



**HAL**  
open science

## Clustering of timed sequences – Application to the analysis of care pathways

Thomas Guyet, Pierre Pinson, Enoal Gesny

### ► To cite this version:

Thomas Guyet, Pierre Pinson, Enoal Gesny. Clustering of timed sequences – Application to the analysis of care pathways. *Data and Knowledge Engineering*, 2025, 156, pp.102401. 10.1016/j.datak.2024.102401 . hal-04849653

**HAL Id: hal-04849653**

**<https://inria.hal.science/hal-04849653v1>**

Submitted on 19 Dec 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Clustering of timed sequences – Application to the analysis of care pathways\*

Thomas Guyet<sup>†</sup>, Pierre Pinson<sup>‡</sup> and Enoal Gesny<sup>†,‡</sup>

<sup>†</sup>. AIstroSight, Inria, Hospices Civils de Lyon, Université Claude Bernard Lyon 1, 56 Bld Niels Bohr, 69603 Villeurbanne, France

thomas.guyet@inria.fr

<sup>‡</sup>. Clinical Research Unit, AP-HP, Cochin Academic Hospital, France.

## Abstract

Improving the future of healthcare starts by better understanding the current actual practices in hospital settings. This motivates the objective of discovering typical care pathways from patient data. Revealing typical care pathways can be achieved through clustering. The difficulty in clustering care pathways, represented by sequences of timestamped events, lies in defining a semantically appropriate metric and clustering algorithms.

In this article, we adapt two methods developed for time series to the clustering of timed sequences: the drop-DTW metric and the DBA approach for the construction of averaged time sequences. These methods are then applied in clustering algorithms to propose original and sound clustering algorithms for timed sequences. This approach is experimented with and evaluated on synthetic and real-world data.

## 1 Introduction

Timed sequences, composed of dated events, are types of data frequently encountered in various application domains such as computer log analysis, life course analysis in social science [8], tourism trajectories [3] or healthcare pathway analysis [18]. A specificity of this type of data is that they have no natural vector representation. In particular, they are formed by two very different dimensions: symbolic and temporal; and, on the other hand, their size generally varies from one example to another (e.g., in a healthcare pathway, some patients have more care events than others). This difficulty in modeling data in a vector space makes it challenging to adapt standard data analysis methods (classification, clustering, or forecasting).

The adaptation of data analysis methods for timed sequences raises the question of formalizing a topology of the objects to be analyzed. Most data analysis methods make the assumption that objects lie in a metric space. This topology assumption provides intuitive ways to generalize sparse and noisy observations. Formalizing such a space for timed sequences involves defining a metric between timed sequences. This problem is very similar to the one encountered in time series analysis (temporal data where the values are numerical, not event-based and generally with a regular and dense temporal sampling). For time series analysis, a large collection of metrics has been proposed and allows them to be analyzed [13, 23]. The challenges reveals to identify a metric that has the appropriate properties for the analysis task one wishes to carry out (e.g., the properties of a distance) and whose semantics will give the expected results. Intuitively, a similarity metric (resp. dissimilarity metric) should produce a high similarity value (resp. low dissimilarity value) when two examples are considered similar by domain experts.

---

\*This article has been published in the Journal of Data & Knowledge Engineering. To cite this article, please use the following reference: *Clustering of timed sequences – Application to the analysis of care pathways*, Thomas Guyet, Pierre Pinson, Enoal Gesny, *Journal of Data & Knowledge Engineering*, 156:102401 (2025) doi:10.1016/j.datak.2024.102401.

For sequential data, few studies have focused on adapting data analysis methods. [10] and [11] are based on tensor representation of sequential data, while [24] or [15] use embedding representations. Nonetheless, only a few methods handle event logs with timestamped events dates [18, 8].

The TraMineR [8] software has been frequently used for epidemiological studies, for instance by [14]. It enables the user to cluster or classify event-based temporal data with different metrics. Nonetheless, this method represents an event-based data as a state-sequence, i.e., a longitudinal sequence which associate at each instant a unique state. One advantage of this tool is to propose an end-to-end analysis of the data. Another end-to-end method for timed sequences is AliClu [18] which combines the definition of a time warping metric with a hierarchical clustering algorithm. Contrary to TraMineR, it does not allow a change of metric.

In addition to these end-to-end methods, some work was dedicated to metrics between temporal sequences [9, 21, 6] or averaging [5, 4, 22]. Despite these proposals, metrics whose semantics can better take into account the temporal dimension are still lacking. In particular, the delays between the occurrence of events are not handled. Most current solutions generally only take into account the sequentiality in the data.

In this work, we extend the drop-DTW metric [7] which was proposed to evaluate distances between sequences taking into account both sequential aspects, by adapting Dynamic Time Warping [20], and by adding support for meaningless events that can be dropped. This metric has the semantics that seem appropriate for our application case – i.e., the clustering of care pathways, but as this metric is not a distance, it does not allow for the construction of an average sequence nor to be used in classical clustering methods such as the K-Means algorithm or hierarchical clustering.

The main contribution of this article is the proposal of an algorithm to compute a mean timed sequence based on drop-DTW. This algorithm is inspired by the DBA method [17] for time series. The construction of an average sequence meets both the needs of clustering algorithms (hierarchical clustering or K-means) and the needs of interpretation of sequence clusters (definition of a representative).

This method is applied on a Electronic Health Records (EHR) dataset in order to discover typical care pathways. In the context of the OPTISOINS project<sup>1</sup>, we were particularly interested in patients who had undergone pulmonary resection surgery for oncological reasons. Care trajectories were reconstructed from data collected in the EHR data warehouse of the Greater Paris Hospital (AP-HP). The care pathway is a timed sequence representing the key stages of managing patient surgery from diagnosis to post-surgery treatments.

The reminder of this article is organized as follows. We first introduce notations related to timed sequences. Then, Section 3 introduces the adaptation of drop-DTW to timed sequences. Section 4 presents the adaptation of the DBA algorithm for averaging timed sequences based on the drop-DTW. This averaging method is required in classical clustering algorithms that are briefly recalled in Section 5. Finally, Sections 6 and 7 present experiments and results on synthetic data and our use case (OPTISOINS project) respectively.

## 2 Times Sequences and Probabilistic Timed Sequences

In this section, we introduce notations for timed sequences and probabilistic timed sequences. The probabilistic version of timed sequences enables to adapt the metrics defined for time series to timed sequences (see next section).

