Unsupervised classification of plethysmography signals with advanced visual representations

Thibaut Germain [1]**, Charles Truong** [1]**, Laurent Oudre** [1,*] **and Eric Krejci** [2]

[1] *Université Paris Saclay, Université Paris Cité, ENS Paris Saclay, CNRS, SSA, INSERM, Centre Borelli, F-91190, Gif-sur-Yvette, France*
[2] *Université Paris Cité, Université Paris Saclay, ENS Paris Saclay, CNRS, SSA, INSERM, Centre Borelli, F-75006, Paris, France*

Correspondence*:
Laurent Oudre
laurent.oudre@ens-paris-saclay.fr

2 ABSTRACT

Ventilation is a simple physiological function that ensures the vital supply of oxygen and the elimination of CO_2. The recording of the airflow through the nostrils of a mouse over time makes it possible to calculate the position of critical points, based on the shape of the signals, to compute the respiratory frequency and the volume of air exchanged. These descriptors only account for a part of the dynamics of respiratory exchanges. In this work we present a new algorithm that directly compares the shapes of signals and considers meaningful information about the breathing dynamics omitted by the previous descriptors. The algorithm leads to a new classification of inspiration and expiration, which reveals that mice respond and adapt differently to inhibition of cholinesterases, enzymes targeted by nerve gas, pesticide, or drug intoxication.

Keywords: Respiration, Breathing, Dynamic Time Warping (DTW), Clustering, Machine Learning

1 INTRODUCTION

Measurement of respiratory function in conscious, spontaneously breathing animals is essential in different settings such as studying drug effects on the respiratory system (Murphy, 2002), monitoring mouse models of human diseases (Willmann et al., 2017) or evaluating airway irritant molecules (Vijayaraghavan et al., 1993). Plethysmography methods are commonly used to record respiration. Several plethysmographs exist (whole-body (WBP, Bartlett and Tenney (1970)), dual-chamber (DCP, Hoymann (2012)), head out-of-body (HOB, Vijayaraghavan et al. (1993)), and the choice between them is based on a trade-off between invasiveness and accuracy of measurement (Bates and Irvin, 2003).

For WBP, the mouse is placed in near-natural conditions: it is a large box where the mouse is not restrained and can move freely. Changes in pressure or flow within the chamber are measured over time, reflecting changes of volume, humidity and temperature of air entering and leaving the lung. Flow or differential pressure recordings allow the computation of the respiratory frequency and the volumes

24 exchanged. Nevertheless, lung function is poorly measured, and the reproducibility of the experiment
25 depends on many environmental parameters (Bates and Irvin, 2003; Bruggink et al., 2022).

26   The DCP consists of two sealed compartments where the animal's head is in one compartment and its
27 body in the other. The mouse is constrained in a tube with the nose pointing into the nasal compartment
28 (respiration is primarily nasal in mice) (Mailhot-Larouche et al., 2018). This device allows for independent
29 monitoring of the nasal airflow and the airflow caused by the thoracic movements of an animal. DCP
30 is a relevant approach to assessing the ventilatory mechanics of the respiratory system, and it provides
31 information on ventilatory and lung function (Hoymann, 2012). As the mouse is constrained, it limits the
32 duration of respiration recording to less than one hour. As an alternative, the HOP uses only the thoracic
33 compartment imposing less constraint on the mouse (Vijayaraghavan et al., 1993). Recordings from DCP
34 and HOP directly reflect the air inhaled and exhaled during respiration. These methods have been used
35 for several decades to monitor changes in mouse respiration caused by airborne chemicals on the airways
36 (Vijayaraghavan et al., 1993) and have been improved to limit air leakage from collar (Bruggink et al.,
37 2022).

38   Respiration airflows are used to compute respiratory cycle descriptors like inspiration/expiration duration,
39 air volume inhaled/exhaled, or respiratory frequency (IXO2 software, emka TECHNOLOGIES, Mailhot-
40 Larouche et al. (2018)). These descriptors are essential to quantify respiratory exchanges; nevertheless,
41 they only reveal part of the information contained in the respiration airflow.

42   Recently, a new machine learning-based method attempted to incorporate the missing information by
43 extracting respiratory cycle patterns from flows recorded in WBP (Sunshine and Fuller, 2021). They used
44 principal component analysis and a clustering algorithm to group common respiratory cycle patterns. The
45 groups reveal variations in the temporal appearance of the flows that are not detectable with standard
46 analysis of respiratory rate and tidal volume. Because each group had physiological significance, they
47 could track significant changes over time. However, the different categories cannot be associated with
48 physiological alterations or adaptations because the signals recorded in WBP mix too many parameters
49 (volume/temperature/humidity).

50   It is well established that DCP or HOP is much more accurate for studying respiratory physiology than
51 WBP. Nonetheless, descriptors inferred from the nasal or thoracic airflow also miss meaningful information
52 to describe breathing dynamics adequately. In this study, we propose a new approach to classify the
53 different respiratory behaviors from signals recorded with DCP or HOP. The method relies on a robust
54 algorithm to identify the beginning of inspiration and expiration phases. As a difference with (Sunshine
55 and Fuller, 2021) which studies the respiration cycle, we independently study inspiration and expiration.
56 Our method is based on machine learning tools for time series. It uses a K-Means clustering algorithm and
57 the well-established Dynamic Time Warping (DTW) distance. DTW compares the shape of time series
58 independently of time fluctuations. This property is particularly interesting for handling inter-individual
59 variability. In contrast to (Sunshine and Fuller, 2021), we assess the similarity between inhalation or
60 exhalation cycles directly from their shapes rather than from learned features. It allows a more robust and
61 interpretable study of respiratory behaviors and dynamics. To evaluate the relevance of this new method,
62 we exploited part of the recordings from a previous experiment of our group (Nervo et al., 2019), where we
63 studied the consequences on respiration of partial deficits of acetylcholinesterase (AChE) and its inhibition.
64 AChE normally destroys acetylcholine (ACh) in synapses of the nervous systems (central and peripheral)
65 and skeletal muscles. Inhibition of this enzyme results in respiratory arrest, which may have multiple
66 origins (Stone, 2018).
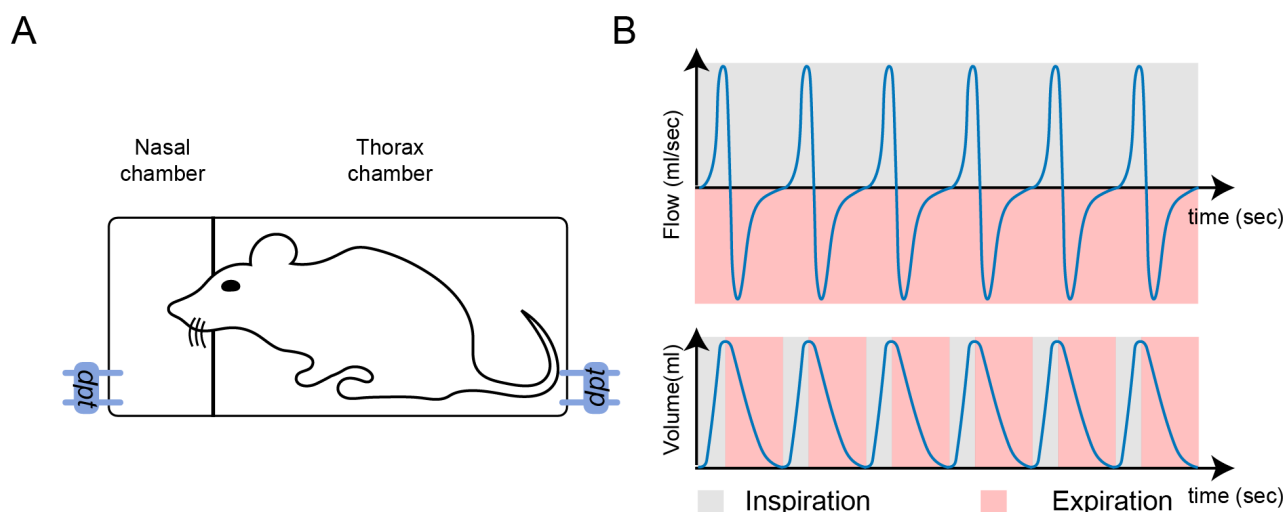
Figure 1. **A**: Illustration of a double-chamber plethysmograph. The term *dpt* stands for differential pressure transducer which measures the pressure in each compartment, the pressure then being converted to flow. **B**: Nasal airflow (top) and lung volume (bottom). During inspiration, airflow is positive (grey) and during expiration, airflow is negative (pink).

2 METHOD

2.1 Background

With a Double Chamber Plethysmograph (DCP), the nasal airflow induced by breathing is tracked through the head compartment (Hoymann, 2012), as illustrated on Figure 1-A. The flow is expressed in $ml.s^{-1}$, and is a positive quantity for inspiration, and a negative quantity for expiration. The detection of inspiration or expiration start times cannot be done accurately from the nasal airflow due to biological phenomena such as coughing or vocalizing (Bates and Irvin, 2003). As an alternative, the lung volume, which is obtained by integration of the nasal airflow, offers more robust properties for such detection (Vijayaraghavan et al., 1993). Intuitively, the lung volume fluctuates successively from being empty (inspiration start time: t_{in}) to being full (expiration start time: t_{out}). These states correspond to local minima and maxima on the volume, which are easy to track with automated procedures. The characterization of breathing phases is illustrated in Figure 1-B.
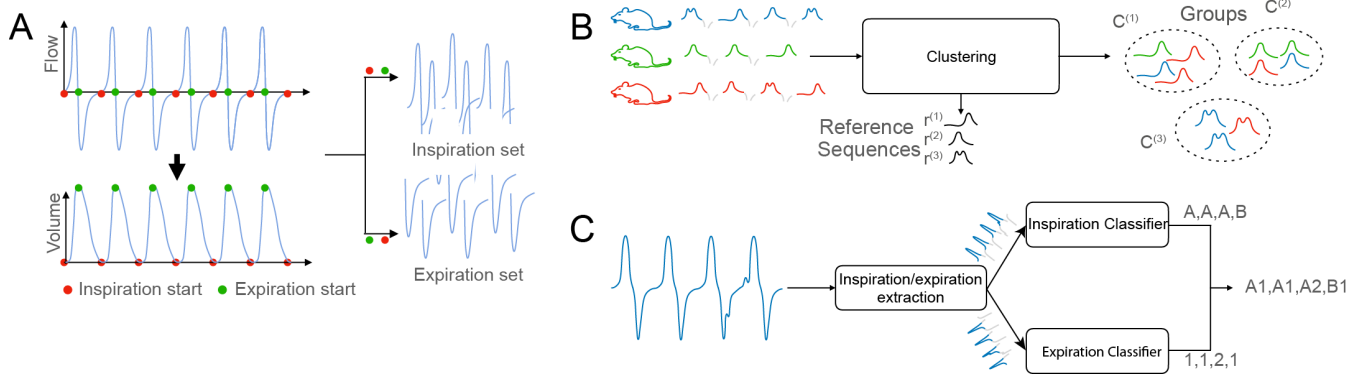
2.2 Overview of the method

The method is composed of three main steps:

1. detection of the respiratory cycles and extraction of the inspiration/expiration sequences,
2. computation of the reference sequences through an unsupervised clustering procedure,
3. characterization and symbolization of recordings based on the extracted reference sequences.

