# DATA MINING 2

# Time Series - Shapelt/Motif Discovery

Riccardo Guidotti

a.a. 2019/2020

Slides edited from Keogh Eamonn’s tutorial

UNIVERSITÀ DI PISA

# Motif

# Time Series Motif Discovery

- Finding repeated patterns, i.e., pattern mining.
- Are there any repeated patterns, of length *m* in the TS?

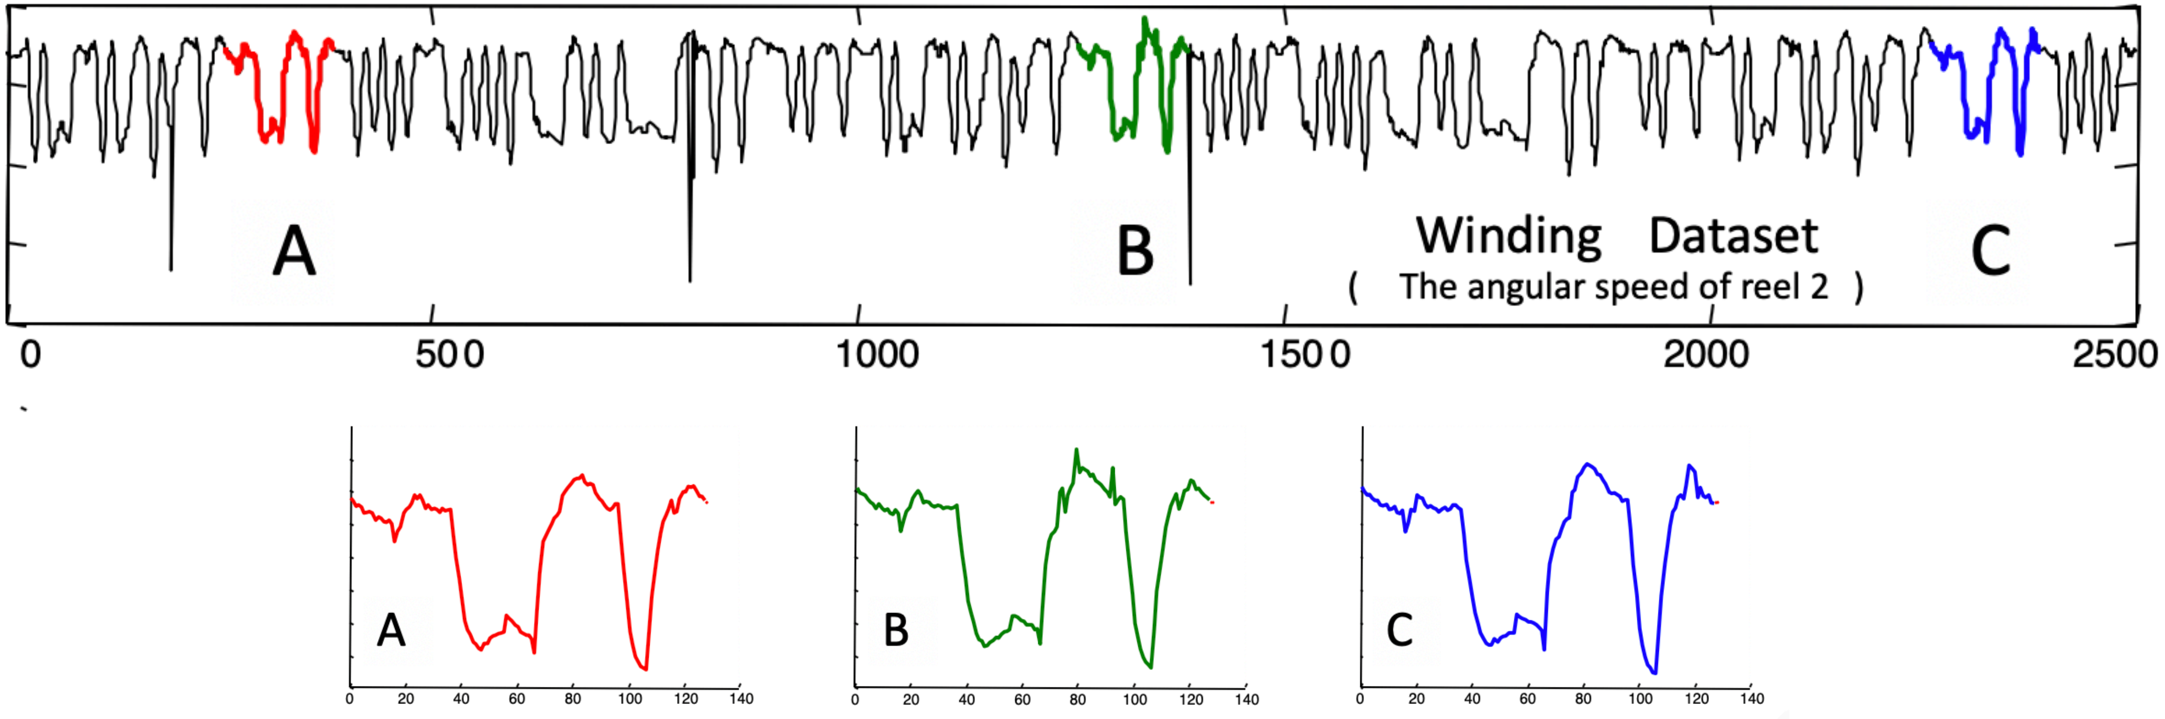

# Why Find Motifs?

- Mining **association rules** in TS requires the discovery of motifs. These are referred to as primitive shapes and frequent patterns.
- Several **TS classifiers** work by constructing typical prototypes of each class. These prototypes may be considered motifs.
- Many **TS anomaly detection** algorithms consist of modeling normal behavior with a set of typical shapes (which we see as motifs), and detecting future patterns that are dissimilar to all typical shapes.

# How do we find Motifs?

- Given a predefined motif length *m*, a brute-force method searches for motifs from all possible comparisons of subsequences.
- It is obviously very slow and computationally expensive.
- The most reference algorithm is based on a hot idea from bioinformatics, random projection* and the fact that SAX allows use to lower bound discrete representations of TSs.
- J Buhler and M Tompa. Finding motifs using random projections. In RECOMB'01. 2001.

# Example of the Motif Discovery Algorithm

- Assume that we have a time series T of length 1,000, and a motif of length 16, which occurs twice, at time $T_1$ and time $T_{58}$.

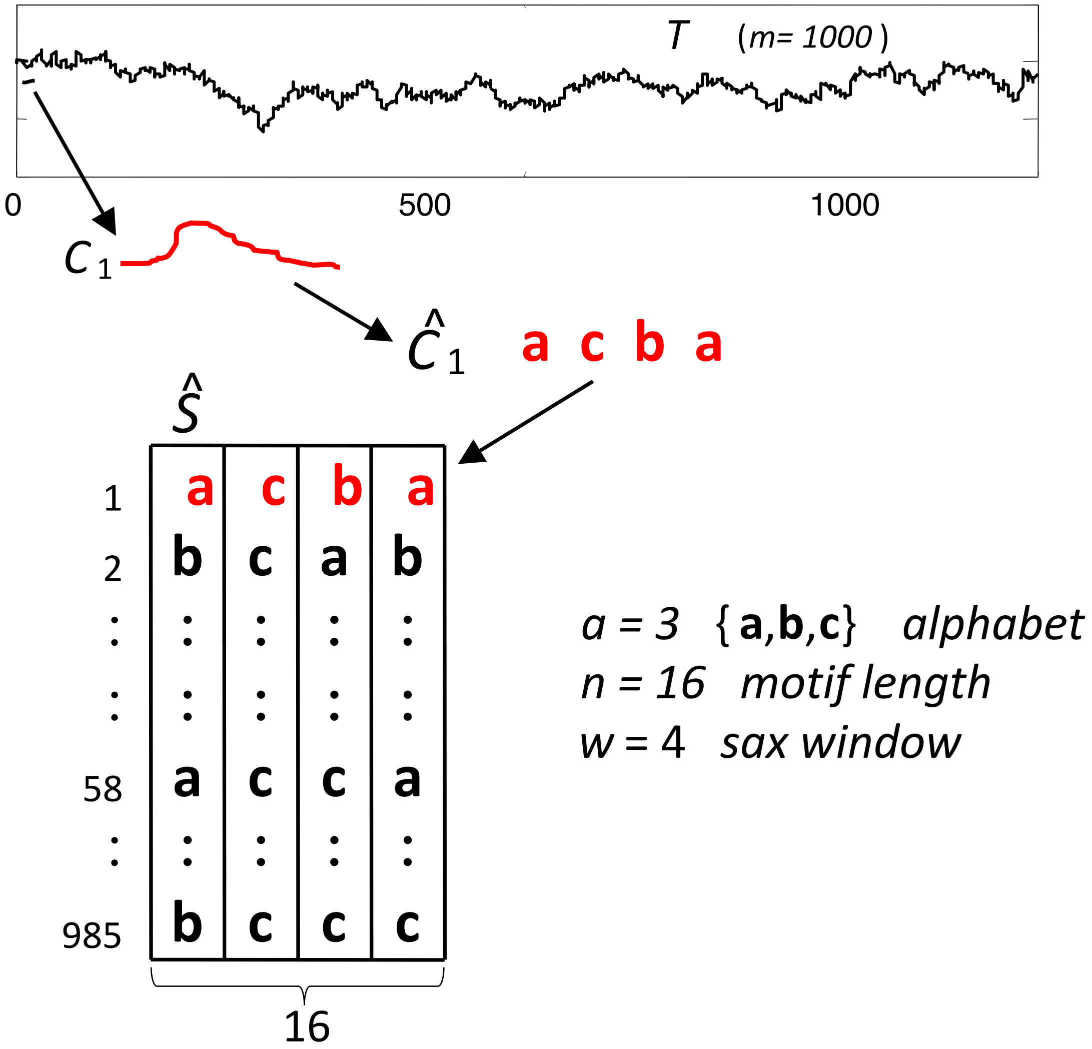

| | | | | |
|---|---|---|---|---|
| 1 | a | c | b | a |
| 2 | b | c | a | b |
| : | : | : | : | : |
| : | : | : | : | : |
| 58 | a | c | c | a |
| : | : | : | : | : |
| 985 | b | c | c | c |

*a* = *3* {**a**,**b**,**c**} *alphabet*

*n* = *16* *motif length*

*w* = 4 *sax window*

# Example of the Motif Discovery Algorithm

- A mask {1,2} was randomly chosen, so the values in columns {1,2} were used to project matrix into buckets.
- Collisions are recorded by incrementing the appropriate location in the collision matrix.

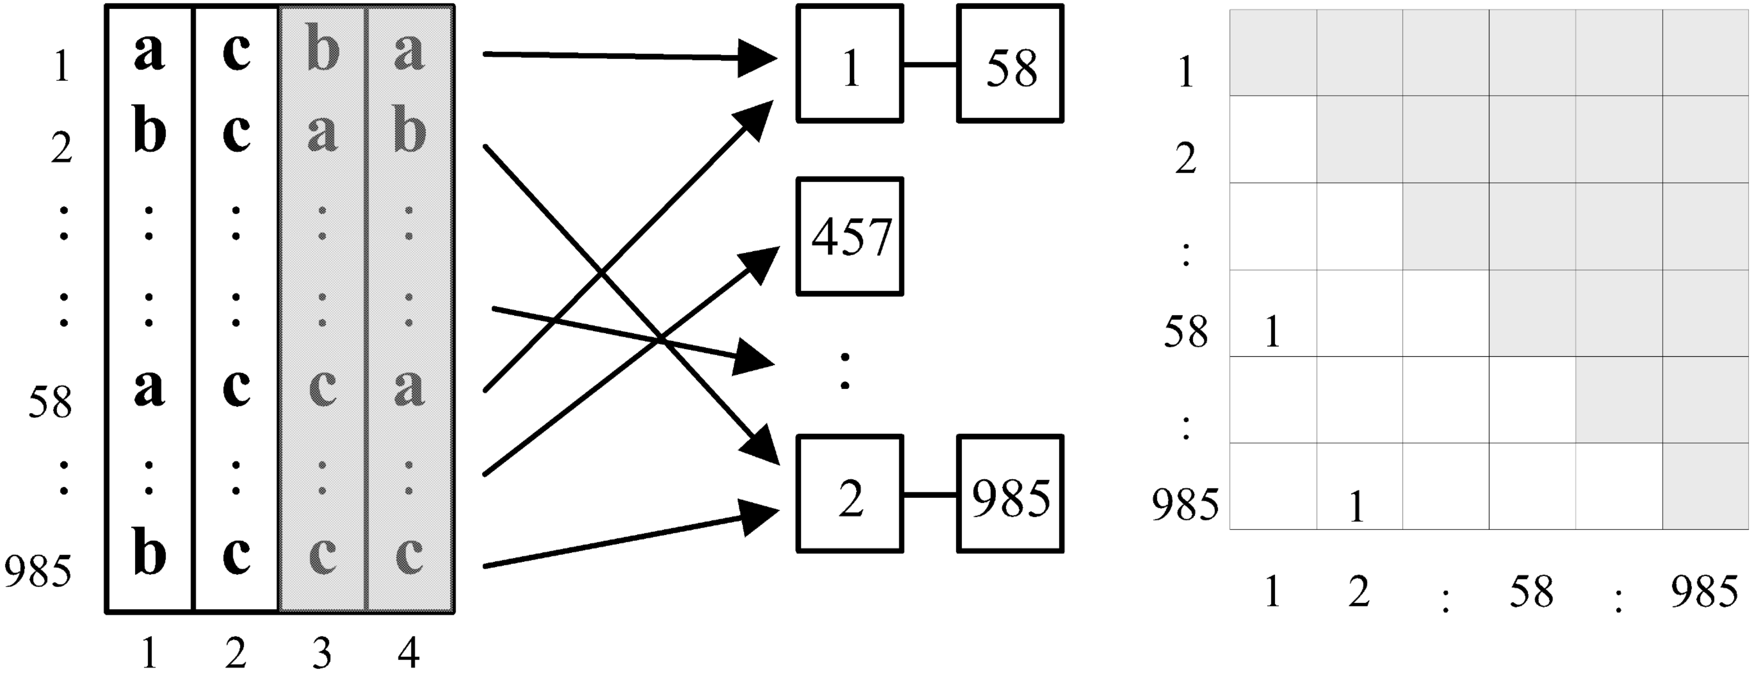

# Example of the Motif Discovery Algorithm

- A mask {2,4} was randomly chosen, so the values in columns {2,4} were used to project matrix into buckets.
- Once again, collisions are recorded by incrementing the appropriate location in the collision matrix.

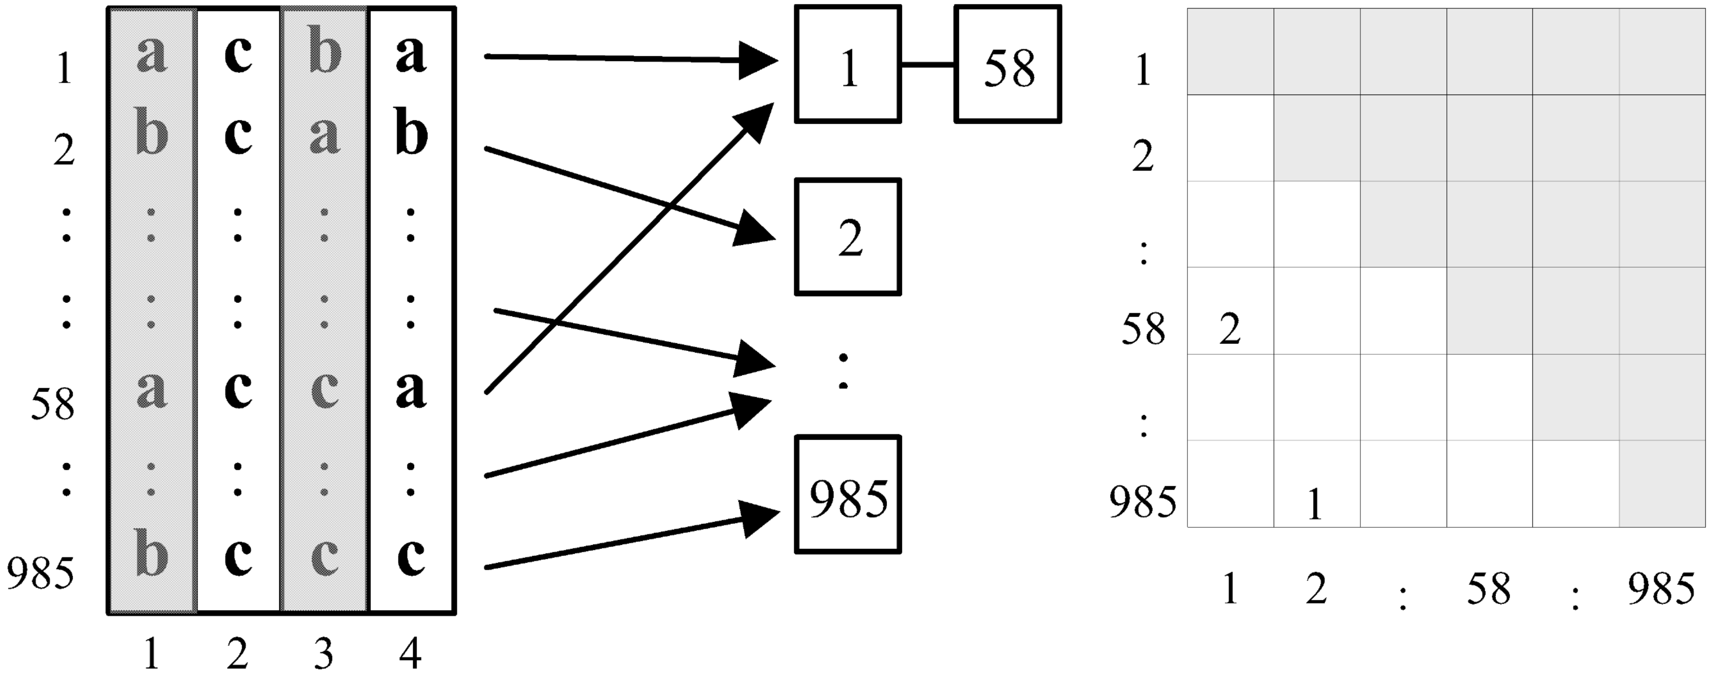

|  | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| 1 | a | c | b | a |
| 2 | b | c | a | b |
| : | : | : | : | : |
| : | : | : | : | : |
| 58 | a | c | c | a |
| : | : | : | : | : |
| 985 | b | c | c | c |

|  | 1 | 2 | : | 58 | : | 985 |
|---|---|---|---|---|---|---|
| 1 |  |  |  |  |  |  |
| 2 |  |  |  |  |  |  |
| : |  |  |  |  |  |  |
| 58 | 2 |  |  |  |  |  |
| : |  |  |  |  |  |  |
| 985 |  | 1 |  |  |  |  |

# Matrix Profile

- The Matrix Profile (MP) is a data structure that annotates a TS and can be exploited for many purposed: e.g. efficient Motif Discovery.
- Given a time series, T and a desired subsequence length, m.

