Identifying low-dimensional trajectories of mechanically-ventilated patient systems: Empirical phenotypes of joint patient+care processes to enhance temporal analysis in ARDS research

J.N. Stroh, Peter D. Sottile, Yanran Wang, Bradford J. Smith, Tellen D. Bennett, Marc Moss, David J. Albers

June 10, 2024

Abstract

Refined management of mechanically ventilation is an obvious target for improving patient outcomes, but is impeded by the nature of data for study and hypothesis generation. The connections between clinical outcomes and temporal development of iatrogenic injuries current lung-protective ventilator settings remain poorly understood. Analysis of lung-ventilator system (LVS) evolution at relevant timescales is frustrated by data volume and multiple sources of heterogeneity. This work motivates, presents, and validates a computational pipeline for resolving LVS systems into the joint evolution of data-conditioned model parameters and ventilator information. Applied to individuals, the workflow yields a concise low-dimensional representation of LVS behavior expressed in phenotypic breath waveforms suitable for analysis. The effectiveness of this approach is demonstrated through application to multi-day observational series of 35 patients. Individual patient analyses reveal multiple types of patient-oriented dynamics and breath behavior to expose the complexity of LVS evolution; less than 10% of phenotype changes related to ventilator settings changes. Dynamics are shown to including both stable and unstable phenotype transitions as well as both discrete and continuous changes unrelated to ventilator settings. At a cohort scale, 721 phenotypes constructed from individual data are condensed into a set of 16 groups that empirically organize around certain settings (positive end-expository pressure and ventilator mode) and structurally similar pressure-volume loop characterizations. Individual and cohort scale phenotypes, which may be refined by hypothesis-specific constructions, provide a common framework for ongoing temporal analysis and investigation of LVS dynamics.

keywords: pulmonary ventilation; patient-ventilator asynchrony; ventilator-induced lung injury; respiratory distress syndrome, patient-specific modeling; knowledge representation

1. Introduction

Critical care often employs mechanical ventilation (MV) to manage patients with disorders such as acute respiratory distress syndrome (ARDS), which is characterized by inflammation and pulmonary edema. Modern respiratory care protocols and technologies [1] emphasize lung-protective strategies [2] to minimize deleterious effects like ventilator-induced lung injury (VILI,[3]). Such strategies rely an understanding of lung physiology to inform MV settings such positive end-expiratory pressure (PEEP), tidal volume, and driving pressure [4, 5, 6]. Technological advances in MV have not eliminated ventilator dyssynchrony (VD), a mismatch in patient-ventilator delivery and respiratory effort timing. VD may play a role in the development and propagation of VILI, a known contributor to mortality in ARDS patients [7]. Reduction in ARDS-related mortality has plateaued in recent decades [8] following a significant curtailment during the two decades prior [9]. Further reduction of mortality motivates the continued study of MV effects as VILI and VD likely support residual negative ARDS outcomes.

Clinical MV observables include airway pressure (p), volume (V), and flow timeseries that record the dynamic interaction between patient lungs and an engineered control apparatus. The underlying data generation process, the human lung-ventilator system (LVS), contains the key components relevant for studying the effects of MV over time from ventilator observations [10]. Non-ventilator aspects of patient care can also affect lung-ventilator dynamics and associated observables. The format (scalar, but millions of data points per patient per day) and multiple sources of patient- and care-specific heterogeneities hinders direct analysis of LVS waveforms (*ibid.*). Additionally, consideration of intra-breath-scale events over the duration of MV requires multi-scale analysis to detect signatures of injuries like VILI development from LVS data.

Notable previous works in machine learning (ML) applications [11, 12] used identify the frequency and occurrence of different ventilator dyssynchronies from LVS data. Features included breath properties such as peak inspiratory pressure, inspiratory-to-expiratory time ratio (I:E), etc. to coarsely characterize waveforms in parameters familiar to practitioners [13, 14]. In another vein of research, analysis of full MV waveform data with attention to patient-ventilator dyssynchrony resolution has focused on hybrid-modeling methods leveraging empirical parameter fitting [15, 16, 10, 17]. MV waveforms often violate mechanistic model assumptions; the hybrid schemes evade this limit by reducing the assumptions through universal model [10] or using high-fidelity behavior-specific models [15]. These data-informed modeling methods transformation waveform data into parametric representations, but differ in assumptions and method. However, LVS behavior over time has not yet been thoroughly investigated. This study presents a temporal framework for combining the parametric approaches together with unsupervised ML. The workflow is applied to MV data (waveform and ventilator documentation) to reveal the structure and complexity of breath evolution of individuals patients from heterogeneous LVS factors. A natural extension to cohort-scale analysis is also developed with an eye toward statistical or learning-based temporal analysis.

The framework presented below extends the analysis of LVS behavior from the breath level to the scale of hours-to-days while jointly considering the context of ventilator settings. It relies on an interpretable and unsupervised segmentation pipeline [18] that supports many methodological choices including how waveform data are represented. Implementing a context-free digitization [10] with limited assumptions about the data, *this work hypothesizes that respiratory behavior or other patient properties may be identified from joint LVS data by separating the influence of changes in MV*. Investigation proceeds by scrutinizing typological breath changes occurring independently of ventilator management within the context of the joint patient-ventilator system. Analysis of ARDS patient data within this perspective demonstrates compact descriptors of LVS evolution, broadly categorizes MV breaths, and proposes sources of heterogeneity needed to further develop the problem domain.

2. Method

The root approach involves analyzing LVS data, including waveforms and ventilator settings, through a computational pipeline that begins with model-based inference of waveform data. The method parametrizes waveform data and identifies patient-specific breath phenotypes similarities between joint waveform and ventilator state descriptors. The evolution of LVSs may be examined through phenotypes when co-labeled according to time.

2.1. Data

Mechanically-ventilated patient data were collected under the University of Colorado Multiple Institutional Review Board (COMIRB, protocol #18-1433). These data include airway pressure, volume, and ventilator settings for 36 patients, all of whom had ARDS diagnoses and substantial risk of VILI. Children, pregnant women, and age-censored elders, and the imprisoned were excluded. Esophageal pressures were recorded but are used in this work; collection imposed additional exclusion criterion (viz. esophageal fistula, variceal bleeding or banding, facial fracture, and recent gastric/esophageal surgery). Source patients include 14 women and 22 men with median[IQR] age 59[25] years; 72% are white, 35% of which identify as Hispanic or Latino. Table 1 summarizes clinical and demographic characteristics of patients. Data total 1.74 million breaths over 71.14 recording-days (median 1.97[1.56] days per patient) recorded at 32 millisecond sampling (31.25 Hz) from Hamilton G5 ventilators (`https://www.hamilton-medical.com`). Adaptive pressure ventilation-controlled and pressure-controlled mandatory ventilation modes (APVCMV and P-CMV, respectively) account for roughly 84% and 10% of breaths, respectively, with the remainder in standby and spontaneous modes. Ventilator management throughout employs the ARDSnet protocols [7].

Table 1: Tabular summary of the patient cohort and associated data. 'Monitored' and 'Recorded' durations denote the number of hours spanned by data and length continuous data contents, respectively. P:F ratio is the PaO_2/FiO_2 ratio at admission, AA = African-American, AI = American-Indian, AK = Alaska, NMB = Neuromuscular Blockade (paralytic)

Detail	*Count*	%	*Median*	*IQR*
Monitored (hrs)			47.0	37.2
Recorded (hrs)			43.1	40
Age (years)	36		58.5	24.5
Gender				
Female	14	38.9	54.5	25.0
Male	22	61.1	58.5	26.0
Race/Ethicity				
White	26	72.2		
Unknown/NA	5	13.9		
Black/AA	3	8.3		
AI or AK Native	1	2.8		
More than one race	1	2.8		
ARDS risk				
Pneumonia	12	33.3		
COVID	11	30.6		
Sepsis	6	16.7		
Other	3	8.3		
Pancreatitis	2	5.6		
Aspiration	2	5.6		
P:F ratio			135.9	81.0
Mortality	9	25.0		
NMB use	9	25.0		

Dyssynchrony labels. An existing supervised ML technique [11] identifies breath-wise VD to enrich LVS evolution context and provide comparison for newly calculated labels. Type-specific VD models each label breaths according to features characterizing dyssynchronous breaths (see *ibid.* SI). VD labels include normal (NL), reverse triggered (RT) with early (RTe) and middle (RTm) subtypes, early flow limited (eFL) with

intermediate (eFLi) and severe (eFLs) subtypes, double trigger (DT) with reverse- (DTr) and patient- (DTp) subtypes, and early vent termination (EVT); breath mechanics of these VDs are described in [12]. The present work ignores VD subtype labels and uses VD labels to identify likely VD occurrence during model time intervals.

