Abdullah Mueen

Eamonn Keogh

Time Series Data Mining Using the Matrix Profile: A Unifying View of Motif Discovery, Anomaly Detection, Segmentation, Classification, Clustering and Similarity Joins

IEEE BigData 2017

To get these slides in PPT or PDF, go to www.cs.ucr.edu/~eamonn/MatrixProfile.html

This tutorial is based on work by:

Chin-Chia Michael Yeh, **Yan Zhu**, **Abdullah Mueen**

Nurjahan Begum, Yifei Ding, Hoang Anh Dau

Diego Furtado Silva, Liudmila Ulanova, Eamonn Keogh

Zachary Zimmerman, Nader S. Senobari, Philip Brisk

Shaghayegh Gharghabi, Kaveh Kamgar

..and others that have inspired us, forgive any omissions.

Your face here

This tutorial is funded by:

- NSF IIS-1161997 II
- NSF IIS 1510741
- NSF 544969
- CNS 1544969
- SHF-1527127
- AFRL FA9453-17-C-0024

Any errors or controversial statements are due solely to Mueen and Keogh

Disclaimer:

Time series is an inherently *visual* domain, and we exploit that fact in this tutorial.

We therefore keep formal notations and proofs to an absolute minimum.

If you want them, you can read the relevant papers [a]

All the datasets used in tutorial are freely available, all experiments are reproducible.

[a] www.cs.ucr.edu/~eamonn/MatrixProfile.html

If you enjoy this tutorial, please check out our other tutorials..

www.cs.ucr.edu/~eamonn/public/SDM_How_to_do_Research_Keogh.pdf

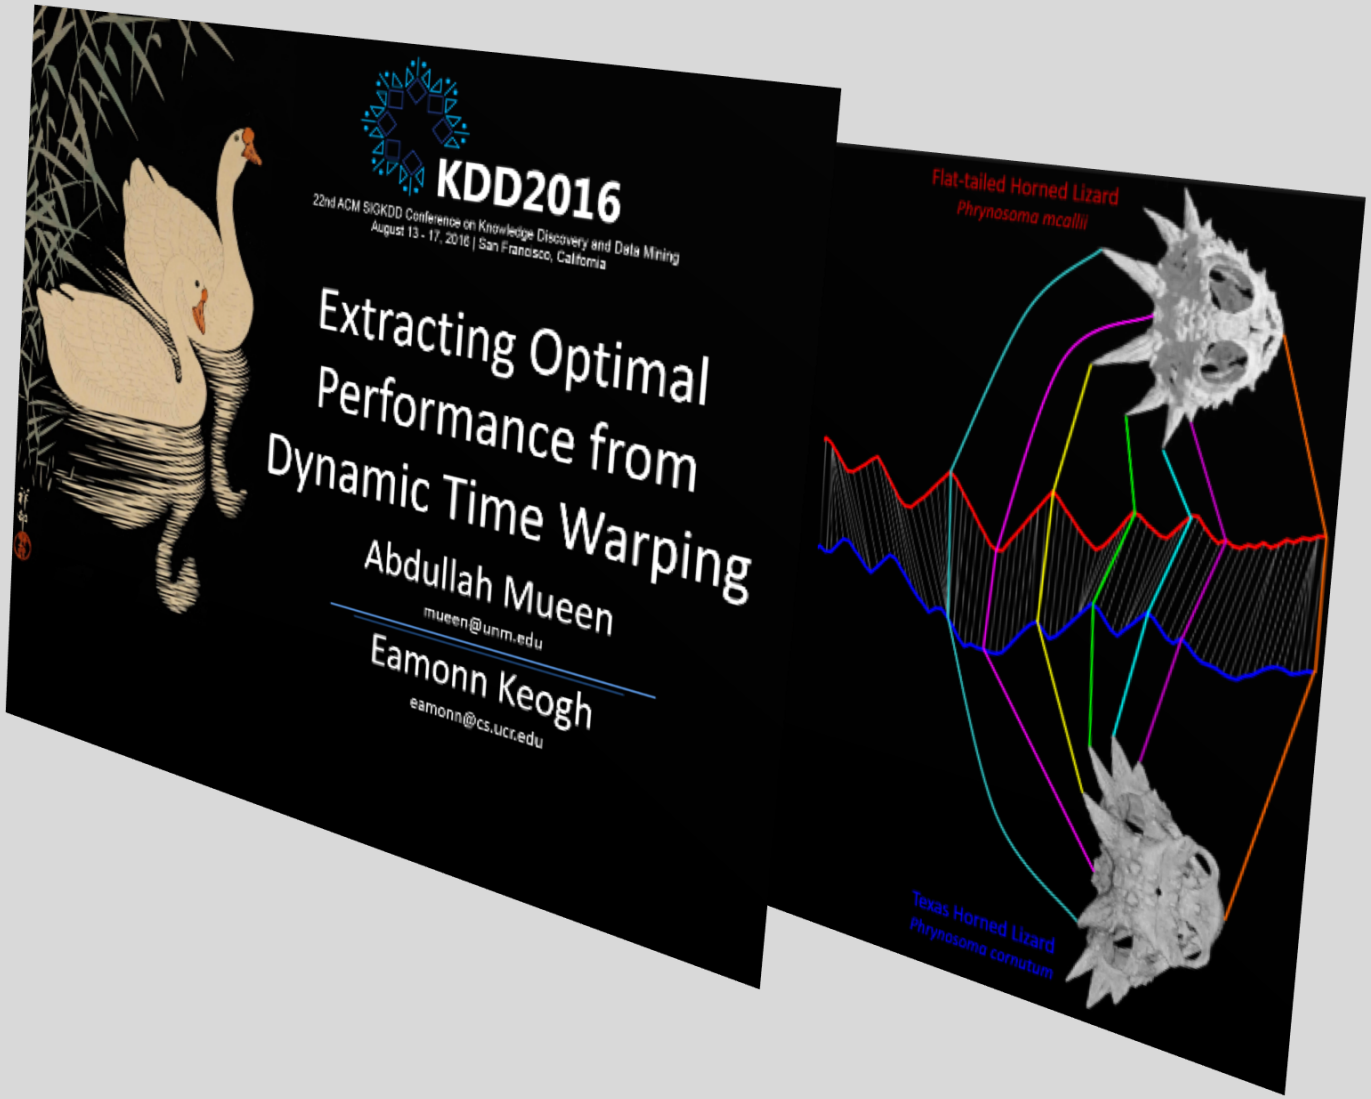

www.cs.unm.edu/~mueen/DTW.pdf

Outline

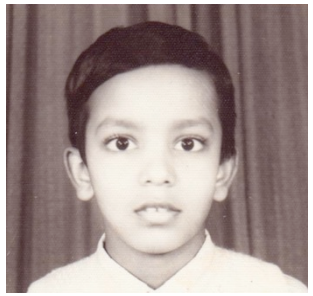

Act 1

- Our Fundamental Assumption
- What is the (MP) Matrix Profile?
- Properties of the MP
- Developing a Visual Intuition for MP
- Basic Algorithms
  - MP Motif Discovery
  - MP Time Series Chains
  - MP Anomaly Discovery
  - MP Joins (self and AB)
  - MP Semantic Segmentation
- From Domain Agnostic to Domain Aware: The Annotation Vector (A simple way to use domain knowledge to adjust your results)
- The "*Matrix Profile and ten lines of code is all you need*" philosophy.
- Break

Act 2

- Background on time series mining
  - Similarity Measures
  - Normalization
- Distance Profile
  - Definition and Trivial Approach
  - Just-in-time Normalization
  - The MASS Algorithm
  - Weighted Distance Profile
  - Distance Profile with Gaps
- Matrix Profile
  - STAMP
  - STOMP
  - GPU-STOMP
- Open problems to solve

Fundamental Assumption: *Conservation is Key*

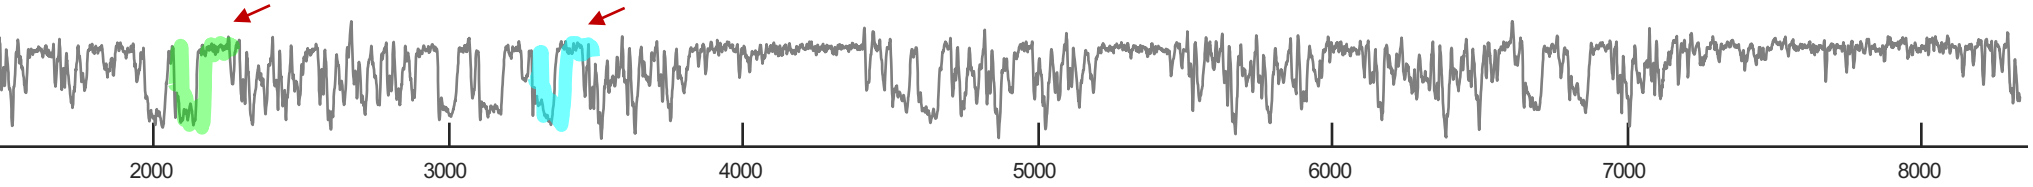

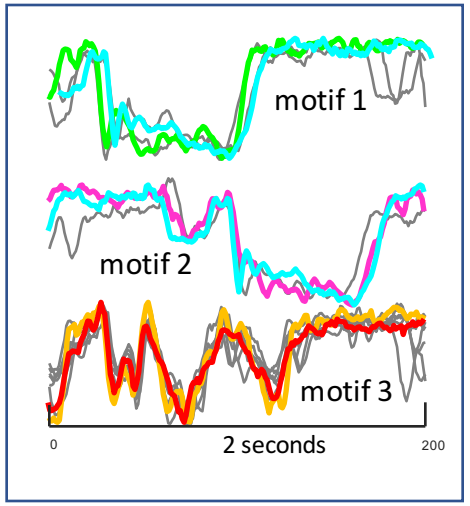

If a pattern is *conserved*, there must be some mechanism that conserves it. This is true in linguistics, music, genetics, literature, religions….

Much of our work asks *what* is conserved in time series, *when* is it conserved, and *why* was an expected conservation not observed…

For discrete strings, *conserved* is easy to define, for example `papa` `=*a*a`. For *time series* it requires a distance function, here we will use Euclidean Distance.

* Bengali: Bābā
* Mandarin : baba
* Polish : tata
* Swahili : baba
* Turkish : baba
* Xhosa: -tata
* Norwegian : papa
* Spanish : papá
* Swahili : baba
* English : papa
* Hindi : papa
* Indonesian : bapa

en.wikipedia.org/wiki/Mama_and_papa

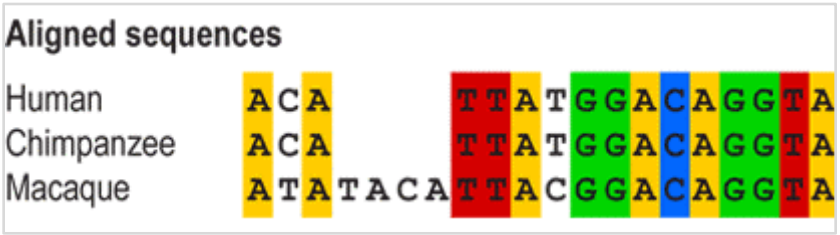

What is the Matrix Profile?

- The Matrix Profile (MP) is a data structure that annotates a time series.
- **Key Claim:** Given the MP, most time series data mining problems are trivial or easy!
- We will show about ten problems that are trivial given the MP, including motif discovery, density estimation, anomaly detection, rule discovery, joins, segmentation, clustering etc. However, *you* can use the MP to solve *your* problems, or to solve a problem listed above, but in a different way, tailored to your interests/domain.

What is the Matrix Profile?

- The Matrix Profile (MP) is a data structure that annotates a time series.
- **Key Claim:** Given the MP, most time series data mining problems are trivial or easy!
- We will show about ten problems that are trivial given the MP, including motif discovery, density estimation, anomaly detection, rule discovery, joins, segmentation, clustering etc. However, *you* can use the MP to solve *your* problems, or to solve a problem listed above, but in a different way, tailored to your interests/domain.
- **Key Insight**: The MP profile has many highly desirable properties, and any algorithm you build on top of it, will inherit those properties.
  - Say you use the MP to create: *An Algorithm to Segment Sleep States*
  - Then, *for free*, you have also created: *An Anytime Algorithm to Segment Sleep States*
    An Online Algorithm to Segment Sleep States
    A Parallelizable Algorithm to Segment Sleep States
    A GPU Accelerated Algorithm to Segment Sleep States
    An Algorithm to Segment Sleep States with Missing Data
    etc.

The Highly Desirable Properties of the Matrix Profile I

- It is **exact**: For motif discovery, discord discovery, time series joins etc., the Matrix Profile based methods provide no false positives or false dismissals.
- It is **simple** and **parameter-free**: In contrast, the more general algorithms in this space that typically require building and tuning spatial access methods and/or hash functions.
- It is **space efficient**: Matrix Profile construction algorithms requires an inconsequential space overhead, just linear in the time series length with a small constant factor, allowing massive datasets to be processed in main memory (for most data mining, *disk is death*).
- It allows **anytime algorithms**: While exact MP algorithms are extremely scalable, for extremely large datasets we can compute the Matrix Profile in an anytime fashion, allowing ultra-fast approximate solutions and real-time data interaction.
- It is **incrementally maintainable**: Having computed the Matrix Profile for a dataset, we can incrementally update it very efficiently. In many domains this means we can effectively maintain exact joins/motifs/discords on streaming data forever.

The Highly Desirable Properties of the Matrix Profile II

- It can **leverage hardware**: Matrix Profile construction is embarrassingly parallelizable, both on multicore processors, GPUs, distributed systems etc.
- It is **free of the curse of dimensionality**: That is to say, It has *time complexity that is constant in subsequence length*: This is a very unusual and desirable property; virtually all existing algorithms in the time series scale poorly as the subsequence length grows.
- It can be constructed in **deterministic time**: Almost all algorithms for time series data mining can take radically different times to finish on two (even slightly) different datasets. In contrast, given *only* the length of the time series, we can precisely predict in advance how long it will take to compute the Matrix Profile. (this allows resource planning)
- It can handle **missing data**: Even in the presence of missing data, we can provide answers which are guaranteed to have no false negatives.
- Finally, and subjectively: **Simplicity and Intuitiveness:** Seeing the world through the MP lens often invites/suggests simple and elegant solutions.

Developing a Visual Intuition for the Matrix Profile

In the following slides we are going to develop a visual intuition for the matrix profile, *without* regard to how we obtain it.

We are ignoring the elephant in the room; the MP *seems* to be much too slow to compute to be practical. We will address this in Part II of the tutorial.

Note that algorithms that use the MP do *not* require us to visualize the MP, but it happens that *just* visualizing the MP can be 99% of solving a problem.

Intuition behind the Matrix Profile: Assume we have a time series *T*, lets start with a synthetic one...

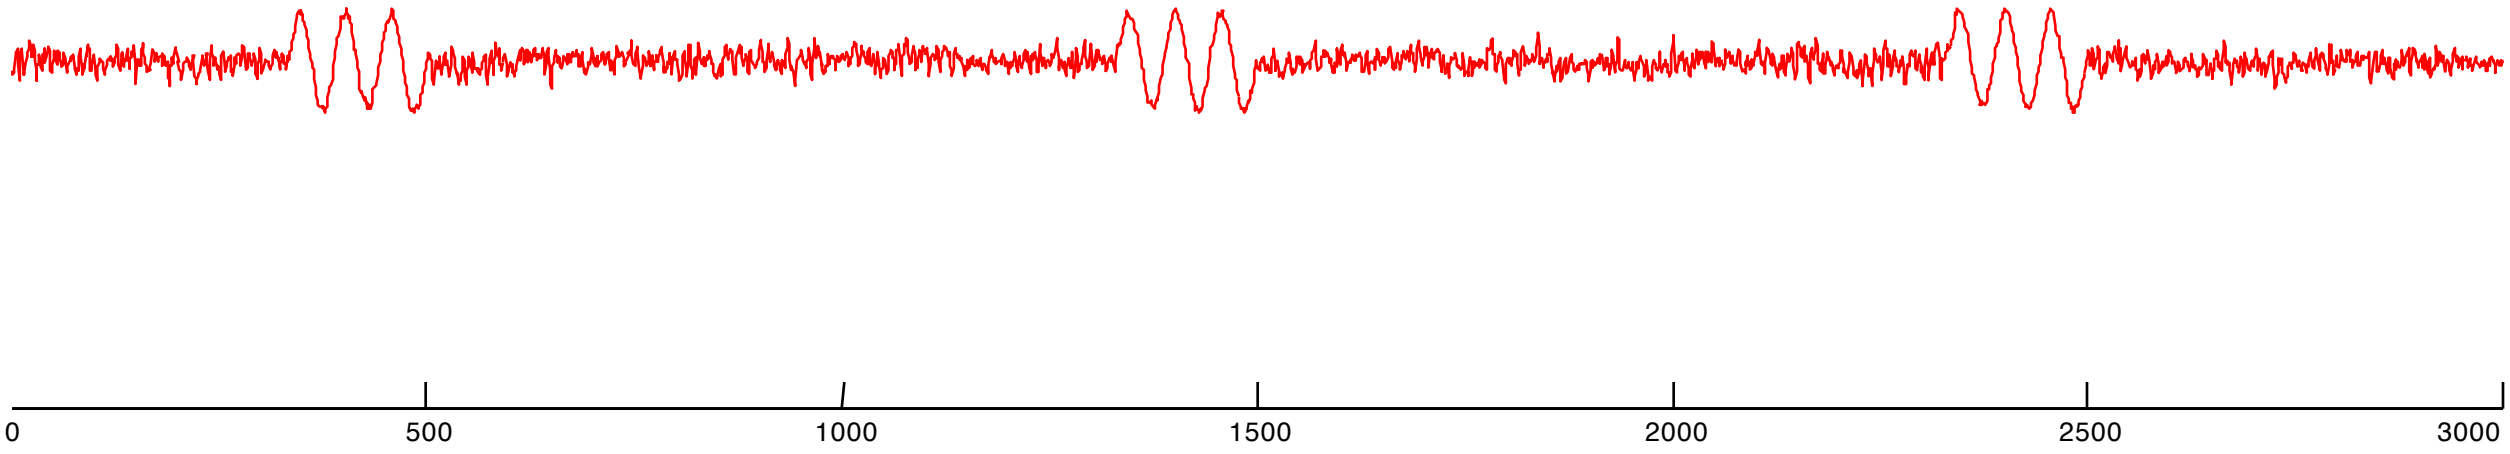

$$|T| = n = 3{,}000$$

Note that for most time series data mining tasks, we are not interested in any *global* properties of the time series, we are only interested in small *local* subsequences, of this length, *m*

These subsequences might be about the length of individual heartbeats (for ECGs), individual days (for social media behavior), individual words (for speech analysis) etc

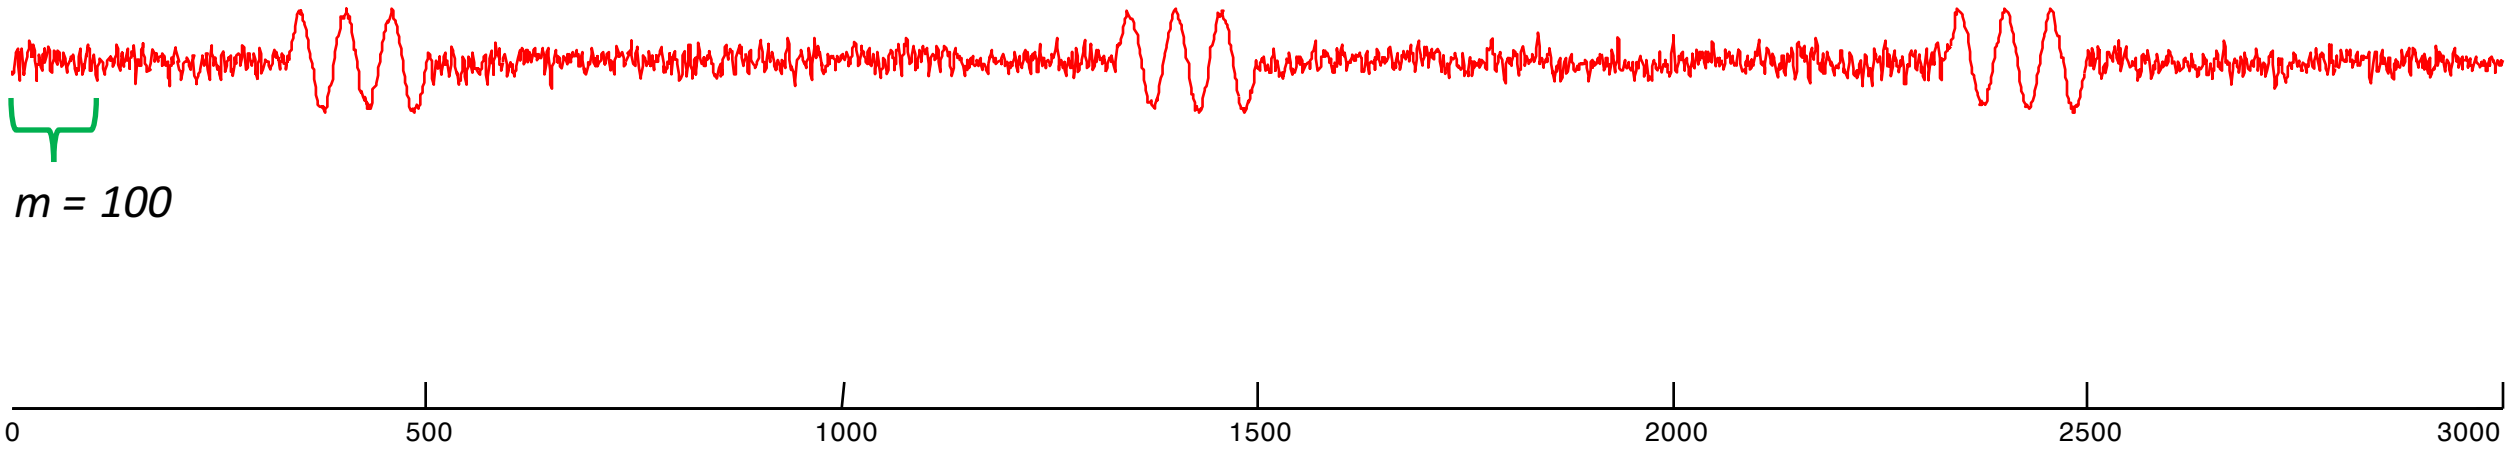

We can created a companion “time series”, called a Matrix Profile or MP.

The matrix profile at the i^{th} location records the distance of the subsequence in T, at the i^{th} location, to its nearest neighbor under z-normalized Euclidean Distance.

For example, in the below, the subsequence starting at 921 happens to have a distance of 177.0 to its nearest neighbor (wherever it is).

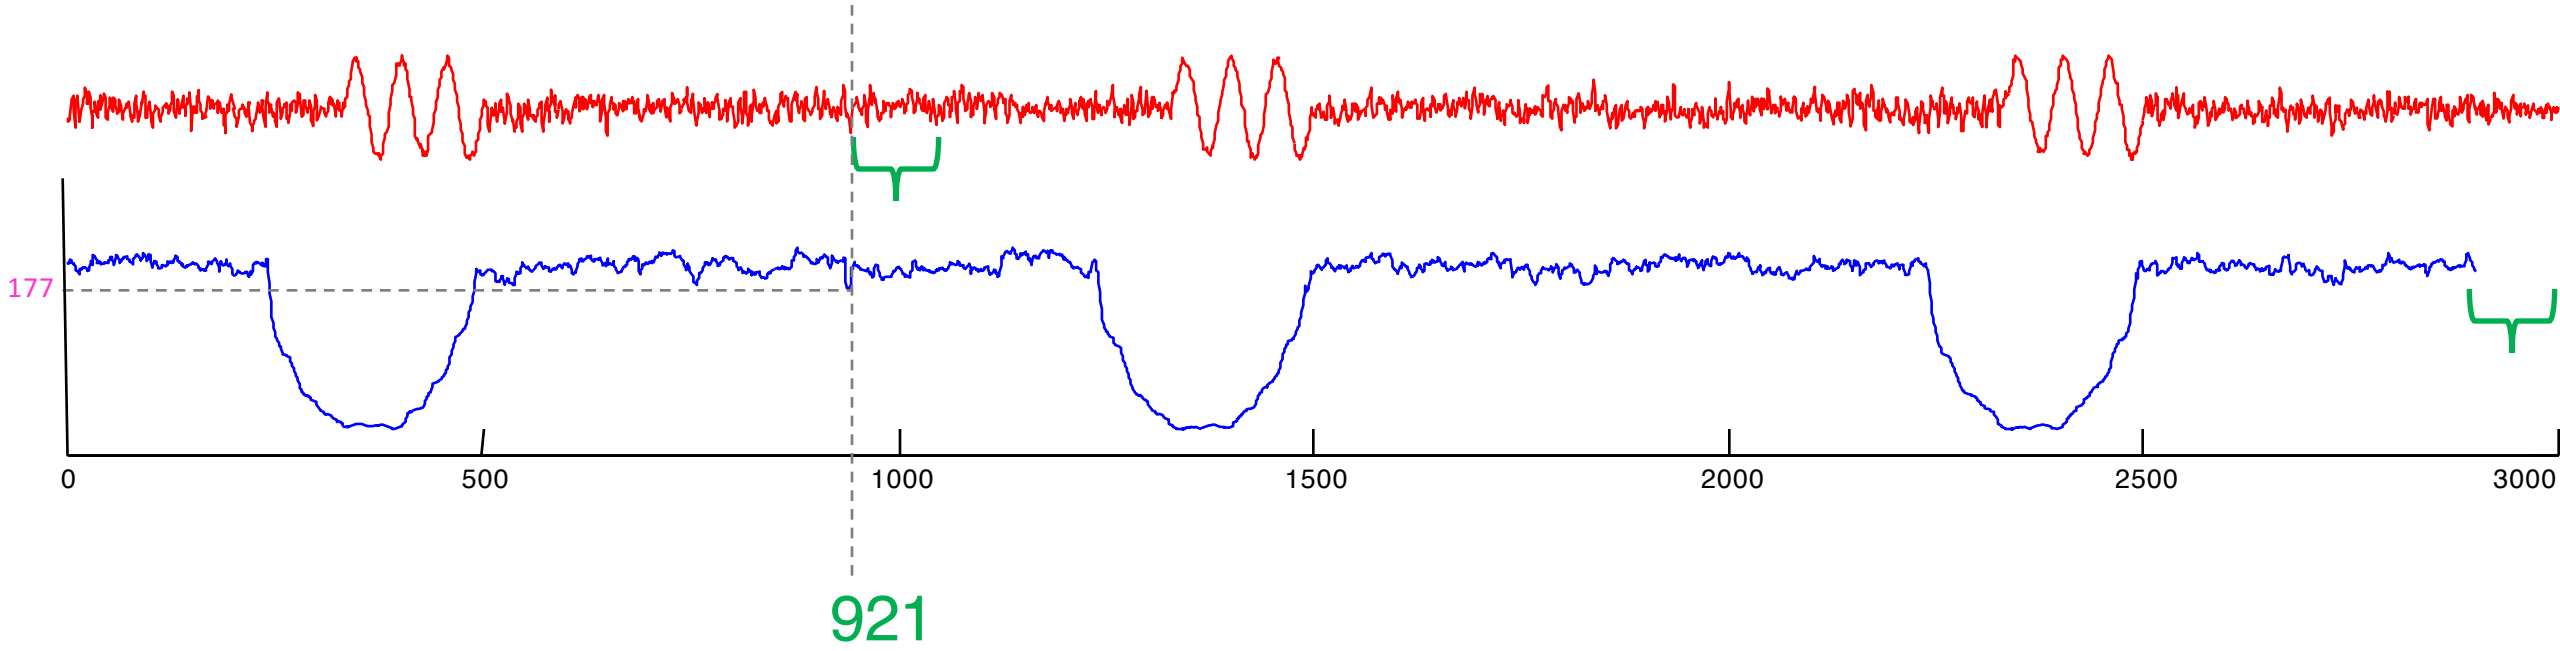

Another example. In the below, the subsequence starting at 378 happens to have a distance of 34.2 to its nearest neighbor (wherever it is).

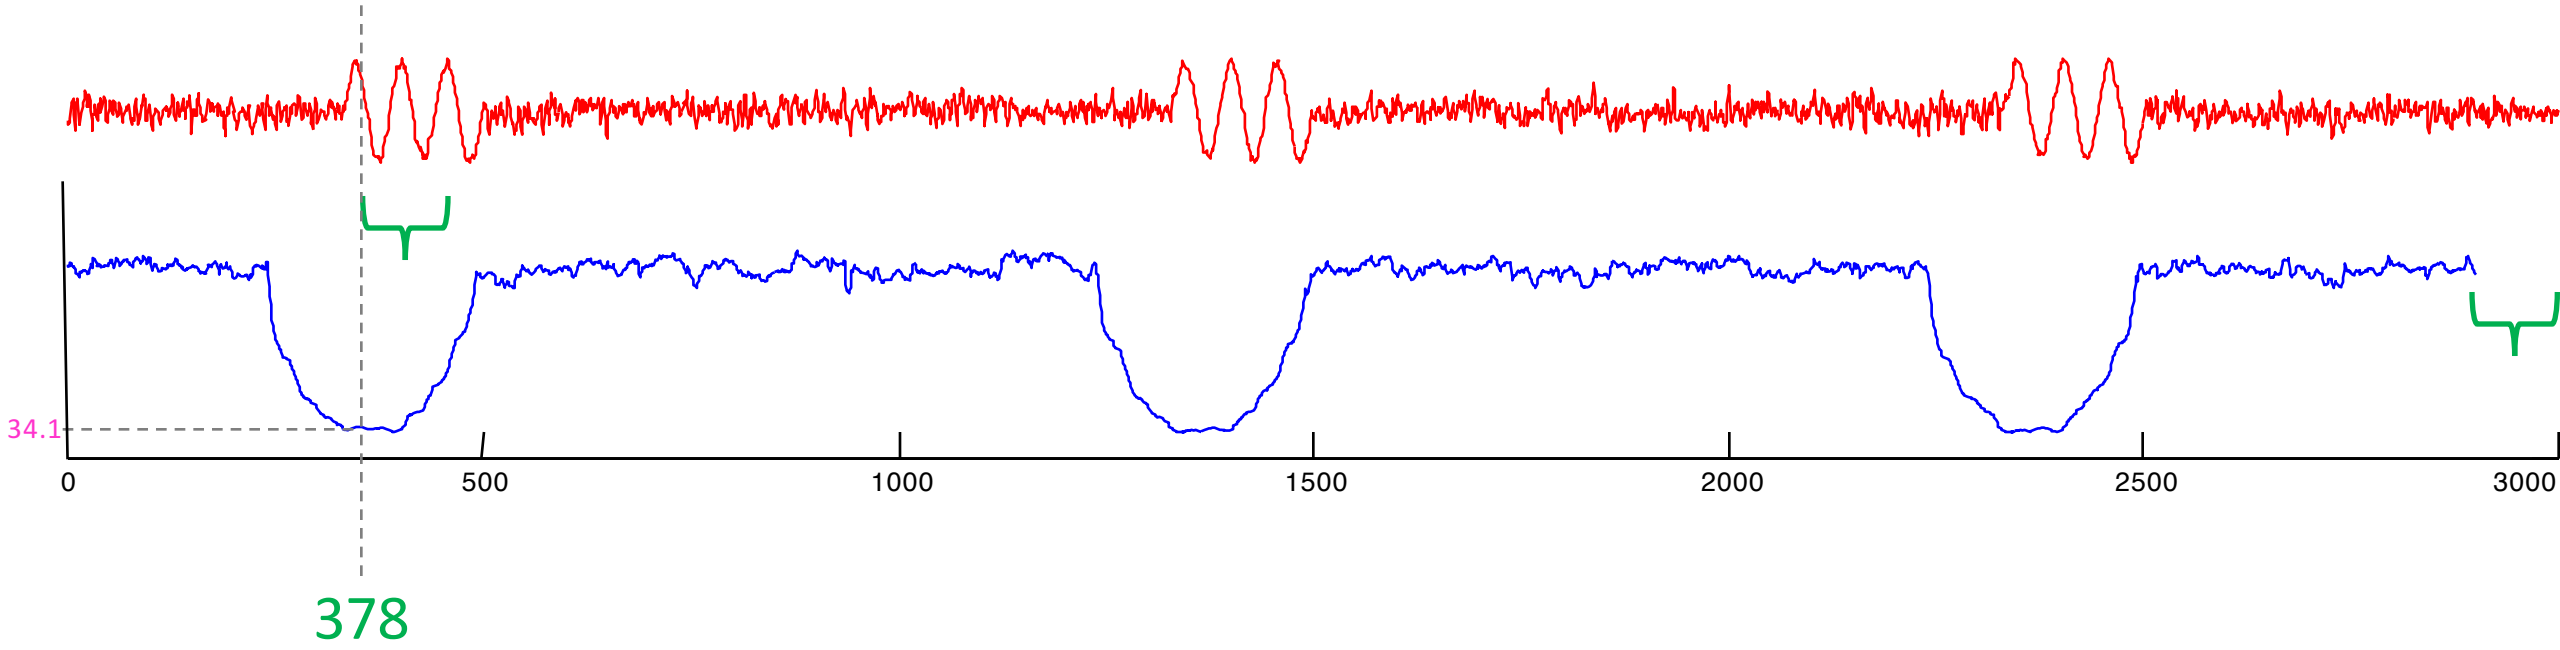

For the rest of this tutorial….

The Matrix Profile is always shown in blue.

The real time series data, is generally shown in red.

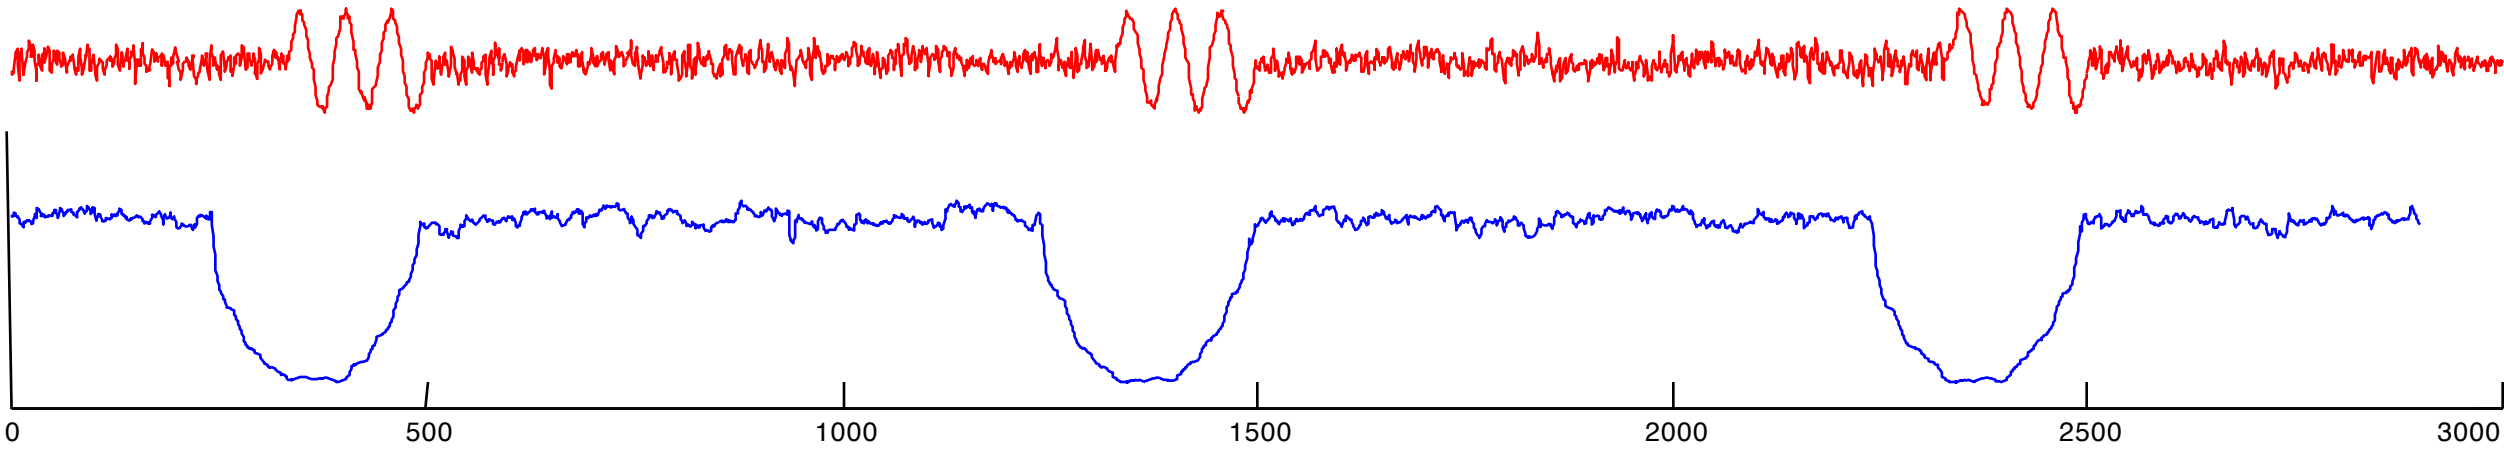

We can create another companion sequence, called a matrix profile index.

The MPI contains integers that are used as pointers. As a practical matter, even 32-bits will let us have a MP of length 2,147,483,647, over two years of data at 60Hz. A 64-bit integer gives us ten billion years at 60Hz)

In the following slides we won't bother to show the matrix profile index, but be aware it exists, and it allows us to find the nearest neighbor to any subsequence in constant time.

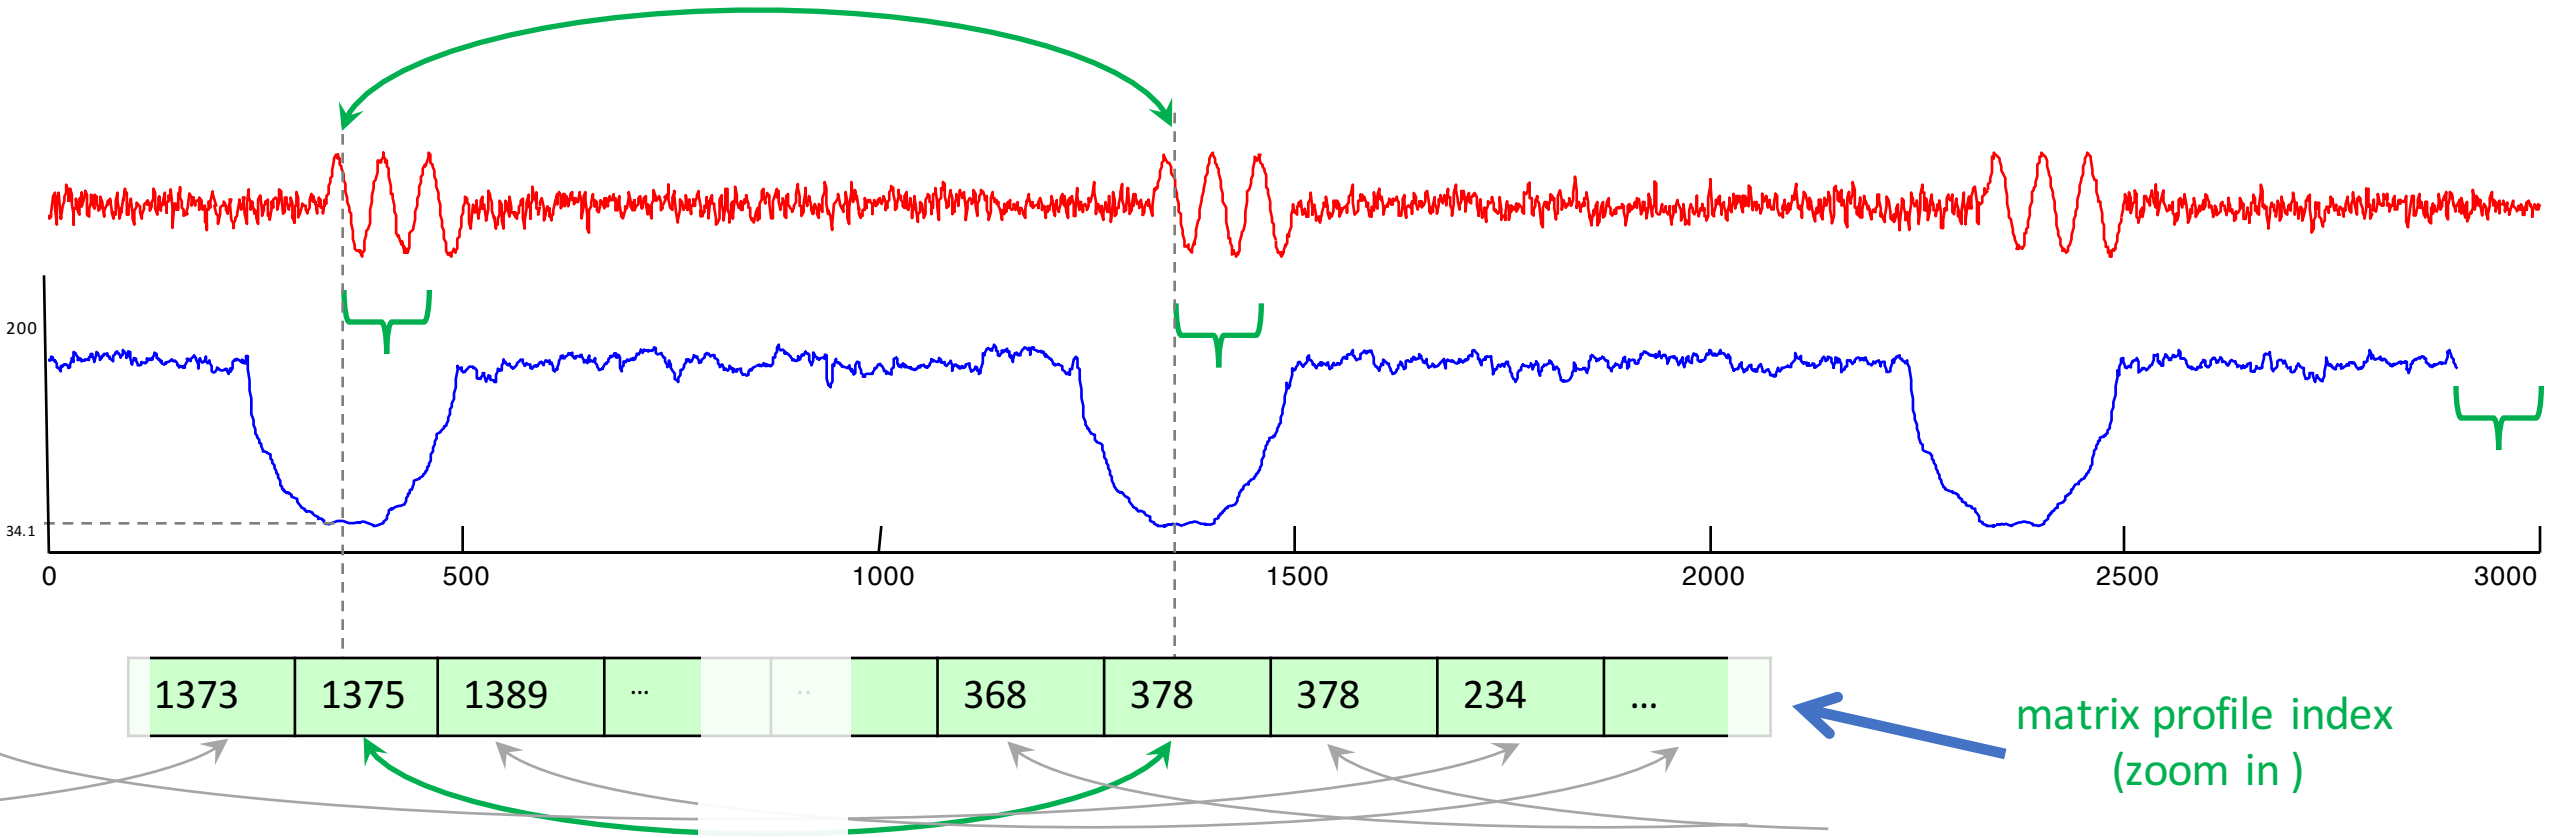

Note that the pointers in the matrix profile index are not necessarily symmetric.

If **A** points to **B**, then **B** may or may not point to **A**

An interesting exception, the two smallest values in the MP must have the same value, and their pointers must be mutual. This is the classic *time series motif*.

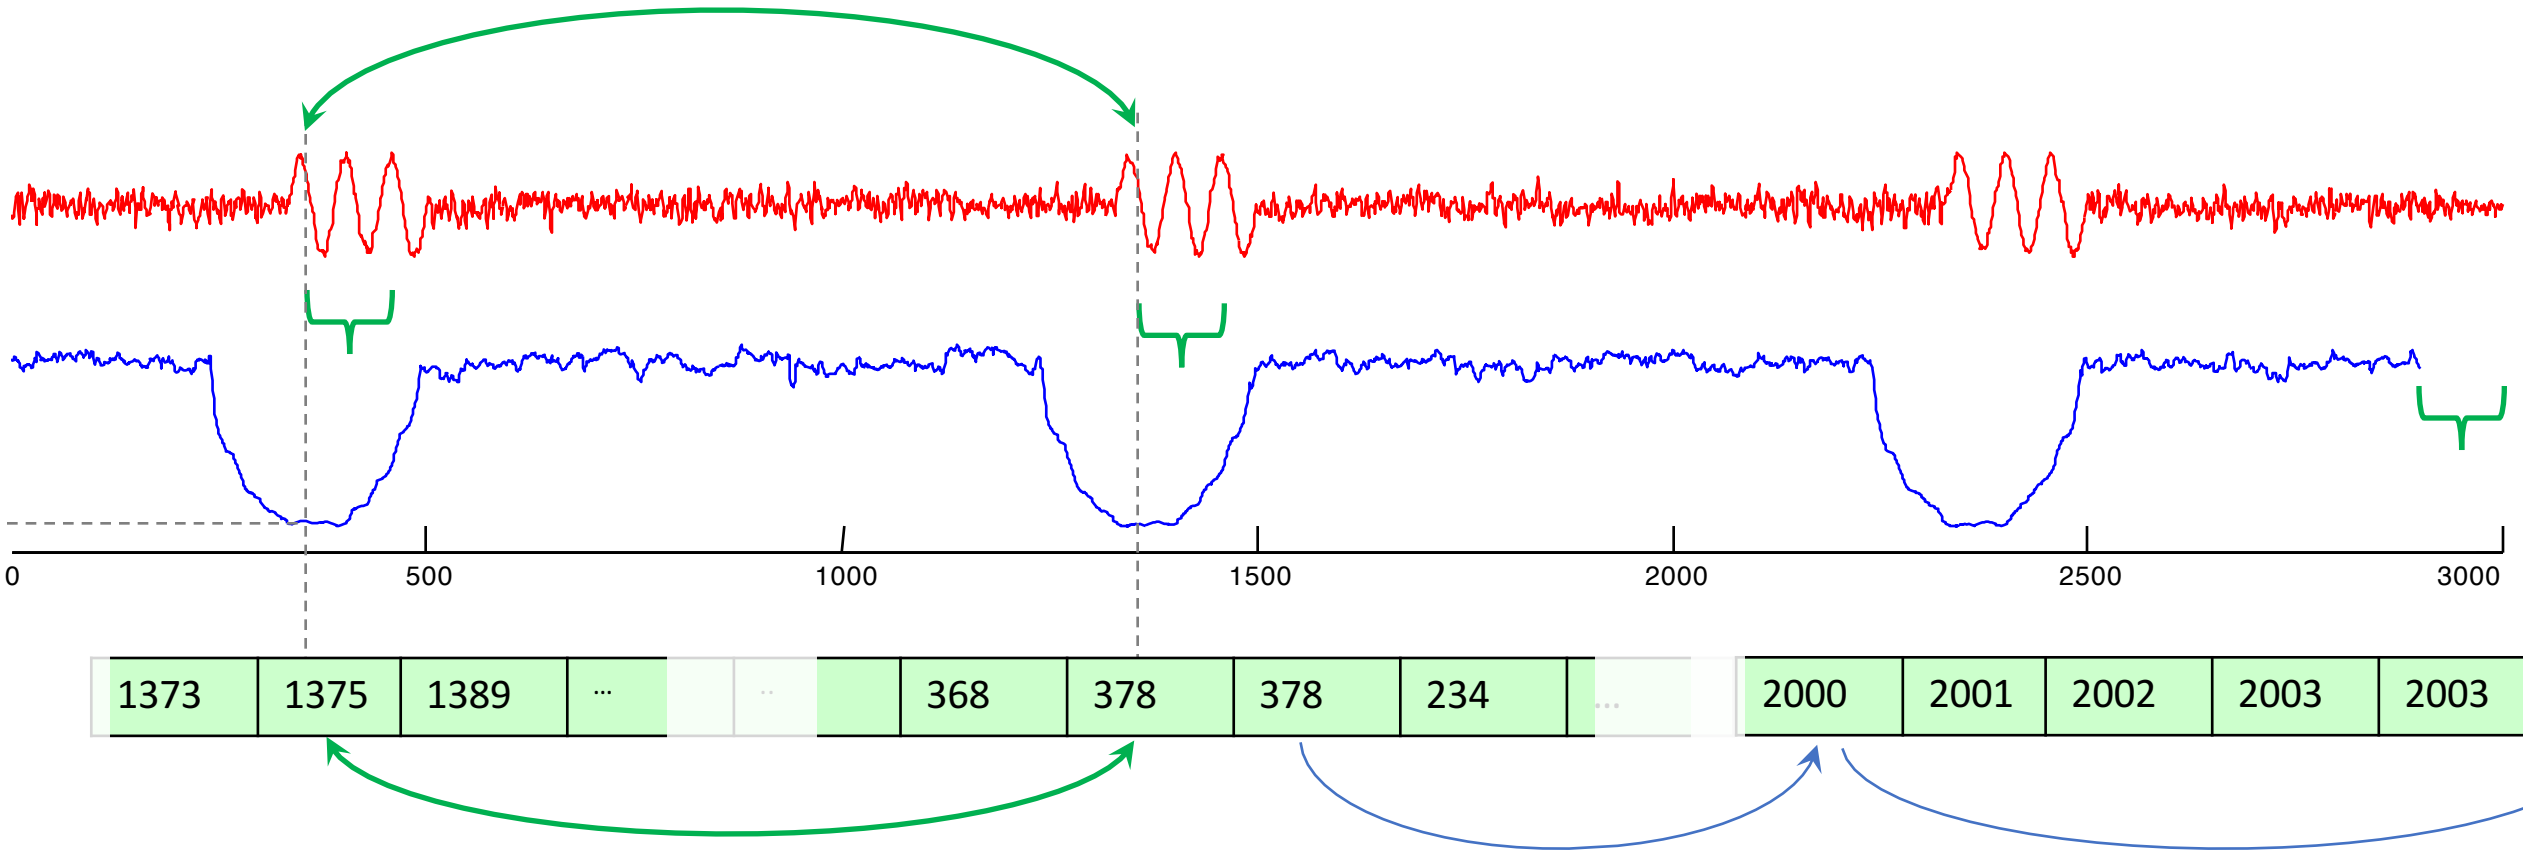

Why is it called the Matrix Profile?

One naïve way to compute it would be to construct a distance matrix of all pairs of subsequences of length *m*.

For each column, we could then "project" down the smallest (*non diagonal*) value to a vector, and that vector would be the Matrix Profile.

