Change-point Detection for Degradation Process Monitoring

Bastien Lhopitallier[1,2], Chrysoula Kosma[1], Charles Truong[3,4], Jerôme Terrier[2], Guillaume Ramelet[2], Laurent Oudre[1]
[1]*Université Paris Saclay, Université Paris Cité, ENS Paris Saclay, CNRS, SSA, INSERM, Centre Borelli,* F-91190, Gif-sur-Yvette, France
[2]*Manufacture française des pneumatiques Michelin,* F-63000 Clermont-Ferrand, France
[3] *The American University of Paris*
[4] *Université Paris Saclay, Université d'Evry Val d'Essonne, CNRS, LaMME,* F-91037, Evry, France

***Abstract*—This study proposes a novel approach to understand the degradation process in monitoring industrial time series through local comparisons with an elastic reference barycenter. In contrast to global approaches, the proposed method identifies specific barycenter points whose deformation trends can capture their local shifts, when compared to the initial time series. Change-point detection is applied to these deformation trajectories to reveal shifts in local regimes, providing area-specific information across the entire sequence. We evaluate the proposed approach on an unsupervised detection task aimed at retrieving the onset and termination of critical deformation events. Synthetic sequences are generated from real samples from a real-world battery dataset. Results demonstrate the advantage of our local approach over global baseline methods.**

***Index Terms*—Time series, change-point detection, degradation monitoring, predictive maintenance**

I. Introduction

Time series data are widely generated in industrial environments to monitor machinery and production processes [1], [2]. Many industrial processes operate in a cyclic manner. For instance, in rotating machinery, a cycle refers to a complete operational phase of the machine or equipment in use, such as a full motor rotation. These processes are monitored using sensors that generate large volumes of data. As a result, time series data derived from assembly lines often require data manipulation software, which apply compression strategies to reduce storage requirements [3]. Through these compression mechanisms, the derived time series exhibit both *(i) cyclic patterns* reflecting repeated operational phases and *(ii) irregular sampling*, induced by the compression step. These properties must be explicitly taken into account when designing machine learning algorithms to automatically solve associated tasks, such as predictive maintenance of degrading machinery [1].

In industrial applications, assessing the state of degradation of a system by quantifying its deviation from normal behavior is often referred to as "prognostics" [2]. Although often associated with failure prediction, our objective is to provide a structural analysis of the degradation dynamics. Specifically, we focus on the detection of degradation stage transition times and on the characterization of the system's progressive departure from its reference regime. In several practical scenarios, degradation is showcased through subtle deformations of the shape of operational cycles. These deformations can reflect early failures, progressive wear, or changes in operating conditions [4]. Comparison of observed cycles with a reference cycle representative of normal behavior forms a direct framework to analyze degradation. However, such deformations often appear progressively across cycles, and not all observed variations reflect degradation (see Figure 1). Some fluctuations arise from noise or operational variability. Consequently, global similarity measures between cycles and a reference may obscure interesting structural changes. A localized analysis is thus required to identify particular regions within each cycle where significant structural changes emerge.

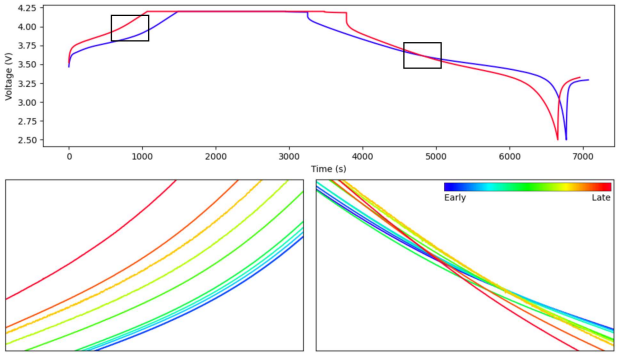

Fig. 1. Top: first (blue) and last (red) recorded cycles of battery `XJTU_2C_battery-8` from the BatteryLife dataset. Bottom: zoomed lifetime cycles (in color gradient). The left region captures battery degradation dynamics (slow deformation speed for early cycles, that increases later), while the right region does not contain meaningful information on deformations.

The retrospective identification of degradation stages in industrial time series has been addressed in the literature through several methodological approaches.