2.2. Model-based waveform parametrization

A waveform parametrization [10] is adopted in this work for convenience. It may easily be substituted for spectrograms [19], data features [11], or other model-based waveform characterization [20]. Briefly, the model transforms waveform data y^{obs} on a continuous time window I into a parameter vector $\mathbf{a}$ under the assumption that PEEP and breath period θ are locally constant. A differential equation models the state variable y (pressure, volume, or another waveform variable of interest):

$$\frac{dy}{dt} + g \cdot (y(t) - y_0) = \varphi(t; \mathbf{a}, \theta) \tag{1}$$

where t is time, g is a smoothing parameter, y_0 is a reference state (such as PEEP when y is pressure, and zero when y is volume), and $\varphi(t)$ is a time-dependent function of local amplitude parameters $a \in \mathbb{R}^M$ and breath period θ. The model, Eq(1), defines a map $\mathbf{a} \mapsto y(t)$ that simulates a state trajectory y from a periodic step function φ modulated by parameters $\mathbf{a}$:

$$\varphi(t; \mathbf{a}, \theta) = \frac{1 - e^{-g\Delta t}}{g} \sum_{i=1}^{M} a_i \left[\!\left[(i-1) \leq \frac{\widehat{t}}{\Delta t} < i \right]\!\right]. \tag{2}$$

where $\widehat{t} := t (\mathrm{mod}\, \theta)$ is the local breath time, which is divided into M equal epochs of width $\Delta t = \theta / M$. The corresponding inverse problem [21] is solvable by data assimilation, mapping the observations to a distribution of parameters that reconstruct the data: $y^{obs} \mapsto \{\mathbf{a}\}$. Optimal parameter distributions for pressure and volume data are generated by applying an ensemble Kalman-like smoother [22]. A moderate resolution model ($M = 28$) is sequentially inferred for each 10-second block of data using 1.6 second windows with 0.8 second overlaps. Details, error analysis, and validation of the parametrization is found in [10].

2.3. Pipeline

The computational pipeline extracts low-dimensional representations of LVS data that effectively encode relevant features of both breath waveform data and the ventilator settings associated with them. The method (depicted in Figure 1) follows [18] using model-inferred parameter distributions to uncover latent similarities within the data. The four stages of application to LVS data focus on changing system representation, followed by a final interpretation.

1.) Waveform parametrization. On each 10-second interval, parameters distributions for continuous pressure (p) and volume (V) waveform data inferred for the model (§2.2) with a moderate resolution ($M = 28$). Specifically, an ensemble of $N_{ens} = 25$ solutions is optimized within a moving sub-window of length 1.6 seconds and 0.8 second overlap. The initial ensemble prior comprises parameter values $a_i = 0.5$, $i = 1..M$ perturbed by 10% white noise. Afterwards, the posterior parameter estimate ($\{\mathbf{a}_i\}\, i = 1..N_{ens}$ above) becomes the prior for next sub-window, with updates using only data not previously assimilated. For

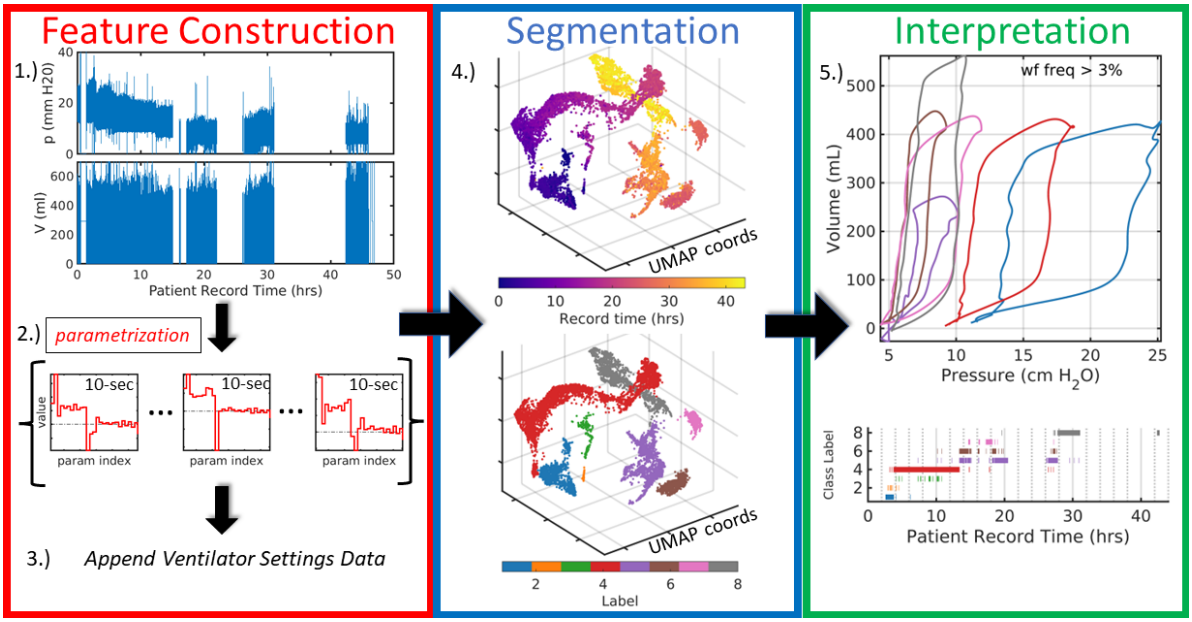

Figure 1: Broad pipeline organization. Raw data (1.) are digitally parametrized (2.) over short windows by the model (§2.2). Distributional parameter estimates are summarized and augmented with the contextual data of ventilator settings (3.) which include information such as ventilator operation mode, PEEP or other baseline pressure, flow and pressure triggers, and minimum mandatory breath rate. Feature vectors, defined by the augmented LVS descriptors, are dimensionally reduced to three dimensions using UMAP (4.) where they can be analyzed based on time ordering (top) and structural similarity via segmentation (bottom). Finally, in (5.), temporal evolution of the system is compactly encoded in the time-ordered LVS descriptor labels and their associated waveform characterizations in an interpretable and explainable way. The process transforms raw data (1.) into a more easily comprehensible form such as (5.).

example, the prior for the second sub-window (0.8–2.4 seconds) is conditioned on data from 0–1.6 seconds, and assimilates data from 1.6–2.4 seconds to initialize the window starting at 1.6 seconds. Parameters associated with late exhalation are not typically informed data until the third sub-window (1.6–3.2 seconds). Excluding these first three estimates, roughly 10 ensemble estimates define an empirical parameter density that captures the waveform data within the 10-second window.

2.) Parameter Distribution Summarization. The empirical distributions of the $2M$-dimensional parameters are reduced to vectors of statistical summaries. This is done to ease comparison of waveform behaviors at different points in time by applying similarity measures to distribution summaries. Descriptor components include mean, quartiles, variance, and mode as well as non-gaussian measures (skewness, kurtosis, and Kolmogorov-Smirnov distance) to capture bimodal or asymmetric parameters distributions characterizing non-stationary LVS behavior. Period and baseline pressure, assumed stationary during inference, are included in parameter summaries to permit accurate waveform reconstruction. The strategy reduces the temporal sampling rate (from 31.25Hz to 0.1Hz) by representing 10-second windows of waveform data statistically through data-informed model parameters. Reduction in the overall data volume is governed by summary window length (under weaker stationarity assumptions) and model resolution (M).

3.) Augmentation. Appending ventilator setting data to each statistical waveform parameter summary contextualizes them in the health-care process. Ventilator settings detail the mode of operation (volume control, pressure control, spontaneous, standby, etc.), targeted quantities (pressure or tidal volume) as

well as various machine settings (tidal volume, trigger thresholds, ramp time, mandatory minimum breath rate, etc). Some ventilator settings, such as PEEP and I:E ratio, are already represented in implicitly in summary descriptors and need not be explicitly included. Other available factors such as ventilator delivery power are not considered here but may be included in other applications. Ventilator mode is a nominal variable represented as set of binary variables using one-hot encoding. Settings summaries reflecting the most frequent breath-level record of each augment waveform descriptors to define LVS feature vectors in subsequent pipeline steps. Ventilator settings change infrequently, and summary errors are therefore rare among estimated intervals.

4.) Cluster Labeling. Segmentation labels groups of LVS descriptors based on content similarity and can be applied at individual patient or aggregated cohort levels. Descriptors are first dimensionally reduced to diminish the computational burden imposed by many large feature vectors and to assist in visual and analytical assessment of label assignment. This is followed by unsupervised segmentation ([23]) where appropriate groupings are identified for the joint LVS descriptors in reduced coordinates.

The Uniform Manifold Approximation and Projection (UMAP,[24, 25]) identifies a low-dimensional projection that preserves the local and global structure of high-dimensional joint LVS descriptors ($\sim$ 400-dimensional for $M = 28$; fixed parameters: neighborhood size 5, minimum distance 0.01, 3 output dimensions). Similarity structure is determined by the the Gower distance, which averages over range-normalized distances for continuous variables and binary difference in categorical variables (ventilator modes). Density-Based Spatial Clustering of Applications with Noise (DBSCAN, [26, 27]) then groups the similarity-organized LVS descriptors based on point densities of the UMAP coordinates. A brief grid-search over hyper-parameters (min. core point neighbors 4–12; neighborhood radius 1.5–5 by 0.5) identifies a grouping with minimum total distance between centroids. Flexible labeling sought to accommodate the feature variations that generally increase with the LVS record length.

Another dimensional reduction option, t-distributed Stochastic Neighborhood Embedding (tSNE,[28], produced similar LVS groupings with the same DBSCAN parameters ([29], SI B.2). Support vector clustering [30, 31] also yielded similar labels but required significantly longer computation time. The k-means and k-medoids [32] methods were considered for efficiency but struggled with the non-convex groupings that typically emerged from the UMAP projection of LVS descriptors.

5.) Phenotype interpretation. Labeled descriptors correspond one-to-one with 10-second data windows and associated information, including the observed waveform data, waveform parameters, ventilator settings, and relative measure of group-wise similarity. Identification and further analysis of group characteristics may be leveraged from these details. For example, applying the model (Eq.(1)) to e.g., median model parameters associated with a label yields pressure or volume waveforms characterizing the central behavior of each group. From this perspective, group labels distinguish phenotypes of the LVS observables for an individual patient.