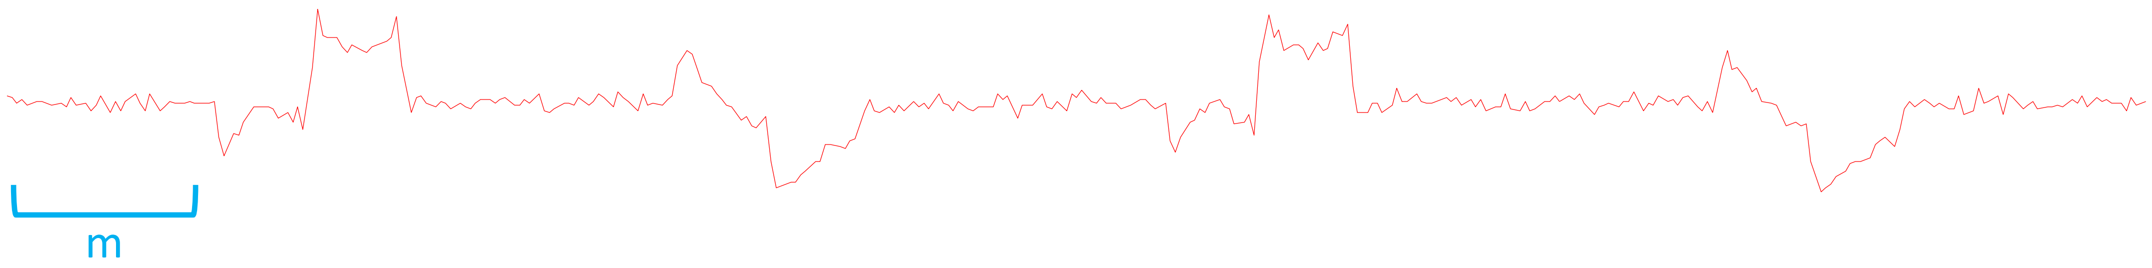

# Matrix Profile

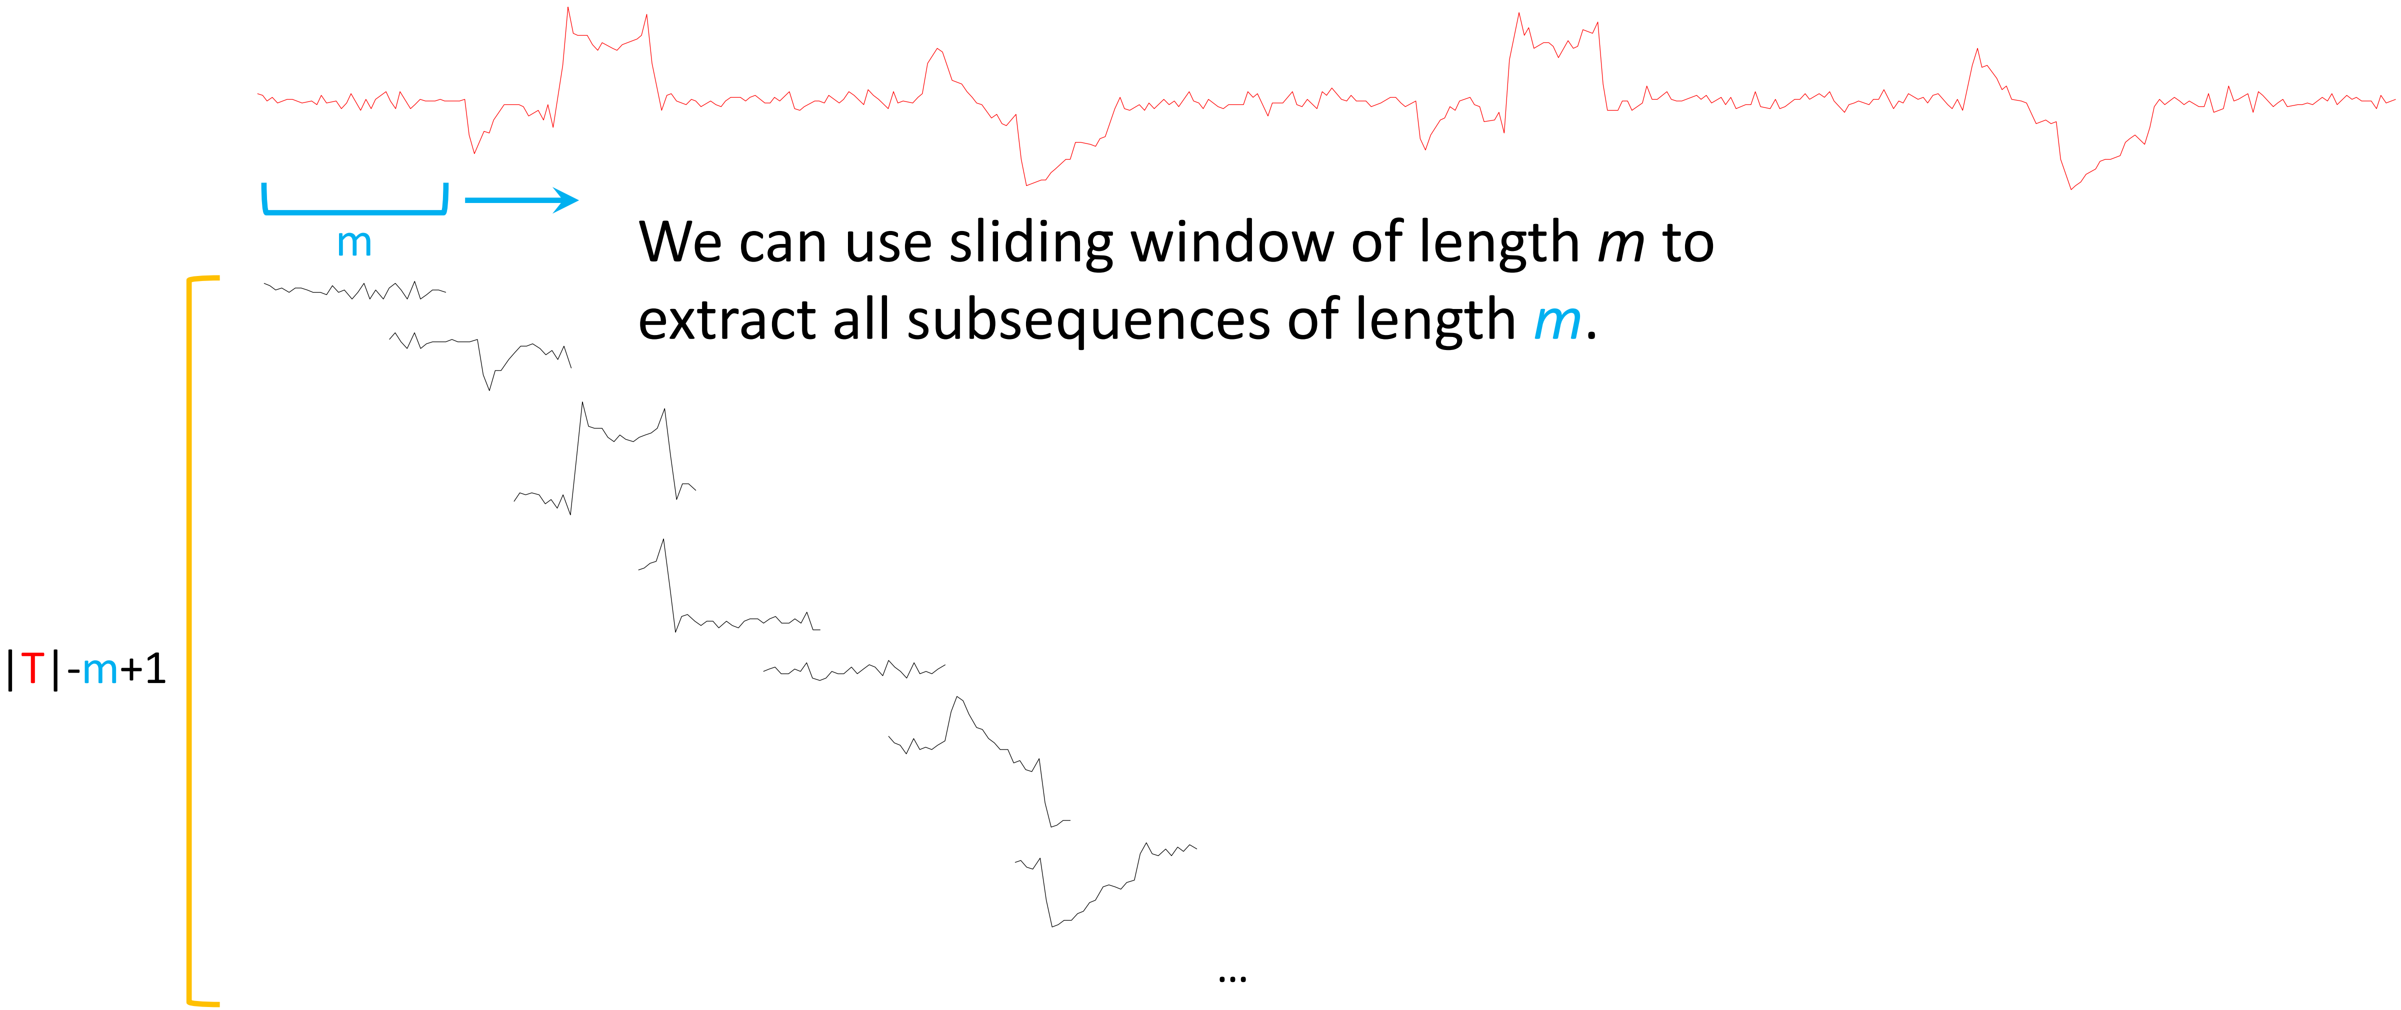

We can use sliding window of length *m* to extract all subsequences of length *m*.

# Matrix Profile

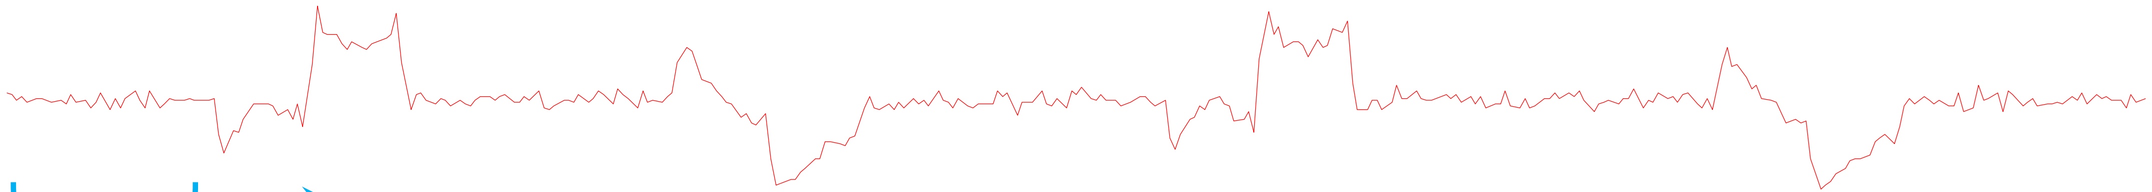

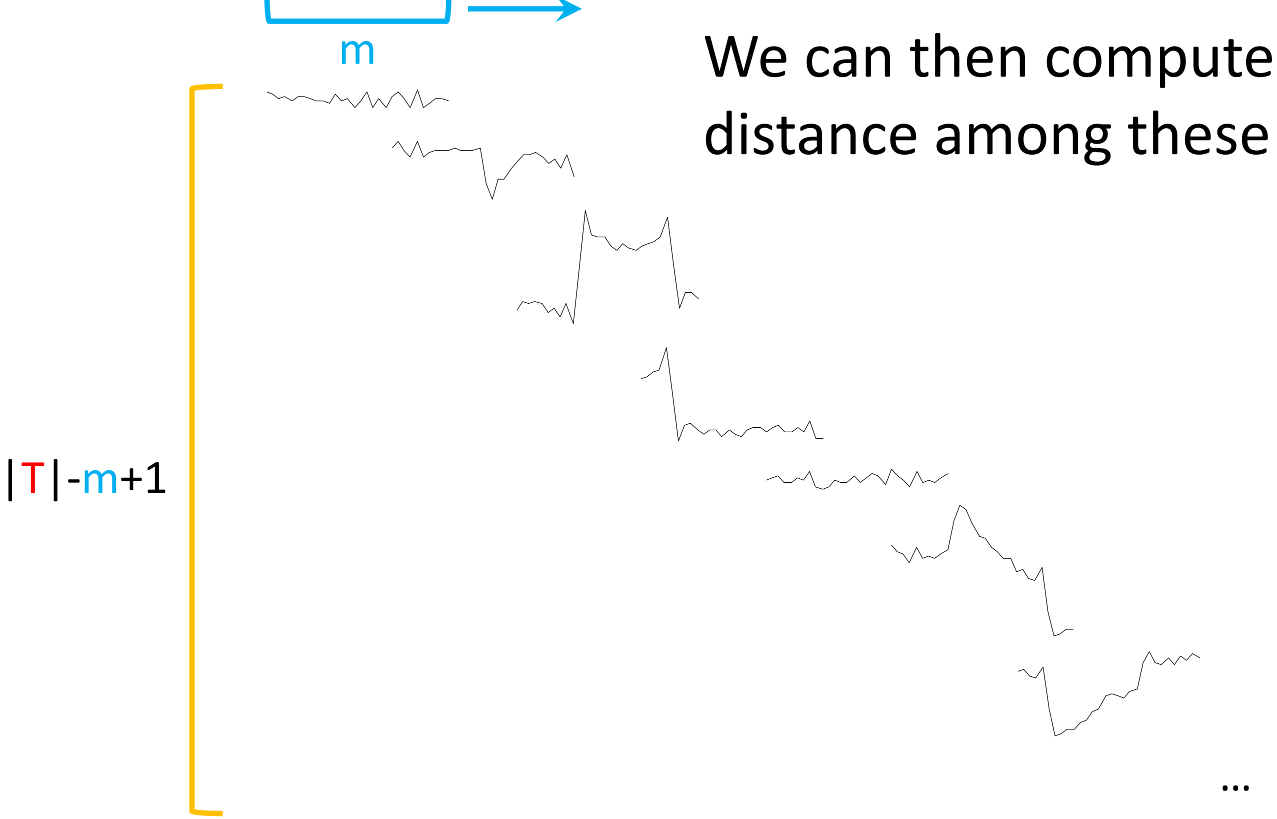

We can then compute the pairwise distance among these subsequences.

| | | | |
|---|---|---|---|
| 0 | 7.6952 | 7.7399 | … |
| 7.6952 | 0 | 7.7106 | … |
| 7.7399 | 7.7106 | 0 | … |
| … | … | … | … |

# Matrix Profile

- For each subsequence we keep only the distance with the closest nearest neighbor.

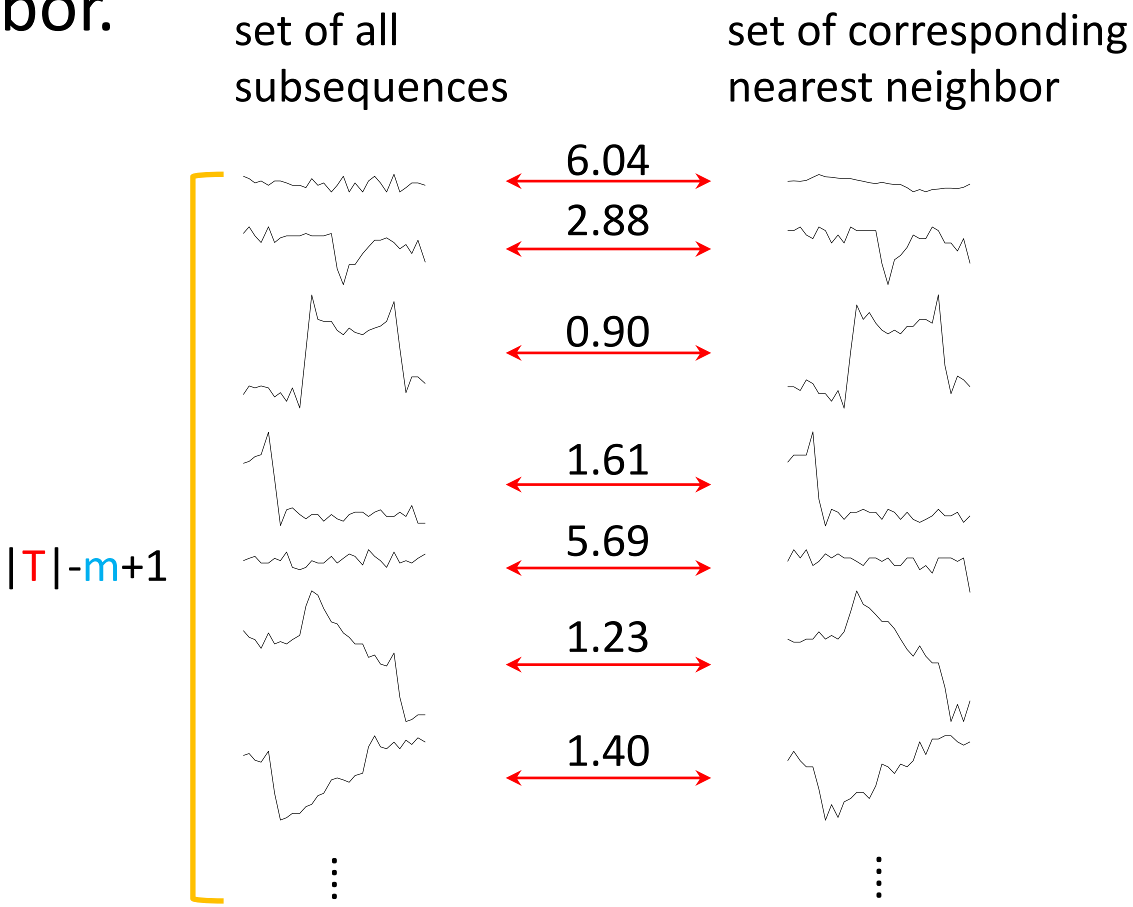

# Matrix Profile

- The distance to the corresponding nearest neighbor of each subsequence can be stored in a vector called **matrix profile P**.

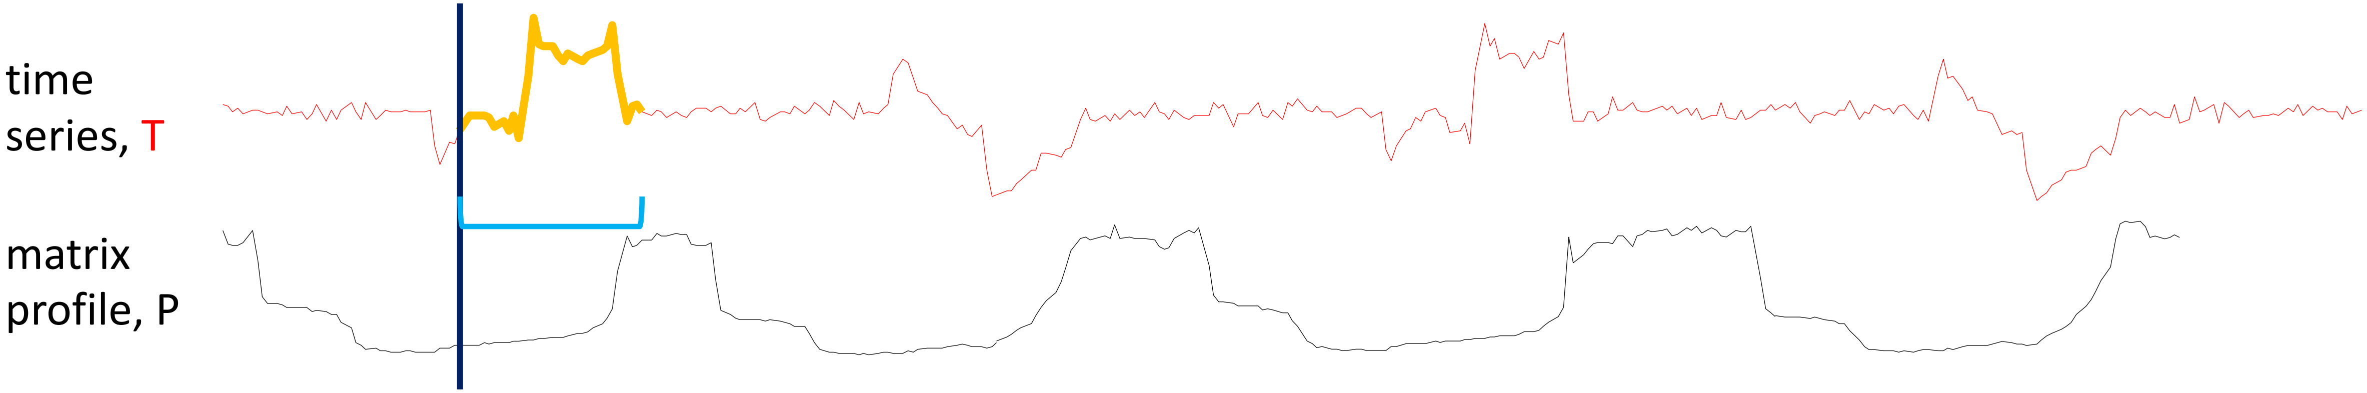

The matrix profile value at location *i* is the distance between $T_i$ and its nearest neighbor

# Matrix Profile

- The index of corresponding nearest neighbor of each subsequence is also stored in a vector called matrix profile index.

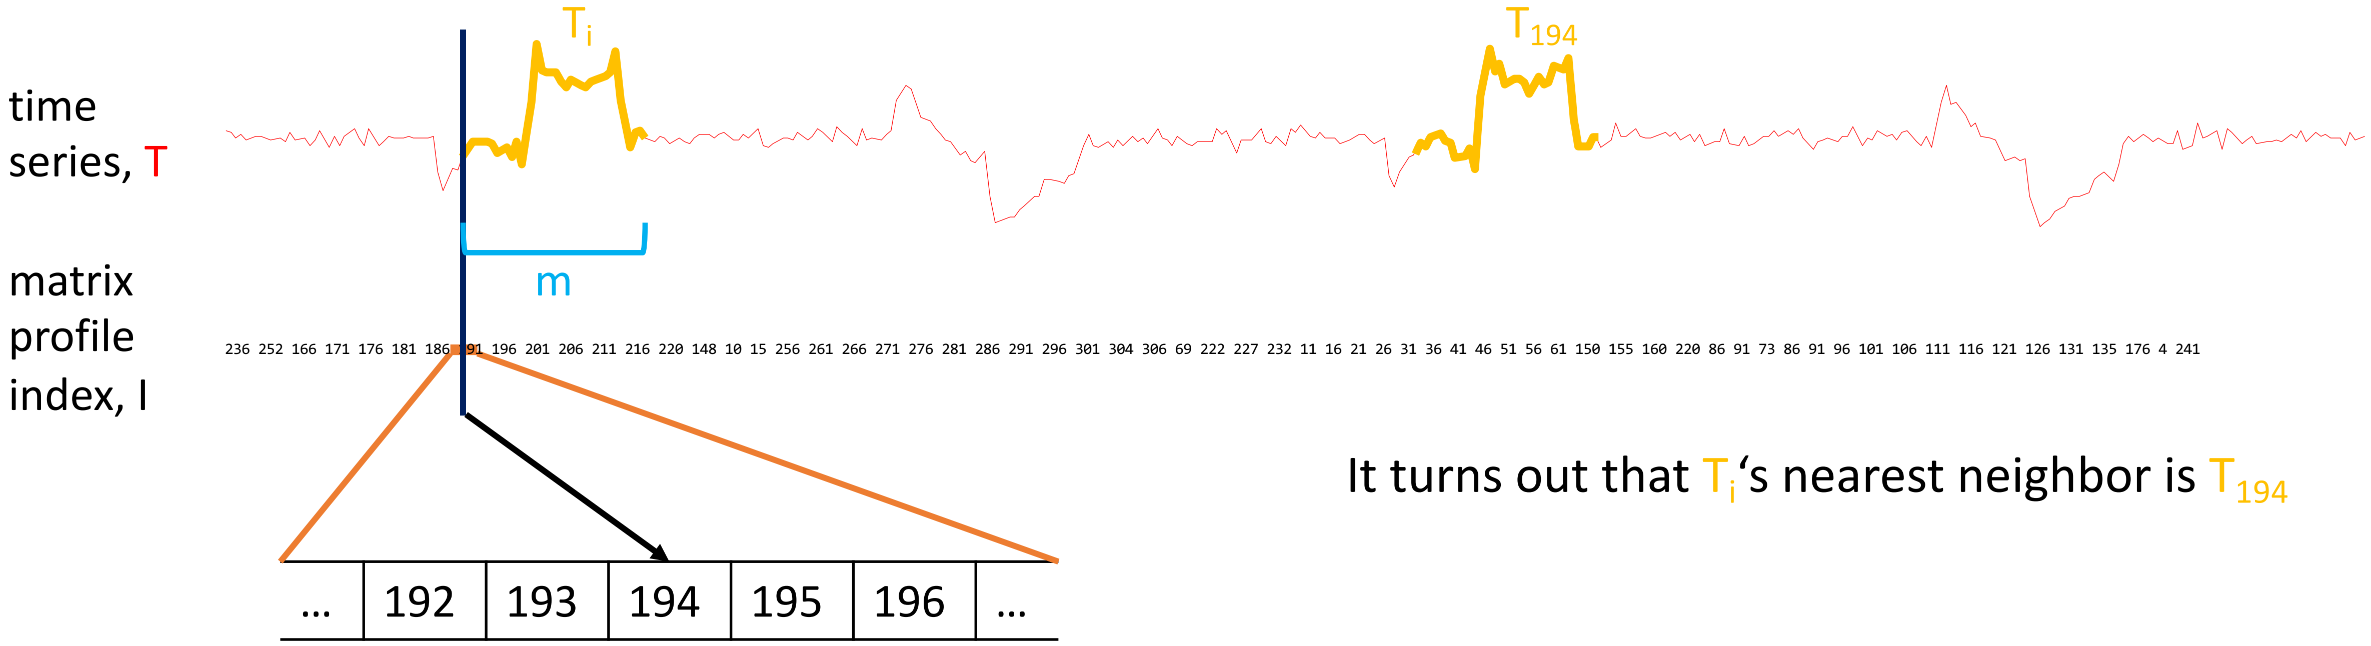

It turns out that $T_i$‘s nearest neighbor is $T_{194}$

The matrix profile value at location *i* is the distance between $T_i$ and its nearest neighbor

# Matrix Profile

- The MP index allows to find the nearest neighbor to any subsequence in constant time.
- Note that the pointers in the matrix profile index are not necessarily symmetric.
- If A points to B, then B may or may not point to A
- The classic TS motif: the two smallest values in the MP must have the same value, and their pointers must be mutual.