A prominent direction focuses on change-point detection [5], which identifies time indices at which the statistical properties of time series change. This family of methods is built upon the assumption that data is generated by piecewise stationary processes with different segments corresponding to distinct regimes. In industrial prognostics, similar approaches

This work has been funded by the Industrial Data Analytics and Machine Learning (IdAML) chair hosted at ENS Paris-Saclay, Université Paris-Saclay.

are used to model multi-phase degradation trajectories, with stage transitions being represented through change-point regression or two-phase stochastic models [6], [7]. However, the effectiveness of change-point detection algorithms strongly depends on the signal representation. For instance, global descriptors can mask critical information if degradation appears through local deformations within certain cycle regions.

An alternative line of work relies on anomaly detection methods, in which degradation is inferred from deviations with respect to normal behavior. A reference list of trajectories representing "normal behavior" is extracted and test cycles are compared against them to quantify abnormality [8]. Elastic similarity measures, including Dynamic Time Warping (DTW) [9] and its differentiable extension (soft-DTW) [10], have been employed to compare degradation trajectories [11]. However, these similarity measures operate at a global level, without localizing degradation-related phenomena. Interestingly, a large DTW score can signal deformation, without indicating which cycle region is responsible for it.

In industrial systems, degradation may progress heterogeneously: some cycle regions remain stable, others fluctuate randomly, and only a subset of changes carry meaningful diagnostic information. Identifying *when the system transitions between degradation stages* and *where in the cycle these transitions manifest*, is therefore a key challenge for degradation monitoring. To address this, we propose an unsupervised framework that combines elastic alignment with change-point detection. We first compute a reference barycenter and elastically align each observed cycle to it, then extract localized deformation trajectories that quantify regional shifts with respect to the reference regime. A critical-region detection step isolates structured deformation areas, and change-point detection pinpoints transition times on the corresponding trajectories. Our proposed framework localizes and quantifies structural changes without requiring stage annotations.

II. Background

This section introduces key methodological components of our proposed framework, including elastic alignment methods, reference barycenter construction, and change-point detection algorithms, later combined for degradation stage identification.

A. Elastic alignment methods

Unlike point-wise distances, elastic measures allow nonlinear time warping, enabling alignment of unequal-length sequences despite local shifts or speed variations.

1) Dynamic Time Warping (DTW): DTW [9] allows elastic matching between two time series $\mathbf{x} = (x_1, \dots, x_{L_1})$ and $\mathbf{y} = (y_1, \dots, y_{L_2})$, not necessarily of equal lengths. DTW computes the minimal cumulative alignment cost over all admissible warping paths.

The cumulative cost matrix $D \in \mathbb{R}^{(L_1+1)\times(L_2+1)}$ for two sequences of lengths L_1 and L_2, is recursively defined as

$$D(i,j) = \delta(x_i, y_j) + \min \begin{cases} D(i-1,j) \\ D(i,j-1) \\ D(i-1,j-1) \end{cases} \tag{1}$$

for $1 \le i \le L_1$ and $1 \le j \le L_2$, with boundary conditions $D(0,0) = 0$ and $D(i,0) = D(0,j) = \infty$. The cost $\delta(x_i, y_j)$ is commonly chosen as the Manhattan or the squared Euclidean distance.

The DTW distance is given by $D(L_1, L_2)$. The optimal alignment is represented by a path $\pi = \{(i_\ell, j_\ell)\}_{\ell=1}^{L} \subset \{1, \dots, L_1\} \times \{1, \dots, L_2\}$, starting at $(1,1)$ and ending at (L_1, L_2), satisfying monotonicity constraints. Each pair $(i,j) \in \pi$ indicates that x_i is aligned with y_j.

2) Soft Dynamic Time Warping (soft-DTW): Soft-DTW [10] is a differentiable variant of DTW, obtained by replacing the hard-minimum in Equation (1) with a soft-min operator parametrized by a smoothing parameter γ:

$$\min_\gamma(a_1, \dots, a_n) = -\gamma \log \left(\sum_{i=1}^{n} e^{-a_i/\gamma} \right), \tag{2}$$

where $\min_\gamma$ converges to the standard minimum as $\gamma \to 0$.