2.4. Phenotypes and Characterizations of LVS data

The phenotyping pipeline identifies windows with empirically similar LVS states, organizing data into discrete categories. The co-evolution of the patient-ventilator system is captured in the temporal progression through these categories and may be analyzed over longer timescales. A stated objective is to reveal changes originating from the patient-side of the system with no corresponding ventilator changes. Such changes

suggest the presence of factors that influence LVS trajectory such as changes in patient expectation and breathing pattern (*e.g.*, patient effort, respiratory drive), lung mechanical function (*e.g.*, VILI progression or recovery from ARDS), or another aspect of physiology.

Pipeline experiments are performed individually on 35 ARDS patient data records. Experiments reducing LVS evolution to categories posits that LVS behavior includes patient-side changes that are detectable from waveform data. The objective of each experiment is to assess whether this is true and whether it is potentially representable in low-dimensional categories. Phenotype evolution is presented in the context of ventilator settings and in relation to classified VD ([11]). Additionally, pressure-volume (pV) characterizations, computed from model parameters nearest to the phenotype center (*viz.* median), provide a familiar synopsis of associated waveform data for each window represented in the data. Such visualizations intend to summarize key features and notable changes defining the LVS trajectory. Subsequent analysis and discussions employ principal component analysis (PCA), an empirical signal factorization based on variance minimization [33, 34]. This tool reveals the degree of LVS variance occurring under during stationarity to investigate non-ventilator temporal changes not identified by segmentation.

2.4.1. Cohort-scale phenotyping

Direct application of the individual pipeline to cohort data is a computationally expensive problem due to the data volume ($O(10^6)$ 10-second intervals of continuous multivariate variables), A simple alternative is to develop cohort-scale meta-labels for the population of individual phenotypes. However, appropriate scaling of volume waveforms is necessary to ensure adequate mixing of patients in feature space, as tidal volume value depends patient anthropometry. Treating both waveform components equally, pressure waveforms are standardized by zeroing on PEEP or baseline pressure and scaling by driving (maximum-minus-baseline) pressure within each window. Feature vectors for cohort clustering are individual phenotype statistics of: baseline pressure, driving pressure, scaled tidal volume, estimated parameters of normalized waveform data, and associated ventilator settings. Segmentation and label assignment proceeds by identically using UMAP-DBSCAN as in the individual case, albeit with different hyper-parameter values (UMAP: neighborhood size 12, minimum distance 1, euclidean metric; DBSCAN: epsilon 2.7, min points 5).

3. Results

The clinical data associated with ARDS patients (Table 1) is an important and practical use-case because such patients may be prone LVS changes related to ventilator dyssynchronies and VILI. This section reports the results of experiments applying the pipeline to individual ARDS patient data records (§3.1) and the assembly of cohort-scale phenotypes (§3.2). Within individual experiments, the temporal structure of LVS data labels is examined for consistency and resolution. Phenotypes aggregated across the cohort produce generalized LVS descriptor characterizations. Sequences of dyssynchrony labels [11] provide additional comparative context for exploring MV states of ventilator settings and waveform characteristics.

3.1. Patient-level Phenotyping

LVS patient data are identified with 20[14] (median[IRQ]) individual phenotypes, totaling 721 across the cohort. Approximately half of these labels correspond to infrequent behaviors each associated with less than 1% of a given patient record. A median of 8[6.5] core clusters each representing more than 3% of the data

account for the remainder of the data of each patient. Reducing label specificity through UMAP-DBSCAN hyper-parameters can eliminate low-occurrence groups, but a high degree of specificity is needed to resolve feature heterogeneity that increases with the total duration of the patient record. As record durations span 0.7–92 hrs (median[IQR] 47[37] hrs), the over-segmentation of shorter LVS records supports the resolution needed for longer records without a more robust optimization of individual segmentation.

Individual phenotype labels capture essential changes in ventilator settings and capture unrelated changes. There is high correspondence between changes in ventilator settings and persistent changes (lasting longer than 30 second) in individual patient phenotype labels (SI Table A.3). Changes in settings are typically (mean($s2l$) $> 60\%$) reflected in label changes, with $\sim$92% of changes in PEEP, MV mode. and V_T inducing label changes. The former assessment is biased by few settings changes in some patients and by counting changes with likely no direct effect on discrete breath behavior (e.g., trigger sensitivity or mandatory breath rates). Label-to-settings change coherence ($l2s$) is considerably lower; less than 10% of label changes are associated with ventilator settings changes. While obviously impacted by the much larger number of changes in labels than settings and the over-specificity discussed above, individual breath phenotypes include important changes LVS behavior, which are broader than changes ventilator properties.

Figures 2 and 3a–d visualize particular aspects of the low-dimensional time-ordered pipeline output for two patients (149 and 114 of Table1, respectively). Their cases are typical of experiments in record length ($\sim$24 hours), number of ventilator settings changes, and number of identified phenotypic breaths. The complexity and heterogeneity of joint patient-care data preclude in this work; SI B provides additional examples.

Example 1:. Figure 2 of patient 149 illustrates the trajectory of their LVS is driven by a progression of PEEP reduction and ventilator mode changes from APVCPM to spontaneous breath support for several hours. Changes in these settings, along with tidal volume, account for primary drivers of LVS behavior in nearly all experiments. Externally labeled VD types show little identified dyssynchrony as most breaths are identified as normal. However, there is also heterogeneous behavior indicated by labels (b) during the period from 4 to 12 hours under stationary ventilator settings. Here, LVS state vacillate between labels #1 and #3 with notably distinct pV characterization (c) during this period. Based on analysis of similar behavior in other experiments, irregularity of delivered tidal volume by the APVCMV mode in response to previous breath pressure is a likely explanation. Parsing LVS evolution is obviously burdened by system heterogeneity even within an individual.

Example 2:. Figure 3 illustrates an analysis of patient #114 whose LVS undergoes multiple changes over a 24-hour data period. The portion during 7–14.5 hours is dominated by normal breaths that spans two labels with similar characterizations as pV loops but differ in mean respiratory rate. The difference is minor (the mean difference is less than 20 milliseconds), although this affects model parameter and could combined via posterior analysis, with small DBSCAN hyper-parameter changes, or coarser period binning. Changes in flow trigger settings occur around 3 hours and reduce the occurrence of eFL near the star of the record, associated with caving in pV loops (label 1, dark blue). Dyssynchronies return when the flow trigger is returned to its initial value, near 15 hours. PEEP and tidal volume targets are also adjusted several times. Brief ventilator changes in ventilator mode around 20 and 23 hours allow spontaneous breathing which have a profoundly different pV characterizations (label 12,tan). The interim period (20.5–22.5 hours) consists of

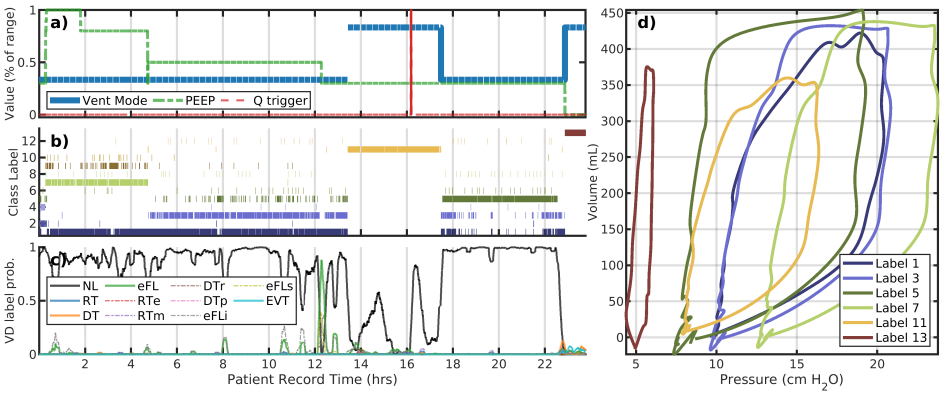

Figure 2: LVS evolution of patient 149. Panels a–c correspond to changes in ventilator settings, segmentation labels, and externally identified VD type, respectively. The horizontal axis for these panels is the patient record time in hours. The panel (d) shows the model image of segmented data median parameters, which characterize the pV loops of breaths with that label (shown with the same color). Evolution of the LVS can be parsed pictorially from these figures. The LVS evolution of patient 149 label#1 is discontinuous in time and occurs under different PEEP values suggesting waveform shapes vary only in baseline pressure. There is a lot of waveform variation present within the largely VD-less evolution, and significant changes in non-ventilator aspects of the LVS. Settings changes (a) are relative values to indicate change occurrences of multiple unlike values.

262 primarily normal breaths (label 13, brown) under the default pressure-control mode.