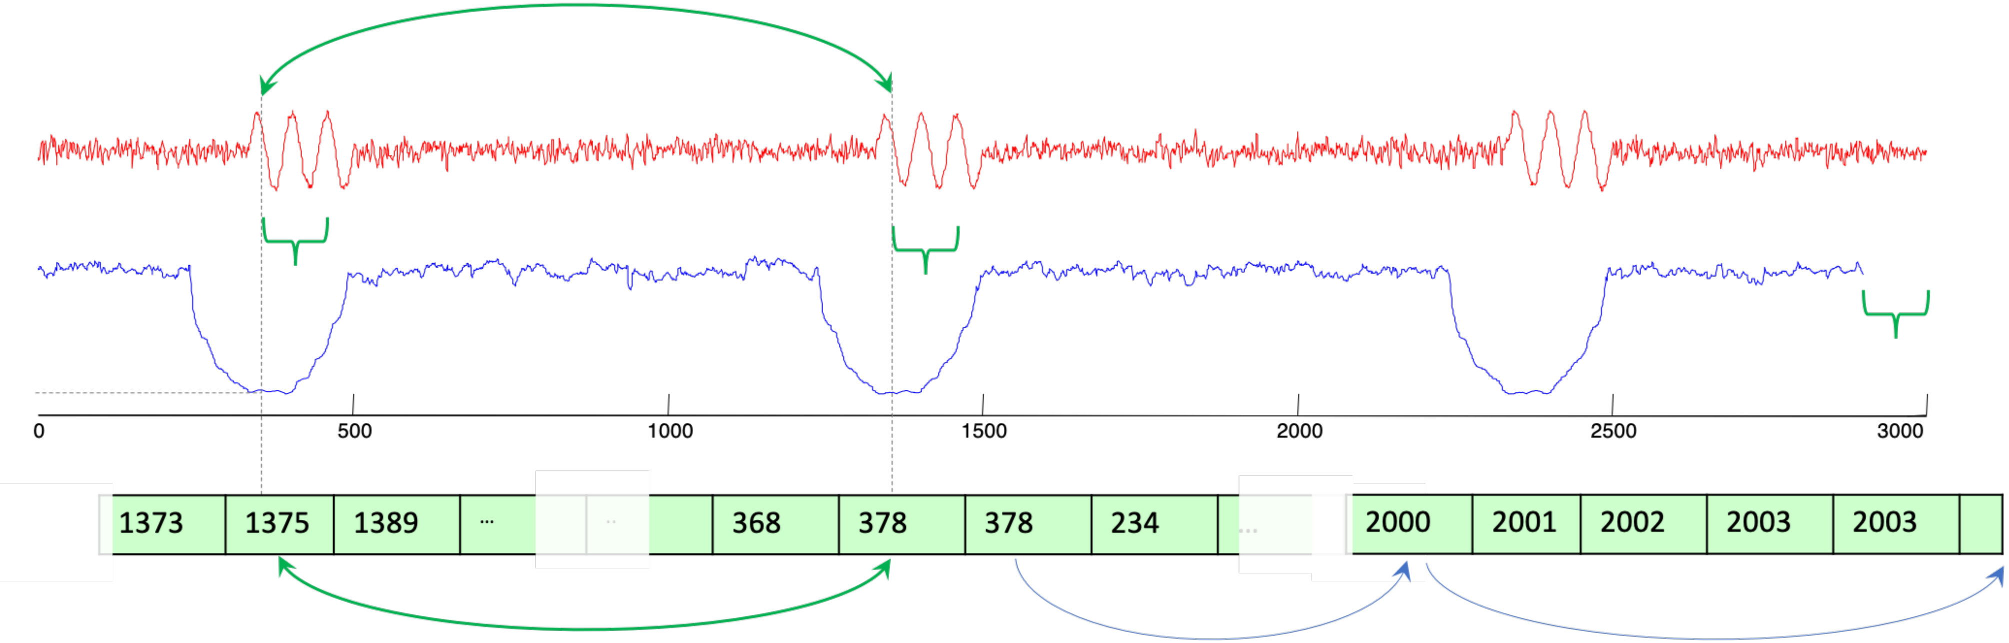

| | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1373 | 1375 | 1389 | ··· | ·· | 368 | 378 | 378 | 234 | ... | 2000 | 2001 | 2002 | 2003 | 2003 |

# How to “read” a Matrix Profile

- For relatively low values, you know that the subsequence in the original TS must have (at least one) relatively similar subsequence elsewhere in the data (such regions are “motifs”)
- For relatively high values, you know that the subsequence in the original TS must be unique in its shape (such areas are anomalies).

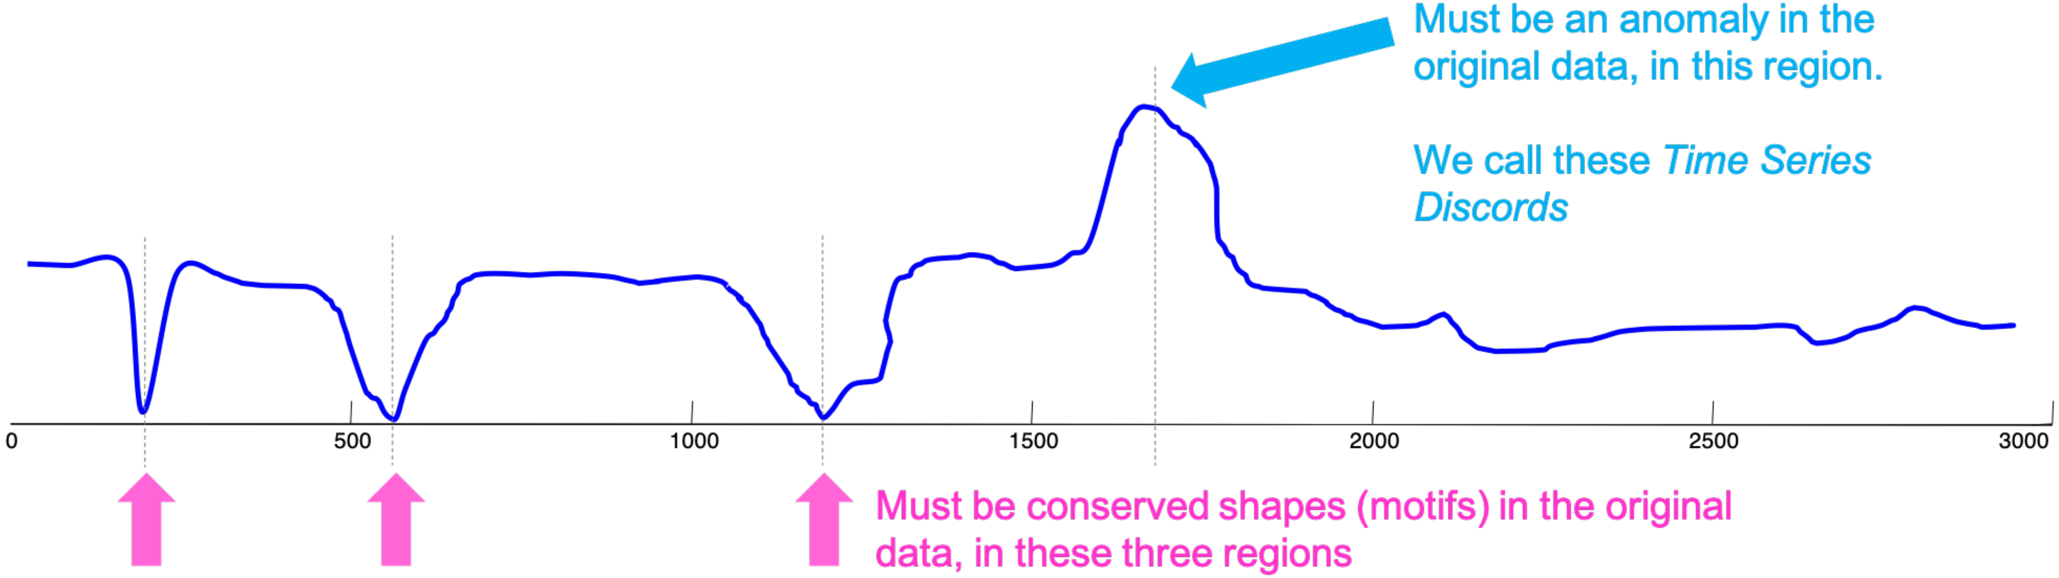

Must be an anomaly in the original data, in this region.

We call these *Time Series Discords*

Must be conserved shapes (motifs) in the original data, in these three regions

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

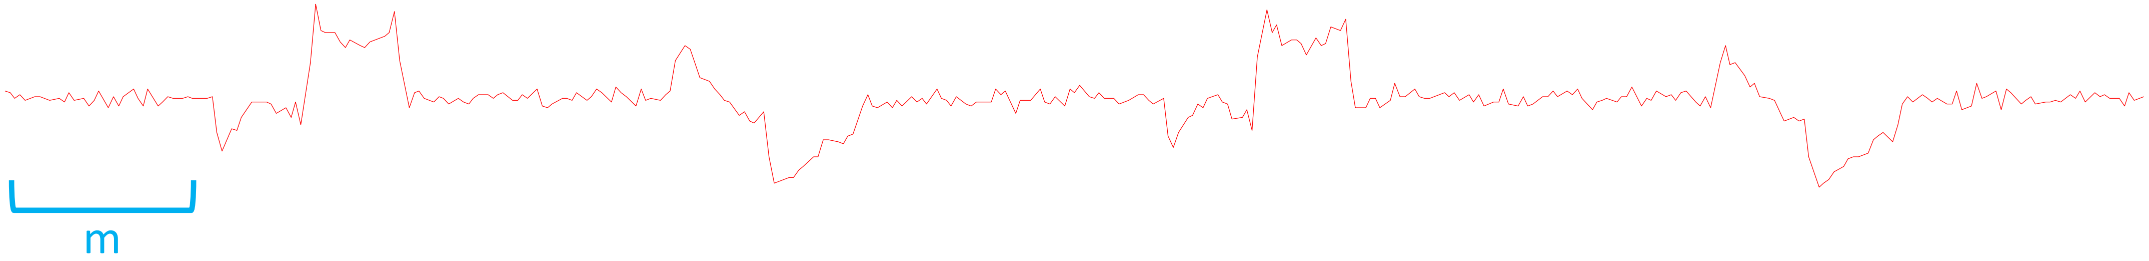

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf |

Matrix profile is initialized as inf vector

This is just a toy example, so the values and the vector length does not fit the time series shown above

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

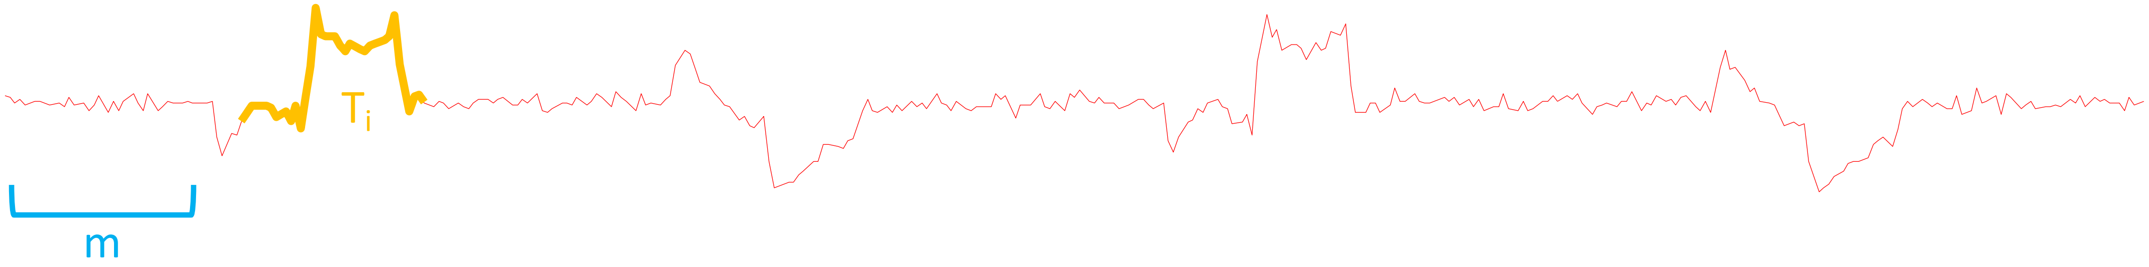

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf |

At the first iteration, a subsequence $T_i$ is randomly selected from T

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

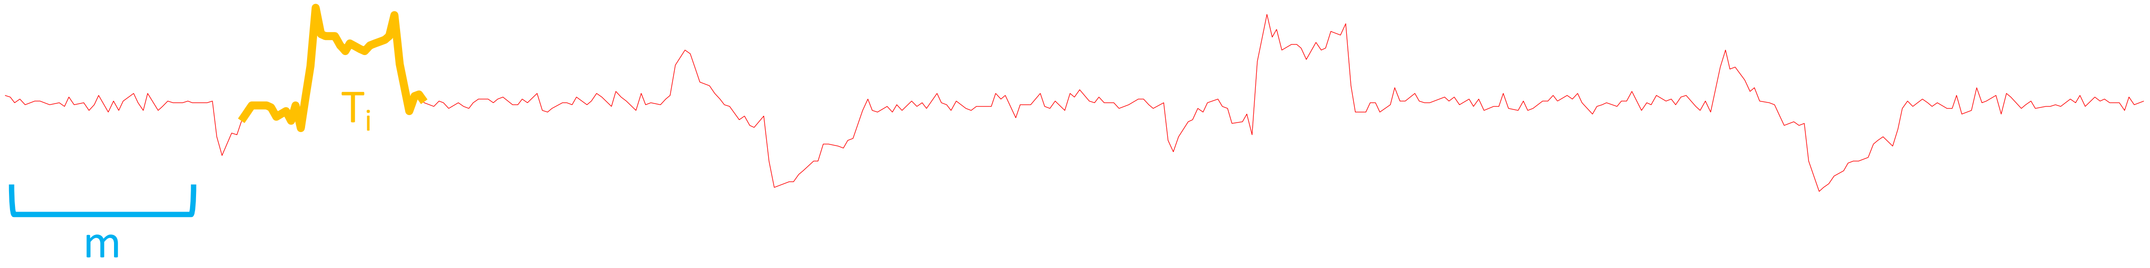

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf |

We compute the distances between $T_i$ and every subsequences from T (time complexity = O(|T|log(|T|)))
We then put the distances in a vector based on the position of the subsequences

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | 2 | 0 | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

The distance between $T_i$ and $T_1$ (first subsequence) is 3

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

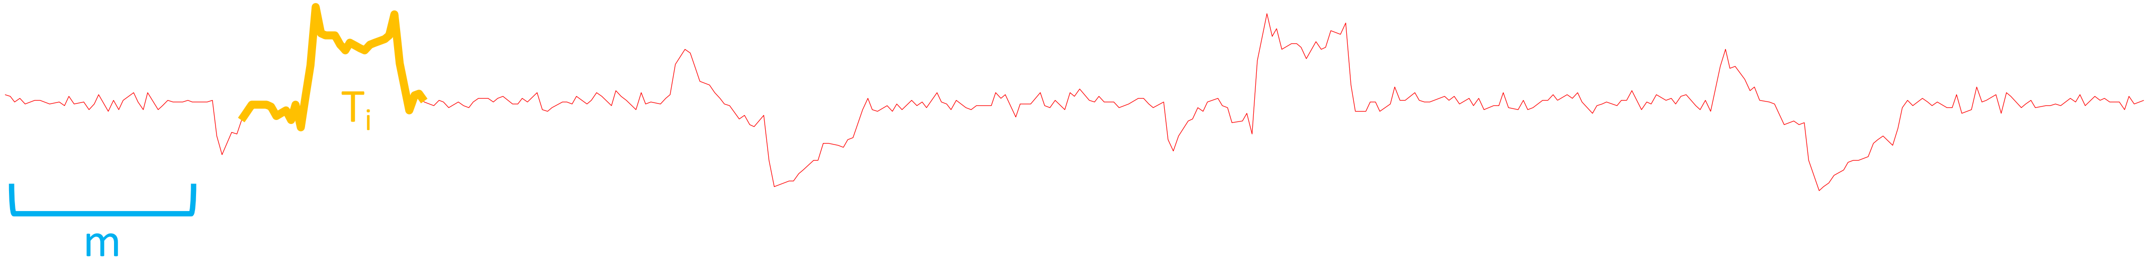

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf |

We compute the distances between $T_i$ and every subsequences from T (time complexity = O(|T|log(|T|)))
We them put the distances in a vector based on the position of the subsequences

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | 2 | 0 | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

Let say $T_i$ happen to be the third subsequences, therefore
the third value in the distance vector is 0

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

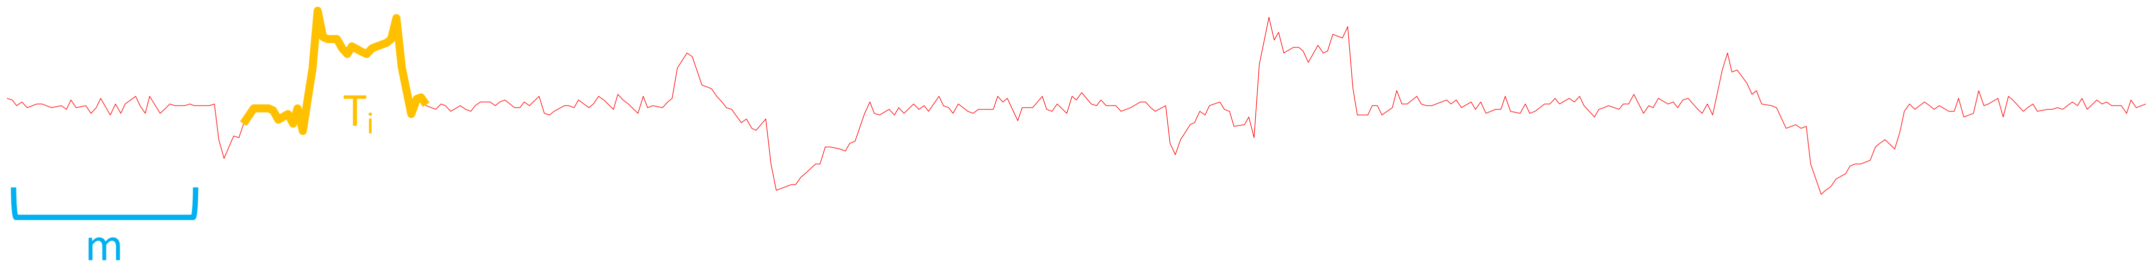

|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf |

min ↕

Matrix profile is updated by apply elementwise minimum to these two vectors

|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | 2 | 0 | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

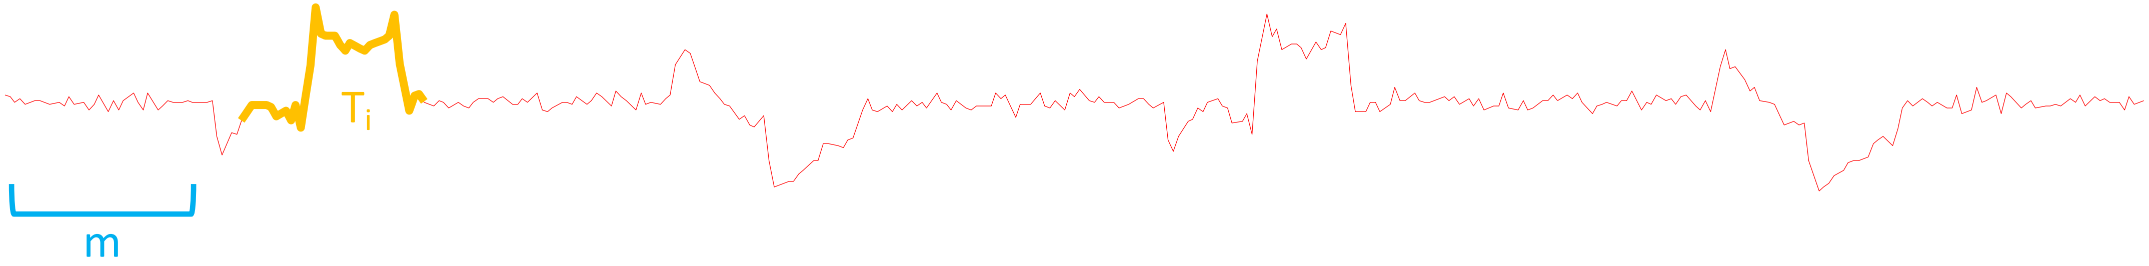

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf | inf |

min

Matrix profile is updated by apply elementwise minimum to these two vectors

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | 2 | 0 | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