Let  $\Sigma$  be a finite set of  $n$  event types. A timed sequence is a sequence of pairs  $S = \langle (s_i, t_i) \rangle$  where  $s_i \in \Sigma$  is an event type and  $t_i \in \mathbb{R}$  is the event’s timestamp.  $S[i] = (s_i, t_i)$  denotes the  $i$ -th event in the sequence.  $|S|$  is the length of the timed sequences, i.e. its number of events. By convention, events in a sequence are ordered by their timestamps. Several events may occur at the same time only if their event types are different, i.e.  $(s_i, t) = (s_j, t) \Rightarrow s_i \neq s_j$ .

We propose embedding the space of timed sequences into a larger space where events are represented by their one-hot encoding, which is a vector representation. Instead of representing a sequence of events, we

<sup>1</sup><https://www.bernoulli-lab.fr/project/optisoins/>




Figure 1: Illustration of the timed sequences of Example 1 (on the left) and the cost matrix for the alignment between  $s_1$  and  $s_2$  according to the drop-DTW (on the right). The alignment is illustrated by the grey bars on the left and by the colored cells in the matrix (on the right). The red-crosses illustrate that  $P$  and  $S$  has been dropped in this alignment, the cells at the corner are dark because they do not satisfy the constraint  $\tau = 3.5$ . Clinical pathway interpretation:  $S$  stands for surgery,  $C$  consultation,  $P$  physiotherapist session and  $R$  radiotherapy session.

represent a sequence of event types distributions:

$$\langle ([d_1^1, \dots, d_1^n], t_1), \dots, ([d_K^1, \dots, d_K^n], t_K) \rangle$$

where  $\mathbf{d} = [d^1, \dots, d^n] \in [0, 1]^n$ , such that  $\sum_i d^i = 1$ .  $\mathbf{d}$  can be seen as a probability distribution of event types. Thus, events can be represented in a vector space and the representation of a timed sequence becomes close to the representation of a (multivariate, irregular) time series.

Example 1 illustrates the representation of a timed sequence as a probabilistic timed sequence. Conversely, most of the probabilistic timed sequences can not be represented by a time sequence.

**Example 1** (Probabilistic representation of a timed sequence). *Let  $\Sigma = \{S, R, C, P\}$  be an ordered set of event types. In the context of clinical care pathways,  $S$  may represent a surgery event,  $C$  a consultation,  $P$  a physiotherapist session, and  $R$  a radiotherapy session. We define three timed sequences:  $s_1 = \langle (S, 1), (C, 2), (R, 4.5) \rangle$ ,  $s_2 = \langle (C, 0), (C, 2), (P, 3), (S, 4), (R, 5) \rangle$ , and  $s_3 = \langle (S, 0), (C, 1), (C, 2), (R, 4) \rangle$  be three synthetic timed sequences illustrated in Figure 1. These sequences can be represented as probabilistic sequences:*

$$\begin{aligned} s_1 &= \langle ([1, 0, 0, 0], 1), ([0, 0, 1, 0], 2), ([0, 1, 0, 0], 4.5) \rangle \\ s_2 &= \langle ([0, 0, 1, 0], 0), ([0, 0, 1, 0], 2), ([0, 0, 0, 1], 3), ([1, 0, 0, 0], 4), ([0, 1, 0, 0], 5) \rangle \\ s_3 &= \langle ([1, 0, 0, 0], 0), ([0, 0, 1, 0], 1), ([0, 0, 1, 0], 2), ([0, 1, 0, 0], 4) \rangle \end{aligned}$$

### 3 Comparing Timed Sequences

This section introduces the drop-DTW for timed sequences. We first remind the definition of DTW and drop-DTW and then we adapt the latter for timed sequences.

#### 3.1 DTW and drop-DTW

Dynamic Time Warping (DTW) [20] computes the optimal alignment between two series subject to certain constraints. Here, we adapt these definitions to the case of timed sequences. Let  $X = \langle (x_1, t_1^x), \dots, (x_N, t_N^x) \rangle$  and  $Z = \langle (z_1, t_1^z), \dots, (z_K, t_K^z) \rangle$  be two input sequences, where  $N, K$  are their respective lengths. An alignment between two sequences is defined as a binary matrix,  $\pi \in \{0, 1\}^{K \times N}$  where  $\pi_{i,j} = 1$  if  $z_i$  is matched to  $x_j$ , and 0 otherwise. Matching an element  $z_i$  to an element  $x_j$  incurs a cost  $C_{i,j}$  between the  $i$ -th and  $j$ -th elements of the sequences. The cost values are user-defined. DTW finds the alignment  $\pi^*$  between sequences  $Z$  and  $X$  that minimizes the global cost:

$$\pi^* = \operatorname{argmin}_{\pi \in \mathcal{P}} \langle \pi, C \rangle = \operatorname{argmin}_{\pi \in \mathcal{P}} \sum_{i \in [K], j \in [N]} \pi_{i,j} C_{i,j},$$

where  $\langle \pi, C \rangle$  is the Frobenius inner product and  $\mathcal{P}$  is the set of all feasible alignments satisfying the following constraints: monotonicity, continuity, and end-point correspondence ( $\pi_{1,1} = \pi_{K,N} = 1, \forall \pi \in \mathcal{P}$ ).

drop-DTW [7] extends the set of feasible alignments,  $\mathcal{P}$ , to those adhering only to the monotonicity constraint. Consequently, unlike DTW, events can be dropped from the alignment process. drop-DTW finds the optimal temporal alignment while allowing for the removal of certain elements. To account for unmatched elements in the alignment cost calculation, [7] introduce a new drop cost:  $\delta \in \mathbb{R}^+$ .<sup>2</sup> The optimal matching can then be defined as follows:

$$\pi^* = \underset{\pi \in \mathcal{P}}{\operatorname{argmin}} \langle \pi, C \rangle + \delta \cdot (P_r(\pi) + P_c(\pi))$$

where  $\overline{\mathcal{P}} \in \{0, 1\}^K$  is the set of binary matrices satisfying only the monotonicity constraint, and  $P_c(\pi) \in \{0, 1\}^N$  is a vector whose  $j$ -th element is 1 if  $\pi_{:,j} = \mathbf{0}$  and 0 otherwise;  $P_r(\pi)$  is defined similarly, but along the rows. The indices stand for rows and columns. Similar to DTW, drop-DTW can be efficiently evaluated using dynamic programming techniques.

