

Artificial intelligence for gait analysis

La Soirée de l'IA

Laurent Oudre

laurent.oudre@ens-paris-saclay.fr

May, 3rd 2023



école
normale
supérieure
paris—saclay

université
PARIS-SACLAY



Université
Paris Cité



Inserm

Contents

1. Introduction
2. Choosing the adequate cost function
3. Search method
4. Choosing the number of change-points
5. Conclusion

The Centre Borelli



The Centre Borelli

Fusion of two labs :


- ▶ The Centre de mathématiques et de leurs applications (CMLA) : applied mathematics for the study of complex phenomena and data
- ▶ The Cognition & Action Group (CognacG) : quantification and study of human and animal behavior

Main scientific questions

How to quantify the human behavior

- ▶ Adventure launched since 2012 : interdisciplinary collaboration between mathematicians, physicians, neuroscientists, engineers, biologists, etc...
- ▶ Implementation of measurement chains “pipelines”, platforms and intelligent tools but also of procedures for analysis, measurement and processing of data
- ▶ Creation of tools for diagnostic assistance, inter-individual comparison and longitudinal follow-up
- ▶ Integration into a clinical environment and interaction between algorithms and medical/neuroscience experts

Study of locomotion



Why study the locomotion ?

- ▶ Most common dynamic human activity
- ▶ Reveals a large number of neurological, orthopedic, rheumatological disorders...
- ▶ Strong influence on daily life: risk of falling, frailty, mobility, dependence...

Study of locomotion



How to study locomotion?

- ▶ Historically : clinical examination by the physician, functional tests, clinical questionnaires

+	Ease of execution, clinician expertise
-	Lack of precision, difficult to compare two sessions

- ▶ Platforms for studying locomotion : instrumented mats, video/optical systems

+	Very precise, extraction of a large number of parameters, objective quantification
-	High cost, difficult to implement

General principles

★ Objective quantification of locomotion

→ Use of sensors and physiological measures

★ Longitudinal follow-up and inter-individual comparison

→ Need for a fixed protocol

★ Experimentation on the field

→ Lightweight sensors and fully automatic device for consultation and routine use

★ Clean and quality data



→ Control of the whole measurement chain, robust and reproducible algorithms

★ Willingness to capture clinician expertise


→ Clinical annotations and metadata

Protocol and sensors


- ▶ Comfort speed protocol: stop (6 sec), walk forward (10 m), turn around, return, stop
- ▶ Four wireless inertial units: left foot, right foot, lower back, head
- ▶ Nine signals per sensor : linear acceleration (3D), angular velocity (3D), magnetic field (3D)




Signals



Healthy subject



Mild neurological disease



Severe neurological disease


Non-stationary signals

→ How to detect the different regimes (walk, U-turn, etc.)?

Notion of stationarity

- ▶ Stationarity: statistical properties of the signal do not change over time
- ▶ Wide-sense stationarity is one of the most common assumption in signal processing
- ▶ In several contexts, some changes occur in the data over time: most real time series are not stationary
- ▶ Change-point detection consists in estimating as precisely as possible the times where these changes occur
- ▶ Two main usecases:
 - ▶ Usecase 1: Retrieve these breakpoints to detect when the changes occurred
 - ▶ Usecase 2: Use the breakpoint information to divide the time series into smaller stationary signals

Problem statement: Change-Point Detection




Change-Point Detection

Given a time series x , retrieve the times (t_1, \dots, t_K) where a significant change occurs

- ▶ Necessitates to estimate both the change-points and the number of changes K
- ▶ Highly depends on the meaning given to *change*

Problem statement




- ▶ Let assume that signal $x[n]$ undergoes abrupt changes at times

$$\mathcal{T}^* = (t_1^*, \dots, t_{K^*}^*)$$

- ▶ Goal: retrieve the number of change-points K^* and their times \mathcal{T}^*
- ▶ Two assumptions: offline segmentation (but can easily be adapted to online setting) [Truong et al., 2020] and known number of changes K (will be discussed later)

Problem statement

$$(\hat{t}_1, \dots, \hat{t}_K) = \underset{(t_1, \dots, t_K)}{\operatorname{argmin}} \sum_{k=0}^K c(x[t_k : t_{k+1}])$$




Cost function $c(\cdot)$

- ▶ Measures the homogeneity of the segments
- ▶ Choosing $c(\cdot)$ conditions the type of change-points that we want to detect
- ▶ Often based on a probabilistic model for the data

Problem solving

- ▶ Optimal resolution with dynamic programming
- ▶ Approximate resolution (sliding windows...)