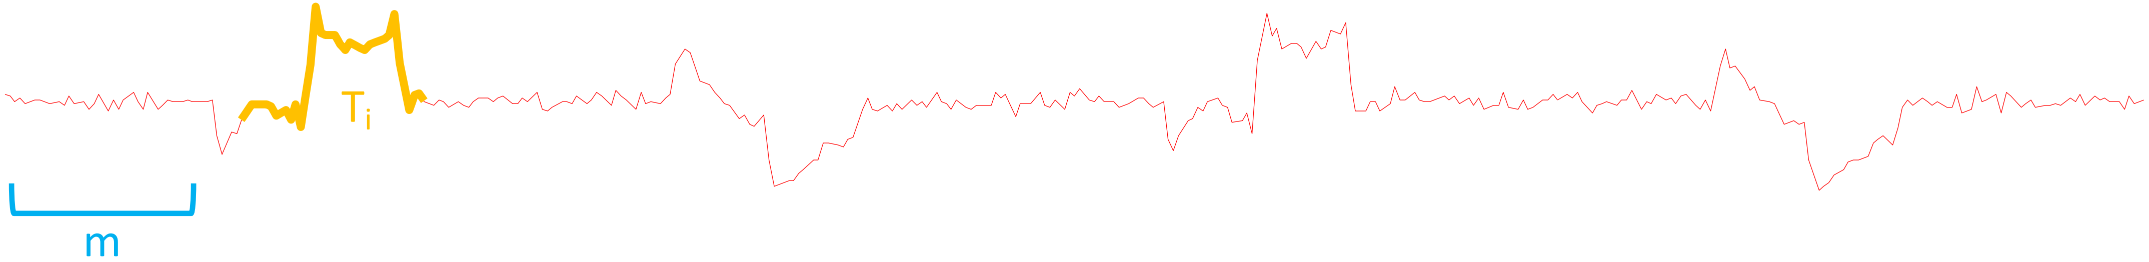

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | 2 | inf | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

After we finish update matrix profile for the first iteration

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | 2 | 0 | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

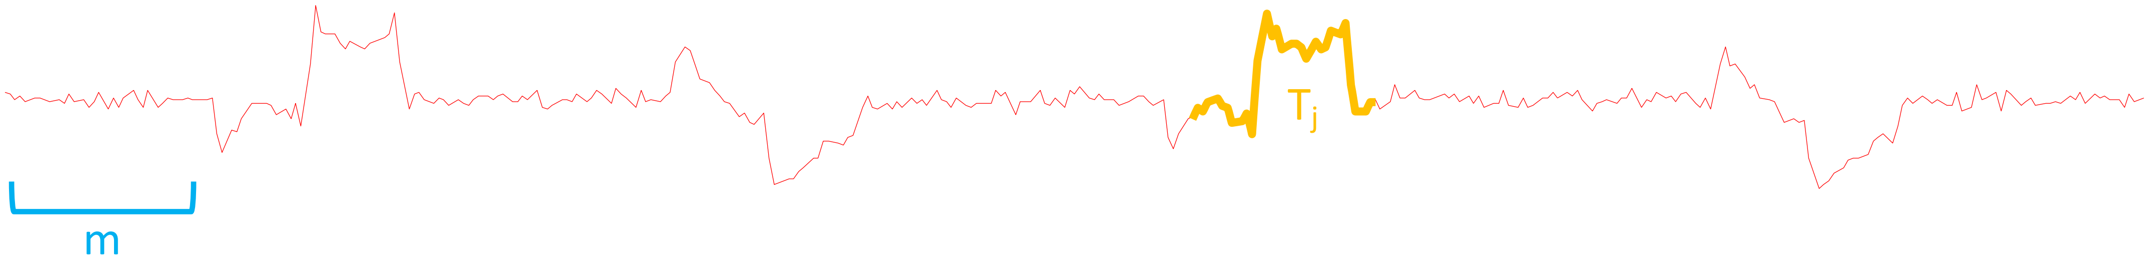

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | 2 | inf | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

In the second iteration, we randomly select another subsequence $T_j$ and it happens to be the 12th subsequences

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

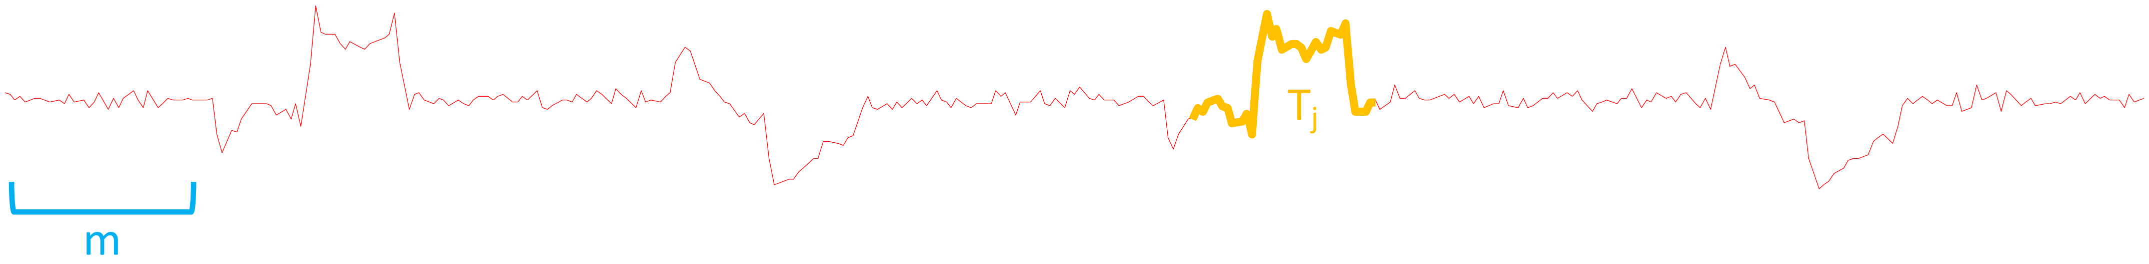

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | 2 | inf | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

Once again, we compute the distance between $T_j$ and every subsequences of T

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 3 | 1 | 4 | 4 | 3 | 6 | 2 | 1 | 5 | 8 | 0 | 2 | 3 | 5 | 9 | 4 | 2 | 2 |

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

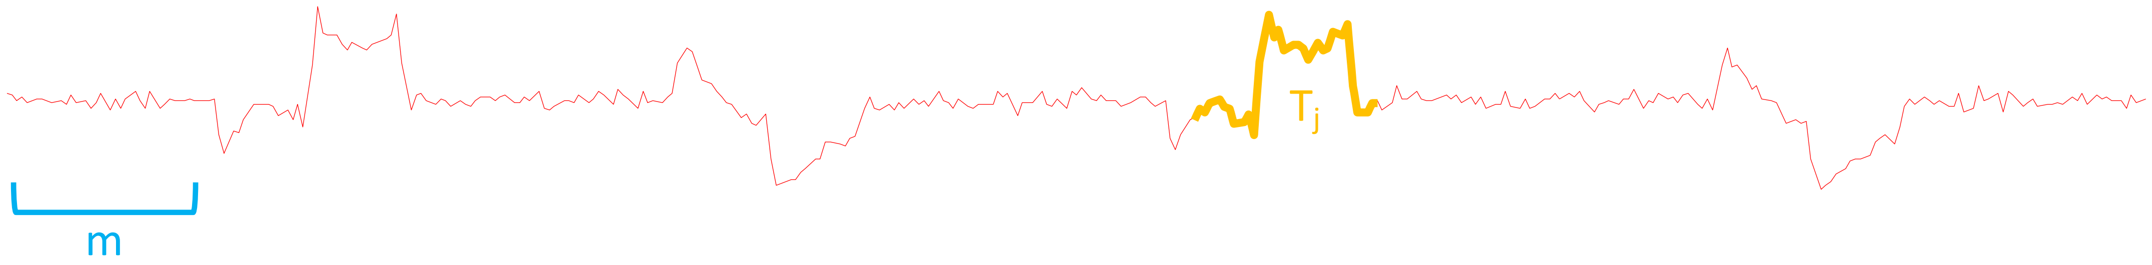

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 3 | 2 | inf | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

min ↕

The same elementwise minimum

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 3 | 1 | 4 | 4 | 3 | 6 | 2 | 1 | 5 | 8 | 0 | 2 | 3 | 5 | 9 | 4 | 2 | 2 |

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

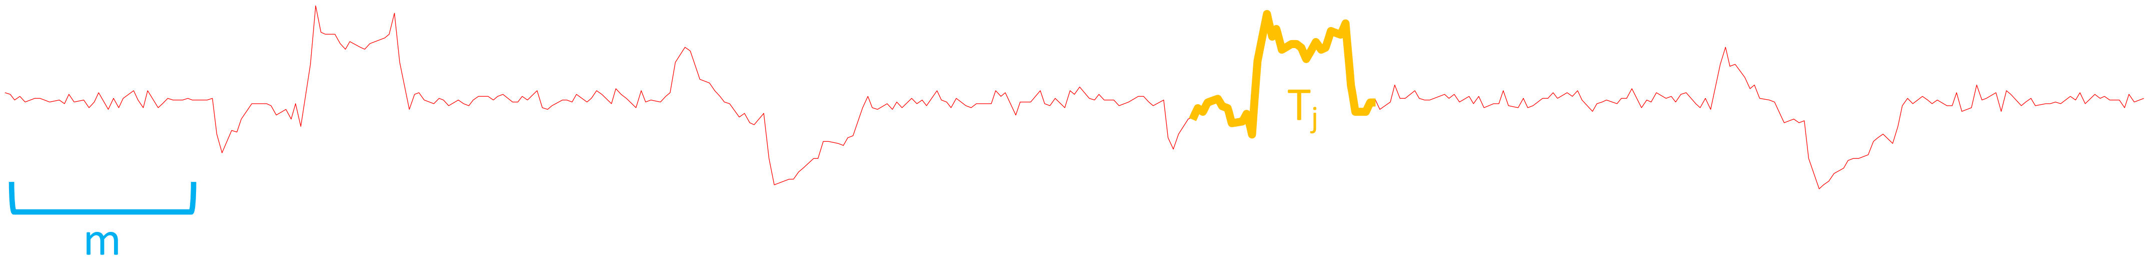

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 2 | inf | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

min ↕

The same elementwise minimum

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 3 | 1 | 4 | 4 | 3 | 6 | 2 | 1 | 5 | 8 | 0 | 2 | 3 | 5 | 9 | 4 | 2 | 2 |

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

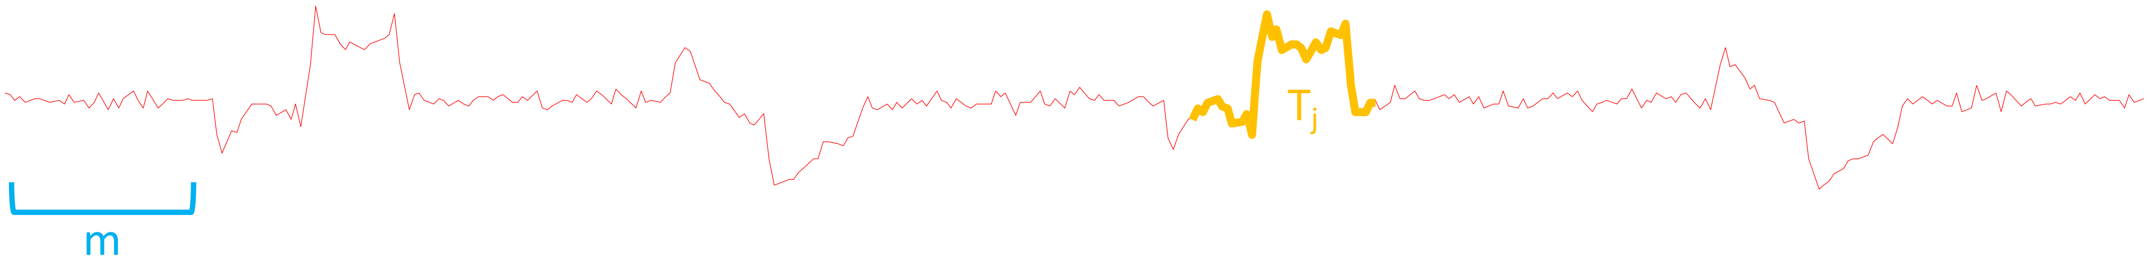

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 2 | inf | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

min ↕

The same elementwise minimum

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 3 | 1 | 4 | 4 | 3 | 6 | 2 | 1 | 5 | 8 | 0 | 2 | 3 | 5 | 9 | 4 | 2 | 2 |

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

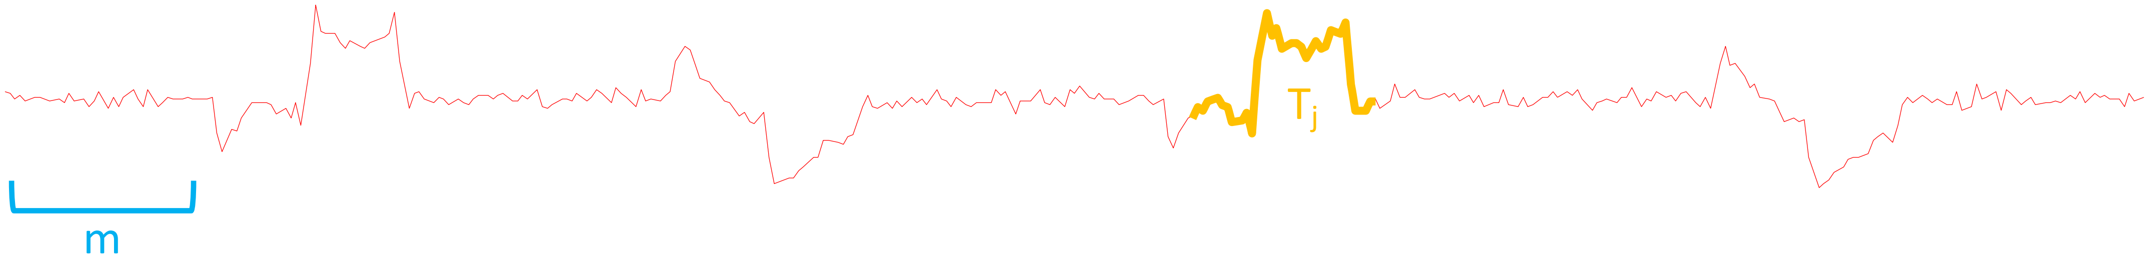

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 2 | 1 | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

min

The same elementwise minimum

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 3 | 1 | 4 | 4 | 3 | 6 | 2 | 1 | 5 | 8 | 0 | 2 | 3 | 5 | 9 | 4 | 2 | 2 |

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

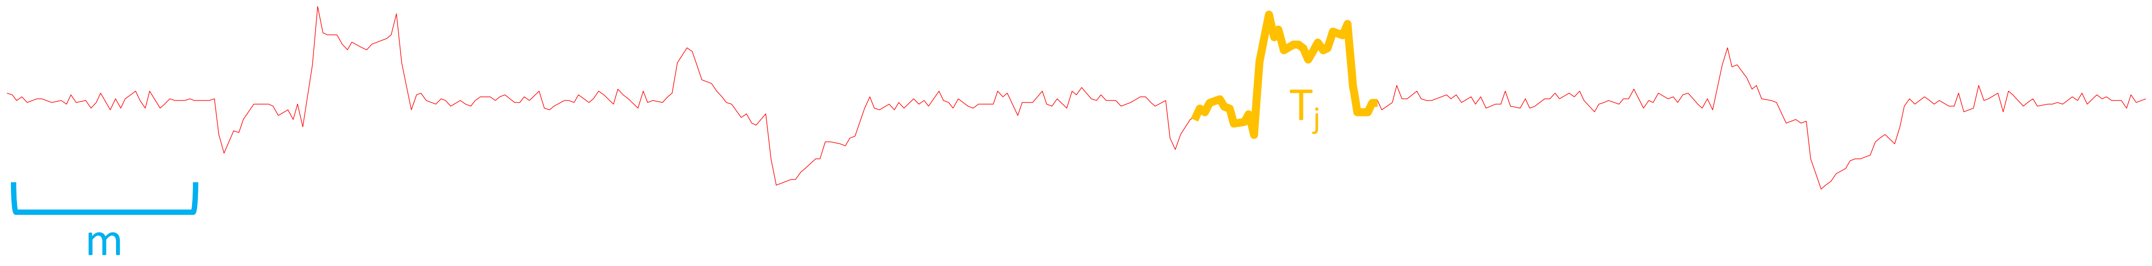

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 2 | 1 | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

min ↕

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 3 | 1 | 4 | 4 | 3 | 6 | 2 | 1 | 5 | 8 | 0 | 2 | 3 | 5 | 9 | 4 | 2 | 2 |

We repeat the two steps (distance computation and update) until we have used every subsequences

# How to Compute Matrix Profile?

- Given a time series, T and a desired subsequence length, m.

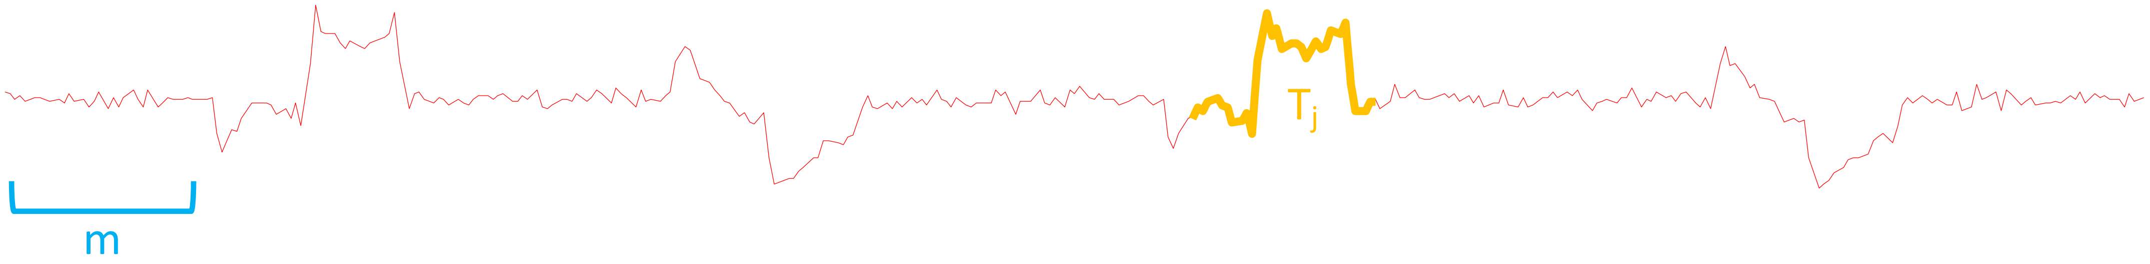

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 2 | 1 | 5 | 3 | 4 | 5 | 1 | 2 | 9 | 8 | 4 | 2 | 3 | 4 | 8 | 6 | 2 | 1 |

min

| | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 3 | 1 | 4 | 4 | 3 | 6 | 2 | 1 | 5 | 8 | 0 | 2 | 3 | 5 | 9 | 4 | 2 | 2 |