Step 1: Detection of the respiratory cycles and extraction of the inspiration/expiration sequences. The first step of the process consists in extracting the respiratory cycles from the input data. Each cycle is composed of two phases: inspiration and expiration. A segmentation algorithm isolates the two periods. Simplistically, given a raw signal s, the first step of our method outputs a set of inspiration sequences

Figure 2. A: Step 1, Detection of the respiratory cycles and extraction of the inspiration/expiration sequences. **B**: Step 2, Computation of the reference sequences. The $C^{(i)}$ denote the clusters and $\mathbf{r}^{(i)}$, the reference sequences. **C**: Step 3, Characterization and symbolization of a recordings.

87 $\{\mathbf{s}_{in}^{(1)}, ..., \mathbf{s}_{in}^{(N_s)}\}$ and a set of expiration sequences $\{\mathbf{s}_{out}^{(1)}, ..., \mathbf{s}_{out}^{(N_s)}\}$, where N_s is the total number of
88 cycles observed in the original signal $\mathbf{s}$. Figure 2-A illustrates the detection and extraction process of
89 inspiration/expiration.

90 **Step 2: Computation of the reference sequences.** The second step consists in computing a small number
91 of reference sequences from the sets of inspiration/expiration sequences. The reference sequences represent
92 groups of sequences with common properties to highlight typical inspiration/expiration behaviors. To that
93 aim, the clustering algorithm K-means is combined with the measure of fit Dynamic Time Warping (DTW),
94 which computes the similarities between sequences of potentially different lengths. The output of this step
95 is a set of inspiration reference sequences $\{\mathbf{r_{in}}^{(1)}, \mathbf{r_{in}}^{(2)}, ...\}$ and a set of expiration reference sequences
96 $\{\mathbf{r_{out}}^{(1)}, \mathbf{r_{out}}^{(2)}, ...\}$. Figure 2-B illustrates the computation process of reference sequences in the case of
97 inspiration.

98 **Step 3: Characterization and symbolization of recordings.** The objective is to automatically
99 characterize a recording $\mathbf{s}'$ using the reference sequences extracted in Step 2. To that end, the signal
100 is first segmented through the procedure described in Step 1. Then, each of the $N_{s'}$ inspiration/expiration
101 sequences present in $\mathbf{s}'$, is assigned a symbol which represents the reference sequence that is closest
102 considering the measure DTW. This procedure results in a symbolic representation of $\mathbf{s}'$, where each
103 respiratory cycle is replaced by a symbol composed of a letter (which specifies the type of inspiration
104 sequence observed) and a number (which specifies the type of expiration sequence observed), Figure 2-C
105 illustrates the process of building a symbolic representation.

106 ## 2.3 Detection of the respiratory cycles and extraction of the inspiration/expiration
107 sequences

108 As mentioned previously, nasal airflow suffers from noise, making current inspiration and expiration
109 phases detection methods unreliable. Inaccurate detection then leads to biased descriptors and eventually
110 to false experimental conclusions. To address this challenge we propose an algorithm that looks for local
111 minima and maxima of the lung volume. Let $\mathbf{s}$ denote a nasal airflow signal.

112   First the lung volume $\mathbf{v}$ is computed from the nasal airflow $\mathbf{s}$. This can be done by robust numerical
113   integration:

$$v_t := \left(\sum_{u=1}^{t} s_u \right) - (\hat{a}t + \hat{b}) \tag{1}$$

114   where $\hat{a}, \hat{b} \in \mathbb{R}$ are such that $\sum_t v_t = 0$ and $\sum_t t v_t = 0$. The affine function $t \rightarrow \hat{a}t + \hat{b}$ removes the linear
115   trend appearing during the integration process.

116   Next, the inspiration start times t_{in} and the expiration start times t_{out} are identified using a peak-
117   searching procedure that detects local minima (respectively maxima), of the nasal volume signal $\mathbf{v}$.
118   To ensure an alternation between inspiration and expiration, the algorithm first searches for all local
119   minima (corresponding to the starts of the inspirations) and then searches for the maximum between two
120   consecutive local minima. The algorithm that detects local minima/maxima is described in Appendix 1.
121   Once all inspiration/expiration start times t_{in} and t_{out} are extracted, the original nasal airflow signal $\mathbf{s}$ is
122   split into a set of inspiration sequences $\{\mathbf{s}_{in}^{(1)}, ..., \mathbf{s}_{in}^{(N_s)}\}$ and a set of expiration sequences $\{\mathbf{s}_{out}^{(1)}, ..., \mathbf{s}_{out}^{(N_s)}\}$,
123   where N_s is the total number of cycles observed in the original signal $\mathbf{s}$.

2.4   Computation of the reference sequences

125   Provided a set of inspiration/expiration sequences, we now aim to compute K reference sequences that
126   represent typical respiratory behaviors. In the following sections, $\mathcal{X} = \{\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(N)}\}$ represents a set of
127   sequences (either inspiration or expiration) of potentially different durations.

2.4.1   Clustering algorithm

129   The K reference sequences from the set $\mathcal{X}$ are computed with the well-known unsupervised clustering
130   procedure called K-Means (Kaufman and Rousseeuw, 2009). This algorithm creates K non-overlapping
131   groups (or clusters) $\{\mathcal{C}^{(1)}, \ldots, \mathcal{C}^{(K)}\}$ of sequences with common properties. Roughly, K-Means is a two-
132   step iterative refinement technique that assigns each sequence to the closest current centroid and then
133   updates each centroid with regard to the new assignments. A centroid is a reference sequence $\mathbf{r}^{(i)}$ which
134   corresponds to the average sequence of the cluster $\mathcal{C}^{(i)}$. Two crucial ingredients of the K-means algorithm
135   are the measure of fit that is used to assign each sequence to a cluster and the procedure used to compute
136   the reference sequence of each cluster. Although most publications usually use the Euclidean distance, it is
137   not possible in our context since the sequences to cluster do not have the same duration. Also, the measure
138   of fit must be invariant to some sequence properties: amplitude offset, amplitude shift, time fluctuation,
139   noise and outliers. Visual representation of each distortions are presented in Figure 3-A.

140   The main steps of our clustering approach are:

141   1. The input sequences are pre-processed by taking the z-normalized sequences.
142   2. The clustering algorithm is initialized using the K-Means++ algorithm (Arthur and Vassilvitskii, 2007).
143   3. Sequences are assigned to a cluster according to the Dynamic Time Warping (DTW) measure of fit.
144   4. The reference sequences are computed using the Batch Stochastic DTW Barycenter Averaging (BS-
145   DBA) procedure.
146   5. Step 2 and 3 are repeated until a certain stopping criterion is met.

147   For our experiment, the stopping criterion corresponds to 10 iterations of steps 2 and 3.
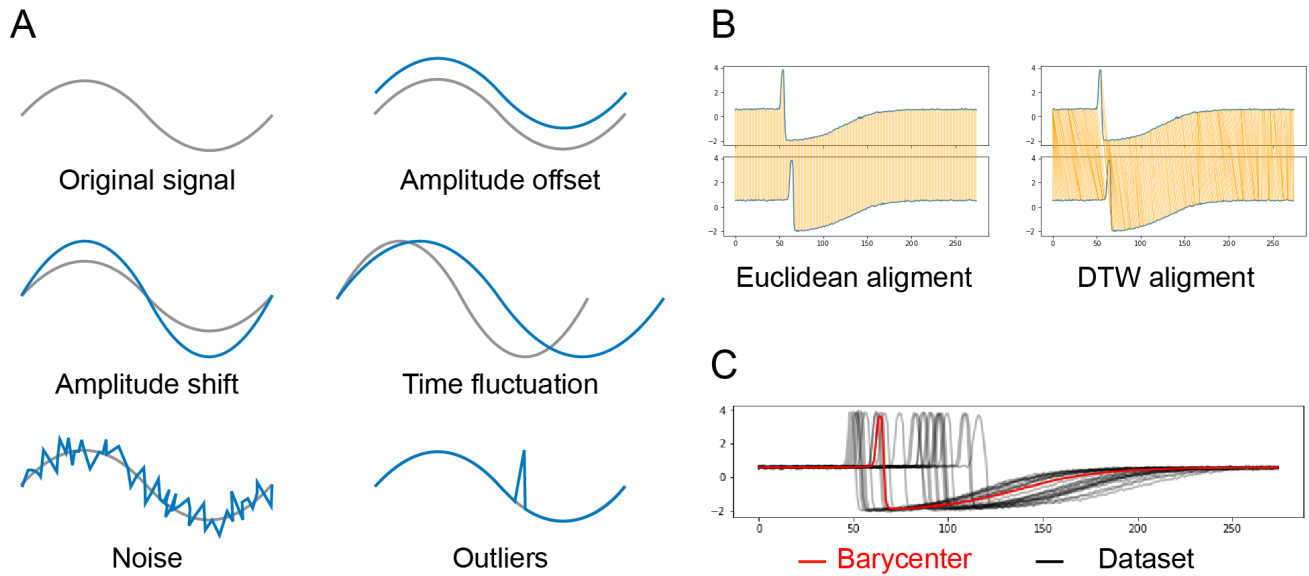
Figure 3. A: Undesirable sequence distortions. The grey line is the original signal and the blue line is the distorted signal. **B**: Difference between Euclidean alignment and DTW alignment. The compared sequences are in blue, and the orange lines represent the point-wise matching between the two sequences in the Euclidean case and the DTW case. **C**: Representation of a barycenter computed through our method BS-DBA. The learned barycenter is in red, and the observations are in black.

2.4.2 Pre-processing

During the pre-processing step, all sequences are first centered to zero mean and scaled to unit variance (z-normalization):

$$\tilde{x}_t = \frac{x_t - \bar{x}}{\sigma_x} \tag{2}$$

where $\bar{x}$ and σ_x are respectively the average and the standard deviation of **x**. The pre-processing step allows being invariant to amplitude offset and amplitude shift.