While in general we could never afford the memory to do this (4TB for just |T|= one million), for most applications the Matrix Profile is the *only* thing we need from the full matrix, and we can compute and store it very efficiently. (as we will see later)

Key:
Small distances are blue
Large distances are red
Dark stripe is excluded

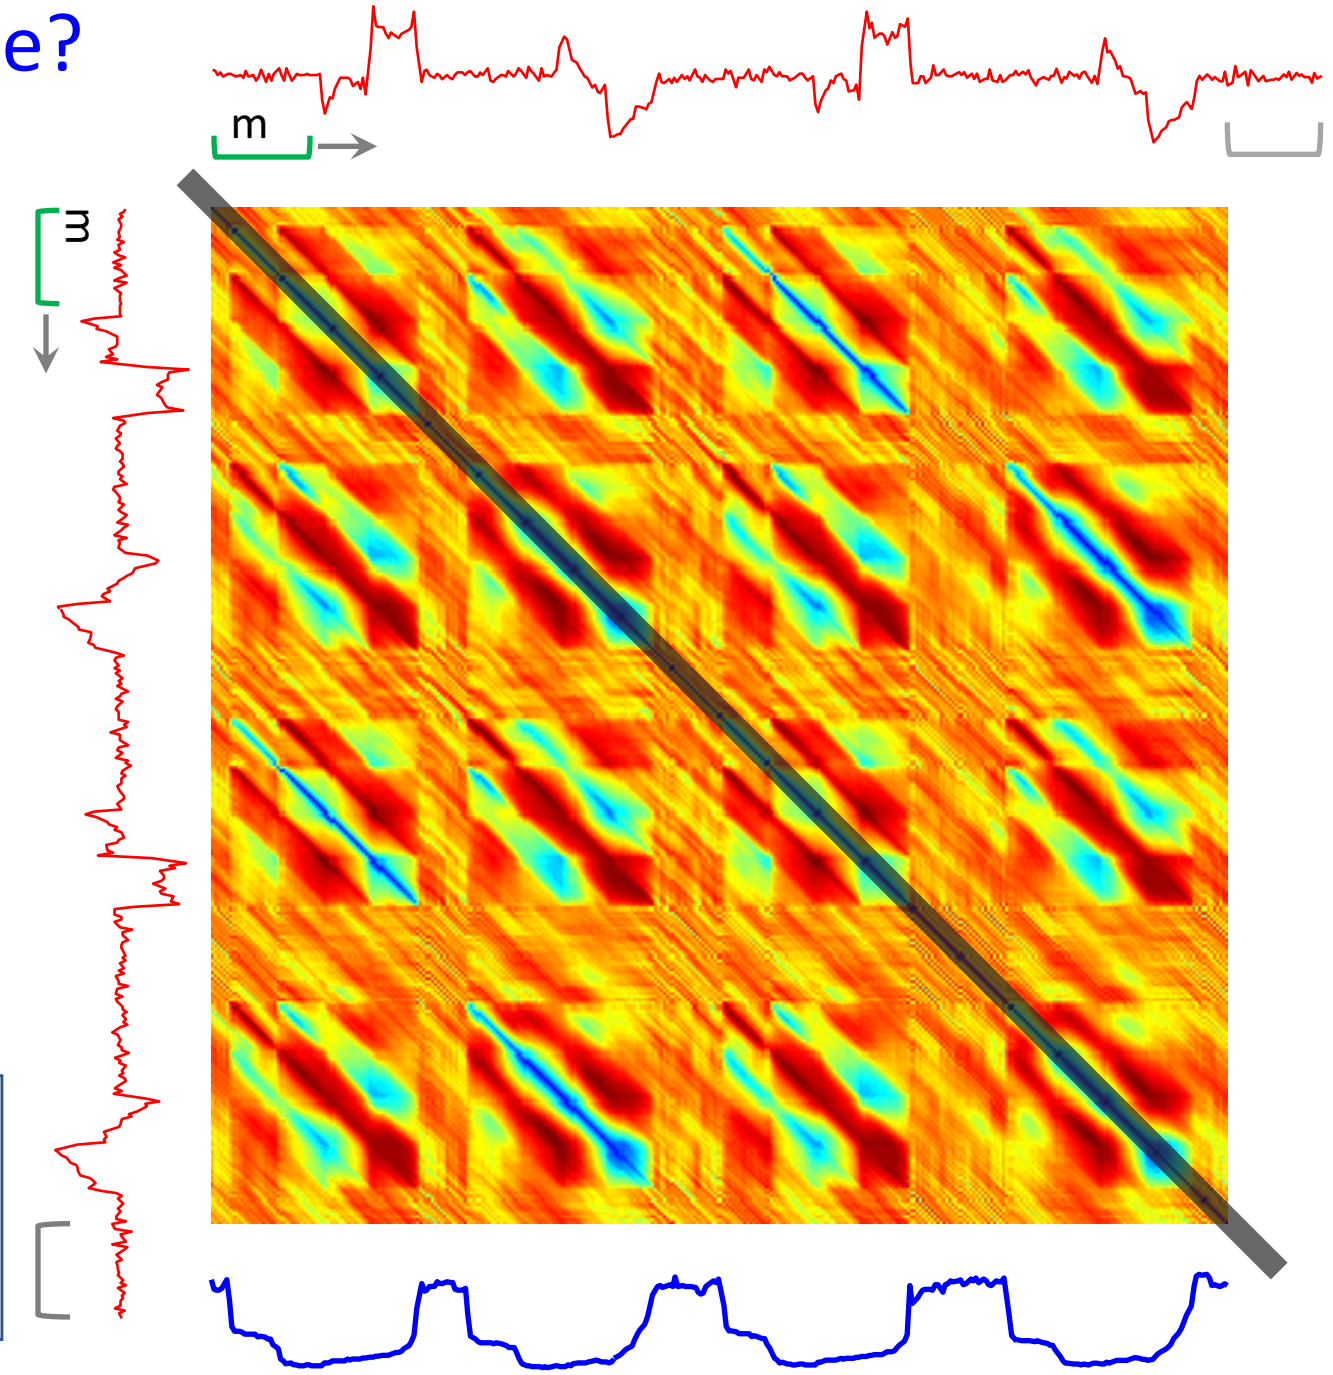

How to “read” a Matrix Profile

Where you see relatively low values, you know that the subsequence in the original time series must have (at least one) relatively similar subsequence elsewhere in the data (such regions are “motifs” or reoccurring patterns)

Where you see relatively high values, you know that the subsequence in the original time series must be unique in its shape (such areas are “discords” or anomalies).

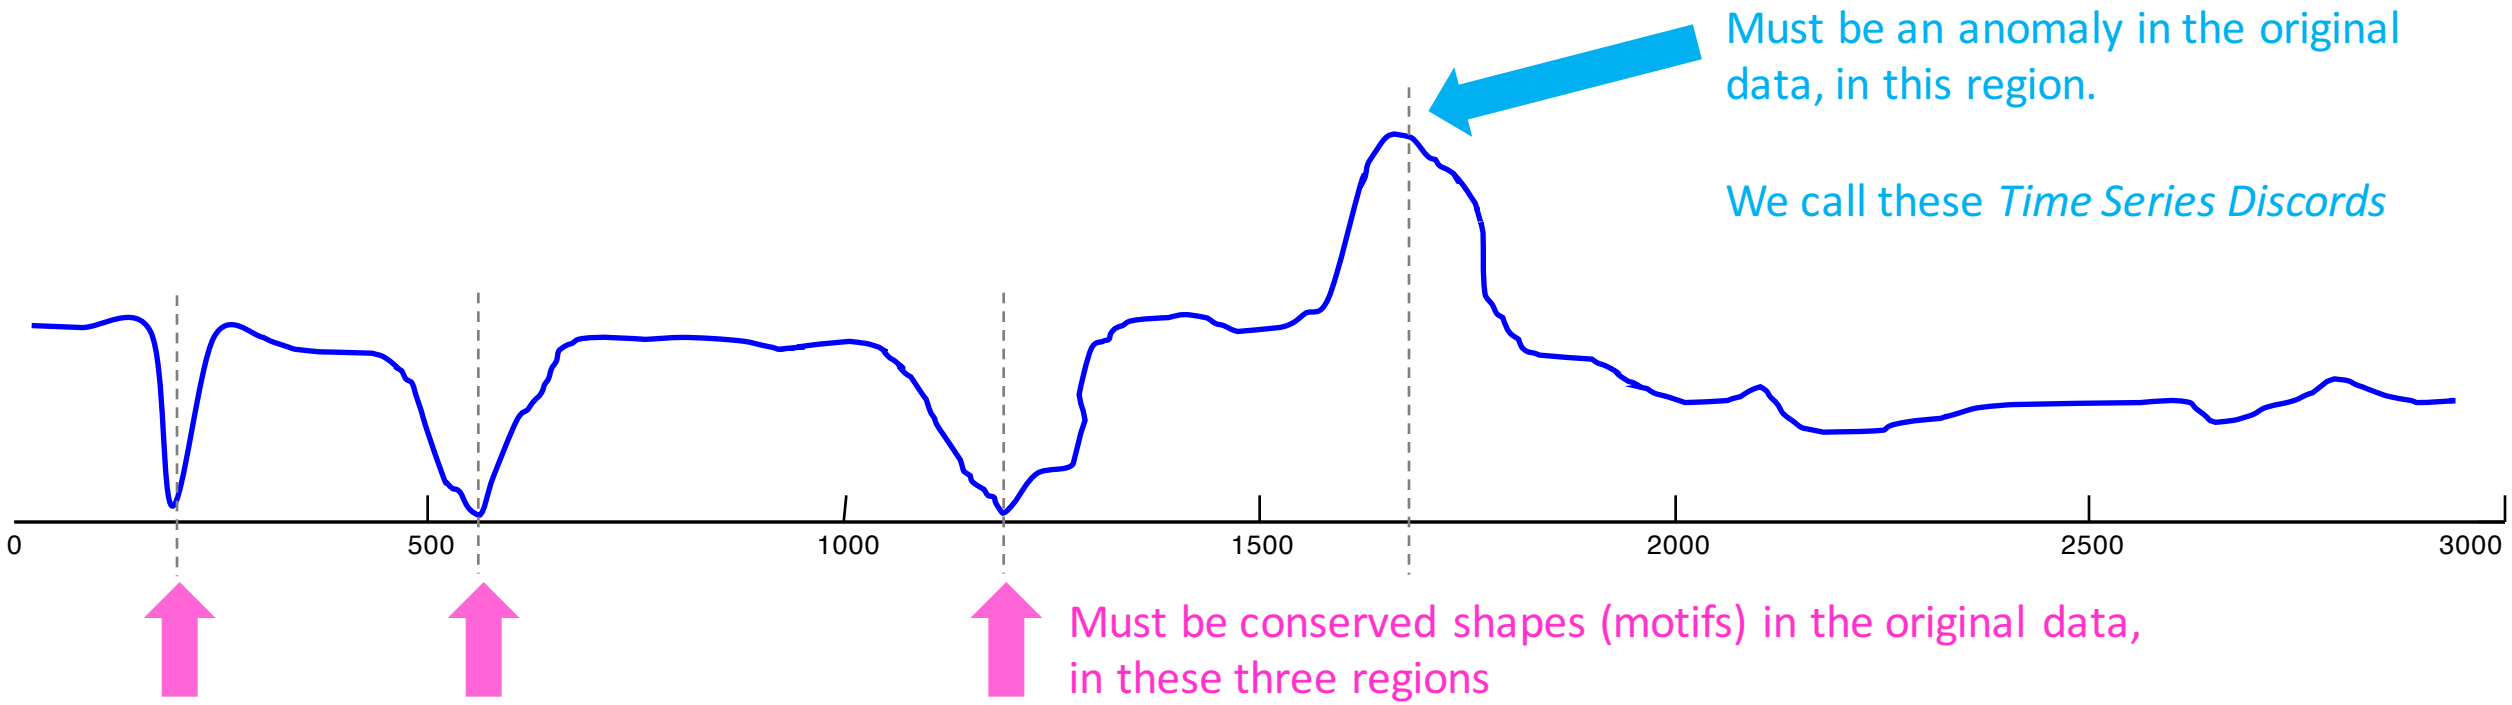

How to “read” a Matrix Profile: Synthetic Anomaly Example

Where you see relatively high values, you know that the subsequence in the original time series must be unique in its shape. In fact, the highest point is *exactly* the definition of Time Series Discord, perhaps the best anomaly detector for time series*

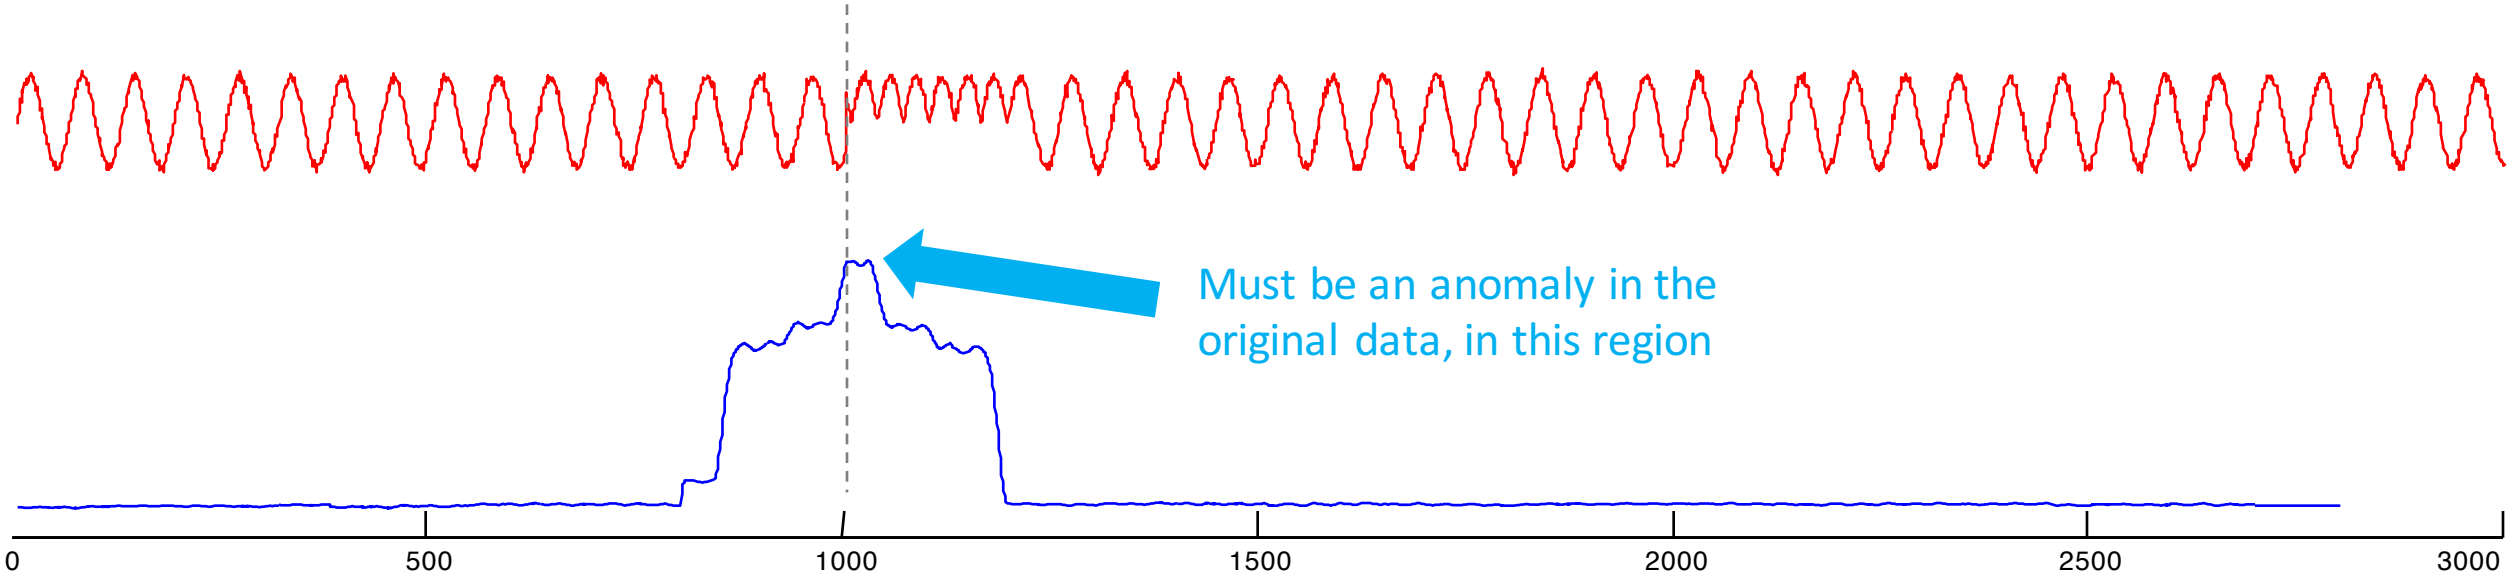

* Vipin Kumar performed an extensive empirical evaluation and noted that “..on 19 different publicly available data sets, comparing 9 different techniques (time series discords) is the best overall technique.”. V. Chandola, D. Cheboli, V. Kumar. Detecting Anomalies in a Time Series Database. UMN TR09-004

How to "read" a Matrix Profile: Synthetic Motif Example

Where you see relatively low values, you know that the subsequence in the original time series must have (at least one) relatively similar subsequence elsewhere in the data.

In fact, the lowest points must be a tieing pair, and correspond exactly to the classic definition of *time series motifs*.

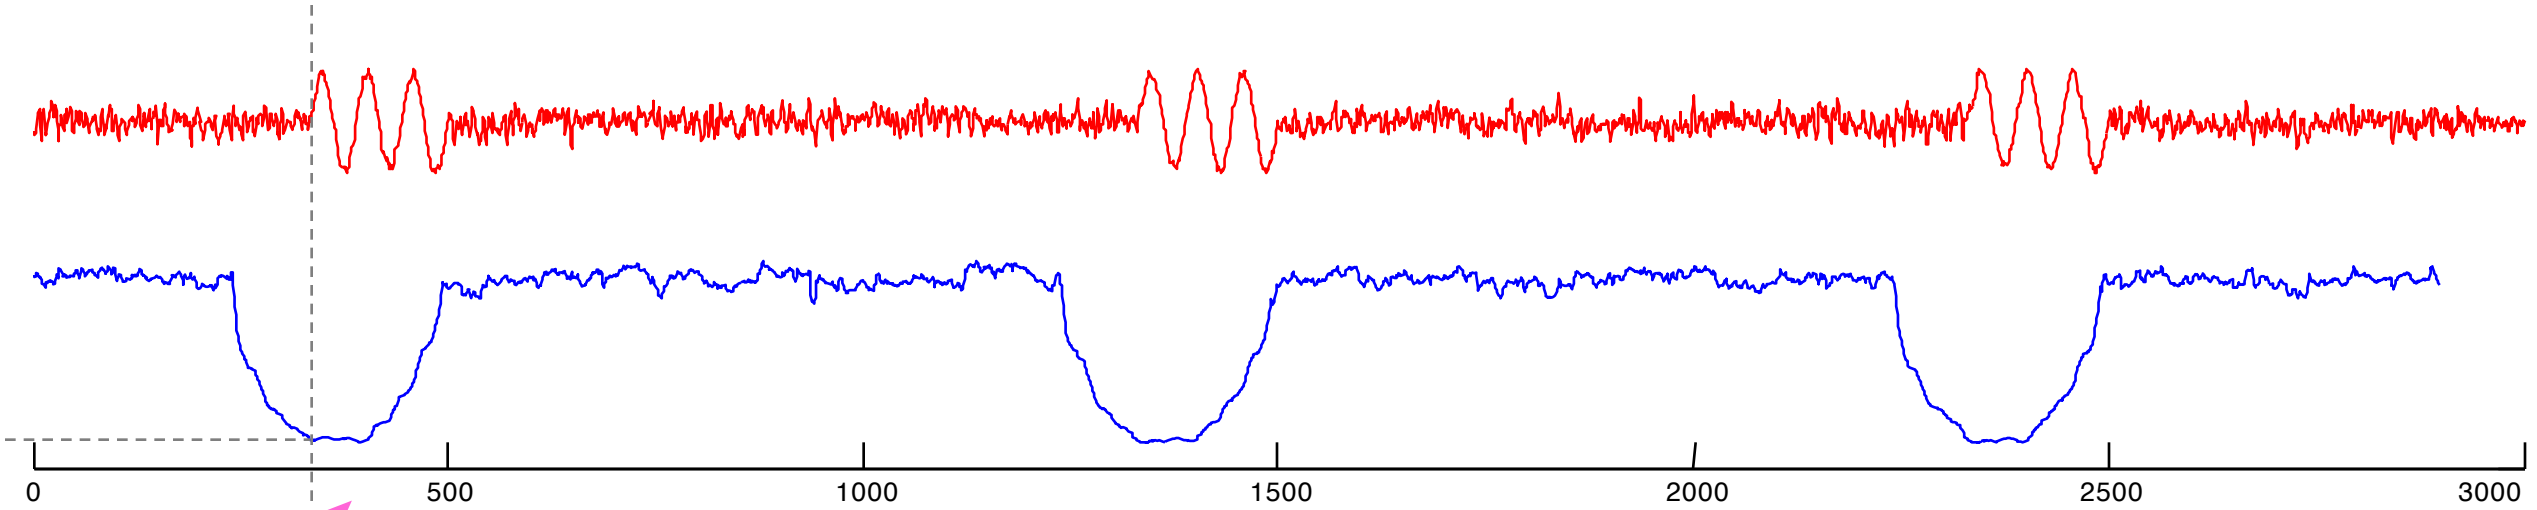

The corresponding subsequence in the raw data at this location, must have a*t least one* similar subsequence somewhere

How to “read” a Matrix Profile:

Now that we understand what a Matrix Profile is, and we have some practice interpreting them on *synthetic* data, let us spend the next five minutes to see some examples on *real* data.

Note that we will typically create algorithms that use the Matrix Profile, without actually having humans *look* at it.

Nevertheless, in many exploratory time series data mining tasks, just looking at the Matrix Profile can give us unexpected and actionable insights.

Ready to begin?

Taxi Example: Part I

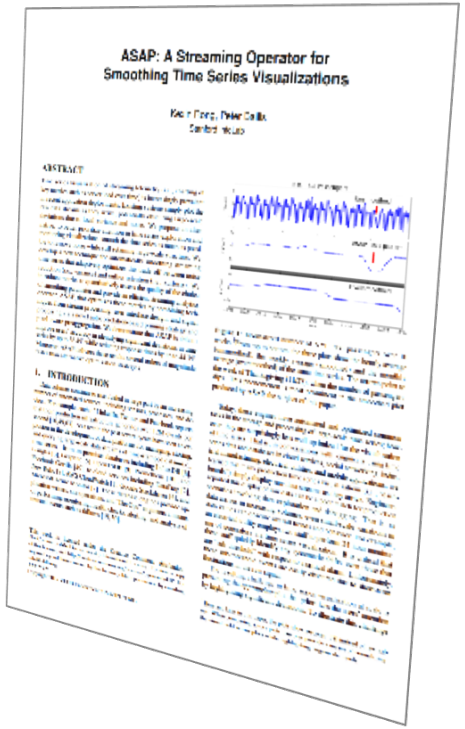

Given a long time series, where should you examine carefully?
The problem is called "Attention Prioritization", a group at Stanford is working on this [a].
However we think that the Matrix Profile can be used for this, "for free".

Below is the data, the hourly average of the number of NYC taxi passengers over 75 days in Fall of 2014.

Lets compute the Matrix Profile for it, we choose a subsequence length corresponding to two days…. (next slide)

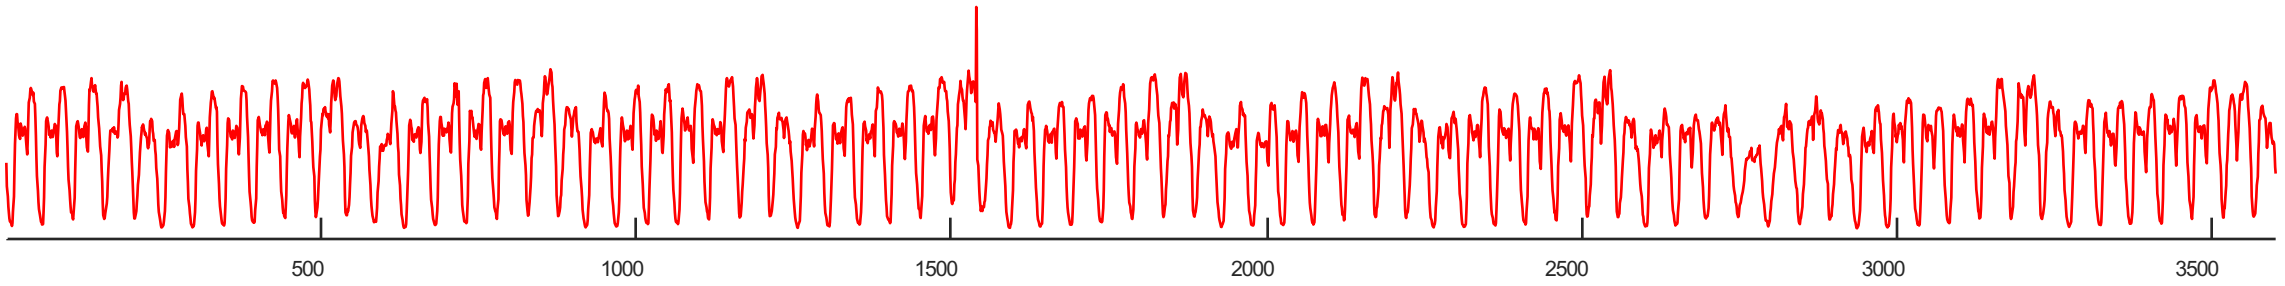

[a] http://futuredata.stanford.edu/ASAP/extended.pdf

Taxi Example: Part II

- The highest value corresponds to Thanksgiving (the uniqueness of Thanksgiving was the only thing the Stanford Team noted)
- We find a secondary peak around Nov 6th, what could it be? Daylight Saving Time! The clock going backwards one hour, gives an *apparent* doubling of taxi load.
- We find a tertiary peak around Oct 13th, what could it be? Columbus Day! Columbus Day is largely ignored in much of America, but still a big deal in NY, with its large Italian American community.

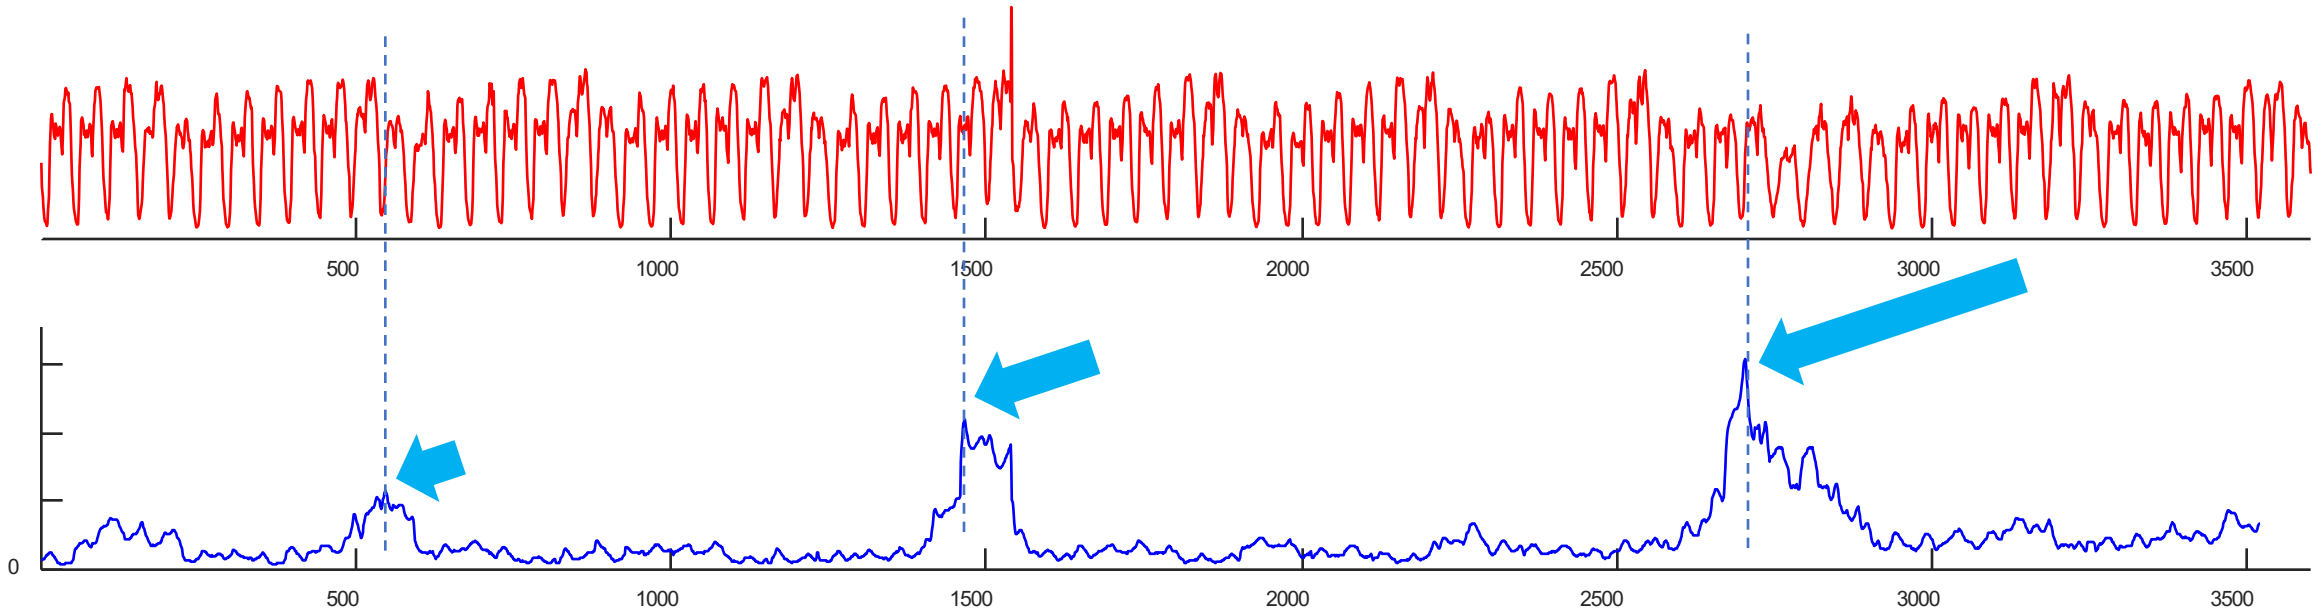

Taxi Example: Part III

- What about the lowest values? (the best motifs)
- They are *exactly* seven days apart, suggesting that in this dataset, there might be a periodicity of *seven* days, in addition to the more obvious periodicity of *one* day.

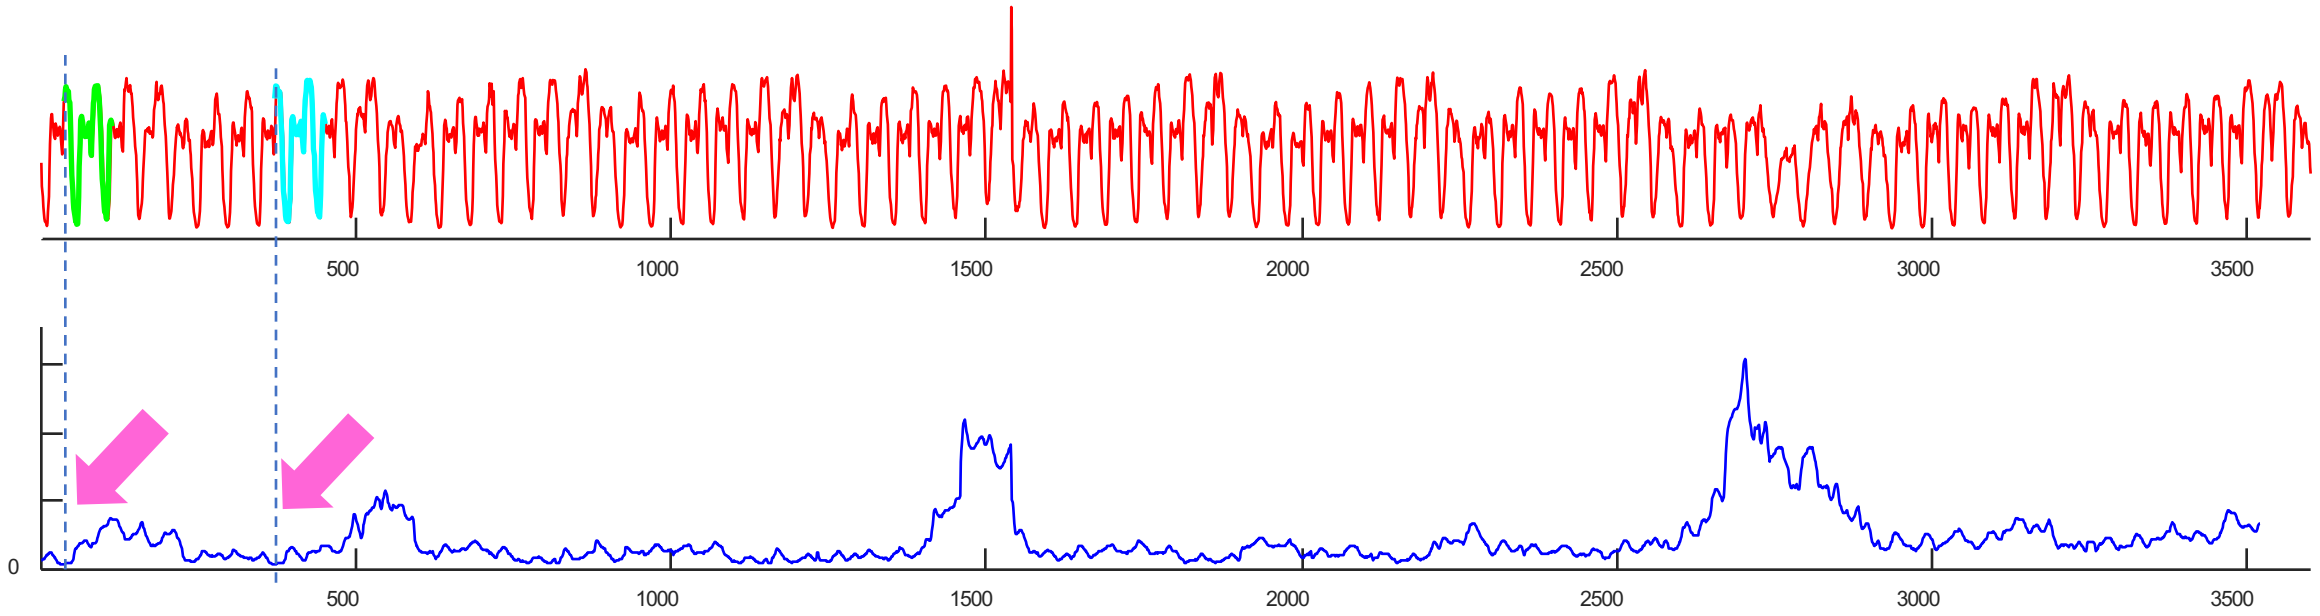

Italy Power Demand (1995 to 1998)

The top motif is a typical work week, starting from Tuesday

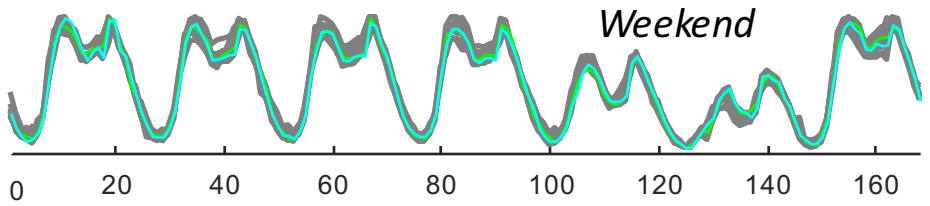

The Taxi example was easy to solve by manual inspection of the *raw* data, but with just an order of magnitude more data, the problem becomes much harder. Lets try a similar, but larger example, Italian Power Demand 1995 to 1998.

Note that the matrix profile is very low *on average*, most weeks are similar to the previous week (*persistence*) or the same week in a different year (*history*).

All the high values can be explained by Italian holidays, most of which fall on different days in consecutive years.

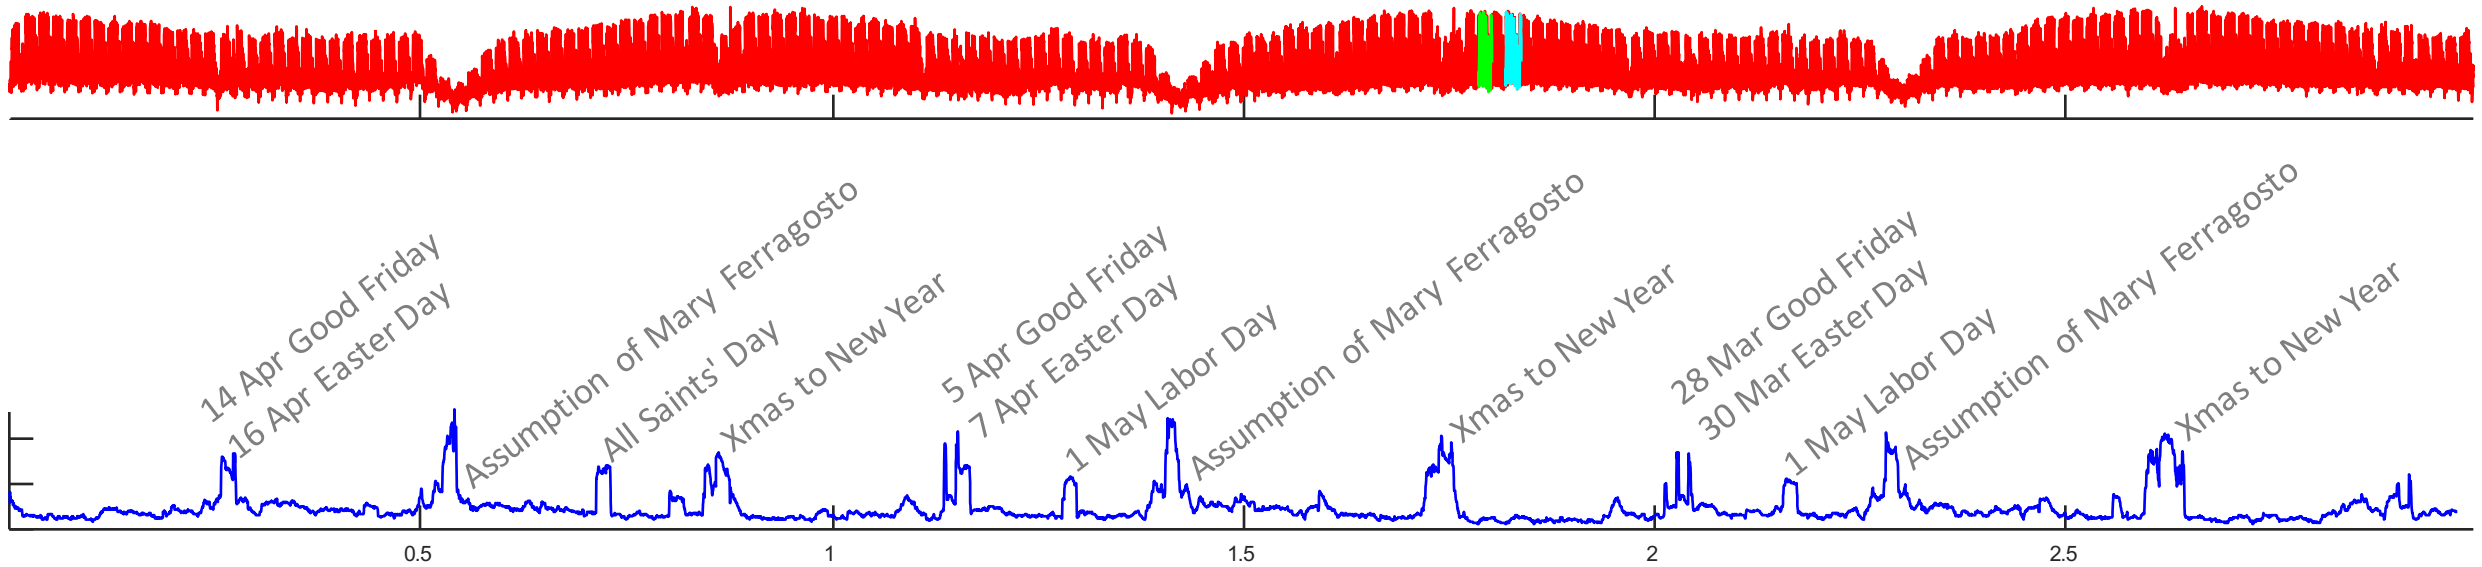

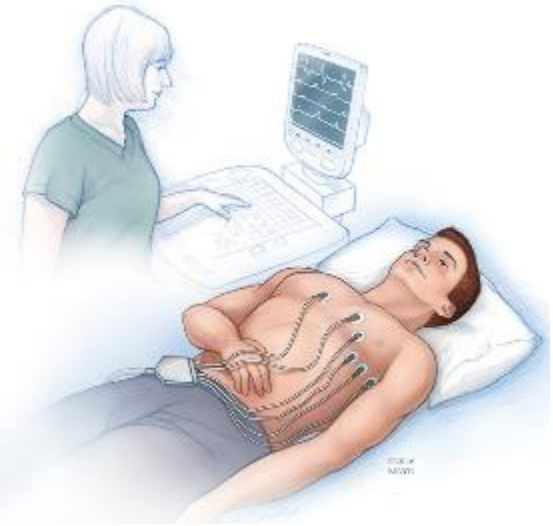

Electrocardiogram

(MIT-BIH Long-Term ECG Database)

In this case there are two anomalies annotated by MIT cardiologists. The Matrix Profile clearly indicates them.

Here the subsequence length was set to 150, but we still find these anomalies if we *half* or *triple* that length.

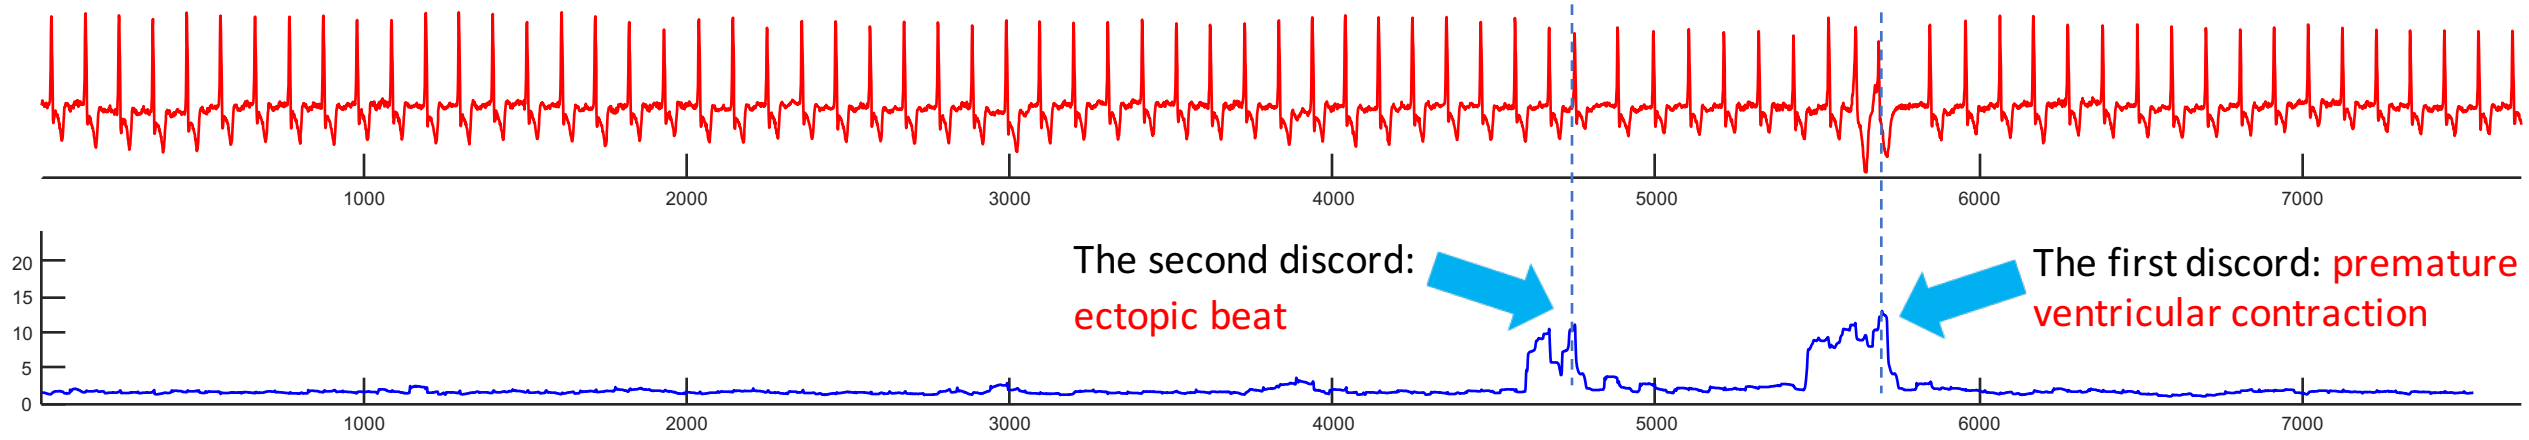

Zebra Finch

(Zebra Finch Vocalizations in MFCC, 100 day old male)

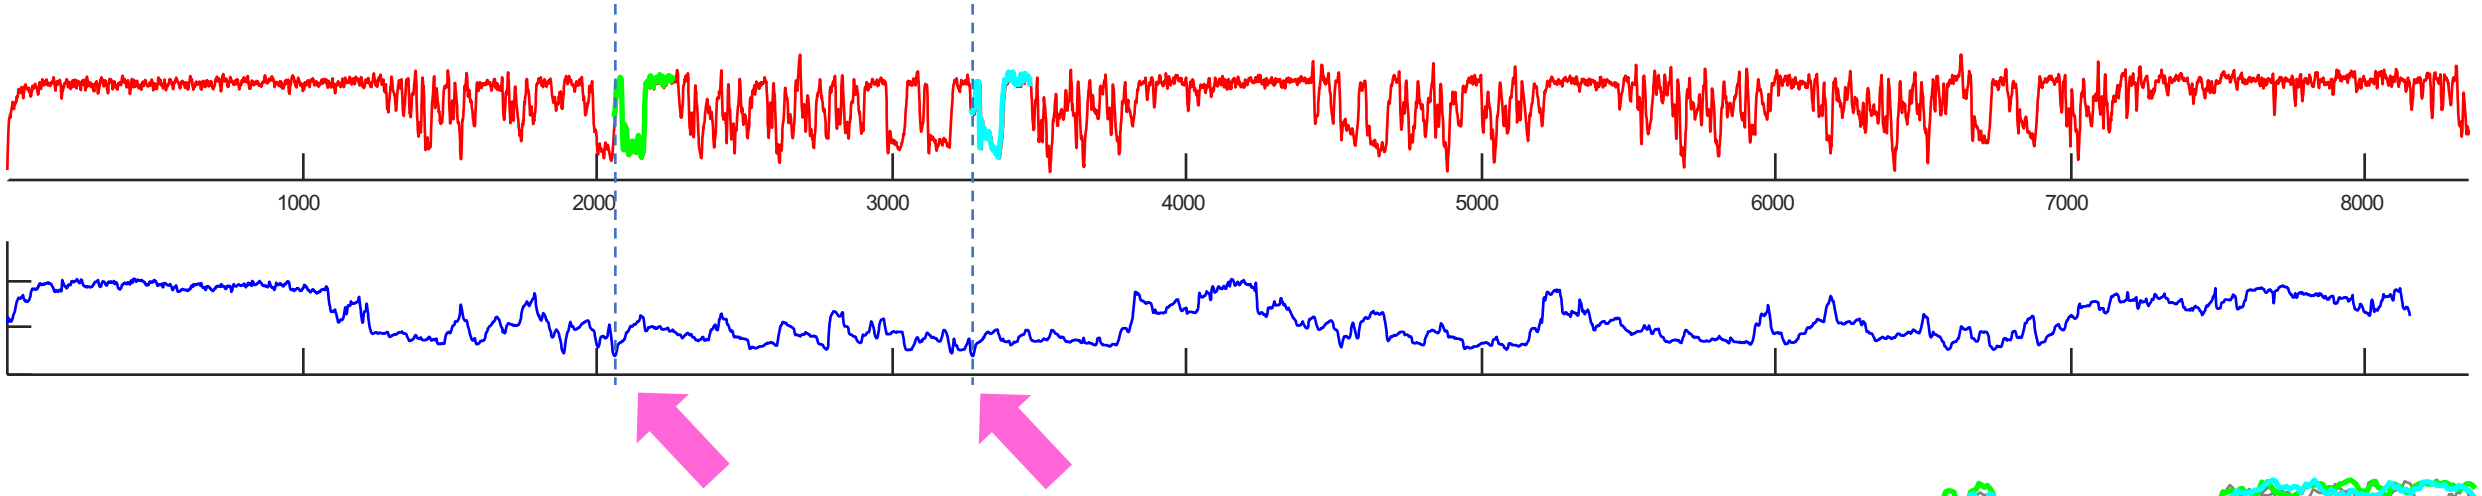

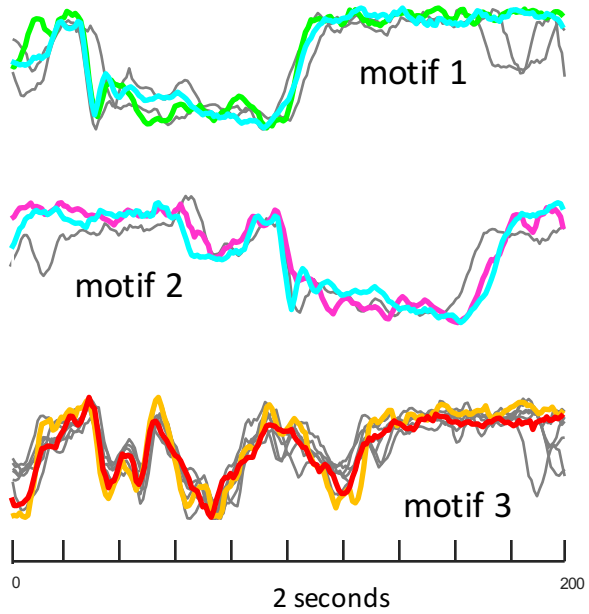

Motif discovery can often surprise you.

While it is clear that this time series is not random, we did not expect the motifs to be so well conserved or repeated so many times. There is evidence of a *vocabulary*, and maybe even a *grammar*…

Seismology

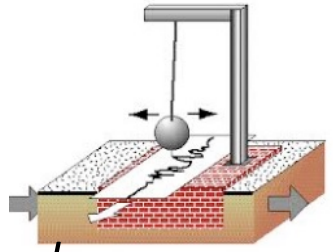

If we see low values in the MP of a seismograph, it means there must have been a *repeated earthquake*.
Repeated earthquakes can happen *decades* apart.
Many fundamental problems seismology, including the discovery of foreshocks, aftershocks, triggered earthquakes, swarms, volcanic activity and induced seismicity, can be reduced to the discovery of these repeated patterns.