**Example 2** (Drops representation in drop-DTW). *In Example 1, the alignment of the sequences  $s_1$  and  $s_2$  leads to skip the  $P$  event and the second  $S$  event (see Figure 1). Then, the drops are represented as follows:  $P_r = [1, 1, 1]$  and  $P_c = [1, 1, 0, 0, 1]$ .*

### 3.2 Adaptation of drop-DTW for timed sequences

The proposed adaptation of drop-DTW focuses on proposing a metric between events and adding temporal constraints to the drop-DTW algorithm. The Dynamic Time Warping algorithm requires defining a cost between two events to compare sequences. To take into account both time and events, we decided to assign weights to a time distance cost  $c_t$ , and an event difference cost  $c_e$ . Having introduced the probabilistic representation of timed sequences, the distance between two events  $e_a = (\mathbf{d}_a = [d_a^1, \dots, d_a^n], t_a)$  and  $e_b = (\mathbf{d}_b = [d_b^1, \dots, d_b^n], t_b)$  is defined by:

$$d(e_a, e_b) = \sqrt{p_t \cdot \|\mathbf{d}_a - \mathbf{d}_b\|_2^2 + p_e \cdot (t_b - t_a)^2}$$

where  $p_t, p_e \in \mathbb{R}_+$  are weights of two terms: a distance between probabilistic events and a time gap. This metric implies that time events are in an Euclidean space of dimension  $n + 1$ . This is of particular interest in the next section, to define average events.

We can now adapt the metric between two timed sequences. We have chosen to use drop-DTW for its ability to handle aberrant or missing events. Algorithm 1 presents the method for calculating the proposed distance between two timed sequences. It takes as input the cost matrix  $C$  that computes the cost of aligning two events using the distance  $d$ . The algorithm finds the best alignment between sequence events to minimize the overall metric. The use of dynamic programming technique ensures the efficiency of the method. For the sake of computing efficiency, we also add constraints on the possible event pairings by allowing matches only between events that are close in time, *i.e.* with a delay between their occurrence lower than a predefined threshold  $\tau$ .

**Example 3** (Illustration of drop-cost usefulness). *Considering the sequences from Example 1, sequences  $s_1$  and  $s_3$  are very similar, but sequences  $s_1$  and  $s_2$  does not seem so different if we ignore some of the events. In the proposed alignment between  $s_1$  and  $s_2$ , the drop-cost was lower than the cost of pairing  $P$  in the sequence  $s_2$  with another event of  $s_1$ . Thus, it was preferable to ignore this event. This is directly related to the definition of the distance between a physiologist session and other types of event. In the case of the second surgery event in  $s_2$ , it was too far in time from the one in  $s_1$  according to the  $\tau$  constraint. Then, against, it was preferable to ignore it in the final alignment.*

<sup>2</sup>In the original approach of Dvornik et al., the cost value depends on the event type. We use a constant value for the sake of simplification.

---

**Algorithm 1:** drop-DTW with proximity constraints

---

**Data:**  $C \in \mathbb{R}^{M \times N}$ : cost matrix;  $\delta \in \mathbb{R}^+$ : drop cost;  $\sigma \in \mathbb{R}^+$ ,  $\tau \in \mathbb{R}^+$ : sequential/temporal proximity constraints

```
1  $D_{0,0}^+ \leftarrow 0$ ;  $D_{i,0}^+ \leftarrow +\infty$ ;  $D_{0,j}^+ \leftarrow +\infty$ ,  $\forall i \in [M], j \in [N]$ ; // Match table
2  $D_{0,0}^- \leftarrow 0$ ;  $D_{i,0}^- \leftarrow i \times \delta$ ;  $D_{0,j}^- \leftarrow j \times \delta$ ,  $\forall i \in [M], j \in [N]$ ; // Drop table
3  $D_{0,0} \leftarrow 0$ ;  $D_{i,0} \leftarrow D_{i,0}^-$ ;  $D_{0,j} \leftarrow D_{0,j}^-$ ,  $\forall i \in [M], j \in [N]$ ; // Optimal solution table
4 for  $i = 1$  to  $N + 1$  do
5   for  $j = \max(1, i - \sigma + 1)$  to  $\min(m + 1, i + \sigma)$  do // Sakoe-Chiba constraints
6     if  $\text{distdate}(i - 1, j - 1) \leq \tau$  then // Temporal constraint
7        $D_{i,j}^+ \leftarrow C_{i-1,j-1} + \min(D_{i-1,j}, D_{i,j-1}, D_{i-1,j-1}^+)$ ; // Alignment cost
8        $D_{i,j}^- \leftarrow \delta + D_{i,j-1}^-$ ; // Drop cost
9        $D_{i,j} \leftarrow \min(D_{i,j}^+, D_{i,j}^-)$ ; // Final cost
10    else
11    |  $D_{i,j} \leftarrow +\infty$ ; // Unsatisfied temporal constraints  $\Rightarrow$  impossible match
12 return  $D_{M,N}$ ,  $D$ ;
```

---

## 4 Construction of an Average Timed Sequence

Computing an average timed sequence consistent with the notion of similarity is essential for clustering algorithms. In the following, we propose a method inspired by the DBA algorithm [16] to create average sequences in terms of drop-DTW, taking into account the possibility of removal. DBA addresses the problem of generating an average time series in terms of DTW – which does not hold distance properties. In this case, the generalized average is defined as a time series that minimizes the global inertia (see Equation 1).

DBA iteratively refines a (randomly initialized) series to minimize its quadratic distance (DTW) with respect to the series to be averaged. At each iteration, the algorithm performs “vertical” centroids of the aligned elements of the time series. The definition of a “vertical” centroid corresponds to averaging series real values. Our approach enables the adaptation to timed sequences thanks to the representation of events as distributions of types.

### 4.1 Average Timed Sequences

Let  $\mathcal{S} = \{s_1, \dots, s_m\}$  be a set of timed sequences, Algorithm 2 details the steps for generating an average timed sequence inspired by the principle of *vertical* averages of DBA [17]. In this algorithm,  $\mathcal{C}^{(k)}$  denotes the current average sequence at the  $k$ -th iteration. At line 1, the current average time sequence,  $\mathcal{C}^{(0)}$  is initialized with a timed sequence chosen from  $\mathcal{S}$ . Then, this current average is processed into the outer loop which repeats  $\text{maxit}$  times the alignment and update steps. The alignment step (lines 6-11) calculates the alignment by drop-DTW between  $\mathcal{C}^{(0)}$  and each of the other sequences. The alignment to one sequence is obtained line 8 from the distance matrix computed the drop-DTW. We refer the reader to the original DTW algorithm to get the details on how to compute an alignment from a distance matrix. Lines 9 to 11 collect, on the one hand, the list of paired events in the sequences of  $\mathcal{S}$  per event id,  $e$ , in  $\mathcal{E}_e$  and, on the other hand, their dates per event id,  $e$ , in  $\mathcal{T}_e$ . These lists are then averaged lines 12-13. This is the update step. This vertical average of events is constructed as the histogram of event distributions. Lastly, line 14 creates a new average sequence which can be refined in a subsequent iteration until convergence. In our implementation, we assume convergence after a fixed number of iterations but a convergence criteria could be implemented to prevent useless iterations.