B. Soft-DTW barycenter averaging (softDBA)

Given a collection of b time series cycles $\{\mathbf{x}^{(k)}\}_{k=1}^{b}$, a barycenter can be defined as a representative sequence $\mathbf{z} = (z_1, \dots, z_m) \in \mathbb{R}^m$, minimizing the average alignment cost.

Leveraging soft-DTW, notated as $\mathrm{sDTW}_\gamma(\mathbf{x}, \mathbf{y})$, of Equation 2 for the similarity measure, the barycenter is derived from solving

$$\tilde{\mathbf{z}} = \arg\min_{\mathbf{z}} \frac{1}{b} \sum_{k=1}^{b} \mathrm{sDTW}_\gamma\big(\mathbf{z}, \mathbf{x}^{(k)}\big), \tag{3}$$

where $\gamma > 0$ is the smoothing parameter. This parametrization enables a differentiable measure with respect to its inputs, suitable for gradient-based optimization. This procedure for barycenter construction is commonly referred to as soft dynamic time warping barycenter averaging (softDBA) [10], originally proposed without the smooth relaxation in [12].

C. Change-point detection with PELT

Change-point detection aims to segment a time series input $\mathbf{x} = (x_1, \dots, x_L)$ into contiguous segments, such that key statistical properties are preserved per segment.

Let $\mathcal{S} = \{t_1, \dots, t_K\}$ denote the set of change-point indices, where both the number K and the locations of change-points must be estimated. The optimal segmentation can be derived by minimizing

$$\min_{\mathcal{S}} \left\{ \sum_{k=1}^{K+1} \mathcal{C}\big(\mathbf{x}_{(t_{k-1}+1):t_k}\big) + \beta\, K \right\}, \tag{4}$$

where $\mathcal{C}(\cdot)$ is a segment-wise cost function, $\beta > 0$ is a penalty parameter controlling the number of detected change-points, $t_0 = 0$ and $t_{K+1} = L$. The Pruned Exact Linear Time (PELT) algorithm [13] efficiently solves Equation 4 with linear computational cost under mild conditions.

III. METHOD

The proposed method consists of four main steps: the computation of a reference barycenter, the construction of localized deformation trajectories, the identification of critical regions, and the detection of change-points in critical deformations.

Let $\mathcal{X} = \{X^{(k)}\}_{k=1}^{N}$ denote the lifetime cycles, with each cycle a sequence of time-value vectors

$$X^{(k)} = \Big(\big(t_1^{(k)}, x_1^{(k)}\big), \ldots, \big(t_{n_k}^{(k)}, x_{n_k}^{(k)}\big) \Big), \tag{5}$$

sampled at possibly irregular time stamps $t_j^{(k)}$. The objective is to detect degradation or deformation events across cycles. Formally, we seek cycle indices

$$\widehat{\mathcal{S}} = \{\hat{k}_1, \ldots, \hat{k}_{\widehat{K}}\} \subset \{1, \ldots, N\},\ 1 \le \hat{k}_1 < \ldots < \hat{k}_{\widehat{K}} \le N, \tag{6}$$

where each $\hat{k}_j$ marks the boundaries of a deformation event and $\widehat{K}$ is the number of detected deformation onsets.

A. Computation of a reference barycenter

We first construct a reference cycle, considering the first b cycles, which are assumed to be representative of normal behavior. The barycenter is iteratively refined using soft-DBA [10], [12], which unlike a simple mean accounts for temporal variability across cycles. Computing a barycenter yields a representative reference cycle by averaging aligned cycles. The smoothing parameter γ of soft-DTW, controls the sharpness of the soft-alignments, offering a tunable trade-off between noise sensitivity and temporal fidelity.

Formally, prior to barycenter computation, all cycles are z-normalized using the mean, standard deviation, and length of the first cycle. This ensures a consistent scaling across cycles while preserving the unsupervised setting.

The barycenter is then initialized by uniformly interpolating the first cycle $X^{(1)}$ to a reduced number of points m, by keeping one point out of every five, resulting in $m = \lceil n_1/5 \rceil$, where n_1 denotes the length of the first cycle. This produces the initial reference sequence of time-value vectors $Z = \big((t_1, z_1), \ldots, (t_m, z_m)\big)$.