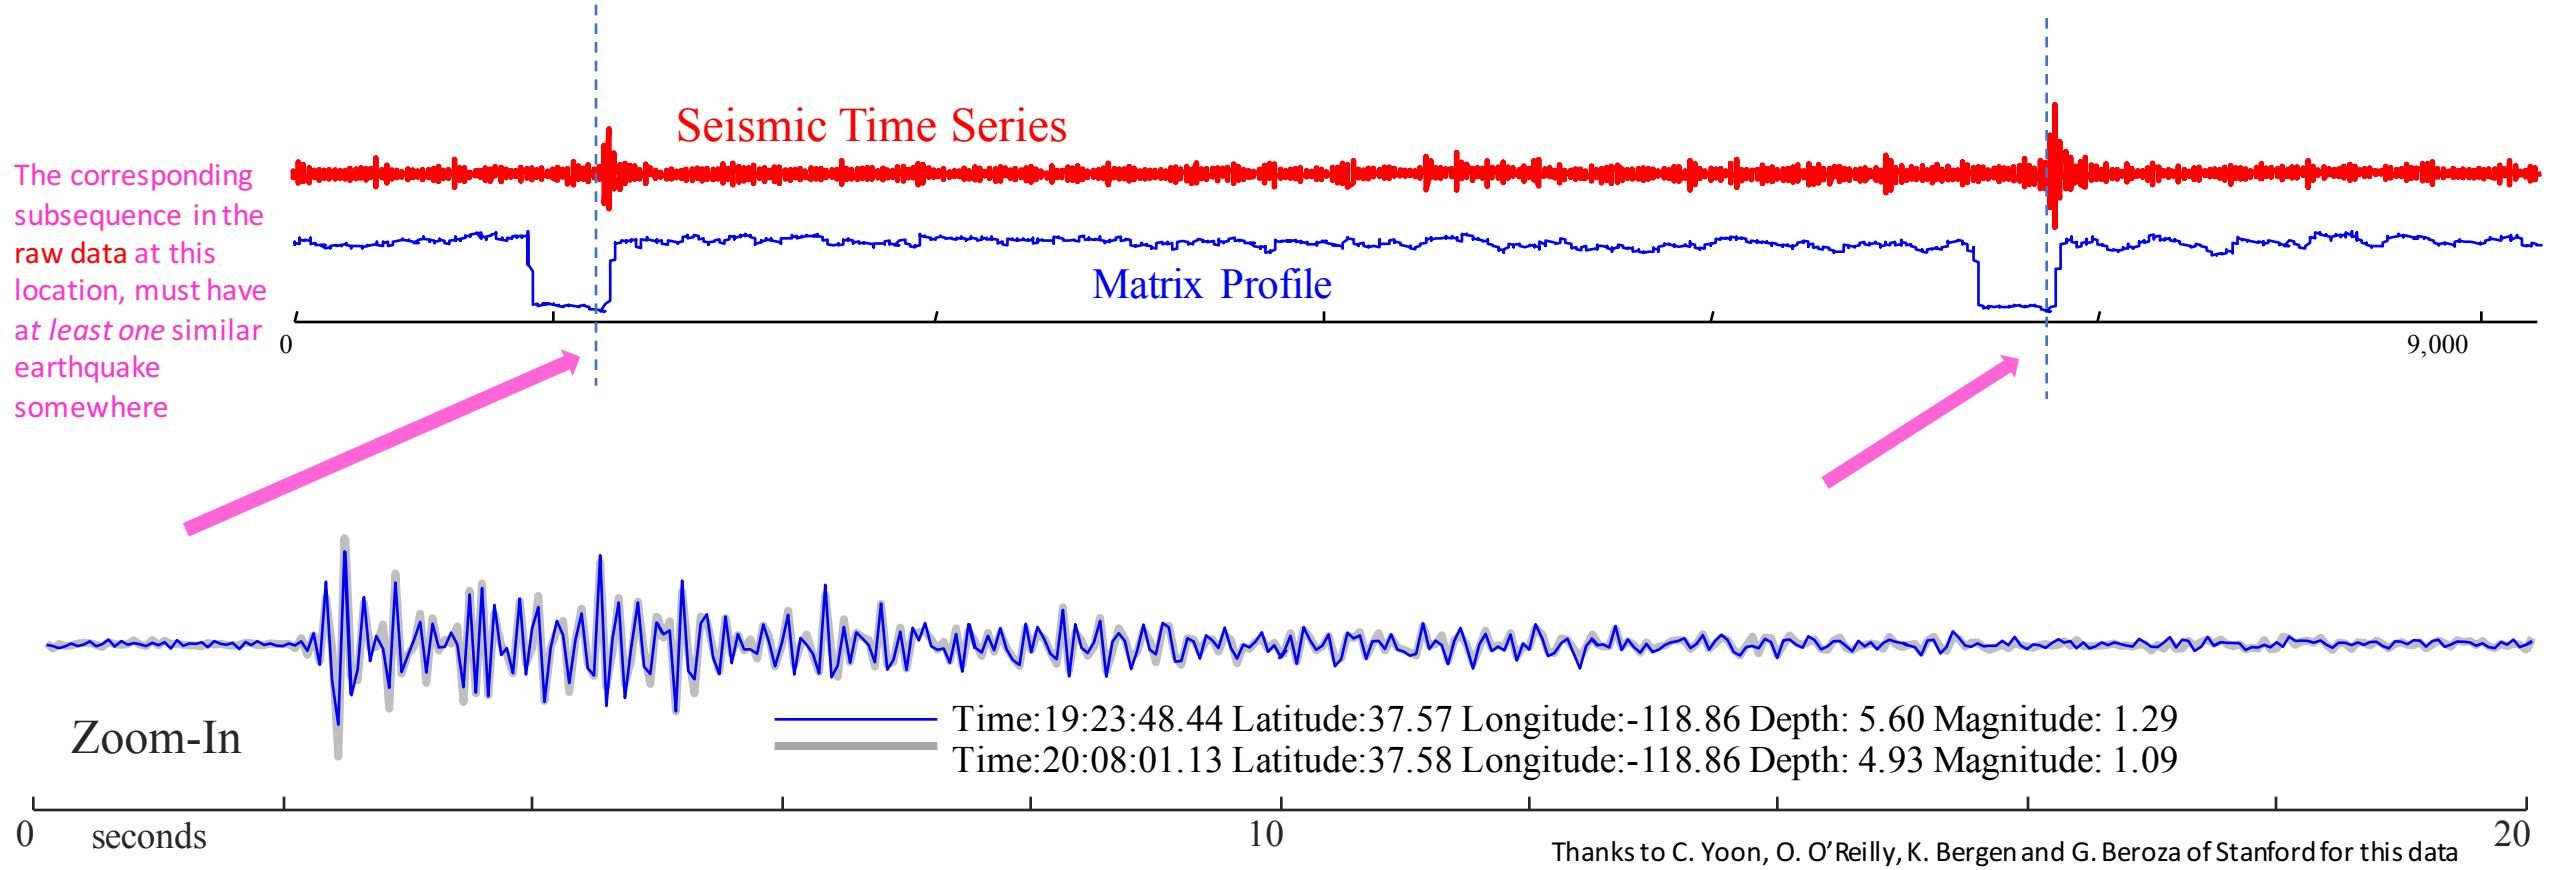

Thanks to C. Yoon, O. O'Reilly, K. Bergen and G. Beroza of Stanford for this data

Chimp DNA

Y-chromosome converted to time series

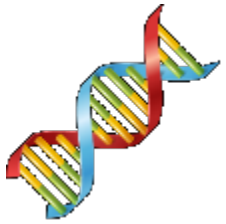

It is possible to convert DNA strings to real-valued time series, in a lossless fashion

```
T1 = 0;
for i = 1 to length(chromosome)
    if chromosome_i = A, then T_i+1 = T_i + 2
    if chromosome_i = G, then T_i+1 = T_i + 1
    if chromosome_i = C, then T_i+1 = T_i - 1
    if chromosome_i = T, then T_i+1 = T_i - 2
end
```

Let us search the Chimp's DNA for repeated structure, of length 60,000…

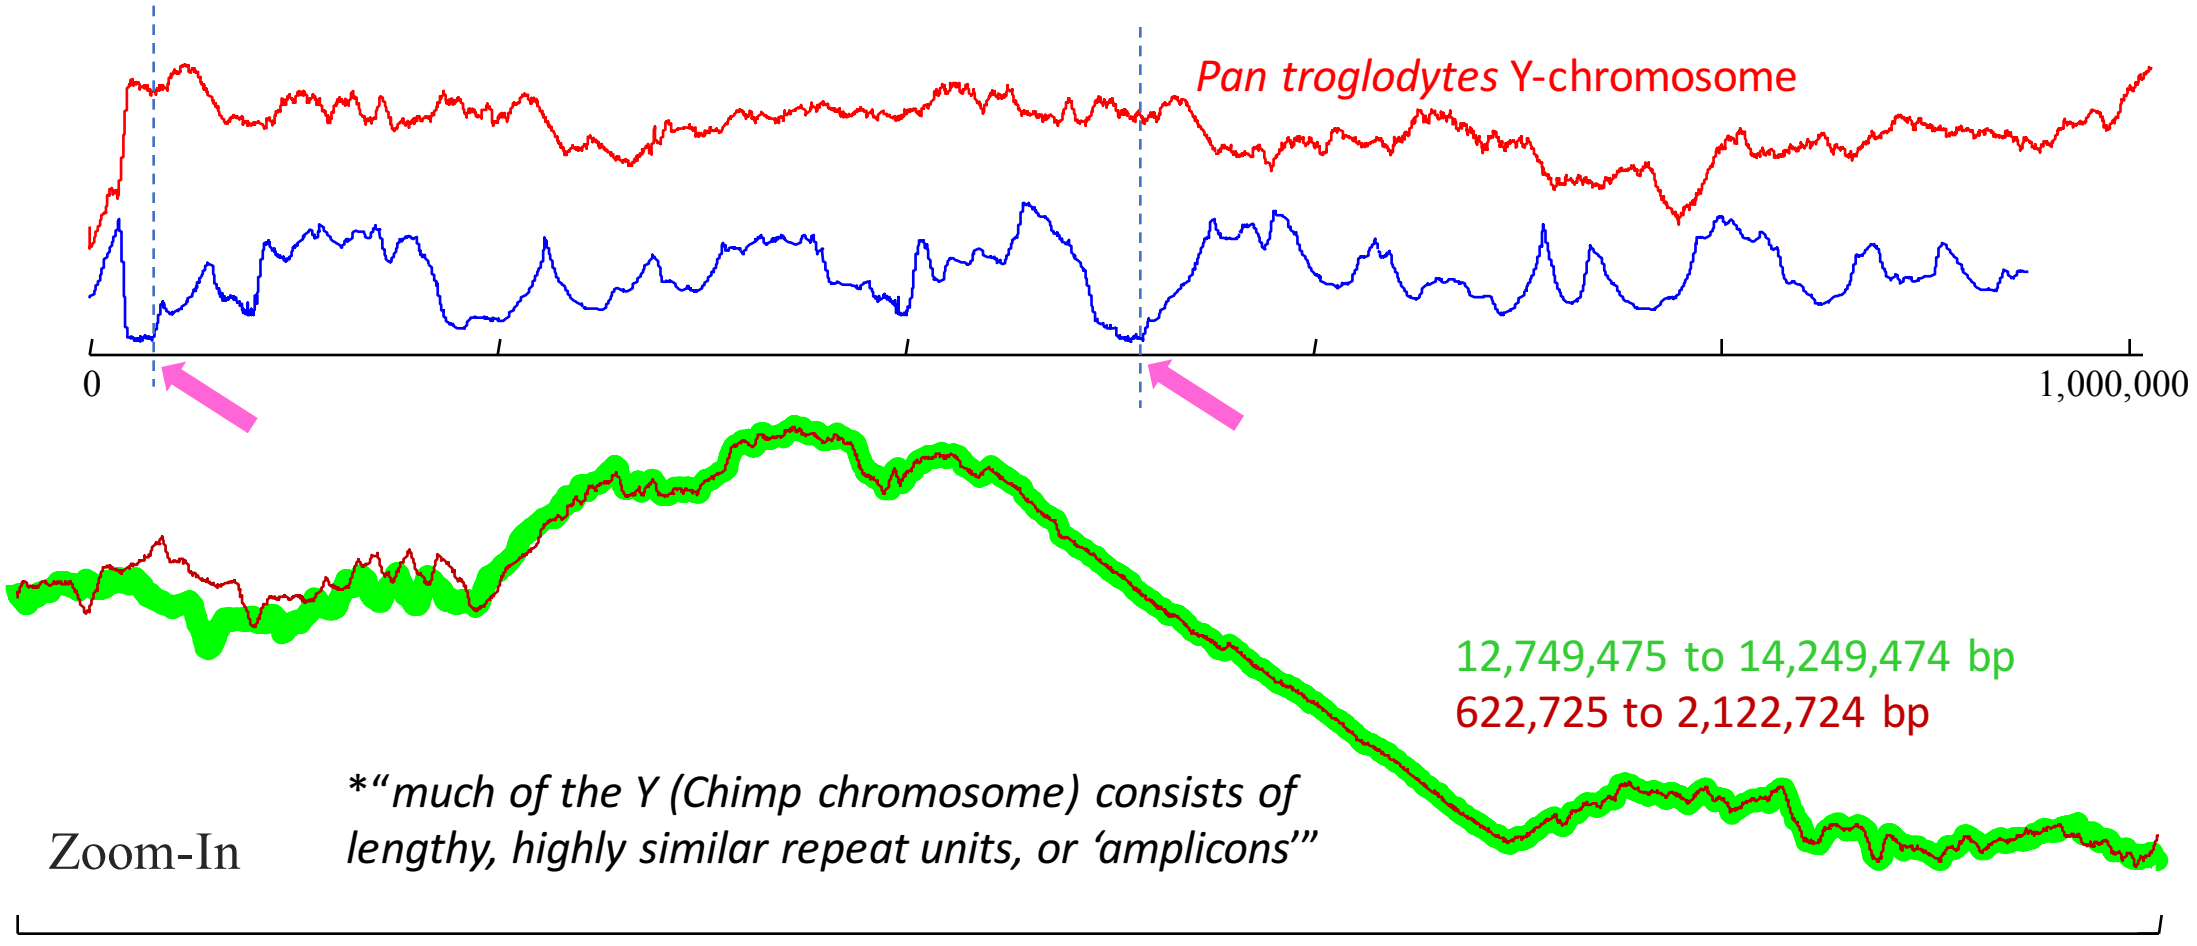

*J. Hughes et al., "Chimpanzee and human Y chromosomes are remarkably divergent in structure" Nature 463, (2010).

Music

While *motifs* usually point to chorus,
discords point to bridges or solos

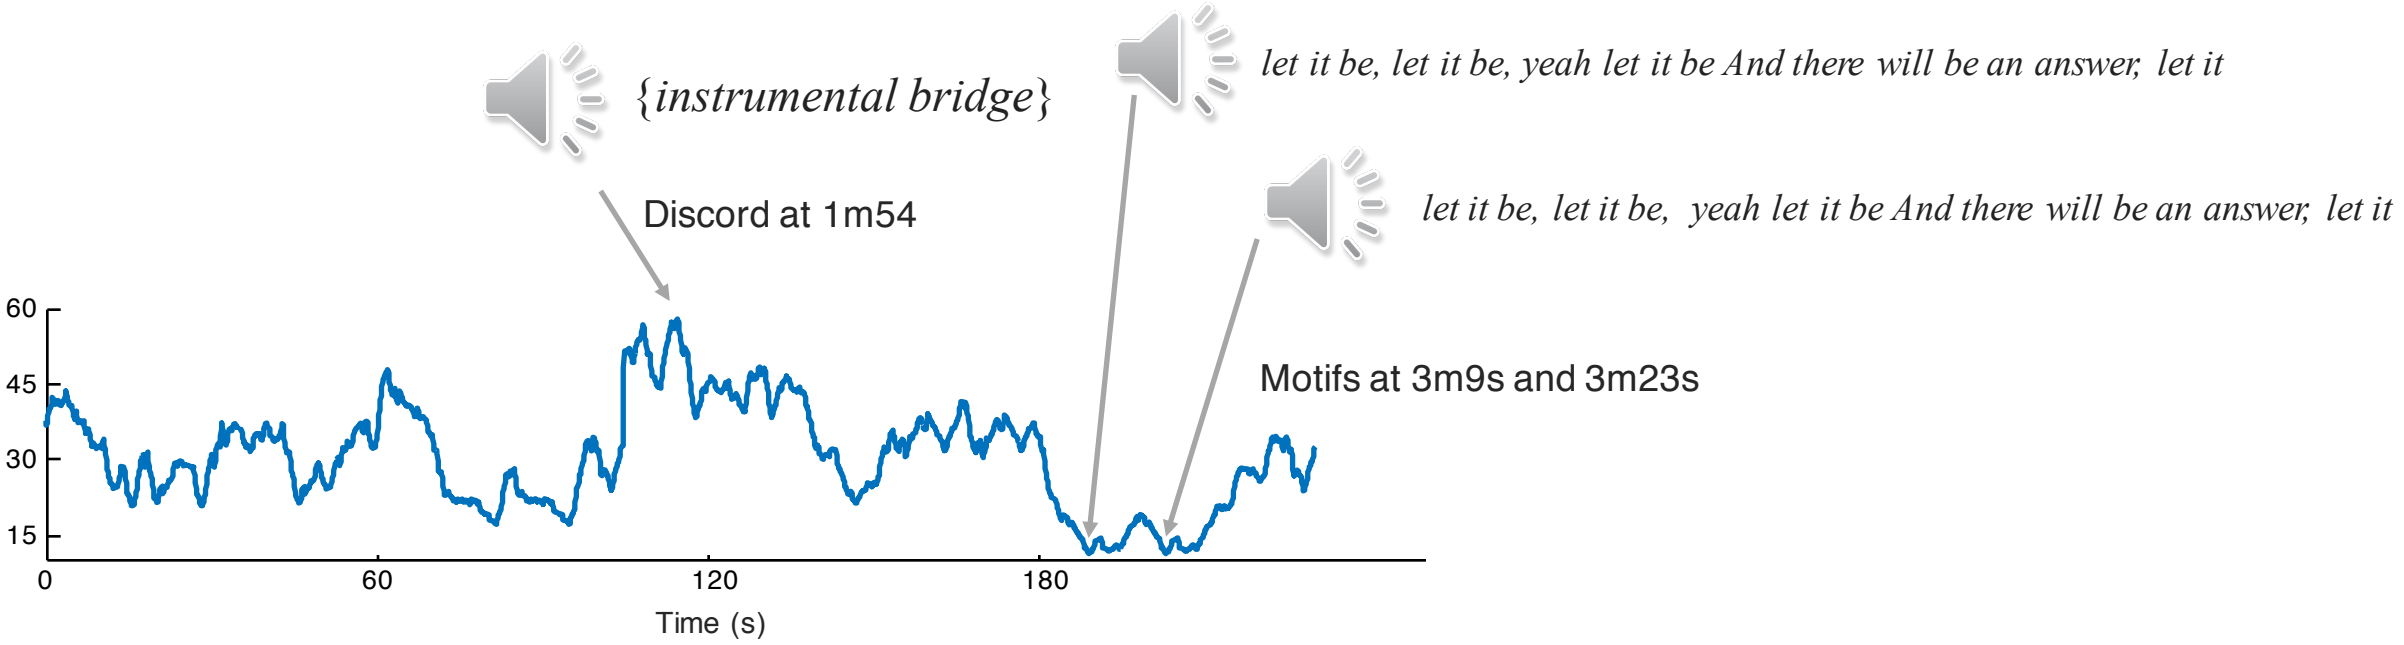

Summary

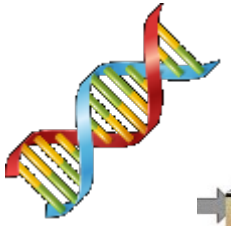
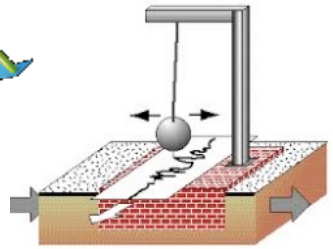

We could do this all day! We have applied the MP to *hundreds* of datasets.

It is worth reminding ourselves of the following:

- The MP can find structure in taxi demand, seismology, DNA, power demand, heartbeats, music and bird vocalization data. However the MP does not "*know*" anything about this domains, it was not tweaked, tuned or adjusted in anyway.
- This domain-agnostic nature of MP is a *great* strength, you typically get very good results "out-of-the-box" no matter what your domain.

The following is also worth stating again:

- We spent time *looking* at the MP to gain confidence that it is doing something useful. However, most of the time, only a higher-level algorithm will look at the MP, with no human intervention or inspection (we will make that clearer in the following slides).

A Minor Visual Mapping Trick

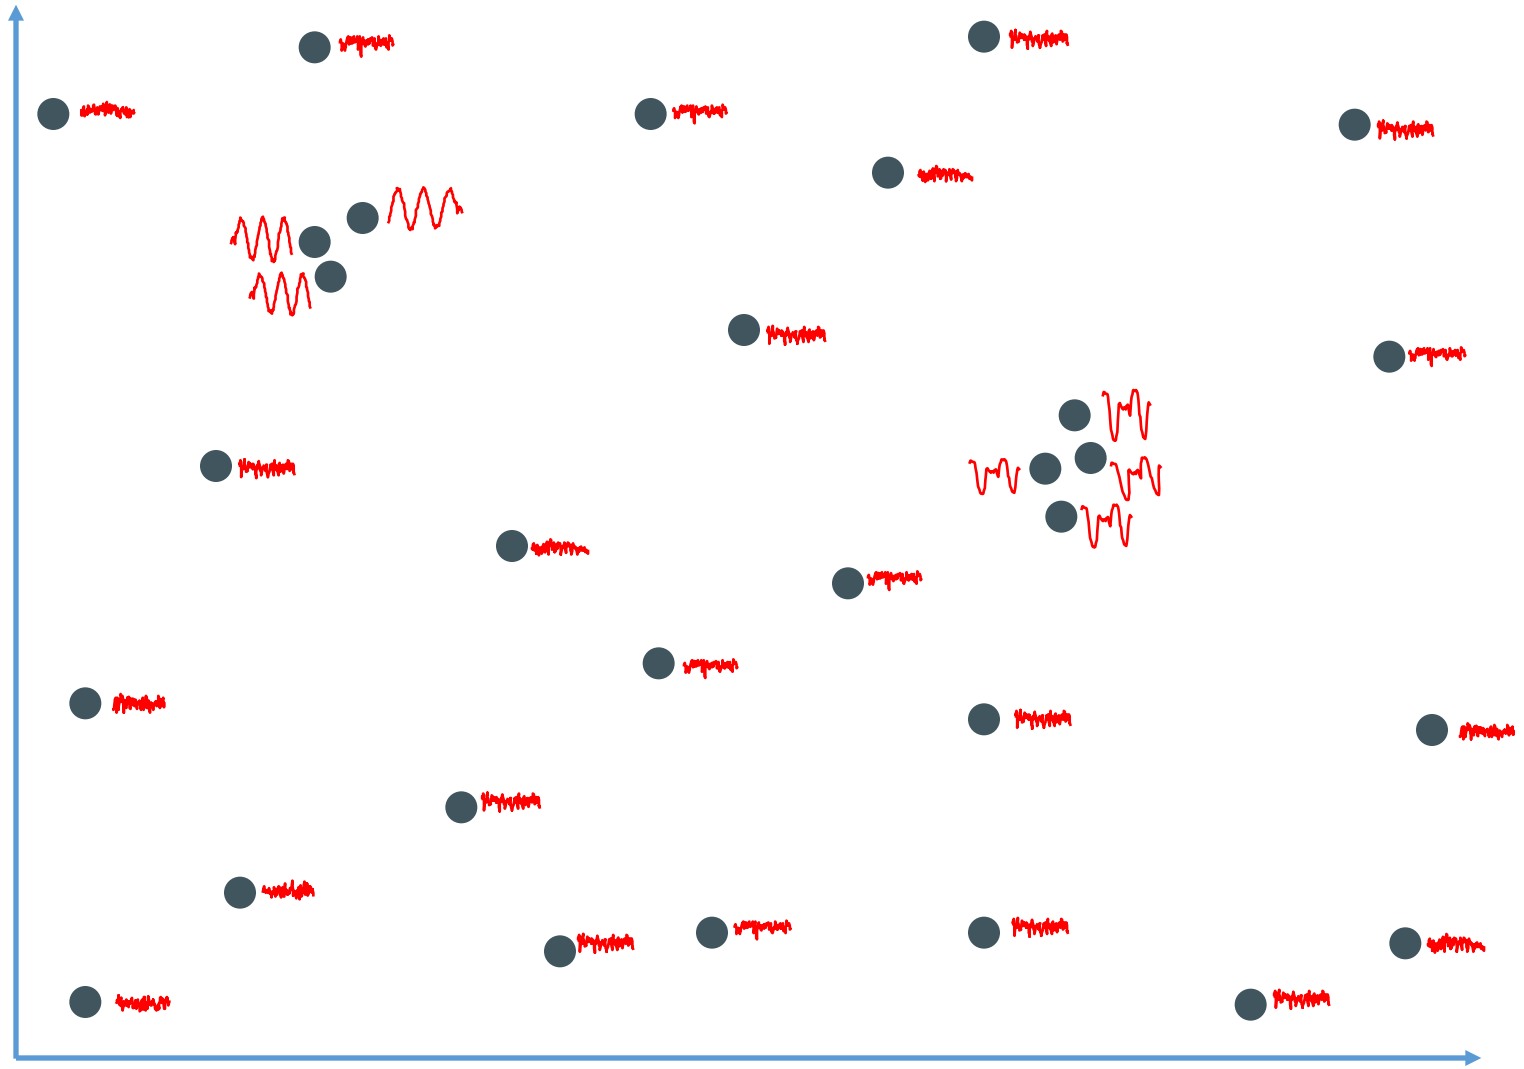

It is sometime useful to think of time series subsequences as points in m-dimensional space.

In this view, dense regions in the *m*-dimensional space correspond to regions of the time series that have a low corresponding MP

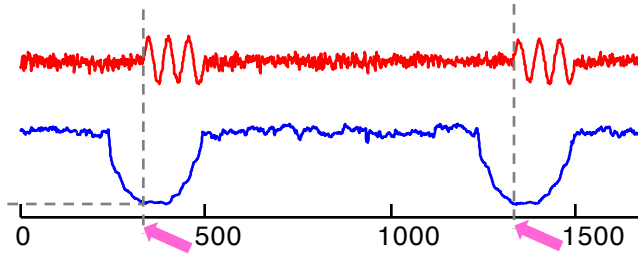

The Top-K Motifs I

Here we show a *sensible* way to extract the top-K motifs. However, there is nothing stopping you from inventing a different way. If you do, the MP will let you compute it in milliseconds.

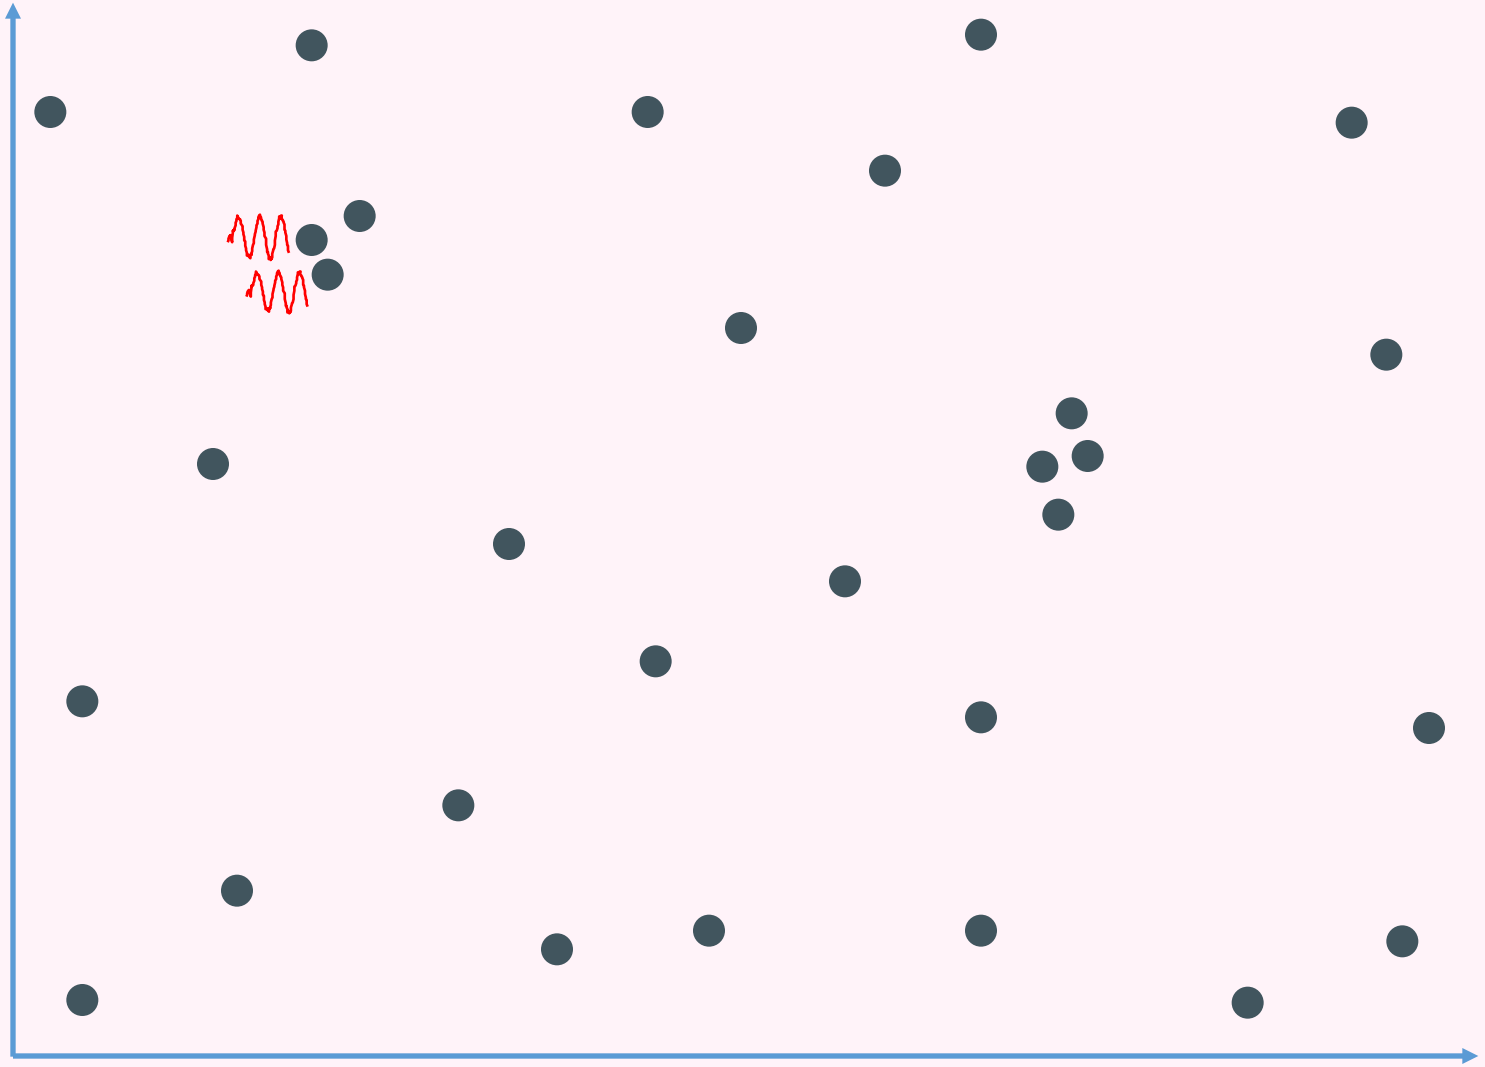

We need a parameter R.

$1 < R <$ (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)

Next slide…

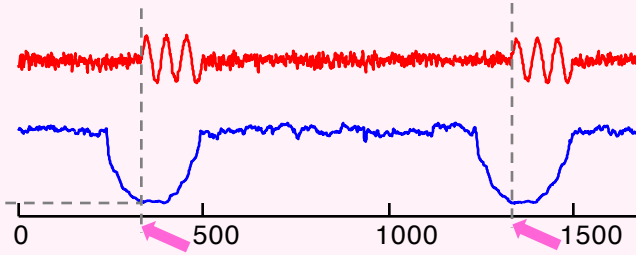

The Top-K Motifs II

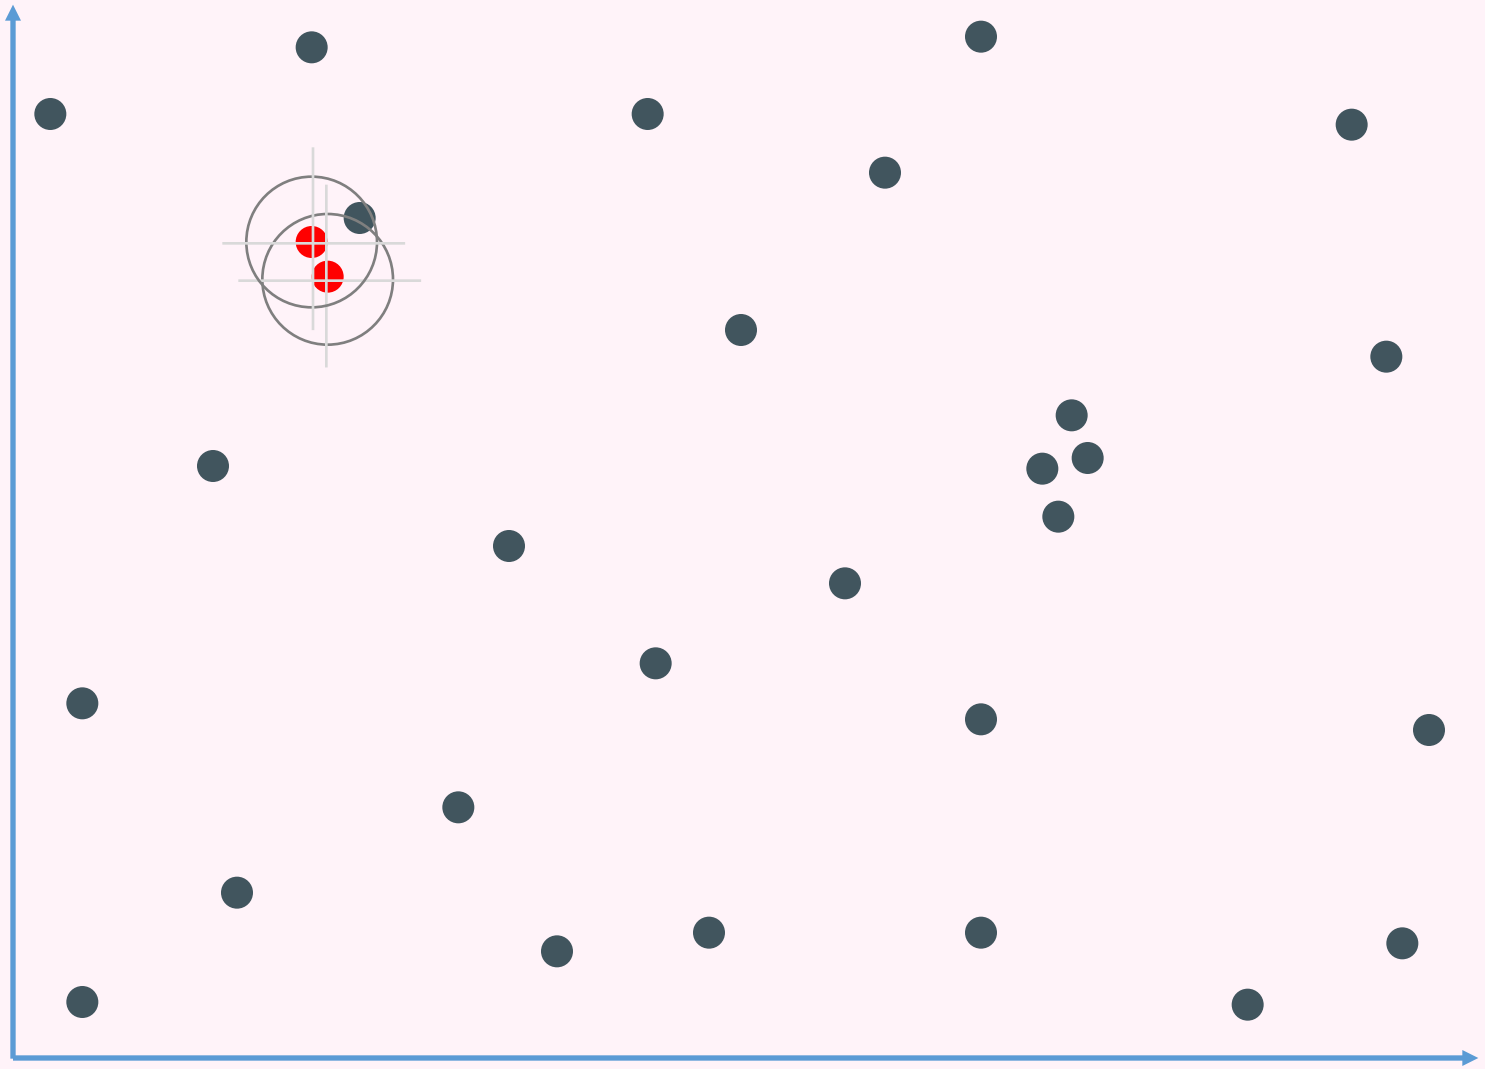

We find the nearest pair of points are D_1 apart.

Lets draw a circle, D_1 *times* R, around both points.

Any points that are within either of these circles, are added to this motif, in this case there is just one…

See next slide…

The Top-K Motifs III

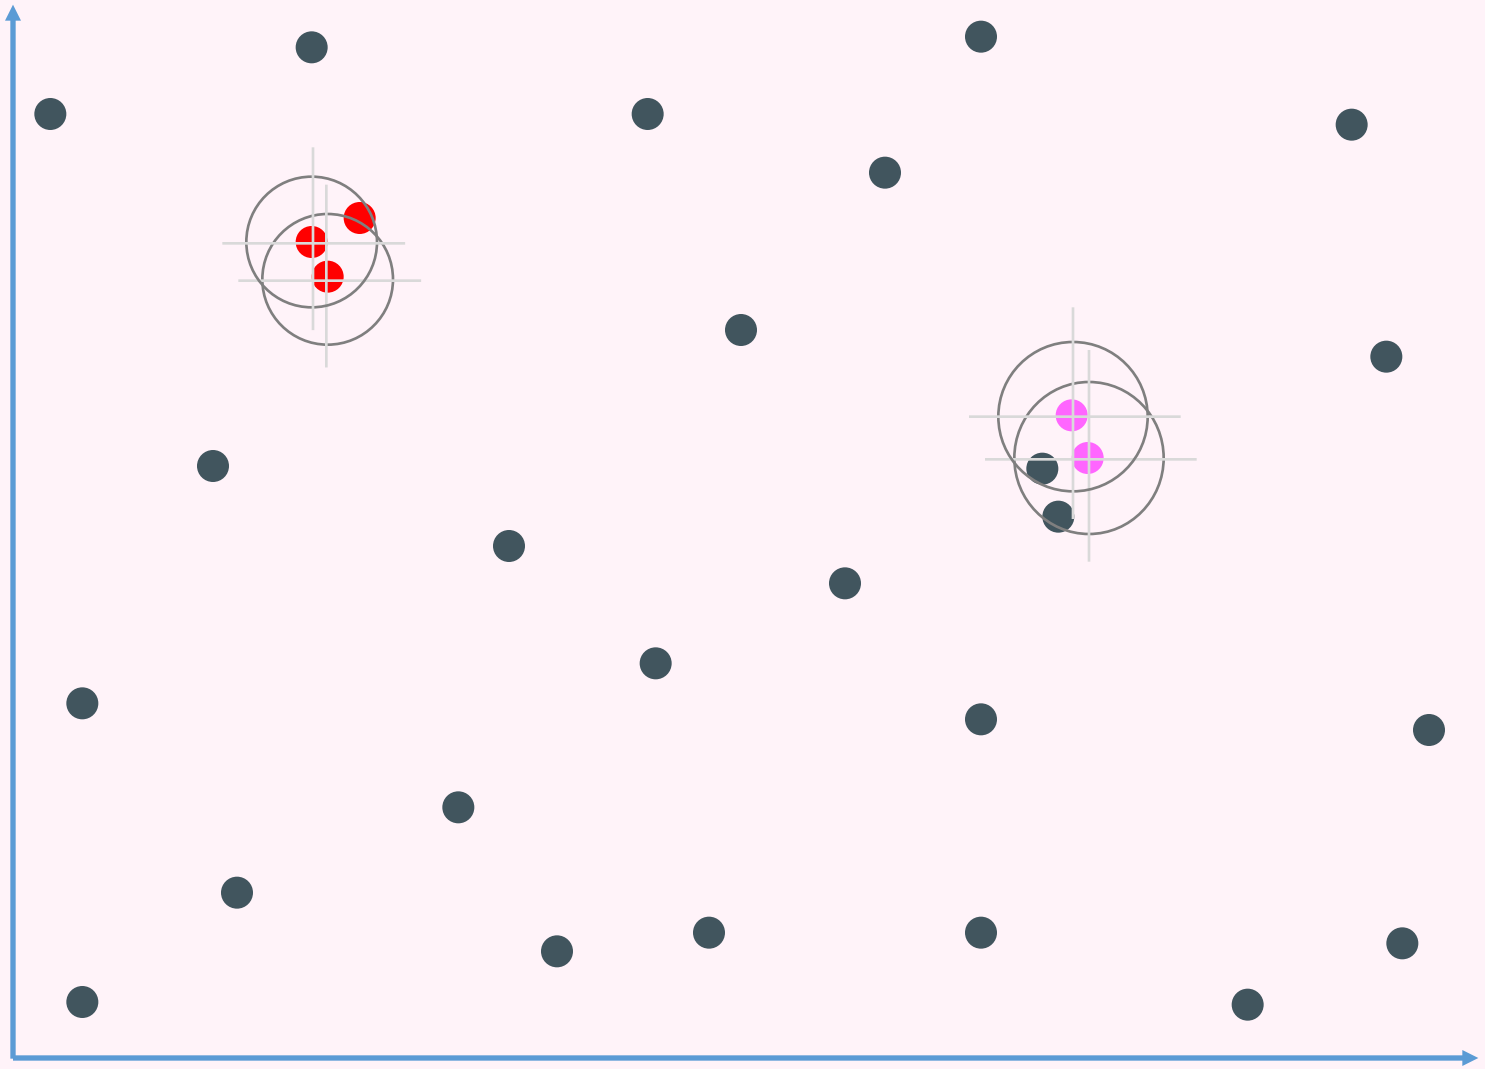

The Top-1 motif has three members, it is done.

Now lets find the Top-2 motif. We begin by finding the nearest pair of points, excluding anything from the top motif.

The nearest pair of points are D_2 apart.
Lets draw a circle D_1 *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… See next slide…

The Top-K Motifs IIII

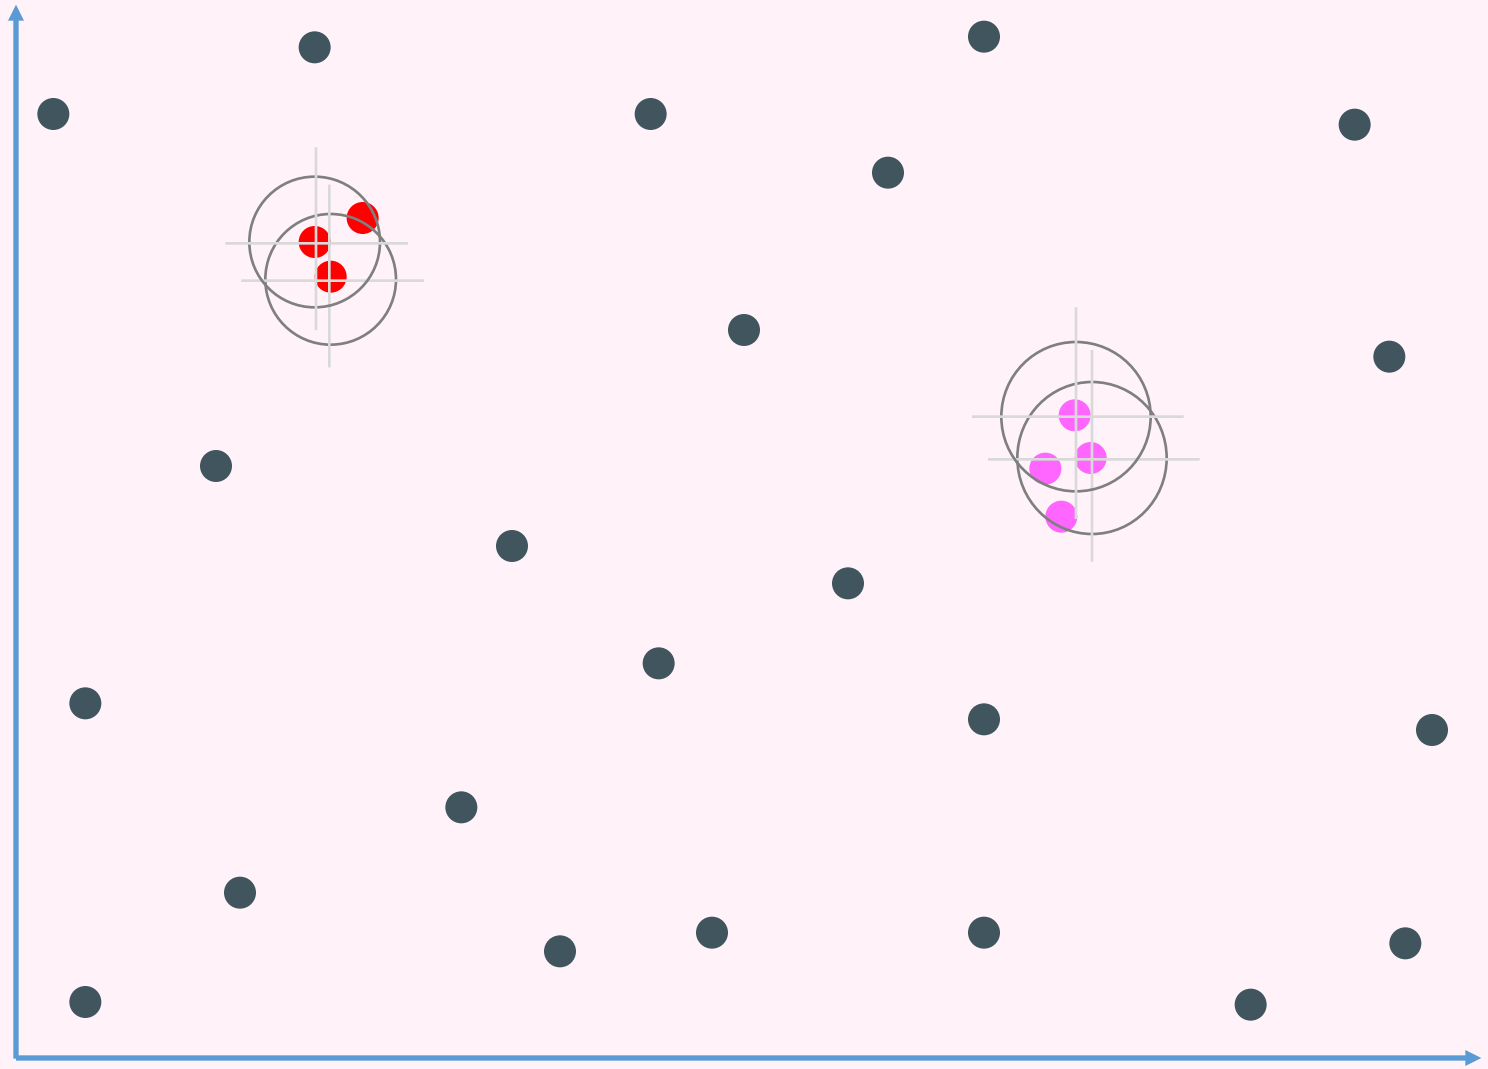

The Top-1 motif has three members, it is done.

Now lets find the Top-2 motif. We begin by finding the nearest pair of points, excluding anything from the top motif.

The nearest pair of points are D_2 apart.
Lets draw a circle D_1 *times* R, around both points.

Any points that are within either of these circles, are added to this motif, in this case there are two, for a total of four items in the Top-2 Motif

The Top-K Motifs V

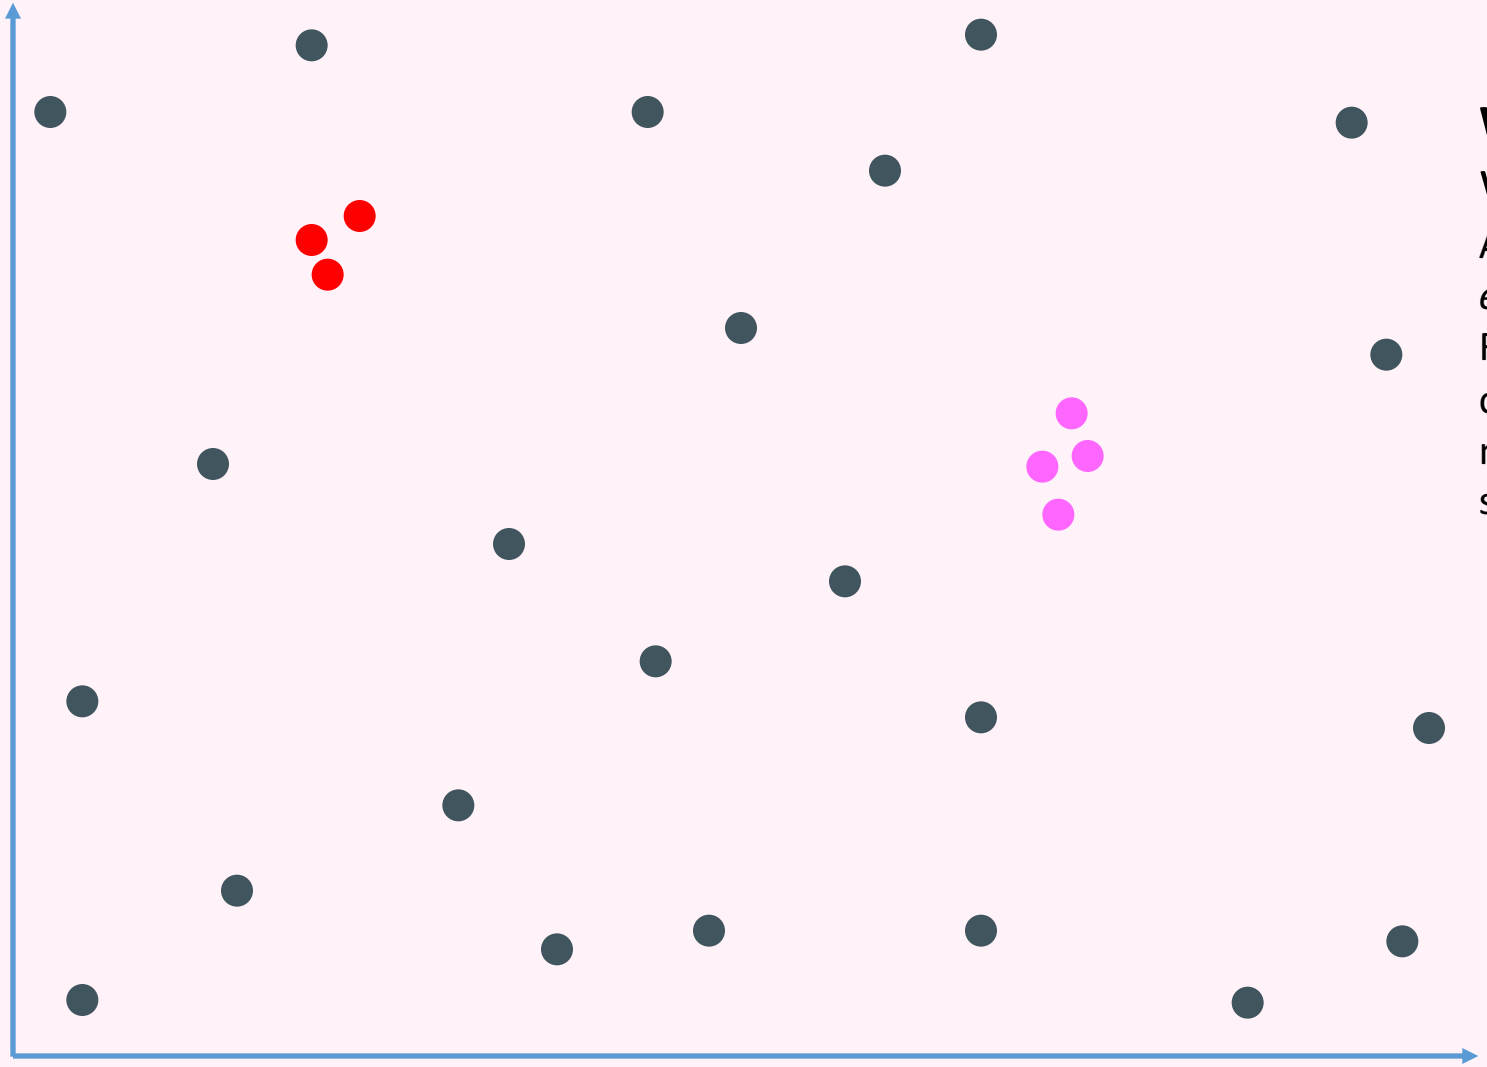

We are done with the Top-2 Motif

Note that we will always have:

$D_1 < D_2 < D_3 \ldots$

When to stop? (what is K?)

We could use MDL etc.

As a practical matter, we can pull out all K, and use *eyeballing* to judge the quality of motifs.

For example, in the below, Motif 1 is stunningly well conserved, Motif 2 is somewhat conserved, Motif 3 may be getting close to random…. So here we would say we have a strong **Top-2** Motifs.