Similarly to DBA, events from probabilistic sequences are represented in a vector space. Thus, the new average sequence is always closer (according to drop-DTW) to the sequences it averages, and the algorithm converges well (see Section 4.2). Additionally, this algorithm preserves the drop-DTW’s ability to allow temporal shifts and handle missing or aberrant events.

---

**Algorithm 2:** TSR: Generation of an Average Timed Sequence
 

---

**Data:**  $\mathcal{S} = \{s_1, s_2, \dots, s_m\}$ : list of timed sequences;  $maxit \in \mathbb{N}$ : maximum number of iterations

```

1  $\mathcal{C}^{(0)} \leftarrow alea(\mathcal{S});$ 
2  $k \leftarrow 0;$ 
3 while  $k < maxit$  do
4    $\mathcal{E} \leftarrow [\emptyset, \dots, \emptyset];$  // sets of length  $|\tau|$ 
5    $\mathcal{T} \leftarrow [\emptyset, \dots, \emptyset];$  // idem
6   for  $i = 1, \dots, m$  do
7      $\mu, D \leftarrow DropDTW(\mathcal{C}^{(k)}, s_i);$  // calculate the distances matrix of drop-DTW between  $\mathcal{C}^{(k)}$  and  $s_i$ 
8      $\pi^* \leftarrow GetAlignement(D);$  // calculate the best alignment from  $D$ 
9     for  $(i_r, i_s) \in \pi^*$  do
10       $\mathcal{E}_{i_r} \leftarrow \mathcal{E}_{i_r} \cup \{s_i[i_s].d\};$  // add event from  $s_i$  aligned at  $i_r$ 
11       $\mathcal{T}_{i_r} \leftarrow \mathcal{T}_{i_r} \cup \{s_i[i_s].t\};$  // add date from  $s_i$  aligned at  $i_r$ 
12    $\tau \leftarrow \{\tau_j \mid \forall j \in [|\mathcal{C}^{(k-1)}|], \tau_j = \frac{1}{|\mathcal{T}_j|} \sum_{t \in \mathcal{T}_j} t, |\mathcal{T}_j| \neq 0\};$  // calculate vertical averages
13    $\mathcal{H} \leftarrow \{h_j \mid \forall j \in [|\mathcal{C}^{(k-1)}|], h_j = \frac{1}{|\mathcal{E}_j|} \sum_{d \in \mathcal{E}_j} d, |\mathcal{E}_j| \neq 0\};$ 
14    $\mathcal{C}^{(k+1)} \leftarrow (\mathcal{H}, \tau);$  // refined average timed sequence
15    $k \leftarrow k + 1;$ 
16 return  $\mathcal{C}^{(k)};$ 

```

---




Figure 2: Illustration of one averaging step. At the top, the alignment of the current average  $s_r$  is computed with each sequence. Then, the “vertical” average of the set of events paired to one  $s_r$  event yields an average probabilistic event in  $s'_r$  (at the bottom).

The choice of the timed sequence in line 1 is crucial for the result. We choose to randomly select a sequence of the longest length. We made this choice because our algorithm decreases the length of the current average timed sequence, but it can not increase it. Note that there are some smarter initialization strategies for k-means like algorithms that could improve the convergence, for instance K-means++ [1], or recent provably good seeding methods [2].

**Example 4** (Iteration of the averaging algorithm). *Figure 2 illustrates one iteration of average calculation for the three sequences in Example 1. It assumes that we already have a current average sequence  $s_r$  (a surgery, then a consultation and two radiotherapy sessions). The figure depicts the alignments by drop-DTW of  $s_r$  with each sequence of Example 1. The gray lines represents the alignments (line 8).  $s'_r$  is the updated average sequence obtained by the lines 11-12 of Algorithm 2. As the C event has been paired with both C and R events, the second probabilistic event of  $s'_r$  combines the two types of events (proportionally). In addition, the first R event of  $s_r$  is removed from  $s'_r$  because it is paired with no event. The other events are purely paired with similar event types, then the probabilistic event is a pure S or R in  $s'_r$  but their time positions have been adjusted with respect to the positions of the new aligned events.*

## 4.2 Convergence of the Algorithm

In this section, we prove the theoretical convergence of the proposed algorithm.

The algorithm 2 iteratively refines an average sequence for a set of sequences,  $\{s_i\}_{i \in [m]}$ . At step  $k$ , the inertia of the average sequence is defined as:

$$I^{(k)} = \frac{1}{m} \sum_{i \in [m]} \text{DropDTW}(\mathcal{C}^{(k)}, s_i) \quad (1)$$

The inertia represents the average error incurred by using  $\mathcal{C}^{(k)}$  as the representative of the set of sequences  $\{s_i\}_{i \in [m]}$ . To demonstrate the convergence, we prove that  $I^{(k+1)}$  is a decreasing suite, i.e.,  $I^{(k+1)} \leq I^{(k)}$ , for all  $k \in \mathbb{N}$ .

The  $i$ th term of the inertia is

$$\text{DropDTW}(\mathcal{C}^{(k)}, s_i) = \min_{\pi} \sum_{(x,y) \in \pi} d(\mathcal{C}^{(k)}[x], s_i[y]) + \delta. (P_r(\pi) + P_c(\pi))$$

where  $d$  is the metric between timed events (see Section 3.2) and  $\pi$  plays the same role as in section 3.1. It represents the alignment between two sequences. In the remainder, we denote by  $\pi_i^{(k)}$  the optimal alignment of sequence  $s_i$  with sequence  $\mathcal{C}^{(k)}$ , i.e., the alignment that minimizes the sum of costs.