There are |T| subsequences and the distance computation is $O(|T|\log(|T|))$

The overall time complexity is $O(|T|^2\log(|T|))$

# Motif Discovery From Matrix Profile

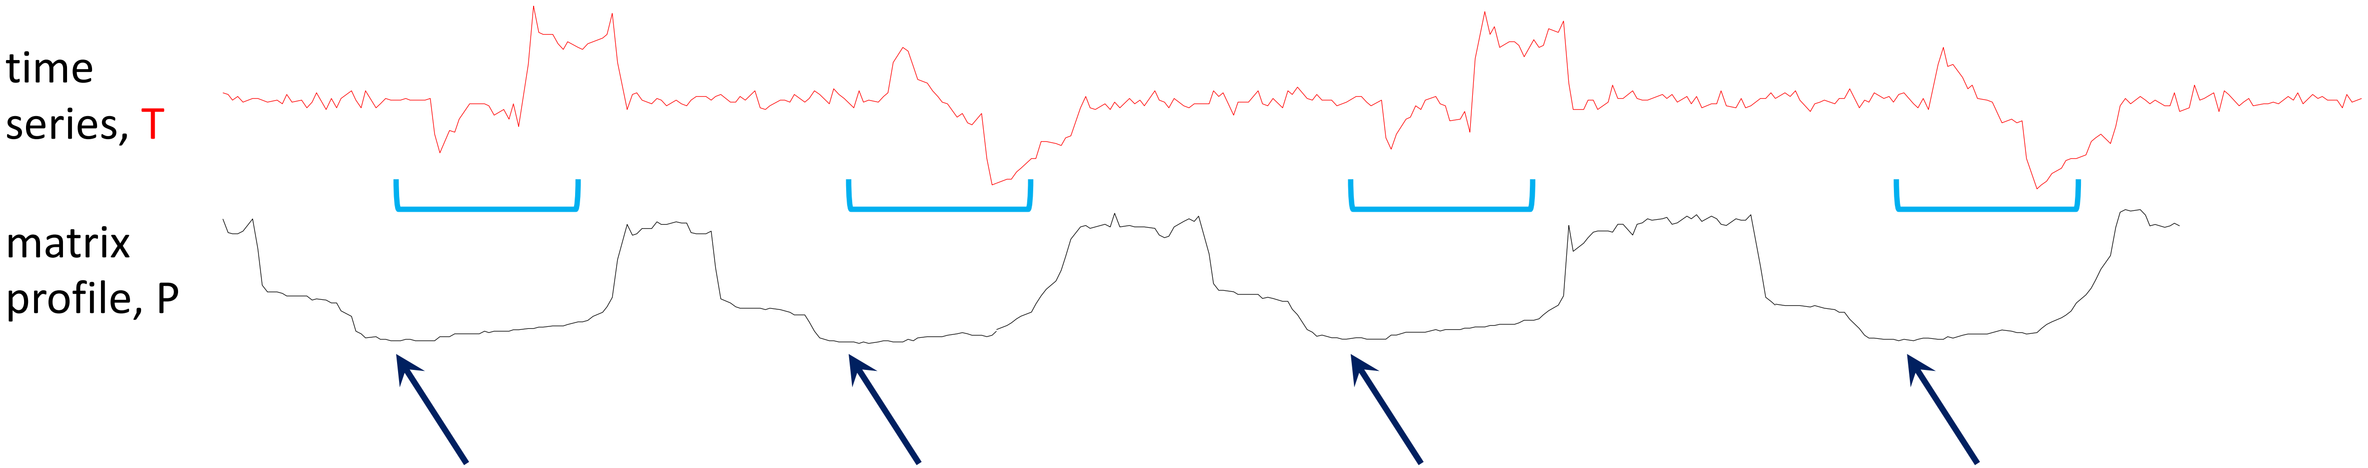

Local minimums are corresponding to motifs

# Motif Discovery From Matrix Profile

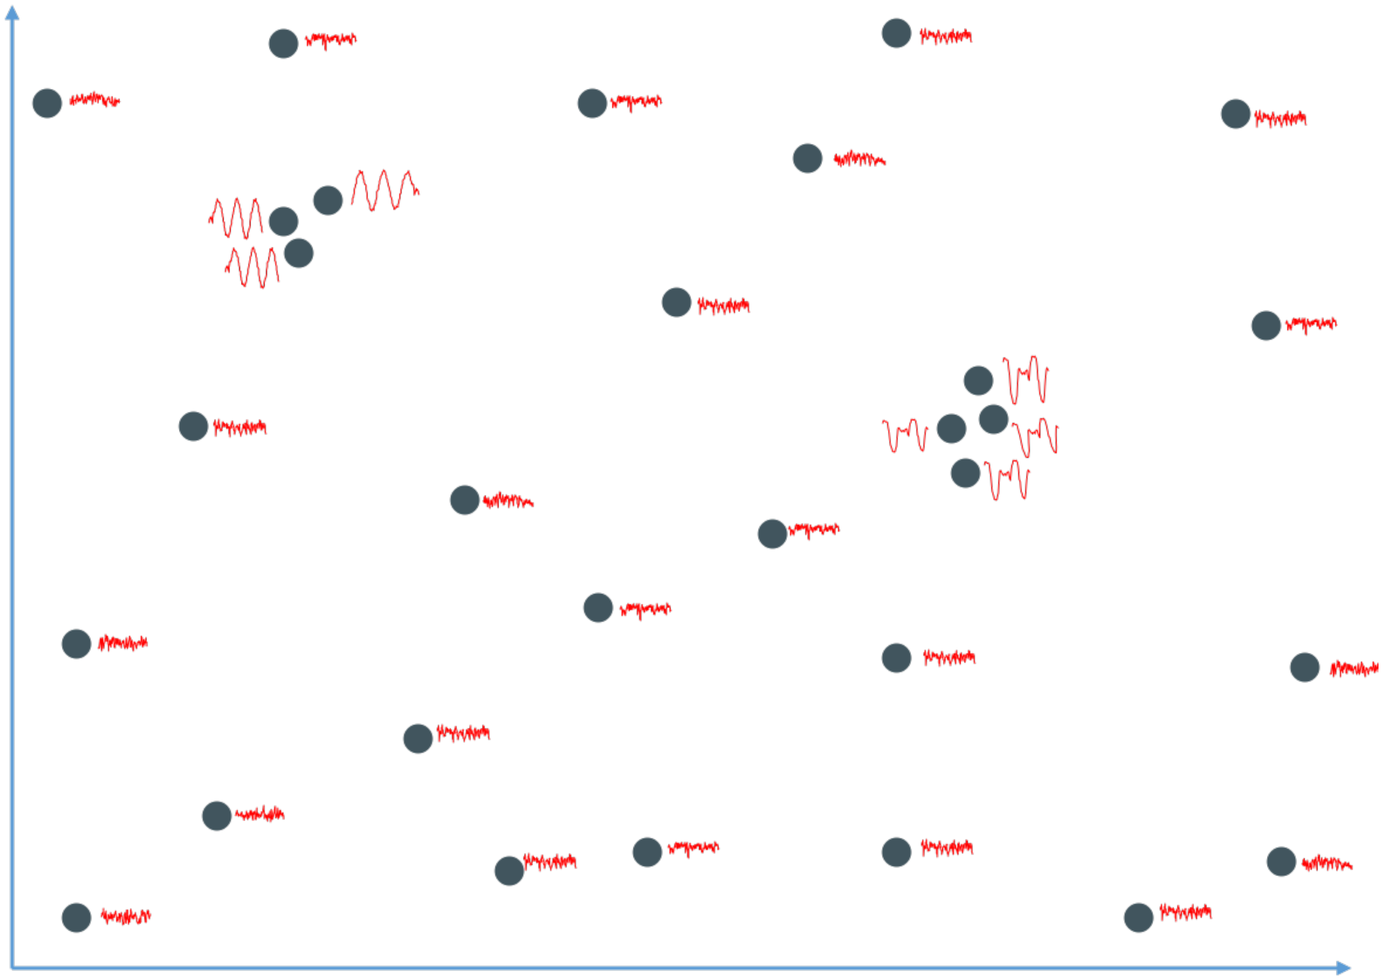

- It is sometime useful to think of time series subsequences as points in m-dimensional space.
- In this view, dense regions in the m-dimensional space correspond to regions of the time series that have a low corresponding MP.

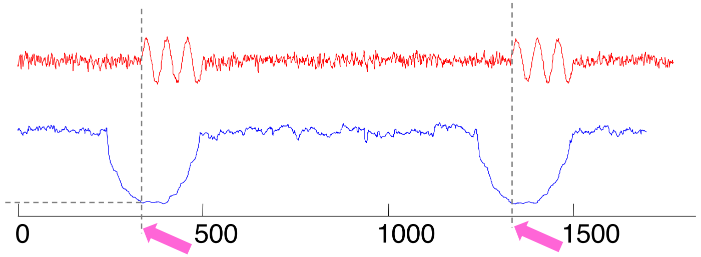

# Top-K Motifs

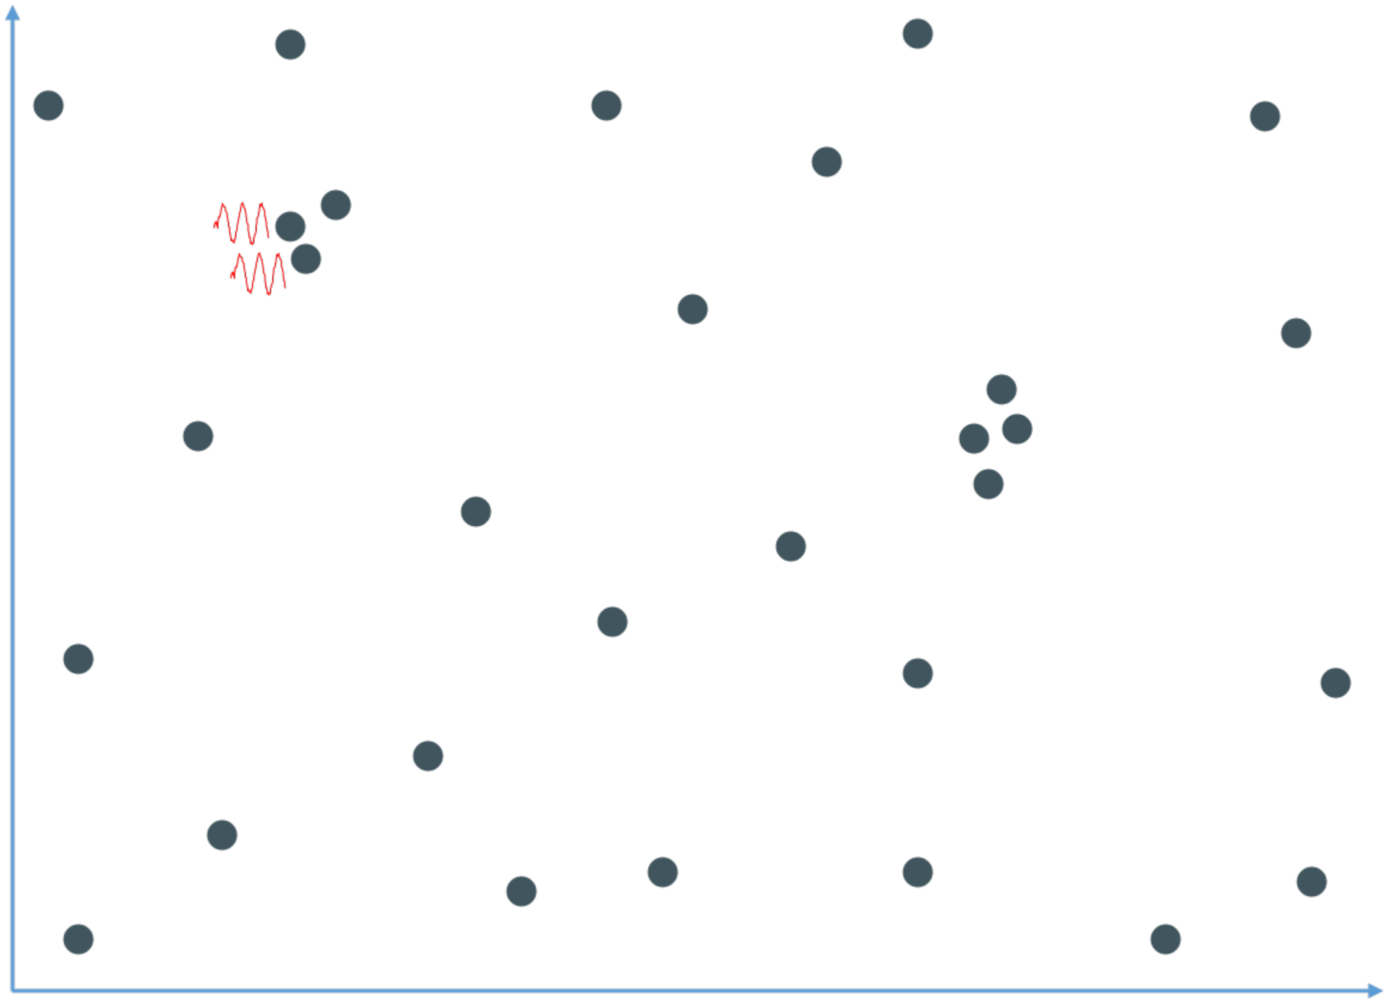

- We need a parameter R.
- 1 < R < (small number, say 3)
- Lets make R = 2 for now.
- We begin by finding the nearest pair of points, the *motif pair*….
- This the pair of subsequences corresponding to lowest pair of values in the MP

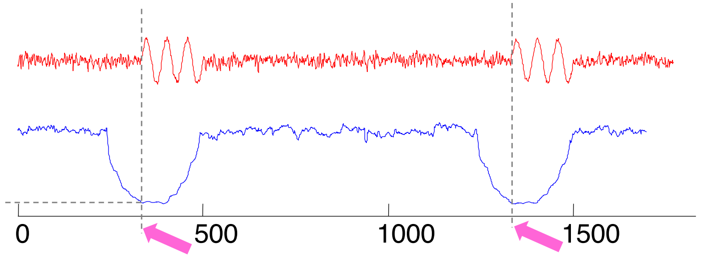

# Top-K Motifs

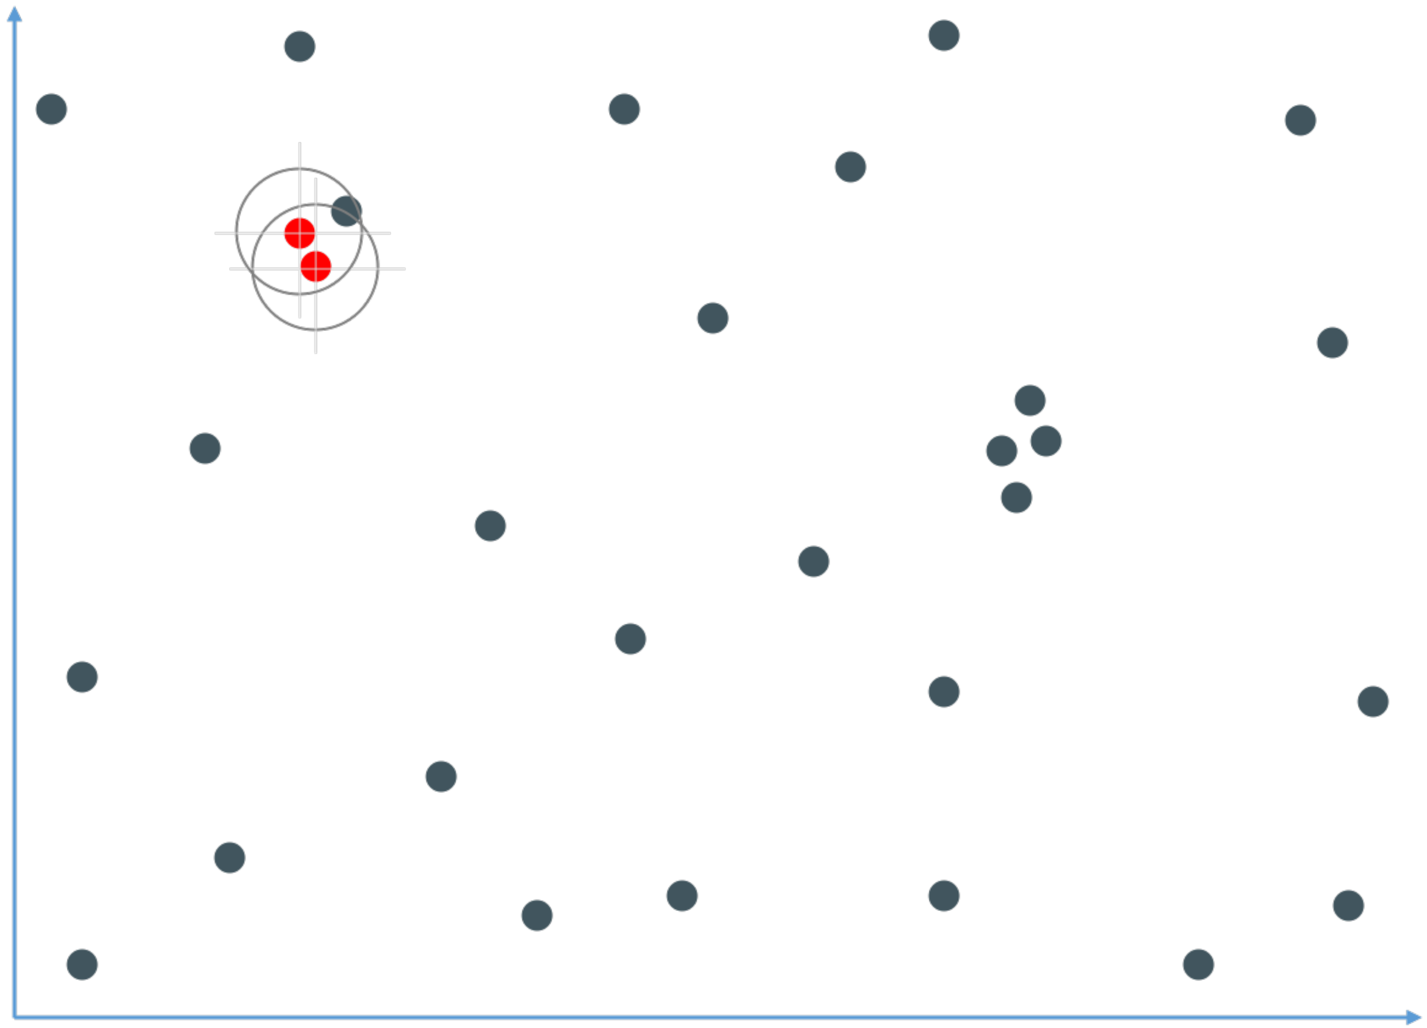

- We find the nearest pair of points are D1 apart.
- Lets draw a circle, D1 times R, around both points.
- Any points that are within either of these circles, are added to this motif, in this case just one.
- The Top-1 motif has three members, it is done.

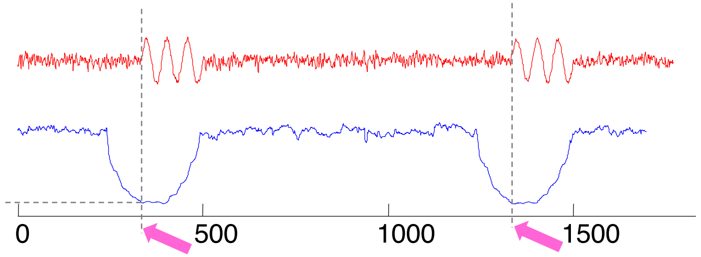

# Top-K Motifs

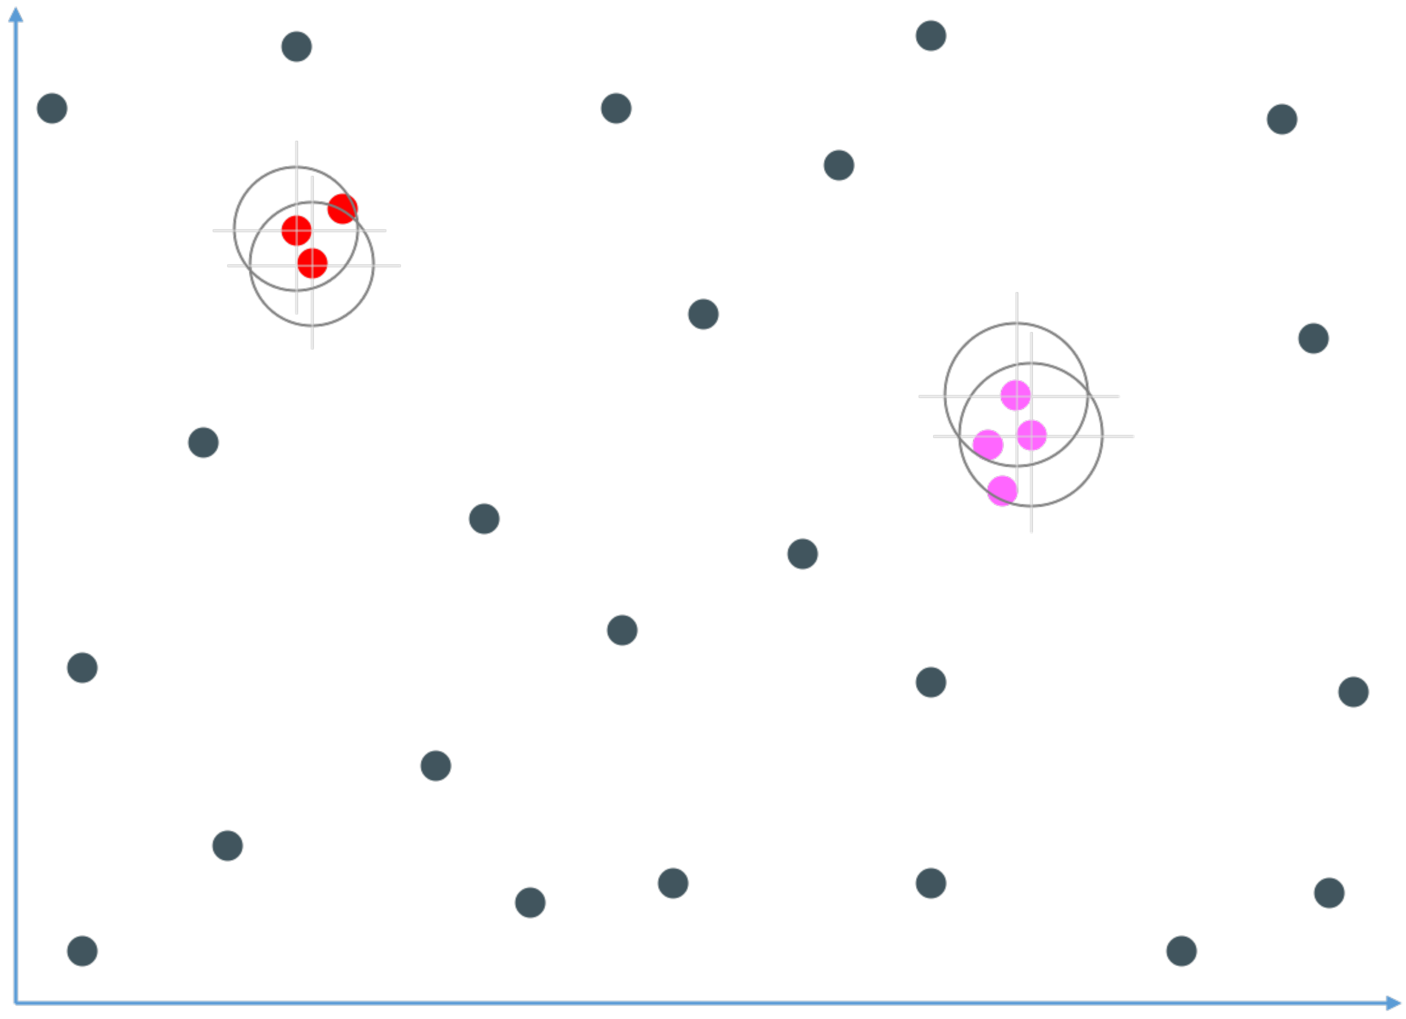

- Now lets find the Top-2 motif. We find the ***nearest pair of points***, excluding anything from the top motif.
- The nearest pair of points are D2 apart.
- Lets draw a circle D1 times R, around both points.
- Any points that are within either of these circles, is added to this motif, in this case there are two for a total of four items in the Top-2 Motif

# Top-K Motifs

- We are done with the Top-2 Motif
- Note that we will always have:
  - $D_1 < D_2 < D_3 \dots D_K$
- **When to stop?** (what is K?)
- We could use MDL or a predefined K.