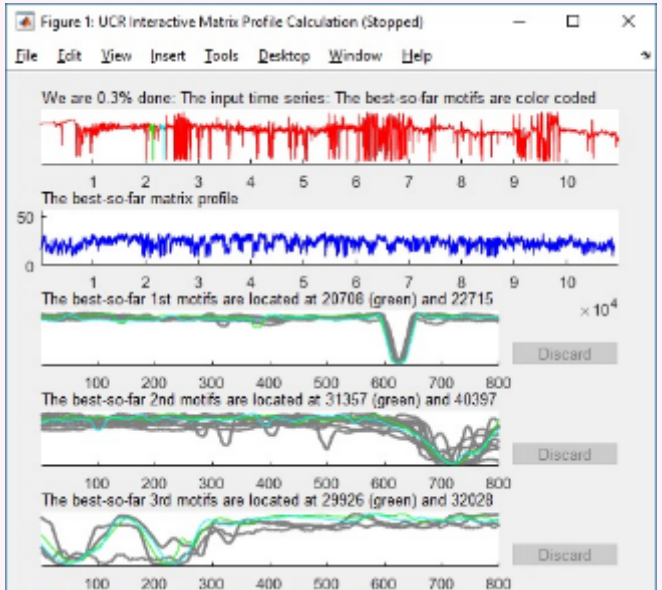

From Motifs to Time Series Chains

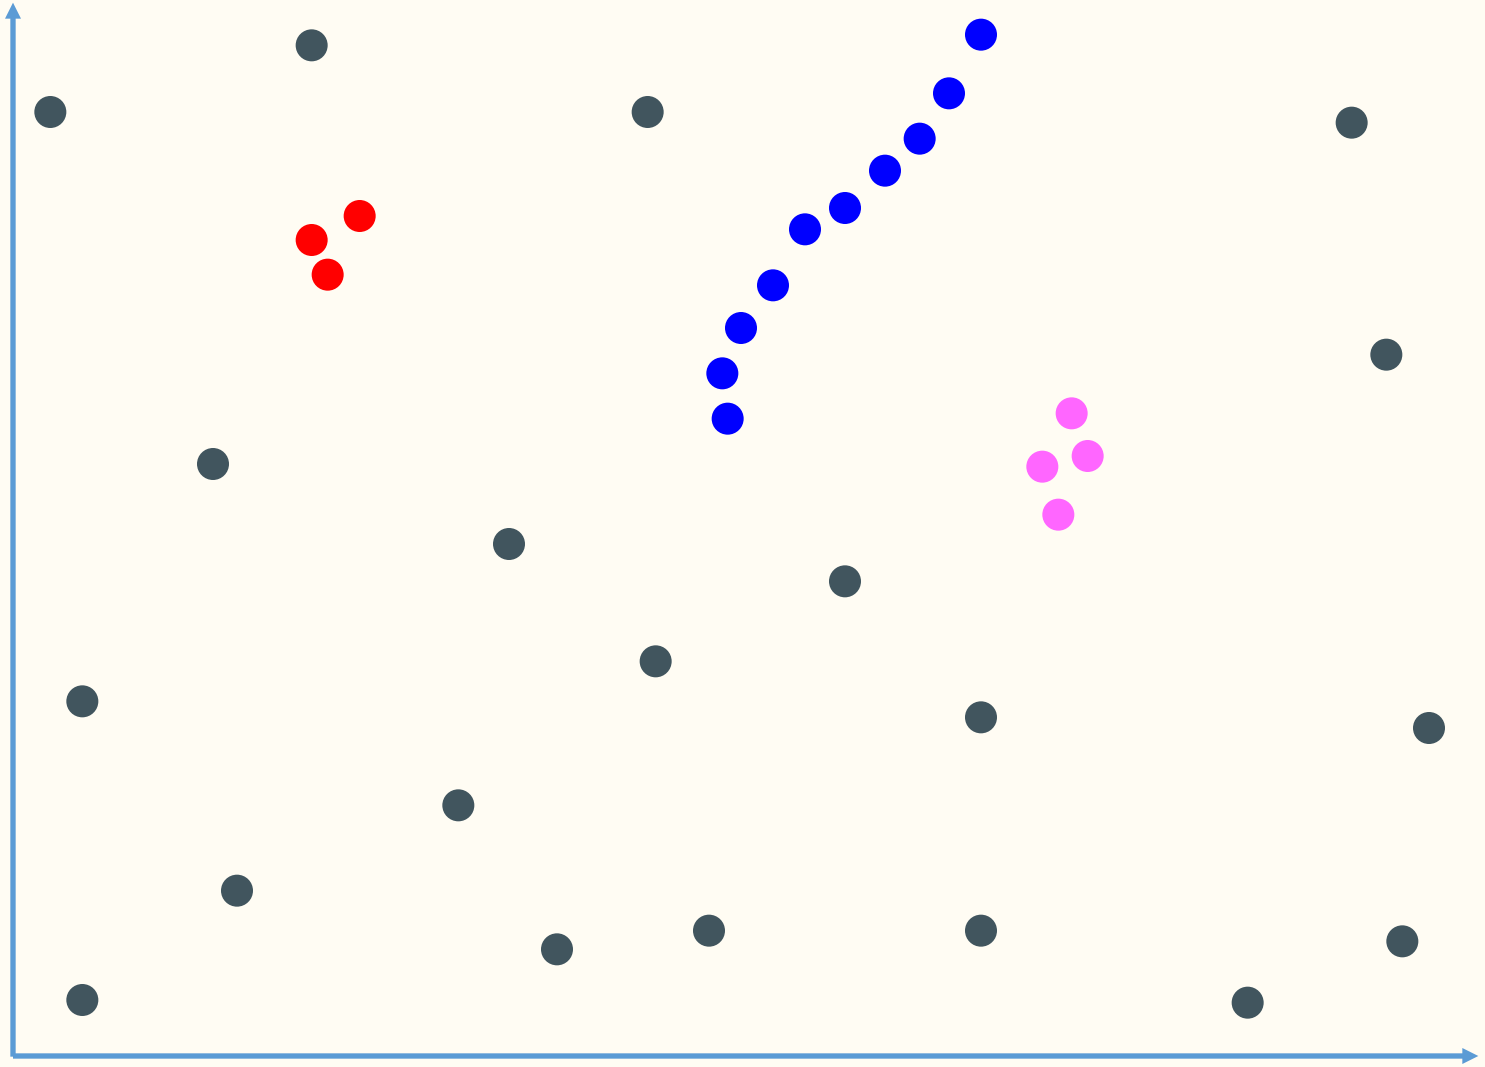

Take a look at the blue ‘subsequences”

They would not from a single motif (but perhaps they could form a *set* of motifs).

From Motifs to Time Series Chains

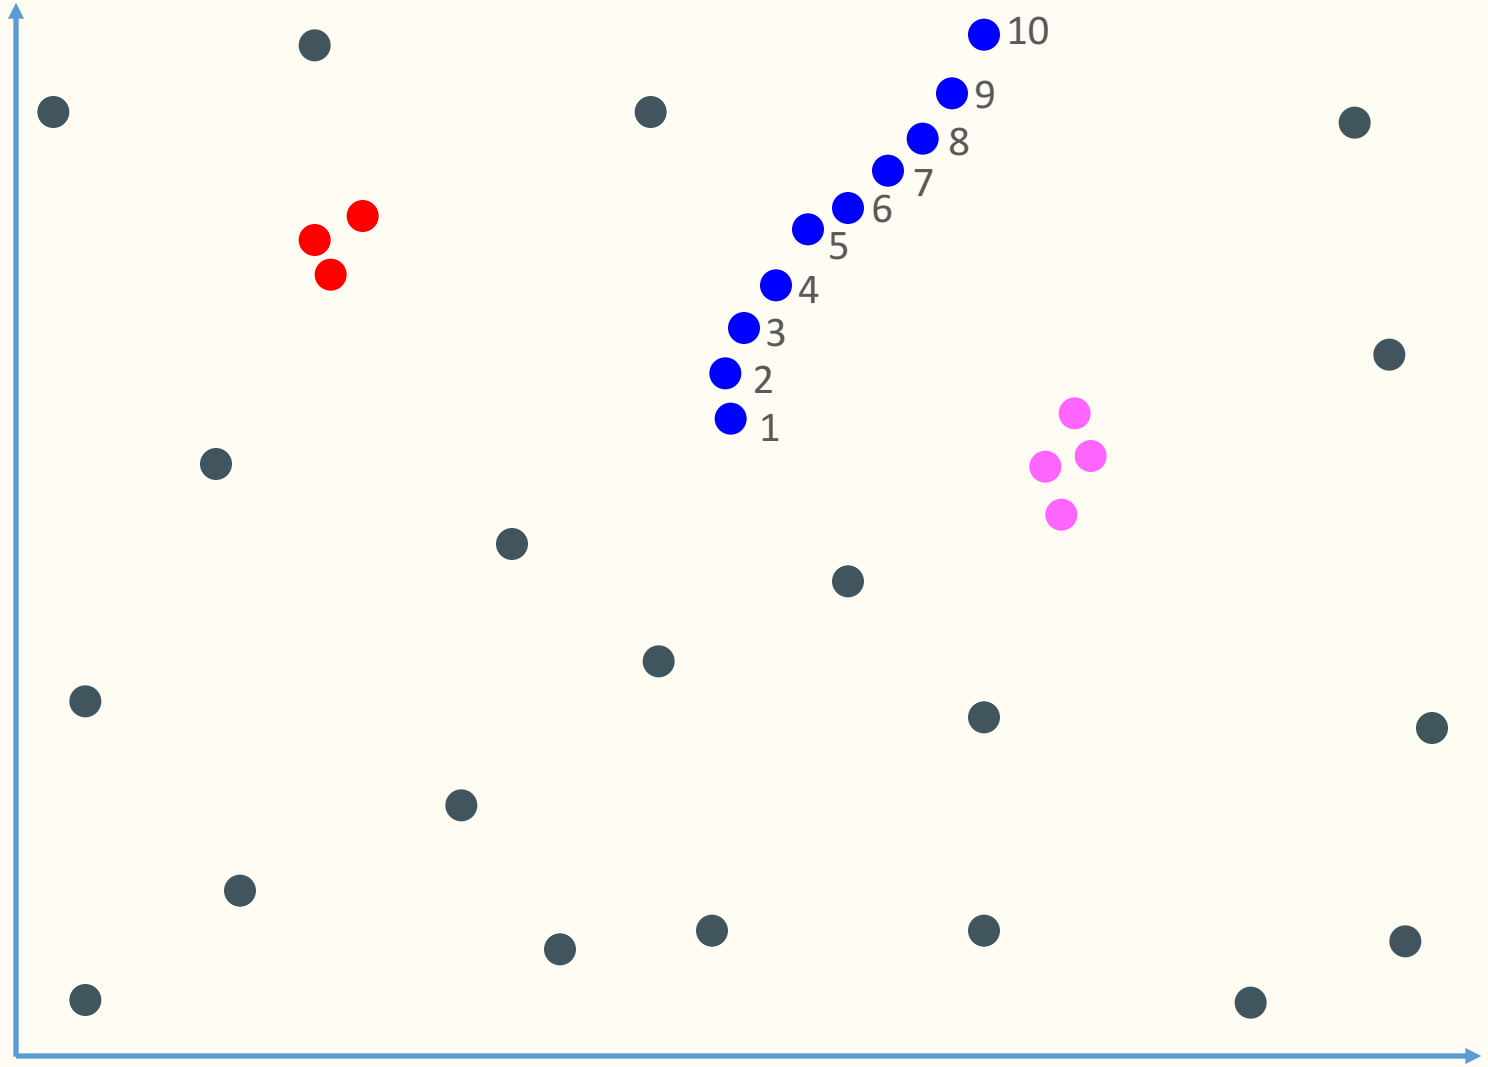

However, if we label them by *arrival* time, you can see that they are *drifting*, or *evolving* in time.

This is *actionable*, for example, where will the 11th item land? Surely just Northeast of the 10th item

We call such pattern chains, with the first item as the *anchor*.

Do such patterns exist in the real world?

Can we find them?

Arterial Blood Pressure

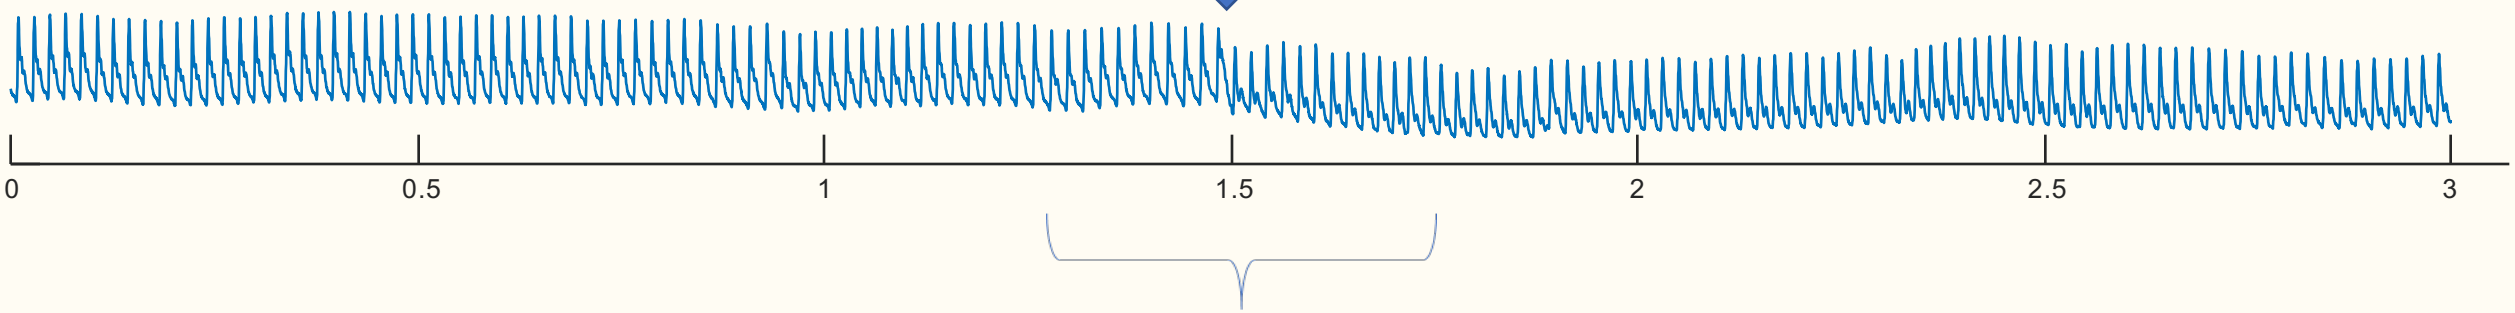

We will zoom-in to here in the next slide

We ran time series chain discovery on the dataset. The only thing we tell it is the length of the subsequence to use (about one heartbeat long).

Zoom In

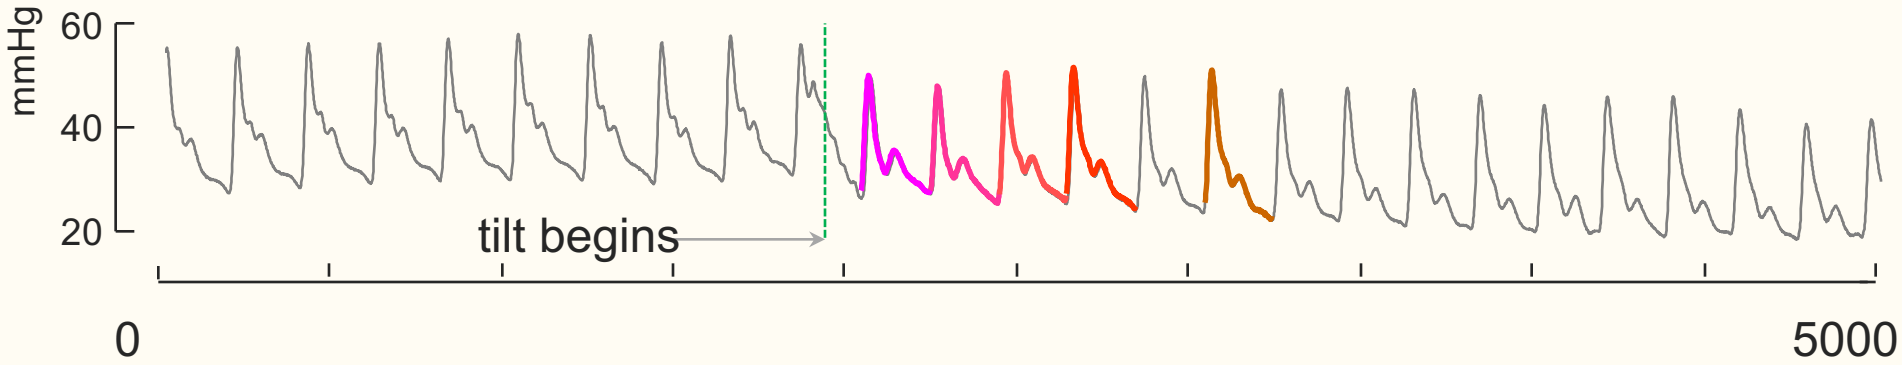

Ads the chain progresses, the depth of the dicrotic notch decreases….

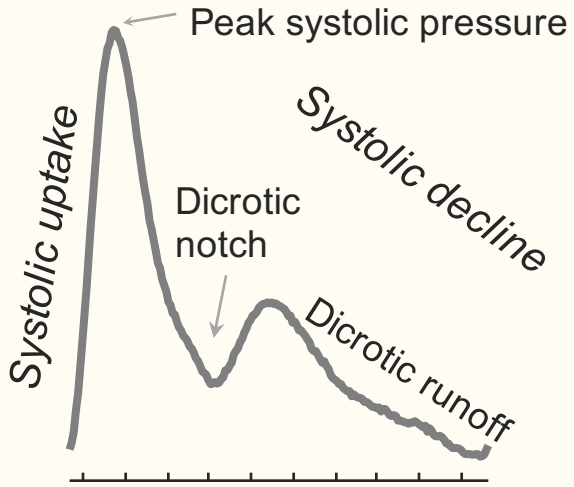

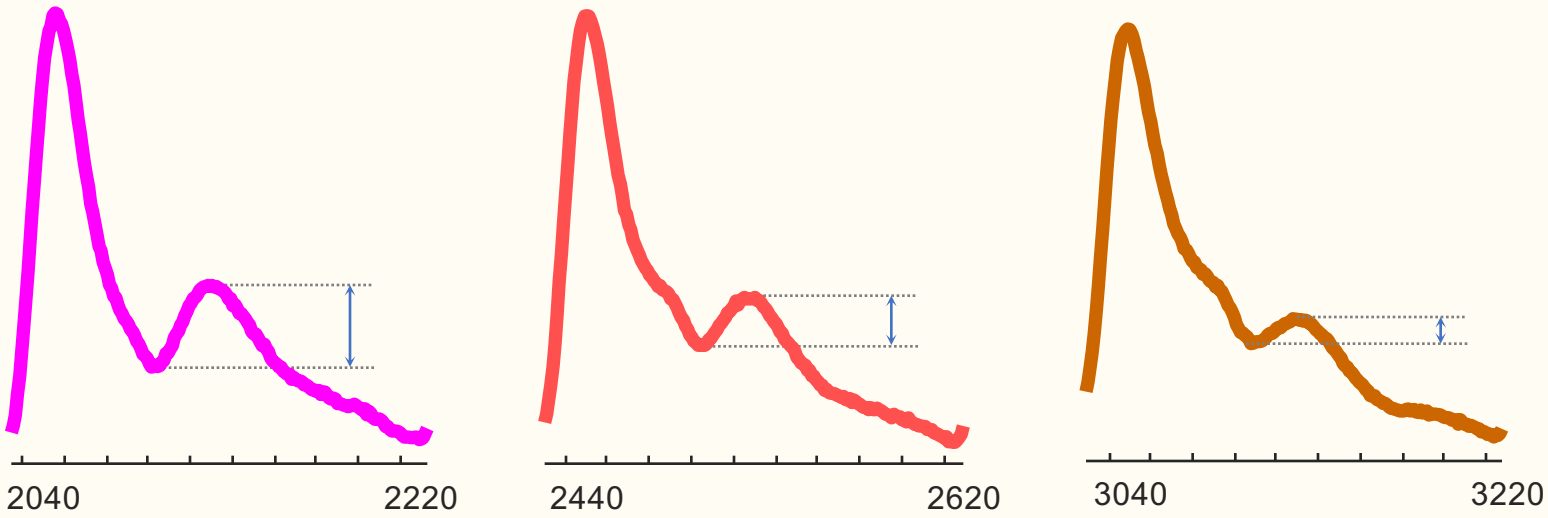

More Time Series Chains

We looked at the google query volume for *Kohl's*, an American retail chain.

The discovered chain shows that over the decade, the bump transitions from a smooth bump covering most of the period between thanksgiving and Xmas, to a more sharply focus bump centered on thanksgiving. This seems to reflect the growing importance of *Cyber Monday*, a marketing term for the Monday after Thanksgiving. The phrase was created by marketing companies to persuade people to shop online. The term made its debut on November 28th, 2005 in a press release entitled "*Cyber Monday Quickly Becoming One of the Biggest Online Shopping Days of the Year*" . Note that this date coincides with the first glimpse of the sharping peak in our chain.

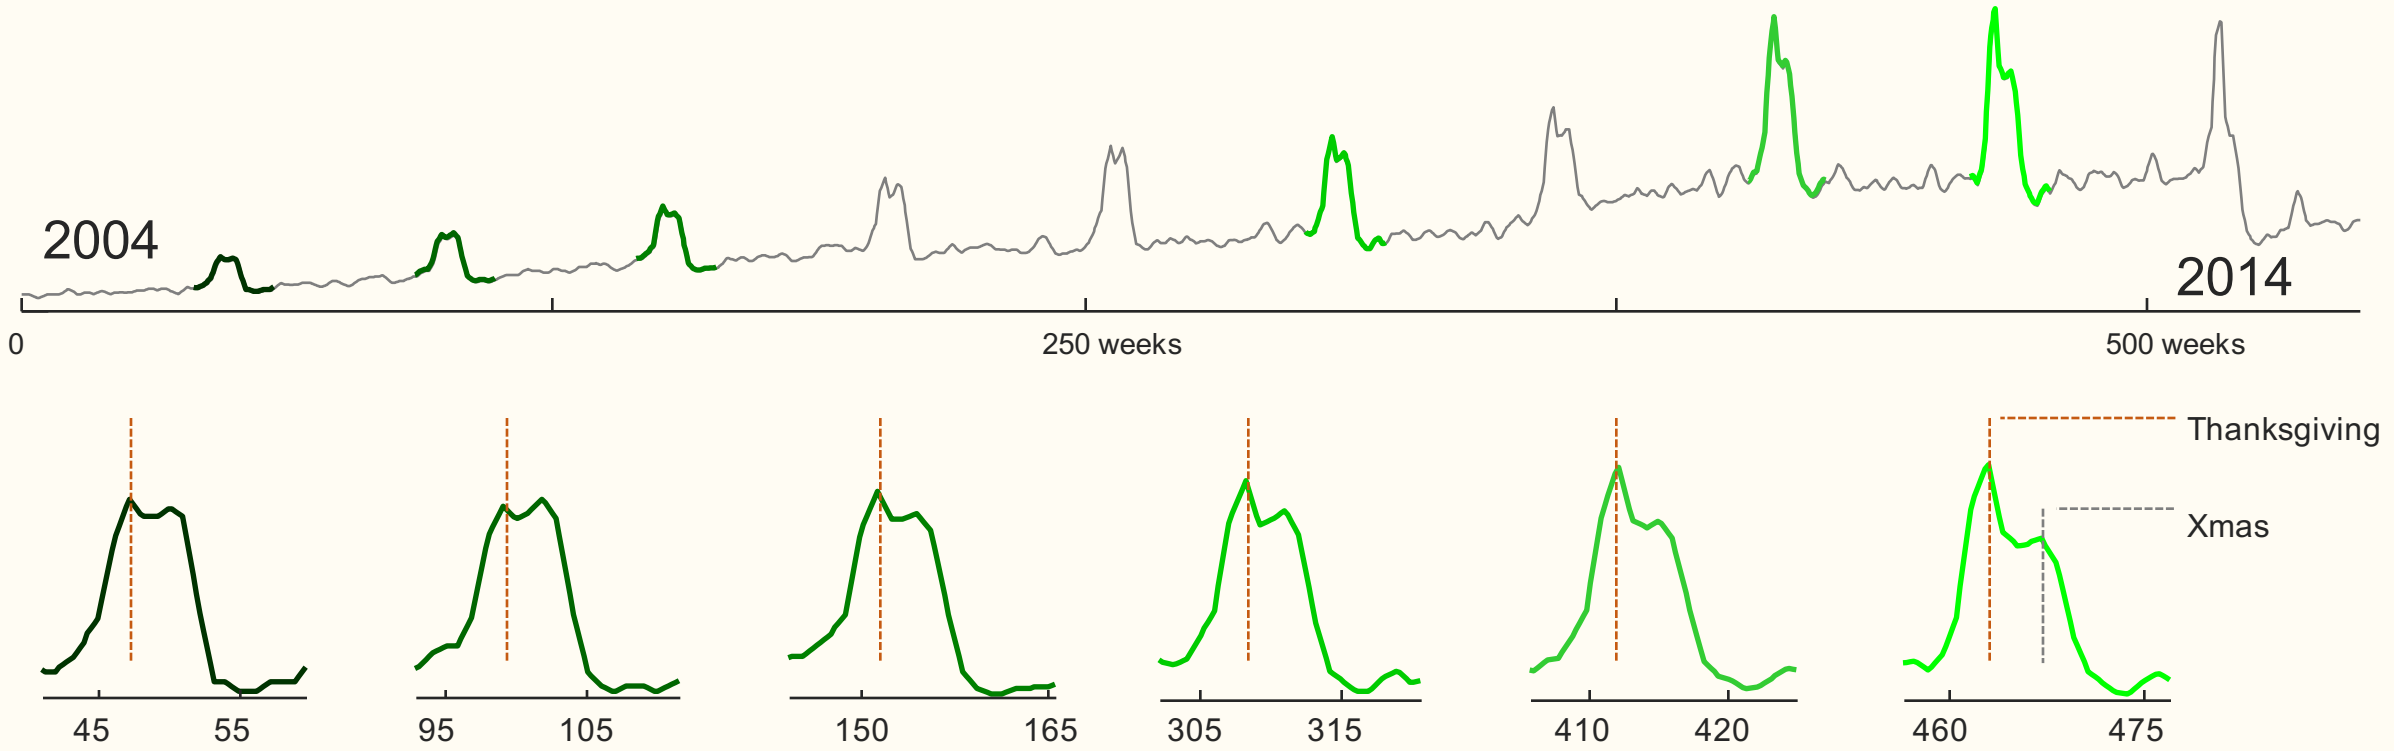

One Last Time Series Chain

Magellanic penguins regularly dive to depths of up to 50m to hunt prey. Penguins have typical body densities for a bird, but just before diving they take a very deep breath that makes them exceptionally buoyant. This positive buoyancy is difficult to overcome near the surface, but at depth, the compression of water pressure cancels it. In order to get to down to their hunting ground below sea level it is clear that "*locomotory muscle workload, varies significantly at the beginning of dives*"*.

The snippet of time series shown in does not suggest much of a change in *stroke-rate*, however penguins are able vary the thrust of their flapping by twisting their wings. The chains we discovered shows this dramatic and evolving sprint downwards leveling off to a comfortable cruise.

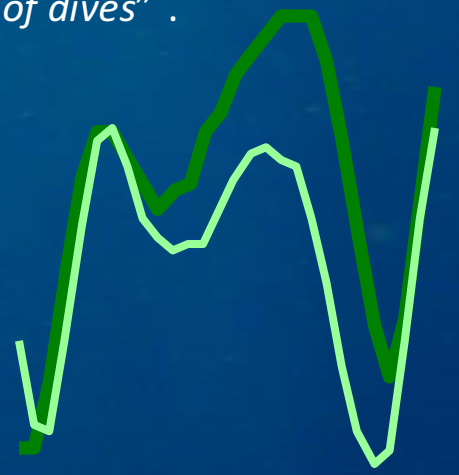

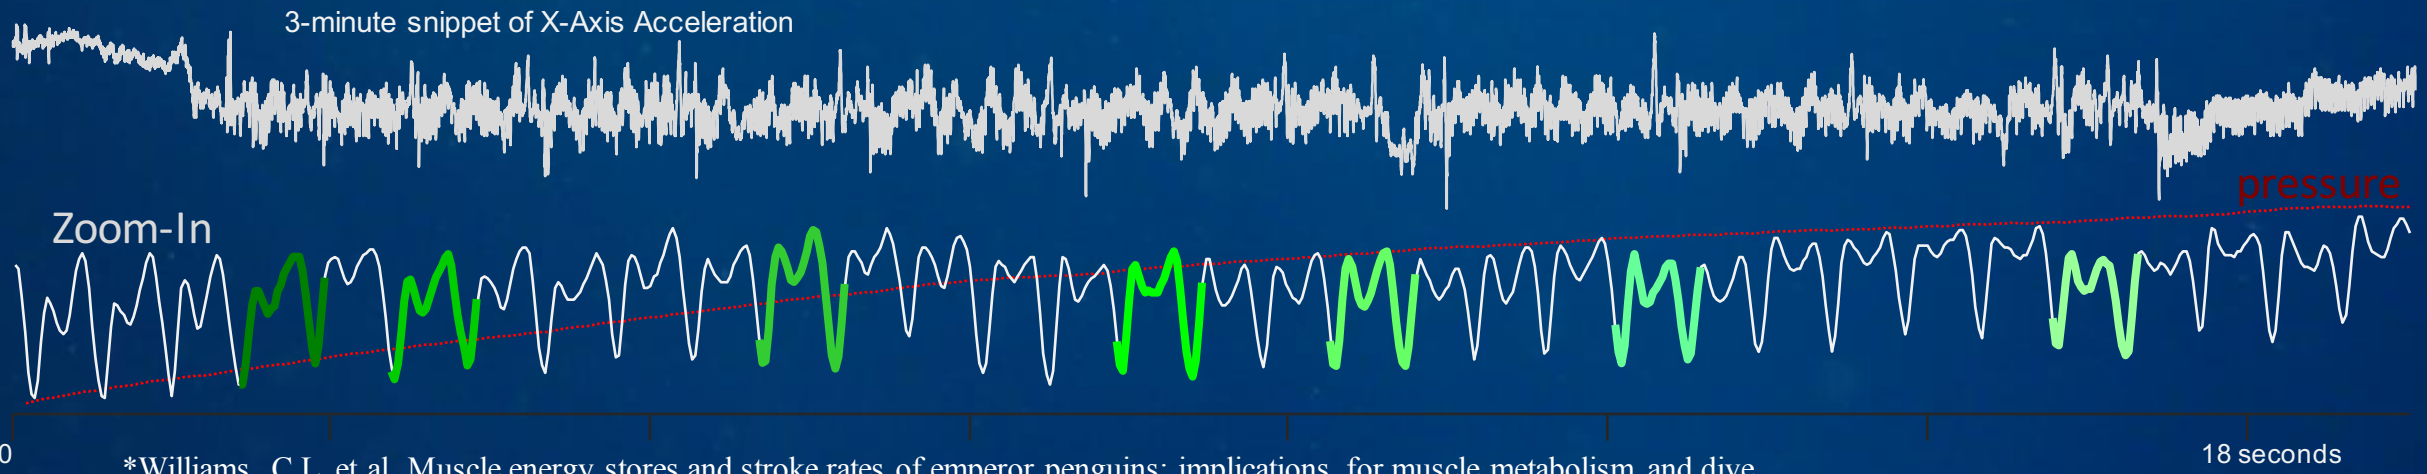

*Williams, C.L. et al. Muscle energy stores and stroke rates of emperor penguins: implications for muscle metabolism and dive performance. *Physiological and Biochemical Zoology.85.2*(2011):120-133

Photo by Paul J. Ponganis

- There are literally 100's of time series anomaly detectors.
- However, many claim that *Time Series Discords* is among the best.

> *..on 19 different publicly available data sets, comparing 9 different techniques (time series **discords**) is the best overall technique among all techniques.* Vipin Kumar*

Vipin Kumar
ACM SIGKDD 2012 Innovation Award Winner

- This is good news for us, because if you compute the matrix profile, you have the discords "for free". In fact, you have *all* the top K-discords, for any K.
- Why are discords so effective? (our subjective opinion)
  - They make no assumptions about the data (so no *wrong* assumptions).
  - They don't need to learn a bunch of parameters, with no parameters to fit, it is hard to *over*fit.
- There is one pathological (but fixable) case where they don't work (next slide)

*https://www.cs.umn.edu/research/technical_reports/view/09-004

The twin freak problem (see next slide)

The definition of a discord is:
The subsequence D that has the maximum distance from its (non-trivial match) nearest neighbor.

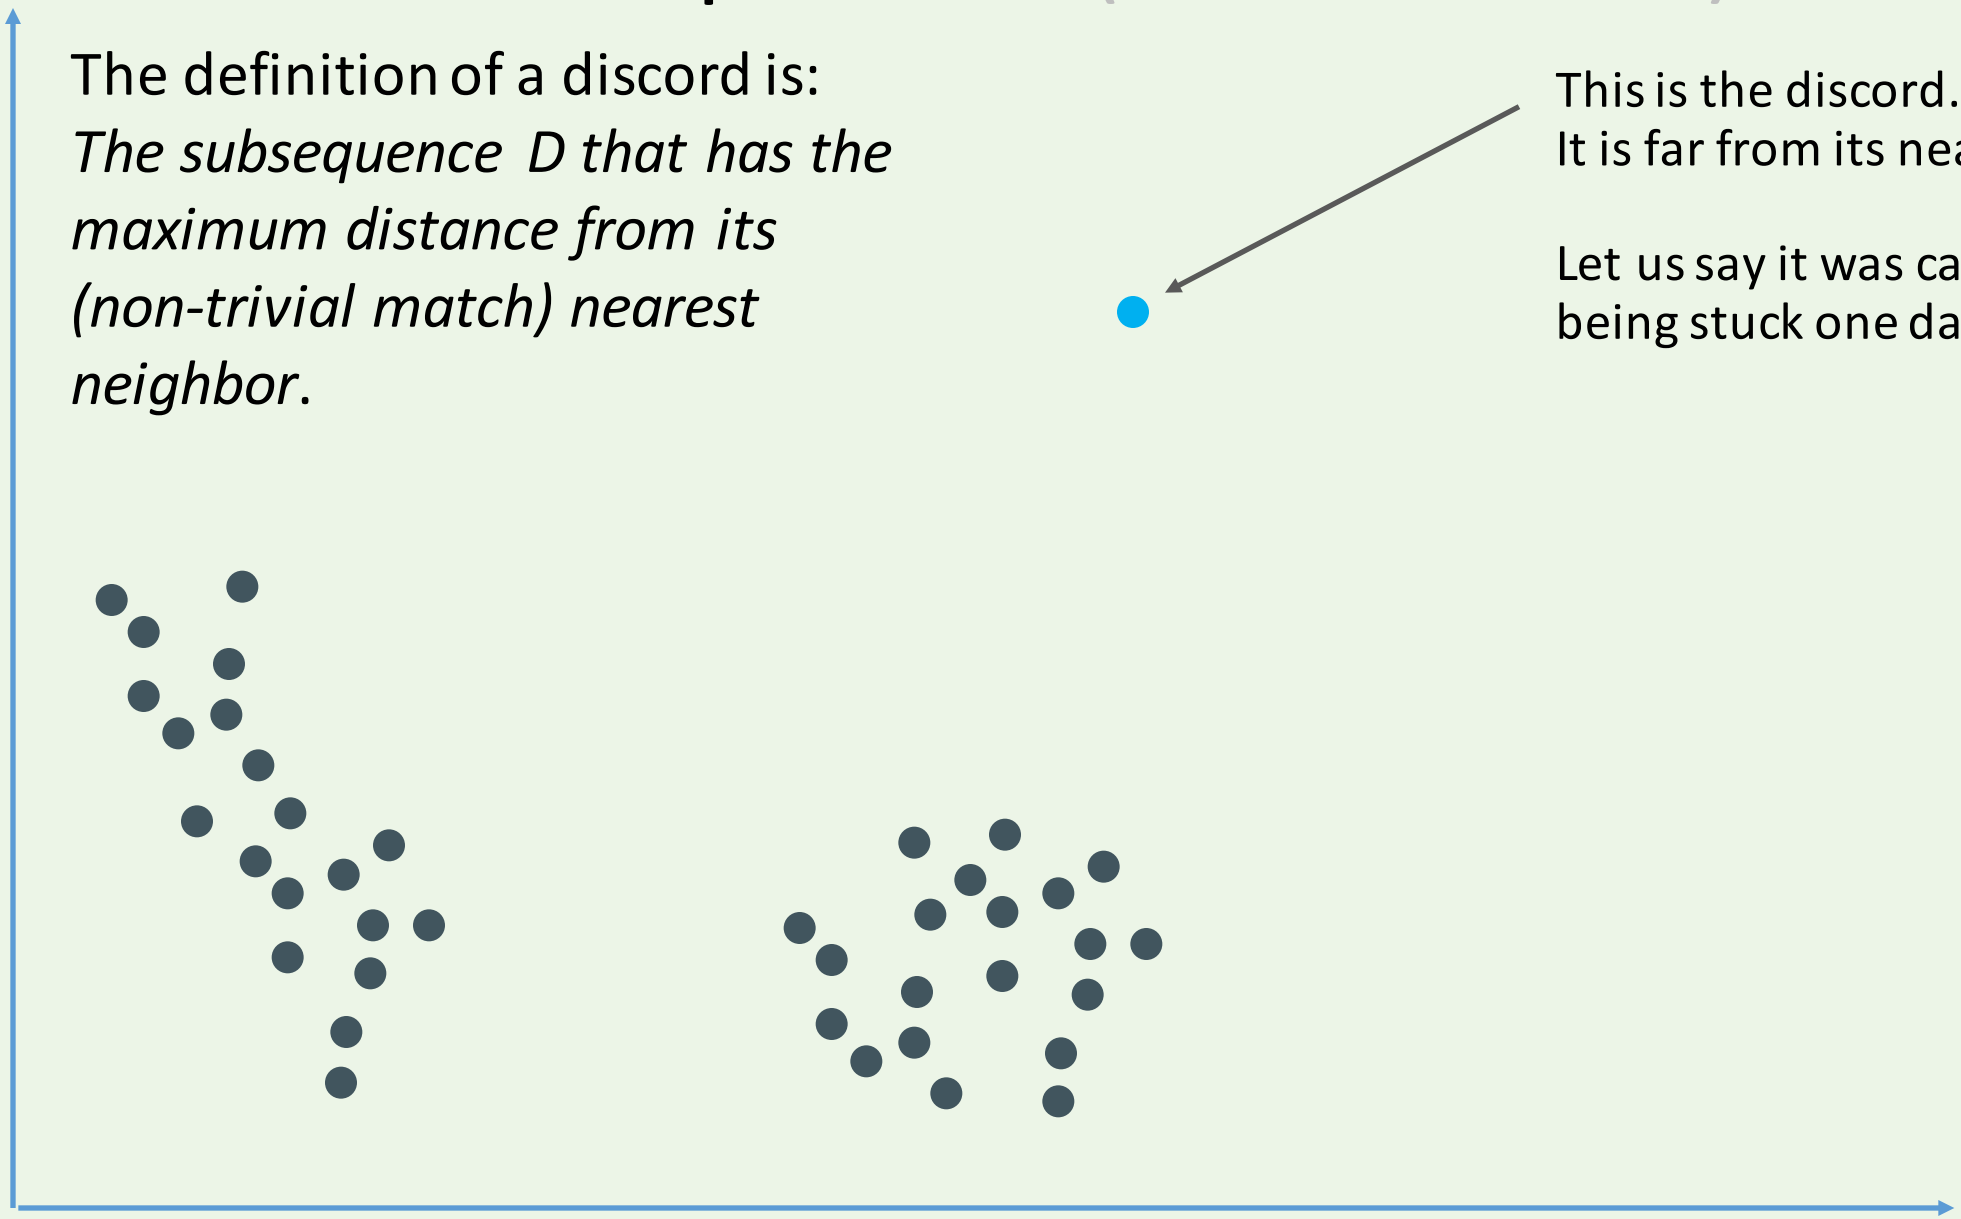

This is the discord.
It is far from its nearest neighbor

Let us say it was caused be a valve being stuck one day..

The twin freak problem

The definition of a discord is:
The subsequence D that has the maximum distance from its (non-trivial match) nearest neighbor.

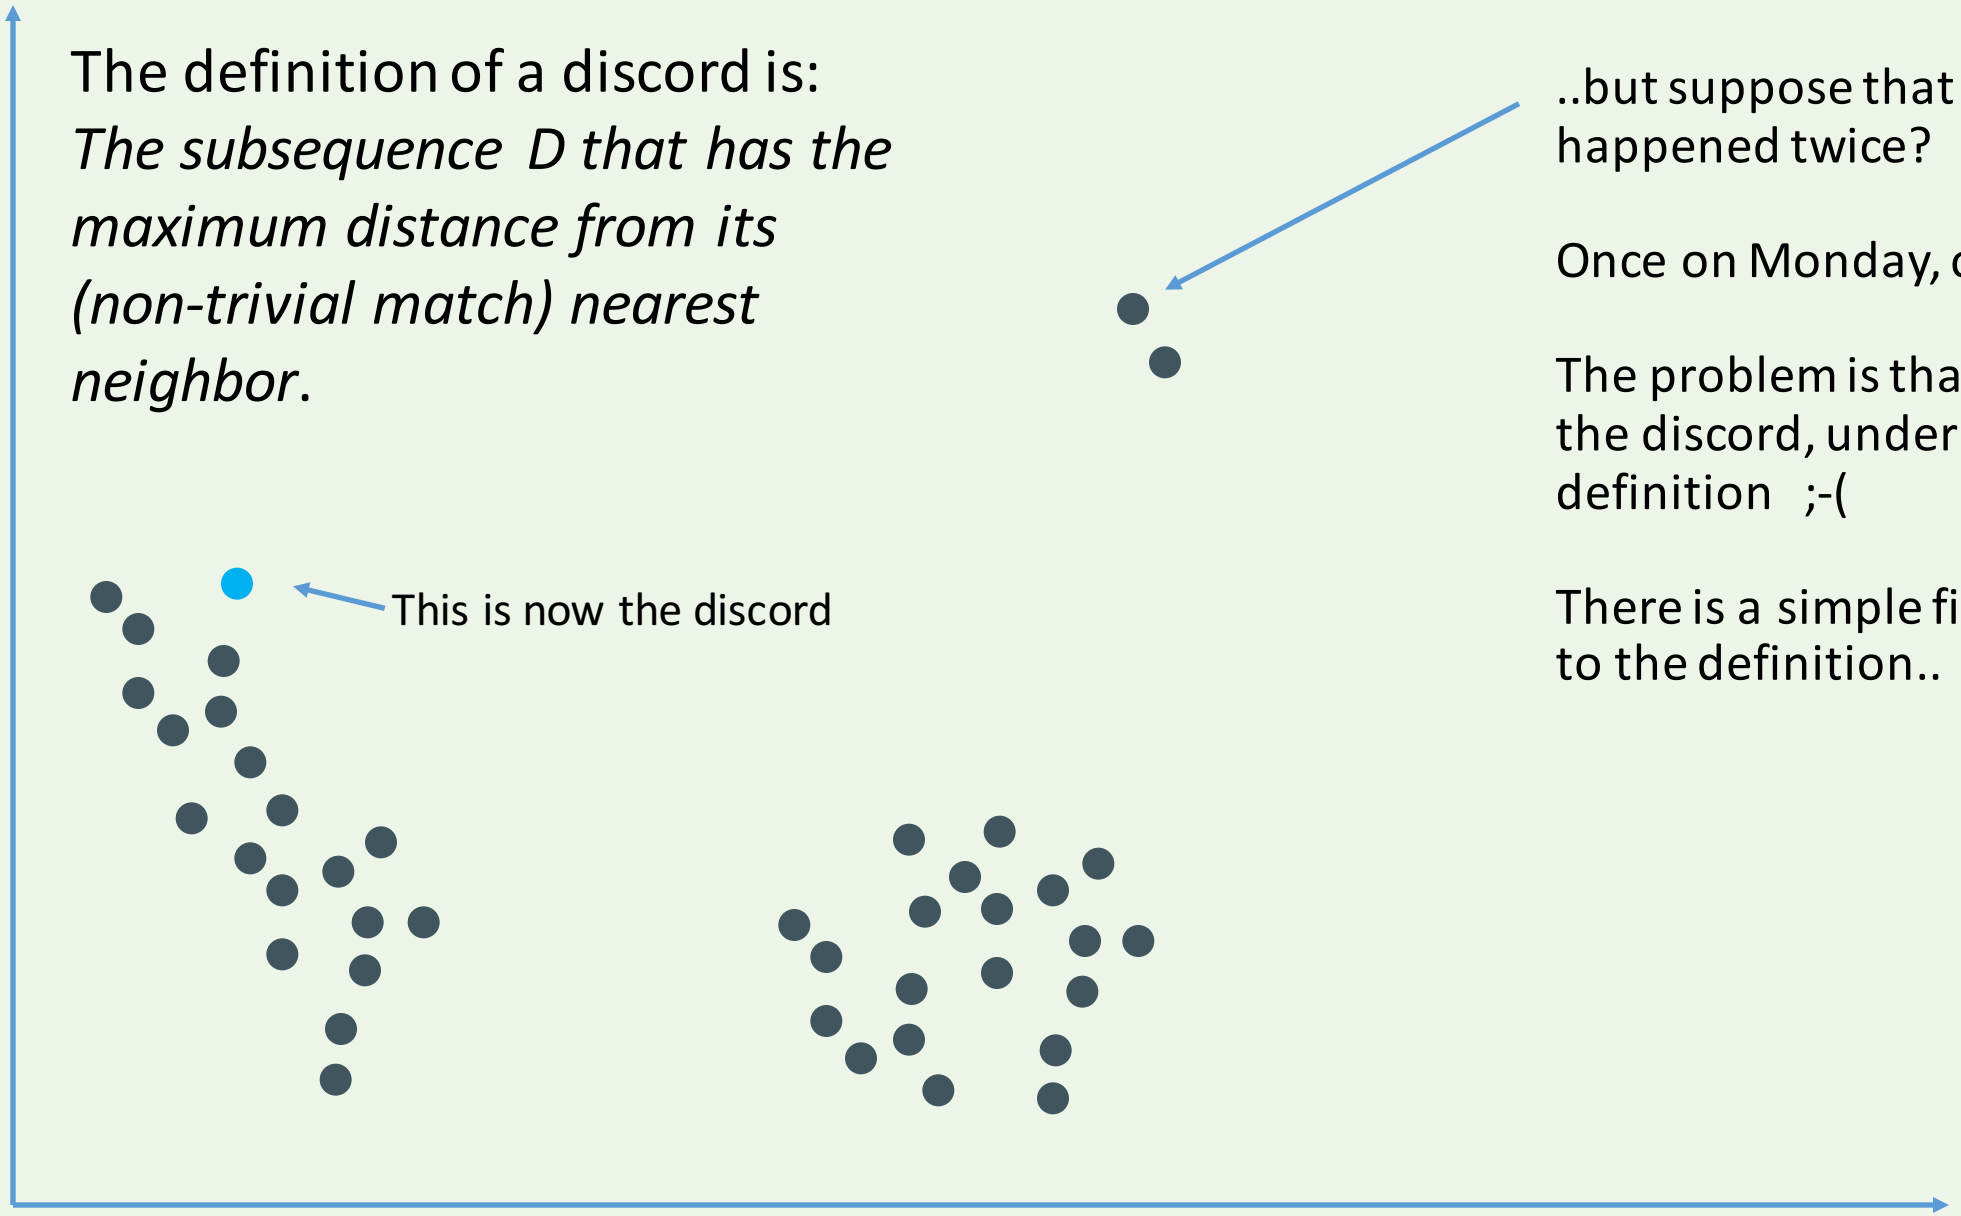

..but suppose that the anomaly happened twice?

Once on Monday, once on Friday…

The problem is that it is no longer the discord, under our classic definition ;-(

There is a simple fix, a minor change to the definition..

The twin freak problem

The new definition of a discord is:
The subsequence D that has the maximum distance from its (non-trivial match) second nearest neighbor.

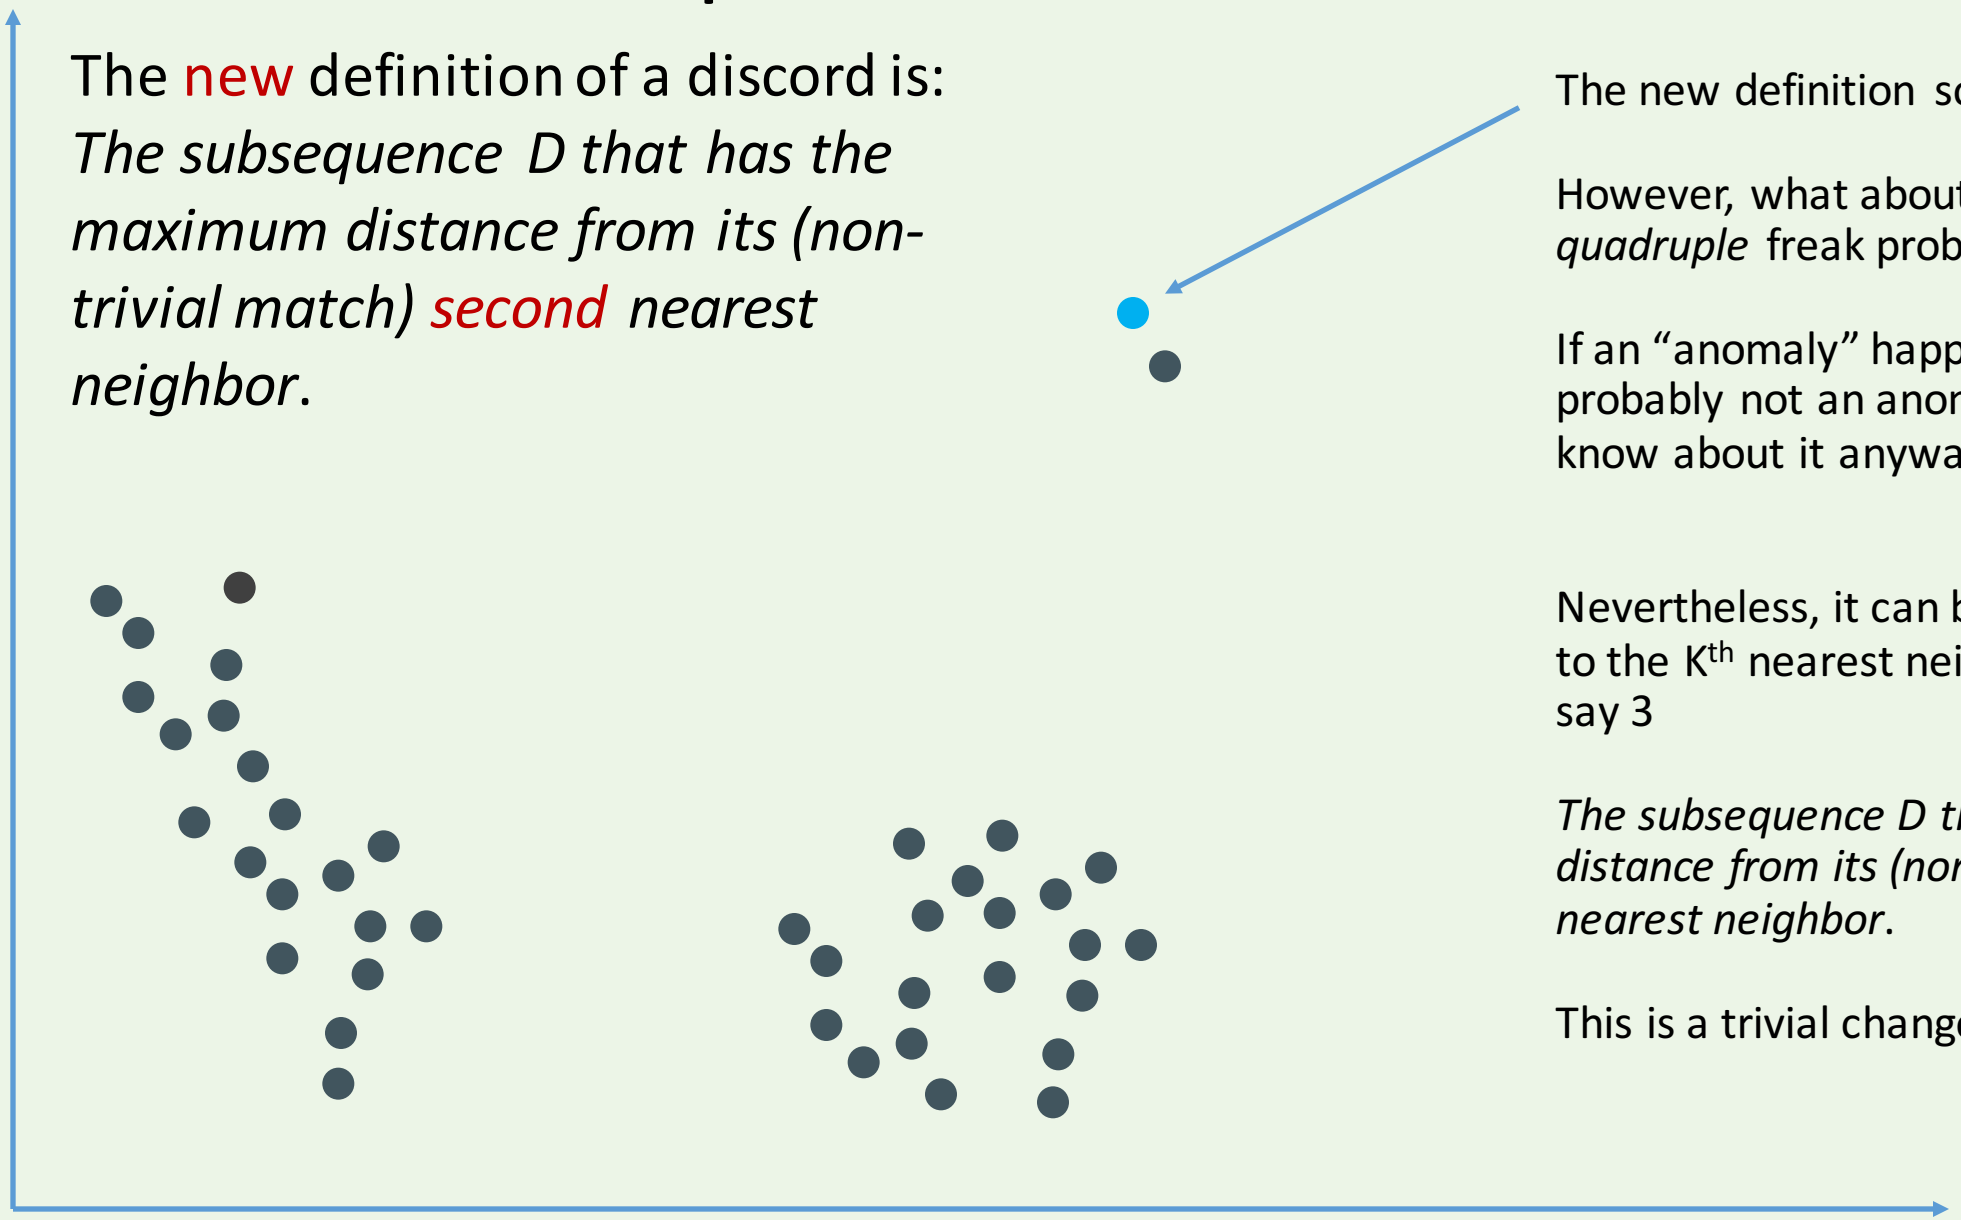

The new definition solves the problem.