2.4.3 Dynamic Time Warping

At each iteration, the K-Means algorithm assigns each sequence to the nearest centroid. The distance is computed using DTW (Berndt and Clifford, 1994). DTW is commonly used in times-series data-mining (Esling and Agon, 2012; Fu, 2011). Intuitively, DTW considers as very similar (the distance is close to zero) two sequences of a given phenomenon occurring at different speeds. This property is particularly interesting for our problem since some mice may breathe in or exhale faster than others. To do this, DTW finds an optimal match between a query sequence and a referent sequence by locally stretching or contracting the time axis of the query sequence. The DTW measure produces the squared Euclidean distance between the aligned time series.

This measure is invariant to temporal fluctuation and can compare sequences of different duration. Considering two sequences $x \in \mathbb{R}^m$ and $y \in \mathbb{R}^n$, the computation of the DTW measure is done in $O(mn)$ in time and space using dynamic programming. In its original form, the DTW measure is sensitive to noise and outliers. Such distortions can lead to pathological alignments with unrealistic time dilations. To avoid

166 such alignments, we use the Sakoe-Chiba constraints which impose that the dilations are smaller than a
167 given duration (Sakoe and Chiba, 1978).

168 Figure 3-B shows the difference between the linear mapping of the Euclidean distance and the nonlinear
169 mapping of the DTW distance. A mathematical definition of DTW is given in Appendix 2.

170 2.4.4 Time-series averaging

171 Finding an average sequence is an important sub-routine of K-Means algorithm. Indeed, the quality of
172 each cluster is highly dependent on the quality of its centroid (Aghabozorgi et al., 2015). At each iteration,
173 all sequences in the data set $\mathcal{X}$ are assigned to their closest centroids $\{\mathbf{r}^{(1)}, \ldots, \mathbf{r}^{(K)}\}$. Then, each centroid
174 is updated by computing the average sequence based on the new assigment.

175 For any set of sequences $\mathcal{X}' \subset \mathcal{X}$, the average sequence, with respect to the DTW, is the solution of the
176 following optimization problem:

$$\arg\min_{\mathbf{y} \in \mathbb{R}^L} \sum_{\mathbf{x}' \in \mathcal{X}'} \text{DTW}^2(\mathbf{y}, \mathbf{x}') \tag{3}$$

177 where $L > 0$ is the average duration of the sequences in $\mathcal{X}'$.

178 Accurately and efficiently solving Expression 3 is not trivial (Niennattrakul and Ratanamahatana, 2007;
179 Jain, 2019). Traditional averaging methods cannot deal with the non-linear mapping between sequences of
180 potentially different duration and several algorithms have tried to solve this issue (Petitjean et al., 2011;
181 Morel et al., 2018). A recent work (Schultz and Jain, 2018) uses the subdifferentiability property of the
182 optimization function to develop a stochastic subgradient descent algorithm (S-DBA). For a trade-off
183 between accuracy and speed, we implemented a batch version of S-DBA called BS-DBA. Figure 3-C
184 illustrates the result of averaging a time series data set using BS-DBA. Appendix 3 provides details on the
185 subdifferentiability of the optimization function and the implementation of BS-DBA.

186 2.5 Characterization and symbolization of recordings

187 From a recording $\mathbf{s}'$, we first perform the segmentation process described in Section 2.3 in order to extract
188 the inspiration/expiration sequences. Then, we use a 1-NN (nearest neighbor) algorithm to assign each
189 sequence to the reference sequence, which is the closest to it, in the sense of the DTW measure.

190 To avoid incoherent symbols, some inspiration/expiration sequences are treated as outliers if their distance
191 to their reference sequence is higher than a threshold. The threshold is different for each reference sequence.
192 It corresponds to the α-quantile of the distance distribution observed within the reference sequence cluster
193 during the learning step. By default we choose the threshold value $\alpha = 0.95$.

194 This procedure yields a symbolic representation of $\mathbf{s}'$, where each respiratory cycle is replaced by a
195 symbol composed of a letter (which specifies the type of inspiration) and a number (which specifies the
196 type of expiration).

197 2.6 Connection with ventilation pattern descriptors

198 Most ventilation pattern descriptors are computed with algebraic formulas based on a cycle segmentation
199 of airflow using IOX2 software from emka TECHNOLOGIES (Mailhot-Larouche et al., 2018). The
200 algorithm to extract inspiration and expiration sequences presented in section 2.3 can be used as a
201 preprocessing step to compute such descriptors with more precision.

202 In the present work and for the purpose of validation, we have used four descriptors:

203 • **Inspiratory/Expiratory Time (Ti/Te, s):** Duration of inspiration/expiration.

204 • **Nasal Inspiratory/Expiratory Volume (NIV/NEV, ml):** Volume of air in/out during
205   inspiration/expiration.

3 DATA AND EXPERIMENT

3.1 Data Origin

207   We applied our methodology to a subset of data from experiments that aimed to understand and evaluate
208 how cholinesterase (ChE) inhibitors affect mice respiration with partial deficit in AChE (Nervo et al., 2019).
209 Acetylcholine (ACh) is a well-known neurotransmitter in the central and peripheral nervous systems. It is
210 also found at the neuromuscular junction (NMJ). ACh in synapses is hydrolyzed by acetylcholinesterase
211 (AChE). ACh is also used by numerous non-neuronal cells to communicate (Grando et al., 2015). Inhibition
212 of ChE changes the dynamic of ACh and thus may modify respiration at different physiological levels.
213 To better understand the mechanisms, we have recorded the nasal and thoracic airflow of mice with
214 different partial AChE deficits induced by injection of physostigmine, an inhibitor of ACHe, using a Double
215 Chamber Plethysmograph (DCP), (Mailhot-Larouche et al., 2018). We recorded nasal and thoracic airflow
216 from control mice (WT mice), PRiMA KO mice (PRiMA mice: AChE deficiency in cholinergic neurons of
217 the brain and peripheral nervous systems (autonomic and enteric)) (Farar et al., 2012)); muscle KO mice
218 (AChE1iRR; absence of AChE in skeletal muscles); ColQ KO mice (ColQ mice: no AChE anchoring in
219 muscles and some tissues) (Bernard et al., 2011).

220   As described in (Nervo et al., 2019), mice of different genotypes were exposed as follows:

221 1. Phase 1: The mouse is placed in a DCP for 15 or 20 minutes to serve as an internal control.

222 2. Phase 2: The mouse is removed from the DCP and injected with physostigmine.

223 3. Phase 3: The mouse is placed back into the DCP, and its nasal flow is recorded for 35 or 40 minutes.

3.2 Experiment

225   In order to test our approach, we have run and evaluated the results of the following experiment:

226 1. Creation of a data set.

227 2. Extraction of training data set for inspiration/expiration.

228 3. Computation of inspiration/expiration referent sequences.

229 4. Symbolization of all signals in the data set.

230   **Step 1: Creation of a data set.** Our data set includes the nasal airflow recording of 32 different mice.
231 Among all recordings available from Nervo et al. (2019), we have selected 8 mice for each genotype:
232 WT, PRiMA, AChE1iRR, ColQ. All mice were exposed to the same inhibitor: physostigmine. All signals
233 were recorded at 2,000Hz and have been down-sampled to 250Hz. By default, the double chamber
234 plethysmograph includes a bandpass filter, whose band limits are 0.250Hz and 35000Hz, which has not
235 been modified.

236   **Step 2: Extraction of training data set for inspiration/expiration.** On average, a mouse's respiratory
237 cycle lasts about 0.3 seconds. The original data set contains approximately 350,000 cycles, and therefore
238 computing reference sequences (section 2.4) from the entire data set would have been time-consuming.
239 Thus, for each recording, we extracted 1800 cycles that were evenly selected in time. This subsampling

corresponded to approximately 36 cycles per minute, resulting in a set of 57,600 cycles that were divided into an inspiration training data set and an expiration training data set.

Step 3: Computation of inspiration/expiration reference sequences. Referent sequences were computed according to the algorithm presented in Section 2.4. The hyperparameters are presented in the following section. The learning is based on the inspiration/expiration training data set.

Step 4: Symbolization of all signals in the data set. All recordings in the original data set are symbolized. The symbolization is based on the reference sequences learned from the training data sets.

3.3 Hyperparameters

The main parameters are presented below. Parameters for respiratory cycle detection have been set based on physician knowledge of the typical respiratory cycles. For the clustering algorithm, the number of clusters has been set arbitrarily and the Sakoe Chiba radius authorizes small dilatation.

- **Respiratory cycle detection (Step 1)**:
  - Prominence : 0.03 ml
  - Window length : 2 s
  - Minimum inspiration/expiration duration : 0.05 s
  - Maximum inspiration/expiration duration : 2 s
- **Clustering algorithm (Step 2, identical settings for inspiration and expiration):**
  - Number of clusters: 5
  - Number of iterations for K-Means: 10
  - Sakoe Chiba radius: 0.01 s
  - Reference sequence length: 0.2 s
- **Symbolization (step 3):**
  - Quantile threshold: 0.95

A python implementation of the method is available on Github [1].