Let us now introduce  $q_i^{(k+1)} = \sum_{(x,y) \in \pi_i^k} d(\mathcal{C}^{(k+1)}[x], s_i[y]) + \delta. (P_r(\pi_i^k) + P_c(\pi_i^k))$ , for all  $i \in [m]$  and  $k \in \mathbb{N}^*$ . Providing  $\pi_i^{k+1}$  is the optimal alignment, we have:

$$\sum_{(x,y) \in \pi_i^{k+1}} d(\mathcal{C}^{(k+1)}[x], s_i[y]) + \delta. (P_r(\pi_i^{k+1}) + P_c(\pi_i^{k+1})) \leq q_i^{(k+1)}$$

It is worth noticing that we retrieve the inertia of the  $(k+1)$ -th average sequence  $I^{(k+1)}$  by summing the left-hand terms over  $i$ , and then:

$$I^{(k+1)} \leq \frac{1}{m} \sum_{j \in [m]} q_j^{(k+1)} \quad (2)$$

Let us not have a look at the sum over  $i$  of the  $q_i^{(k+1)}$ . This sum can be rewritten as follows:

$$\begin{aligned} \frac{1}{m} \sum_{j \in [m]} q_j^{(k+1)} &= \frac{1}{m} \sum_{j \in [m]} \sum_{(x,y) \in \pi_j^{k+1}} d(\mathcal{C}^{(k+1)}[x], s_j[y]) + \delta. (P_r(\pi_j^{k+1}) + P_c(\pi_j^{k+1})) \\ &= \sum_x \frac{1}{m} \sum_{j \in [m]} \sum_{\{y | (x,y) \in \pi_j^{k+1}\}} d(\mathcal{C}^{(k+1)}[x], s_j[y]) + \delta. (P_r(\pi_j^{k+1}) + P_c(\pi_j^{k+1})) \end{aligned}$$

Let us now have a look at the construction of  $\mathcal{C}^{(k+1)}[x]$  in the algorithm (see lines 11-13). It is the barycenter of all elements  $s_j[y]$  that have been associated with  $\mathcal{C}^{(k)}[x]$  by the alignment  $\pi_j^{k+1}$ . Thanks to the property of barycenters in an Euclidean space, we have for any  $\mathbf{v}$ :

$$\frac{1}{m} \sum_{j \in [m]} \sum_{\{y | (x,y) \in \pi_j^k\}} d(\mathbf{v}, s_j[y]) \geq \frac{1}{m} \sum_{j \in [m]} \sum_{\{y | (x,y) \in \pi_j^k\}} d(\mathcal{C}^{(k+1)}[x], s_j[y]) \quad (3)$$

Considering that:

$$I^{(k)} = \sum_x \frac{1}{m} \sum_{j \in [m]} \sum_{\{y | (x,y) \in \pi_j^k\}} d(\mathcal{C}^{(k)}[x], s_j[y]) + \delta. (P_r(\pi_j^k) + P_c(\pi_j^k))$$



and applying inequation 3 with  $\mathbf{v} = \mathcal{C}^{(k)}[x]$ , it ensues:

$$\begin{aligned}
I^{(k)} &\geq \sum_x \frac{1}{m} \sum_{j \in [m]} \sum_{\{y | (x,y) \in \pi_j^k\}} d(\mathcal{C}^{(k+1)}[x], s_j[y]) + \delta \cdot (Pr(\pi_j^k) + Pc(\pi_j^k)) \\
&\geq \frac{1}{m} \sum_{j \in [m]} q_j^{(k+1)}
\end{aligned} \tag{4}$$

In the case  $\{y | (x, y) \in \pi_j^k\} = \emptyset$ , this means that for all  $s_i$ , there is an empty row at line  $x$  of the alignment with  $\mathcal{C}^{(k+1)}$ . For instance, the first  $R$  event in Figure 2 is assigned to none event. Then, this event contributes to the inertia only via a drop-cost. Removing such an event from  $\mathcal{C}^{(k+1)}$  turns out to remove an empty row for all sequences  $s_j$ , and thus subtracts one to all  $Pr(\pi_j^{k+1})$ . Then, it also decreases the true values of  $q_j^{(k+1)}$ .

Finally, with equations 2 and 4, we conclude that  $I^{(k+1)} \leq I^{(k)}$ . This proves that Algorithm 2 actually refines the average timed sequence and converges. Nevertheless, it does not guarantee to find the global minimum of inertia. It only ensures to find a local minimum.

### 4.3 Parameters

The proposed method for averaging time sequences has two types of parameters: the parameters of the drop-DTW metric ( $\tau \in \mathbb{R}$ ,  $p_e \in [0, 1]$ ,  $p_t \in [0, 1]$ ,  $\delta \in \mathbb{R}$ ) and the parameters of the Algorithm 2 (maximum iterations for drop-DTW and iterative averaging).

The proposed metric is mainly parameterized by the ratio  $\frac{p_t}{p_e}$  and the drop-cost  $\delta$ . We explicitly explain here how to determine these values relative to the unique threshold value  $\tau$ . Thus, the expert only has one parameter to set up:  $\tau$ , which can be seen as a threshold of indifference delay. Above this delay, two events cannot be considered as similar, regardless of their types. An expert can set up  $\tau$  based on this intuition.

The ratio  $\frac{p_t}{p_e}$  can be inferred from the delay  $\tau$  by the following reasoning. For events located within the time interval  $[t - \tau, t + \tau]$ , the importance of the event type should be greater than the importance of time. We must then choose  $p_e$  to be at least  $\tau^2$  times greater than  $p_t$ . Without lost of generality, the value of  $p_t$  can be set at 1 by default. The drop-cost is also derived from  $\tau$ , i.e., the delay from which there is no longer any interest in associating two events. Thus, we propose to define the drop-cost as:

$$\delta = \frac{1}{\tau^2}(\tau + t_{max}) + 1$$

where  $t_{max}$  is the maximum time distance to consider.

Note that this suggested values are only a guideline to initiate a set up for the metric. Nonetheless, it has to be adjusted and tested on real data.

## 5 Clustering of Timed Sequences

The proposed metric between timed sequences and the averaging technique are used to adapt two classical clustering techniques: a hierarchical ascendant clustering (HAC) and a K-means algorithm [12]. These two algorithms require to predefine the expected number of clusters,  $K$ .<sup>3</sup>

HIERASTISEQ denotes the HAC algorithm for timed sequences. The principle of the algorithm is to iteratively aggregate the two closest timed sequences according to the drop-DTW. A group of two timed sequences is then averaged using Algorithm 2, and the average timed sequence replaces the two former ones. The algorithm terminates when it reaches the expected number of clusters.

The  $K$ -means algorithm also proceeds to refine iteratively  $K$  clusters.  $K$ -means algorithm has two main steps: first the sequences of each cluster are averaged (Algorithm 2). The computed averages are the

<sup>3</sup>Model selection based on the inertia of the clustering may be applied to select the optimal number of clusters.

	<i>Predicted</i>	
Real	0	15
Real	15	15

	<i>Predicted</i>	
Real	15	0
Real	0	30

Table 1: Confusion matrices with supplementary events. On the left:  $p_t = \frac{1}{9}$ ,  $p_e = 1, \delta = +\infty$ ; On the right:  $p_t = \frac{1}{9}$ ,  $p_e = 1, \delta = 4$

representatives of their clusters. Then, new clusters are made by assigning optimally each sequence to one of the  $K$  representatives.