# Anomaly Discovery From Matrix Profile

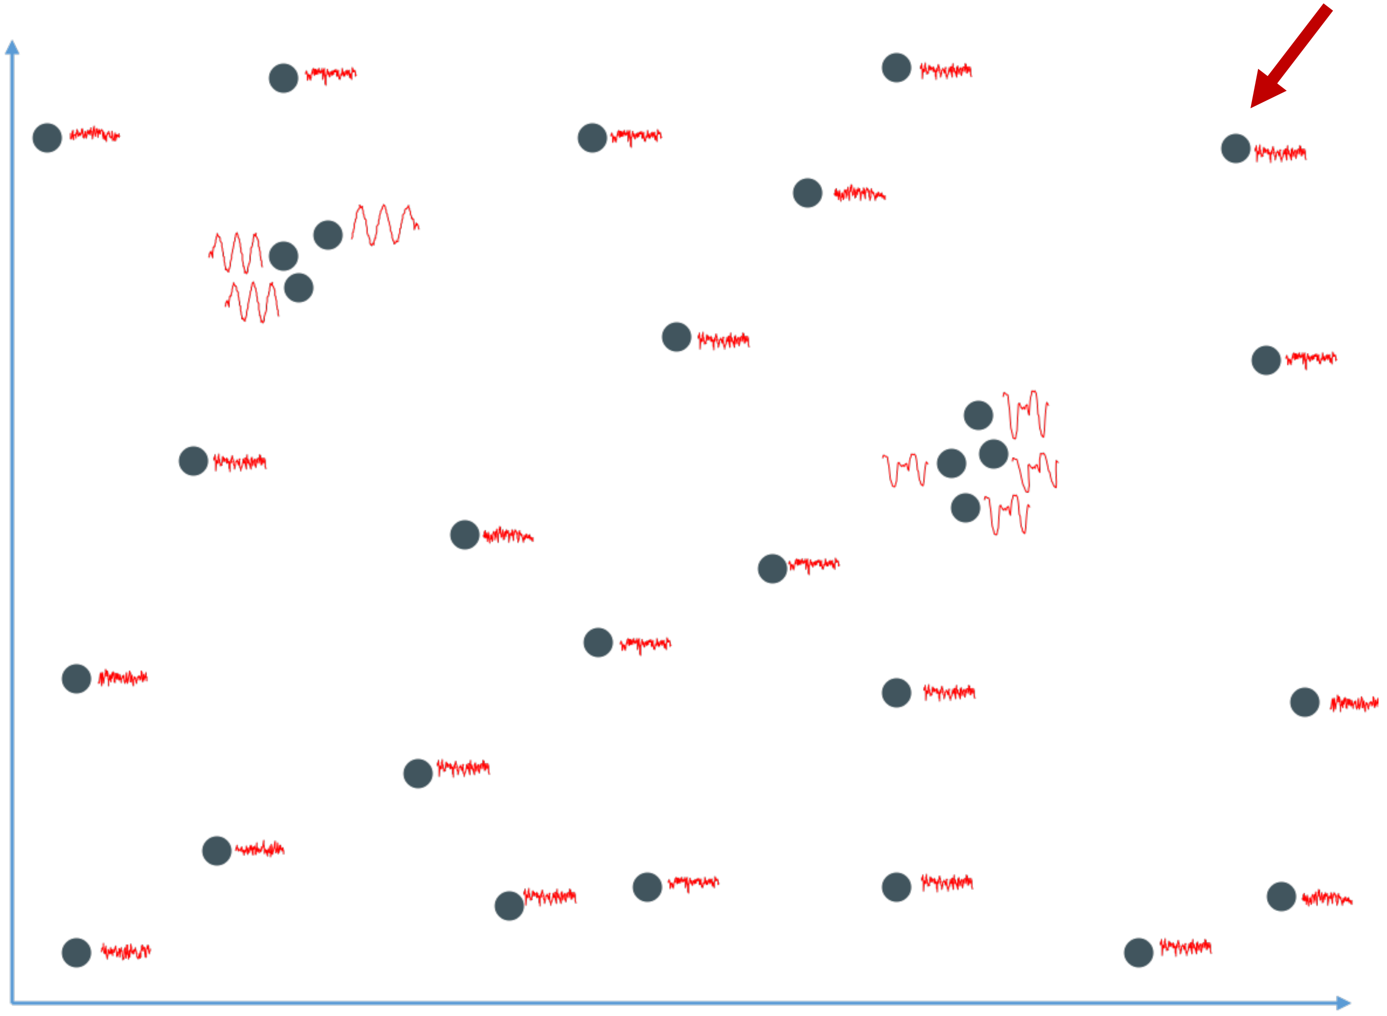

- We need a parameter E of subseqeunces to exclude in the vicinity of the anomaly.
- Lets make E = 2 for now.
- We begin by finding the subsequence with the highest distance in the MP
- This corresponding to biggest anomaly

# Top-K Anomaly

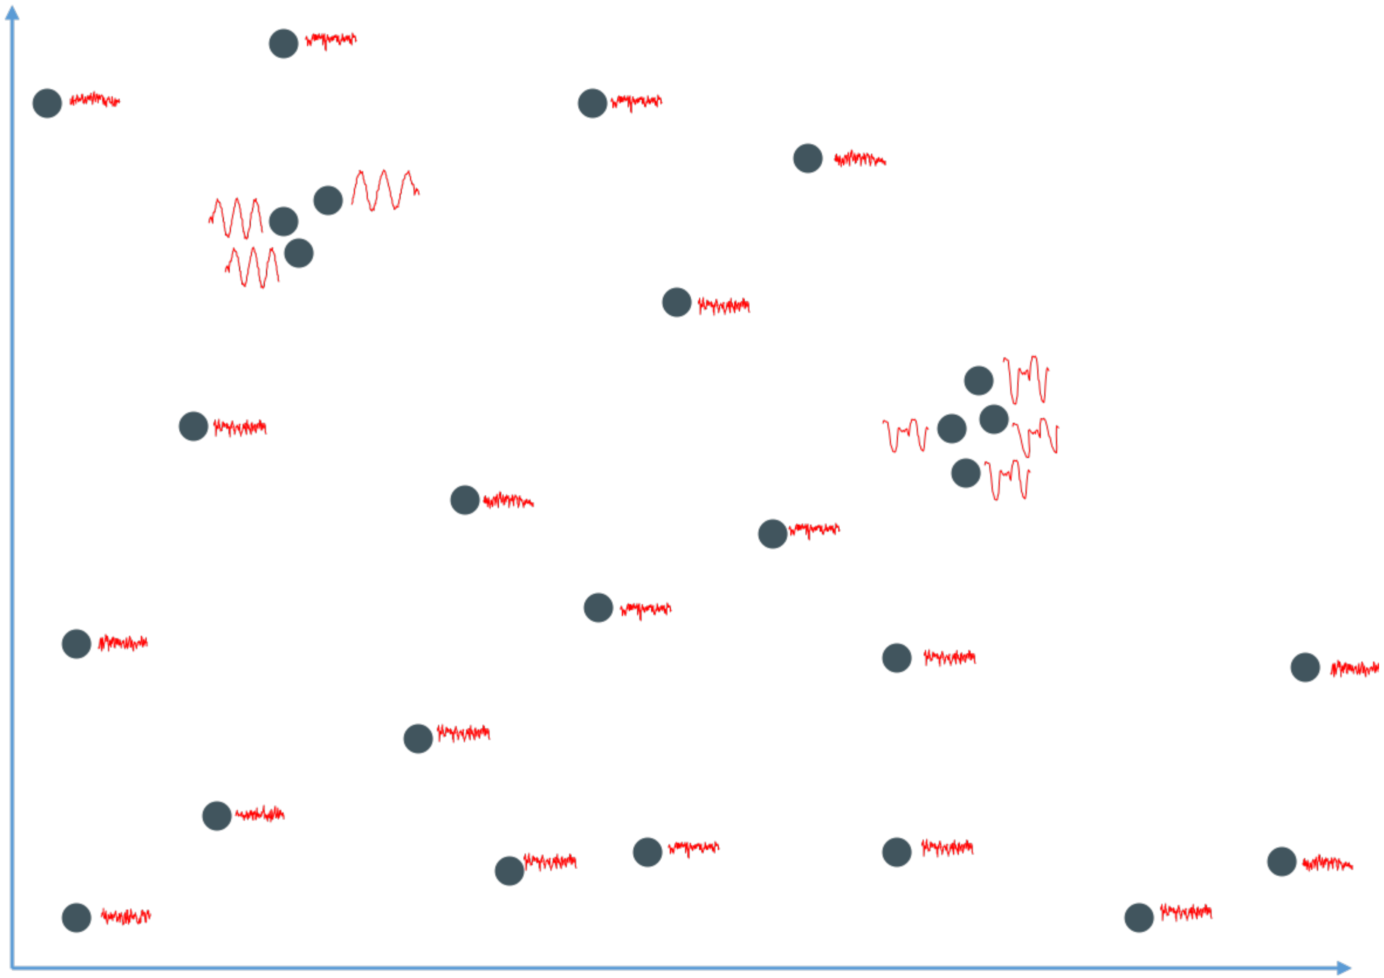

- Then we look for the E closest subsequences to the anomaly.
- We remove all of them.
- We can use a predefined K or the MDL to stop.

# Shapelet

# Time Series Classification

- Given a set $X$ of $n$ time series, $X = \{x_1, x_2, \ldots, x_n\}$, each time series has m ordered values $x_i = < x_{t1}, x_{t2}, \ldots, x_{tm} >$ and a class value $c_i$.
- The objective is to find a function $f$ that maps from the space of possible time series to the space of possible class values.
- Generally, it is assumed that all the TS have the same length $m$.

# Shapelet-based Classification

1. Represent a TS as a vector of distances with representative subsequences, namely  shapelets.
2. Use it to as input for machine learning classifiers.

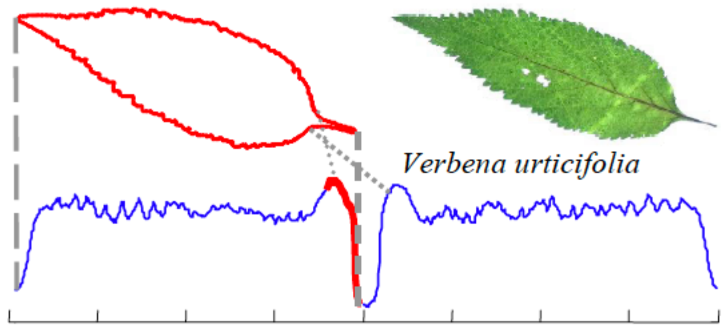

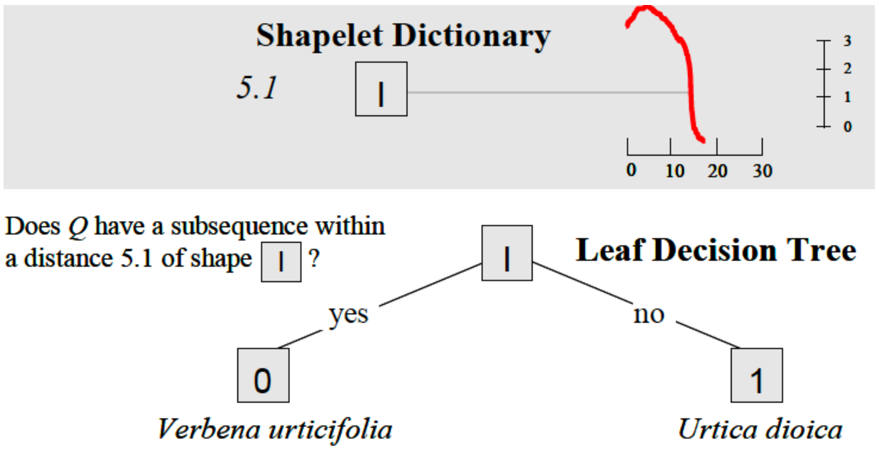

Does $Q$ have a subsequence within a distance 5.1 of shape I ?

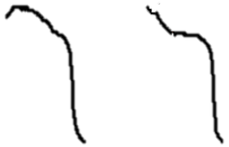

| | |
|---|---|
| 3.2 | 8.7 |
| 1.4 | 7.9 |
| 6.7 | 4.2 |
| 9.2 | 3.4 |

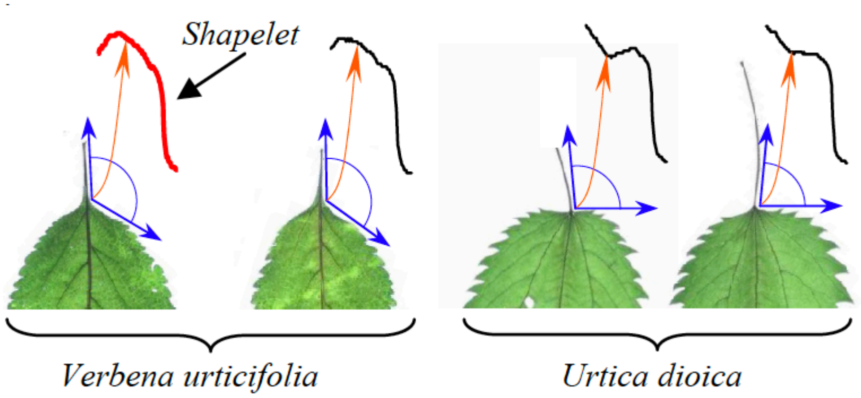

# Time Series Shapelets

- Shapelets are TS subsequences which are maximally representative of a class.
- Shapelets can provide interpretable results, which may help domain practitioners better understand their data.
- Shapelets can be significantly more accurate/robust because they are *local features*, whereas most other state-of-the-art TS classifiers consider *global features*.