4 RESULTS

In this section, we summarized the complete pipeline of our method. It is composed of three main steps:

1. The first step consists in detecting of the respiratory cycles and extracting the inspiration/expiration sequences from the input data, Figure 2-A.
2. The second step consists in computing a small number of reference sequences from the sets of inspiration/expiration sequences. The reference sequences represent groups of sequences with common properties to highlight typical inspiration/expiration behaviors, Figure 2-B.
3. The third step consists in simplifying a recording using a symbolization based on the reference sequences extracted in Step 2, Figure 2-C.
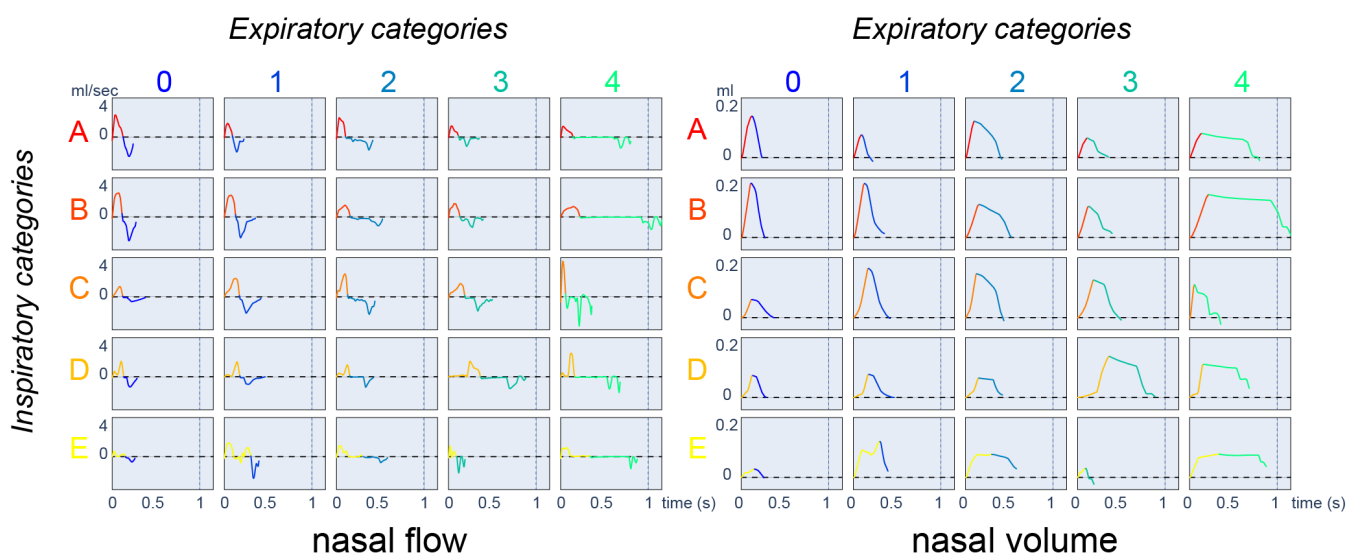
Figure 4. Respiratory cycle map displays with nasal airflow ($ml.s^{-1}$) on the left and nasal volume (ml) on the right. Positive flow corresponds to inspiration and negative flow corresponds to expiration.

4.1 Categorization of the respiratory cycles

We first aim to categorize breathing cycles, inspirations, and expirations. The limits of inspiration and expiration are unambiguously defined from the volume obtained by integrating the flow (Vijayaraghavan et al., 1993). We define a referent cycle as the association between a referent inspiration and a referent expiration. Considering K_1 referent inspirations and K_2 referent expirations, there exist $K_1 K_2$ referent cycles. In order to compare them, we develop a map where each row corresponds to a referent inspiration and each column to a referent expiration. Each referent cycle is represented by an actual cycle selected as follows:

- Among the identically labeled cycles in the training database, we select the cycle whose cumulative DTW distance (DTW distance to the referent inspiration + DTW distance to the referent expiration) is the smallest.

- The respiratory cycle map can be displayed using either the nasal airflow or the nasal volume. In any case, the inspiration/expiration phases are matched, accordingly, to their attributed colors. For inspiration, the color scale goes from red to yellow; for expiration, it goes from blue to green.

- Inspiration/expiration referent sequences are ordered in increasing order according to the average duration observed in each group. Therefore, as the number/letter increases, the average inspiration/expiration duration is longer. Visually, lighter colors (yellow/green) correspond to longer duration.

In our experiment, we set the number of inspiration and expiration referent sequences to 5, as presented in Figure 4. Short duration cycles (`A0`, `A1`, `B0`, `B1`) are characterized by a nasal airflow of sinusoidal shape. All 25 of the resulting classes are used in the following sections to visualize and compare the respiratory cycles of mice of different strains before and after physostigmine injection.
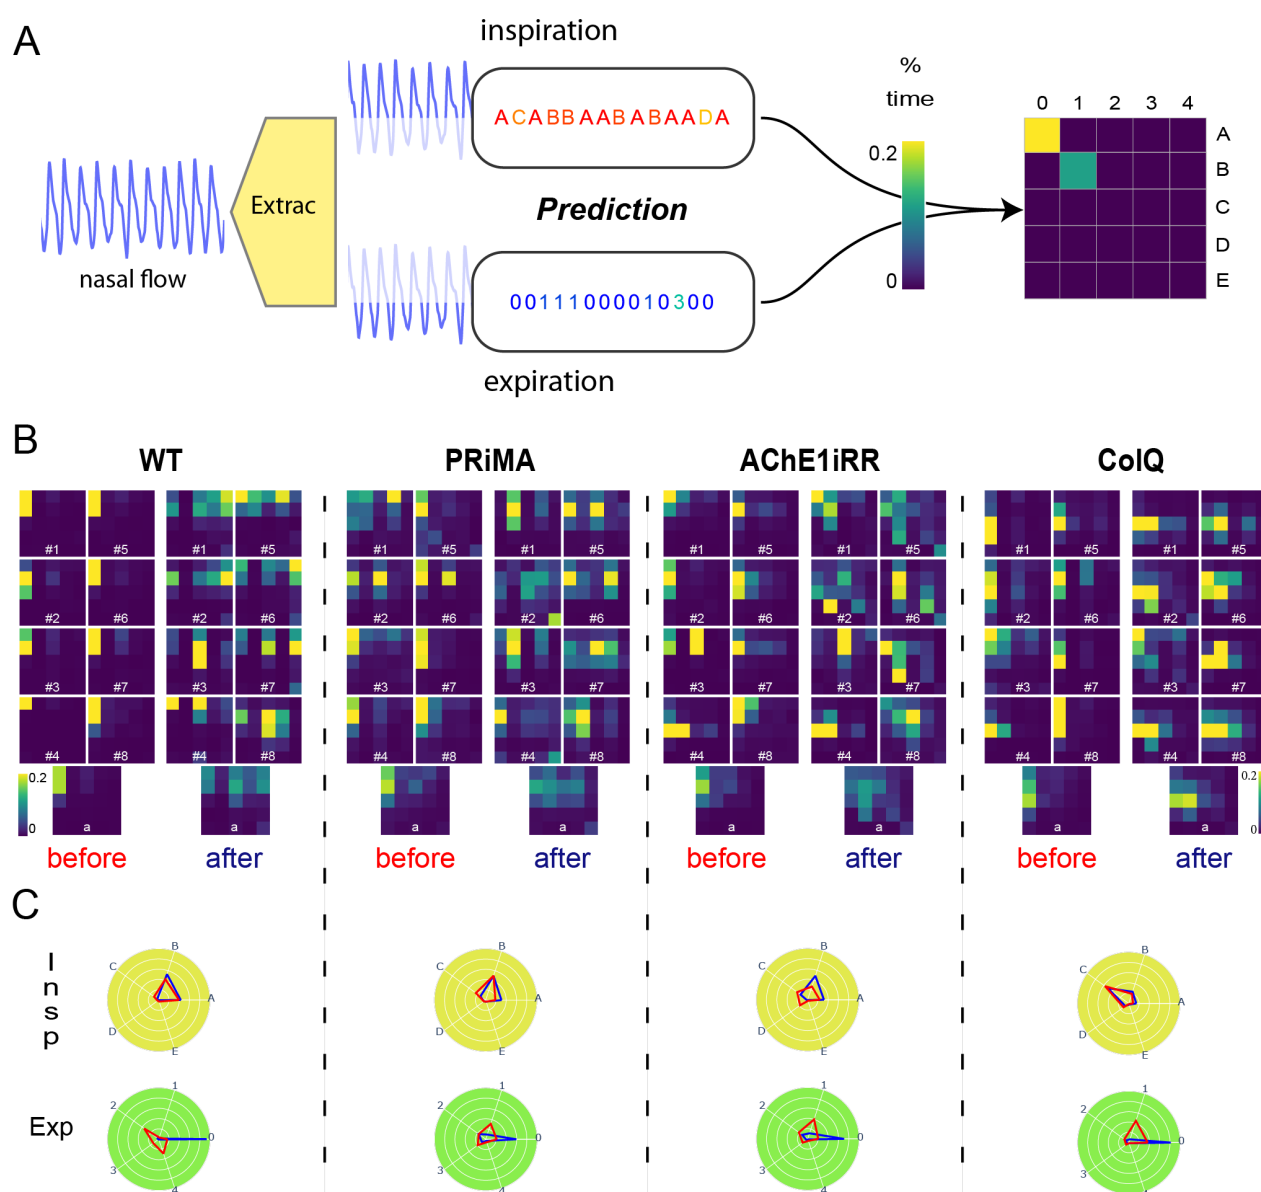
Figure 5. **(A)**: Respiratory Cycle map (RC map) built-up process. **(B)**: Respiratory RC maps: All RC maps are truncated at the threshold value of 20%. RC maps are grouped by genotype: WT, PRIMA, AChE1iRR, ColQ. For each genotype, the two left columns and the two right collumns gathered RC maps respectively before and after physostigmine injection. Numbers on RC maps correspond to the mouse id. The bottom line corresponds to the average RC maps observed per genotype before and after drug injection. **(C)**: Average reference sequence polar plots: Polar plots are grouped by genotype. Inspirations are on the top, and expirations are on the bottom. The values on each angular axis correspond to the average percentage of time assigned to the associated reference sequence. The blue polygon corresponds to the values observed before injection, and the red polygon corresponds to the values observed after the injection.

4.2 Distribution of respiratory cycle categories

In order to study the importance of each reference cycle for a given symbolic representation, we introduce a new visualization of the histogram that takes the form of a heat map. The respiratory cycle map (RC map) corresponds to a heat map where rows are inspiration symbols and columns are expiration symbols as presented in Figure 5-A. Thus, each cell corresponds to a referent cycle, and its value is set to the percentage of time assigned to that specific referent cycle. To ease the study of less frequent referent cycles, we use a thresholded version of the respiratory RC map where all reference sequences that represent more than 20 % of the total duration are assigned to the threshold value of 20 %. A RC map provides a quick understanding of the dominant respiratory behavior of a mouse. In addition, RC maps can be aggregated over a population, allowing comparisons of a mouse's respiratory behavior to the average behavior.

In Figure 5-B, RC maps are grouped by genotype: WT, PRIMA, AChE1iRR, ColQ. For each genotype, the two left columns gathered RC maps before injection, and the two left columns gathered RC maps after injection. The bottom line corresponds to the average RC maps observed per genotype before and after drug injection.

In addition, we have created two conjoint polar plots, one for inspiration and one for expiration. Each angular axis corresponds to a referent sequence, and the value on each axis is equal to the percentage of time assigned to that specific referent sequence. These values are linked together to form a polygon. As for RC maps, the visualization can be done at the individual level or aggregated over a group of mice. This representation complements RC maps as it decorrelates inspiration from expiration, easing the study of both mechanisms independently as presented in Figure 5-C.