However, what about the *triple* freak, or *quadruple* freak problem etc.…

If an "anomaly" happens many times, it is probably not an anomaly, and we probably know about it anyway.

Nevertheless, it can be useful to generalize to the K^{th} nearest neighbor, for a small K, say 3

The subsequence D that has the maximum distance from its (non-trivial match) K nearest neighbor.

This is a trivial change/addition to the MP

We have already seen examples of time series discords (although we did not explicitly call them that) so we will not revisit this here.

Discords are simply high values in the Matrix Profile.

There *are* many other algorithms to find discords. But why bother with them, when the Matrix Profile gives you them for free?

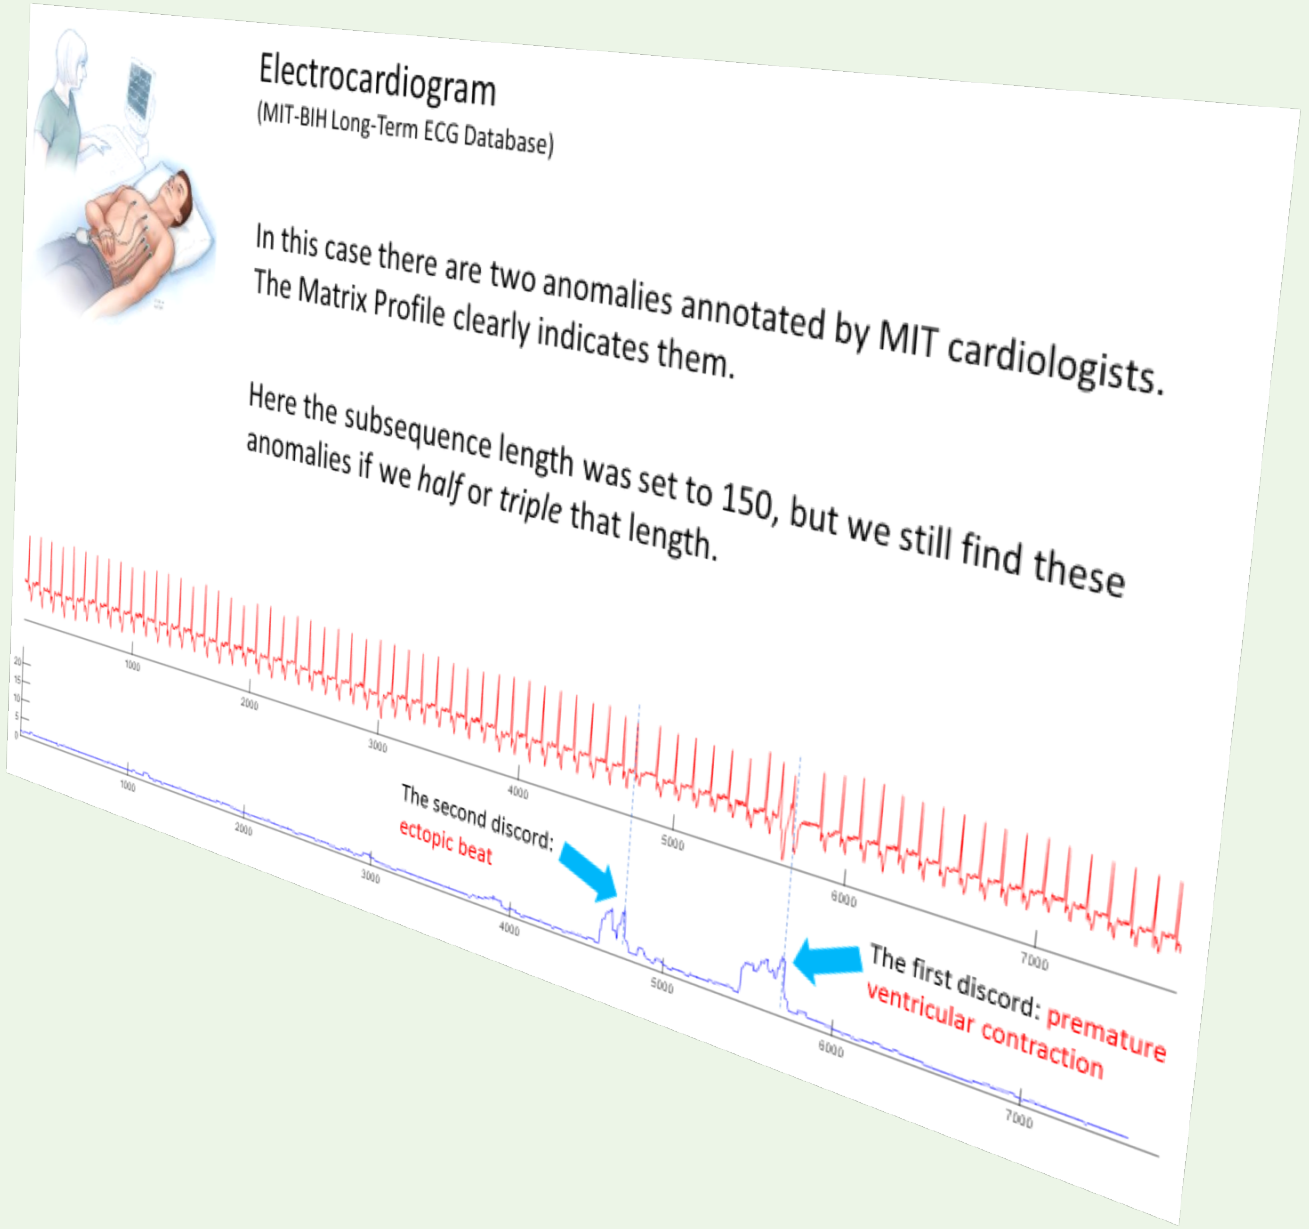

Generalizing to Joins

- We can think of the MP as a type of similarity self-join. For every subsequence in T_A, we join it with its nearest (non- trivial) neighbor in T_A, or $\mathbf{J}_{T_A T_A}$ *or* $T_A \bowtie 1\text{nn}\ T_A$
- This is also known as all-pairs-similarity-search (or similarity join).
- However, we can genialize to an **AB-**similarity join. For every subsequence in **A**, we join can it with its nearest neighbor in **B**, or $\mathbf{J}_{T_A T_B}$ *or* $T_A \bowtie 1\text{nn}\ T_B$
- Note that in general: $\mathbf{J}_{T_A T_B} \neq \mathbf{J}_{T_B T_A}$
- Note that **A** and **B** can be radically different sizes
- We may be interested in:
  - What is *conserved* between two time series (the join motifs)
  - What is *different* between two time series (the join discords)
- The tricks for understanding and reading join-based MPs are the same as before, we will see some examples to make that clear.

Generalizing to Joins

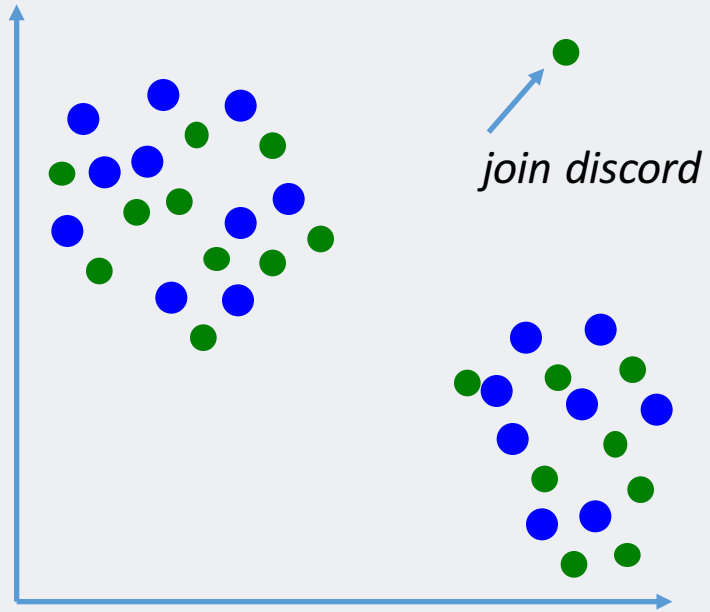

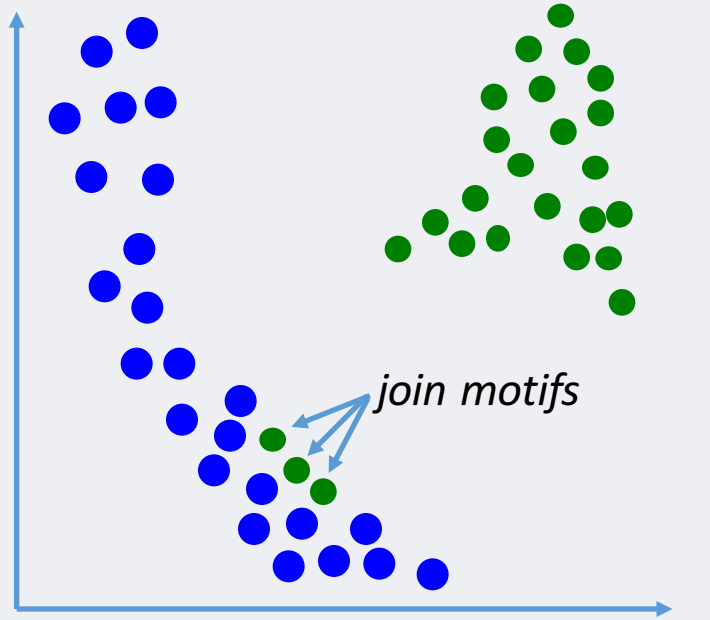

Two scenarios of interest: we do a $\mathbf{J}_{T_A T_B}$…

1) **The Golden Batch**: Here we have two time series that we think should be about the same. But when we join them, there is a *join discord*, a subsequence that appears only in only in **A**, but not in **B**, but why? (*spoken word* example below)

2) **The Suspicious Similarity:** Here we have two time series that we have *no* reason to think should be the same. But when we join them, there are *join motifs*, some subsequences from **A** appear in **B**, but why? (*music* example below. Another example would result from "meter swapping")

Now let us consider the join of two time series.

Assume we have two time series T_A and T_B ...

Note that they can be of different lengths

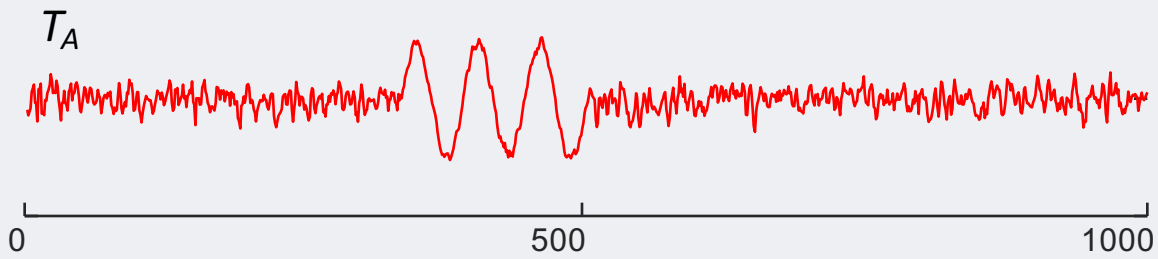

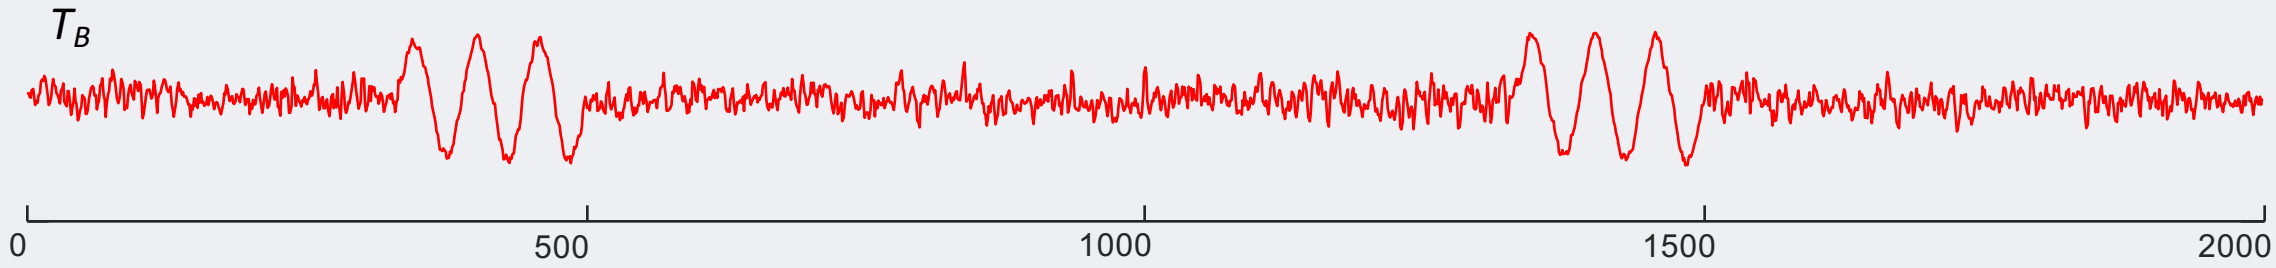

$| T_A | = 1{,}000$

$| T_B | = 2{,}000$

As before, we are not interested in any *global* properties of the time series, we are only interested in small *local* subsequences, of this length, *m*

These subsequences might be about the length of individual heartbeats (for ECGs), individual days (for social media behavior), individual words (for speech analysis) etc.

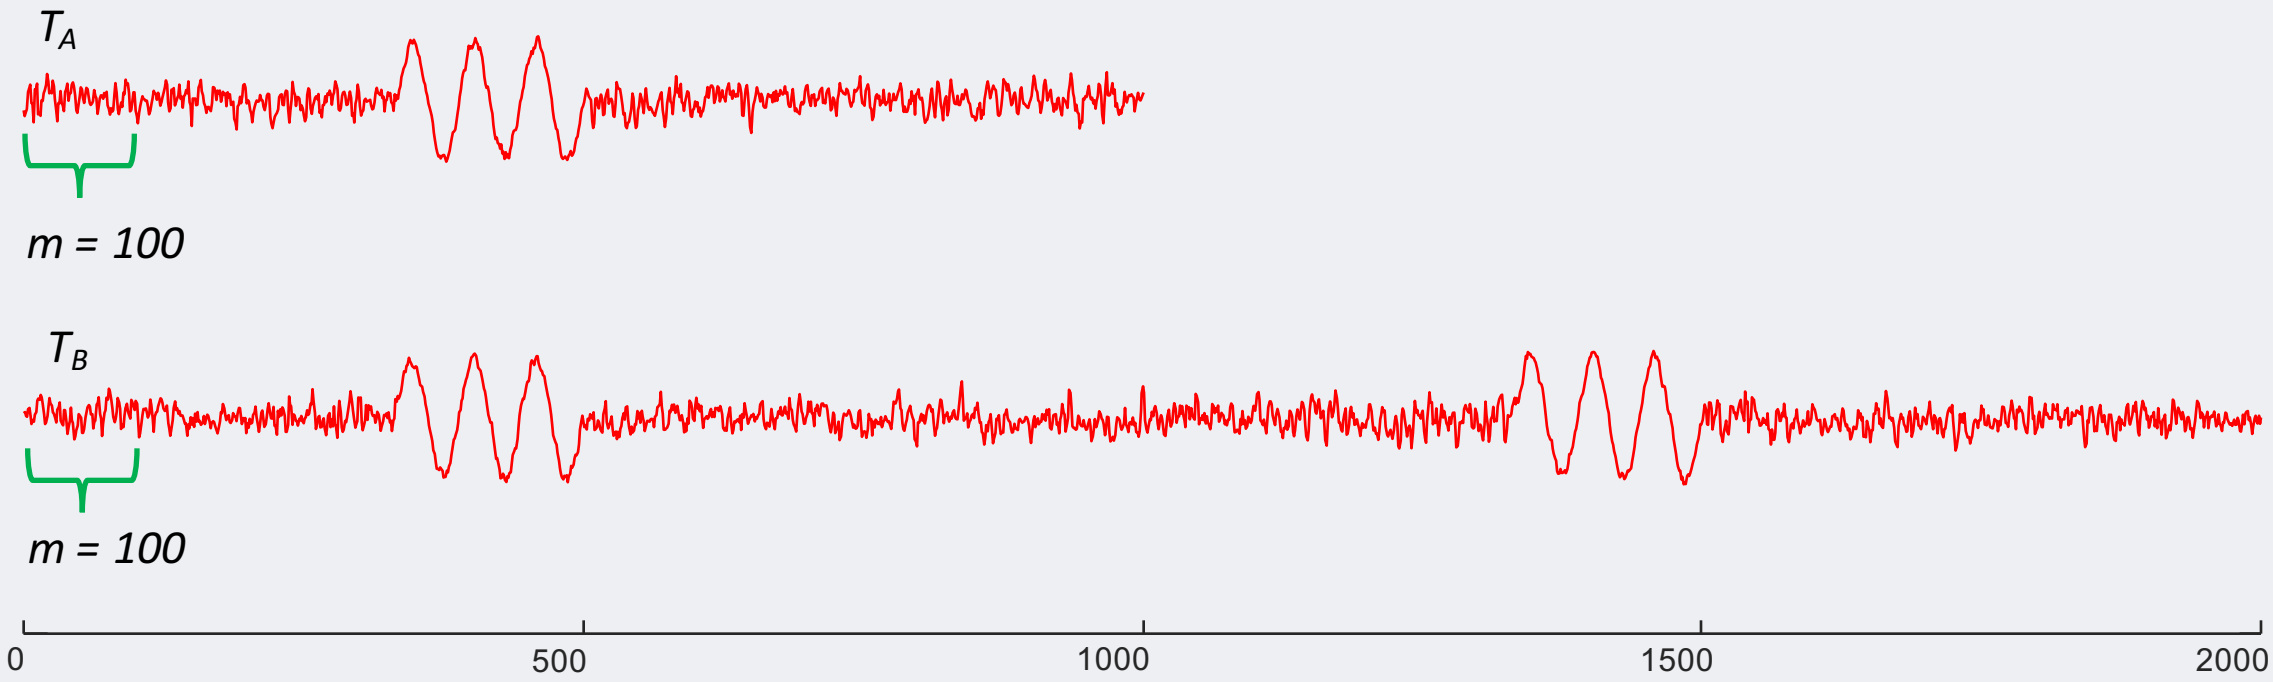

We can create a companion Matrix Profile of T_A.

For every subsequence in T_A, we look for its nearest neighbor in T_B.

The Matrix Profile at the i^{th} location records the distance of the subsequence in T_A, at the i^{th} location, to its nearest neighbor in T_B, under z-normalized Euclidean Distance.

The Matrix Profile is almost the same length as T_A, it is shorter by just m

For example, in the below, the subsequence of length 100 starting at 362 happens to have a distance of 1.24 to its nearest neighbor (wherever it is) in T_B.

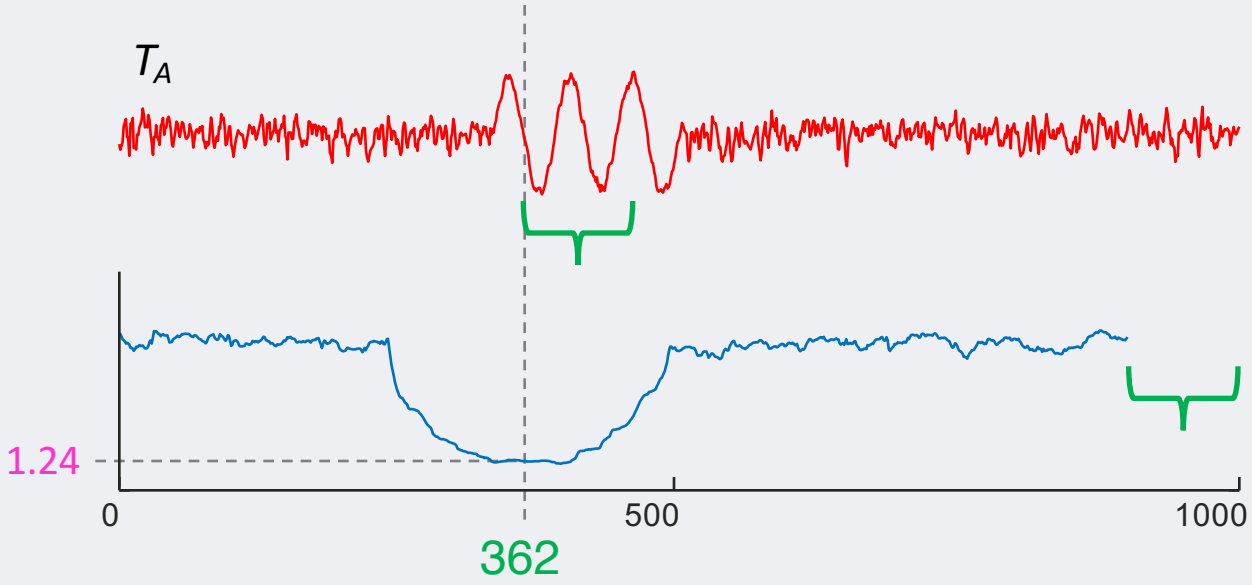

Informally: how far is each subsequence in T_A, from its nearest neighbor in T_b?

Recall that the Matrix Profile at the i^{th} location records the distance of the subsequence in T_A, at the i^{th} location, to its nearest neighbor in T_B, under z-normalized Euclidean Distance.

However, it does not tell us where the *location* of the nearest neighbor in T_B. To store this information, we can create another companion sequence, called a matrix profile index.

The green arrow points from the subsequence of T_A starting at 362 to its nearest neighbor in T_B. The nearest neighbor locates at 359 of T_B.

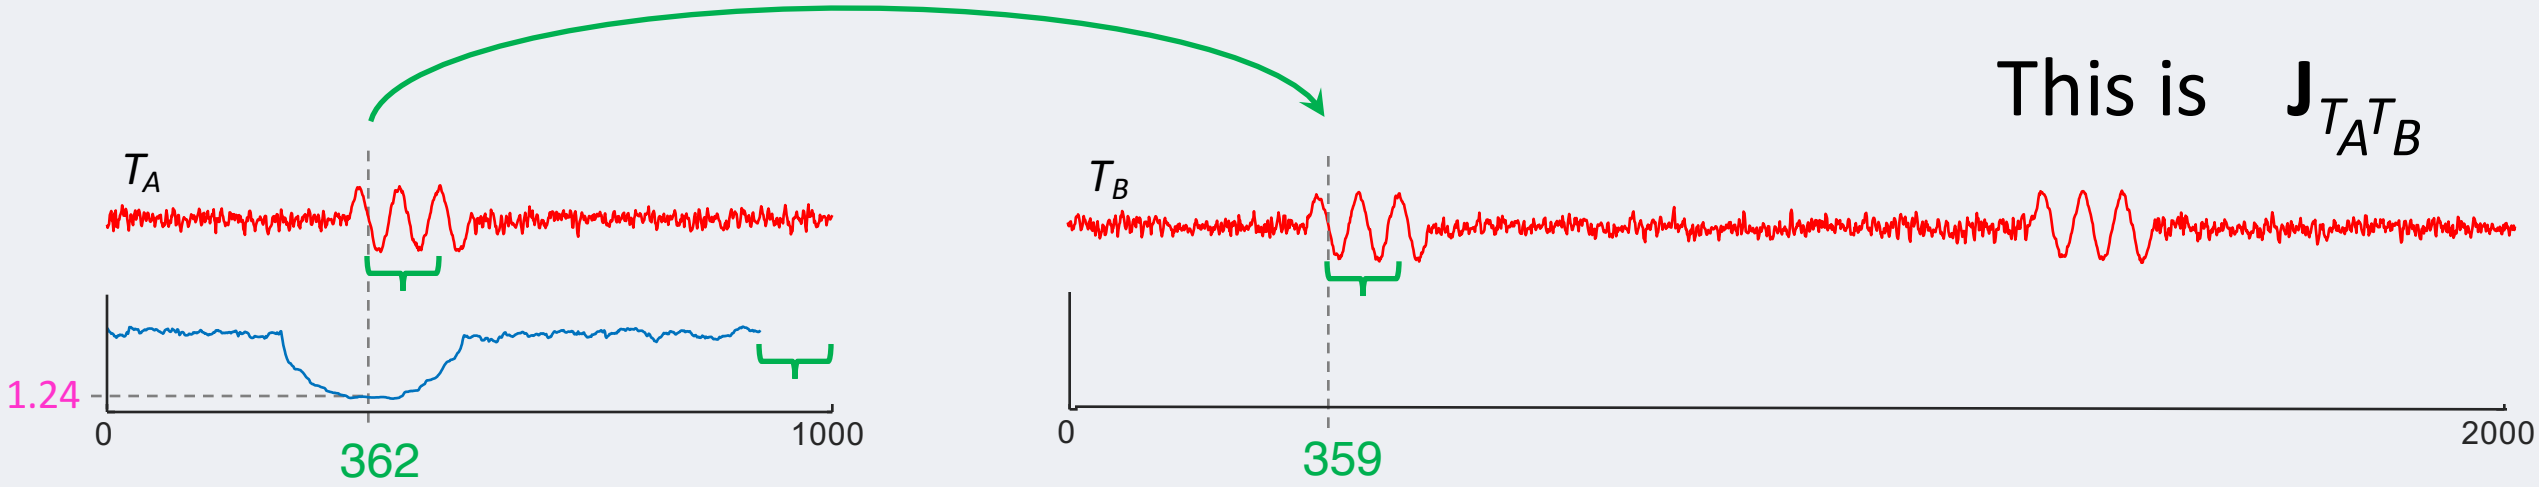

357	359	1401	...

matrix profile index (zoom in)

Informally: For each subsequence in T_A, point to its nearest neighbor in T_b

We can reverse the direction of the join…

Here the matrix profile index tell us the *location* of the nearest neighbor of each subsequence in T_B.

The green arrow points from the subsequence of T_B starting at 1340 to its nearest neighbor in T_A. The nearest neighbor locates at 395 of T_A.

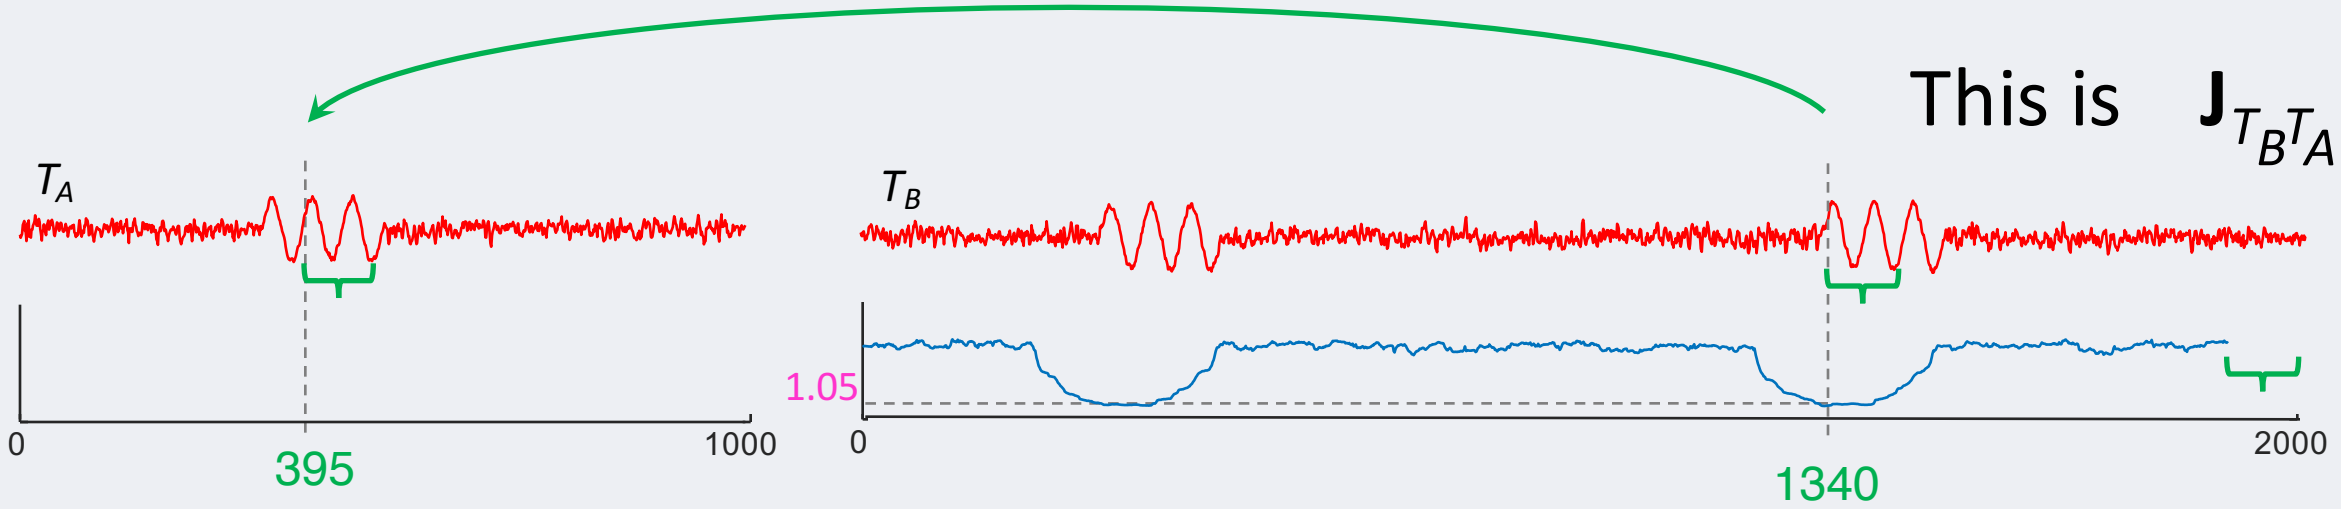

matrix profile index
(zoom in)

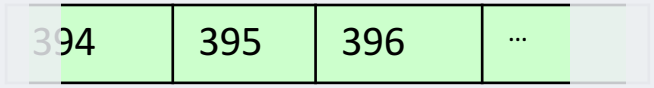

Music I (join case)

Can you see any common structure between the two time series below?

Hint, it is probably about this length

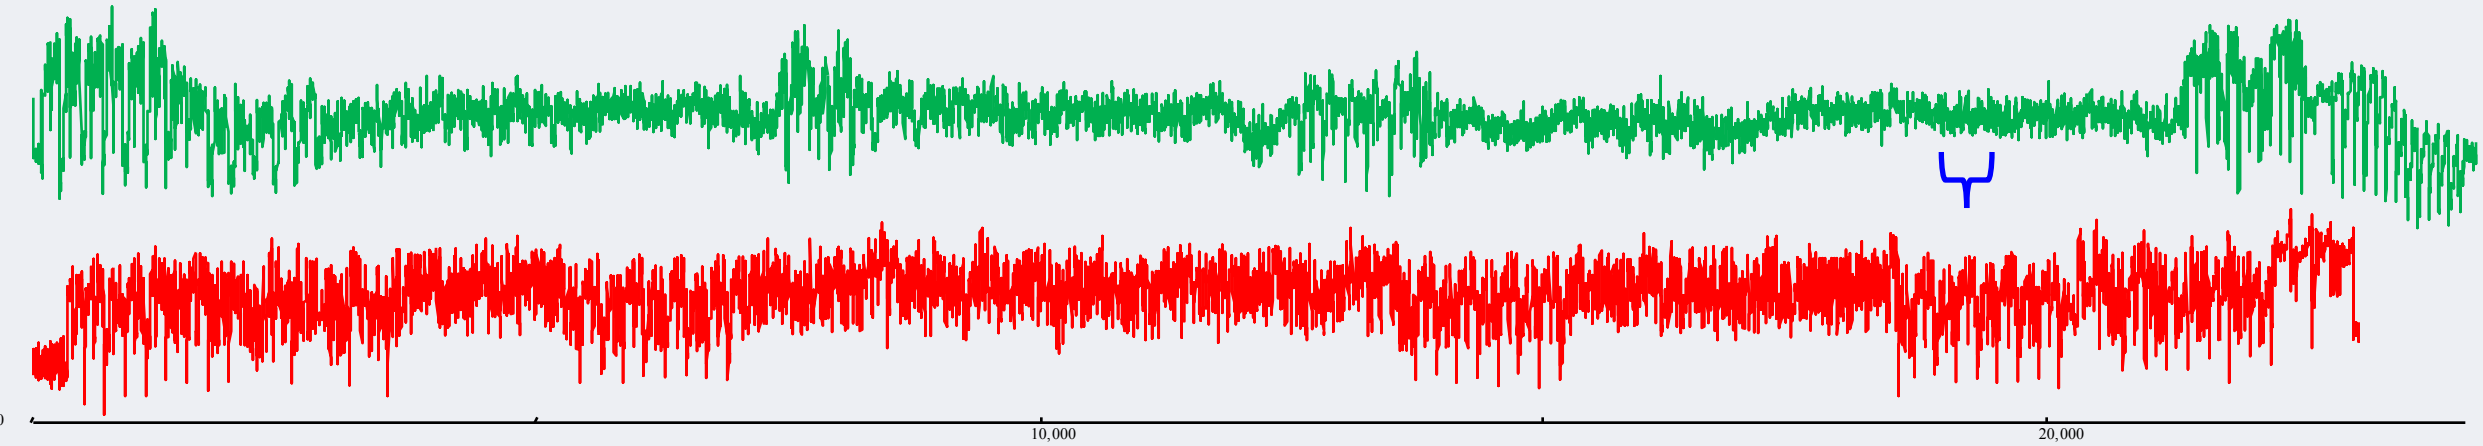

Music II (join case)

The data is the 2nd MFCC of two songs, *Under Pressure* and *Ice Ice Baby*

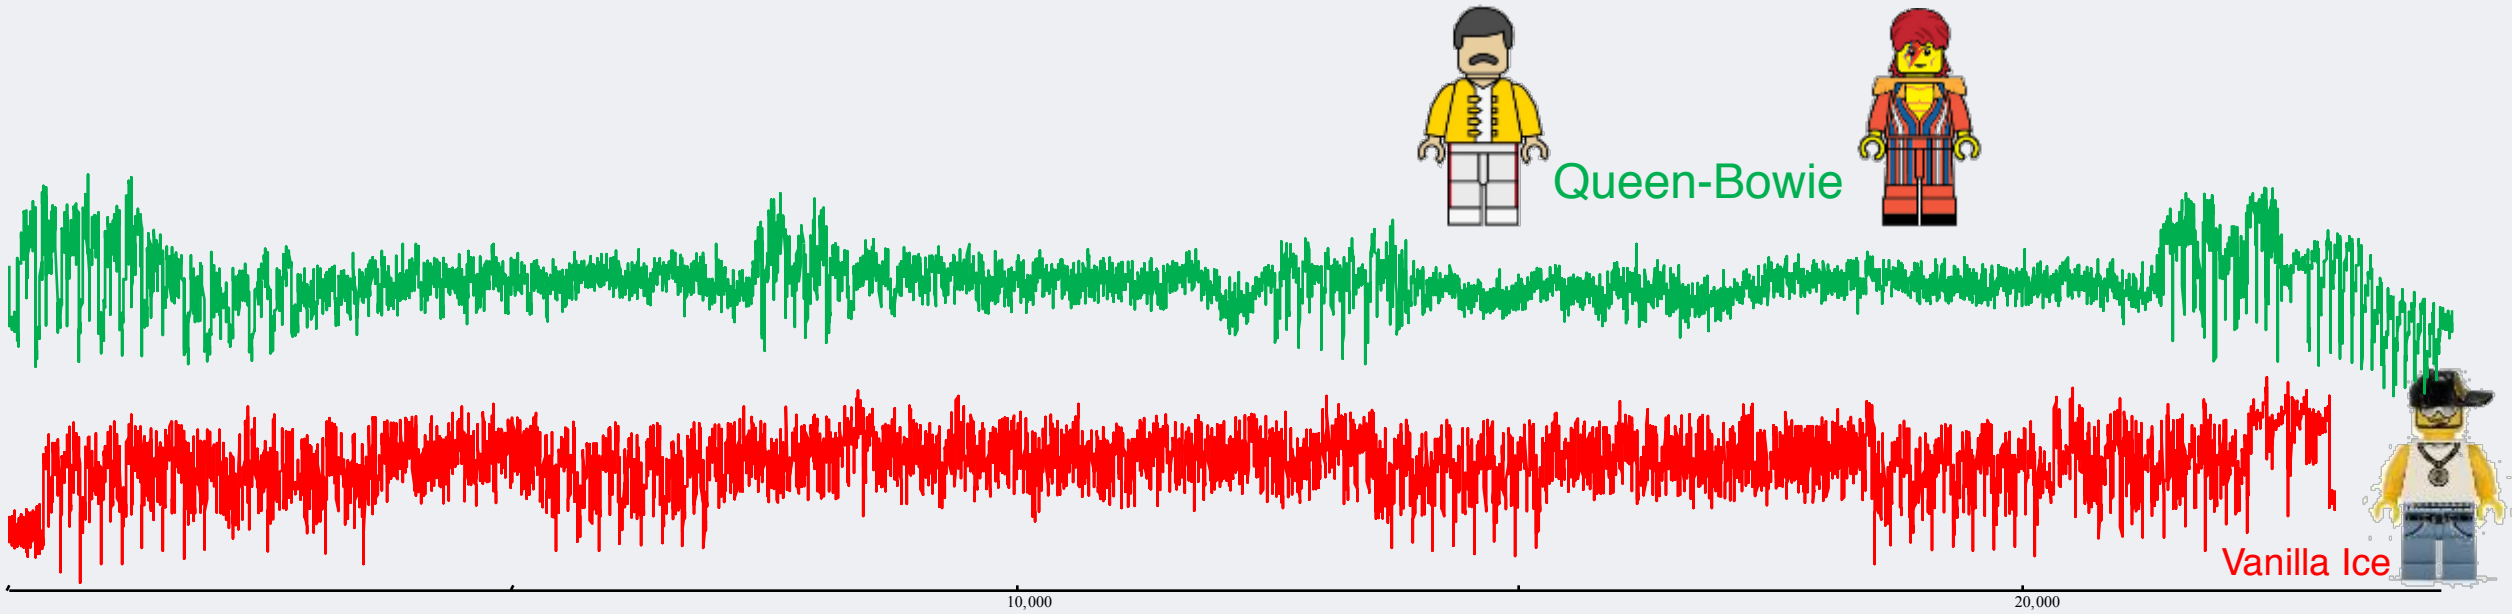

A zoom-in of the best conserved region between the two time series (the similarity join)

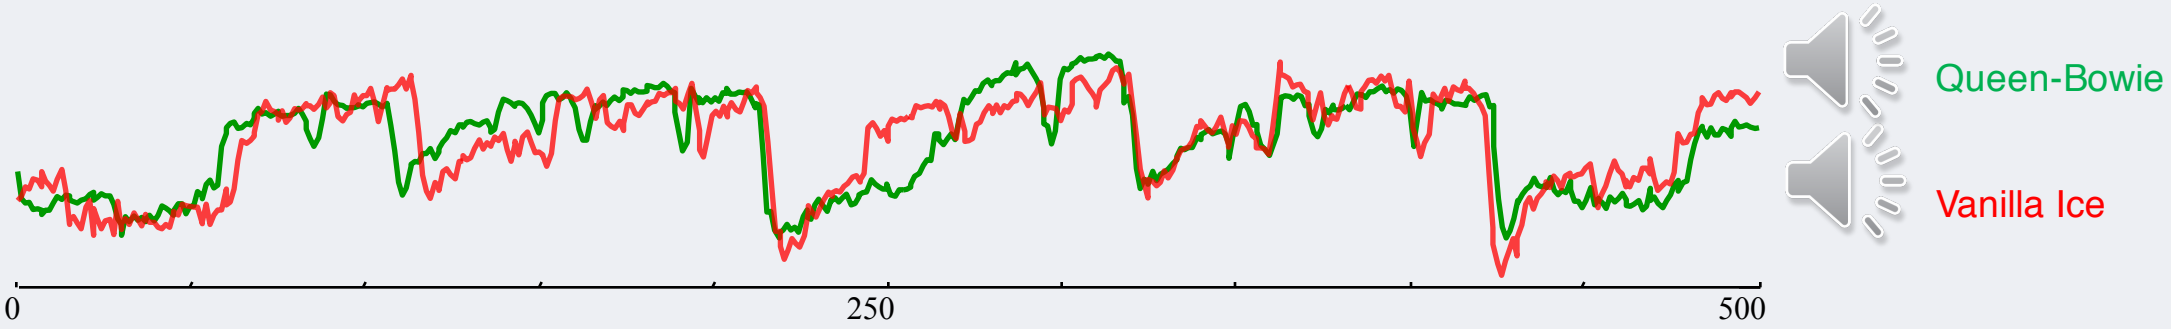

SPOKEN WORD

In the previous example we asked you to find "*common structure between the two time series*" Now I am going to ask you the opposite question.
What is *different* between the two time series?

Hint, it is probably about this length

Hint, it cant be the regions in the matching boxes, since they have matches…

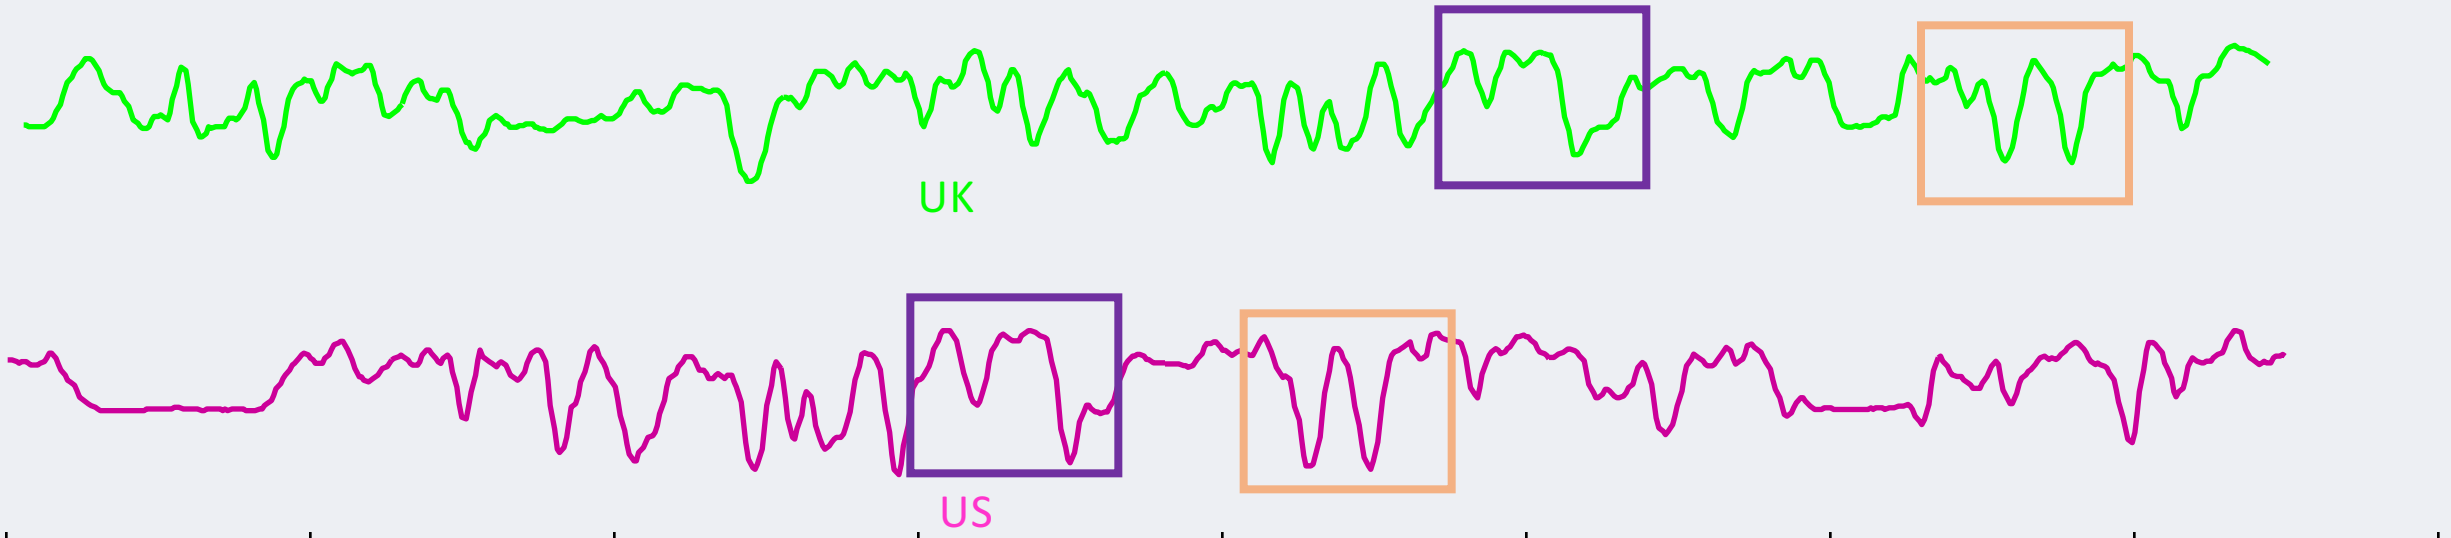

SPOKEN WORD II

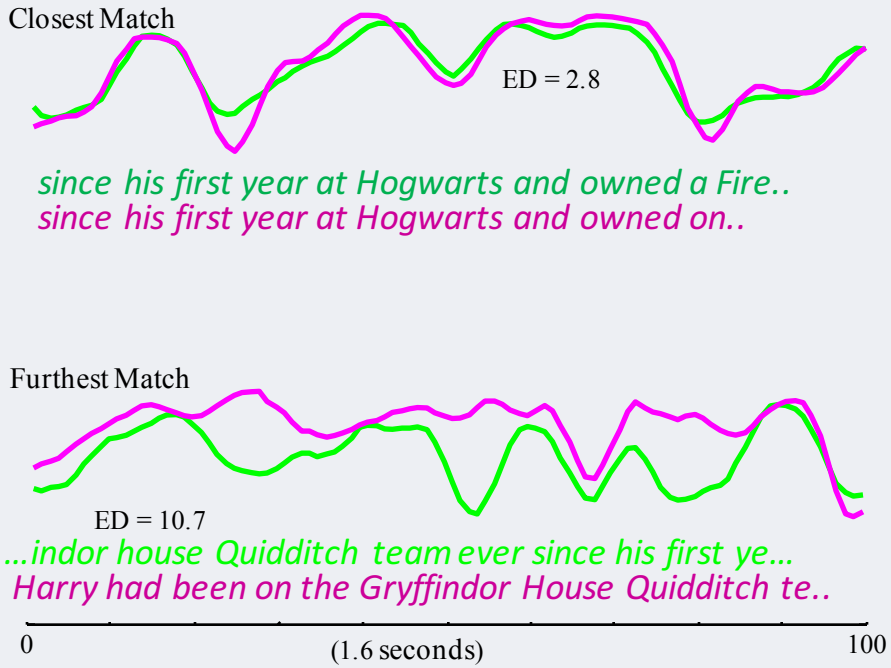

Here the difference is due to a unique phrase that only appears in the USA version of the Harry Potter books.

UK version : *Harry was passionate about Quidditch. He had played as Seeker on the Gryffindor house Quidditch team ever since his first year at Hogwarts and owned a Firebolt, one of the best racing brooms in the world...*

USA version : ***Harry had been on the Gryffindor House Quidditch team** ever since his first year at Hogwarts and owned one of the best racing brooms in the world, a Firebolt.*

DNA (join case)

We consider two strains of Legionella bacteria, *L. pneumophila Paris* and *L. pneumophila Lens*, which consist of 3,503,504 and 3,345,567 bp respectively. We consider a subsequence length of 100,000

However, we flipped one of the time series “backwards”, before computing the join.

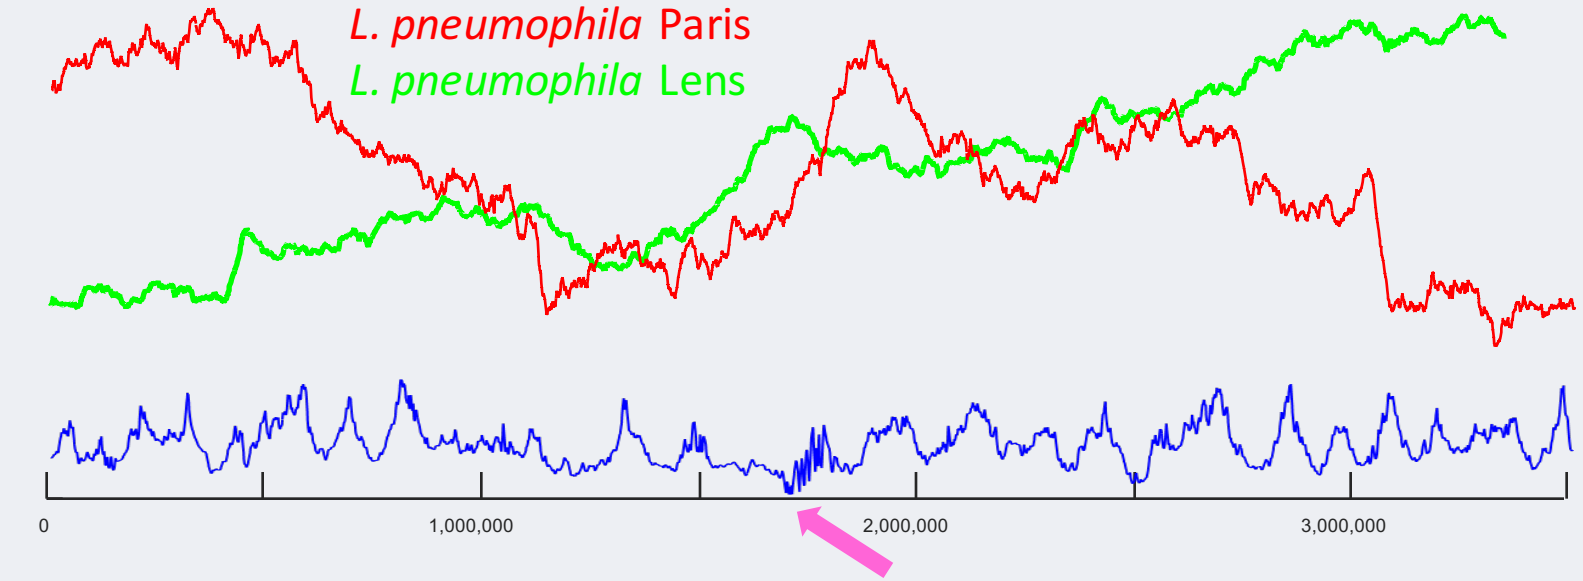

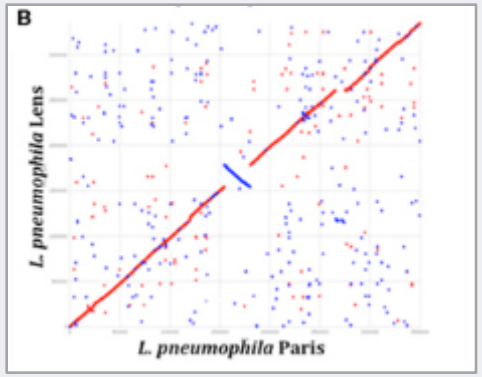

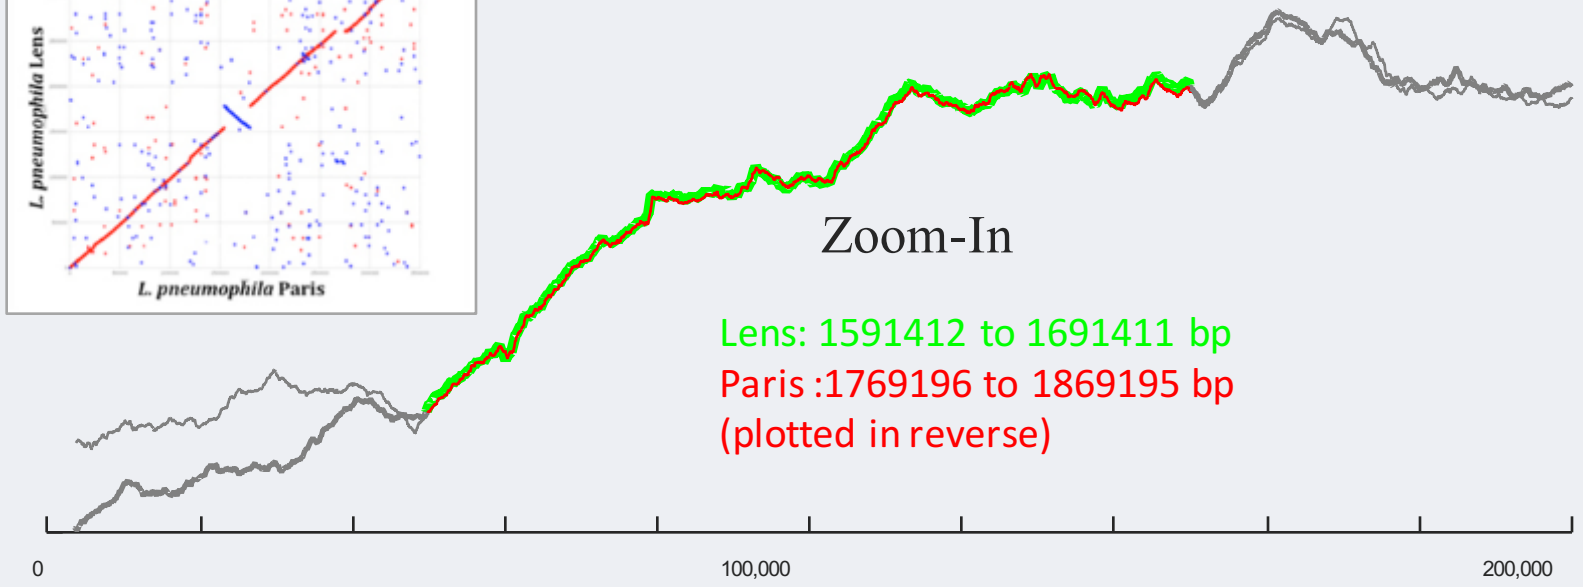

Laura Gomez-Valero et al. Comparative and functional genomics of *Legionella* identified eukaryotic like proteins as key players in host–pathogen interactions

Time Series Semantic Segmentation

Sometime the system we are monitoring changes regimes, can we detect such changes?