## 6 Experiments and Results on Synthetic Data

In this section, we present experiments on synthetic data. We conducted several tests to evaluate the effectiveness of the method under various scenarios and to study the effect of the parameters.

### 6.1 Effects of the $\frac{p_t}{p_e}$ Ratio

For this first experiment, we created two model timed sequences that differ in one or more of the dimensions we want to study: sequence length, event nature, and event dates. Therefore, there are  $8 = 2^3$  possible variations. A dataset consists of 15 sequences for each of the 8 possible variations and the base sequence (thus  $135 = 9 \times 15$  sequences in total). The dates of these timed sequences are obtained and made random by adding uniform noise to the date. The objective is to evaluate the accuracy of identifying the 9 categories of sequences.

We tested two sets of parameters: first, we chose the ratio  $\frac{p_t}{p_e} = \frac{1}{9}$ , with  $\delta = +\infty$ . With this configuration, the cost of an event is more important than the time cost for 3 days around the event. The second set of parameters used was  $\frac{p_t}{p_e} = \frac{1}{400}$ , still with  $\delta = +\infty$ . This time, the difference cost between events is more important than the time cost over the first 20 days.

In the case of the first configuration, we observe that all clusters generated by HIERASTISEQ were correctly identified (Kappa score of 1). However, in the case of the second configuration, we observe a degradation of results with a Kappa score of 0.62. This degradation was expected because the maximum time distance observed between events is around 20. Consequently, the notion of time is practically never taken into account by the metric.

This experiment thus shows that the choice of parameters is crucial but can be made by taking into account the distance between events of the sequences.

### 6.2 Effect of Drop-Cost

We can then test the effect of the drop-cost on a synthetic dataset, with and without drop-cost. We consider 3 models of sequences. We generate 15 concrete sequences from each of them by randomizing the date of the events<sup>4</sup>. The sequences of the first and the third types are similar, except that the sequences of the third type have an aberrant event.

1.  $\langle (D, 0.0)(E, 3.0)(F, 5.0)(D, 7.0)(D, 10.0)(E, 12.0)(F, 15.0) \rangle$
2.  $\langle (E, 2.0)(A, 4.0)(D, 8.0)(C, 12.0) \rangle$
3.  $\langle (D, 0.0)(E, 3.0)(F, 5.0)(D, 7.0)(D, 10.0)(E, 12.0)(F, 15.0)(D, 1500.0) \rangle$

Intuitively, we want to cluster together all sequences of types 1 and 3. We apply HIERASTISEQ without drop-cost and obtain the confusion matrix in Table 1 on the left, with a Kappa score of  $-0.5$  (disagreement). We observe that the algorithm associated instances of sequence 1 with instances of sequence 2.

<sup>4</sup>We added a centered normal random noise with a standard deviation of 1.

	<i>Predicted</i>	
Real	0	15
Real	15	15

	<i>Predicted</i>	
Real	15	0
Real	0	30

Table 2: Confusion matrices with missing events. On the left:  $p_t = \frac{1}{9}$ ,  $p_e = 1, \delta = +\infty$ ; On the right:  $p_t = \frac{1}{9}$ ,  $p_e = 1, \delta = 4$

Adding drop-cost coinciding with the distances between the sequences allows us to obtain the confusion matrix in Table 1 on the right, with a Kappa score of 1 (total agreement). The drop cost allows neglecting the last event considered aberrant, and the results match expectations.

We also conducted a similar experiment with missing values. This time, the model sequences are as follows:

1.  $\langle (D, 0.0), (E, 3.0), (F, 5.0), (D, 7.0) \rangle$
2.  $\langle (E, 2.0), (A, 4.0), (E, 7.0), (D, 8.0), (D, 10.0), (E, 12.0) \rangle$
3.  $\langle (D, 0.0), (E, 2.0), (E, 3.0), (F, 5.0), (D, 7.0), (F, 9.0), (A, 13.0) \rangle$

We then notice that all events dated from sequence 1 are present in sequence 3, but they are all different from sequence 2. We would therefore like to cluster sequences from 1 and 3 together. Keeping the same parameters as before. Without drop-cost, we obtain the confusion matrix in Table 2 on the left. This result corresponds to expectations because the sequences from 2 and 3 are then associated as they are closer in time. However, with  $\delta = 4$ , we obtain the confusion matrix in Table 2. The table on the right, as expected, shows that the events furthest in time are no longer associated.

The tests on synthetic data lead to the expected results, thus validating the effectiveness of clustering in the situations described. We now continue studying the method with an experiment on real data to evaluate its ability to identify relevant clusters.

## 7 Application to Real-World Care Pathways

In this section, we apply the proposed method to real-world care pathways and compare our results with those of Sequence State Analysis (SSA) [8]. SSA, particularly the TraMineR tool, is commonly used for categorizing care pathways [19, 14].

This experiment is conducted on the data warehouse of the Greater Paris Hospital (AP-HP). We identified a population of 3 311 patients having a stay with a lung resection surgery. We collected events that occurred within 90 days around a lung resection surgery as well as the date of death if applicable. Events further away are clinically less interesting to study. For clinical reasons, we also decided to keep only the first occurrence of chemotherapy for each patient. We made the same choice for radiotherapy and immunotherapy.

The parameters chosen are  $K = 5$  clusters,  $\tau = 7$ ,  $t_{\max} = 16$ , and  $\delta = 1$ . For choosing the number of clusters, different values were tested. By increasing the number of clusters, it is mainly small-sized clusters that split. Therefore, there may not necessarily be an interest in increasing this number.  $\tau$  is set to 7 days because, in our dataset, it may be interesting to associate the same types of events that occurred within a week of each other. Finally, based on the dataset, we consider that there is no interest in pairing two events more than 16 days apart.

For this experiment, we used the servers from the Paris University Hospital made of AMD EPYC 7713 processors. Our algorithms are neither parallelized nor distributed. However, the implementation of soft-DTW by [23] can utilize multiple cores. During our experiments, we requested only one core from the servers.