4.3 Time line representation of respiratory cycle categories (bar codes)

Previous representations give an overview of the respiratory behavior of a mouse or a population. Nonetheless, they do not offer insights into the temporal evolution of a mouse's respiratory behavior when facing a stressor. This evolution can be read from the symbolic representation with proper visualization.

To that aim, we construct a respiratory bar code for each mouse that includes the time information, as presented in Figure 6-B. The respiratory bar code is composed of two lines, the upper line represents the inspirations, and the lower line representing the expirations. The central white area corresponds to the period of inhibitor injection, and the light grey area corresponds to unpredictable cycles. Each line is composed of rectangles whose color refers to the associated reference sequence and whose length is proportional to the duration of the associated respiratory cycle.

Figure 6-B presents respiratory bar codes of all mice in the data set. They are gathered by genotype, and mouse identification numbers are on the left of the bar codes. For each genotype, the left section corresponds to bar codes before injection and the right section to bar codes after injection.

4.4 Statistical analysis of respiratory cycle categories

RC maps provide visual comprehension of the heterogeneity in breathing behaviors and changes due to the presence of a stressor. In complement to the visual presentation, we provide a statistical analysis that compares the breathing behaviors between genotypes and the breathing responses to the presence of a stressor.

[1] `https://github.com/thibaut-germain/DCP_Clustering`
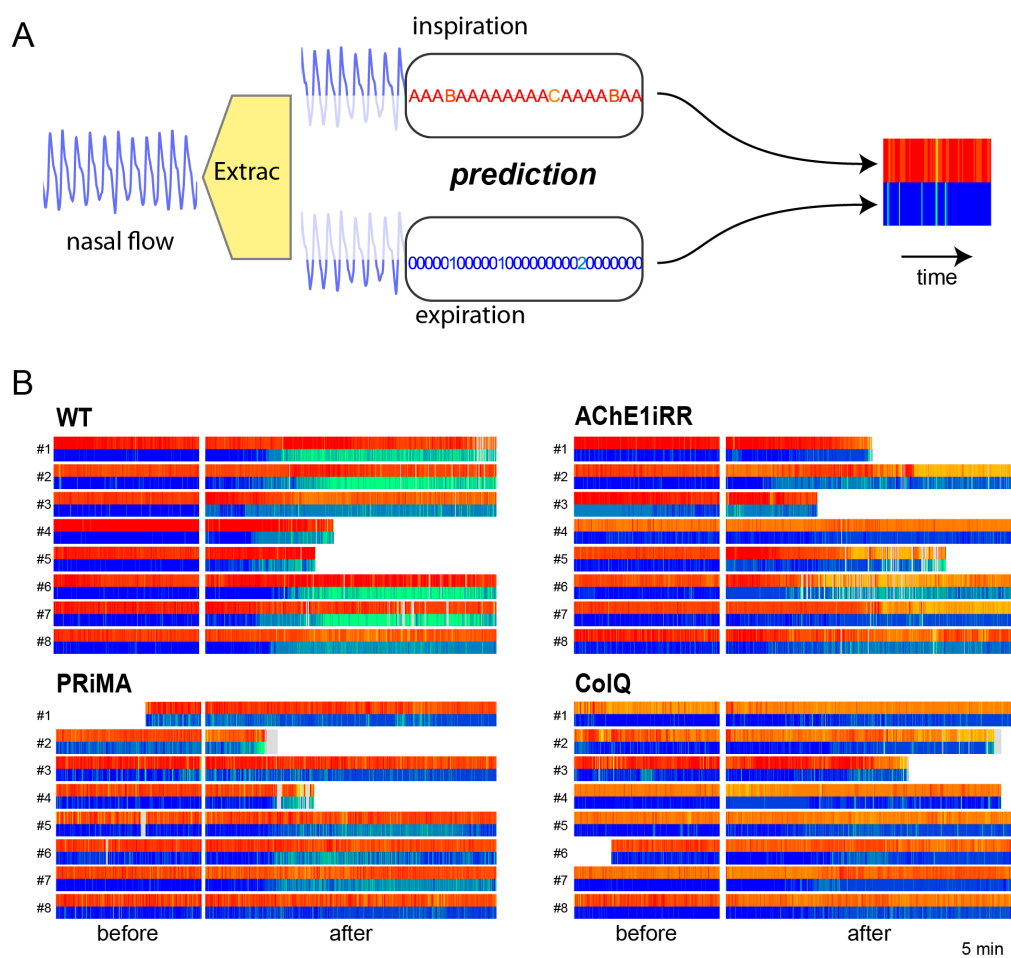
Figure 6. Respiratory bar codes: Respiratory bar codes are gathered by genotype: (top,left): WT, (bottom,left): PRIMA, (top,right): AChE1iRR, (bottom,right): ColQ. Numbers to the left of bar codes correspond to the mouse id. For each genotype, the left section corresponds to barcodes before drug injection and the right section to bar codes after injection. Grey areas in bar codes like mouse PRIMA-2 correspond to unpredictable cycles. Some experiments were shorter than others resulting in shorter bar codes.

332 The first statistical test compares the respiratory cycle distribution of AChE-deficient mice (PRIMA,
333 AChE1iRR, ColQ) with that of control mice (WT). The null hypothesis is that the cohort of AChE-deficient
334 mice has the same respiratory cycle distribution as the cohort of control mice. The alternative hypothesis is
335 different respiratory cycle distributions.

336 The second statistical test compares the distribution of respiratory cycles for each genotype before and
337 after drug injection. For the cohort of a given genotype, the null hypothesis is to have the same distribution
338 of respiratory cycles before and after drug injection. The alternative hypothesis is different respiratory
339 cycle distributions.

340 In both cases, we implemented a multiple testing scheme with a false discovery rate (FDR) correction of
341 5%, performing a Mann-Whitney U test for each type of respiratory cycle. Application of this test gives
342 a map where each cell represents a type of respiratory cycle, with the row corresponding to the type of
343 inspiration and the column to the type of expiration. A cell is colored black if the unit null hypothesis is
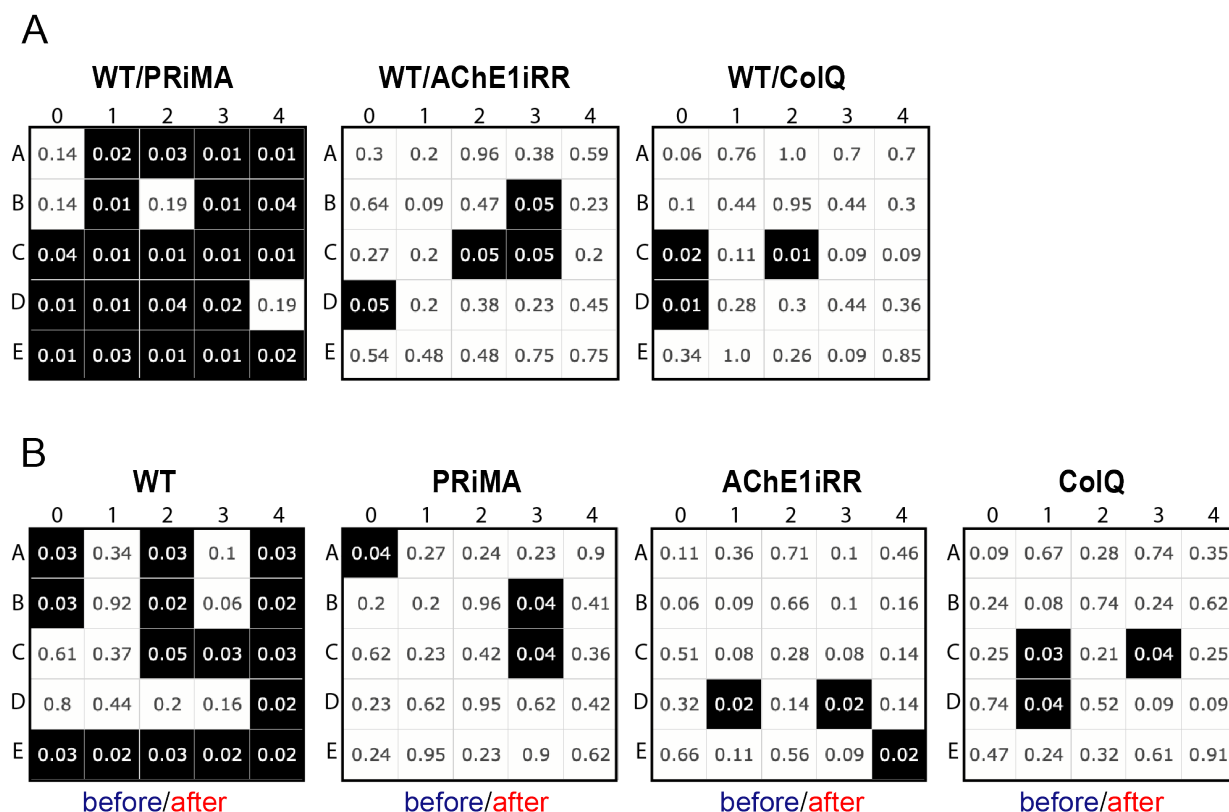
Figure 7. Multiple testing scheme with a false discovery rate (FDR) correction of 5%, performing a Mann-Whitney U test for each type of respiratory cycle. A cell is colored black if the unit null hypothesis is rejected after FDR correction and includes the corrected p-value of the associated unit test. **(A):** Statistical tests comparing the distribution of respiratory cycles of control (WT) and AChE-deficient (PRIMA, AChE1iRR, COLQ) mice before drug injection. **(B):** Statistical tests comparing the distribution of respiratory cycles before and after drug injection for each genotype.

rejected after FDR correction. In each cell, we also displayed the corrected p-value of the associated unit test.

All tests are rejected, Figure 7-A, and the number of unit tests rejected at 5% is for WT vs. PRIMA: 21, WT vs. AChE1iRR: 4, WT vs. COLQ: 3. Similarly, all tests are rejected, Figure 7-B, and the number of unit tests rejected at 5% is for WT: 15, PRiMA: 3, AChE1iRR: 3, COLQ: 3.

5 DISCUSSION

This paper presents a new method to compare and quantify cyclic signals that may be particularly appropriate for biological investigations, such as respiratory signals. Rather than comparing cycles based on the ventilation descriptors, cycles' shapes are compared to shape representations of most typical cycles. We will discuss the contributions and limitations of this new strategy by analyzing a part of recordings previously published (Nervo et al., 2019).