PigInternalBleedingDatasetArtPressureFluidFilled

..sedated pig, *has bleeding induced…*

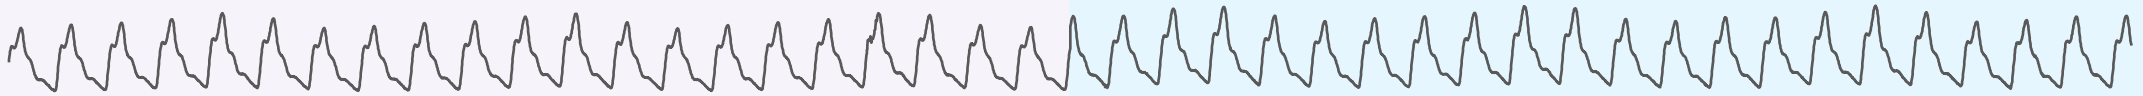

PulsusParadoxusSP02

..regular beat, *then Pulsus Paradoxus..*

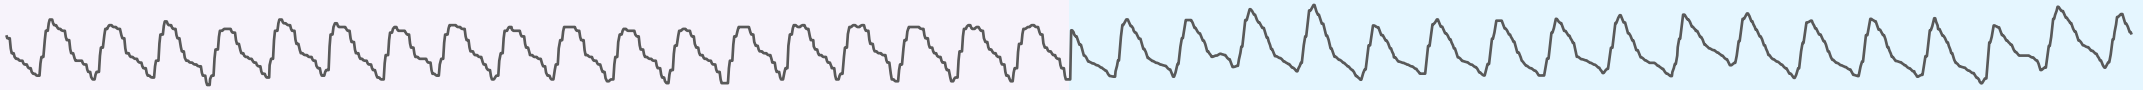

SuddenCardiacDeath2

very irregular beat, *then ventricular tachyarrhythmia*

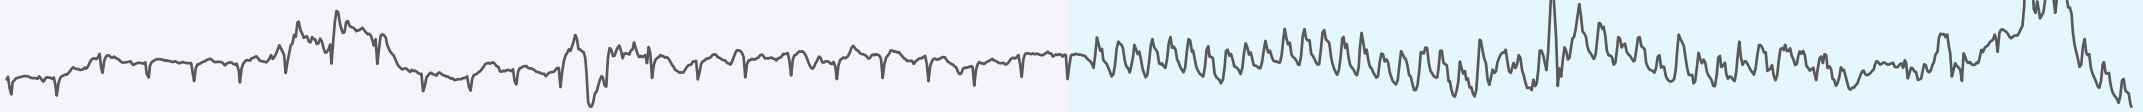

RoboticDogActivityX

..walking, *then playing…*

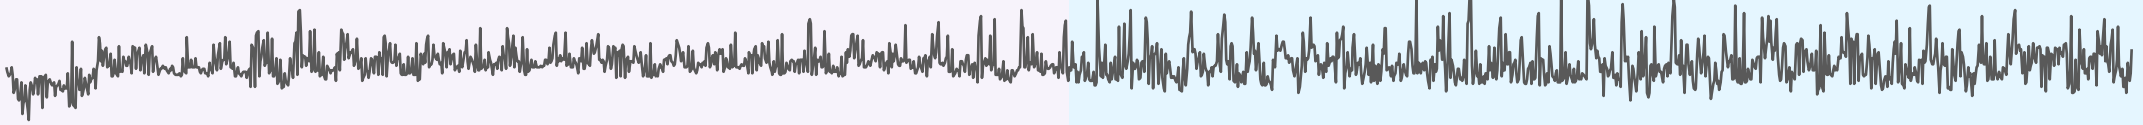

TiltECG

..lying horizontal, *titling begins …*

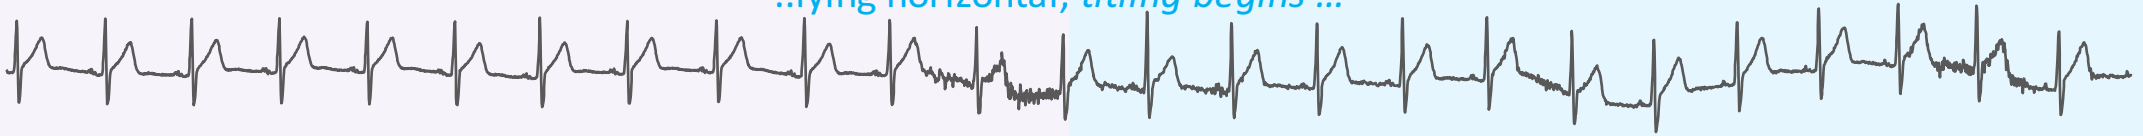

FLOSS: Matrix Profile Segmentation

What do we want in a Semantic Segmentation Algorithm?

- **Handle fast online, or huge batch data**
- **Domain agnosticism**. It would be nice to have a single algorithm the works on all kinds of data
- **Parameter-Lite**. (Tuning parameters almost guarantees overfitting).
- **High accuracy**.
- **Able to report *degree* of segmentation or confidence** (A binary decision is too brittle in most cases).

- **Claim:** We can do all this by looking at the Matrix Profile Index…

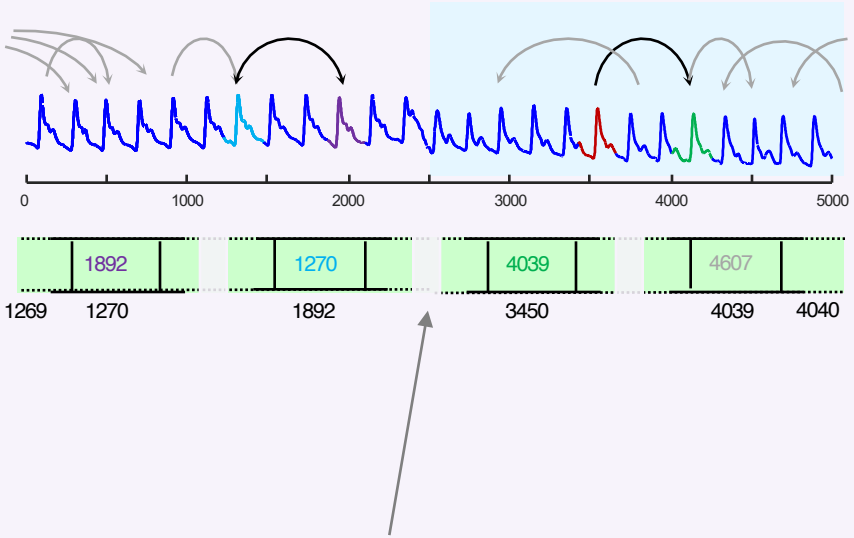

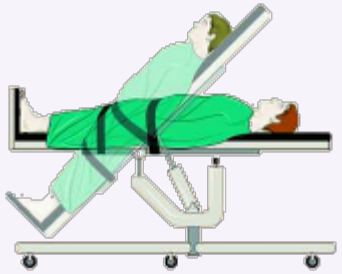

Key Observation

Recall that the Matrix Profile Index has pointers (*arrows*, *arcs*) that point to the nearest neighbor of each subsequence.

If we have a change of behavior, say from *walk* to *run*, we should expect very few arrows to span that change of behavior.

- Most *walk* subsequences will point to another *walk* subsequence
- Most *run* subsequences will point to another *run* subsequence
- Rarely, if ever, will a *run* point to a *walk*.

So, if we slide across the Matrix Profile Index, and count how many arrows cross each particular point, we expect to find few that span the change of behavior.

Lets try this, next slide…

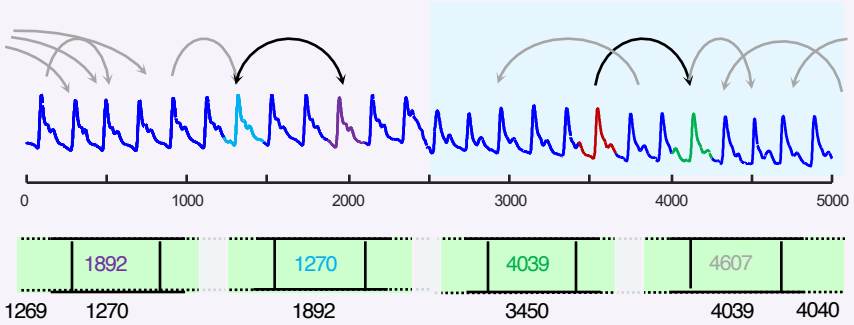

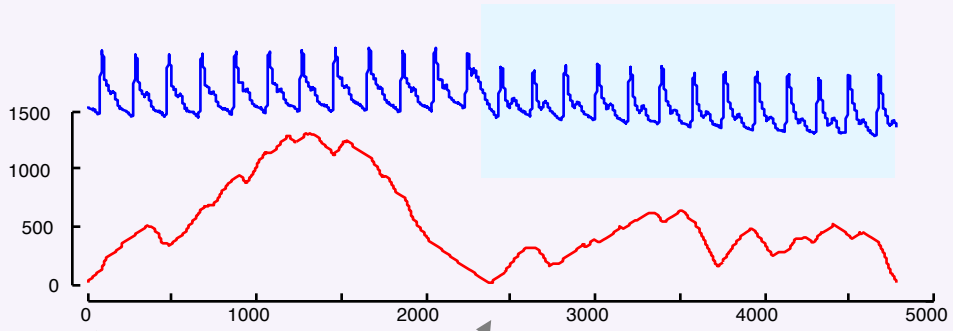

The arc count here is almost zero!

But the beginning (and end) are also near zero.

This works!

If we use the sliding arc count to produce an arc-curve, we find it is near zero at the point of system transition. This low value signals the location of the system change.

There is one flaw. The arc-curve, tends to be low near the beginning and end of the time series, just because there are fewer arcs that *could* cross at those locations.

What we can do is calculate what the arc-curve would look like if there was *no* system transition, and use that to *correct* the arc-curve.

If there was no system transition structure, the arc-curve would be a inverted parabola, with a height ½ the time series length.

Lets try this, next slide…

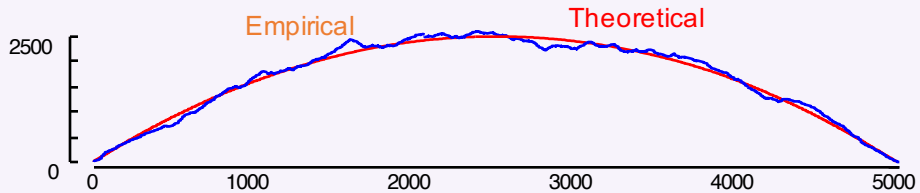

The number of arcs that cross a given index, if the links are assigned randomly

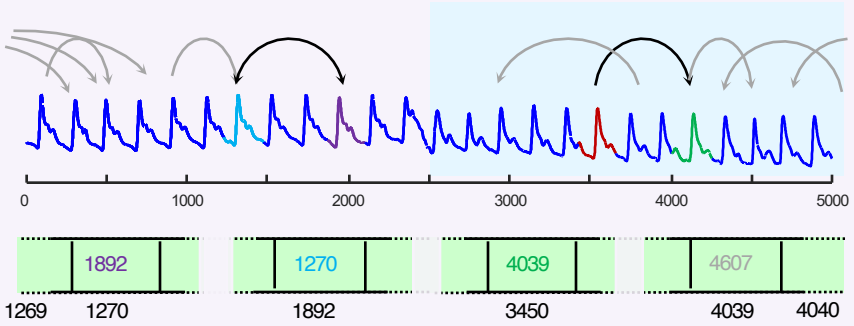

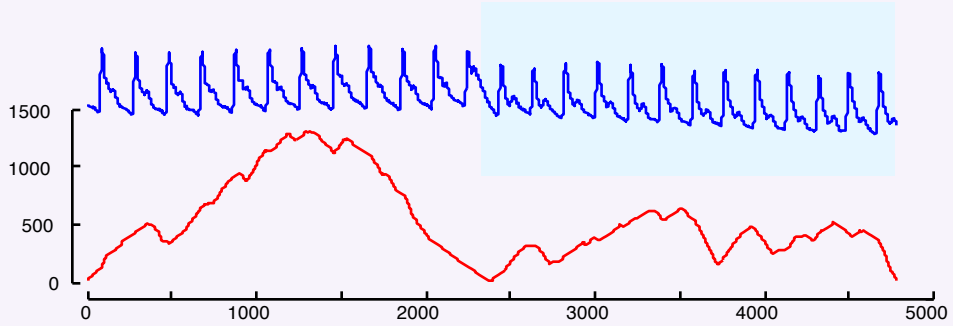

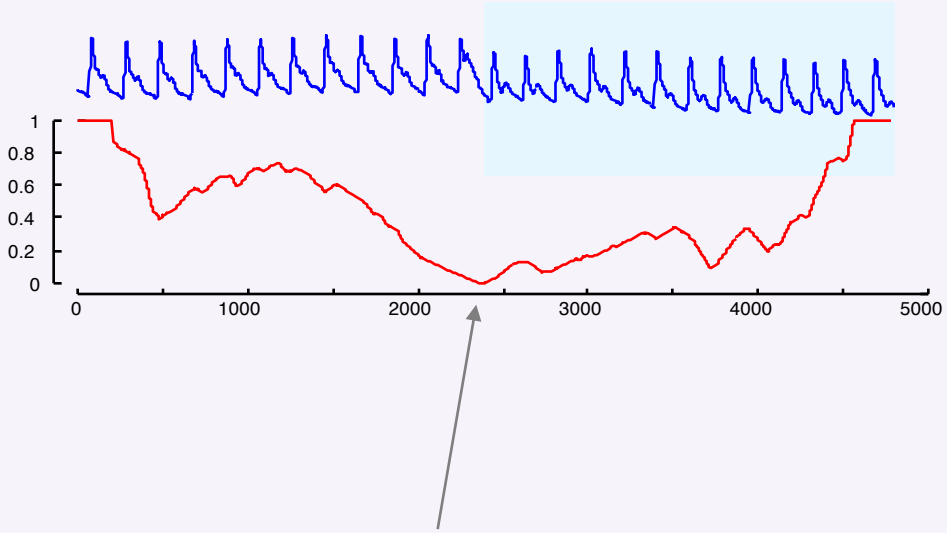

The corrected arc-curve here is almost zero

This works even better!

The corrected arc-curve minimizes in the right place, and does not have spurious minimizations at the beginning and end.

How robust is the corrected arc-curve? Lets add some distortions to the data, and see what it does.

Lets try this, next slide…

FLOSS is very robust to the data's properties

Consider the following distortions

- Downsampling from the original 250 Hz to 125 Hz (red).
- Reducing the bit depth from 64-bit to 8-bit (blue).
- Adding a linear trend of ten degrees (cyan).
- Adding 20dB of white noise (black).
- Smoothing, with MATLAB's default settings (pink).
- Randomly deleting 3% of the data, and filling it back in with simple linear interpolation (green).

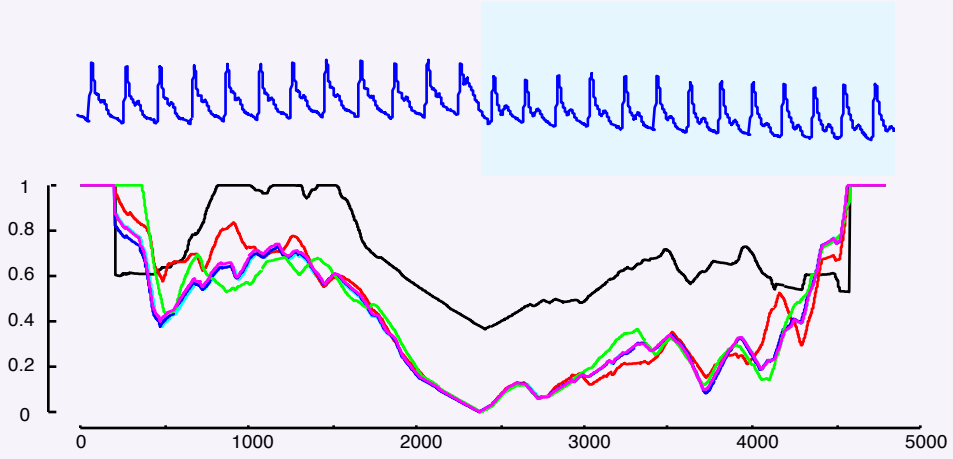

Most distortions make almost no difference. The only one that does move the CAC significantly was adding noise. But even then we **still** find the correct segmentation, and we added a *lot* of noise

We added a *lot* of noise

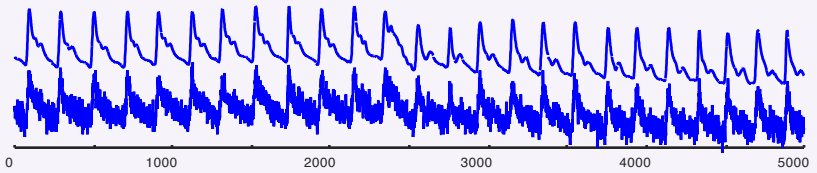

FLOSS is very robust to its only parameter

The CAC has a single parameter, the subsequence length *m*.

But we can typically change it by an order of magnitude, and get very good results.

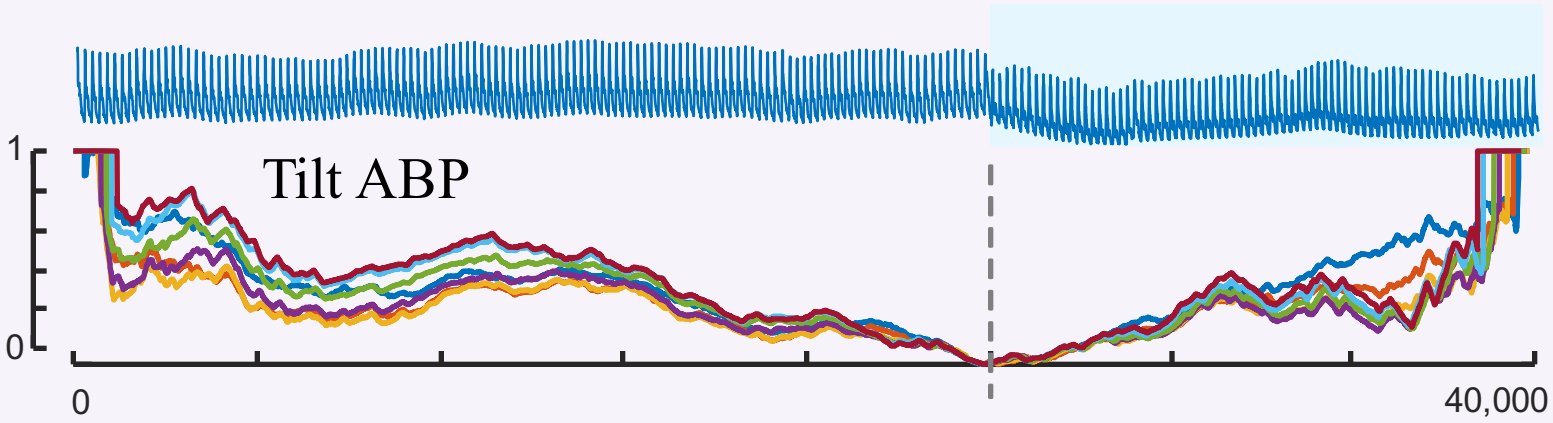

The CAC computed for:

(*top*) TiltABP with *m* = {100, 150, 200, 250, 300, 350, 400}

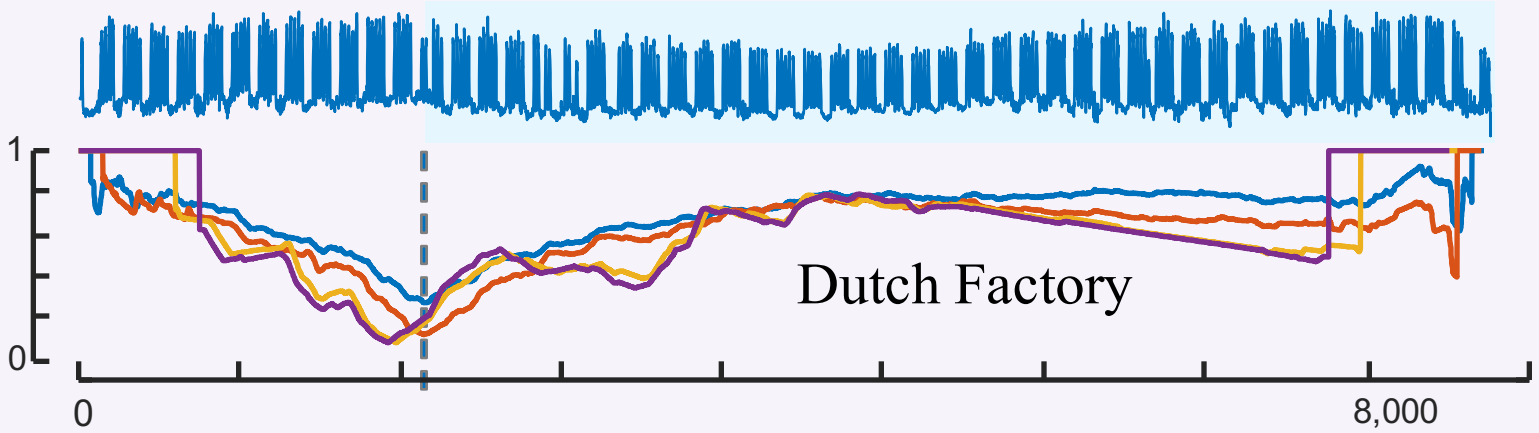

(*bottom*) DutchFactory for *m* = {25, 50, 200, 250}. Even for this huge range of values for *L*, the output of FLOSS is essentially unchanged.

Great Barbet

(Psilopogon virens)

One individual Great Barbet sings…, *….another takes over…,* *…yet another takes over*

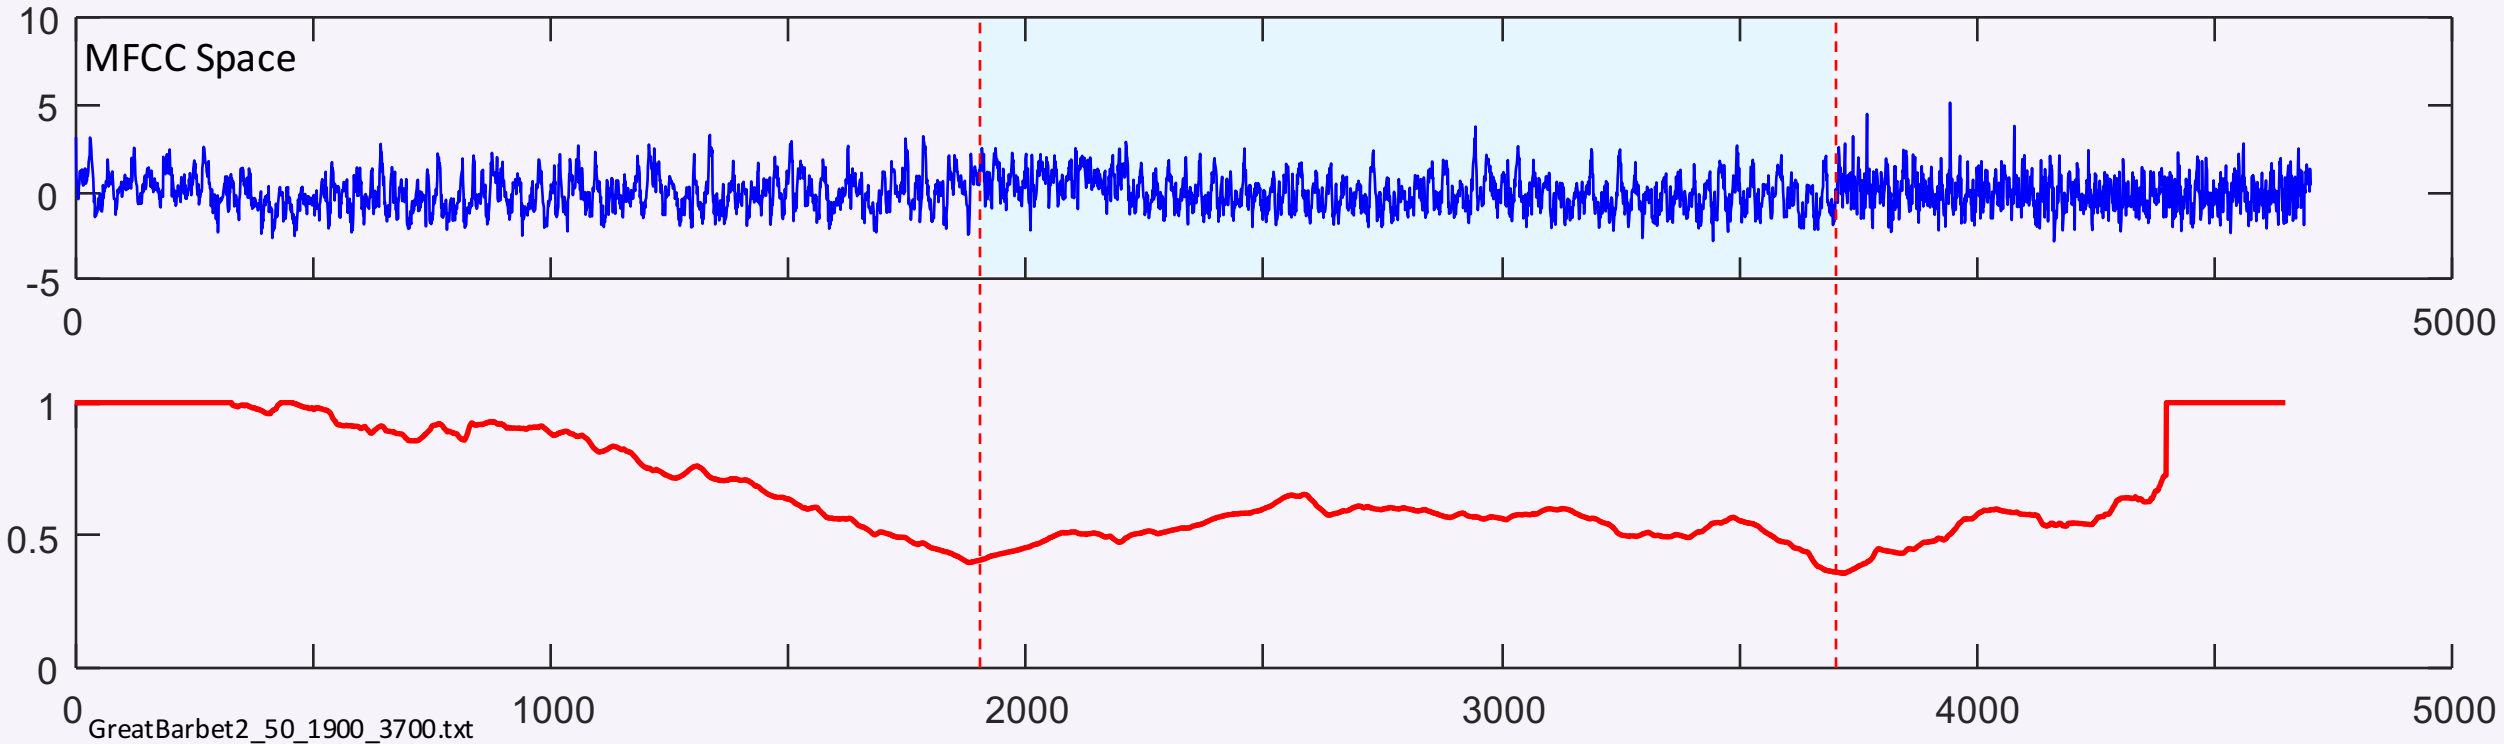

Asian citrus psyllid
(Diaphorina citri)

This dataset was hand annotated by an entomologist. The insect changes its feeding behavior at about time 1,800.

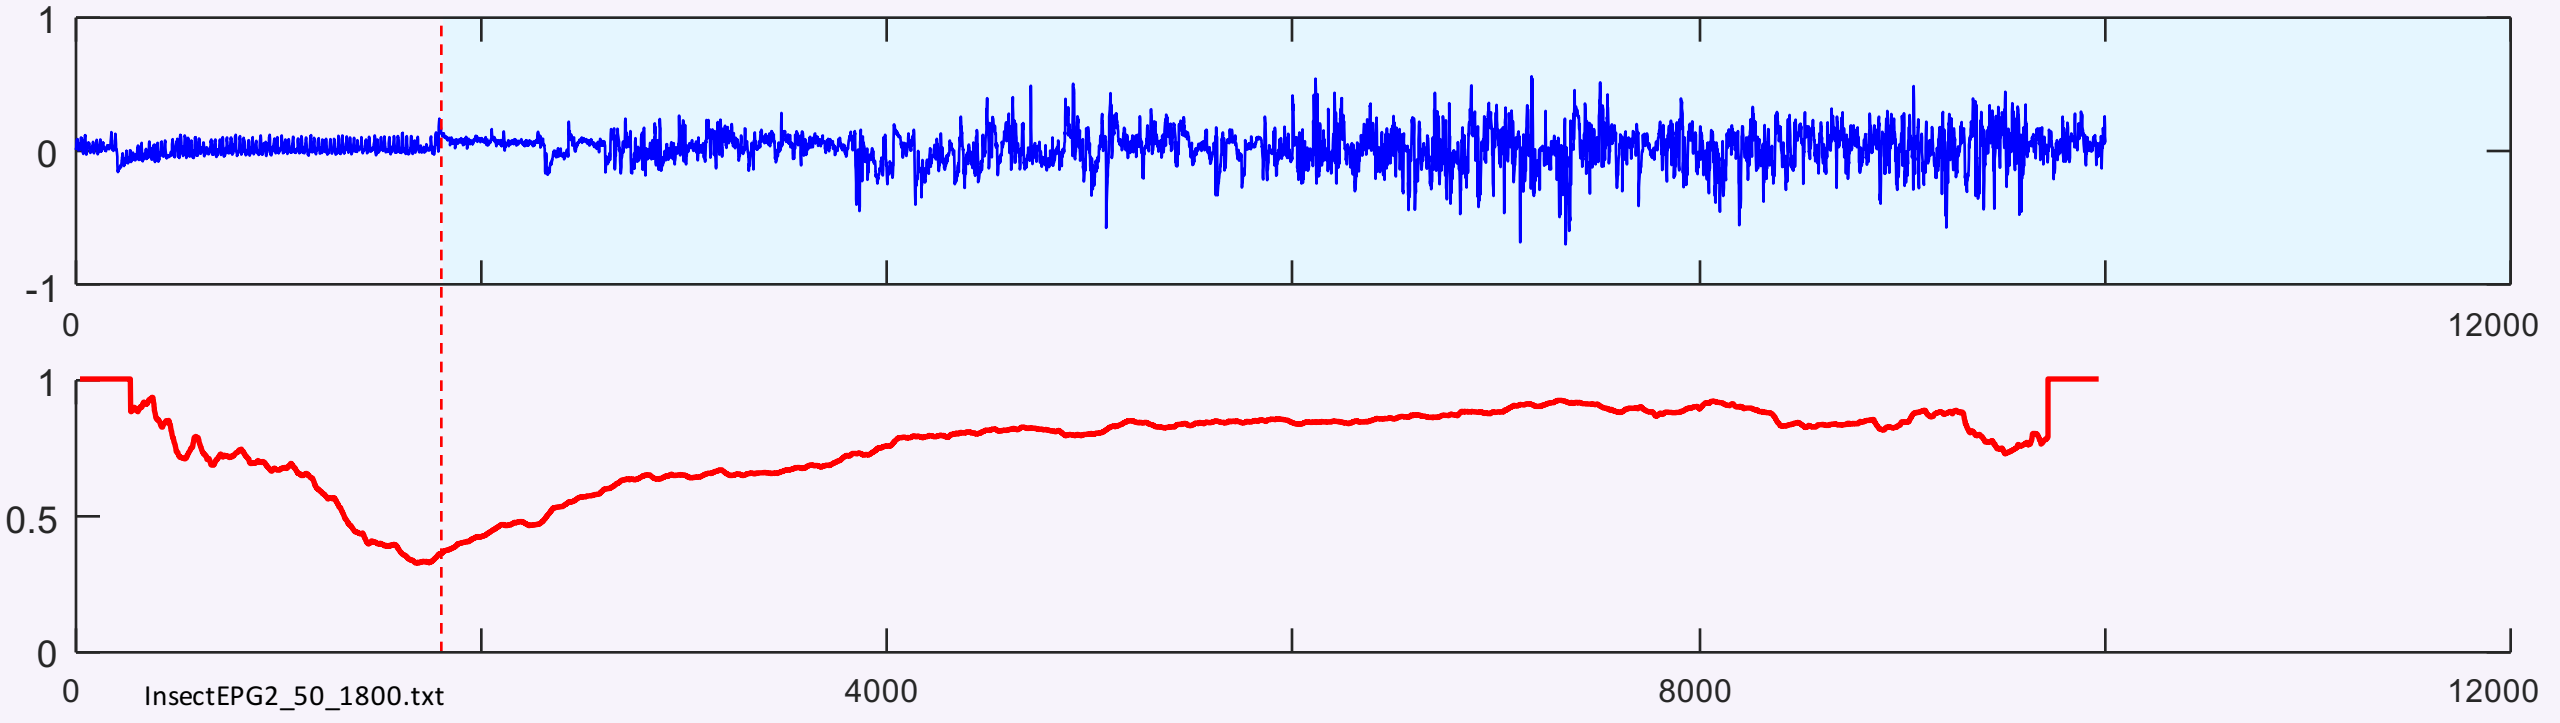

Pulsus Paradoxus is often visually apparent in the $SP0_2$ trace. Here we deliberately ignore this fact, and look *only* in the ECG trace, which is normally considered as *not* predictive of Pulsus Paradoxus.

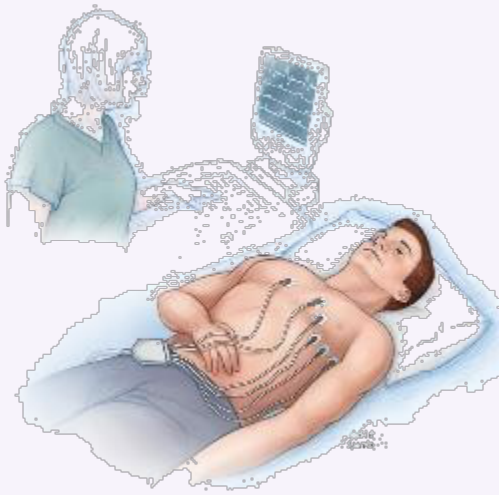

Note that the clinician that annotated this data was *in the room* at the time and may have had access to information that is simply not available in this signal.

Pulsus paradoxus (PP), also paradoxic pulse or paradoxical pulse, is an abnormally large decrease in systolic blood pressure and pulse wave amplitude during inspiration.
See also https://www.youtube.com/watch?v=7AXIYQK5BBM

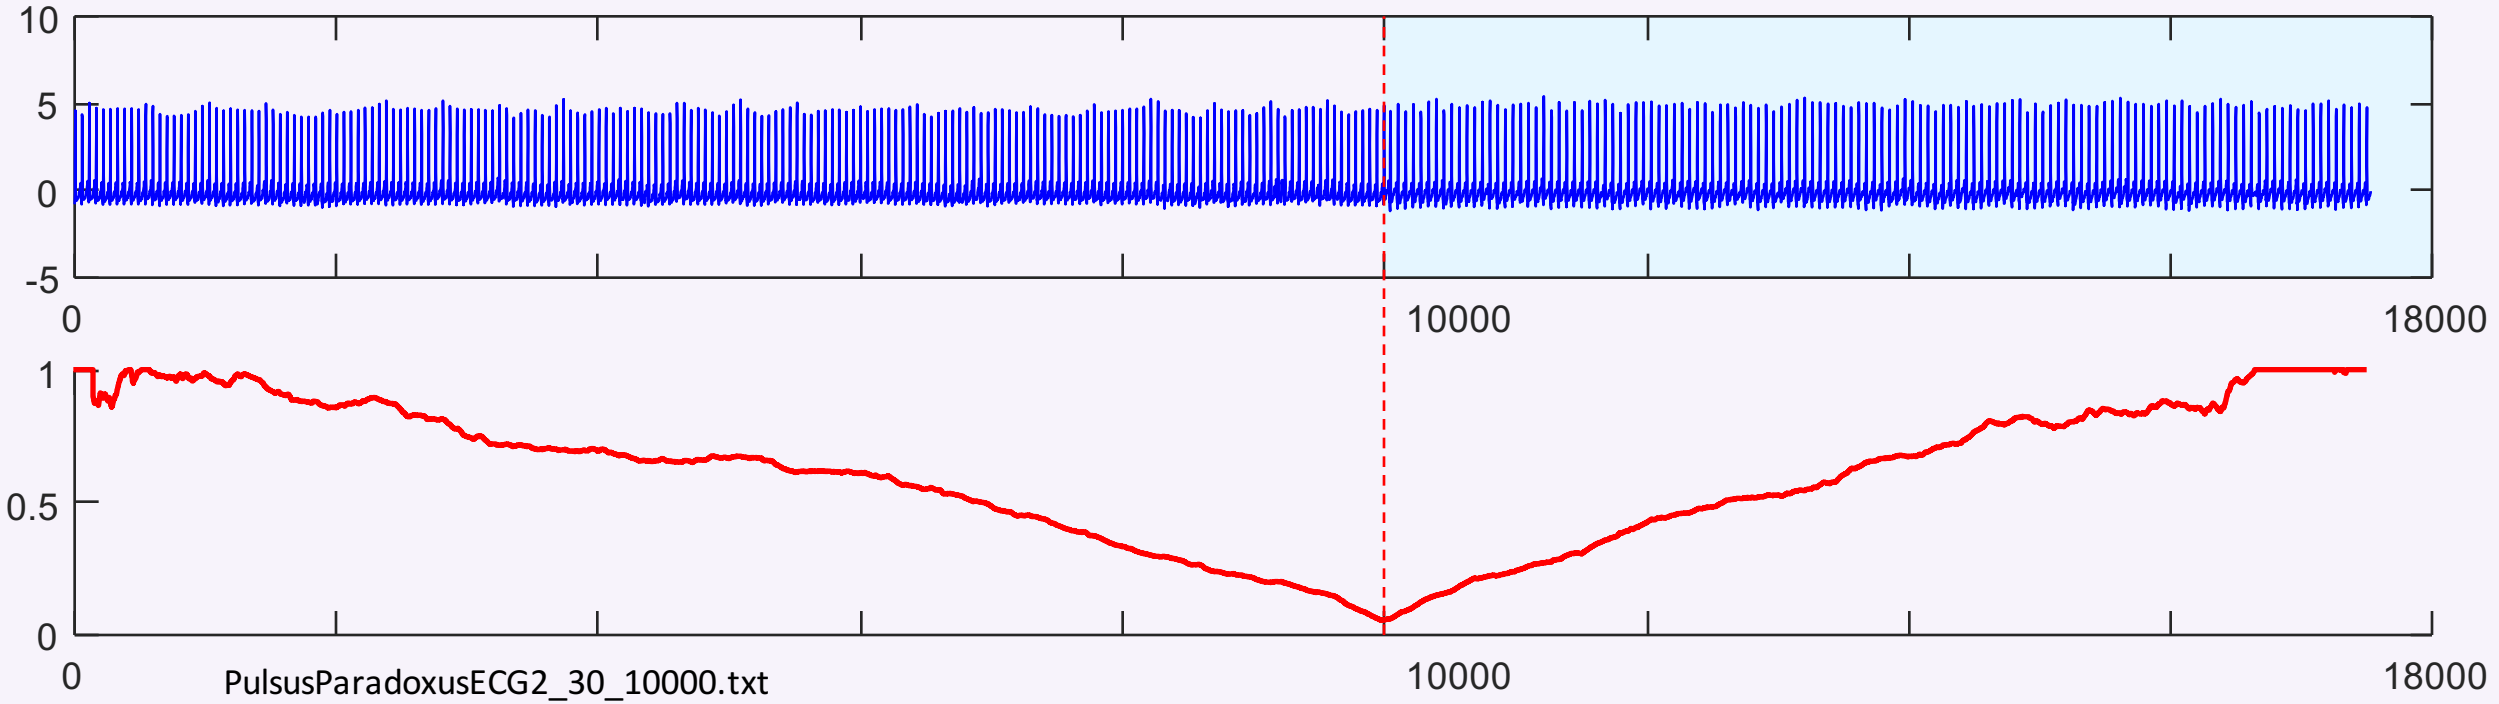

Summary for Time Series Segmentation

The Matrix Profile allows a simple algorithm, FLUSS/FLOSS, for time series segmentation.

Because it is built on the MP, it inherits all the MPs properties

- It is incrementally computable, i.e. *online* (This variant is called FLOSS)
- It is fast (at least 40 times faster than real-time for typical accelerometer data)
- It is domain agnostic (But you can use the AV to add domain knowledge, see below)
- It is parameter-lite (only one parameter, and it is not sensitive to its value)

- It has been tested on the largest and most diverse collection of time series ever considered for this problem, and in spite *of* (or perhaps, *because of*) its simplicity, it is state-of-the-art. Better than rival methods, and better than *humans* (details offline).

From Domain Agnostic to Domain Aware*

- The great strength of the MP is that is *domain agnostic*. A single black box algorithm works for taxi demand, seismology, DNA, power demand, heartbeats, music, bird vocalizations....
- However, in a handful of cases, there is a need to, or some utility in, incorporating some domain knowledge/constraints.
- There is a simple, generic and elegant way to do this, using the *Annotation Vector* (AV).
- In the following slides we will show you the *annotation vector* in the context of motif discovery, but you can use it with any MP algorithm.
- We will begin by showing you some examples of spurious motifs that can be discovered in particular domains, then we will show you how the AV mitigates them.

*Hoang Anh Dau and Eamonn Keogh. Matrix Profile V: A Generic Technique to Incorporate Domain Knowledge into Motif Discovery. KDD'17, Halifax, Canada.

Motivating Example 1:
Stop-word Motif Bias

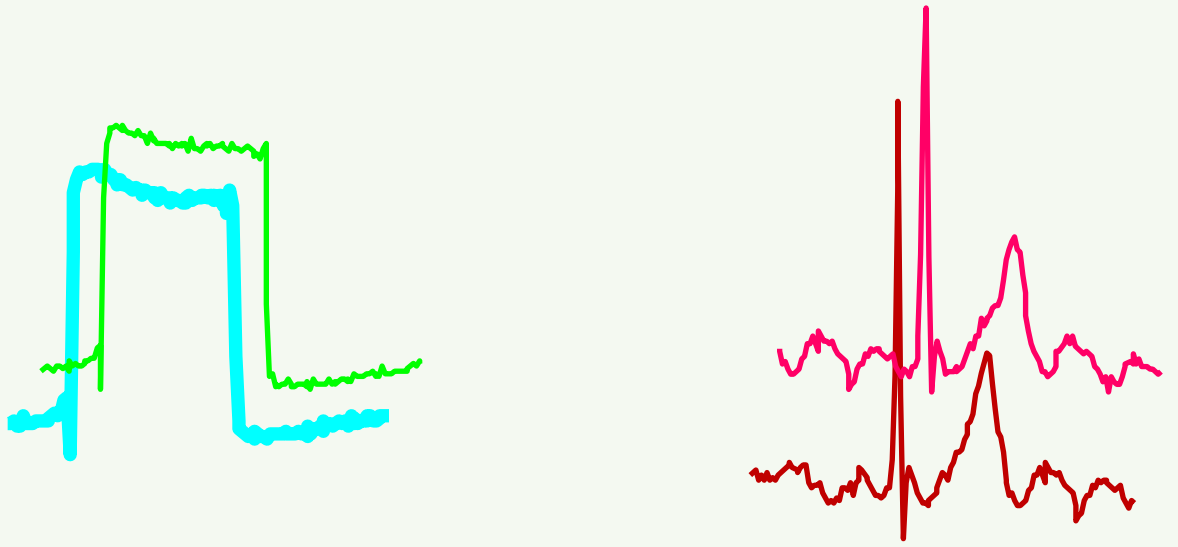

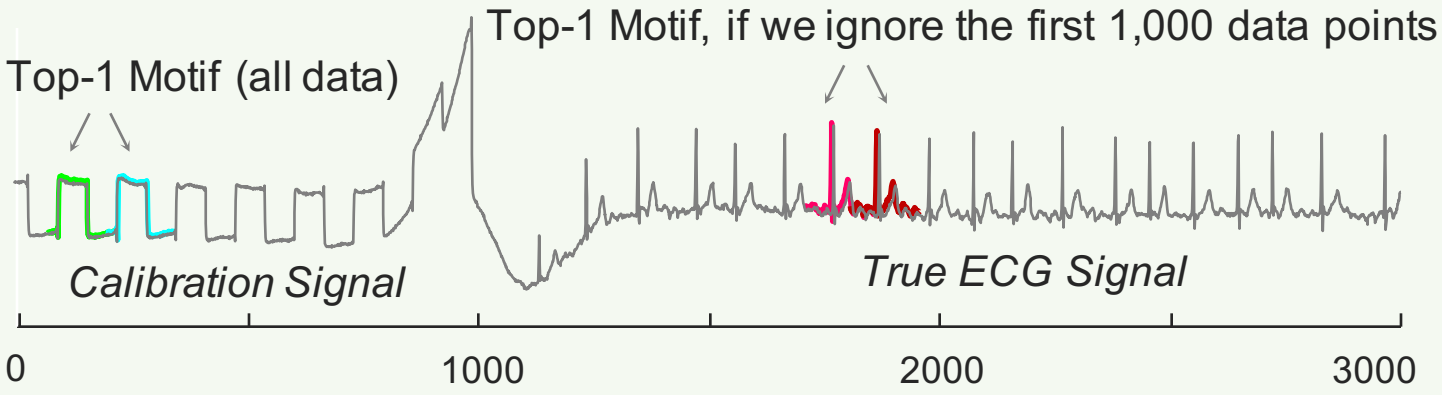

In some medical datasets, the true motifs may be "swamped" by more frequent, but biologically meaningless patterns. Much like the stop words "*and*" and "*the*" in text mining.

Here the approximate square wave is just a calibration signal, sent when the sensor has weak contact with the skin. It is a frequent, but spurious motif.

A snippet of ECG data from the LTAF-71 Database. The top motifs come from regions of the calibration signal because they are much more similar than the motifs discovered if we search only data that contains true ECGs.

Motivating Example 2: Simplicity Bias

Euclidean Distance has a bias toward simple shapes.

"*Pairs of complex objects, even those which subjectively may seem very similar to the human eye, tend to be further apart under* (Euclidean) *distance than pairs of simple objects.*" [1]

Surprisingly, the top-motif does not correspond to the motion artifacts, but to simple regions of "drift"

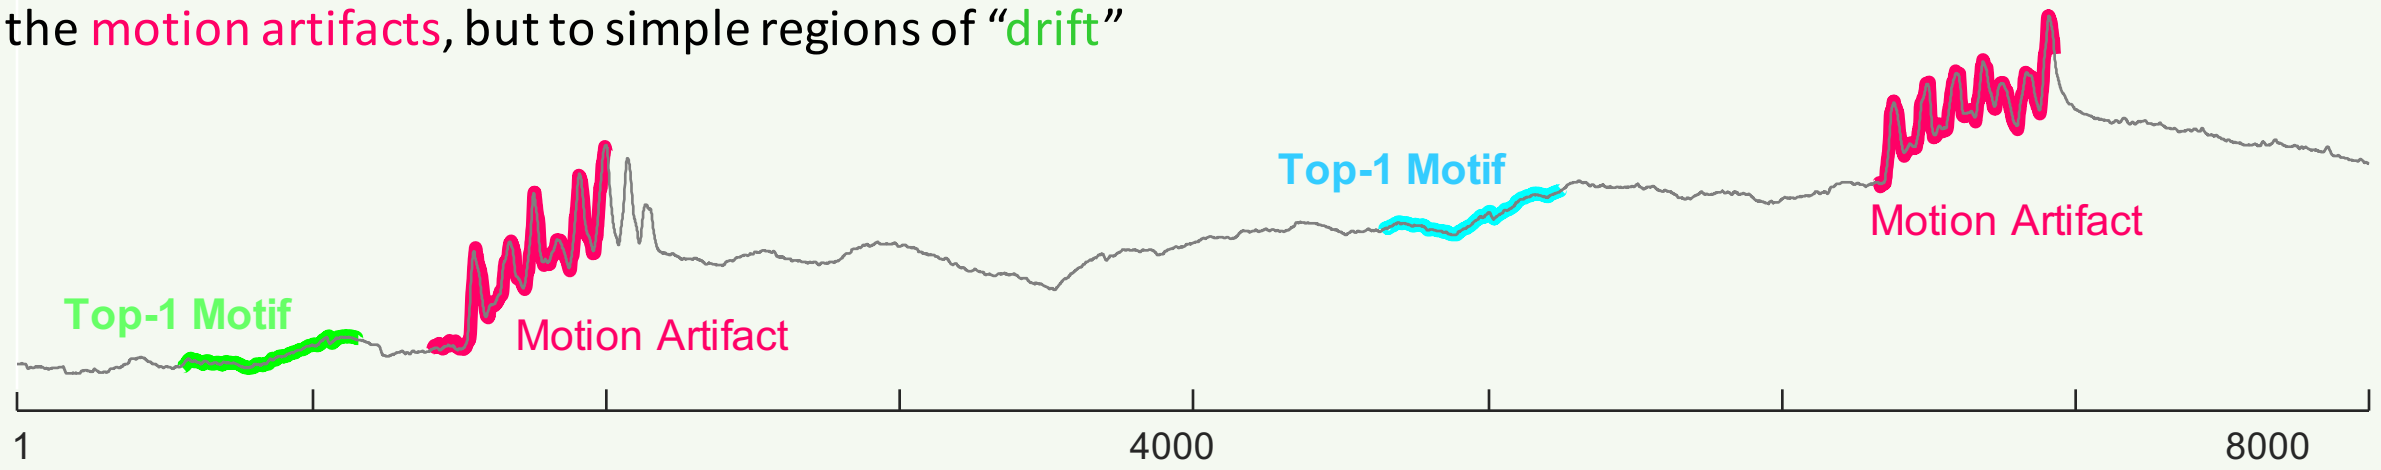

A snippet of ECG time series in which two motion artifacts were deliberately introduced by the attending physician.

[1] Batista et al. "*CID: an efficient complexity-invariant distance for time series*." Data Mining and Knowledge Discovery, 2014

Motivating Example 3:
Actionability Bias

In many cases a domain expert wants to find not simply the best motif, but regularities in the data which are *exploitable or actionable* in some domain specific ways.

> *"I want to find motifs in this web-click data, preferably occurring on or close to the weekend."*
>
> *"I want to find motifs in this oil pressure data, but they would be more useful if they end with a rising trend."*

Such queries can be almost seen as a hybrid between motif search and similarity search.