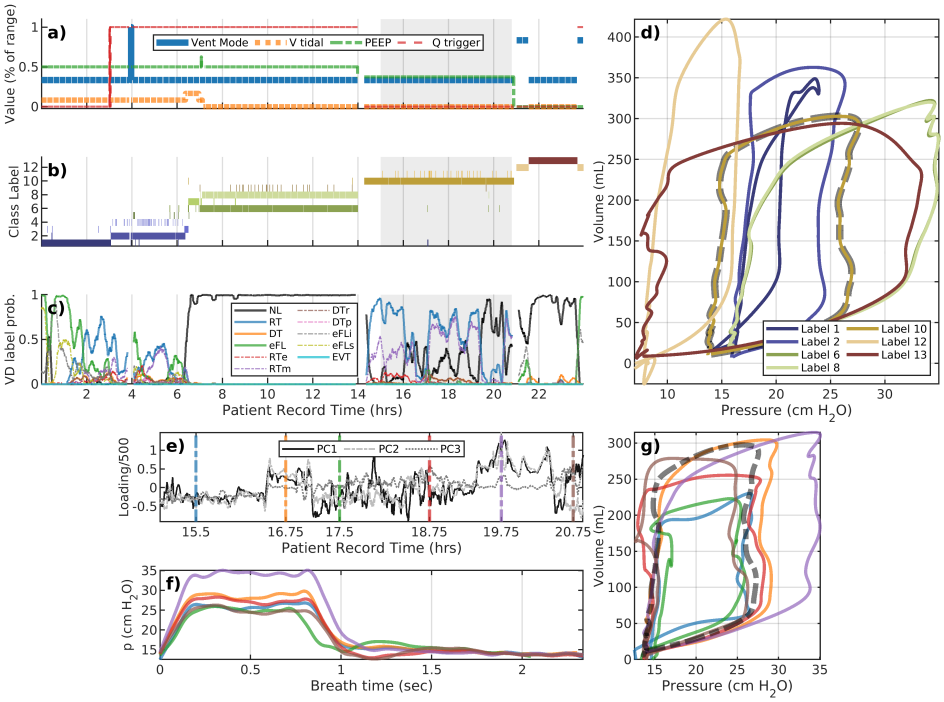

Figure 3: A representative example: patient #114. The upper plot layout is the same as the previous figure. The lower plot examines the variability, and a shortcoming of low resolution segmentation to capture changes that may be highly diverse at a local level. The mean– the dashed black line – coincides with the golden pV loop (cluster #10) in the upper plot. The many distinct breath subtypes identified are more similar than to other main types in the upper plot; as a result, they are grouped together at this choice of hyper-parameters.

Intra-label variability: A closer look at Label 10 of patient #114. Breath phenotype analysis of patient #114 (Fig.3) indicates no ventilator setting changes during the record interval 15–21 hours. Although one phenotypic breath dominates this period (d, dashed outline), various ML-labeled dyssynchronies (c) intimate more variability. Principal components during this interval (e) reveal structural waveform changes (f,g) that are not clearly identified as sub-phenotypes. While pressure characterizations (f) suggest the differences are largely attributed to pressure plateau pressure, full characterization also indicates ∼35% variability in tidal volume (g) as well. The *continuous* LVS evolution through these subtypes – and their comparative differences to the other types – leads to their collective identity. SIB.8 demonstrates a case where intra-label variability may be discretely resolved.

3.2. Cohort-scale phenotyping

Secondary segmentation (§2.4.1) of 721 individually-identified LVS phenotypes generates 16 groups of systemic behaviors. Figure 4 presents the coordination of labels and statistical summaries of data properties in dimensionally reduced form. Although there remains inherent variability, label-partitioned data have consistent properties and ventilator settings. Importantly, groups mix patients (b) while separating PEEP (c), with exceptions for specific, rare ventilation modes (d) that include few patients. Table 2 and Fig 4 quantitatively validate labeling of original data in the general settings. Specifically, labels consistently align with structured properties of the LVS data. Fig 5 shows the associated non-dimensional waveform characterizations; PEEP, tidal volume, and peak pressure features are used to normalize these data across patients.

Granularity of cohort meta-characterization depends on UMAP-DBSCAN hyper-parameters. Although UMAP representation was robust, cohort labeling was sensitive to the neighborhood size (SI C) due to the relatively small population of phenotypes. Selected parameters aimed to maximize the number of phenotypes easily communicating waveform characterizations in an array of figures; the results are qualitatively similar for nearby parameters. Table 2 summarizes the occurrence and properties the 16 cohort phenotypes.

3.3. Synthesis

Empirical segmentation analysis of patient+care data from MV patients indicates that changes in PEEP and ventilation mode are define the primary organization of group identities, followed by changes in tidal volume and those of non-ventilator origin (patient behavior or unattributed factors). Changes in these settings, along with multiple types of observed intra-patient variability, reveal that joint consideration of both patient and care processes is needed to understand the evolution of LVS systems. Analyzing patient state through breath data, especially for VILI detection and to track ARDS progression, requires considering ventilator settings. LVS behavior is often variable during periods of ventilator stationarity. Variability is sometimes reflected in label changes (Fig2b, hours 5–12, 17–18, 22–23; Fig3b, hours 3–14). When it is not, explanations include insufficient label granularity to resolve identifiable subgroups (*e.g.*, SI Figs B.6e–g and B.8) and/or continuous changes in breath behavior (*e.g.*, Fig1, step 4, label #4 in red; Fig 3e–g). Nevertheless, the space of joint breaths is greatly reduced through expression in the empirical phenotypes, with labeled data reflecting a hierarchical organization based on key vent settings (PEEP and mode) followed by waveform properties. Additionally, cohort scale analysis is also possible by segmenting the batch of individual phenotypes. Such a categorization provides a coarse -but scalable and unified- basis for analyzing the evolution of LVSs in terms of their consistent statistical properties.

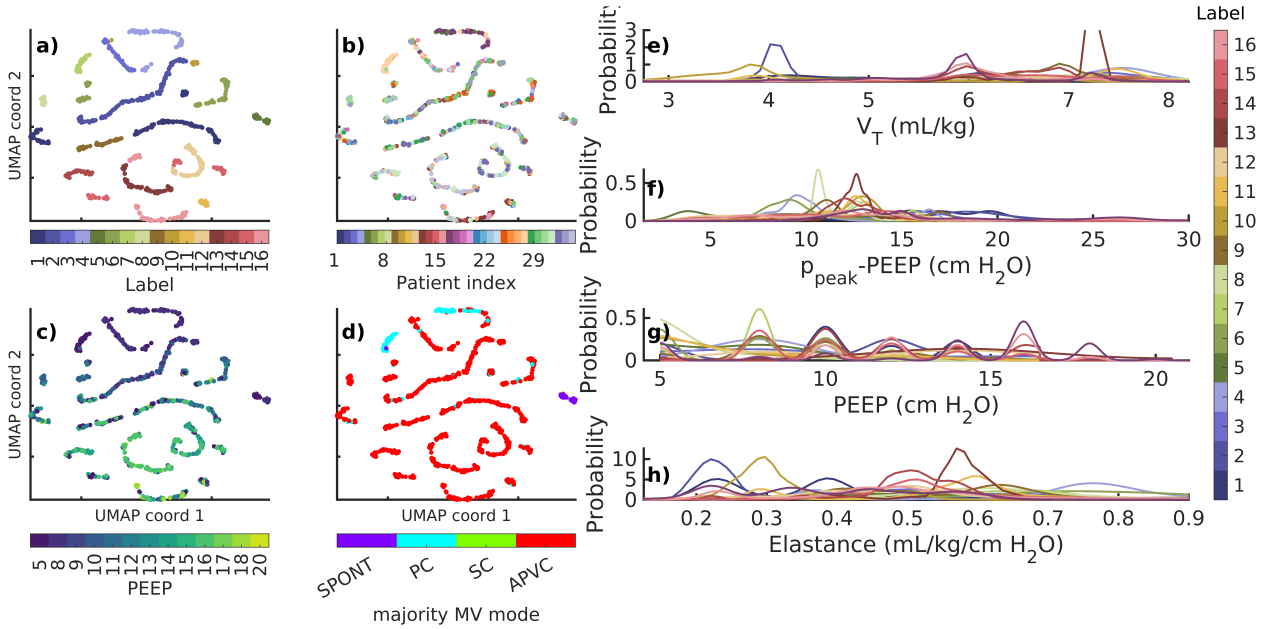

Figure 4: Membership and data properties associated with 721 phenotypes. Points in panels (a–d) correspond to individual phenotypes in UMAP coordinates. Labels (a) mix patients (b) while defining empirical partitions of other factors of patient data (c–h). Groupings separate PEEP (c,g) and ventilator modes (d), which are arguably among the most important ventilator feature elements. Structured distributional separation occurs for continuous breath variables such as tidal volume (e), driving pressure (f), and elastance ($V_T/(p_{\max} - p_{\text{base}})$). PEEP (c) and ventilator mode (d) of UMAP labels identify the median value of each individual phenotype; probability densities (e–h) are computed from original data and colored according to panel (a). *Modes: spontaneous (SPONT), Pressure controlled (PC), Synchronized controlled (SC), and Adaptive Pressure Volume Controlled (APVC)*

Table 2: Cohort label properties. Columns identify: cohort-level label, the contained percentage of 10-second windows, the number of contained patients (N_{pat}), the number of contained individual phenotypes (N_{pheno}), the median[IQR] of baseline pressures (p_{base}, typically PEEP) and pressure change ($\Delta p := p_{\text{peak}} - p_{\text{base}}$) in cm H_2O, the median[IQR] of tidal volumes (V_T) in mL/kg, and the dominant associated ventilator mode. Values are determined from breath-level data aggregated over individual phenotypes with a given cohort-level label. *=5–10% SPONT, **=10–20% SPONT, ***= 40% SPONT