| | |
|---|---|
| Verbena | 0.87 |
| Urtica | 0.34 |

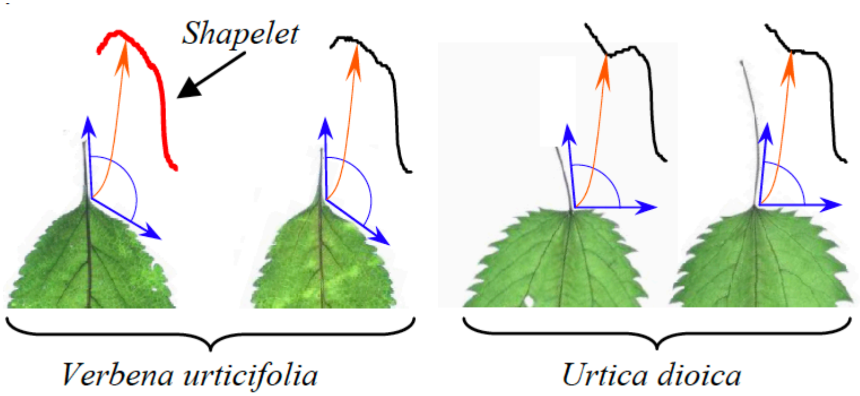

# Extract Subsequences of all Possible Lengths

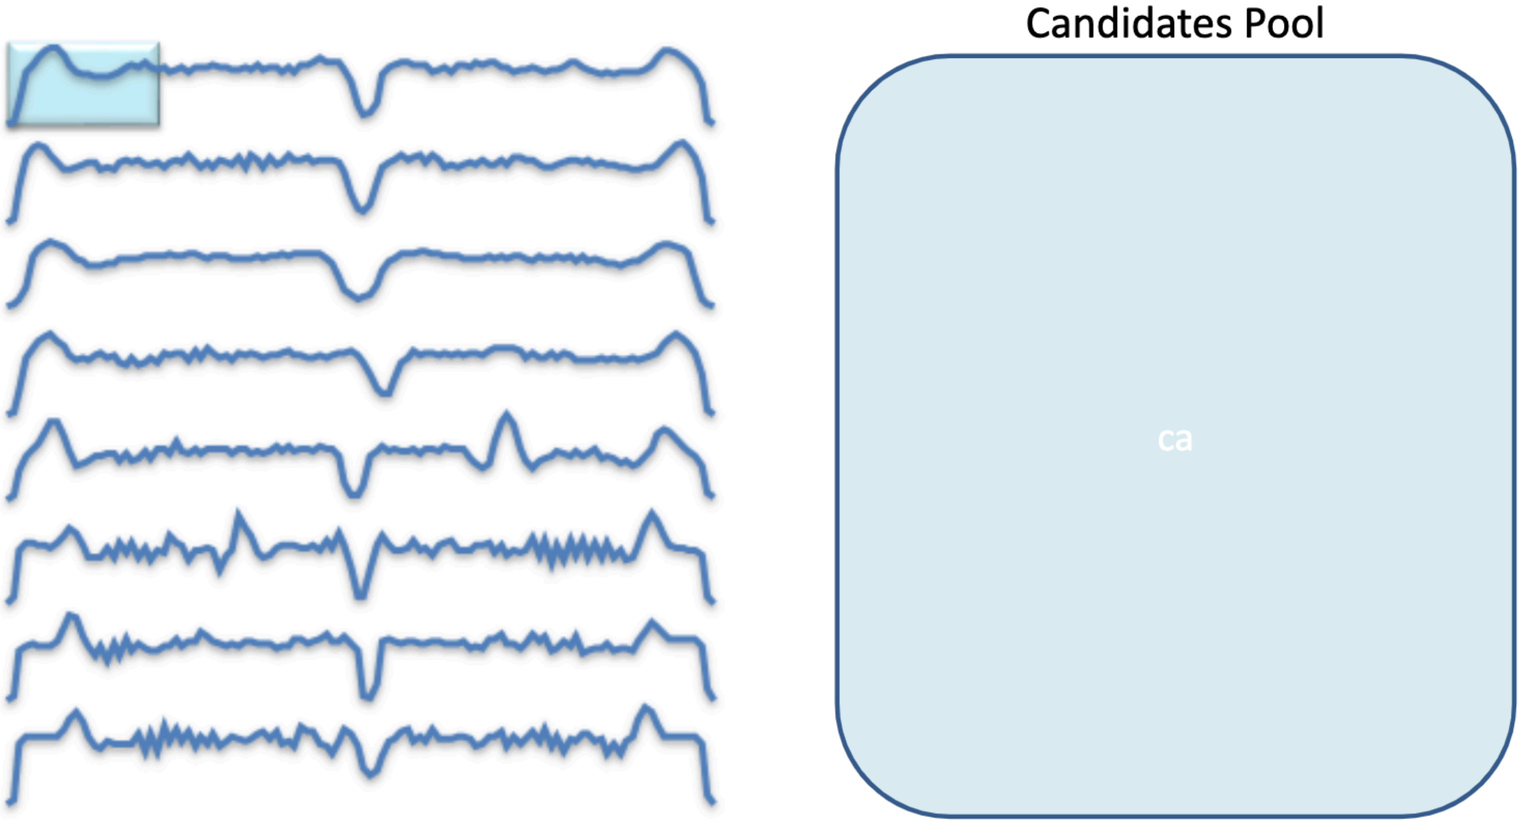

Candidates Pool

ca

# Extract Subsequences of all Possible Lengths

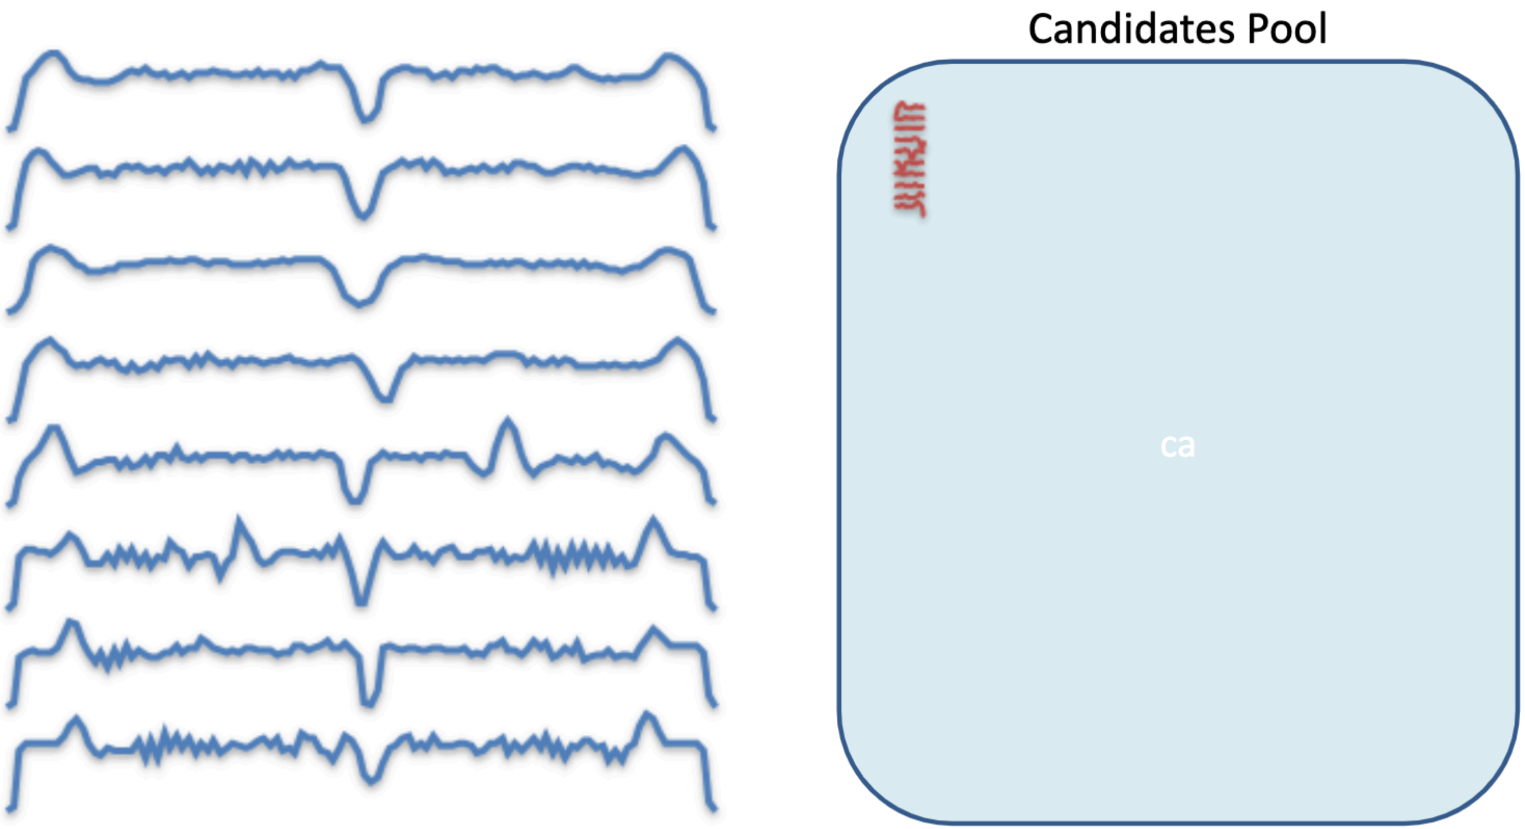

# Extract Subsequences of all Possible Lengths

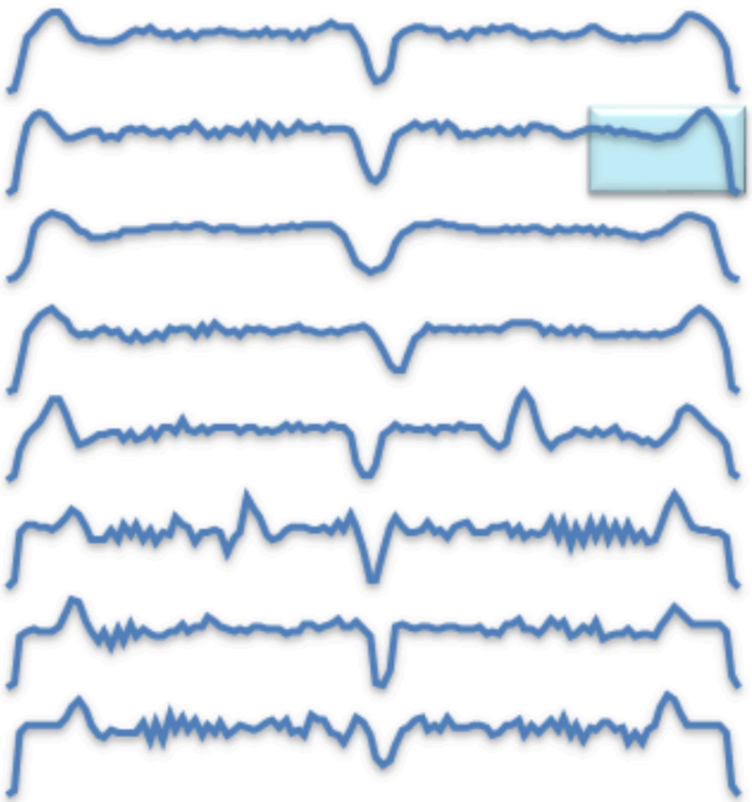

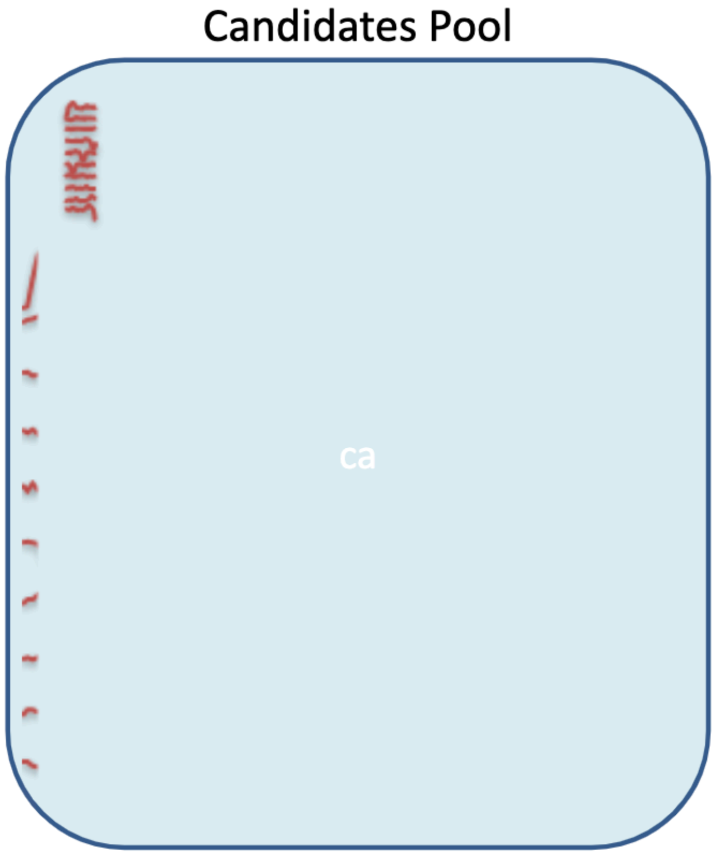

# Extract Subsequences of all Possible Lengths

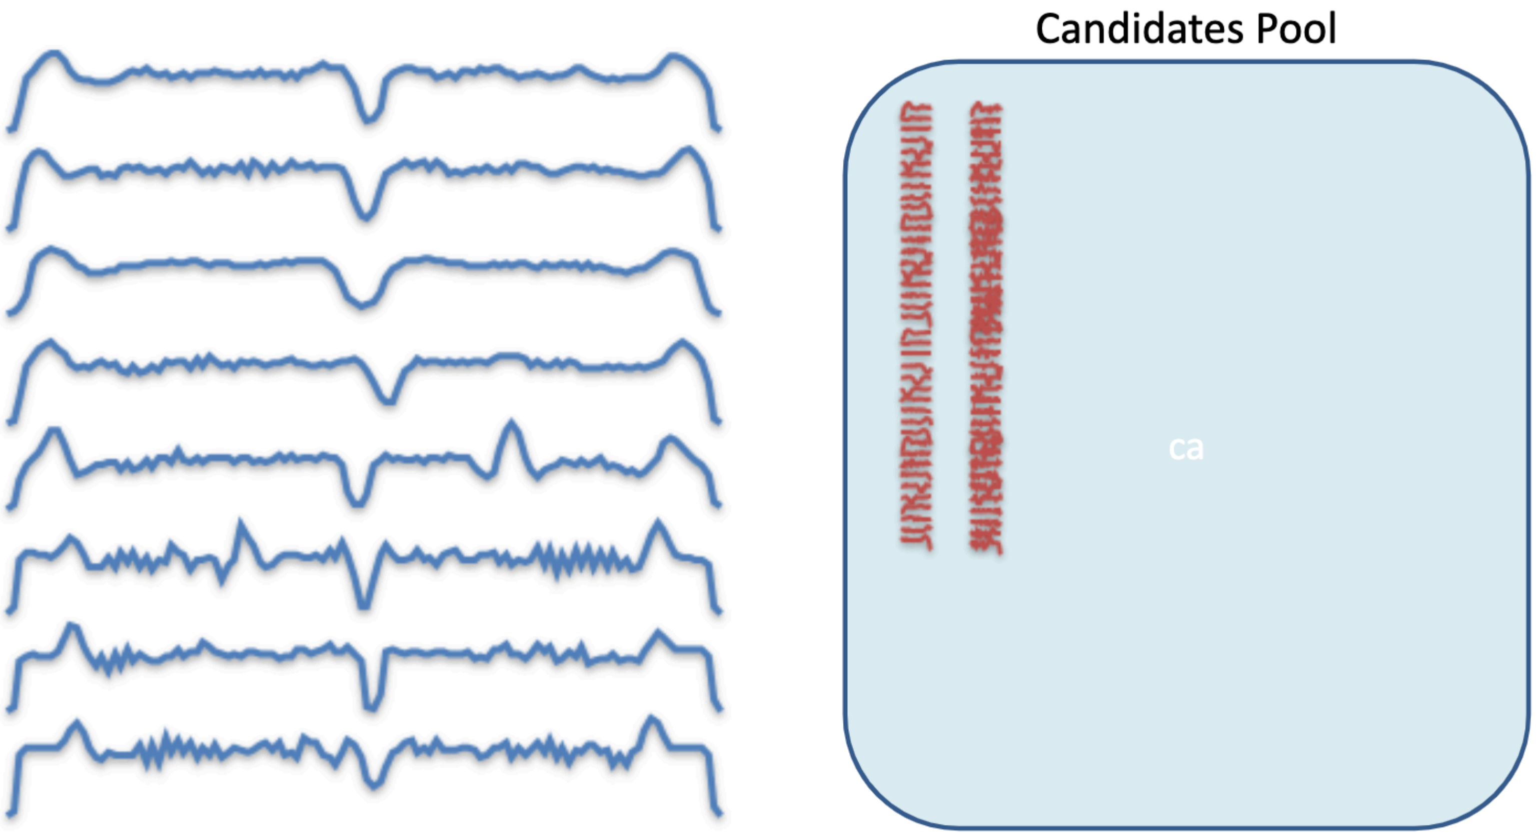

# Extract Subsequences of all Possible Lengths

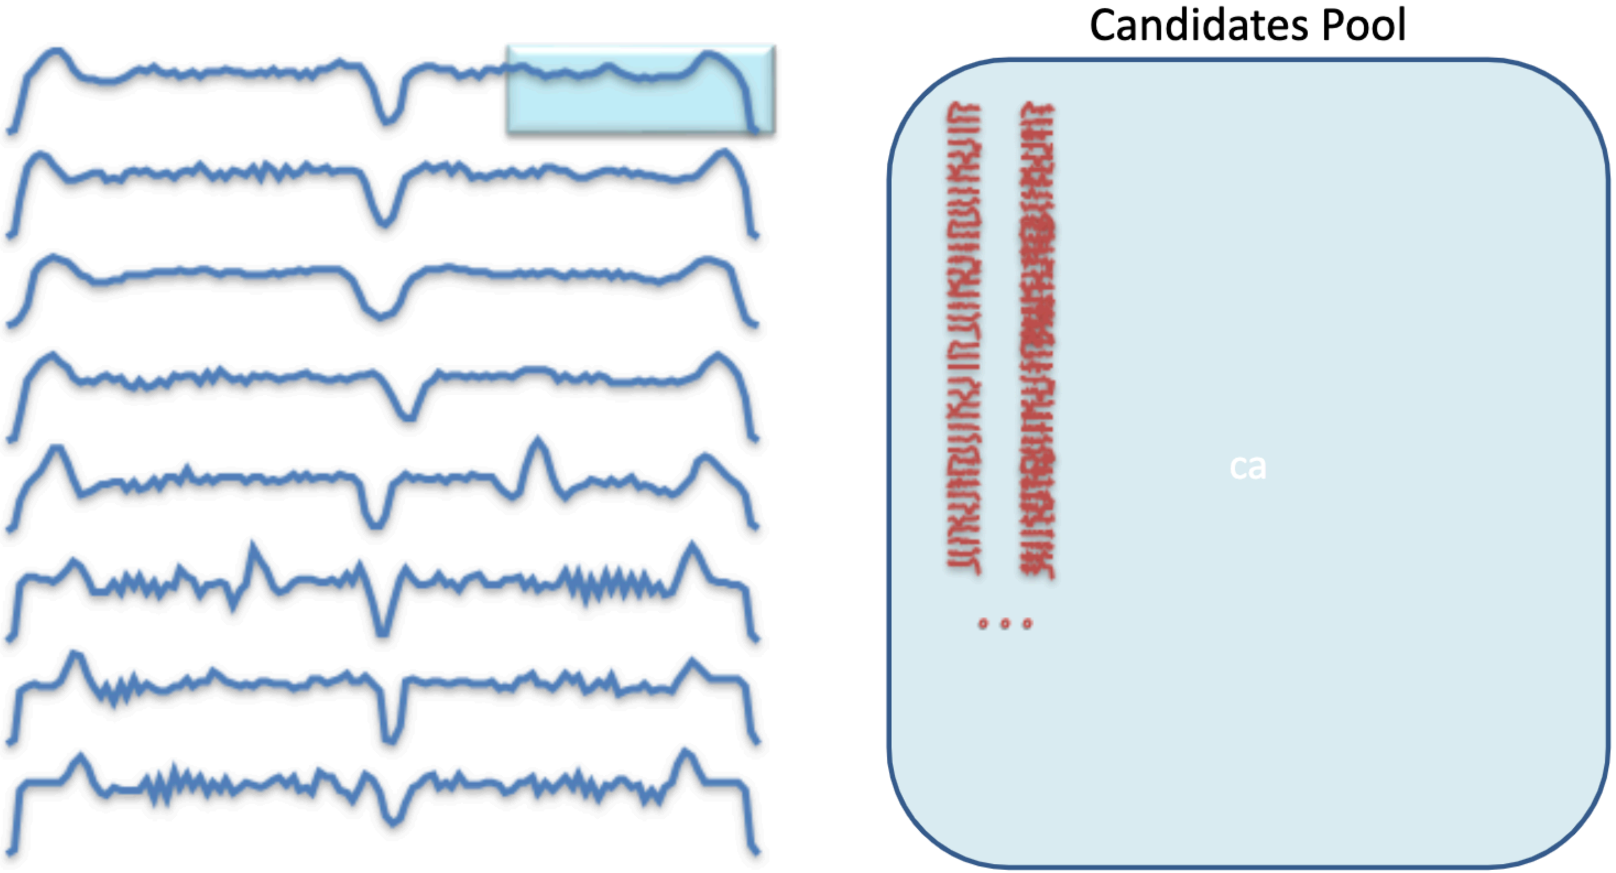

# Distance with a Subsequence

- Distance from the TS to the subsequence *SubsequenceDist(T, S)* is a distance function that takes time series *T* and subsequence *S* as inputs and returns a nonnegative value *d*, which is the distance from *T* to *S*.
- *SubsequenceDist(T, S) = min(Dist(S, S')), for S'* $\in S_T^{|S|}$
- where $S_T^{|S|}$ is the set of all possible subsequences of *T*
- Intuitively, it is the distance between *S* and its best matching location in *T*.

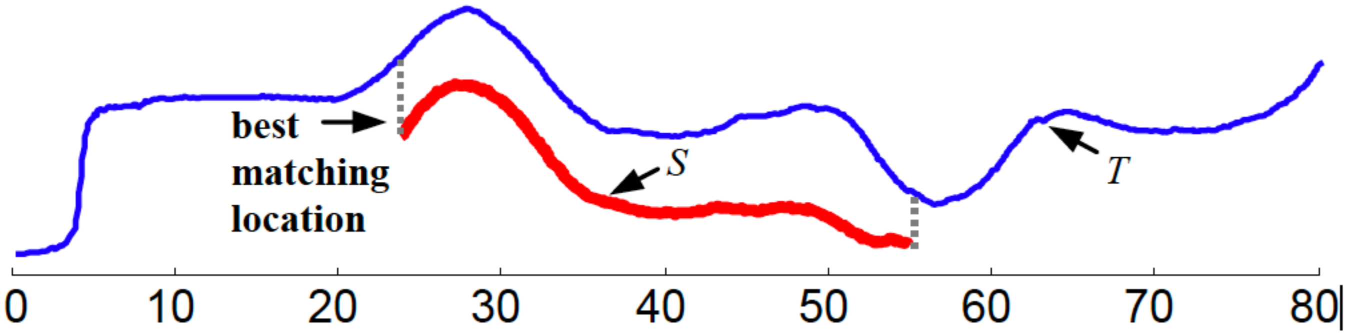

# Testing The Utility of a Candidate Shapelet

- Arrange the TSs in the dataset *D* based on the distance from the candidate.
- Find the optimal split point that maximizes the information gain (same as for Decision Tree classifiers)
- Pick the candidate achieving best utility as the shapelet

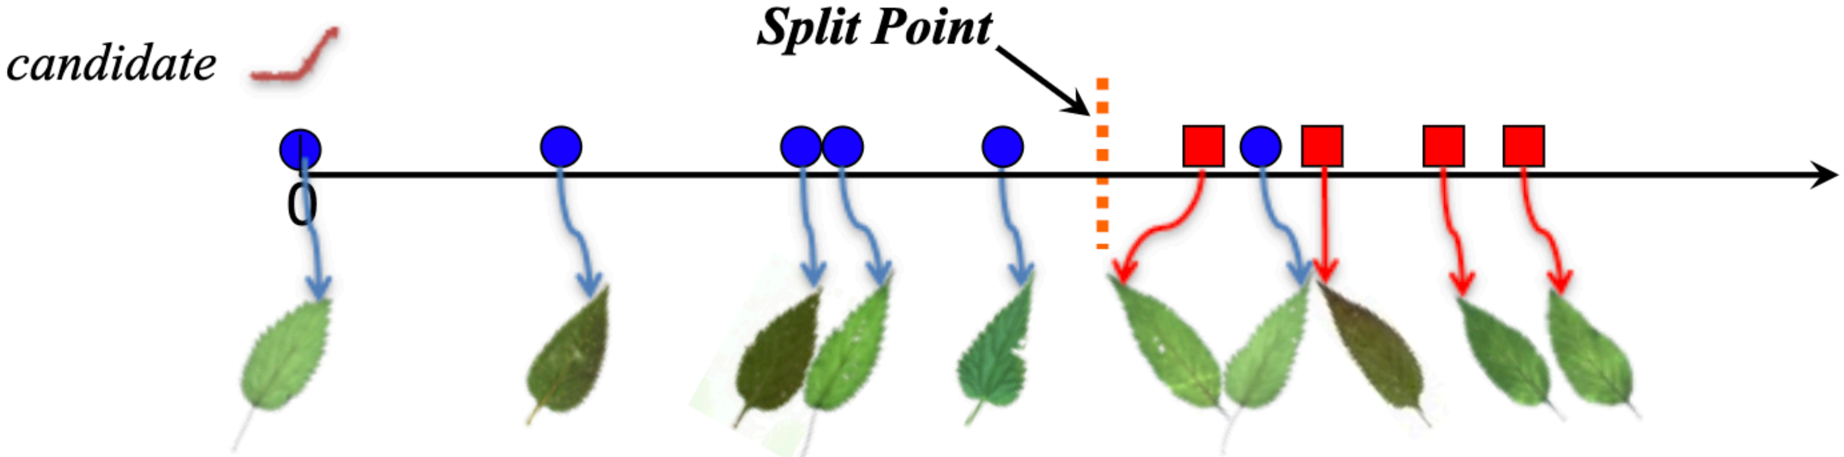

# Entropy

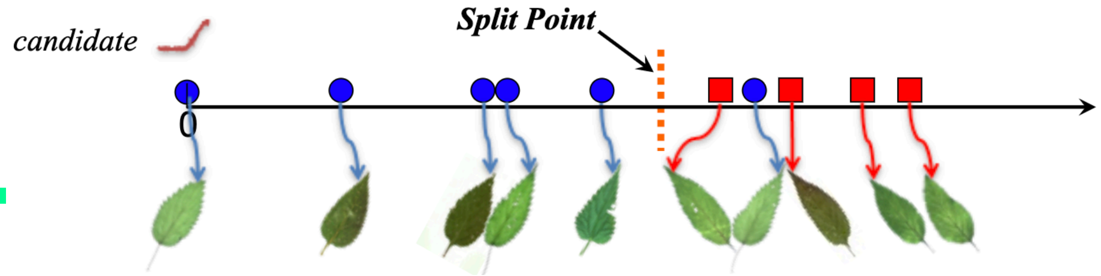

- A TS dataset *D* consists of two classes, A and B.
- Given that the proportion of objects in class A is *p(A)* and the proportion of objects in class B is *p(B),*
- The **Entropy** of D is: $I(D) = -p(A)log(p(A)) -p(B)log(p(B))$.
- Given a strategy that divides the *D* into two subsets $D_1$ and $D_2$, the information remaining in the dataset after splitting is defined by the weighted average entropy of each subset.
- If the fraction of objects in $D_1$ is $f(D_1)$ and in $D_2$ is $f(D_2)$,
- The total entropy of *D* after splitting is $\hat{I}(D) = f(D_1)I(D_1) + f(D_2)I(D_2)$.

# Information Gain

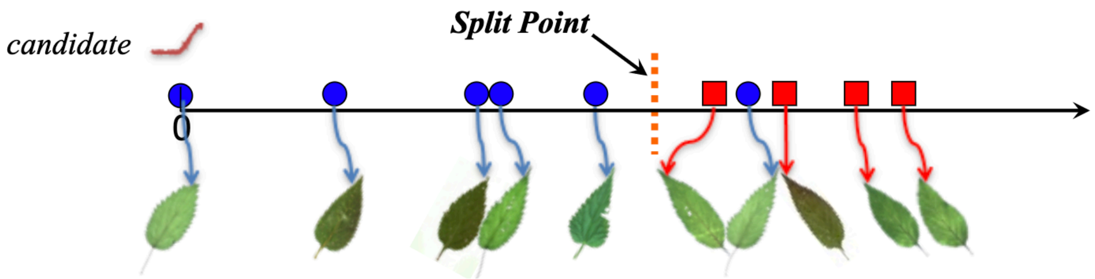

Split point distance from shapelet = 5.1

- Given a certain split strategy *sp* which divides *D* into two subsets $D_1$ and $D_2$, the entropy before and after splitting is $I(D)$ and $\hat{I}(D)$.
- The **information gain** for this splitting rule is:
- $Gain(sp) = I(D) - \hat{I}(D) =$
- $= I(D) - f(D_1)I(D_1) + f(D_2)I(D_2)$.
- We use the distance from *T* to a shapelet *S* as the splitting rule *sp*.

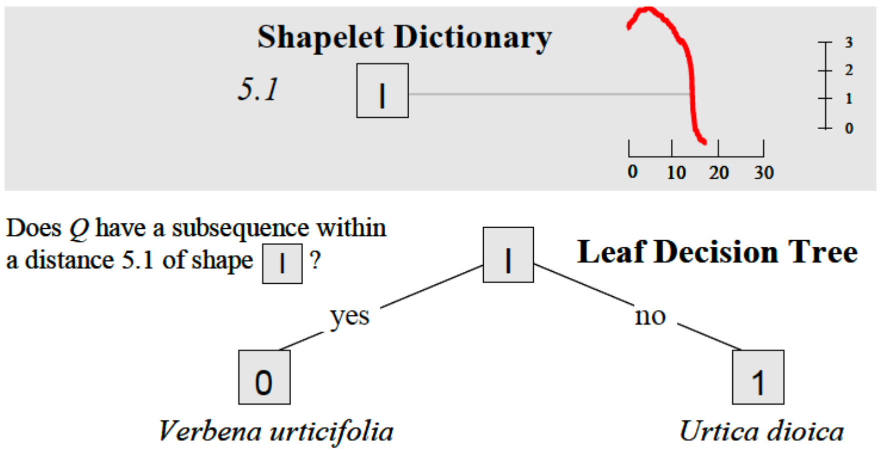

Does $Q$ have a subsequence within a distance 5.1 of shape I ?