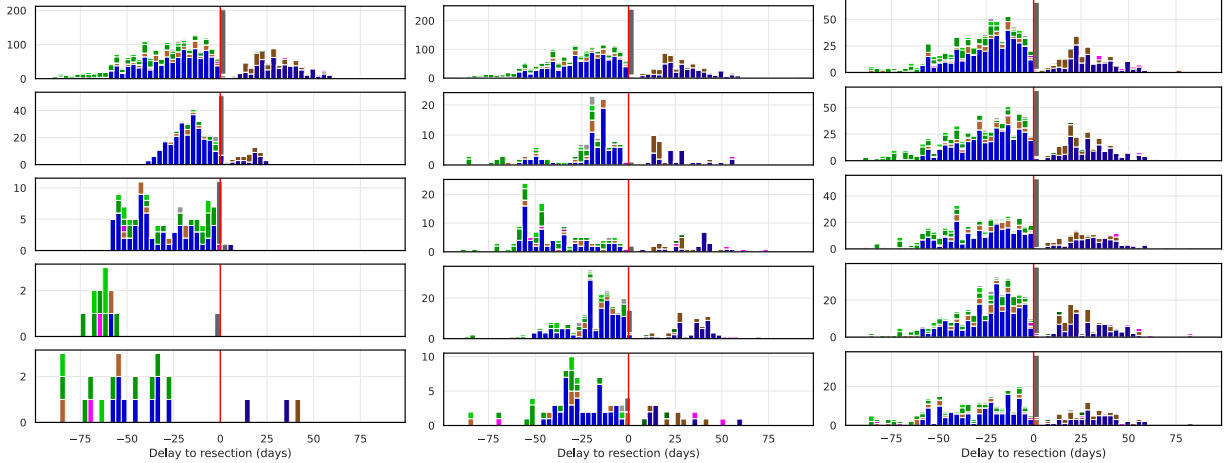


Figure 3: Clusters obtained by the clustering methods, from left to right: HIERASTiSEQ, K-means and TraMineR. Clusters are ordered row-wise by their size. Each cluster is represented by the histogram of event types over time. Each color corresponds to a specific type of event (see Figure 4 for color legend). The higher the bar, the more events there are at the given time across all patients. The vertical red line indicates the index-date of the resection (0 delay).

## 7.1 Comparison with TraMineR

In this section, we selected randomly 500 patients among the cohort of patients to compare the three clustering approaches: HIERASTiSEQ, K-means with drop-DTW and TraMineR. Regarding the computing time of HIERASTiSEQ and TraMineR, this comparison is conducted only on a small subset of patients. TraMineR requires converting timed sequences into time series. As TraMineR is sensitive to sparse series, we increased the time sampling from 1 to 3 days for calculations, and we added the empty event corresponding to the absence of data over a period of 3 days. TraMineR uses the hierarchical clustering (ward criteria) and on optimal matching metric between state sequences.

Figure 3 corresponds to the  $K = 5$  clusters for the three methods. Each cluster is represented as an histogram of events over time. One can first notice that TraMineR provides clusters that look very similar to each others. It is difficult to identify specificities for each cluster. Contrary to the other two methods, based on our drop-DTW metric, the sizes of the clusters are more homogeneous. For HIERASTiSEQ and K-means, there is one large cluster and several small clusters that gather patient specific timed events. The two large clusters (cluster 1 for HIERASTiSEQ and cluster 1 for K-means) looks very similar in distribution of events. HIERASTiSEQ identifies two clusters with specific behaviors: the cluster 3 has no post-operative events. This may correspond to patients that had a surgery in the hospital but their follow-up in another medical structure; the cluster 2 are patients with short delays between pre- or post-operative care events around the surgery. The clusters 3 and 4 identified by K-means also gather patients with a late pre-operative consultation (wrt the resection). According to the distributions, these patients have two pre-operative consultations: in cluster 3, the first pre-operative consultation occurs earlier than in cluster 4; and cluster 4 holds patients having a late post-operative consultation.

We conclude from this experiment that the use of drop-DTW yields clusters of patients with more specific timed sequences than TraMineR. For this dataset, the latter method does not seem to be able to reveal interesting clusters of patients. We performed a quantitative (confusion matrices) and qualitative analysis. The analysis of confusion between HIERASTiSEQ and TraMineR shows that there is no correspondence between the clusters. This suggests that the metrics are semantically different. In this case, the semantics of drop-DTW on timed sequence seems more meaningful, but it does not mean that it would be the case for all datasets.

It is worth noting that results of K-means with drop-DTW may significantly differ across different runs




Figure 4: Clusters of care-pathways identified by K-means with drop-DTW on 3311 patients. Each row is a cluster with the representation of the barycenter on the left and the representation of the events distribution over time on the right. Each color corresponds to a specific type of event.

due to the randomness of the cluster seeds. HIERASTiSEQ and TraMineR are more deterministic, but the iterative computation of drop-DTW or optimal matching may introduce some randomness. Nonetheless, the hierarchical clustering strategy of HIERASTiSEQ and TraMineR is resource demanding. This prevents applying them on the entire dataset. This was possible only with K-means and takes only a few minutes.

## 7.2 Cluster Analysis of K-means with drop-DTW

In this section, we present the results of the K-means clustering with drop-DTW metric applied on the entire cohort. Figure 4 illustrates the  $K = 5$  clusters. Each cluster is depicted with its barycenters and the distribution of events.

Let us first comment on the clusters based on their event distributions. Three main clusters have been identified (clusters 1, 2 and 3 and share almost the same size (about 1000 patients). The two other clusters are very small (numbered 4 and 5). These small clusters of patients seem to be outliers. Similarly to the previous results, the K-means identified clusters of patients with different delays between the pre-operative consultations and resection. The cluster 1 corresponds to patients having a small delay between these events and possibly two pre-operative consultations. These patients possibly have a late stage of the disease and need a quick surgery. We made this hypothesis because their survival at 6 months is significantly lower than the two others clusters. The cluster 2 and cluster 3 differ mainly in the delay between the surgery and the post-operative consultations. The cluster 3 also holds patients without exams the day of the resection.

The same observations can be made based on the analysis of the barycenters. We remind that a barycenter is a probabilistic timed sequence. Each event is represented by a distribution of event's types. On the figure, each bar is an event located at its time of occurrence. A bar stacks the probabilities per event's type. We remark that most of the events have one main type. The event at the 10th day for cluster 3 is a mixture of mainly two events types: post-operative consultation or post-operative multi-disciplinary consultation meeting (MCM). This possible confusion may be explained by the imprecision of the records. MCM are identified by consultation reports. If the report is not correctly tagged as an MCM, then it is simply a

consultation. Then, the flexibility of the metric makes it somehow robust to possible labeling errors in the data.

We can also notice that there is a clear alignment between the barycenters' events and the event distribution. This illustrates that the averaging method, based on drop-DTW provides a meaningful average of a collection of complex timed sequences. More specifically, comparing the clusters through the decay between consultations and surgery leads to the same conclusions as before.

Lastly, one can note that there are a lot of repeated events in the barycenter. Furthermore, the presence of repeated events in the barycenter can be misleading, as it could improve either indicate a repetition of events or alternative delays for a single event due to possible drops.