Label	Total%	N_{pat}	N_{pheno}	p_{base}	Δp	V_T	$\Delta p/V_T$	MV mode
1	15.5	23	101	10	12.1[3.7]	6.3[1.0]	1.9[0.6]	APVCMV
2	13.8	22	101	12	14.2[3.3]	6.0[1.1]	2.3[1.0]	APVCMV
3	11.4	11	52	8	12.7[4.1]	7.9[1.3]	1.5[0.4]	PCMV*
4	8.4	12	37	14	12.2[6.9]	5.9[0.2]	2.0[1.3]	APVCMV*
5	7.3	11	32	12	15.1[13.6]	5.9[0.1]	2.7[2.4]	APVCMV
6	6.9	17	58	11	13.1[2.9]	6.2[1.3]	2.1[0.5]	APVCMV
7	6.3	8	49	14	12.6[2.9]	6.2[1.3]	1.9[0.7]	APVCMV
8	6.2	11	49	16	13.4[2.3]	6.0[0.6]	2.2[0.6]	APVCMV
9	6.2	9	51	16	15.9[6.0]	5.9[2.8]	2.6[2.4]	APVCMV
10	4.1	11	34	8	9.7[2.7]	6.8[1.5]	1.6[0.4]	APVCMV
11	3.7	5	25	5	10.7[0.2]	6.5[0.7]	1.7[0.2]	PCMV**
12	3.4	14	22	5	8.9[4.1]	7.0[2.4]	1.2[0.8]	APVCMV***
13	2.5	6	10	8	11.1[1.4]	6.6[1.0]	1.7[0.2]	APVCMV
14	1.7	11	27	14	13.3[2.9]	6.0[1.3]	2.0[0.8]	APVCMV
15	1.5	5	14	10	13.5[1.9]	6.5[0.3]	2.1[0.4]	APVCMV
16	1.1	10	16	14	21.3[9.5]	5.6[1.8]	3.7[3.1]	APVCMV

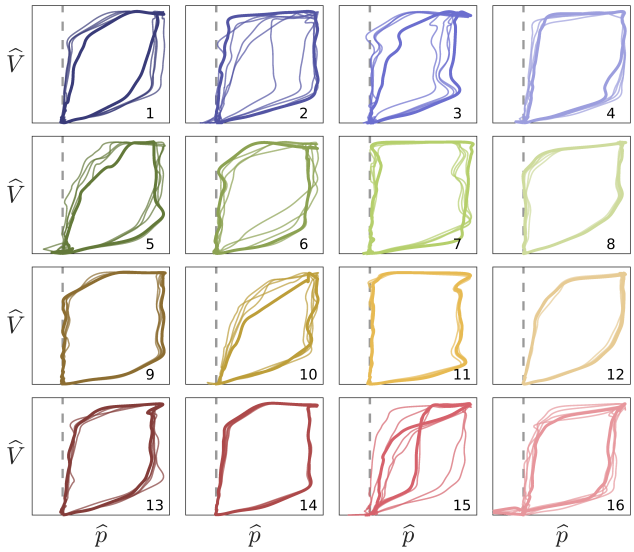

Figure 5: Non-dimensional waveform shapes. Pressure-volume traces correspond to median (bold) and nearby (thin) window characterizations of each cohort phenotype. Labels and colors correspond to Fig4a. Vertical and horizontal scales axes correspond to $\widehat{V} := V/V_T$ and $\widehat{p} := (p(t) - p_{\text{base}})/(p_{\text{peak}} - p_{\text{base}})$, respectively, per Fig4e–g The dashed line indicates baseline pressure. Cohort phenotypes differentiate waveform shape characteristics and pressure-volume coordination in conjunction with associated dimensional properties. Intra-group variation is naturally high given the low specificity of each type.

4. Discussion

This study presents a framework for extracting meaningful, low-dimensional characterizations of lung-ventilator system (LVS) states from ventilator records and observable data of mechanically ventilated medical patients. Research into ARDS and VILI involves studying patient-ventilator interactions from data with a high degree of heterogeneity. Temporal analysis of highly heterogeneous LVS data is required to disentangle iatrogenic effects and changes in patient dynamics from changes in ventilator settings. This work facilitates the analysis of LVS behavior and its changes from continuous MV data, aiming to generate hypotheses about care improvement.

The phenotyping pipeline is built to flexibly handle different representations or analyses of ventilator waveform data, commonly including airway pressure and volume (or flow). The process involves aggregating segmented analyses of individual patient data over short (10-second) intervals and empirically identifying clusters of similar states. Consequently, the observable LVS data is reduced to a small set of patient-level phenotypes, making it discrete and more manageable compared to continuous joint high-resolution waveforms and breath-wise ventilator settings. LVS description employed a generic model to transform waveform data rather than a mechanistic model which would condition segmentation on modeled physiology.

Experiments conducted on clinical data of 35 patients with strong ARDS risks, including 8 2020 patients with COVID-19, found the automated phenotyping process sufficient to discern between changes in the ventilator and the patient components of the LVS system. Individual LVS phenotypes were primarily determined by ventilator setting changes, given that changes in mode, PEEP, and tidal volume can profoundly affect waveform shapes as well. However, temporal changes in phenotype uncoordinated with ventilator changes were also present in all patients with more than 12 hours of data, revealing changes in the patient side of the LVS system. Not all such changes were captured by the naive segmentation based on uniformly weighted

data-derived features. Unlike the rapid and instantaneous transitions related to MV setting changes, patient-side changes exhibit a variety of behaviors including: continuous but non-monotonic progression, transient behavior, and alternation between both similar and non-similar breath patterns. The investigation identified pervasive and varied non-ventilator changes occurring within the LVS system, which was a primary goal of this work. Progressive changes suggest effects on lung physiology related to VILI and ARDS. Others suggest multi-breath scale volatility potentially related to dyssynchrony between the patient and adaptive control mechanisms. These behaviors could be detected through principal component analysis of data over intervals of static MV, but additional EHR data required to adequately explain them are presently unavailable.

4.1. Validation and Interpretation

Clinical validation formally requires benchmarking computed phenotypes against documented patient conditions [35]. Target biomarkers of breath behavior do not currently exist, and patient outcomes are extremely unlikely to relate directly to the respiratory data of initial snippets of longer encounters. When such targets are established, labeling only necessary and sufficient LVS variables under hyper-parameters objectively optimized for a given purpose. Instead, analytical validation used a naive system representation and hyper-parameter choice to investigate the structure of intra-label variability and to demonstrate label consistency in relation to changes in PEEP, ventilator mode, and tidal volume. Not all such changes induced label changes; waveform shape similarity under different ventilator settings was sufficient to preserve grouping. Coordination between settings and label changes (SI Table A.3) indicates that the phenotypings, which did not include physiological information (model or label assumptions), are more granular than ventilator setting stratification. Intra-label variability indicated the presence of potential label subtypes, suggesting that hierarchical or multi-stage clustering may be important for future applications. Although 10-second window scalar phenotypes are directly incomparable to breath-wise vector types of ventilator dyssynchrony (VD), changes in label-described behavior strongly coordinate with changes in VD type. Notably, both phenotype variability analysis and VD labels identified similar temporal patterns in many cases including those presented (in text in SI) without VD knowledge informing LVS descriptors. Additionally, esophageal pressures were not encoded into phenotypes but are required to confirm certain VD types [12].

The need for a framework to develop hypotheses about temporal effects and outcome validation targets for MV from retrospective cohort data motivated the methodology presented in this work. For specific choice of hyper-parameters applied to individual phenotypes, the $\sim$1.5M breaths reduce to a small set of 16 pV loop shapes Fig5 and distributional statistics that include MV settings (Table2, Fig4). Cohort labels demonstrably partition data into consistent groups, with an expected high degree of variability given the heterogeneity of LVS behaviors. Signs of dyssynchrony are apparent in these median pV shapes such as ineffective triggering (sub-baseline pressures in 5,15, and 16) and flow limitation (inspiratory coving in 3, 11, and 15). This indicates that some of the cohort scale phenotypes, while broader and less specific than VD types, center on elements of dyssynchronous behavior. Including VD labels or other physiological information in feature descriptors may better align phenotypes with VD labels in applications targeting LVS specific behaviors.

4.2. Limitations and Improvements

Combining data assimilation-based parametrization with unsupervised learning ([18]) overcomes primary shortcomings of existing approaches. In particular, the mechanism-free encoding of waveform data into

parameters with a priori definition circumvents patient- and care-dependent heterogeneity which strongly limit physiological model use in this domain ([10]). This approach digitizes waveform data into statistical distributions under 10-second (3–4 breath) stationarity without imposing other physiological assumptions that limit generalizability of mechanistic models. Although physiological information is not incorporated here, other analyses of waveform data can easily augment or replace LVS the waveform descriptors used in this work.

An important limitation of this work and its clinical application regards the dependence on hyper-parameters and a distance function used in dimensional reduction and group labeling. Fixed UMAP and narrow DBSCAN parameter search ranges do not adequately account for individual record length, internal waveform heterogeneity, or the number of unique ventilator settings that affect individual phenotype resolution. Presented examples clearly showed segmentation was insufficient resolve certain changes while also generating smaller, low occurrence phenotypes. To capture LVS behavior optimally with this method, phenotype resolution depends on the LVS descriptors used, the length and variability of data records, and the application target. This optimization is important for clinical use but lies beyond the present scope focused on obtaining low-dimensional representation of LVS evolution.

Select data sources limit the applicability and strength of conclusions based on empirical LVS phenotypes. The pipeline ignored esophageal pressure data, which are essential to confirm certain dyssynchronies. Analysis avoided these data because their rarity limits generalizability, they require high model resolution to resolve, and their inconsistencies (gaps, calibration) prevent continuous time characterization. The waveform parametrization also relies on ventilator-identified breath rate, so the pipeline lacks the flexibility needed to identify double-triggered VD events that occur over multiple ventilator cycles. Most importantly, the analysis omitted extra-LVS influences on pV waveforms such as patient sedation, neuromuscular blockade use, posture, and airway secretions as these data are not available. These patient-state factors undoubtedly impact observations and must be included to properly vet phenotypes identified under ventilator stationarity (cf. Fig3e–g and SI FigsB.6e–g).