The reference barycenter $\tilde{Z}$ is then obtained by solving Equation (3) using gradient-based optimization, over a subset of the first b cycles $\{X^{(k)}\}_{k=1}^{b}$, represented as time-value pairs, and using Z for the barycenter initialization.

B. Construction of deformation trajectories

In this step, we convert the alignment between each cycle and the reference barycenter into localized deformation trajectories, that quantify how specific regions of the cycle shift over successive cycles. While the reference barycenter $\tilde{Z}$ is computed using soft-DTW in the joint time-value space, the deformation trajectories are extracted from a classical DTW alignment performed on the values only.

Therefore, for each cycle k, we compute a standard DTW alignment of Equation (1) between the barycenter values $\tilde{\mathbf{z}} = (\tilde{z}_1, \ldots, \tilde{z}_m)$ and the cycle values $\mathbf{x}^{(k)} = (x_1^{(k)}, \ldots, x_{n_k}^{(k)})$. This produces an optimal alignment path $\pi^{(k)} \subset \{1, \ldots, m\} \times \{1, \ldots, n_k\}$.

a) Localized temporal deformation: We next proceed to the construction of the deformation trajectories that form the basis for the subsequent detection of significant structural changes. From the alignment path $\pi^{(k)}$, we define, for each barycenter point $i \in \{1, \ldots, m\}$, the associated set

$$L_i^{(k)} = \{\, j \mid (i, j) \in \pi^{(k)} \,\}.$$

Using the timestamps of the aligned points, we compute a local temporal deformation measure as

$$\tau_i^{(k)} = t_i - \frac{1}{\big|L_i^{(k)}\big|} \sum_{j \in L_i^{(k)}} t_j^{(k)}. \tag{7}$$

This quantity measures the deviation between the reference time t_i and the mean time of the cycle points aligned to it. If $|\tau_i^{(k)}|$ is large, it indicates that the portion of the cycle around the index i had to be significantly displaced in time to align with the reference, showing a local deformation. The sign of $\tau_i^{(k)}$ specifies the direction of the shift, with a positive value indicating that the aligned region in cycle k occurs earlier than in the reference, and vice-versa.

b) Final deformation trajectories: For each cycle k, we therefore produce global degradation trajectories $\boldsymbol{\tau}^{(k)} = \big(\tau_1^{(k)}, \ldots, \tau_m^{(k)}\big)$, which describe the temporal shifts of all reference regions within that cycle.

Conversely, for a fixed barycenter point i, we denote by $\boldsymbol{\tau}_i = \big(\tau_i^{(1)}, \ldots, \tau_i^{(N)}\big)$ the local deformation trajectory associated with the i-th barycenter point throughout all cycles. This representation disentangles the spatial position within a cycle (index i) from its temporal evolution across cycles (index k), enabling the analysis of region-specific deformation dynamics. The top plot of Figure 3 displays deformation trajectories capturing key degradation dynamics.

C. Identification of the critical regions

Interestingly, not all barycenter points carry meaningful degradation information. Some correspond to stable or noisy cycle regions, i. e., regions that do not include structured information about ongoing deformations, as illustrated in Figure 1. While certain parts (e. g., middle-right regions) appear locally distorted, their deformation trajectories $\boldsymbol{\tau}_i$ do not exhibit a clear trend across cycles. In contrast, other regions (e. g., beginning of the cycles) show structured shifts, with monotonic or piecewise-linear evolution in deformation trajectories. Analyzing $\boldsymbol{\tau}_i$, thus provides a direct characterization of how each region deforms across cycles, and trajectories with structured, trend-like behavior possibly reflect degradation mechanisms.

We therefore aim to identify a subset of "critical" barycenter points that are most informative for degradation stage detection, leveraging the deformation trajectories.

1) Segmentation-based relevance criterion: Next, we assess to what extent each deformation trajectory $\boldsymbol{\tau}_i$ can be modeled as a piecewise-linear signal across cycles. For each deformation trajectory $\boldsymbol{\tau}_i$, we compute a segmentation cost C_i, using the PELT algorithm of Equation (4), with a piecewise-linear segment model and fixed β to ensure comparability. Under

this model, each segment is approximated by a linear function of the cycle index. Trajectories dominated by noise require many short segments to be adequately fitted, producing higher segmentation costs. Conversely, trajectories with structured, approximately monotonic, or trend-like behavior are modeled with fewer linear segments, leading to lower costs.