5.1 Inspiration and expiration classes fit respiratory physiological control

The classes learned with the new approach represent various respiratory profiles that carry biological meaning. We illustrate some respiratory profiles through their classes in Figure 8. The last 15 minutes before physostigmine injection represents mice's baseline breathing behaviors. The control mice (WT)
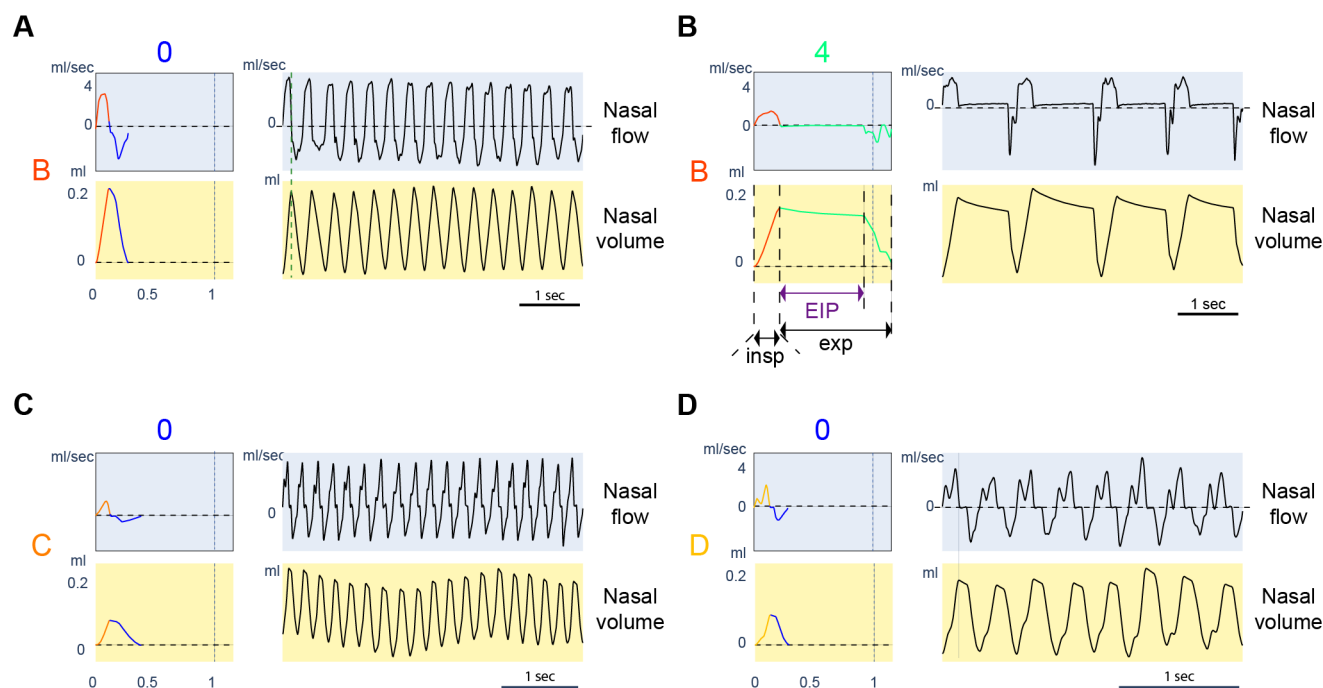
Figure 8. Examples of typical respiratory behaviors. For each panel, the left column represents the referent cycle, and the right column is an extract from a recording of up to 5 seconds where the reference cycle is repeated continuously. Charts with a blue background are expressed in nasal airflow, and charts with a yellow background are expressed in nasal volume. **(A)**: Referent cycle B0. **(B)**: Referent cycle B4. Inspiration, expiration, and end-inspiratory pause (EIP) duration are illustrated. **(C)**: Referent cycle C0. **(D)**: Referent cycle D0.

358 breathe with cycles of type A0 and B0. Figure 8-A shows 5 consecutive seconds of a raw signal with
359 respiratory cycles of B0. After injection of physostigmine, the inspiration classes (A and B) are not changed
360 for the control mice (WT), as shown with the polar plot (Figure 5-C). However, the expiration class
361 changes from type 0 to type 2, 3, 4. Raw signals of 5 consecutive seconds of classes B4 are presented in
362 Figure 8-B. The profile of these classes shows a long pause when the lungs are inflated. They correspond to
363 post-inspiratory pauses. They were analyzed in Nervo et al. (2019), and the authors quantified the duration
364 of these pauses. The new approach captures significant respiratory behaviors making previous results
365 apparent with the new representation: for control mice (WT), post-inspiratory pauses appear after inhibitor
366 injection.

367 The approach also presents details about the inspiration dynamic of ColQ mice. Indeed, the cycles
368 of ColQ mice before injection are grouped into types C0 and D0, which we present in Figure 8-(C,D).
369 Inspiratory classes C and D are characterized by a nasal airflow that enters in two phases. The two phases in
370 class D are distinctive. Compared to D, the separation between phases is less visible in C. The ColQ mouse
371 is a model of congenital myasthenic syndrome with AChE deficit at the neuromuscular junctions. This
372 mouse shows an impairment of motor control, which could be reflected during the motor control required
373 for a smooth inspiration.

374 Bar codes (Figure 5-B) also validate inspiration and expiration classes. A bar code represents the
375 symbolization of a raw signal as a timeline where inspirations and expirations are colored accordingly
376 to their classes. Bar codes reveal the dynamics of respiratory behaviors and their changes. For example,
377 inspiration classes for control mice (WT) after physostigmine injection are almost unchanged. On the
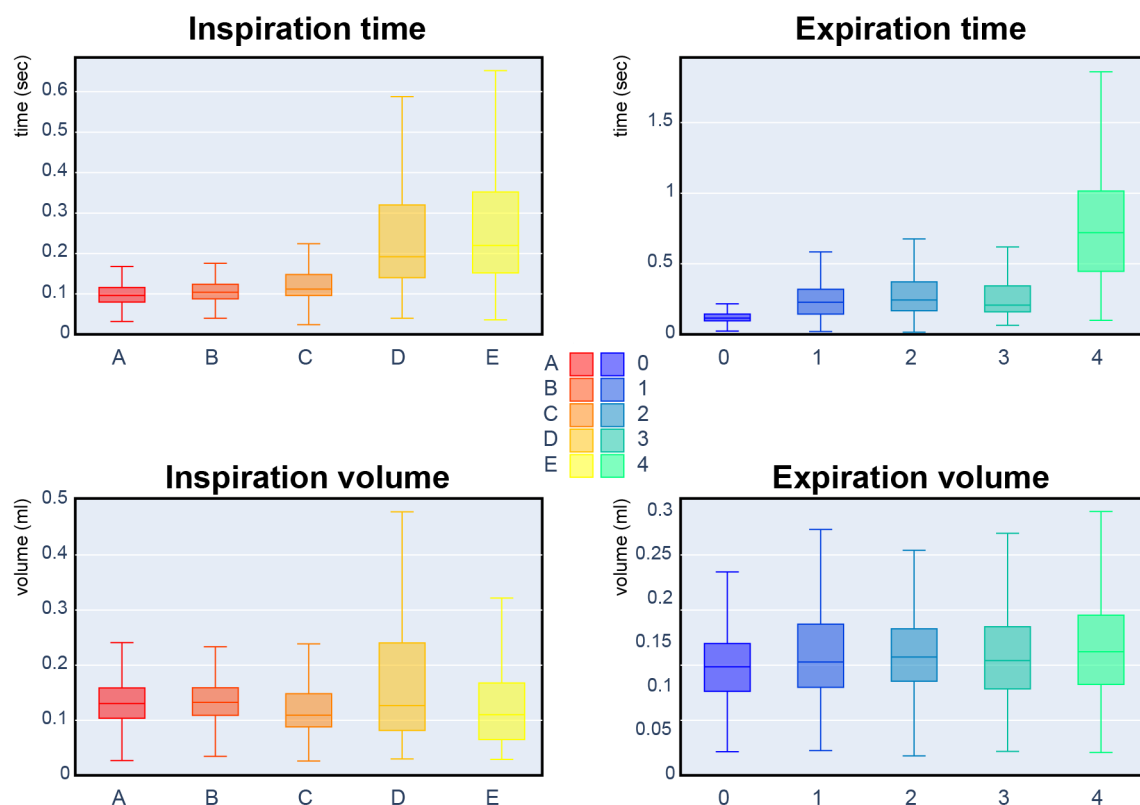
Figure 9. Box plots of the respiratory cycle descriptors: inspiration/expiration time and inspiration/expiration volume. Each box plot represents a referent sequence. A box represents the first quartile (Q1), median, and third quartile (Q3). The lower whisker corresponds to the minimum value observed, and the upper whisker is above the third quartile by 1.5 interquartile range (IQR: Q3-Q1).

contrary, their expiration classes change significantly after a latent period. This dynamic is consistent with results in Nervo et al. (2019) where the mean frequency per minute of respiratory cycle decreases after the injection of physostigmine for control mice (WT). The frequency decrease corresponds to an increase in the duration of the post-inspiration pauses per min. Through the bar codes, it is possible to visualize the appearance of expiration classes 3 and 4 after injection with remarkable precision.

The inspiration and expiration classes have been constructed without prior knowledge of mice's breathing behaviors. Nonetheless, the classes present differences that can be interpreted in terms of physiological modifications. For instance, some of the expiration classes represent post-inspiratory pauses. New inspiration classes have also been described, probably related to the motor controls dynamics during the active ventilation phase.

5.2 Classes reveal heterogeneity: an observation masked by the averaging of algebraic descriptors

Analyses on a small cohort can be biased if individual responses are heterogeneous. Unfortunately, it is often difficult to recognize this heterogeneity through some descriptors. The new symbolization, based on typical inspiration/expiration, the visualization and the quantification tools we proposed, offer perspectives on this critical issue in biology. For example, it is apparent on individual RC maps and bar codes that control mice (WT) present homogeneous respiration; the respiration cycle types are A0 and B0. After

395 injection of AChE inhibitor, the RC maps and bar codes of control mice WT-1,2,6,7 show that they follow
396 the same evolutionary dynamics. Nevertheless, mice WT-3,8 present different dynamics, and mice WT-4,5
397 died during the experiment. Thus, we can conclude that mice adapt differently to cholinesterase inhibition
398 by physostigmine. In addition, the tests highlight changes that are significantly different.

399   We proposed in Boudinot et al. (2009) and Nervo et al. (2019) that mice with partial AChE deficiency
400 were remarkably adapted to AChE deficit in the brain, autonomic nervous systems, and muscles. Indeed,
401 the most frequent respiratory cycles before injection are composed with the inspiration of type A, B, C and
402 the expiration of kind 0, 1, 2. Looking at Figure 9, these reference sequences share similar duration and
403 volume. Therefore, it is impossible to differentiate the genotypes based on inspiration/expiration duration
404 or volume.