Key Insight/Claim

- We could try to have different definitions of motifs, for different domains/people/preferences/situations.
- However, we might have to devise a new search algorithm for each one, and maybe some such algorithms could be hard to speed up.
- That would mean having to abandon our nice, fast, one-size-fits-all matrix profile.
- Instead, we can do the following:
  - Use our one-size-fits-all matrix profile algorithm to find the basic matrix profile
  - Then use a domain dependent function, the Annotation Vector, to “nudge” the matrix profile to better suit the individual desired domains/people/preferences/situations

The annotation vector framework

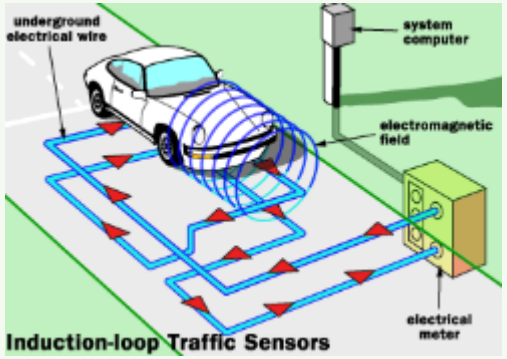

The annotation vector (AV) is a time series consisting of real-valued numbers between [0 - 1].

A lower value indicates the subsequence starting at that index is less desirable, and therefore should be biased *against*.

Conversely, higher values mean the corresponding subsequences should be favored *for* the potential motif pool.

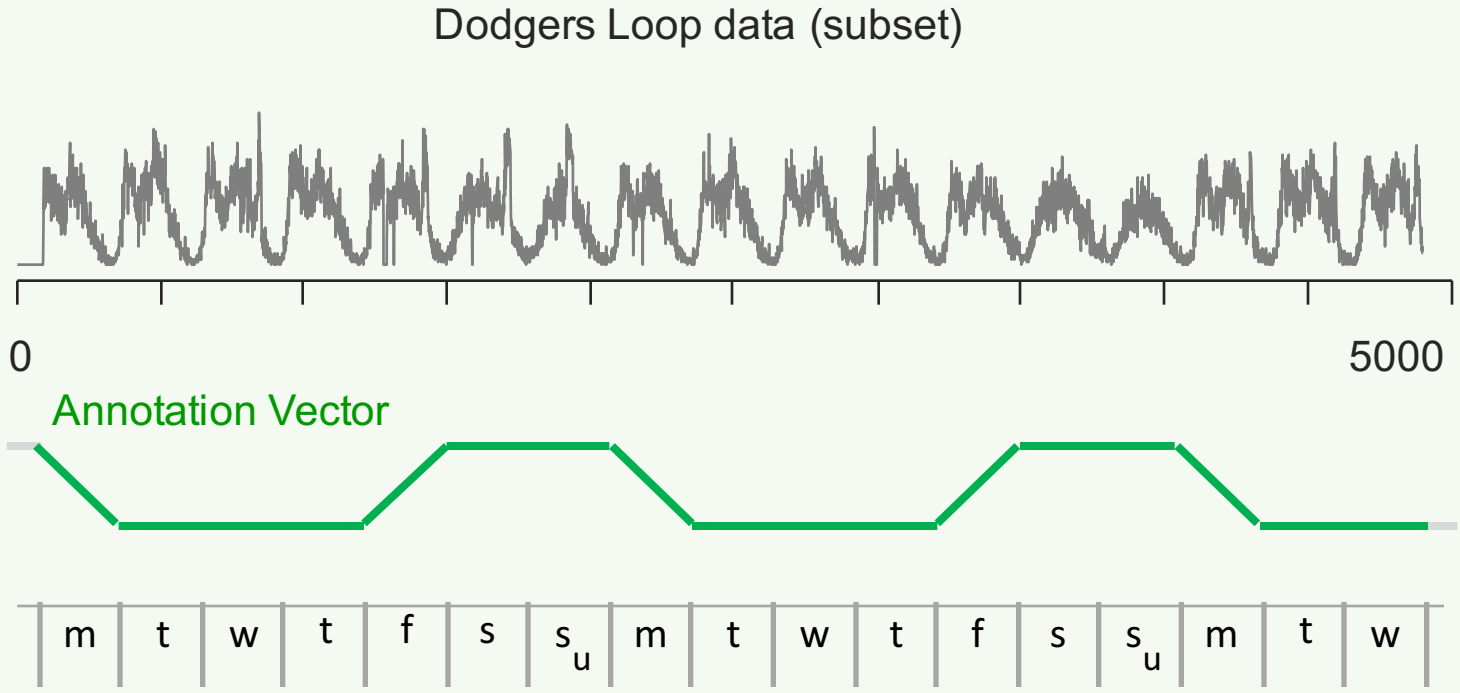

top) Seventeen days from the Dodger Loop dataset.

bottom) The AV that encodes a preference for motifs occurring on or near the weekend.

The annotation vector framework

Combines the matrix profile (MP) with the annotation vector (AV) to produce a new "adjusted" matrix profile.

We refer to this as the "Corrected" MP (CMP), as it correctly incorporates the contextual bias for the problem.

$$CMP_i = MP_i + (1 - AV_i) * \max(MP)$$

If $AV_i = 0 \rightarrow CMP_i = MP_i + \max(MP)$: raises MP value in order to remove the subsequence from potential motif pool
If $AV_i = 1 \rightarrow CMP_i = MP_i$: retains original MP value to allow the motifs that best balance the fidelity of conservation with the user's constraints to rise to the top

This only leaves the question of how do we create such an AV for our domain of interest?

Key Claim: For most problems, a domain expert can design an appropriate AV with 5 minutes of introspection, and implement it in 2 or 3 lines of code or an excel script.

Case study: Stop-word motif bias

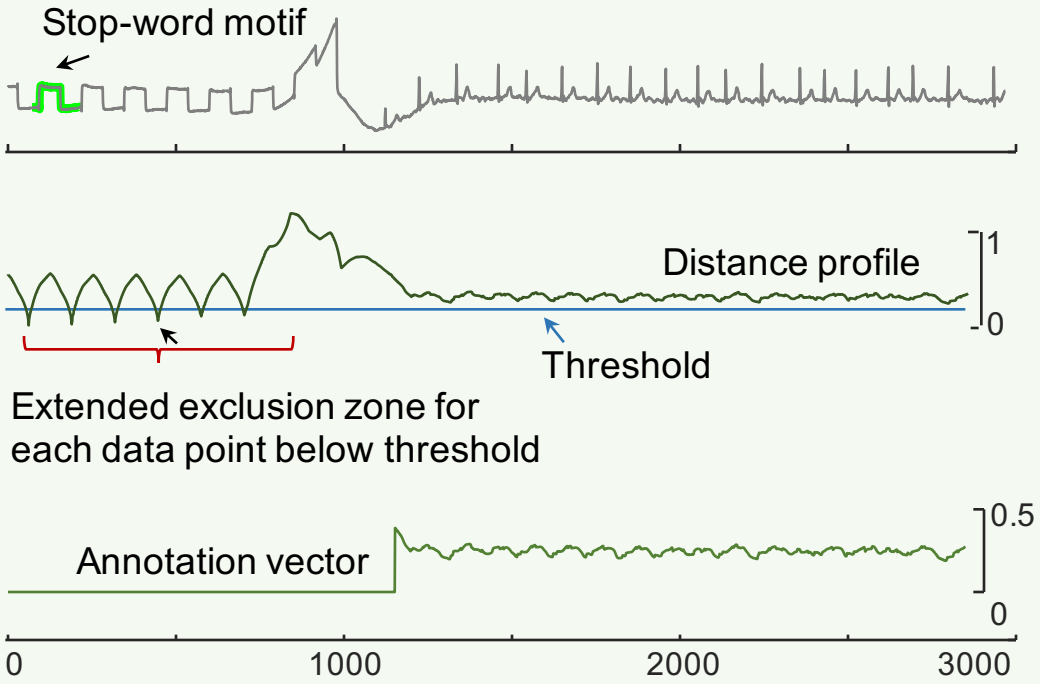

top) We annotated a single stop-word from the LTAF-71 dataset. *middle*) The stop-word distance profile to the entire dataset was thresholded to create an exclusion zone, which was used to create an AV (*bottom*).

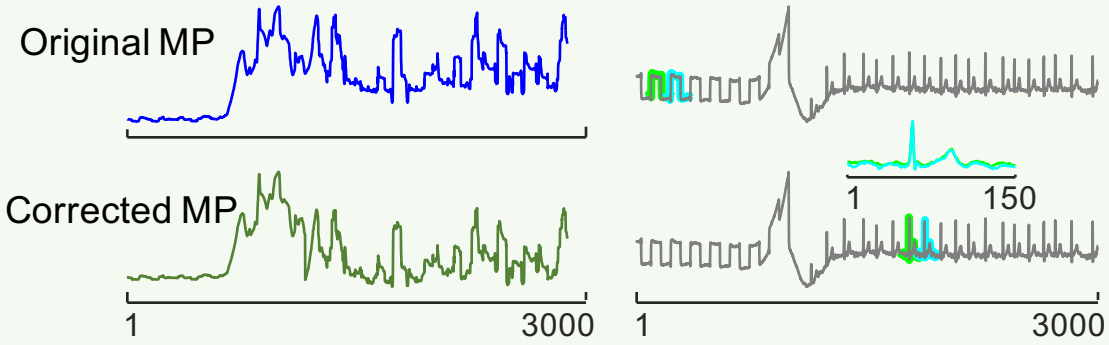

By correcting the MP to bias away from stop-word motifs, we can discover medically meaningful motifs.

Case study: Actionability bias (i)

Suppressing motion artifacts

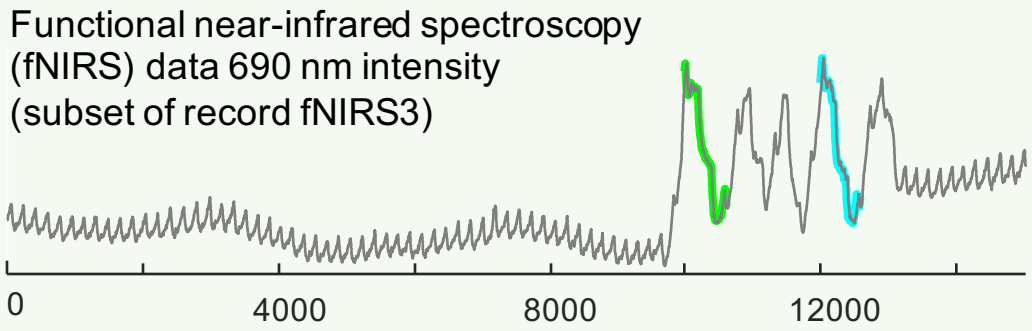

A snippet of fNIRS searched for motifs of length 600. The motifs correspond to an atypical region, which (using external data, see Fig. 7 below) we know is due to a sensor artifact.

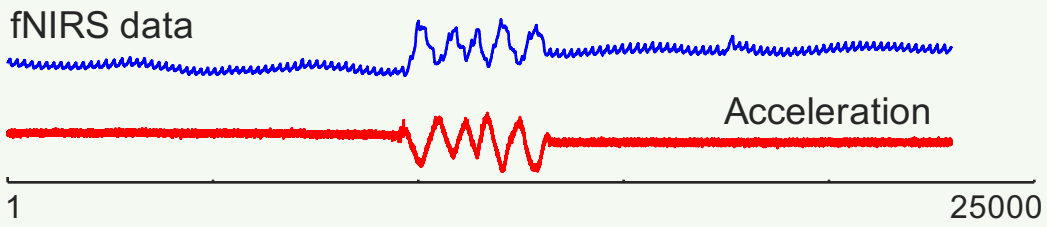

The synchronization between the fNIRS data and accompanying accelerator data.

How to make the AV

- Slides a window of length m across the acceleration time series.
- Compares the STD of each subsequence with the mean of all the subsequences' STDs, and assign the corresponding AV value to be either 0 or 1

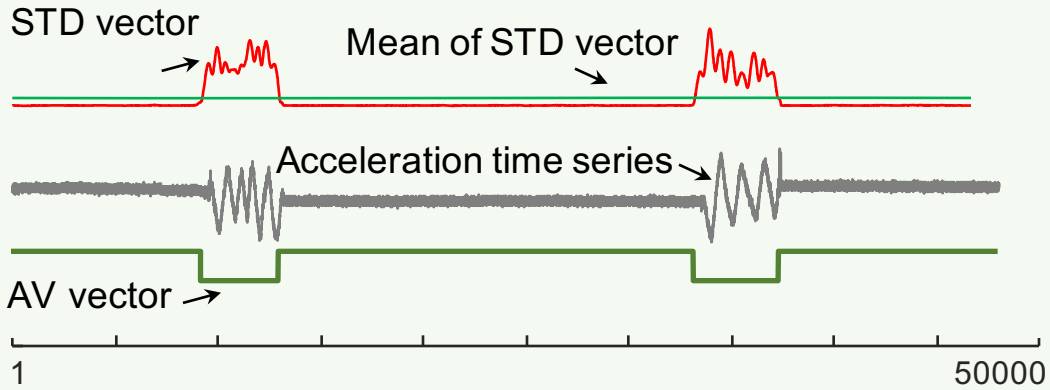

Points above the mean of all subsequences' standard deviation are well aligned with regions of motion artifacts. The corresponding AV values for these points are 0 and 1 for the rest.

Case study: Actionability bias (ii)

Suppressing motion artifacts

(*top to bottom*) Motifs in fNIRS data discovered using classic motif search tend to be spurious motion artifacts, because the matrix profile is minimized by the highly conserved but specious patterns.

If we use an AV to correct the MP, then that CMP allows us to find medically significant motifs.

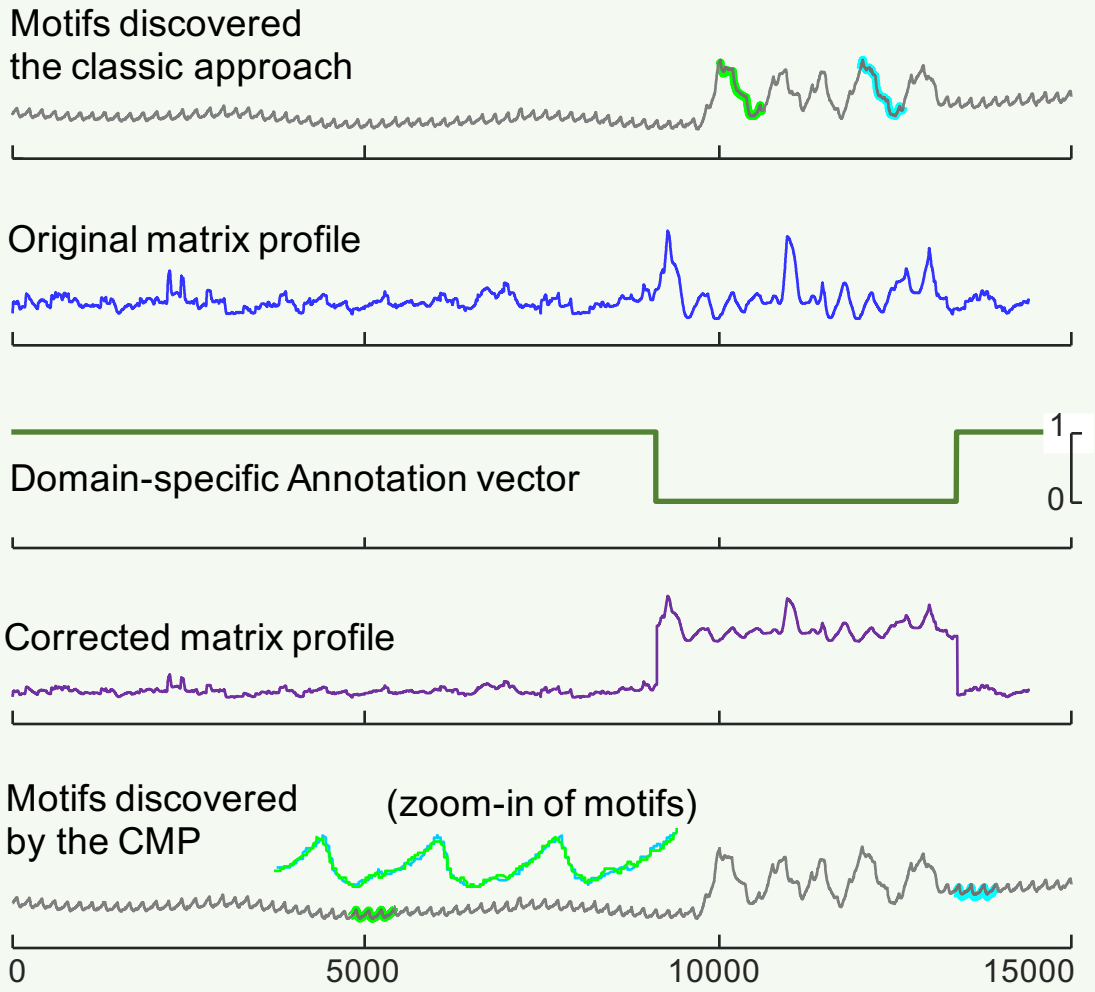

Case Study: Actionability Bias

Suppressing hard-limited artifacts

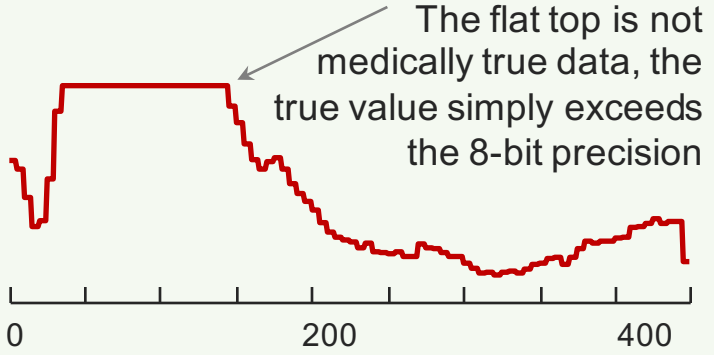

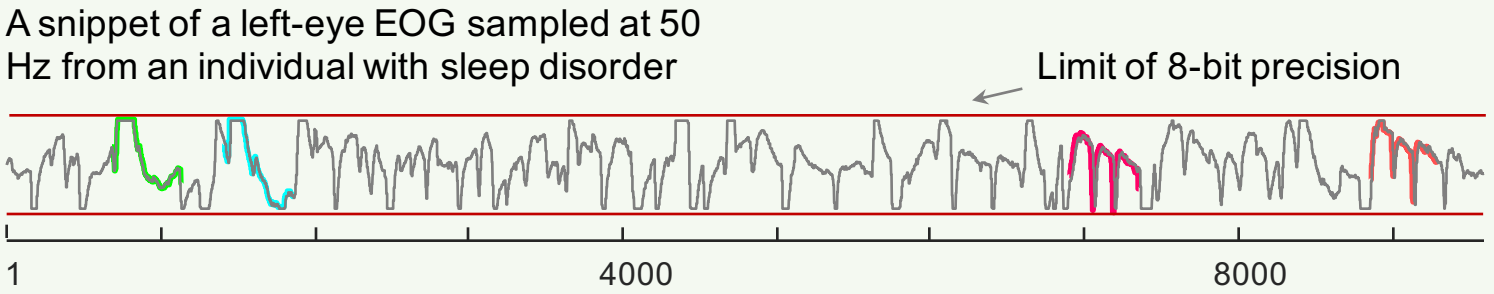

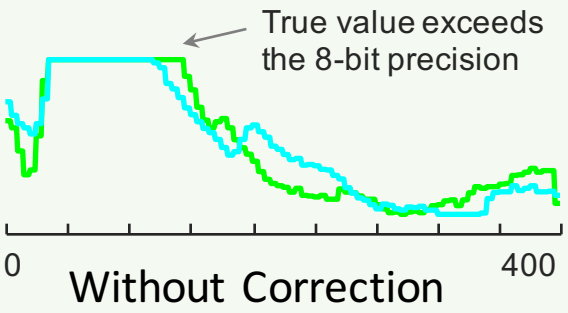

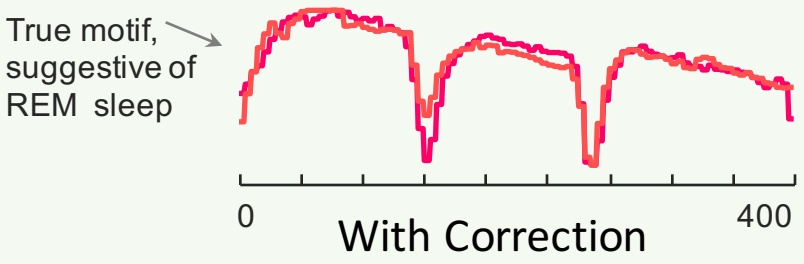

How to make the AV

- record the maximum and minimum values of the time series (the constant values touching the red bars)
- slide a window across the time series to extract subsequences
- count the number of constant values (from being hard-limited above or below) in each subsequence
- This number over the subsequence length is used as the bias function.
- This take 3 minutes, and five lines of MATLAB.

Motifs that are discovered using classic motif search tend to include hard-limited data (*left*), because the matrix profile is minimized by having long constant regions. By creating an AV to correct the MP, we can find true motifs corresponding to ponto-geniculo-occipital waves (*right*)

Case study: Simplicity bias

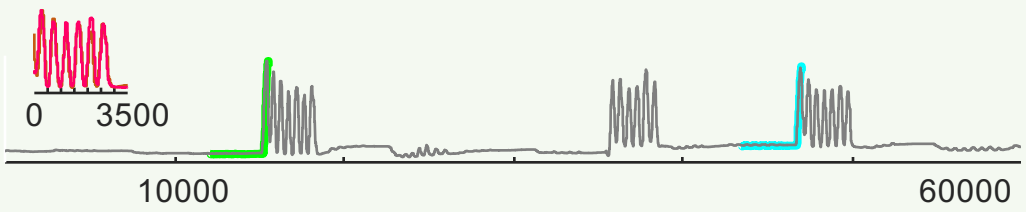

A short snippet of a time series of the flexion of a subject's little finger. Subjectively, most people would expect that two occurrences of consecutive multiple flexions to be the top motif (inset). Instead we find the simple "ramp-up" pattern.

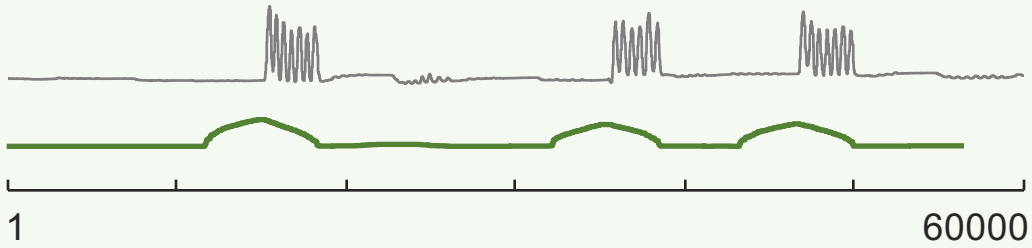

The complexity measure shown in parallel to the raw data. We simply normalize this complexity vector to be in range [0 - 1] to obtain the final AV.

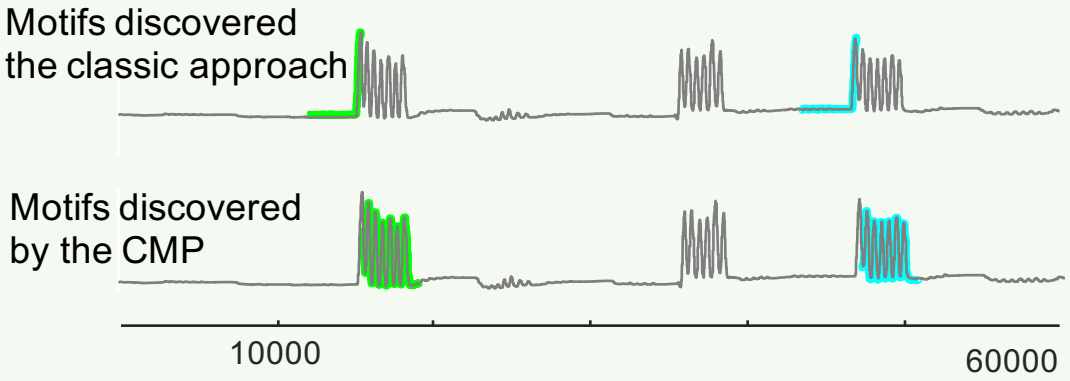

By correcting the matrix profile with an AV based on complexity measure, we discover the true motifs of finger flexion pattern.

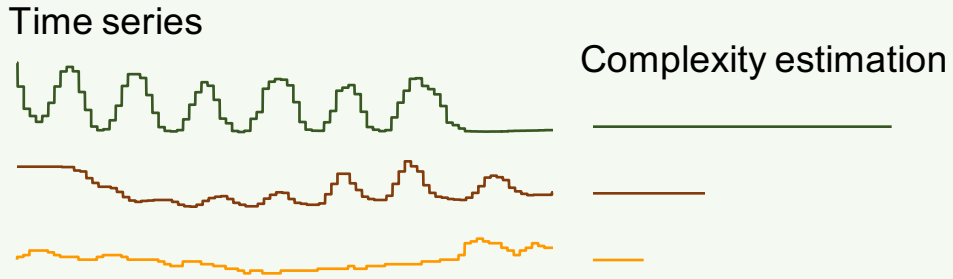

A visual intuition of the complexity estimation of three time series subsequences of different complexity levels.

Summary of the last ten minutes: annotation vector

Most of the time, the plain vanilla MP is going to be all you need to find motifs/discords/chains etc. for you data.

In some cases, you may get spurious results. That is to say, mathematically correct results, but not what you want/need/expect for your domain.

In those cases, you can just invent a simple function to suppress the spurious motifs, code it up as an annotation vector in a handful of lines of code, use it to "correct" the MP, and then run the motifs/discords/chains algorithms as before.

Once you invent an AV, say AV_{diesel_engine} or $AV_{Turkish_folk_music}$, you can reuse it on similar datasets, share it with a friend, publish it etc.

The "*Matrix Profile and ten lines of code is all you need*" philosophy

Key Idea:

- We should think of the Matrix Profile as a black box, a *primitive*.
- As we will later see, in most cases we can think of it as being obtained essentially *for free*.
- We claim that given this primitive, and *at most* ten lines of additional code, we can reproduce the results of hundred of papers.
- This suggests that other people, may be able to take this primitive, add ten lines of code, and do amazing things that have not occurred to us. We look forward to seeing what you come up with!
- In the meantime, lets see an example of: *with ten lines of additional code, we can reproduce the results of a published paper*….

Motifs under Uniform Scaling

We took two exemplars from the same class from the MALLET dataset, and imbedded them into a random walk dataset. Even without the color-coded clue brushed onto the data by the Matrix Profile discovery tool, the repeated pattern is visually obvious.

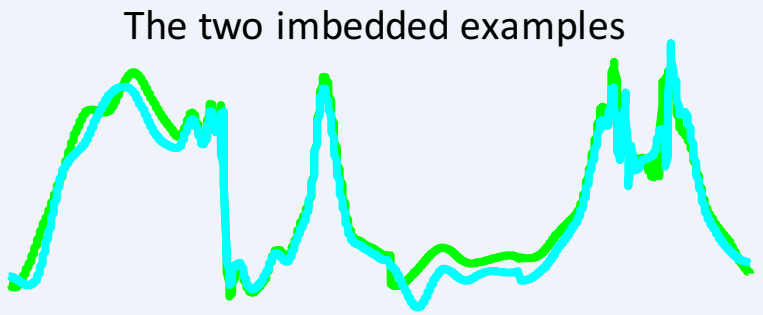

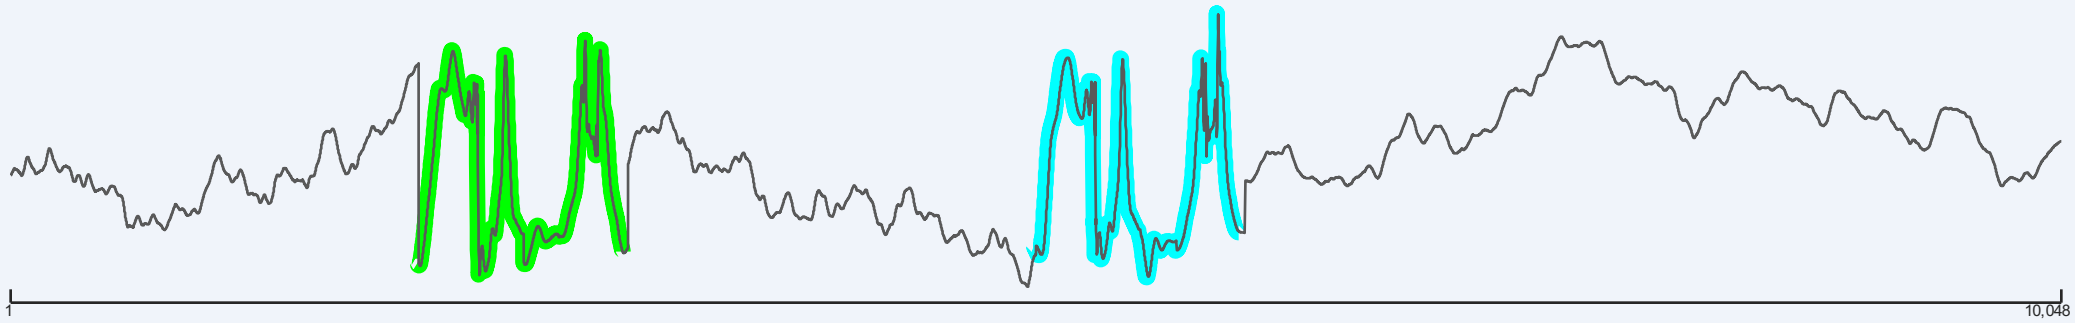

Suggestion: Toggle back and forth with last slide

We stretched the left half of the time series by just 5%, and now the pair of imbedded patterns are no longer the top-1 motif, an unexpected and disquieting result.

There is *one* paper* that offers a solution, but it is approximate, complicated, has lots of parameters, is slow..

5% stretching means the shapes begin to go out of phase, accumulating more and more error…

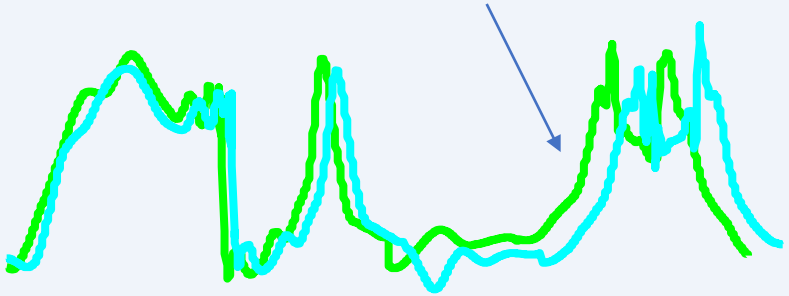

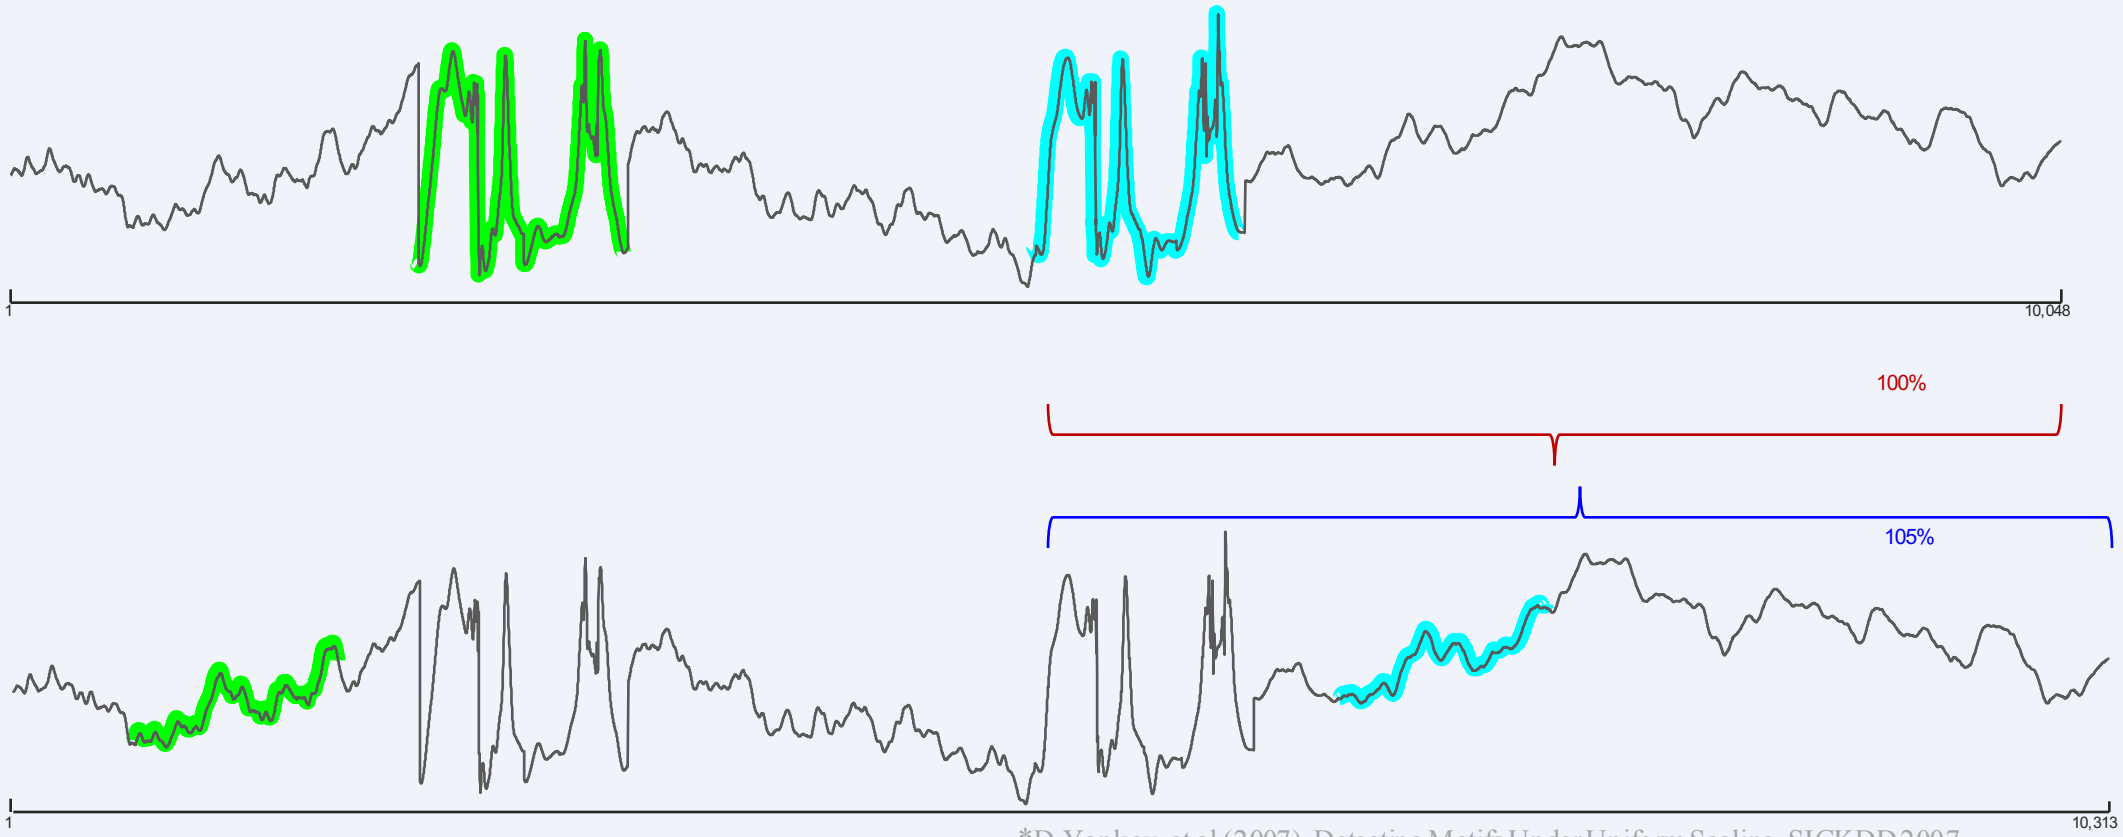

*D. Yankov, et al (2007). Detecting Motifs Under Uniform Scaling. SIGKDD 2007.

This issue is easy to fix with our "*Matrix Profile and ten lines of code is all you need*" philosophy.

For example. Suppose you suspect that there are motifs in your dataset, that differ in length by 164%

Take the original dataset T, and copy a *stretched version* of it into T2, simply by using:

```
T2 = T(1: 100/164: end); % Unofficial matlab way to resample
```

Now call:

```
[JMP, JMPindex] = computeMatrixProfileJoin(T,T2,500);
```

The resulting Matrix Profile will discover the motifs with the appropriate uniform scaling invariance.

This issue is easy to fix with our "*Matrix Profile and ten lines of code is all you need*" philosophy.

For example. Suppose you suspect that there are motifs in your dataset, that differ in length by 164%

Take the original dataset T, and copy a *stretched version* of it into T2, simply by using:

```
T2 = T(1: 100/164: end); % Unofficial matlab way to resample
```

Now call:

```
[JMP, JMPindex] = computeMatrixProfileJoin(T,T2,500);
```

The resulting Matrix Profile will discover the motifs with the appropriate uniform scaling invariance.
We did this for the electric power demand example below…

What if you don't know the scaling factor?
It is trivial to search over all possibilities.

```
for scale_factor = 101 : 170
  T2 = T(1: 100/scale_factor:  end);
  [JMP, JMPindex] = computeMatrixProfileJoin(T,T2,500);
  <trivial code to record best motifs omitted>
end
```

So, with the *Matrix Profile and ten lines of code,* we can reproduce the contribution of a SIGKDD published research effort.

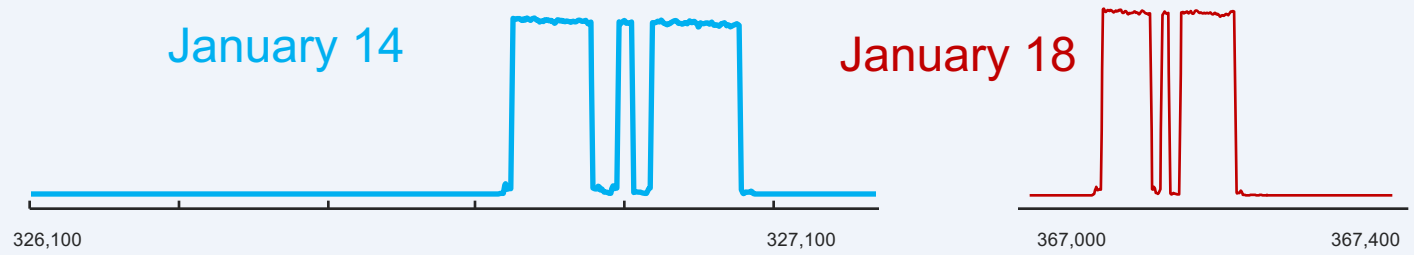

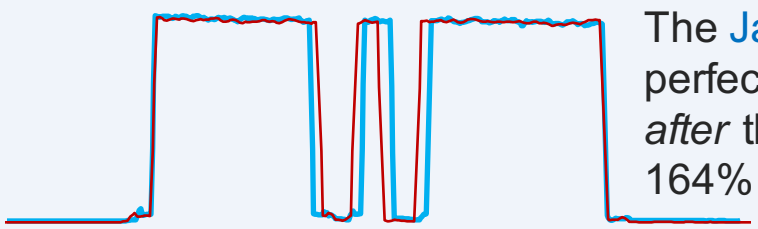

The January 14th pattern is a near perfect match the January 18th pattern, *after* the latter is uniformly stretched to 164% of its original length.

The Great Divorce

- In the last hour or so we have discussed the many wonderful things you can do with the Matrix Profile, *without* explaining how to compute it!
- This divorce is *deliberate*, and very useful.
- We can completely separate the fun, interesting part of time series data mining, from the more challenging backend part.
- The divorce lets us use appropriate computational resources. For example, exploring a million datapoints in real-time with approximate *anytime* matrix profiles on a desktop, but then outsourcing to a GPU or the cloud, when we need exact answers, or we need to explore a billion datapoints.
- All will be revealed, after the coffee break…

KDD2017
Der Bayrisch krieg
Kitzpichel
Radenberg
Küefstain
Break
© Biblioteca Nacional de España

http://www.cs.ucr.edu/~eamonn/MatrixProfile.html

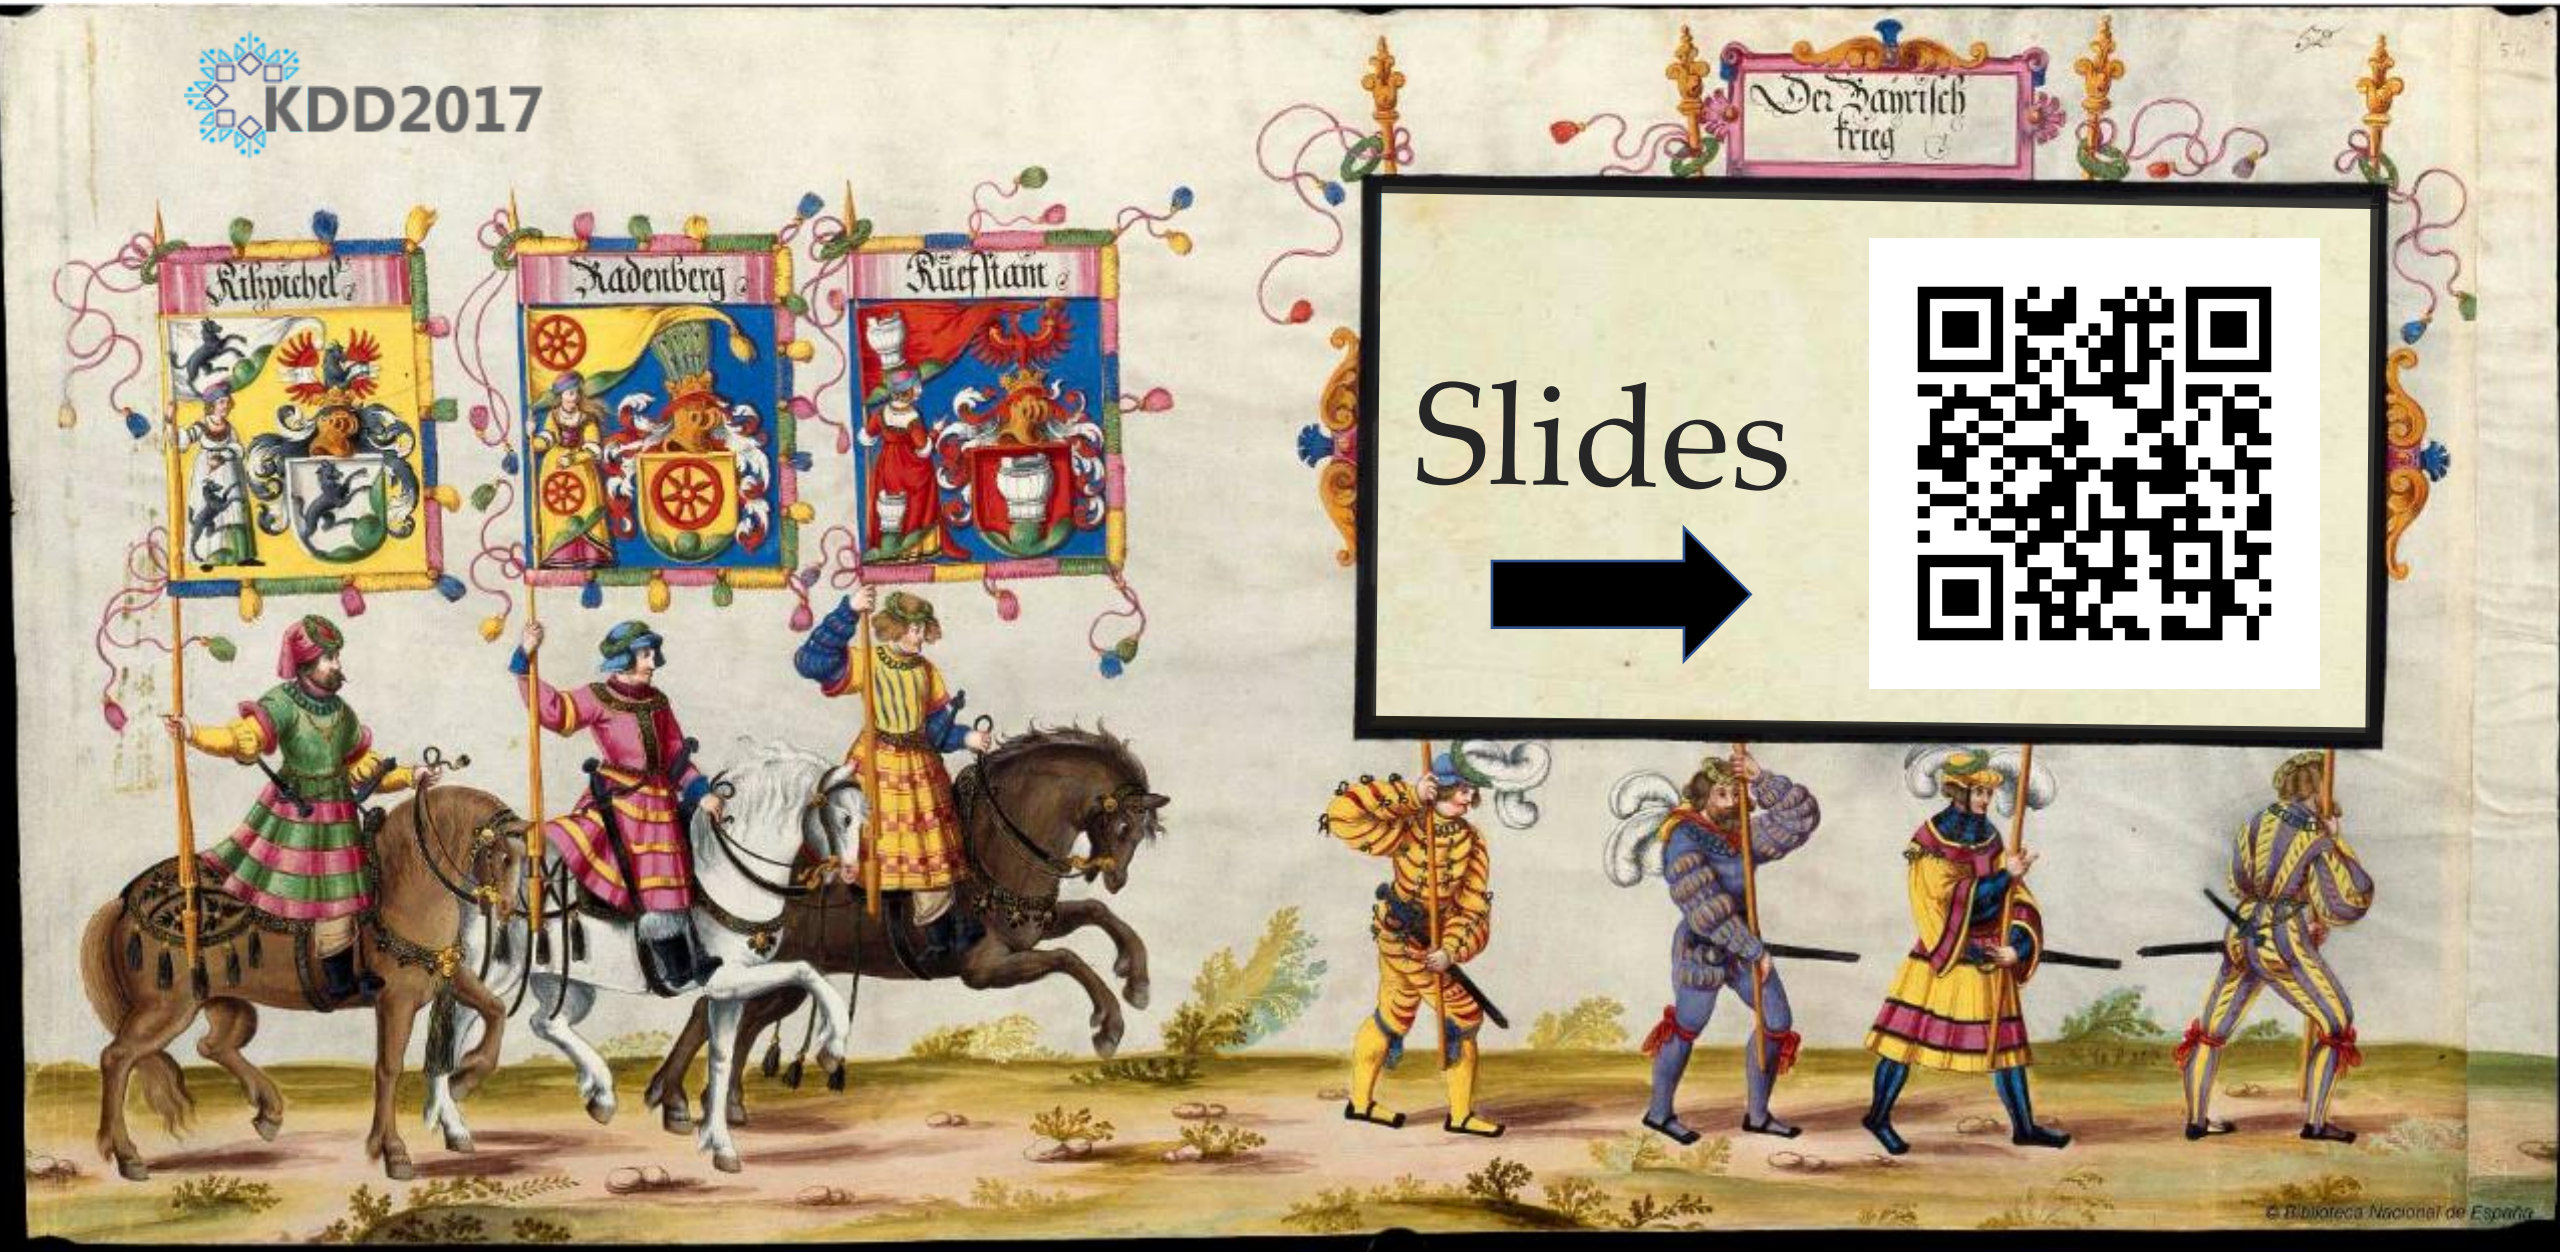

Reconvene

- In the first half, we have seen many wonderful things you can do with the Matrix Profile, *without* explaining how to compute it!
- In this half, we will describe algorithms to compute matrix profile, optimization techniques for scalability, portability to modern hardware, approximation to gain speed, and extension to special cases.
- We embed MATLAB scripts in slides that can reproduce charts and numbers, and explain algorithms.
- Slides are text heavy to facilitate offline readers. In the presentation, focus is on figures.
- Let's begin…

Outline

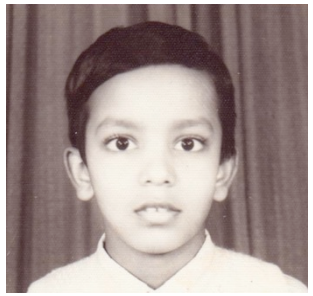

Act 1

- Our Fundamental Assumption
- What is the (MP) Matrix Profile?
- Properties of the MP
- Developing a Visual Intuition for MP
- Basic Algorithms
  - MP Motif Discovery
  - MP Time Series Chains
  - MP Anomaly Discovery
  - MP Joins (self and AB)
  - MP Semantic Segmentation
- From Domain Agnostic to Domain Aware: The Annotation Vector (A simple way to use domain knowledge to adjust your results)
- The "*Matrix Profile and ten lines of code is all you need*" philosophy.
- Break

Act 2

- Background on time series mining
  - Similarity Measures
  - Normalization
- Distance Profile
  - Brute Force Approach
  - Just-in-time Normalization
  - The MASS Algorithm
  - Extensions of MASS
- Matrix Profile
  - STAMP
  - STOMP
  - GPU-STOMP
  - SCRIMP
- Open problems to solve

What are Time Series? 1 of 2

A time series is a collection of observations made sequentially in time.

More than most types of data, time series lend themselves to *visual* inspection and intuitions…

For example, looking at the numbers in this blue vector tells us nothing.
But after *plotting* the data, we can recognize a heartbeat, and possibly even diagnose this person's disease.

When the observations are uniformly sampled, the index of observation can replace the time of observation. In the rest of the tutorial, we assume time series are vectors.

The observations may have a unit.