Contents

1. Introduction

2. Choosing the adequate cost function

3. Search method

4. Choosing the number of change-points

5. Conclusion

Cost function

$$(\hat{t}_1, \dots, \hat{t}_K) = \underset{(t_1, \dots, t_K)}{\operatorname{argmin}} \sum_{k=0}^K c(x[t_k : t_{k+1}])$$

Convention : $t_0 = 0, t_{K+1} = N$

- ▶ Function $c(\cdot)$ is characteristic of the notion of *homogeneity*
- ▶ The most common cost functions are linked to parametric probabilistic models: in this case change-points are defined as changes in the parameters of a probability density function [Basseville et al., 1993]
- ▶ Non parametric cost functions can also be introduced when no model is available


Change in mean

The most popular is indubitably the L2 norm [Page, 1955]

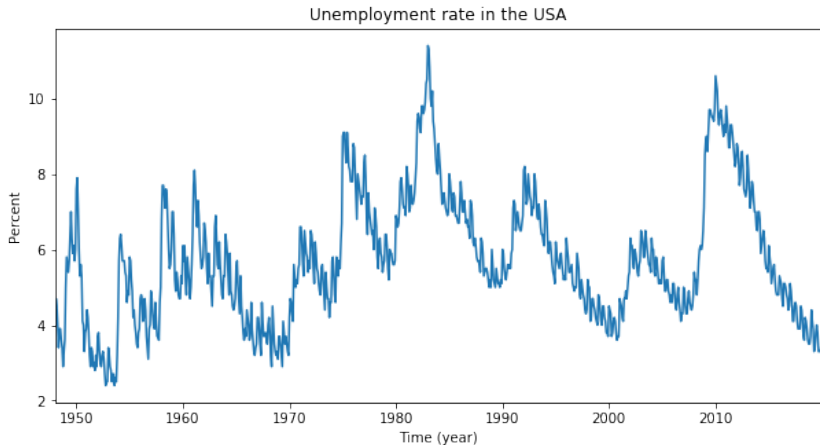
$$c_{L_2}(x[a : b]) = \sum_{n=a+1}^b \|x[n] - \mu_{a:b}\|_2^2$$

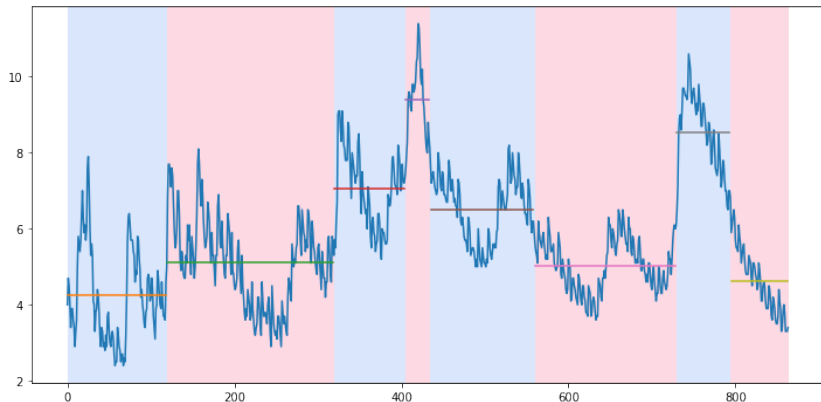
where $\mu_{a:b}$ is the empirical mean of the segment $x[a : b]$.

- ▶ Particular case of c_{ML} with Gaussian model with fixed variance
- ▶ Allows to detect changes in mean




Example



Example: Change-Point Detection with c_{L_2} 

$$K = 7$$

Example: Change-Point Detection with c_{L_2} 


$$K = 12$$

Change in slope and intercept

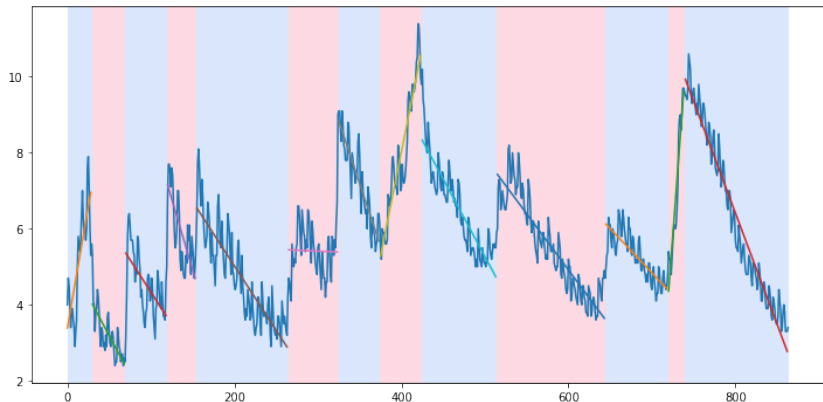
Change in slope and intercept can be handled in the general context of piecewise linear regression

$$c_{\text{linear}}(x[a : b]) = \min_{\alpha} \sum_{n=a+1}^b \left\| x[n] - \sum_{i=1}^M \alpha_i \beta_i[n] \right\|_2^2$$

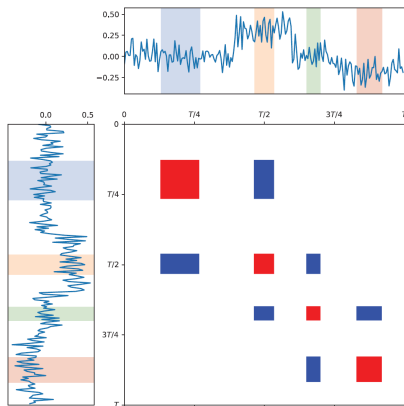
- ▶ Functions $\beta_1[n], \dots, \beta_M[n]$ are covariate functions and we seek for changes in the regression parameters
- ▶ Allows to detect changes in trend, seasonality, etc... [Bai et al., 1998]
- ▶ For slope and intercept, we choose $\beta_1[n] = 1$ and $\beta_2[n] = n$