## 8 Conclusion

We proposed a method for clustering timed sequences inspired by existing techniques for time series. Thus, the use of timed sequences allows co-occurring events, irregular time sampling, and sparse events. All these advantages are very interesting from a clinical point of view for describing care pathways. We showed that by using drop-DTW, the clustering algorithm is less sensitive to missing and aberrant values when drop-DTW is used. The result is interesting for application to real-world care pathways. Then, we tested our method on real data from the OPTISOINS study for our approach compared to TraMineR. Based on the initial results, drop-DTW seems to be clinically promising for identifying average pathways.

However, the large number of parameters is a first weakness of the method. Regarding the number of parameters, our effort focused on a guide for choosing parameter values, which the experiments tend to validate. In perspective, a more in-depth study of the effect of parameters still needs to be carried out. Additionally, we have to validate the clinical interest of average pathways by clinicians.

Finally, the number of computations of a drop-DTW required during the clustering of timed sequences may be very high and may lead to long computation times for large datasets. Some strategies that could make use of distributed resources (grid computing) or fast approximate of the metric could significantly improve the practical usability of our method.

**Acknowledgments** The authors thanks gratefully Dr. Yen-Lân NGuyen from Department of Anaesthesia and Intensive Care Medicine, AP-HP, Cochin Academic Hospital, for her time and expertise she provided during this project. Part of the research presented in this article is funded by the AP-HP Foundation, as part of the AI-RACLES Chair, and has received approval from the Scientific and Ethical Committee of the AP-HP CDW (CSE-20-01-OPTISOINS).

## References

- [1] David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. In *Proceedings of the ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035, 2007.
- [2] Olivier Bachem, Mario Lucic, Hamed Hassani, and Andreas Krause. Fast and provably good seedings for k-means. *Advances in neural information processing systems*, 29, 2016.
- [3] Guillaume Bernard, Cyril Faucher, and Karel Bertet. Towards reconstruction of human trajectories in indoor environments. In *EKAW 2018*, volume 2262, pages 37–40. CEUR-WS. org, 2018.
- [4] Philippe Besnard and Thomas Guyet. *Chronicles: Formalization of a Temporal Model*. Springer Briefs in Computer Science, 2023.
- [5] Yuanzhe Chen, Panpan Xu, and Liu Ren. Sequence synopsis: Optimize visual summary of temporal event data. *IEEE transactions on visualization and computer graphics*, 24(1):45–55, 2017.

- [6] Paul Compagnon, Grégoire Lefebvre, Stefan Duffner, and Christophe Garcia. Sequence metric learning as synchronization of recurrent neural networks. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2021.
- [7] Mikita Dvornik, Isma Hadji, Konstantinos G Derpanis, Animesh Garg, and Allan Jepson. Drop-DTW: Aligning common signal between sequences while dropping outliers. *Advances in Neural Information Processing Systems (NIPS)*, 34:13782–13793, 2021.
- [8] Alexis Gabadinho, Gilbert Ritschard, Nicolas S. Müller, and Matthias Studer. Analyzing and visualizing state sequences in  $R$  with TraMineR. *Journal of Statistical Software*, 40(4):10, 2011.
- [9] Damien Garreau, Rémi Lajugie, Sylvain Arlot, and Francis Bach. Metric learning for temporal sequence alignment. *Advances in neural information processing systems*, 27, 2014.
- [10] Dominique Gay, Romain Guigourès, Marc Boullé, and Fabrice Clérot. TESS: temporal event sequence summarization. In *Proceedings of the International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–10, 2015.
- [11] E. Audureau H. Sebia, T. Guyet. SWoTTeD: An extension of tensor decomposition to temporal phenotyping. *Machine Learning*, page to appear, 2024.
- [12] John A Hartigan. *Clustering algorithms*. John Wiley & Sons, Inc., 1975.
- [13] Christopher Holder, Matthew Middlehurst, and Anthony Bagnall. A review and evaluation of elastic distance functions for time series clustering. *Knowledge and Information Systems*, 66(2):765–809, 2024.
- [14] Nolwenn Le Meur, Cécile Vigneau, Mathilde Lefort, Saïd Lebbah, Jean-Philippe Jais, Eric Daugas, and Sahar Bayat. Categorical state sequence analysis and regression tree to identify determinants of care trajectory in chronic disease: Example of end-stage renal disease. *Statistical Methods in Medical Research*, 28(6):1731–1740, 2019.
- [15] Dang Nguyen, Wei Luo, Tu Dinh Nguyen, Svetha Venkatesh, and Dinh Phung. Sqn2vec: Learning sequence representation via sequential patterns with a gap constraint. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 569–584. Springer, 2018.
- [16] François Petitjean and Pierre Gançarski. Summarizing a set of time series by averaging: From steiner sequence to compact multiple alignment. *Theoretical Computer Science*, 414(1):76–91, 2012.
- [17] François Petitjean, Alain Ketterlin, and Pierre Gançarski. A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognition*, 44(3):678–693, 2011.
- [18] Kishan Rama, Helena Canhão, Alexandra M. Carvalho, and Susana Vinga. AliClu - Temporal sequence alignment for clustering longitudinal clinical data. *BMC Medical Informatics and Decision Making*, 19(1):289, 2019.
- [19] Jonathan Roux, Olivier Grimaud, and Emmanuelle Leray. Use of state sequence analysis for care pathway analysis: The example of multiple sclerosis. *Statistical Methods in Medical Research*, 28(6):1651–1663, 2019.
- [20] Hiroaki Sakoe and Seibi Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing*, 26(1):43–49, 1978.
- [21] Sven Schneider, Maria Maximova, and Holger Giese. Probabilistic metric temporal graph logic. In *Proceedings of the International Conference on Graph Transformation*, pages 58–76. Springer, 2022.
- [22] Nikolaj Tatti and Jilles Vreeken. The long and the short of it: summarising event sequences with serial episodes. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 462–470, 2012.

- [23] Romain Tavenard, Johann Faouzi, Gilles Vandewiele, Felix Divo, Guillaume Androz, Chester Holtz, Marie Payne, Roman Yurchak, Marc Rußwurm, Kushal Kolar, et al. Tslern, a machine learning toolkit for time series data. *The Journal of Machine Learning Research*, 21(1):4686–4691, 2020.
- [24] Xiufan Yu, Karthikeyan Shanmugam, Debarun Bhattacharjya, Tian Gao, Dharmashankar Subramanian, and Lingzhou Xue. Hawkesian graphical event models. In *Proceedings of the International Conference on Probabilistic Graphical Models*, volume 138, pages 569–580, 2020.