405   The present study shows that the distributions of inspiration and expiration classes on AChE1iRR mice
406 are similar. AChE1iRR mice do not have AChE in skeletal muscle. These mice show a high homogeneity
407 of adaptation despite muscle weakness. In contrast, PRiMA mice, which have AChE deficiency in the brain
408 and autonomic nervous systems, adapted well to AChE inhibition, but showed heterogeneous respiratory
409 behavior. The heterogeneity is apparent in inspiration and expiration classes, which suggests the possibility
410 of different respiratory behaviors to cope with AChE deficit in the nervous system. The cohort of ColQ
411 mice also presents heterogeneity in respiratory behaviors, specifically for inspiration. As discussed, the
412 inspiration of ColQ mice is characterized by types C and D. In contrast, the inspiration of other genotypes
413 is characterized by types A and B. While ColQ and AChE1iRR mice have similar AChE deficiency in
414 neuromuscular junctions, AChE1iRR mice adapt better than ColQ mice which also have AChE deficit in
415 other tissues. This result suggests that AChE deficit in skeletal muscle is insufficient to affect these mice's
416 inspiration.

417   If the respiratory adaptations are different, it is not surprising that the consequences of the injection of
418 physostigmine are so variable. Visualization of inspiration and expiration classes, either in RC maps or bar
419 code, makes it possible to account for this diversity. After injection of physostigmine, the changes tend to
420 affect inspiration in AChE1iRR and ColQ mice, whereas expiration is more affected in WT and PRiMA
421 mice.

422   In summary, representing respiratory cycles by classes sharing similar shapes reveals a diversity of
423 unsuspected respiratory behaviors that were not identifiable with descriptors deduced from the airflow. This
424 rich information is synthesized in graphical representations highlighting how mice respond differently to
425 cholinesterase deficits or inhibition.

5.3   Inspiration and expiration classes evoke distinct biological processes.

427   Inspiration and expiration classes are defined without prior knowledge of underlying biological processes.
428 Inspiration classes A and B represent a regular inspiration phase, while classes C and D represent an
429 inspiration phase with a more or less significant pause. The pauses in category C are very short and always
430 during inspiration; they probably correspond to a motor impairment during lung inflation (the main action
431 of the diaphragm, a powerful muscle) or by a fine control of the glottis. The longer pauses of class D may
432 occur during the air inflow and are probably similar, in nature, to class C. In contrast, the long pauses of
433 Class E correspond to a sort of pause before the air enters the lungs. From a physiological point of view,
434 these pauses could correspond to a delay in the glottis's active opening, which is required to allow air to
435 enter into the trachea. Two situations can lead to the glottis remaining closed: the cessation of muscle
436 contractions that control the glottis opening or the spasm (cramp) of the muscles that control the closing
437 of the glottis. Expiration class 0 represents a regular and probably passive phase of expiration. Classes

438 `2`,`3` and `4` start with a post-inspiratory pause whose duration increases progressively from category 2
439 to category 4. These post-inspiratory pauses are well described in the literature and appear in different
440 physiological conditions. They appear when it is necessary to increase the air pressure in the lungs (short
441 pauses) or as reflexes (long pauses), such as those resulting from inhaling molecules that irritate the upper
442 airways (Dutschmann et al., 2014).

443 From these results we can conclude that inspiration and expiration classes learned from a subset of
444 recordings selected from (Nervo et al., 2019) carry interpretable physiological meaning. It is important
445 to note that these classes are specific to the experiment. For instance, applying our method to a set of
446 signals presenting bronchoconstrictions will likely lead to classes differentiating the severity/variety of
447 constrictions in a finer way than using the EF50 metric (Glaab and Braun, 2021).

5.4 Limits and future work

449 Our approach to analyzing respiratory signals is based on learning typical inspirations and expirations,
450 called reference sequences. Currently, the number of referent sequences is arbitrarily set by the user. By
451 doing so, the user chooses the degree of detail incorporated in the symbolization: adding referent sequences
452 divides typical breathing behaviors into subgroups with minor variations. In that manner, reference
453 sequences carry meaningful physiological information for the user. Nevertheless, choosing a good number
454 of reference sequences can be complicated and time-consuming without knowledge of respiratory behavior.
455 In such cases, several heuristics based on mathematical criteria exist to define the number of reference
456 sequences automatically (Kodinariya and Makwana, 2013). In any case, these heuristics can be used as a
457 starting point to properly define the number of clusters in light of the experiment objective.

458 In the current work, we limited ourselves to static descriptors of reference sequences (RC map, polar plot)
459 and visual interpretation of the breathing behavior evolution over time (bar plot). Nevertheless, breathing
460 behavior dynamics can also be quantified using the proposed symbolization of recordings and applying
461 symbolic dynamics theory (Lind and Marcus, 2021; Morse and Hedlund, 1938). Symbolic dynamic theory
462 has been developed to study how a system's configurations change over time and how similar initial states
463 can grow dissimilar.

464 During the experiment, we only symbolized recordings included in the training data set. By doing so, we
465 guaranteed that the most common behaviors present in the recordings were taken into account during the
466 learning step. We do not recommend symbolizing on other recordings as some typical behavior might be
467 negelcted. In future work, we would like to investigate the use of a hierarchical clustering algorithm on a
468 large data set composed of recordings with various experimental conditions. By doing so, we would like to
469 create a universal referential of typical behaviors usable across experiments that can adapt to the level of
470 detail required by selecting a symbolization directly from the hierarchy.

5.5 Prospective use

472 In this work, we have limited ourselves to plethysmograph signals recorded with DCP, but the method and
473 it can easily be extended to head-out plethysmography (Bruggink et al., 2022) as well as to other biological
474 systems. Indeed, our approach relies on accurate segmentation of plethysmograph signals, allowing relevant
475 studies of inspiration and expiration. Any biological system which results in the recording of a cyclic
476 signal can use our approach with proper segmentation. For instance, in the case of electrocardiogram
477 signals, we could combine our approach with a heartbeat detection algorithm (Zong and Jiang, 2003;
478 Pan and Tompkins, 1985) to detect and represent patterns of diseases like arrhythmias, heart attacks,

479 cardiomyopathy, and coronary heart disease. Then, the symbolization of these signals could offer insightful
480 information about the underlying dynamics of such diseases.

6 ACKNOWLEDGMENTS

REFERENCES

483 Aghabozorgi, S., Shirkhorshidi, A. S., and Wah, T. Y. (2015). Time-series clustering - a decade review.
484   *Information Systems* 53, 16–38. doi:10.1016/j.is.2015.04.007
485 Arthur, D. and Vassilvitskii, S. (2007). K-means++: The advantages of careful seeding. In *Proceedings of*
486   *the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (USA: Society for Industrial and
487   Applied Mathematics), SODA '07, 1027–1035
488 Bartlett, D. and Tenney, S. (1970). Control of breathing in experimental anemia. *Respiration physiology*
489   10, 384–395
490 Bates, J. H. and Irvin, C. G. (2003). Measuring lung function in mice: the phenotyping uncertainty principle.
491   *Journal of applied physiology* 94, 1297–1306
492 Bernard, V., Girard, E., Hrabovska, A., Camp, S., Taylor, P., Plaud, B., et al. (2011). Distinct localization
493   of collagen q and prima forms of acetylcholinesterase at the neuromuscular junction. *Molecular and*
494   *Cellular Neuroscience* 46, 272–281
495 Berndt, D. J. and Clifford, J. (1994). Using dynamic time warping to find patterns in time series. In
496   *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining* (AAAI
497   Press), AAAIWS'94, 359–370
498 Boudinot, E., Bernard, V., Camp, S., Taylor, P., Champagnat, J., Krejci, E., et al. (2009). Influence of
499   differential expression of acetylcholinesterase in brain and muscle on respiration. *Respiratory physiology*
500   *& neurobiology* 165, 40–48
501 Bruggink, S., Kentch, K., Kronenfeld, J., and B. Renquist, B. (2022). A leak-free head-out plethysmography
502   system to accurately assess lung function in mice. *Journal of Applied Physiology* 133, 104–118
503 Dutschmann, M., Jones, S., Subramanian, H., Stanic, D., and Bautista, T. (2014). The physiological
504   significance of postinspiration in respiratory control. *Progress in brain research* 212, 113–130
505 Esling, P. and Agon, C. (2012). Time-series data mining. *ACM Computing Surveys* 45, 1–34. doi:10.1145/
506   2379776.2379788
507 Farar, V., Mohr, F., Legrand, M., d'Incamps, B. L., Cendelin, J., Leroy, J., et al. (2012). Near-complete
508   adaptation of the prima knockout to the lack of central acetylcholinesterase. *Journal of neurochemistry*
509   122, 1065–1080
510 Fu, T. (2011). A review on time series data mining. *Engineering Applications of Artificial Intelligence* 24,
511   164–181. doi:10.1016/j.engappai.2010.09.007
512 Glaab, T. and Braun, A. (2021). Noninvasive measurement of pulmonary function in experimental mouse
513   models of airway disease. *Lung* 199, 255–261
514 Grando, S., Kawashima, K., Kirkpatrick, C., Kummer, W., and Wessler, I. (2015). Recent progress in
515   revealing the biological and medical significance of the non-neuronal cholinergic system. *International*
516   *immunopharmacology* 29, 1–7
517 Hoymann, H. (2012). Lung function measurements in rodents in safety pharmacology studies. *Frontiers in*
518   *Pharmacology* 3, 156. doi:10.3389/fphar.2012.00156

519 Jain, B. (2019). Revisiting inaccuracies of time series averaging under dynamic time warping. *Pattern*
520   *Recognition Letters* 125, 418–424

521 Kaufman, L. and Rousseeuw, P. J. (2009). *Finding groups in data: an introduction to cluster analysis*, vol.
522   344 (John Wiley & Sons)

523 Kodinariya, T. and Makwana, P. (2013). Review on determining number of cluster in k-means clustering.
524   *International Journal* 1, 90–95

525 Lind, D. and Marcus, B. (2021). *An introduction to symbolic dynamics and coding* (Cambridge university
526   press)

527 Mailhot-Larouche, S., Deschênes, L., Lortie, K., Gazzola, M., Marsolais, D., Brunet, D., et al. (2018).
528   Assessment of respiratory function in conscious mice by double-chamber plethysmography. *Journal of*
529   *visualized experiments: JoVE*