The similarity metric used in dimensional reduction also requires deeper considerations. Individual experiments, as well as cohort phenotype construction, weight LVS descriptor dimensions equally to identify an empirical data segmentation uninformed by prior knowledge. Uniform components weights are sub-optimal, as e.g., mid-expiratory parameter variance should realistically have less impact than PEEP on LVS state category. An optimal weighting strategy requires targeted apportioning, but objective criteria are currently unknown and are likely to depend on the downstream use of phenotypes. Identifying appropriate relative weightings of waveforms data, settings information, and extra-LVS factors (posture, sedation, etc.) for segmentation is ongoing work.

4.3. Concluding Remarks

This work continues to develop a flexible operationalization of lung-ventilator systems for analyzing patient-ventilator interactions and breath types over extended timescales. It advances the study of VILI and its connection VD by distilling patient-ventilator dynamics embedded in data into discrete phenotypic classes that can be analyzed over time. The research identified system changes unattributable to care-side ventilator in order to being isolating patient-side dynamics. Assessing this type of variability is an essential first step in temporal analysis of MV patient data within the context of applied care. Ongoing work toward

formulating hypotheses about system trajectories related to intervention and outcome motivated cohort-scale phenotypes, providing a shared low-dimensional basis for LVS comparison. This translates LVS evolution, and questions related to protocols governing its control, into forms representable by symbolic dynamics [36, 37, 38] that can be used to examine patterns arising within patient cohorts.

Acknowledgments

This work is supported by National Heart Lung and Blood Institute awards 5R01HL151630 "Predicting and Preventing Ventilator-Induced Lung Injury" (JNS, BJS, DJA) and K23HL145011 "The Detection, Quantification, and Management of Ventilator Dyssynchrony" (PDS). Big-ups as always to Meg Rebull for local administrative support.

Declarations of Interest

None. The authors have no conflicts of interest to disclose.

Declaration of Generative AI and AI-assisted Technology Use

The author used Chat-GPT 3.5 to suggest rephrasings of complex sentences in early drafts. No generative AI output was used directly. No other AI-assisted technologies were employed.

References

[1] Eddy Fan, Jesus Villar, and Arthur S Slutsky. Novel approaches to minimize ventilator-induced lung injury. *BMC medicine*, 11(1):1–9, 2013.

[2] Gerard F Curley, John G Laffey, Haibo Zhang, and Arthur S Slutsky. Biotrauma and ventilator-induced lung injury: clinical implications. *Chest*, 150(5):1109–1117, 2016.

[3] Arthur S Slutsky and V Marco Ranieri. Ventilator-induced lung injury. *New England Journal of Medicine*, 369(22):2126–2136, 2013.

[4] Roy G Brower and Gordon D Rubenfeld. Lung-protective ventilation strategies in acute lung injury. *Critical care medicine*, 31(4):S312–S316, 2003.

[5] Nicola Petrucci and Walter Iacovelli. Lung protective ventilation strategy for the acute respiratory distress syndrome. *Cochrane Database of Systematic Reviews*, (3), 2007.

[6] Yuda Sutherasan, Maria Vargas, and Paolo Pelosi. Protective mechanical ventilation in the non-injured lung: review and meta-analysis. *Annual Update in Intensive Care and Emergency Medicine 2014*, pages 173–192, 2014.

[7] Acute Respiratory Distress Syndrome Network. Ventilation with lower tidal volumes as compared with traditional tidal volumes for acute lung injury and the acute respiratory distress syndrome. *New England Journal of Medicine*, 342(18):1301–1308, 2000.

[8] Shea E Cochi, Jordan A Kempker, Srinadh Annangi, Michael R Kramer, and Greg S Martin. Mortality trends of acute respiratory distress syndrome in the United States from 1999 to 2013. *Annals of the American Thoracic Society*, 13(10):1742–1751, 2016.

[9] Marc Moss and David M Mannino. Race and gender differences in acute respiratory distress syndrome deaths in the United States: an analysis of multiple-cause mortality data (1979–1996). *Critical care medicine*, 30(8):1679–1685, 2002.

[10] JN Stroh, Bradford J Smith, Peter D Sottile, George Hripcsak, and David J Albers. Hypothesis-driven modeling of the human lung-ventilator system: A characterization tool for acute respiratory distress syndrome research. *Journal of Biomedical Informatics*, page 104275, 2022.

[11] Peter D Sottile, David Albers, Carrie Higgins, Jeffery Mckeehan, and Marc M Moss. The association between ventilator dyssynchrony, delivered tidal volume, and sedation using a novel automated ventilator dyssynchrony detection algorithm. *Critical Care Medicine*, 46(2):e151, 2018.

[12] Peter D Sottile, David Albers, Bradford J Smith, Marc M Moss, et al. Ventilator dyssynchrony–detection, pathophysiology, and clinical relevance: A narrative review. *Annals of Thoracic Medicine*, 15(4):190, 2020.

[13] G Schmidt. Ventilator waveforms: clinical interpretation. *Principles of Critical Care*, 427:443, 2005.

[14] Elizabeth Emrath. The basics of ventilator waveforms. *Current pediatrics reports*, 9:11–19, 2021.

[15] Deepak K Agrawal, Bradford J Smith, Peter D Sottile, and David J Albers. A damaged-informed lung ventilator model for ventilator waveforms. *Frontiers in physiology*, 12, 2021.

[16] Cong Zhou, J Geoffrey Chase, Qianhui Sun, Jennifer Knopp, Merryn H Tawhai, Thomas Desaive, Knut Möller, Geoffrey M Shaw, Yeong Shiong Chiew, and Balazs Benyo. Reconstructing asynchrony for mechanical ventilation using a hysteresis loop virtual patient model. *BioMedical Engineering OnLine*, 21(1):1–20, 2022.

[17] Yuhong Chen, Kun Zhang, Cong Zhou, J Geoffrey Chase, and Zhenjie Hu. Automated evaluation of typical patient–ventilator asynchronies based on lung hysteretic responses. *BioMedical Engineering OnLine*, 22(1):102, 2023.

[18] Y Wang, JN Stroh, George Hripcsak, Cecilia C Low Wang, Tellen D Bennett, Julia Wrobel, Caroline DerNigoghossian, Scott Mueller, Jan Claassen, and DJ Albers. A methodology of phenotyping ICU patients from EHR data: high-fidelity, personalized, and interpretable phenotypes estimation. *in review Journal of Biomedical Informatics*, xx(x):xxx, *in review* 2023.

[19] Cheolhyeong Park and Deokwoo Lee. Classification of respiratory states using spectrogram with convolutional neural network. *Applied Sciences*, 12(4):1895, 2022.

[20] Deepak K Agrawal, Bradford J Smith, Peter D Sottile, George Hripcsak, and David J Albers. Quantifiable identification of flow-limited ventilator dyssynchrony with the deformed lung ventilator model. *Computers in Biology and Medicine*, page 108349, 2024.

[21] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005.

[22] Pavel Sakov, Geir Evensen, and Laurent Bertino. Asynchronous data assimilation with the EnKF. *Tellus, Series A: Dynamic Meteorology and Oceanography*, 62(1):24–29, 2010.

[23] Mahamed GH Omran, Andries P Engelbrecht, and Ayed Salman. An overview of clustering methods. *Intelligent Data Analysis*, 11(6):583–605, 2007.

[24] Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

[25] Connor Meehan, Stephen Meehan, and Wayne Moore. Uniform manifold approximation and projection (UMAP v4.2). *MATLAB Central File Exchange*, 2022. `https://www.mathworks.com/matlabcentral/fileexchange/71902`.

[26] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, volume 96, pages 226–231, 1996.

[27] Erich Schubert, Jörg Sander, Martin Ester, Hans Peter Kriegel, and Xiaowei Xu. DBSCAN revisited, revisited: why and how you should (still) use DBSCAN. *ACM Transactions on Database Systems (TODS)*, 42(3):1–21, 2017.

[28] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008.

[29] Dmitry Kobak and George C Linderman. Initialization is critical for preserving global data structure in both t-SNE and UMAP. *Nature biotechnology*, 39(2):156–157, 2021.

[30] Asa Ben-Hur, David Horn, Hava T Siegelmann, and Vladimir Vapnik. Support vector clustering. *Journal of machine learning research*, 2(Dec):125–137, 2001.

[31] Jaewook Lee and Daewon Lee. Dynamic characterization of cluster structures for robust and inductive support vector clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(11):1869–1874, 2006.

[32] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.

[33] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.

[34] C Radhakrishna Rao. The use and interpretation of principal component analysis in applied research. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 329–358, 1964.

[35] Jennifer C Goldsack, Andrea Coravos, Jessie P Bakker, Brinnae Bent, Ariel V Dowling, Cheryl Fitzer-Attas, Alan Godfrey, Job G Godino, Ninad Gujar, Elena Izmailova, et al. Verification, analytical validation, and clinical validation (v3): the foundation of determining fit-for-purpose for biometric monitoring technologies (biomets). *npj digital Medicine*, 3(1):55, 2020.

[36] José M Amigó, Karsten Keller, and Valentina A Unakafova. Ordinal symbolic analysis and its application to biomedical recordings. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(2034):20140091, 2015.

[37] Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. 2nd edition, 2021.

[38] Yoshito Hirata and José M Amigó. A review of symbolic dynamics and symbolic reconstruction of dynamical systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 33(5), 2023.

[39] Ville Satopaa, Jeannie Albrecht, David Irwin, and Barath Raghavan. Finding a "kneedle" in a haystack: Detecting knee points in system behavior. In *2011 31st international conference on distributed computing systems workshops*, pages 166–171. IEEE, 2011.