Before computing the costs, we apply a median filter (with window size 3) to each deformation trajectory to reduce the effect of isolated outliers that could affect the PELT results, and we z-normalize each series independently.

2) Relevance scoring and region selection: We then define a relevance score for each barycenter point as

$$r_i = \frac{1}{\sqrt{C_i}}. \tag{8}$$

so that trajectories with lower segmentation cost (i. e., better approximated by a piecewise-linear model) receive higher relevance scores.

Constant deformation trajectories produce $C_i = 0$ and in this case, we set $C_i = \max_j C_j$ to penalize the absence of relevant information. To make the scores comparable across barycenter points, we also normalize them as $\tilde{r}_i = r_i / \max_j r_j$.

Our objective is to select the barycenter points with the highest relevance scores. However, neighboring points often correspond to the same deformation region, which may result in multiple adjacent high scores. To avoid redundant selections, we apply a peak detection procedure to the normalized relevance signal $\tilde{r} = (\tilde{r}_1, \ldots, \tilde{r}_m)$, defined as

$$\mathcal{I} = \{i \mid \tilde{r}_i \text{ is a peak with prominence} \geq \lambda\},$$

where the prominence of a peak is defined as the vertical distance between the peak and its highest surrounding minimum, and $\lambda \in (0, 1)$ is a minimum prominence threshold.

D. Final change-point detection

Finally, the selected deformation trajectories $\{\boldsymbol{\tau}_i\}_{i \in \mathcal{I}}$ are jointly used to identify deformation events across cycles. The change-point detection objective of Equation (4) is applied to the multivariate signal of the selected trajectories, indexed by the cycle number k. This produces the final deformation transition indices $\widehat{\mathcal{S}}$ defined in Equation (6).

In this final step, the segment-wise cost $\mathcal{C}(\cdot)$ corresponds to a joint piecewise-linear model across all selected trajectories, and the penalty parameter β controls the number of detected change-points (selected via grid search, as outlined in Section IV-A). Unlike the previous univariate use of PELT for region relevance scoring, this multivariate segmentation jointly exploits correlations across degradation-sensitive regions.

IV. Experimental Results

A. Experimental setup

1) Datasets: We construct synthetic sequences of cycles derived from a real cycle extracted from the BatteryLife dataset [14]. This dataset contains voltage, current, and capacity measurements, collected during charge-discharge cycles for different battery types. In terms of our study, only the voltage cycles are used for the experiments.

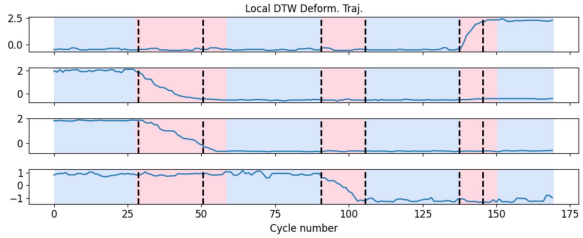

Fig. 2. Example deformation trajectories for our method. PELT segmentation in black dotted lines, true change-points via background color changes.

TABLE I
Performance comparison of considered methods.

Method	Metrics		
	F1-score	Precision	Recall
Stat. Features	$0.389_{\pm 0.149}$	$0.318_{\pm 0.135}$	$0.532_{\pm 0.213}$
Global Eucl. Traj.	$\underline{0.484}_{\pm 0.191}$	$\underline{0.523}_{\pm 0.290}$	$\underline{0.570}_{\pm 0.227}$
Global DTW Traj.	$0.349_{\pm 0.134}$	$0.351_{\pm 0.249}$	$0.564_{\pm 0.284}$
Local DTW Deform. Traj.	$\mathbf{0.738}_{\pm 0.211}$	$\mathbf{0.787}_{\pm 0.242}$	$\mathbf{0.748}_{\pm 0.211}$

2) Synthetic deformation generator: To generate a synthetic sequence, we first randomly select a battery from the BatteryLife dataset and extract its first cycle.We then insert three random deformation events per sequence, over randomly chosen and mutually disjoint intervals of cycles (i. e., two deformation events cannot occur during the same cycles, see Figure 2). For each deformation event, the affected axis (time or amplitude), shift direction (positive or negative), and total magnitude of the shift, are randomly selected and applied gradually, across the selected cycle interval. To make the task more complex, we add random, non-temporally consistent degradations and Gaussian white noise (30 dB). For each experiment, parameters are selected via grid search on 30 sequences of cycles, that are subsequently evaluated on 1000 held-out test sequences. All sequences consist of 170 cycles.