530 Morel, M., Achard, C., Kulpa, R., and Dubuisson, S. (2018). Time-series averaging using constrained
531   dynamic time warping with tolerance. *Pattern Recognition* 74, 77–89

532 Morse, M. and Hedlund, G. (1938). Symbolic dynamics. *American Journal of Mathematics* 60, 815–866

533 Murphy, D. J. (2002). Assessment of respiratory function in safety pharmacology. *Fundamental & clinical*
534   *pharmacology* 16, 183–196

535 Nervo, A., Calas, A., Nachon, F., and Krejci, E. (2019). Respiratory failure triggered by cholinesterase
536   inhibitors may involve activation of a reflex sensory pathway by acetylcholine spillover. *Toxicology* 424,
537   152232

538 Niennattrakul, V. and Ratanamahatana, C. A. (2007). Inaccuracies of shape averaging method using
539   dynamic time warping for time series data. In *Proceedings of the International conference on*
540   *computational science* (Springer), 513–520

541 Pan, J. and Tompkins, W. J. (1985). A real-time qrs detection algorithm. *IEEE transactions on biomedical*
542   *engineering* , 230–236

543 Petitjean, F., Ketterlin, A., and Gançarski, P. (2011). A global averaging method for dynamic time warping,
544   with applications to clustering. *Pattern recognition* 44, 678–693

545 Sakoe, H. and Chiba, S. (1978). Dynamic programming algorithm optimization for spoken word recognition.
546   *IEEE transactions on acoustics, speech, and signal processing* 26, 43–49

547 Schultz, D. and Jain, B. (2018). Nonsmooth analysis and subgradient methods for averaging in dynamic
548   time warping spaces. *Pattern Recognition* 74, 340–358

549 Stone, R. (2018). How to defeat a nerve agent 359, 23

550 Sunshine, M. D. and Fuller, D. D. (2021). Automated classification of whole body plethysmography
551   waveforms to quantify breathing patterns. *Frontiers in Physiology* , 1347

552 Vijayaraghavan, R., Schaper, M., Thompson, R., Stock, M., and Alarie, Y. (1993). Characteristic
553   modifications of the breathing pattern of mice to evaluate the effects of airborne chemicals on the
554   respiratory tract. *Archives of toxicology* 67, 478–490

555 Willmann, R., Gordish-Dressman, H., Meinen, S., Rüegg, M., Yu, Q., Nagaraju, K., et al. (2017).
556   Improving reproducibility of phenotypic assessments in the dyw mouse model of laminin-α2 related
557   congenital muscular dystrophy. *Journal of neuromuscular diseases* 4, 115–126

558 Zong, W. and Jiang, G. B. M. D. (2003). A robust open-source algorithm to detect onset and duration of
559   qrs complexes. In *Computers in Cardiology, 2003* (IEEE), 737–740

1 PEAK DETECTION ALGORITHM

560 In the following, let $\mathbf{x} = (x_1, \dots x_T) \in \mathbb{R}^T$ be a univariate sequence.

561   Peak detection algorithm consists in finding all local maxima and removing those whose prominence
562 is below a prominence threshold. A data point x_t is considered as a local maximum if $x_t > x_{t-1}$ and
563 $x_t > x_{t+1}$. Considering a user-defined window size w_{len} and a local maximum x_u, its prominence is
564 computed as follows:

565 1. **Finding left prominence:** Denote s_u the last time point such as the sequence $(x_{l_u}, ..., x_{u-1})$, where
566     $l_u = \max(0, u - \lfloor w_{len}/2 \rfloor)$, intersects the horizontal line $y = x_u$. If there is no intersection, $s_u = l_u$.
567   Left prominence is defined as: $pl_u = x_u - \min(x_{s_u}, ..., x_{u-1})$.

568 2. **Finding right prominence:** Denote e_u the first time point such as the sequence $(x_{u+1}, ..., x_{r_u})$, where
569     $r_u = \min(T, u + \lfloor w_{len}/2 \rfloor)$, intersects the horizontal line $y = x_u$. If there is no intersection, $e_u = r_u$.
570   Right prominence is defined as: $pr_u = x_u - \min(x_{u+1}, ..., x_{e_u})$.

571 3. **Set prominence:** Prominence of the local maximum x_u is defined as $p_u = \max(pl_u, pr_u)$

572 Figure 10 illustrates the prominence computation. Considering a minimum prominence p_{min}, only local
573 maxima with a prominence greater than p_{min} are considered as peaks. This method is implemented in
574 Python Scipy package[2].
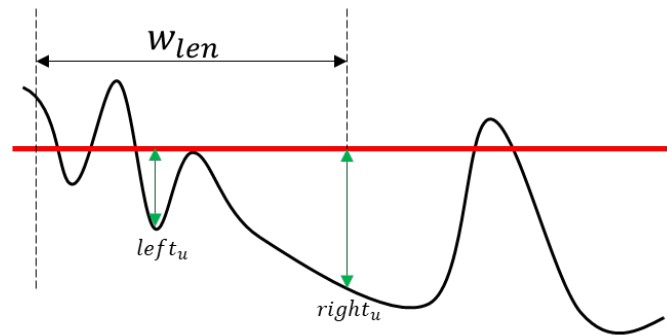
Figure 10. The prominence is $left_u$. There is no intersection between the horizontal line and the curve on the right side, the right search space is bounded by the user defined window size.

2 Dynamic Time Warping (DTW)

575 Considering two sequences $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{y} \in \mathbb{R}^n$, DTW is solution of the optimization problem:

$$\text{DTW}(\mathbf{x}, \mathbf{y}) = \min_{W \in \mathcal{W}_{mn}^{\alpha}} \sqrt{\sum_{i=1}^{m} \sum_{j=1}^{n} W_{ij} |x_i - y_j|^2} \qquad (4)$$

576 where $\mathcal{W}_{mn}^{\alpha} = \{W \in \mathcal{W}_{mn} | W_{ij} = 0 \text{ if } |i - j| > \alpha\}$ are all the warping matrix which respect Sakoe-Chiba
577 constraint parametrized by $\alpha > 0$. A warping matrix $W \in \{0, 1\}^{m \times n}$ is a monotonic non-linear mapping
578 between $\mathbf{x}$ and $\mathbf{y}$ such that x_i is mapped to y_j if $W_{ij} = 1$. W path starts from the upper-left corner (1,
579 1) and connects the lower-right corner (m, n) using only $\rightarrow, \downarrow, \searrow$ moves. A warping matrix minimizing
580 Equation 4 is called an optimal warping matrix.

---

[2] https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.find_peaks.html

3 BATCH STOCHASTIC SUBGRADIENT DTW BARYCENTER AVERAGING (BS-DBA)

581 In the following, let $\mathcal{X} = \{\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(N)}\}$ and $F_X(\mathbf{y}) = \frac{1}{N} \sum_{i=1}^{N} \mathrm{DTW}^2(\mathbf{y}, \mathbf{x}^{(i)})$ where $\mathbf{y} \in \mathbb{R}^L$. Then, the
582 subgradient of F_X at point $\mathbf{y}$ is:

$$\nabla F_X(\mathbf{y}) = \frac{2}{N} \sum_{u=1}^{N} \left(V^{(u)}\mathbf{y} - W^{(u)}\mathbf{x}^{(u)} \right) \tag{5}$$

583 where $W^{(u)} \in \mathcal{W}_{Ln}^{\alpha}$ is the optimal warping matrix between $\mathbf{y}$ and $\mathbf{x}^{(u)}$ as defined in Appendix 2 and $V^{(u)}$
584 is a diagonal matrix in $\mathbb{N}^{L \times L}$ such that:

$$V_{i,i}^{(u)} = \sum_{j=1}^{n} W_{ij}^{(u)} \tag{6}$$

585   BS-DBA is presented in Algorithm 1. If the initialization sequence $\mathbf{y}_{ini}$ is not given, it is set to a vector of
586 size L sampled from an uniform distribution on $[0, 1]$. The learning rate scheduler η is taken from Schultz
587 and Jain (2018):

$$\eta^{(t)} = \begin{cases} \eta^{(t-1)} - (\eta_0 - \eta_1)/\beta & \text{if } 1 \leq t \leq \beta \\ \eta_1 & \text{otherwise} \end{cases} \tag{7}$$

588 where $\eta_0 = \eta^{(0)} = 0.05$, $\eta_1 = 0.01$, n_b is the batch size and $\beta = \lfloor N/n_b \rfloor + 1$ is the number of iteration for
589 one epoch. The learning rate only decreases during the first epoch then it remains fix to η_1. The algorithm
590 stopping criteria is the total number of iterations.

Algorithm 1 BS-DBA

Parameters: n_{epochs} the number of epochs, L the length of the averaging time-series, n_b the size of a batch, n_{it} the number of iterations, η the learning rate scheduler

Inputs: $X = (\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(N)})$ a set of time-series, $\mathbf{y}_{ini}$ (optional) the starting sequence

Output: $\mathbf{y}^*$ the best average time-series

> **if** $\mathbf{y}_{ini}$ is given **then**
>> $\mathbf{y}^{(0)} \leftarrow \mathbf{y}_{ini}$
>
> **else**
>> *Initialize* $\mathbf{y}^{(0)} \in \mathbb{R}^L$
>
> **end if**
> *Initialize best solution* $y^* \leftarrow y^{(0)}$
> **for** $epoch = 1, \ldots, n_{epochs}$ **do**
>> *Batches* $\leftarrow$ *randomly partition* X *in batches of size* n_b
>> **for** $batch \in batches$ **do**
>>> **for** $\mathbf{x}^{(k)} \in batch$ **do**
>>>> $P^{(k)} \leftarrow$ *Optimal warping path between* $\mathbf{y}^{(t-1)}$ *and* $\mathbf{x}^{(k)}$
>>>> $W^{(k)} \leftarrow$ *Warping Matrix of* $P^{(k)}$
>>>> $V^{(k)} \leftarrow$ *Valence Matrix of* $P^{(k)}$
>>>
>>> **end for**
>>> *Update temporal solution:*

$$\mathbf{y}^{(t)} \leftarrow \mathbf{y}^{(t-1)} - \eta^{(t)} \frac{2}{n_b} \sum_{k=1}^{n_b} \left(V^{(k)} \mathbf{y} - W^{(k)} \mathbf{x}^{(k)} \right)$$

>>> *Update best solution such that:* $\mathbf{y}^* = \min\left(F_X(\mathbf{y}^*), F_X(\mathbf{y}^{(t)}) \right)$
>>> **if** $t \geq n_{it}$ **then**
>>>> **break**
>>>
>>> **end if**
>>
>> **end for**
>
> **end for**