SI A. Label changes in relation to vent settings changes

Table A.3 shows that most vent settings changes are accompanied by changes in labels. However, very few phenotype label changes correspond to changes in vent settings. *Over 64% of ventilator settings changes are identified in label changes.* A larger proportion (>87%) are identified when limited to PEEP, tidal volume, and model changes, which induce significant waveforms changes compared with other settings such as mandatory breath rate (set_rate). *Few (8%) identified changes in label, however, are directly associated in time with ventilator settings changes.*

Table A.3: N_s indicate the number of ventilator settings changes in set_PEEP, set_ptrigger, set_qtrigger, set_rate, set_fio2, set_ie, set_flowpat, set_mode, set_vt, and vt_set. N_l indicates the number of persistent label changes, counting those lasting longer than 30 seconds, to omit isolated transient changes and variability occurring as mixed-breath types (*e.g.*, FigB.7 during 11–14 hours, characterized by both alternation between labels #10 and #13 labels and changes in ML-identified VD type.) Column 's2l' indicates the percentage of vent settings changes that occur with a label change within 100 secondsColumn 'l2s' indicate the number of label changes that occur within 100 seconds of vent changes.

ID	N_s	s2l (%)	N_l	l2s (%)	ID	N_s	s2l (%)	N_l	l2s (%)
101	45	73.33	66	16.67	129	11	45.45	138	3.62
102	1	-	1	-	130	14	85.71	296	6.42
103	2	100.00	35	8.57	131	10	70.00	25	20.00
104	25	36.00	356	1.97	133	37	78.38	154	18.83
105	1	100.00	46	4.35	134	76	5.26	175	0.57
107	99	56.57	288	9.72	135	104	86.54	695	6.47
108	20	75.00	342	2.92	136	222	59.01	590	6.78
110	3	33.33	73	1.37	137	137	54.01	166	14.46
111	23	78.26	328	7.01	138	10	60.00	38	21.05
112	59	45.76	177	6.21	139	24	70.83	451	3.99
113	14	78.57	260	6.15	140	47	100.00	629	1.75
114	48	29.17	50	24.00	141	73	45.21	431	7.42
115	0	-	4	0	143	37	62.16	713	2.66
116	83	80.72	370	12.70	144	83	77.11	421	9.03
117	13	76.92	1000	1.40	145	50	62.00	380	8.42
119	51	98.04	265	8.30	146	220	42.27	650	9.38
120	57	40.35	246	8.94	149	39	25.64	296	4.73
123	18	88.89	460	3.91	150	12	75.00	424	2.12
					mean	52	64.57	325	8.00

SI B. Individual Experiments, continued

This supplement continues illustrated examples of §3.1.

Figure B.6 panels a–d illustrate the analysis of Patient #103 whose data consists of 7 record hours with one simple ventilator setting change. Only ventilator PEEP (a) is changed while there are three primary behaviors identified (b,d). The reduction of PEEP occurs about 2 hours following a rise in early flow limited breaths (eFL, panel c). This PEEP change (from 8 to 5 cm H_2O) shifts peak pressure from 16 to 12 cm H_2O for about an hour, at which time higher esophageal pressures returns. These breaths are identified as normal (NL) [11]. Increased specificity may be pursued by local segmentation or other dimensional reduction methods.

A closer look at label 1 of patient #103:. The first principal component loadings (panel e, black) for LVS descriptors over the first 5-hour period track the sequence of normal and eFL VD labels (f, shown as 5 minute statistics for clarity). Within the same breath phenotype (label 1), the *sign* of the component loading statistically the eFL VD labels (AUROC=0.8718); high positive values are associated with eFL breaths (f,g; green) where pressure maxima proceed volume maxima. These LVS variations result from changes in the patient component, as there is no change of ventilator settings. Note that direct correlation between continuous loading values on 10 second windows and statistical breath-wise binary VD label is not well-defined while binary-to-binary comparison is.

The patient #113 (Figure B.7) dataset is nearly twice as long with again only one PEEP change occurring after 10.5 hours of the 15.6 hour record. Breaths are stably identified as normal-type until about 8 hours,

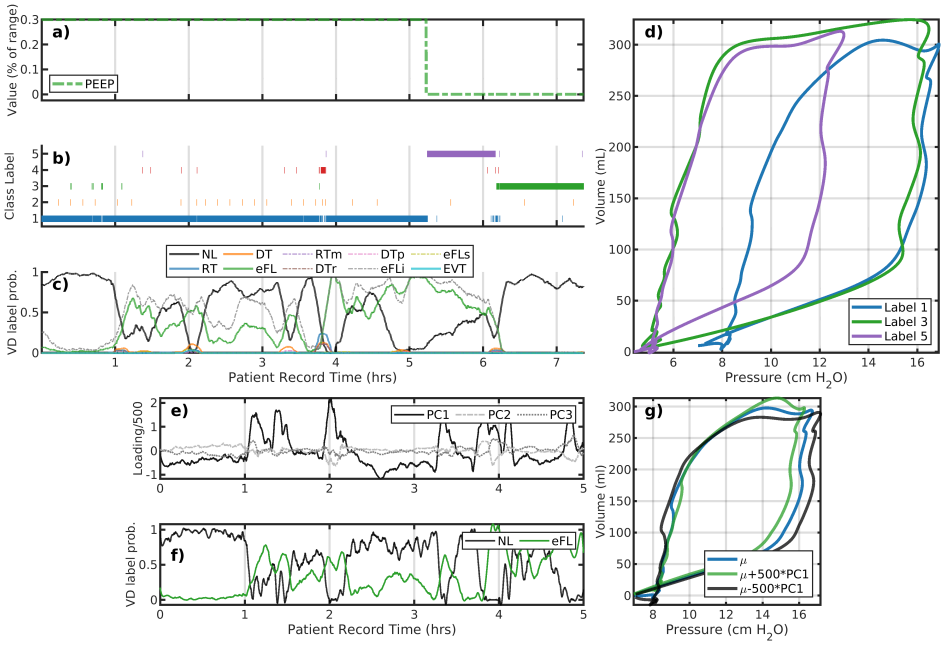

Figure B.6: Analysis of patient #103 LVS data (a–d) and the initial a 5-hour interval (e–g). Panels a–c correspond to changes in ventilator settings, segmentation labels, and identified VD type, respectively. The horizontal axis for these panels is the patient record time in hours. The panel (d) shows the model image of segmented data median parameters, which characterize the pV loops of breaths with that label (shown with the same color). Evolution of the LVS can be parsed pictorially from these figures. Large positive variations in the first principal component loading (e, black) for the initial 5-hour period align with VD labels indicating eFL type breaths (f) for this period. Specifically, this suggests discrimination of breaths shapes (g) can be differentiated using qualitatively criterion on local loadings or other segmentation.

occupying two cluster-identified similar breath shapes. This is followed briefly by eFL breaths and a transition to a new characterization (label 8, light green) for about 30 minutes. In the following period (9–14 hours), breaths are characterized by lower pressure maxima (label 10, gold); these are associated/identified with reverse-trigger breaths (primarily RTm) and waveforms featuring pronounced inspiratory pressure drop. The reduction in PEEP slightly increases the incidence of normal breaths during 11–14 hours although this results in the more frequent appearance of shallow breaths (label 13, red).

SI B.1. Intracluster normal and eFL in p111, label2

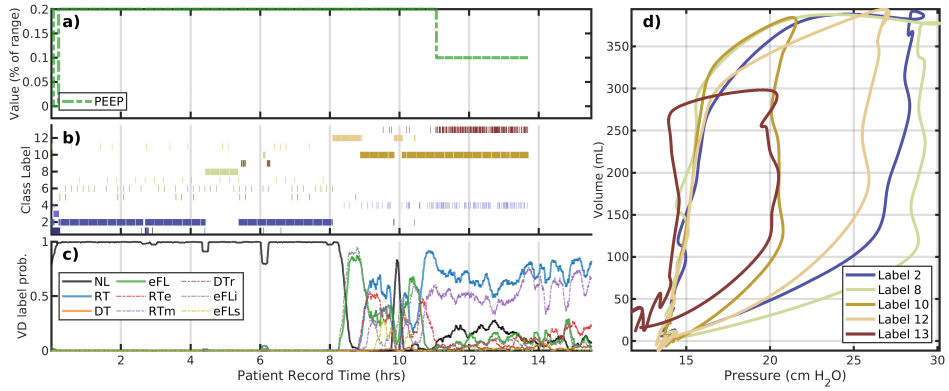

Figure B.7: The patient #113 evolution also includes only PEEP changed. The layout is the same as panels a–d of the previous figure. Under constant ventilator settings, breaths undergo transition several times including intervals of VD prior to PEEP change around 10.5 hours. A 1-hour long shift from label 2 to 8 occurs around 8 hours during which breaths decrease peak pressure and includes an increase in eFL and RT VD occurrence. After the PEEP change, breaths remain highly dyssynchronous and primarily centered around the characterization with label 10.