3) Baselines: We compare the proposed method (Local DTW Deform. Traj. in results tables) with three baseline approaches. The first approach (Stat. Features) uses simple cycle-level statistical features, including the length, mean, and standard deviation of each cycle. The second (Global Eucl. Traj.) and third (Global DTW Traj.) approaches rely on global similarity measures, using the Euclidean distance and the DTW distance, respectively, between each cycle and the reference barycenter. For the Euclidean distance, cycles are resampled to match the length of the barycenter prior to computation. The reference barycenter is computed using the first 10 cycles of the sequences (which are free of deformation events but not of random deformations).

4) Evaluation Metrics: We use the standard precision, recall and F1-score, with a margin $M = 5$.

B. Results

We next present results, including: (i) a comparison of our method to baselines, (ii) a robustness analysis under increasing levels of random deformations, and (iii) a qualitative evaluation on a real sequence of cycles from BatteryLife dataset.

TABLE II
ROBUSTNESS ANALYSIS TO RANDOM SHIFTS (IN F1-SCORE).

	Probability of random shift		
Method	$p = 0.1$	$p = 0.5$	$p = 1.0$
Stat. Features	$0.389_{\pm 0.149}$	$0.388_{\pm 0.132}$	$0.389_{\pm 0.124}$
Global Eucl. Traj.	$\underline{0.484}_{\pm 0.191}$	$\underline{0.434}_{\pm 0.172}$	$\underline{0.412}_{\pm 0.163}$
Global DTW Traj.	$0.349_{\pm 0.134}$	$0.336_{\pm 0.118}$	$0.340_{\pm 0.113}$
Local DTW Deform. Traj.	$\mathbf{0.738}_{\pm 0.211}$	$\mathbf{0.720}_{\pm 0.216}$	$\mathbf{0.686}_{\pm 0.213}$

1) Methods comparison: Results are reported in Table I. The proposed method outperforms all baselines in both precision and recall, achieving the best F1-score, demonstrating a balanced trade-off between the two. These results indicate that our approach is capable of accurately recovering the start and end indices of the simulated deformation events. The grid search analysis further showed that increasing the PELT penalty parameter β, as well as increasing the minimum peak prominence (for our method), improves precision at the expense of recall and vice-versa. Beyond a certain threshold, however, the penalty becomes too restrictive, and no change-points are detected. Finally, the Euclidean distance is supposed to be unable to detect time shifts. Yet, in our setting, cycles are uniformly resampled to match the barycenter length, which evenly distributes shifted points along the cycle, enabling partial detection of time shifts for the Euclidean-based method.

An example of deformation trajectories is displayed in Figure 2. Our proposed method successfully selects the most informative deformation trajectories and, through their analysis, retrieves all events, though with some boundary inaccuracies.

2) Robustness analysis: In this experiment, we evaluate the robustness of the method by varying the probability of random single-cycle shifts (Table II). Our proposed method demonstrates the strongest robustness to additional random deformations. We can see that even in cases of significant random deformations, our method continues to perform very well. Selecting critical zones in the barycenter therefore allows us to identify the important areas.

3) Visual inspection of a real example: The results of our method on a real battery sample (of Figure 1), are illustrated in Figure 3. The bottom plot shows that the selected barycenter points are correctly located within critical zones associated with battery degradation. The deformation trajectories of the selected barycenter points also provide a clear visualization of the variations in degradation speed, highlighted in Figure 1. This confirms our method's ability to capture both spatial localization and temporal evolution of degradation.