# Problem

- The total number of candidate is

$$\sum_{l=MINLEN}^{MAXLEN} \sum_{T_i \in D} (|T_i| - l + 1)$$

- For each candidate you have to compute the distance between this candidate and each training sample
- For instance
  - 200 instances with length 275
  - 7,480,200 shapelet candidates

# Speedup

- Distance calculations form TSs to shapelet candidates is expensive.
- Reduce the time in two ways
- Distance Early Abandon
  - reduce the distance computation time between two TS
- Admissible Entropy Pruning
  - reduce the number of distance calculations

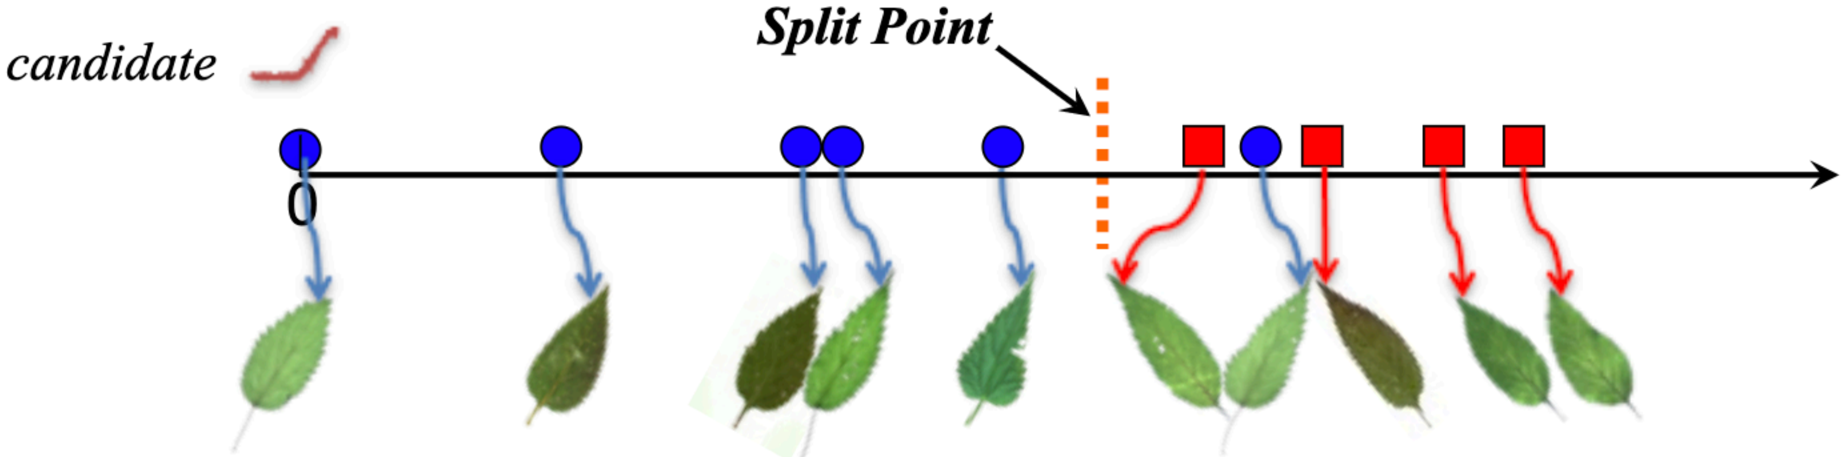

# Distance Early Abandon

- We only need the minimum distance.

- Method
  - Keep the best-so-far distance
  - Abandon the calculation if the current distance is larger than best-so-far.

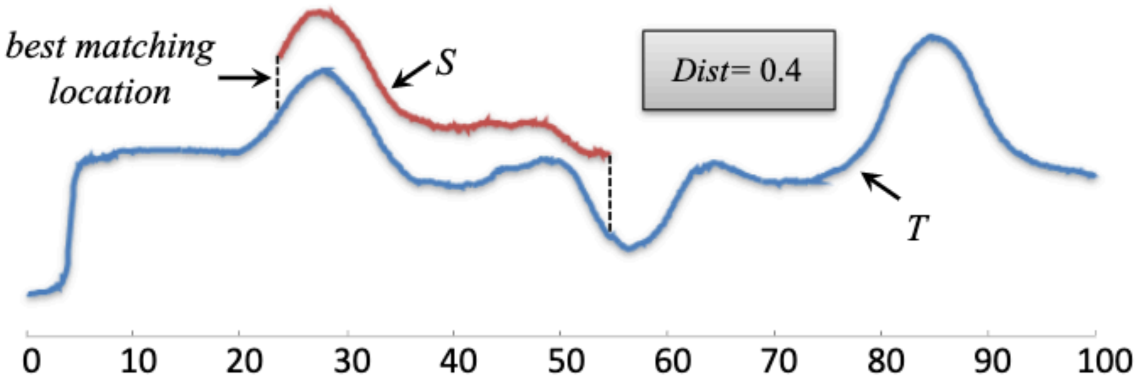

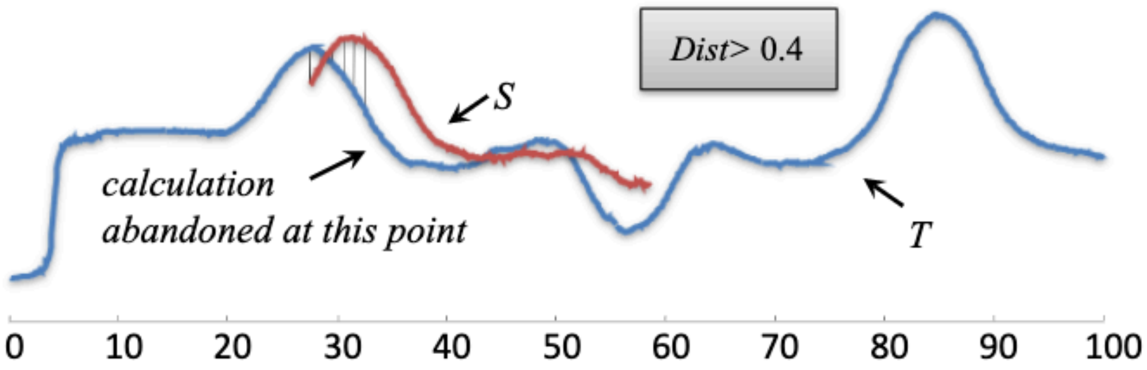

# Admissible Entropy Prunining

- We only need the best shapelet for each class
- For a candidate shapelet
  - We do not need to calculate the distance for each training sample
  - After calculating some training samples, the upper bound of information gain < best candidate shapelet
  - Stop calculation
  - Try next candidate

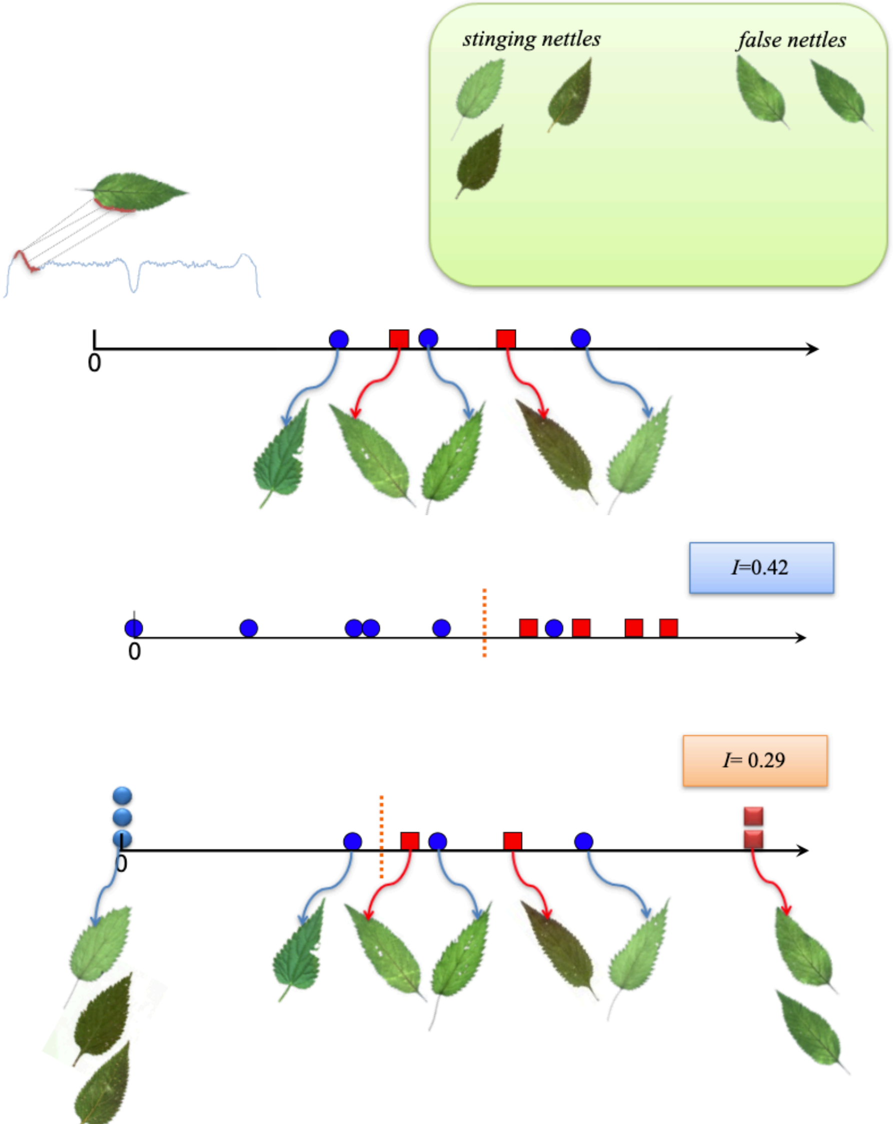

# An Alternative Way for Extracting Shapelets

- The minimum distances (M) between Ts and Shapelets can be used as predictors to approximate the TSs label (Y) using a linear model (W):

$$\hat{Y}_i \quad = \quad W_0 + \sum_{k=1}^{K} M_{i,k} W_k, \qquad \forall i \in \{1, \ldots, I\}$$

- A logistic regression loss can measure the quality of the prediction:

$$\mathcal{L}(Y, \hat{Y}) \quad = \quad -Y \ln \sigma(\hat{Y}) - (1 - Y) \ln \left(1 - \sigma(\hat{Y})\right)$$

- The objective is to minimize a regularized loss function accross all the instances (I) :

$$\underset{S,W}{\operatorname{argmin}} \; \mathcal{F}(S, W) = \underset{S,W}{\operatorname{argmin}} \sum_{i=1}^{I} \mathcal{L}(Y_i, \hat{Y}_i) + \lambda_W ||W||^2$$

- We can find the optimal shapelet for the objective function in a NN fashion by updating the shapelets in the minimum direction of the objective, hence the first gradient. Similarly, the weights can be jointly updated towards minimizing the objective function.

# Motif/Shapelet Summary

- A **motif** is a repeated pattern/subsequence in a given TS.

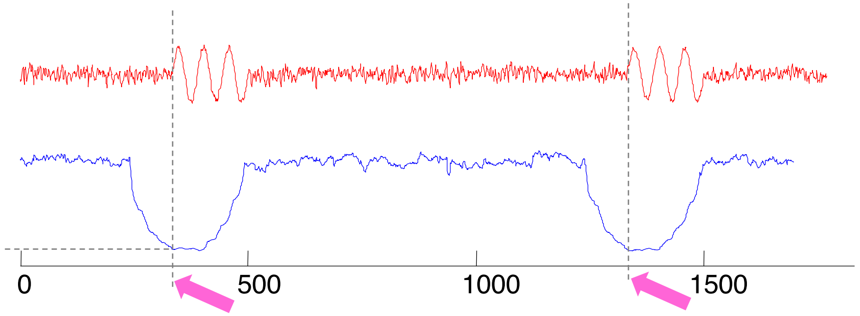

- A **shapelet** is a pattern/subsequence which is maximally representative of a class with respect to a given dataset of TSs.

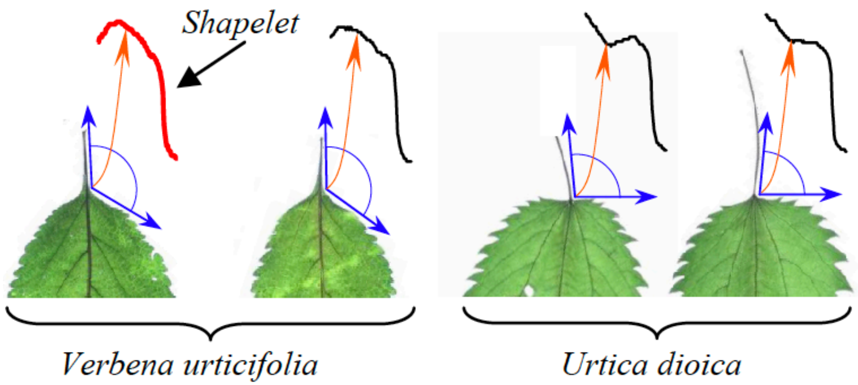

# References

- Matrix Profile I: All Pairs Similarity Joins for Time Series:  A Unifying View that Includes Motifs, Discords and Shapelets. Chin-Chia Michael Yeh et al. 1997
- Time Series Shapelets: A New Primitive for Data Mining. Lexiang Ye and Eamonn Keogh. 2016.
- Josif Grabocka, Nicolas Schilling, Martin Wistuba, Lars Schmidt-Thieme (2014): Learning Time-Series Shapelets, in Proceedings of the 20th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2014
- Deep learning for time series classication: a review. Hassan Ismail Fawaz et al. 2019.

Matrix Profile I: All Pairs Similarity Joins for Time Series:
A Unifying View that Includes Motifs, Discords and Shapelets

Chin-Chia Michael Yeh, Yan Zhu, Liudmila Ulanova, Nurjahan Begum, Yifei Ding,
Hoang Anh Dau, †Diego Furtado Silva, ‡Abdullah Mueen, and Eamonn Keogh
University of California, Riverside, †Universidade de São Paulo, ‡University of New Mexico
{myeh003, yzhu015, lulan001, nbegu001, yding007, hdau001}@ucr.edu, diegofsilva@icmc.usp.br, mueen@unm.edu, eamonn@cs.ucr.edu

***Abstract*— The *all-pairs-similarity-search* (or *similarity join*) problem has been extensively studied for text and a handful of other datatypes. However, surprisingly little progress has been made on similarity joins for time series subsequences. The lack of progress probably stems from the daunting nature of the problem. For even modest sized datasets the obvious nested-loop algorithm can take months, and the typical speed-up techniques in this domain (i.e., indexing, lower-bounding, triangular-inequality pruning and early abandoning) at best produce one or two orders of magnitude speedup. In this work we introduce a novel scalable algorithm for time series subsequence all-pairs-similarity-search. For exceptionally large datasets, the algorithm can be trivially cast as an anytime algorithm and produce high-quality approximate solutions in reasonable time. The exact similarity join algorithm computes the answer to the *time series motif* and *time series discord* problem as a side-effect, and our algorithm incidentally provides the fastest known algorithm for both these extensively-studied problems. We demonstrate the utility of our ideas for many time series data mining problems, including motif discovery, novelty discovery, shapelet discovery, semantic segmentation, density estimation, and contrast set mining.**

***Keywords—Time Series; Similarity Joins; Motif Discovery***

I. INTRODUCTION

The *all-pairs-similarity-search* (also known as *similarity join*) problem comes in several variants. The basic task is this: *Given a collection of data objects, retrieve the nearest neighbor for each object*. In the text domain the algorithm has applications in a host of problems, including community discovery, duplicate detection, collaborative filtering, clustering, and query refinement [1]. While virtually all text processing algorithms have analogues in time series data mining, there has been surprisingly little progress on Time Series subsequences All-Pairs-Similarity-Search (TSAPSS).

We believe that this lack of progress stems not from a lack of interest in this useful primitive, but from the daunting nature of the problem. Consider the following example that reflects the needs of an industrial collaborator. A boiler at a chemical refinery reports pressure once a minute. After a year, we have a time series of length 525,600. A plant manager may wish to do a similarity self-join on this data with week-long subsequences (10,080) to discover operating regimes (summer vs. winter or light distillate vs. heavy distillate etc.) The obvious nested loop algorithm requires 132,880,692,960 Euclidean distance computations. If we assume each one takes 0.0001 seconds, then the join will take 153.8 days. The core contribution of this work is to show that we can reduce this time to 6.3 hours, using an off-the-shelf desktop computer. Moreover, we show that this join can be computed and/or updated incrementally. Thus we could maintain this join essentially forever on a standard desktop, even if the data arrival frequency was much faster than once a minute.

Our algorithm uses an ultra-fast similarity search algorithm under z-normalized Euclidean distance as a subroutine, exploiting the overlap between subsequences using the classic Fast Fourier Transform (FFT) algorithm.

Our method has the following advantages/features:

- It is *exact*, providing no false positives or false dismissals.
- It is simple and parameter-free. In contrast, the more general metric space APSS algorithms require building and tuning spatial access methods and/or hash functions.
- Our algorithm requires an inconsequential space overhead, just O(*n*) with a small constant factor.
- While our *exact* algorithm is extremely scalable, for extremely large datasets we can compute the results in an anytime fashion, allowing ultra-fast *approximate* solutions.
- Having computed the similarity join for a dataset, we can incrementally update it very efficiently. In many domains this means we can effectively maintain exact joins on streaming data forever.
- Our method provides *full* joins, eliminating the need to specify a similarity *threshold*, which as we will show, is a near impossible task in this domain.
- Our algorithm is embarrassingly parallelizable, both on multicore processors and in distributed systems

Time Series Shapelets: A New Primitive for Data Mining

Lexiang Ye
Dept. of Computer Science & Engineering
University of California, Riverside, CA 92521
lexiangy@cs.ucr.edu

Eamonn Keogh
Dept. of Computer Science & Engineering
University of California, Riverside, CA 92521
eamonn@cs.ucr.edu

**ABSTRACT**

Classification of time series has been attracting great interest over the past decade. Recent empirical evidence has strongly suggested that the simple nearest neighbor algorithm is very difficult to beat for most time series problems. While this may be considered good news, given the simplicity of implementing the nearest neighbor algorithm, there are some negative consequences of this. First, the nearest neighbor algorithm requires storing and searching the entire dataset, resulting in a time and space complexity that limits its applicability, especially on resource-limited sensors. Second, beyond mere classification accuracy, we often wish to gain some insight into the data.

In this work we introduce a new time series primitive, *time series shapelets*, which addresses these limitations. Informally, shapelets are time series subsequences which are in some sense maximally representative of a class. As we shall show with extensive empirical evaluations in diverse domains, algorithms based on the time series shapelet primitives can be interpretable, more accurate and significantly faster than state-of-the-art classifiers.

**Categories and Subject Descriptors**

H.2.8 [Database Management]: Database Applications – Data Mining

**General Terms**

Algorithms, Experimentation

**1. INTRODUCTION**

While the last decade has seen a huge interest in time series classification, to date the most accurate and robust method is the simple nearest neighbor algorithm [4][12][14]. While the nearest neighbor algorithm has the advantages of simplicity and not requiring extensive parameter tuning, it does have several important disadvantages. Chief among these are its space and time requirements, and the fact that it does not tell us anything about *why* a particular object was assigned to a particular class.

In this work we present a novel time series data mining primitive called *time series shapelets*. Informally, shapelets are time series subsequences which are in some sense maximally representative of a class. While we believe shapelets can have many uses in data mining, one obvious implication of them is to mitigate the two weaknesses of the nearest neighbor algorithm noted above.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
*KDD'09*, June 29–July 1, 2009, Paris, France
Copyright 2009 ACM 978-1-60558-495-9/09/06….$5.00.

Because we are defining and solving a new problem, we will take some time to consider a detailed motivating example. Figure 1 shows some examples of leaves from two classes, *Urtica dioica* (stinging nettles) and *Verbena urticifolia*. These two plants are commonly confused, hence the colloquial name "false nettle" for *Verbena urticifolia*.

**Figure 1: Samples of leaves from two species. Note that several leaves have the insect-bite damage**

Suppose we wish to build a classifier to distinguish these two plants; what features should we use? Since the intra-variability of color and size within each class completely dwarfs the inter-variability between classes, our best hope is based on the shapes of the leaves. However, as we can see in Figure 1, the differences in the global shape are very subtle. Furthermore, it is very common for leaves to have distortions or "occlusions" due to insect damage, and these are likely to confuse any global measures of shape. Instead we attempt the following. We first convert each leaf into a one-dimensional representation as shown in Figure 2.

**Figure 2: A shape can be converted into a one dimensional "time series" representation. The reason for the highlighted section of the time series will be made apparent shortly**

Such representations have been successfully used for the classification, clustering and outlier detection of shapes in recent years [8]. However, here we find that using a nearest neighbor classifier with either the (rotation invariant) Euclidean distance or Dynamic Time Warping (DTW) distance does not significantly outperform random guessing. The reason for the poor performance of these otherwise very competitive classifiers seems to be due to the fact that the data is somewhat noisy (i.e. insect bites, and different stem lengths), and this noise is enough to swamp the subtle differences in the shapes.

This is the author's version of an article published in Data Mining and
authenticated version is available online at: https://doi.org/10.1007/s

arXiv:1809.04356v4 [cs.LG] 14 May 2019

**Deep learning for time series classificatio**

**Hassan Ismail Fawaz[1] · Germain Forestier[1,2] · Jonathan W**
**Lhassane Idoumghar[1] · Pierre-Alain Muller[1]**

**Abstract** Time Series Classification (TSC) is an important and challe
With the increase of time series data availability, hundreds of TSC a
Among these methods, only a few have considered Deep Neural Net
task. This is surprising as deep learning has seen very successful appli
have indeed revolutionized the field of computer vision especially w
architectures such as Residual and Convolutional Neural Networks.
data such as text and audio can also be processed with DNNs to reac
for document classification and speech recognition. In this article,
the-art performance of deep learning algorithms for TSC by presen
most recent DNN architectures for TSC. We give an overview of the
applications in various time series domains under a unified taxonon
provide an open source deep learning framework to the TSC communi
of the compared approaches and evaluated them on a univariate TS
archive) and 12 multivariate time series datasets. By training 8,73
time series datasets, we propose the most exhaustive study of DNN

**Keywords** Deep learning · Time series · Classification · Review

**1 Introduction**

During the last two decades, Time Series Classification (TSC) has been considered as one of the most challenging problems in data mining (Yang and Wu, 2006; Esling and Agon, 2012). With the increase of temporal data availability (Silva et al., 2018), hundreds of TSC algorithms have been proposed since 2015 (Bagnall et al., 2017). Due to their natural temporal ordering, time series data are present in almost every task that requires some sort of human cognitive process (Längkvist et al., 2014). In fact, any classification problem, using data that is registered taking into account some notion of ordering, can be cast as a TSC problem (Cristian Borges Gamboa, 2017). Time series are encountered in many real-world applications ranging from electronic health records (Rajkomar et al., 2018) and human activity recognition (Nweke et al., 2018; Wang et al., 2018) to acoustic scene classification (Nwe et al., 2017) and cyber-security (Susto et al., 2018). In addition, the diversity of the datasets' types in the UCR/UEA archive (Chen et al., 2015b; Bagnall et al., 2017) (the largest repository of time series datasets) shows the different applications of the TSC problem.

✉ H. Ismail Fawaz
E-mail: hassan.ismail-fawaz@uha.fr

[1] IRIMAS, Université Haute Alsace, Mulhouse, France
[2] Faculty of IT, Monash University, Melbourne, Australia