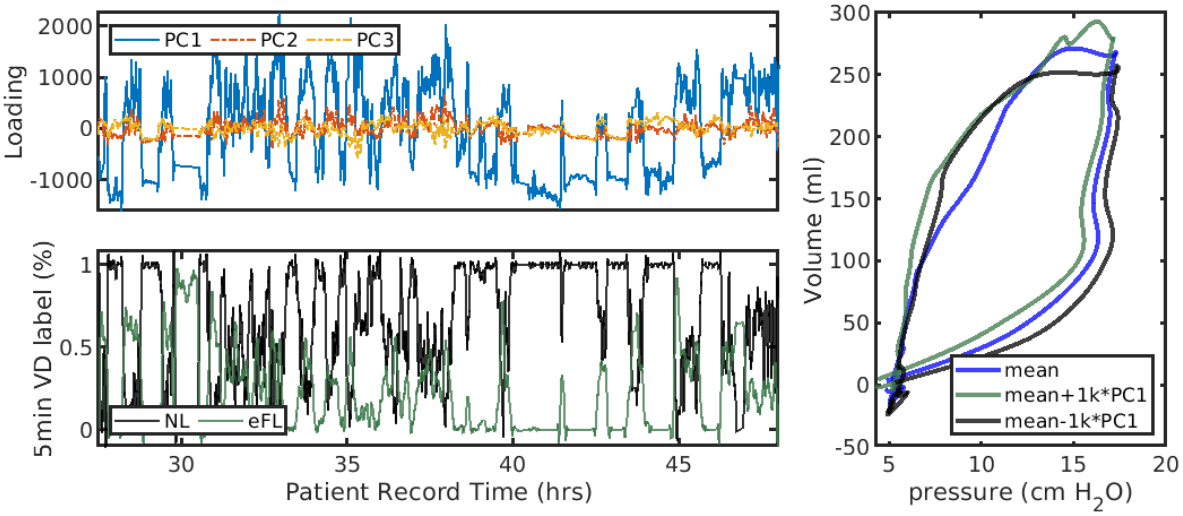

Figure B.8: The sign of PC1 loading roughly divides the VD classes in p111, label2. A threshold for the PC1 loading at zero roughly separates NL and eFL labels by 34%/65% and 85%/14%, respectively, with NL labels strongly associated with negative loadings. The optimal threshold (∼0.05) offers only subtle improvement. The right panel illustrates low fidelity changes in the cluster median pV loop (blue) when modified by these negative (black, more associated with NL) and positive (green, eFL) loadings. Note that this involves comprising 10-second properties (representing typically ∼3–4 breaths) to breath-wise labels, and some representation errors thus arise from summarizing binary VD labels over all breaths intersecting a 10-second analysis window.

SI B.2. Qualitiative equivalence of labels via tSNE & UMAP

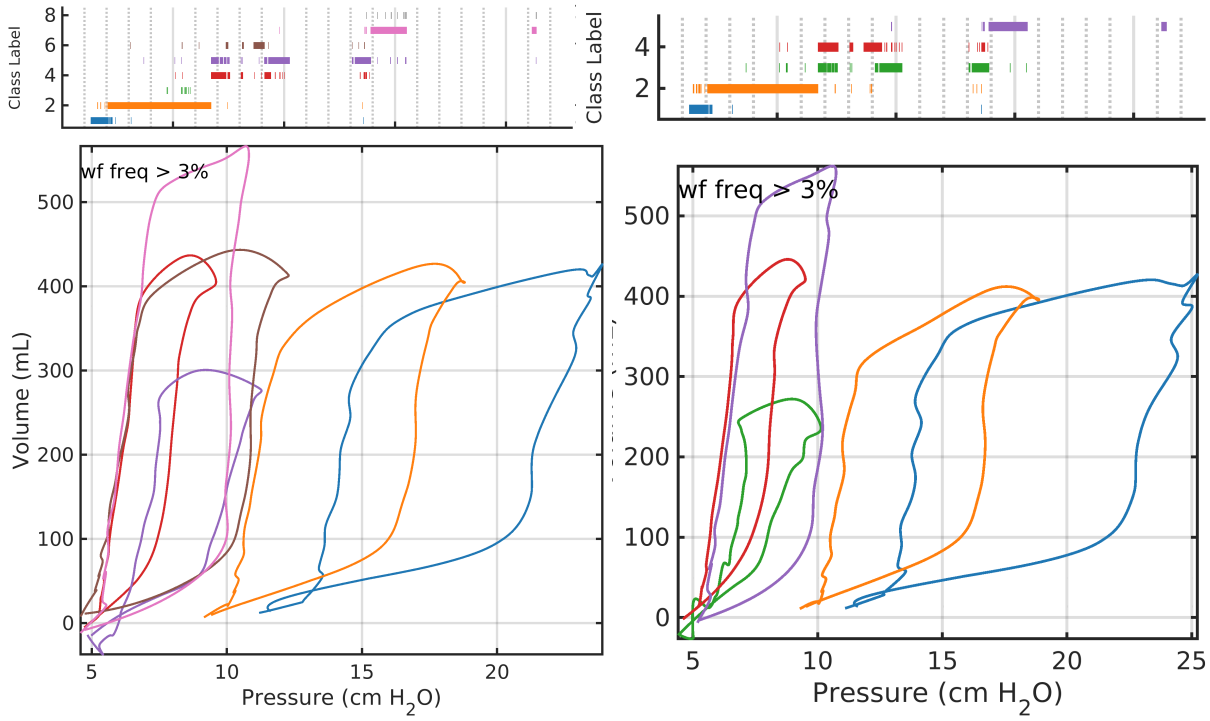

Figure B.9: Patient 101 clustering using tSNE (left) and UMAP (right) feature reduction stages, as an example. Temporal evolution of the LVS is qualitatively similar regardless of whether UMAP (neighbor size=5, minimum distance=0.01) or tSNE (exaggeration=20, perplexity=50, 5000 iterations) projection is used. DBSCAN parameter must also be adjusted as coordinate scales differ between the projections. For the plot shown, DBSCAN hyper-parameters (N_{pts}, ε) are (10,4) following tSNE and (4,1.5) following UMAP. Mild variations in pV characterizations result from medians of different point distributions.

SI C. Influence of Hyperparameter choices on cohort phenotypes

For each of the 721 individual phenotypes, feature vectors defined by the 5-number summaries of period, PEEP, maxima of volume and pressure, ventilator settings, and estimated parameters of range-normalized waveform were assembled from the population of LVS windows with a given label. Ventilator mode was represented as a vector of percentages of each mode rather than a vector of binary categories, which eliminated the need for the Gower distance. UMAP applied to these cohort feature vectors with the scaled-euclidean metric produced a relatively stable point configuration across various hyper-parameter choices; 12 point neighborhoods (2% of data) with a minimum distance of 1 unit were adopted as values. Identified groups were more sensitive to DBSCAN labeling hyper-parameters. Figure C.10a shows the possibilities of different groupings based on the search neighborhood size (ε). Subsequent results in the main section employ a hyper-parameter choice at the 'knee-point' [39] to balance generalizability and specificity. A more specific labeling ($\varepsilon = 2.5$) shown in panel b, is qualitatively similar to that of main text.

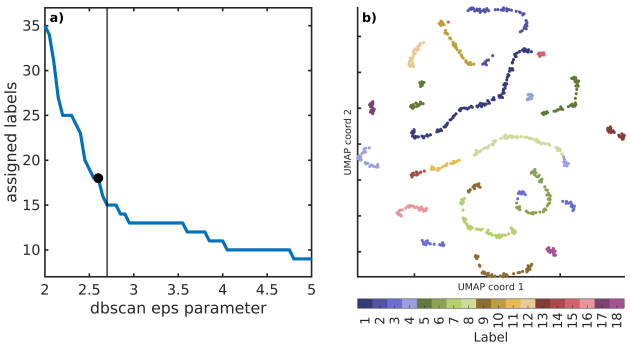

Figure C.10: DBSCAN search radius (ε) *v.* the number of identified groups. The black line indicates $\varepsilon = 2.7$ selected for cohort clustering. Choices of $\varepsilon \in [2.67, 2.82]$ yield equivalent results following increased granularity of groupings at ε lower values.

Table C.4: The equivalent of Table2 for the alternate choice of hyper-parameter $\varepsilon = 2.5$

Label	Total%	N_{pat}	N_{pheno}	p_{min}	p_{drive}	V_{max}	dp/dV	MV mode
1	15.5	11	101	8	12.7[4.1]	7.9[1.3]	1.5[0.4]	APVCMV
2	11.4	16	52	12	13.3[3.6]	6.2[2.1]	2.3[1.2]	PCMV*
3	8.4	5	37	10	13.5[1.9]	6.5[0.3]	2.1[0.4]	APVCMV*
4	7.7	8	45	14	12.6[2.9]	6.2[1.3]	1.9[0.7]	APVCMV
5	6.9	11	58	16	13.4[2.3]	6.0[0.6]	2.2[0.6]	APVCMV
6	6.3	7	49	11	16.6[12.5]	5.9[0.7]	3.4[2.4]	APVCMV
7	6.2	6	49	8	11.1[1.4]	6.6[1.0]	1.7[0.2]	APVCMV
8	6.2	23	56	10	12.1[3.7]	6.3[1.0]	1.9[0.6]	APVCMV
9	6.2	10	51	14	21.3[9.5]	5.6[1.8]	3.7[3.1]	APVCMV
10	4.1	12	34	14	12.2[6.9]	5.9[0.2]	2.0[1.3]	APVCMV*
11	4.0	14	20	5	8.9[4.1]	7.0[2.4]	1.2[0.8]	APVCMV
12	3.7	16	25	12	14.3[3.2]	6.0[0.4]	2.3[0.9]	PCMV**
13	3.4	17	22	11	13.1[2.9]	6.2[1.3]	2.1[0.5]	APVCMV***
14	3.3	11	12	8	9.7[2.7]	6.8[1.5]	1.6[0.4]	APVCMV
15	2.5	8	10	12	15.1[12.5]	5.9[0.1]	2.4[2.2]	APVCMV
16	1.7	11	27	14	13.3[2.9]	6.0[1.3]	2.0[0.8]	APVCMV
17	1.5	9	14	16	15.9[6.0]	5.9[2.8]	2.6[2.4]	APVCMV
18	1.1	5	16	5	10.7[0.2]	6.5[0.7]	1.7[0.2]	APVCMV