V. CONCLUSION

In this work, we introduced a novel method to identify deformation variations across cycles using local comparisons with an elastic reference barycenter. Experiments demonstrated that tracking specific barycenter points reveals local deformations more effectively, compared to global distances or simple statistical features. Our approach can serve as a tool to better understand degradation mechanisms in manufacturing processes and support their refinement. Future work includes

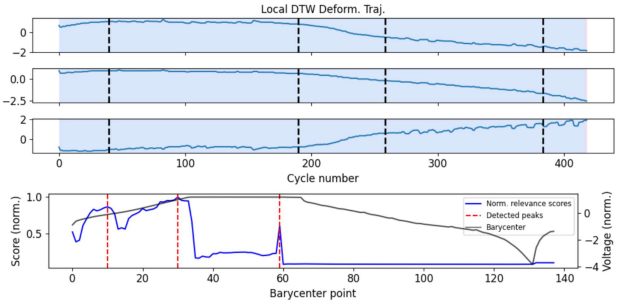

Fig. 3. Top: Deformation trajectories for our method on real lifetime voltage cycles (sample of Fig. 1). Bottom: Superposition of normalized relevance scores (blue), computed barycenter (grey) and selected barycenter points (red).

improved selection of barycenter points that remain robust under increasing noise and replacement of isolated points with regions that are properly-weighted in change-point detection.

REFERENCES

[1] N. Li, Y. Lei, X. Liu, T. Yan, and P. Xu, "Machinery health prognostics with multimodel fusion degradation modeling," *IEEE Transactions on Industrial Electronics*, vol. 70, no. 11, pp. 11764–11773, 2022.

[2] Y. Che, X. Hu, X. Lin, J. Guo, and R. Teodorescu, "Health prognostics for lithium-ion batteries: mechanisms, methods, and prospects," *Energy & Environmental Science*, vol. 16, no. 2, pp. 338–371, 2023.

[3] N. F. Thornhill, M. S. Choudhury, and S. L. Shah, "The impact of compression on data-driven process analyses," *Journal of Process Control*, vol. 14, no. 4, pp. 389–398, 2004.

[4] J. Lyu, R. Ying, N. Lu, and B. Zhang, "Remaining useful life estimation with multiple local similarities," *Engineering Applications of Artificial Intelligence*, vol. 95, p. 103849, 2020.

[5] C. Truong, L. Oudre, and N. Vayatis, "Selective review of offline change point detection methods," *Signal Processing*, vol. 167, p. 107299, 2020.

[6] S. J. Bae, T. Yuan, S. Ning, and W. Kuo, "A bayesian approach to modeling two-phase degradation using change-point regression," *Reliability engineering & system safety*, vol. 134, pp. 66–74, 2015.

[7] R. Wang, M. Zhu, X. Zhang, and H. Pham, "Lithium-ion battery remaining useful life prediction using a two-phase degradation model with a dynamic change point," *Journal of Energy Storage*, vol. 59, p. 106457, 2023.

[8] T. Wang, J. Yu, D. Siegel, and J. Lee, "A similarity-based prognostics approach for remaining useful life estimation of engineered systems," in *2008 international conference on prognostics and health management*, pp. 1–6, IEEE, 2008.

[9] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE transactions on acoustics, speech, and signal processing*, vol. 26, no. 1, pp. 43–49, 1978.

[10] M. Cuturi and M. Blondel, "Soft-dtw: a differentiable loss function for time-series," in *International conference on machine learning*, pp. 894–903, PMLR, 2017.

[11] L. Huang, L. Gong, Y. Chen, D. Li, and G. Zhu, "Trajectory similarity matching and remaining useful life prediction based on dynamic time warping," *Mathematical Problems in Engineering*, vol. 2022, no. 1, p. 5344461, 2022.

[12] F. Petitjean, A. Ketterlin, and P. Gançarski, "A global averaging method for dynamic time warping, with applications to clustering," *Pattern recognition*, vol. 44, no. 3, pp. 678–693, 2011.

[13] R. Killick, P. Fearnhead, and I. A. Eckley, "Optimal detection of changepoints with a linear computational cost," *Journal of the American Statistical Association*, vol. 107, no. 500, pp. 1590–1598, 2012.

[14] R. Tan, W. Hong, J. Tang, X. Lu, R. Ma, X. Zheng, J. Li, J. Huang, and T.-Y. Zhang, "BatteryLife: A Comprehensive Dataset and Benchmark for Battery Life Prediction," in *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, (Toronto ON Canada), pp. 5789–5800, ACM, Aug. 2025.