Example: Change-Point Detection with C_{linear} 

$$K = 7$$

Example: Change-Point Detection with C_{linear}  $K = 12$

Supervised change-point detection



- ▶ Supervised approach based on metric learning [Truong et al., 2019]
- ▶ Use a few annotated examples to learn an adequate cost function
- ▶ Samples belonging to adjacent regimes form a set of constraints that can be used for metric learning
- ▶ Allows to relax the need for an off-the-shelf cost function

- ▶ C. Truong and L. Oudre. Supervised change-point detection with dimension reduction, applied to physiological signals. In NeurIPS Workshop on Learning from Time Series for Health, 2022.
- ▶ C. Truong, L. Oudre and N. Vayatis. Supervised kernel change point detection with partial annotations. In Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP), pages 3147-3151, Brighton, UK, 2019.

Contents

1. Introduction
2. Choosing the adequate cost function
- 3. Search method**
4. Choosing the number of change-points
5. Conclusion

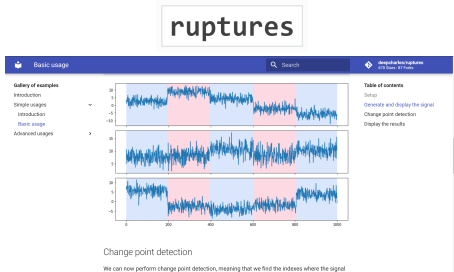
Search method

$$(\hat{t}_1, \dots, \hat{t}_K) = \underset{(t_1, \dots, t_K)}{\operatorname{argmin}} \sum_{k=0}^K c(x[t_k : t_{k+1}])$$

Convention : $t_0 = 0, t_{K+1} = N$

- ▶ Several methods can be used to solve this problem with a fixed K
- ▶ Optimal resolution with dynamic programming: find the true solution of the problem (but costly : Complexity of $\mathcal{O}(N^2)$)
- ▶ Approximated resolution with windows: test for one unique change-point on a window (necessitates some extra parameters and less precise)

The ruptures package



<https://centre-borelli.github.io/ruptures-docs/>

- ▶ Package in Python implementing most of the offline approaches for change-point detection
- ▶ Many cost functions, resolution algorithms and parametrizations
- ▶ More than 7 Million downloads!

C. Truong, L. Oudre, N. Vayatis. Selective review of offline change point detection methods. Signal Processing, 167:107299, 2020

Contents

1. Introduction
2. Choosing the adequate cost function
3. Search method
4. Choosing the number of change-points
5. Conclusion

Finding the number of change points

- ▶ In all previously described algorithms, the number of change-point K was supposed to be known
- ▶ In practice, this parameter is difficult to set: as such, the total cost $\mathcal{V}(\mathcal{T}, x)$ will always decrease when K increases...
- ▶ Three solutions
 - ▶ Use heuristics by testing several values of K
 - ▶ Use a penalized formulation of the CPD problem to seek for a compromise between reconstruction error and complexity

$$(\hat{t}_1, \dots, \hat{t}_K) = \underset{(t_1, \dots, t_K), K}{\operatorname{argmin}} \sum_{k=0}^K c(x[t_k : t_{k+1}]) + \beta K$$

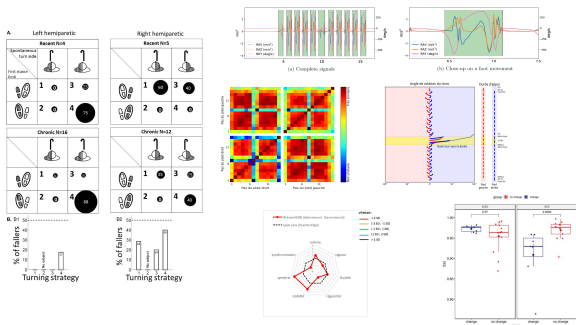
- ▶ Use supervised approaches from annotated signals [Truong et al., 2018]

- ▶ C. Truong and L. Oudre. Supervised change-point detection with dimension reduction, applied to physiological signals. In NeurIPS Workshop on Learning from Time Series for Health, 2022.
- ▶ C. Truong, L. Oudre and N. Vayatis. Penalty Learning for Changepoint Detection. In Proceedings of the European Signal Processing Conference (EUSIPCO), pages 1614-1618, Kos Island, Greece, 2017.

Contents

1. Introduction
2. Choosing the adequate cost function
3. Search method
4. Choosing the number of change-points
5. Conclusion

Use for biomedical research



- ▶ Study of U-turn for post-stroke patients [Barrois-Müller et al., 2017]
- ▶ Step analysis for multiple sclerosis patients [Vienne-Jumeau et al., 2020]
- ▶ Comparison of gait exercises through pattern matching techniques [Vienne-Jumeau et al., 2019]

Thank you for your attention