9] D. Colombo, G. Cammarota, M. Alemani, et al., Efficacy of ventilator waveforms observation in detecting patient-ventilator asynchrony, *Crit. Care Med.* 39 (2011), <https://doi.org/10.1097/CCM.0b013e318225753c>.
- [10] A.W. Thille, P. Rodriguez, B. Cabello, et al., Patient-ventilator asynchrony during assisted mechanical ventilation, *Intensive Care Med.* 32 (2006), <https://doi.org/10.1007/s00134-006-0301-8>.
- [11] C. de Haro, A. Ochagavia, J. López-Aguilar, et al., Patient-ventilator asynchronies during mechanical ventilation: current knowledge and research priorities, *Intensive Care Med. Exp.* (2019) 7.

- [12] H. Yonis, F. Gobert, R. Taponnier, C. Guérin, Reverse triggering in a patient with ARDS, *Intensive Care Med.* 41 (2015), <https://doi.org/10.1007/s00134-015-3702-8>.
- [13] J.R. Beitler, S.A. Sands, S.H. Loring, et al., Quantifying unintended exposure to high tidal volumes from breath stacking dyssynchrony in ARDS: the BREATHE criteria, *Intensive Care Med.* 42 (2016), <https://doi.org/10.1007/s00134-016-4423-3>.
- [14] L. Blanch, A. Villagra, B. Sales, et al., Asynchronies during mechanical ventilation are associated with mortality, *Intensive Care Med.* 41 (2015) 633–641, <https://doi.org/10.1007/s00134-015-3692-6>.
- [15] M. Dres, N. Rittayamai, L. Brochard, Monitoring patient-ventilator asynchrony, *Curr. Opin. Crit. Care* (2016) 22.
- [16] Q. Mulqueeny, P. Ceriana, A. Carlucci, et al., Automatic detection of ineffective triggering and double triggering during mechanical ventilation, *Intensive Care Med.* 33 (2007) 2014–2018, <https://doi.org/10.1007/s00134-007-0767-z>.
- [17] L. Blanch, B. Sales, J. Montanya, et al., Validation of the Better Care® system to detect ineffective efforts during expiration in mechanically ventilated patients: a pilot study, *Intensive Care Med.* 38 (2012) 772–780, <https://doi.org/10.1007/s00134-012-2493-4>.
- [18] L. Zhang, K. Mao, K. Duan, et al., Detection of patient-ventilator asynchrony from mechanical ventilation waveforms using a two-layer long short-term memory neural network, *Comput. Biol. Med.* 120 (2020), <https://doi.org/10.1016/j.combiomed.2020.103721>.
- [19] Q. Pan, L. Zhang, M. Jia, et al., An interpretable 1D convolutional neural network for detecting patient-ventilator asynchrony in mechanical ventilation, *Comput. Methods Prog. Biomed.* 204 (2021), <https://doi.org/10.1016/j.cmpb.2021.106057>.
- [20] Q. Pan, M. Jia, Q. Liu, et al., Identifying patient–ventilator asynchrony on a small dataset using image-based transfer learning, *Sensors* 21 (2021), <https://doi.org/10.3390/s21124149>.
- [21] E. Mireles-Cabodevila, M.T. Siuba, R.L. Chatburn, A taxonomy for patient-ventilator interactions and a method to read ventilator waveforms, *Respir. Care* 67 (2022) 129–148, <https://doi.org/10.4187/RESPCARE.09316>.
- [22] A. Sardá-Espinoza, Time-series clustering in R using the dtwclust, *Package. RJ* 11 (2019) 22, <https://doi.org/10.32614/RJ-2019-023>.
- [23] Moral P Del, Nowaczyk S, S. Pashami, Why is multiclass classification hard? *IEEE Access* 10 (2022) 80448–80462, <https://doi.org/10.1109/ACCESS.2022.3192514>.
- [24] Y. Chen, K. Zhang, C. Zhou, et al., Automated evaluation of typical patient–ventilator asynchronies based on lung hysteretic responses, *Biomed. Eng. Online* 22 (2023), <https://doi.org/10.1186/s12938-023-01165-0>.
- [25] J.Y. Adams, M.K. Lieng, B.T. Kuhn, et al., Development and validation of a multi-algorithm analytic platform to detect off-target mechanical ventilation, *Sci. Rep.* 7 (2017), <https://doi.org/10.1038/s41598-017-15052-x>.
- [26] P.D. Sottile, D. Albers, C. Higgins, et al., The association between ventilator dyssynchrony, delivered tidal volume, and sedation using a novel automated ventilator dyssynchrony detection algorithm*, *Crit. Care Med.* 46 (2018), <https://doi.org/10.1097/CCM.0000000000002849>.
- [27] M. Kyo, T. Shimatani, K. Hosokawa, et al., Patient–ventilator asynchrony, impact on clinical outcomes and effectiveness of interventions: a systematic review and meta-analysis, *J. Intensive Care* 9 (2021), <https://doi.org/10.1186/s40560-021-00565-5>.