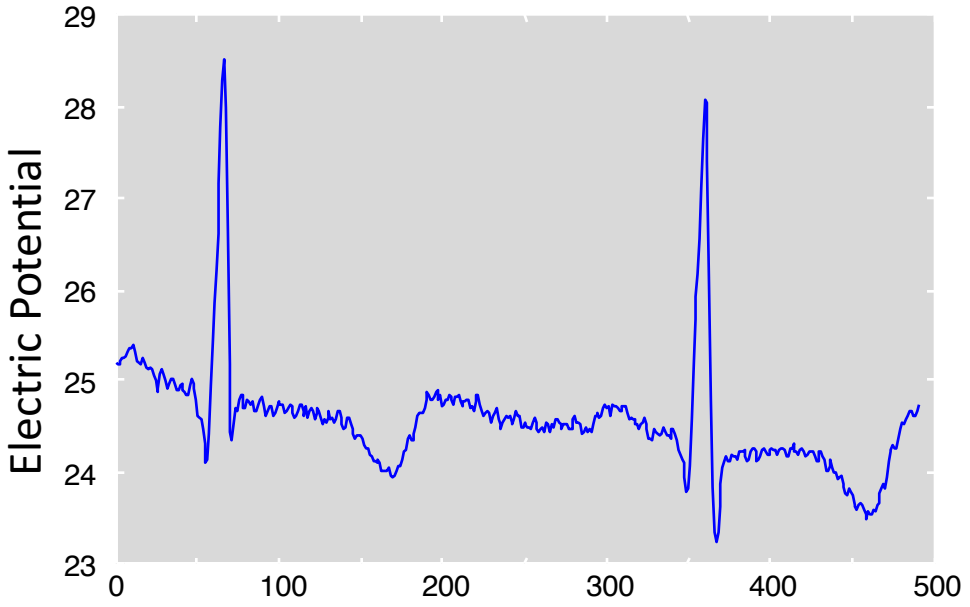

25.350
25.350
25.400
25.400
25.325
25.225
25.200
25.175
..
24.625
24.675
24.675
24.675

What are Time Series? 2 of 2

As an aside… (not the main point for today)

Many types of data that are not *true* time series can be fruitfully transformed into time series, including DNA, speech, textures, core samples, ASCII text, historical handwriting, novels and even *shapes*.

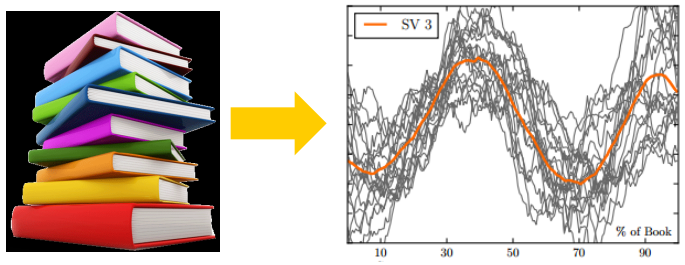

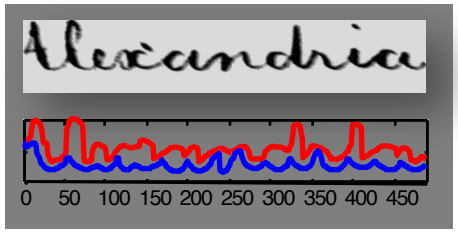

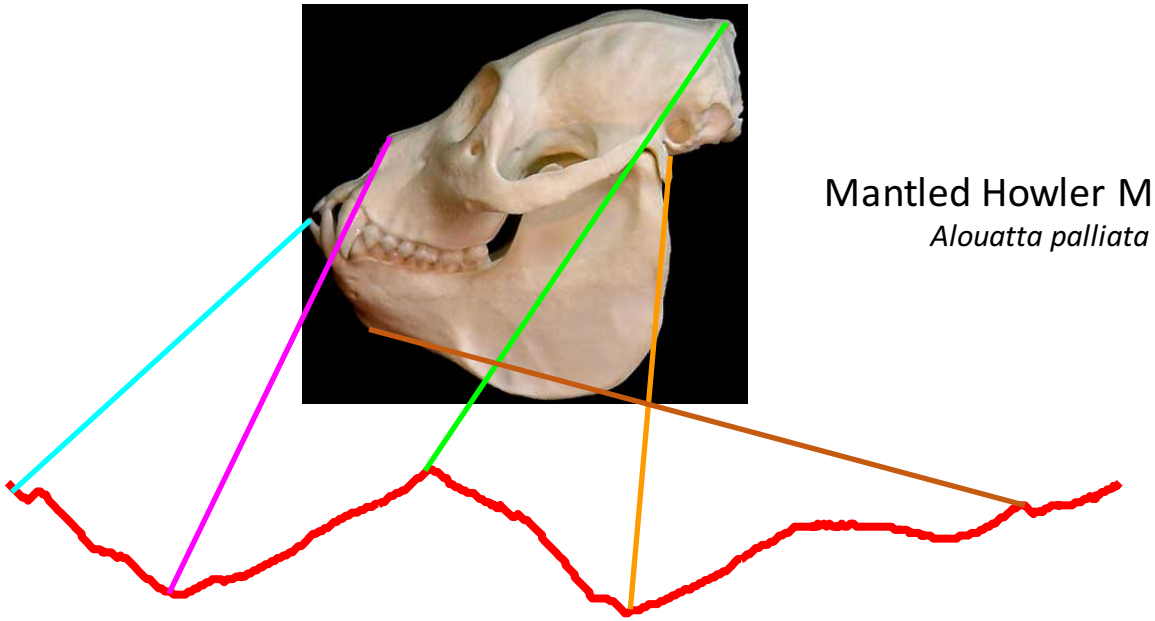

Mantled Howler Monkey
Alouatta palliata

Similarity Measures for Time Series

- A similarity measure compares two time series and produces a number representing their similarity
  - A distance measure is the opposite of similarity measure
- Lockstep Measures
  - Euclidean Distance
  - Correlation Coefficient
  - Cosine Similarity
- Elastic Measures
  - Dynamic Time Warping
  - Edit Distance
  - Longest Common Subsequence

Euclidean Distance Metric

Given two time series

$\mathbf{x} = x_1 \ldots x_n$

and

$\mathbf{y} = y_1 \ldots y_n$

their z-Normalized Euclidean distance is defined as:

$$\hat{x}_i = \frac{x_i - \mu_x}{\sigma_x} \qquad \hat{y}_i = \frac{y_i - \mu_y}{\sigma_y}$$

```
function y = zNorm(x)
y = (x-mean(x))/std(x,1);
```

$$d(x, y) = \sqrt{\sum_{i=1}^{n} (\hat{x}_i - \hat{y}_i)^2}$$

```
function d = EuclideanDistance(x,y)
d = sqrt(sum((x-y).^2));
```

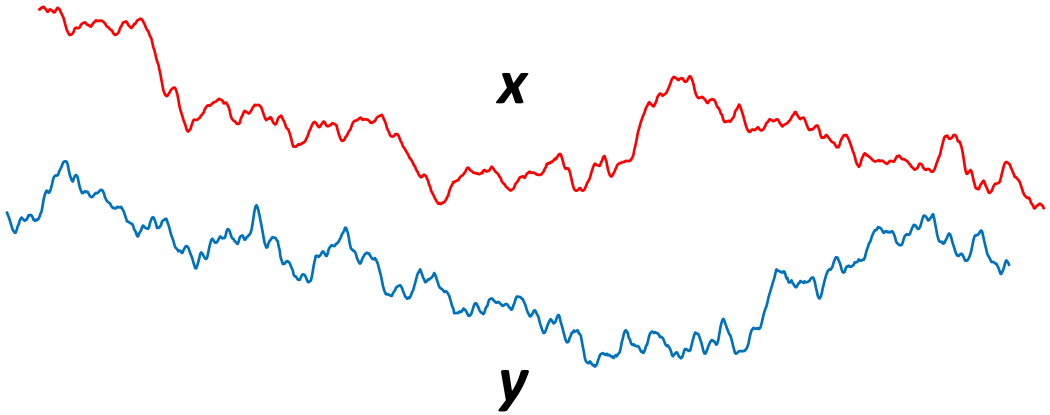

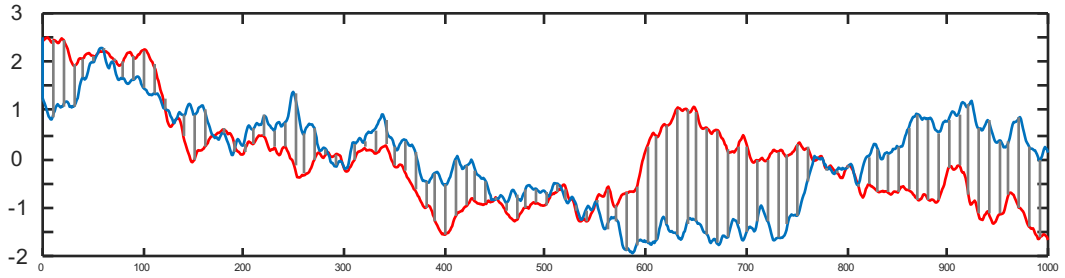

Pearson's Correlation Coefficient

- Given two time series $\boldsymbol{x}$ and $\boldsymbol{y}$ of length m.
- Correlation Coefficient:

$$corr(\boldsymbol{x}, \boldsymbol{y}) = \frac{E[(x - \mu_x)(y - \mu_y)]}{\sigma_x \sigma_y} = \frac{\sum_{i=1}^{m} x_i y_i - m\mu_x \mu_y}{m\sigma_x \sigma_y}$$

  - Where $\mu_x = \frac{\sum_{i=1}^{m} x_i}{m}$ and $\sigma_x^2 = \frac{\sum_{i=1}^{m} x_i^2}{m} - \mu_x^2$

- Sufficient Statistics:

$$\sum_{i=1}^{m} x_i y_i \quad \sum_{i=1}^{m} x_i \quad \sum_{i=1}^{m} y_i \quad \sum_{i=1}^{m} x_i^2 \quad \sum_{i=1}^{m} y_i^2$$

The sufficient statistics can be calculated in one linear scan. Given the sufficient statistics, correlation coefficient is a constant operation. Note the use of the dot product, which is the key component of many lockstep measures.

Relationship with Euclidean Distance

$$d(\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}) = \sqrt{2m(1 - corr(\boldsymbol{x}, \boldsymbol{y}))}$$

- Correlation coefficient does not obey triangular inequality, while Euclidean distance does
- Maximizing correlation coefficient can be achieved by minimizing normalized Euclidean distance and vice versa
- Correlation coefficient is bounded between -1 and 1, while z-normalized Euclidean distance is bounded between zero and a positive number dependent on m

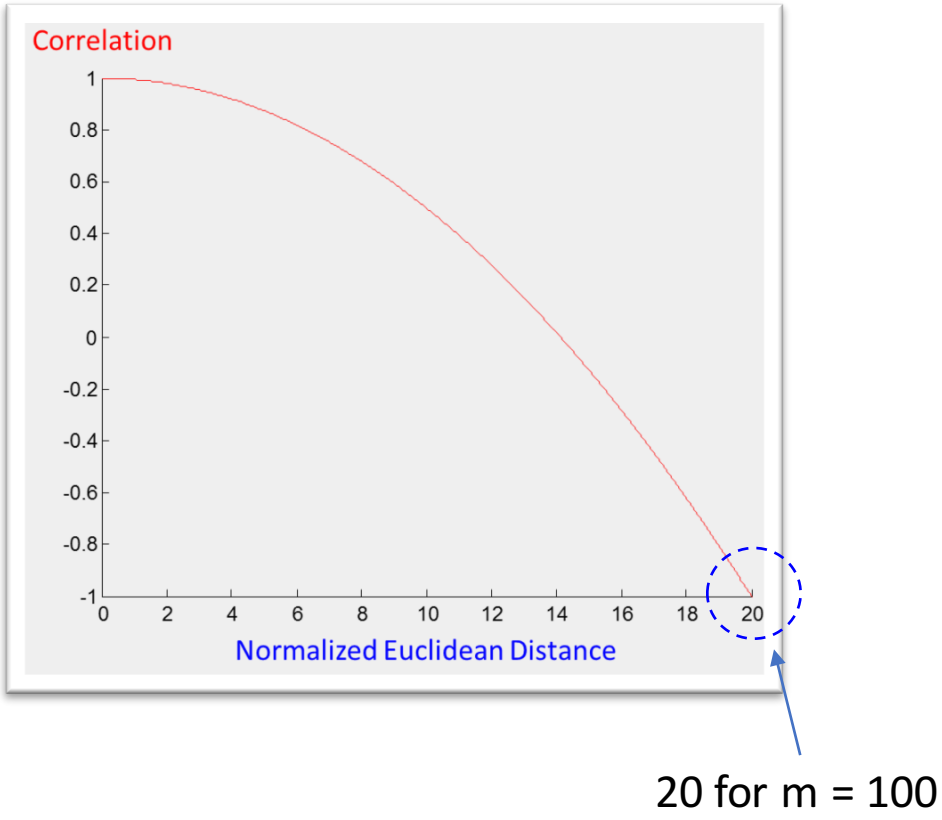

Abdullah Mueen, Suman Nath, Jie Liu: Fast approximate correlation for massive time-series data. SIGMOD Conference 2010: 171-182

Working Formula

$$d(\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}) = \sqrt{2m(1 - \frac{\sum_{i=1}^{m} x_i y_i - m\mu_x \mu_y}{m\sigma_x \sigma_y})}$$

- We will use the above z-Normalized Euclidean distance as the similarity measure for the rest of the presentation
- We claim calculating Matrix Profile for Correlation Coefficient and Cosine Similarity is trivial given an algorithm for z-Normalized Euclidean distance

The Importance of z-Normalization and correlation 1 of 2

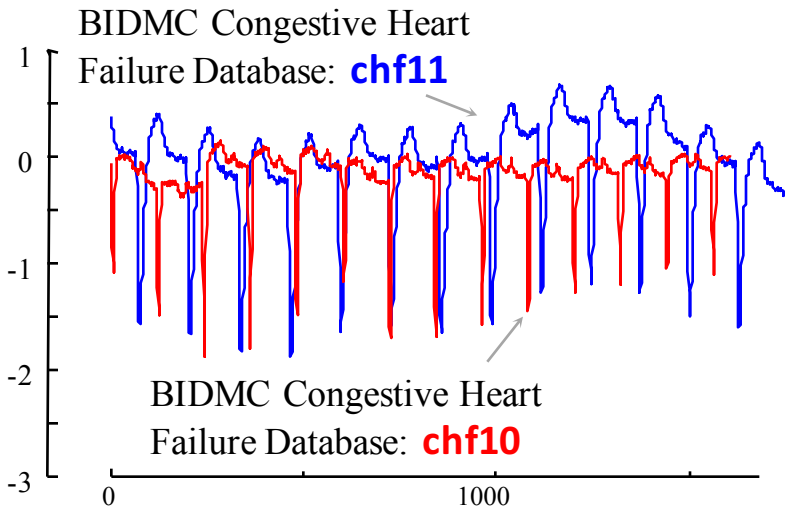

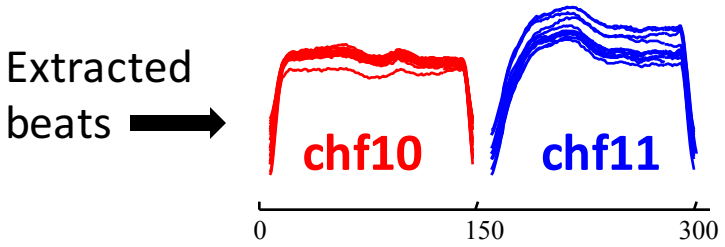

Essentially all datasets must have *every* subsequence z-normalized.

There are a handful of occasions where it does not make sense to z-normalize, but in those cases, similarity search does not make sense either.

In this example, we begin by extracting heartbeats from two unrelated people.

Even without normalization, it happens that both sets have almost the same mean and standard deviation. Given that, do we need to bother to normalize them? (next slide)

The Importance of z-Normalization and correlation 2 of 2

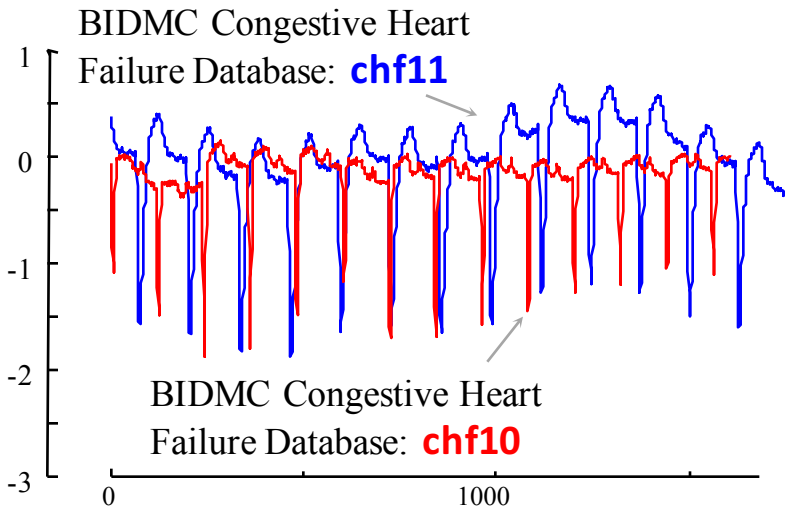

Surprisingly z-normalizing can be a computational bottleneck, but later we will show you how to fix that.

Without normalization, the results are *very* poor, some blue heartbeats are closer to red heartbeats than there are to another blue beat .

With normalization, the results are perfect.

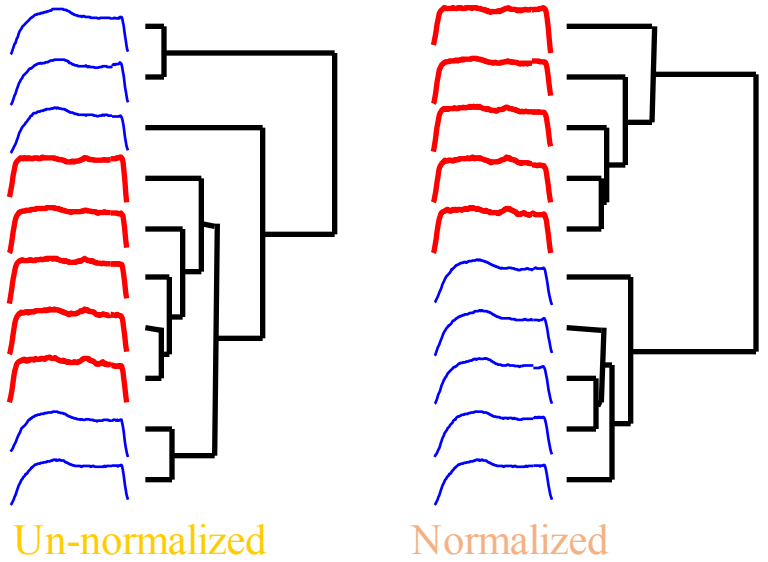

In this example, we extracted heartbeats from two different time series, and clustered them with and without normalization.

Outline

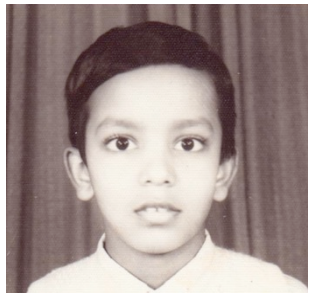

Act 1

- Our Fundamental Assumption
- What is the (MP) Matrix Profile?
- Properties of the MP
- Developing a Visual Intuition for MP
- Basic Algorithms
  - MP Motif Discovery
  - MP Time Series Chains
  - MP Anomaly Discovery
  - MP Joins (self and AB)
  - MP Semantic Segmentation
- From Domain Agnostic to Domain Aware: The Annotation Vector (A simple way to use domain knowledge to adjust your results)
- The "*Matrix Profile and ten lines of code is all you need*" philosophy.
- Break

Act 2

- Background on time series mining
  - Similarity Measures
  - Normalization
- Distance Profile
  - Brute Force Approach
  - Just-in-time Normalization
  - The MASS Algorithm
  - Extensions of MASS
- Matrix Profile
  - STAMP
  - STOMP
  - GPU-STOMP
  - SCRIMP
- Open problems to solve

Distance Profile

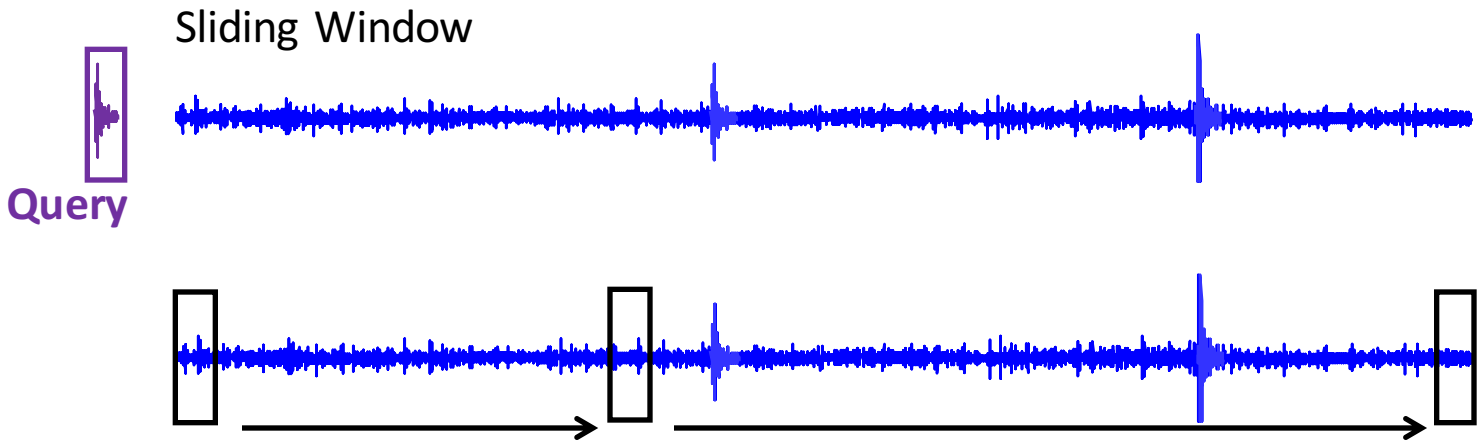

Compute the z-normalized Euclidean distance between **Query** and each window (subsequence) in the time series. We would obtain a vector like this:

d_1	d_2	…	d_{n-m+1}

$\longrightarrow$ D

d_i is the distance between the i^{th} subsequence and the query.

Recall, n is the length of the blue time series and m is the length of the query

The Brute Force Algorithm

- Scan the time series with a sliding window
- Z-Normalize the window
- Calculate Euclidean distance between window and the query

```
d(1:n) = 0;
Q = zNorm(query);
for i = 1:n-m+1
    d(i) = sqrt(sum((zNorm(T(i:i+m-1))-Q).^2));
end
```

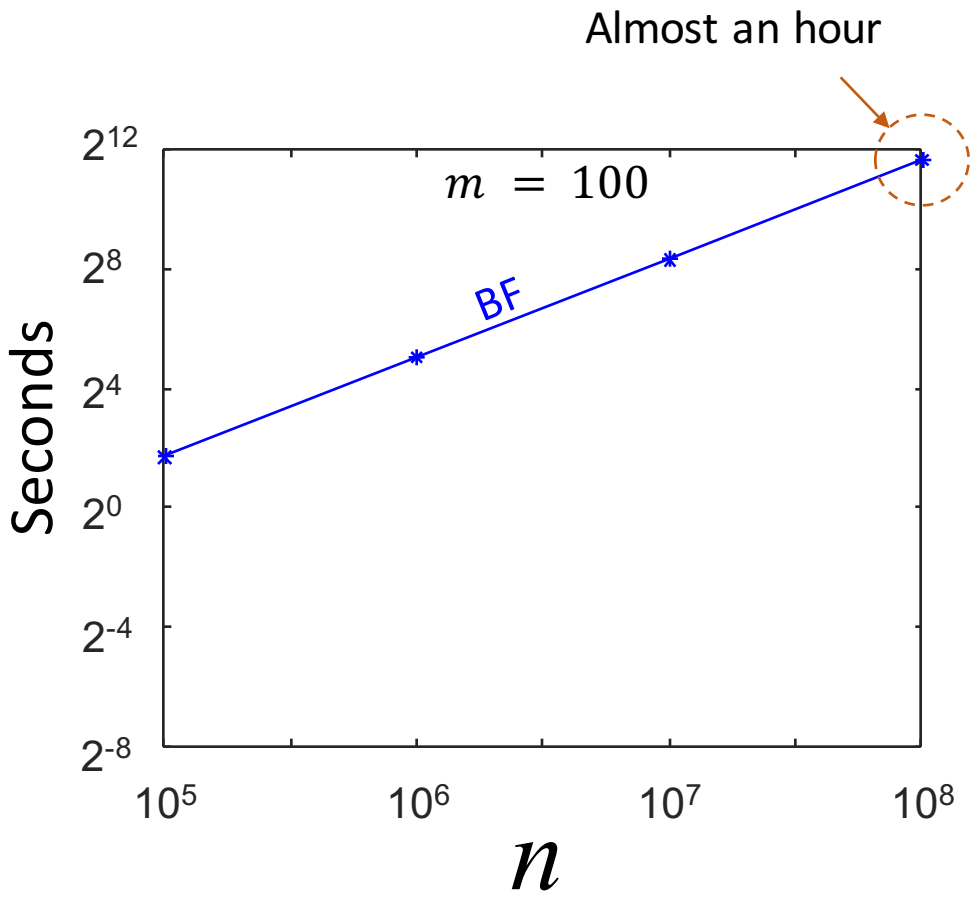

- How long does this algorithm take?
- The time complexity is $O(nm)$ in the average and worst cases. More precisely the window is scanned two times in each iteration of this algorithm. One scan is for z-normalization, and the other scan is for distance calculation.
- Note that we cannot use any early abandoning or pruning as we need all the distances.

Just-in-time Normalization (1 of 3)

- Can we skip the z-normalization scan in each iteration?
- Yes, if we have the means, standard deviations and the dot product to calculate the distances.
- z-normalized sequence has zero mean and one standard deviation.
- Let's assume $\boldsymbol{y}$ is the z-normalized query, and $\boldsymbol{x}$ is the time series (T), therefore, $\mu_y = 0$ and $\sigma_y = 1$

Working Formula

$$d(\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}) = \sqrt{2m\left(1 - \frac{\sum_{i=1}^{m} x_i y_i - m\mu_x\mu_y}{m\sigma_x\sigma_y}\right)}$$

$$d(\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}) = \sqrt{2\left(m - \frac{\sum_{i=1}^{m} x_i y_i}{\sigma_x}\right)}$$

Thanawin Rakthanmanon, et al. Searching and mining trillions of time series subsequences under dynamic time warping. KDD 2012: 262-270

Just-in-time Normalization (2 of 3)

- Can we skip the z-normalization scan in each iteration?
- The standard deviations of moving windows of a fixed size can be calculated in one linear scan.
  - In one pass, calculate cumulative sums of x and x^2 and store
  - Subtract two cumulative sums to obtain the sum over any window
  - Use the sums to calculate the standard deviations of all windows in linear time
- In 2016, MATLAB has introduced a function, `movstd`, that does the above.

$$d(\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}) = \sqrt{2(m - \frac{\sum_{i=1}^{m} x_i y_i}{\sigma_x})}$$

$$C = \sum x \qquad C^2 = \sum x^2$$

$$S_i = C_{i+m} - C_i \qquad S_i^2 = C_{i+m}^2 - C_i^2$$

$$\sigma_i = \sqrt{\frac{S_i^2}{m} - \left(\frac{S_i}{m}\right)^2}$$

Thanawin Rakthanmanon, et al. Searching and mining trillions of time series subsequences under dynamic time warping. KDD 2012: 262-270

Just-in-time Normalization (3 of 3)

- Can we skip the z-normalization scan in each iteration?

```
d(1:n) = 0;
Q = zNorm(query);
S = movstd(T,[0 m-1]);
for i = 1:n-m+1
    d(i) = sqrt(2*(m-sum(T(i:i+m-1).*Q)/S(i)));
end
```

- Still the worst and average cost is $O(nm)$, however, the window is scanned only once per iteration for the dot product.
- Speedup is more than 2X, due to removal of function calls

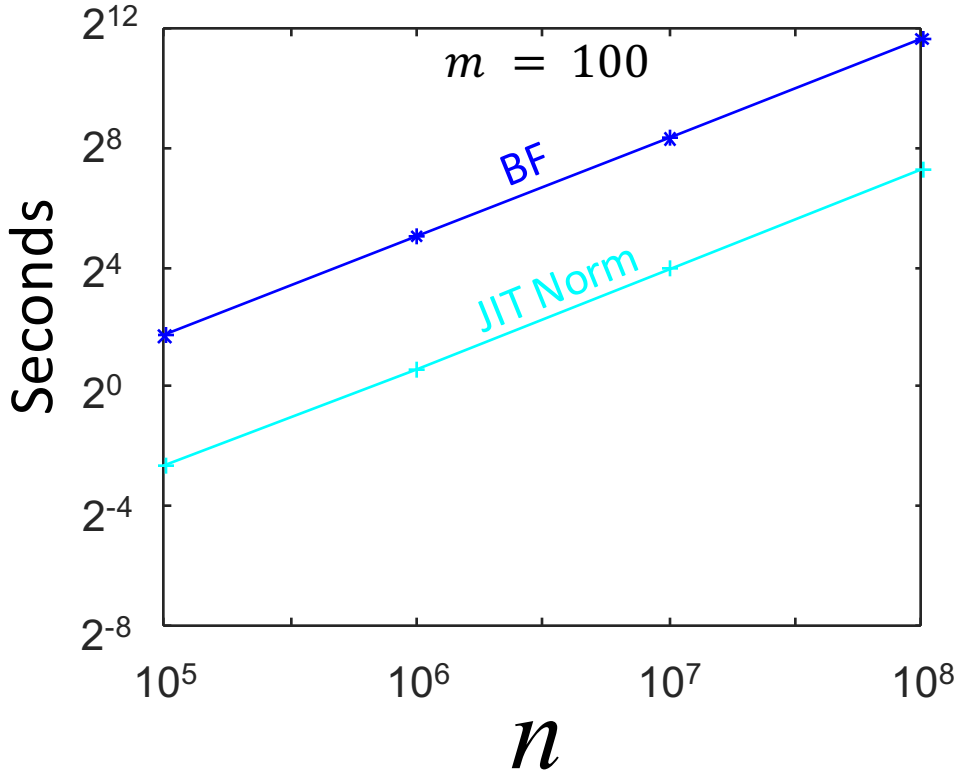

Thanawin Rakthanmanon, et al. Searching and mining trillions of time series subsequences under dynamic time warping. KDD 2012: 262-270

Mueen's Algorithm for Similarity Search (MASS) (1 of 9)

- Can we improve the just-in-time Normalization algorithm?
- MASS uses a convolution based method to calculate sliding dot products in $O(n \log n)$, in addition to just-in-time z-normalization technique
- **Convolution:** If x and y are vectors of polynomial coefficients, convolving them is equivalent to multiplying the two polynomials.
- We use convolution to compute all of the sliding dot products between the query and sliding windows.

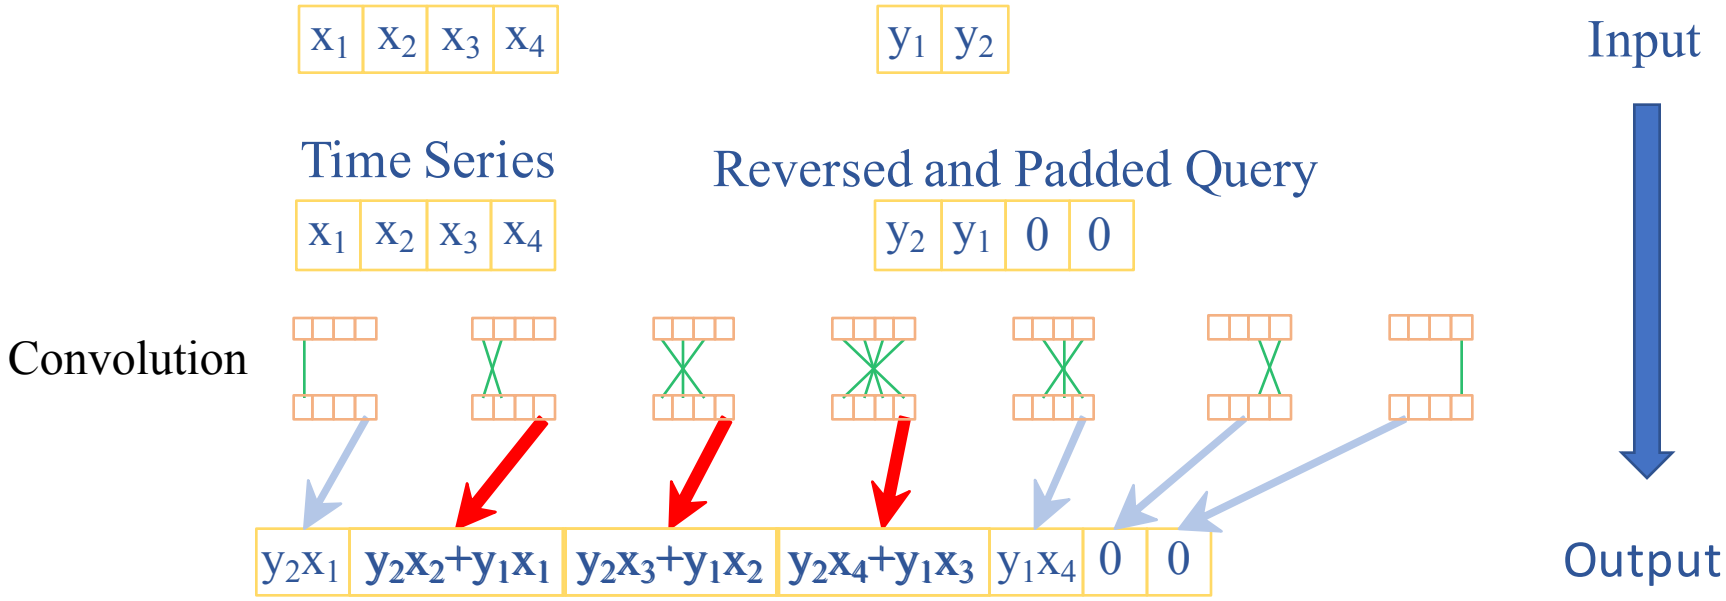

Mueen's Algorithm for Similarity Search (MASS) (2 of 9)

MASS 1.0

```
d(1:n) = 0;
Q = zNorm(query);
Stdv = movstd(T,[0 m-1]);
Q = Q(end:-1:1);        %Reverse the query
Q(m+1:n) = 0;           %pad zeros
dots = conv(T,Q);
dist = 2*(m-(dots(m:n))./Stdv));
dist = sqrt(dist);
```

The loop has been replaced by the following three lines.

```
d(1:n) = 0;
Q = zNorm(query);
S = movstd(T,[0 m-1]);
for i = 1:n-m+1
    dd(i) = sqrt(2*(m-sum(T(i:i+m-1).*Q)/S(i)));
end
```

Vectorized working formula

- Computational cost, $O(n \log n)$, does not depend on the query length (m), thus, free of curse of dimensionality.
- There is no loop. Only known mathematical and built-in MATLAB functions are used.

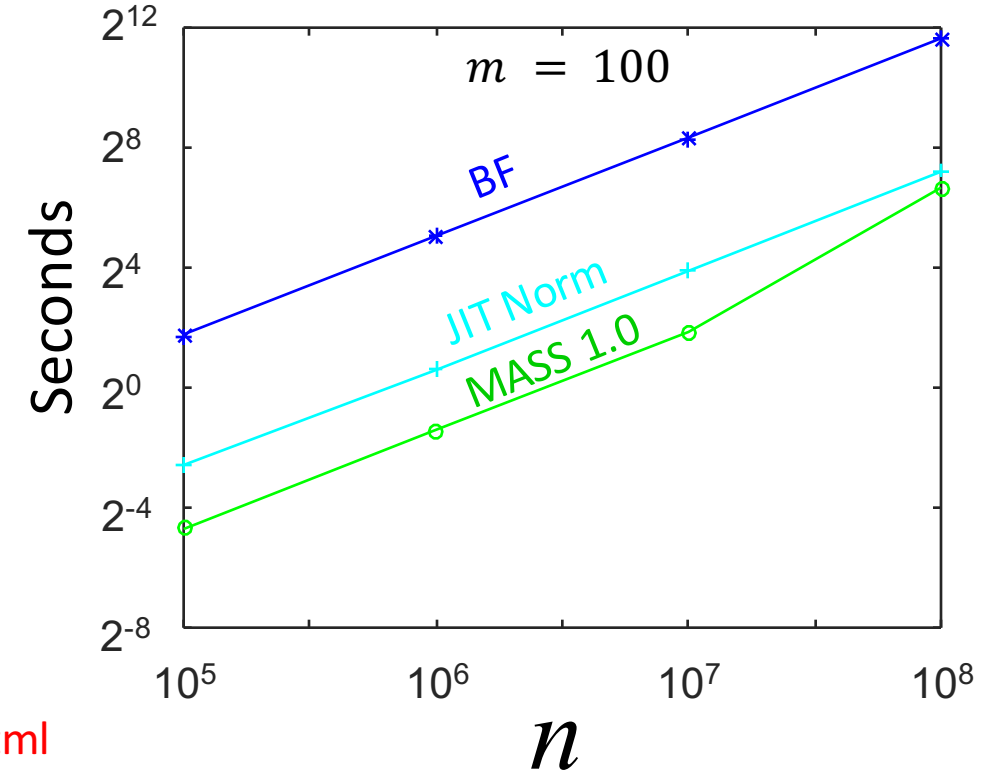

http://www.cs.unm.edu/~mueen/FastestSimilaritySearch.html

Mueen's Algorithm for Similarity Search (MASS) (3 of 9)

- Can we improve MASS 1.0?
- Note that convolution doubles the size of the input vectors in the output.
- MASS uses only half of the output of convolution and throws away the remaining half.
- Can we compute just the necessary half? Let's see what happens inside convolution.
- *Convolution in time domain is multiplication in frequency domain.*
- `conv(x,y) = ifft( fft( double&pad(x)) . fft(double&pad(y))`

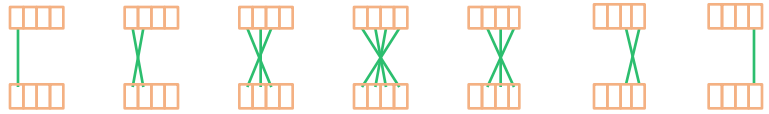

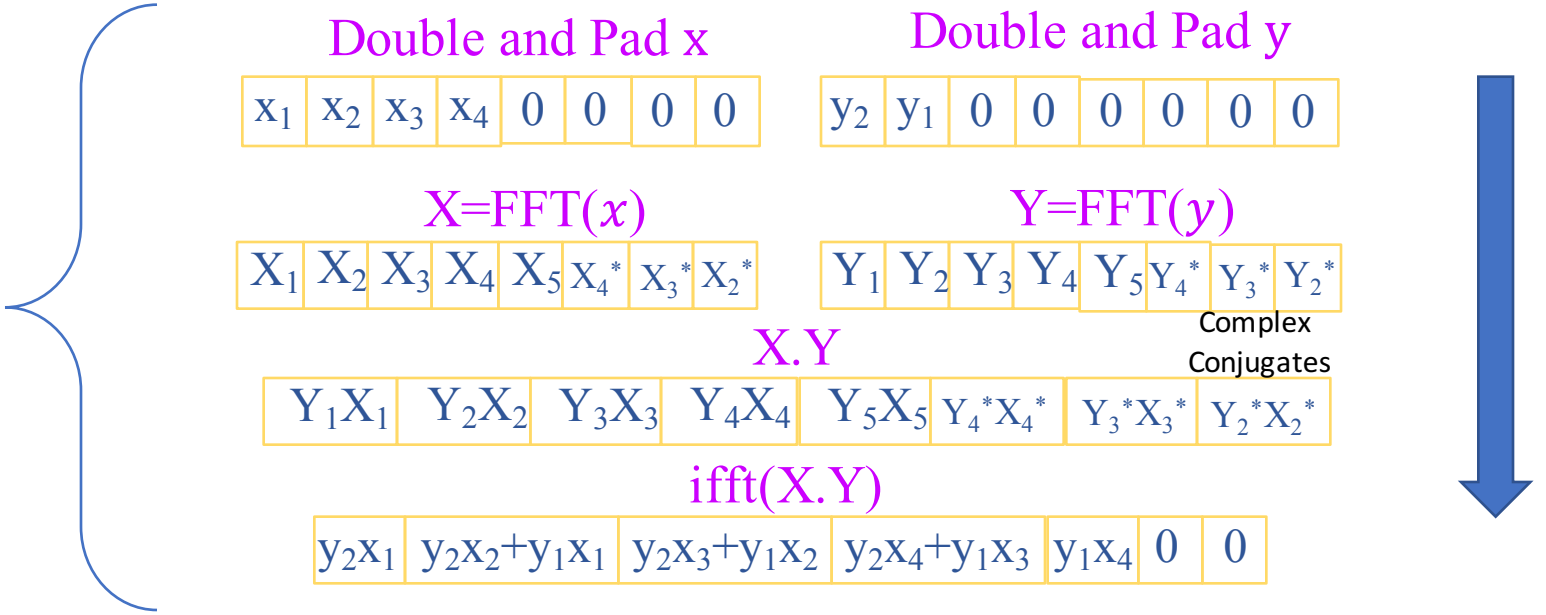

Mueen's Algorithm for Similarity Search (MASS) (4 of 9)

- Can we improve MASS 1.0?
- If we do not double x and y, we obtain a half convolution
- `half conv(x,y)= ifft ( fft(x) . fft(y))`

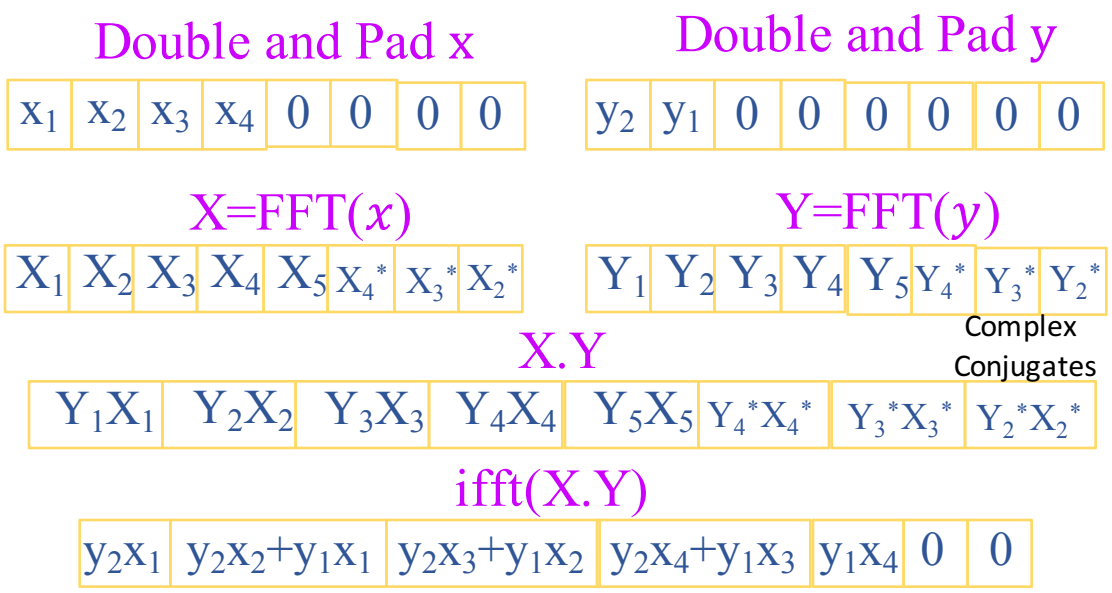

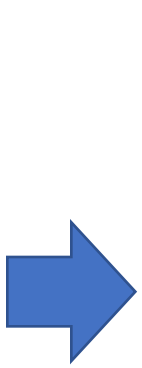

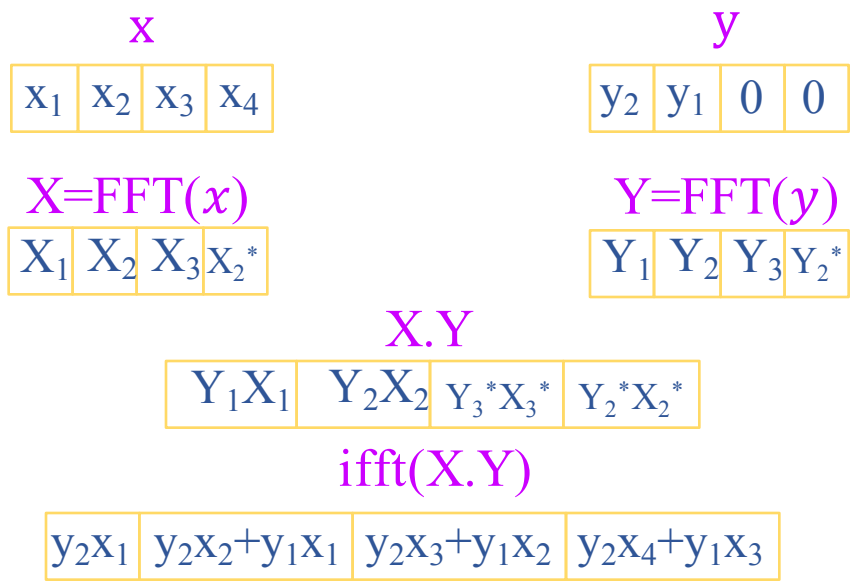


```
conv(x,y) = ifft( fft( double&pad(x)) .
                  fft(double&pad(y) )
```

```
half conv(x,y)= ifft ( fft(x) . fft(y))
```

Mueen's Algorithm for Similarity Search (MASS) (5 of 9)

- Can we improve MASS 1.0?
- Half convolution adds a constraint, $n > \frac{m}{2}$. The constraint is not limiting because the original assumption is $n \gg m$.

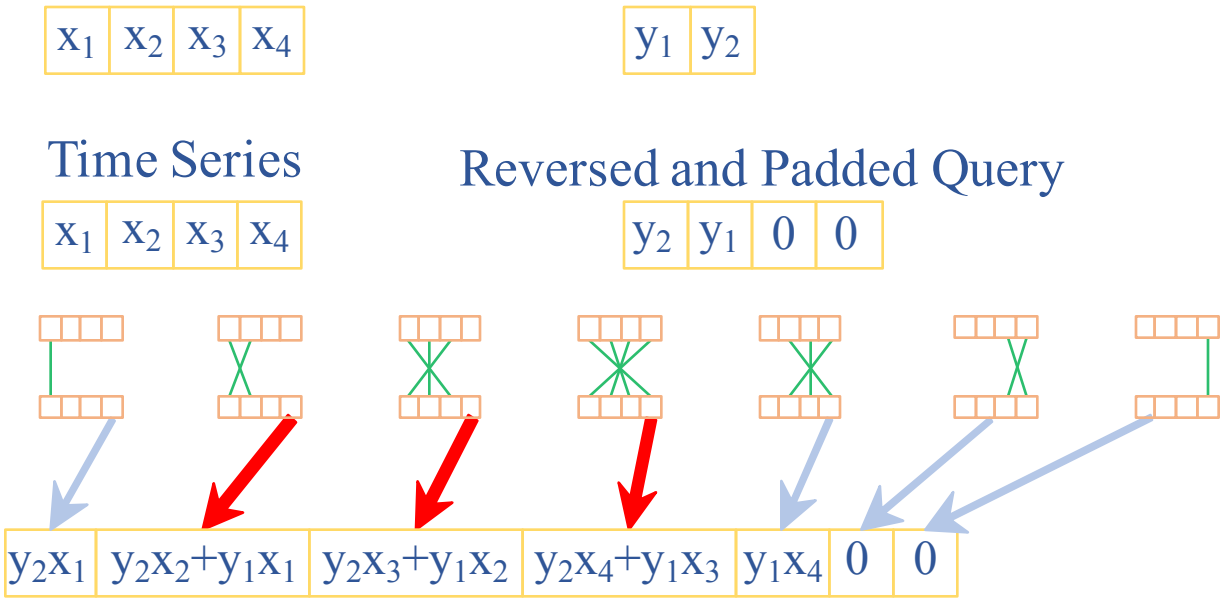

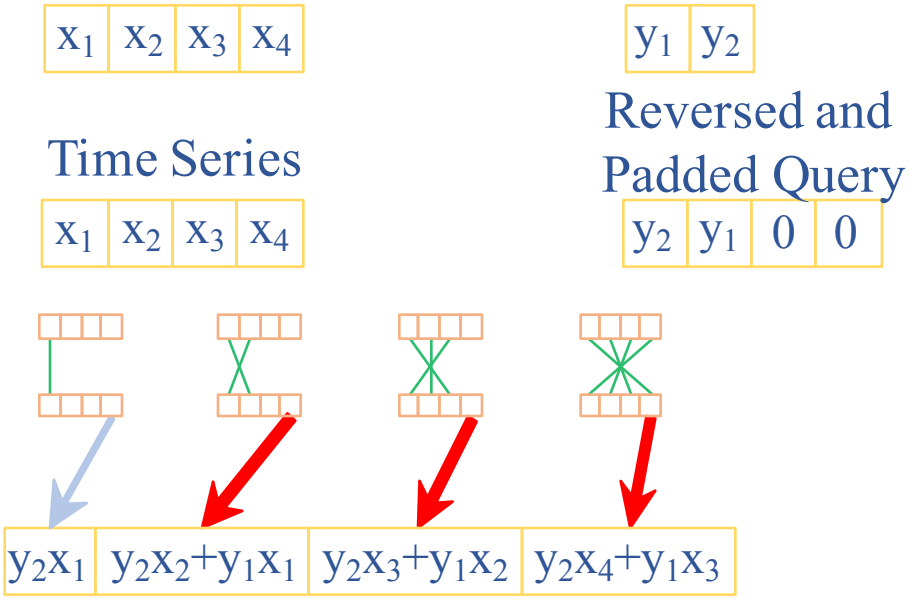


```
conv(x,y) = ifft( fft( double&pad(x)) .
                  fft(double&pad(y) )
```

```
half conv(x,y)= ifft ( fft(x) . fft(y))
```

Mueen's Algorithm for Similarity Search (MASS) (6 of 9)

MASS 2.0

```
d(1:n) = 0;
Q = zNorm(query);
Stdv = movstd(T,[0 m-1]);
Q = Q(end:-1:1);        %Reverse the query
Q(m+1:n) = 0;           %pad zeros
dots = ifft( fft(T).* fft(Q) );
dist = 2*(m-(dots(m:n))./Stdv));
dist = sqrt(dist);
```

The `conv(T,Q)` has been replaced, no doubling of sizes

- Computational cost is still $O(n \log n)$, does not depend on the query length (m), thus, free of curse of dimensionality.
- fast Fourier transform (fft) is used as a subroutine

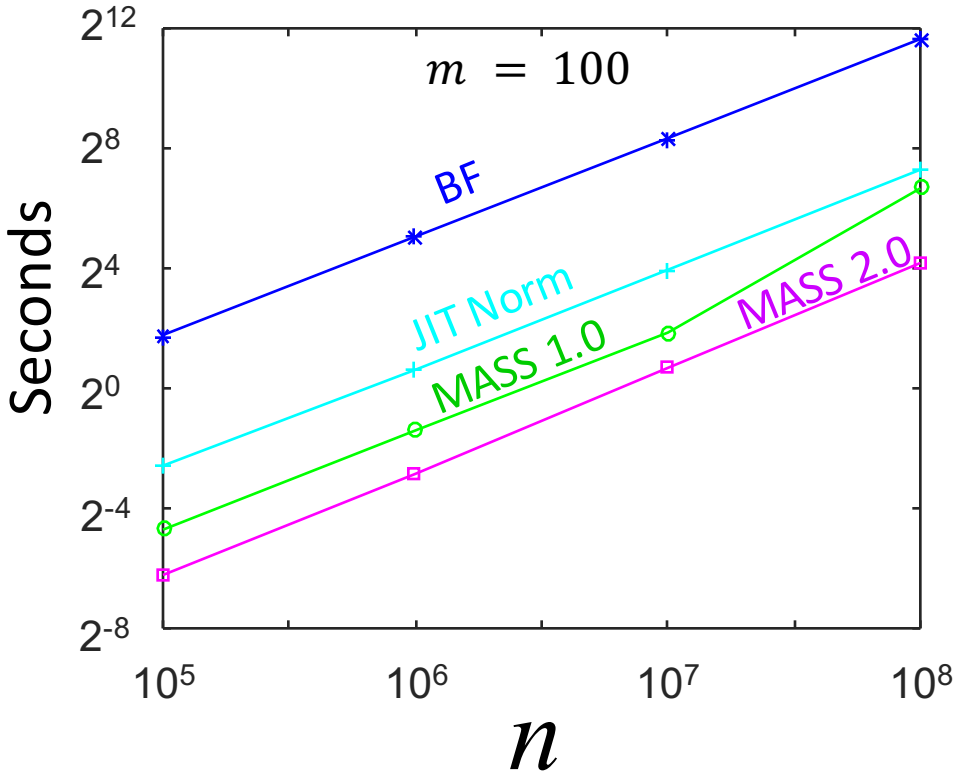

Mueen's Algorithm for Similarity Search (MASS) (7 of 9)

- Can we improve MASS 2.0?
- This idea is a rediscovery of a known result, *Overlap-add*, *Overlap-save* [a]
- Instead of processing all *n* samples at once, we can (carefully) process them *piecewise,* where sizes of the pieces are powers of two.
- The pieces must overlap exactly m-1 samples to preserve continuity of the distance profile.
- The last piece can be of size that is not a power of two.

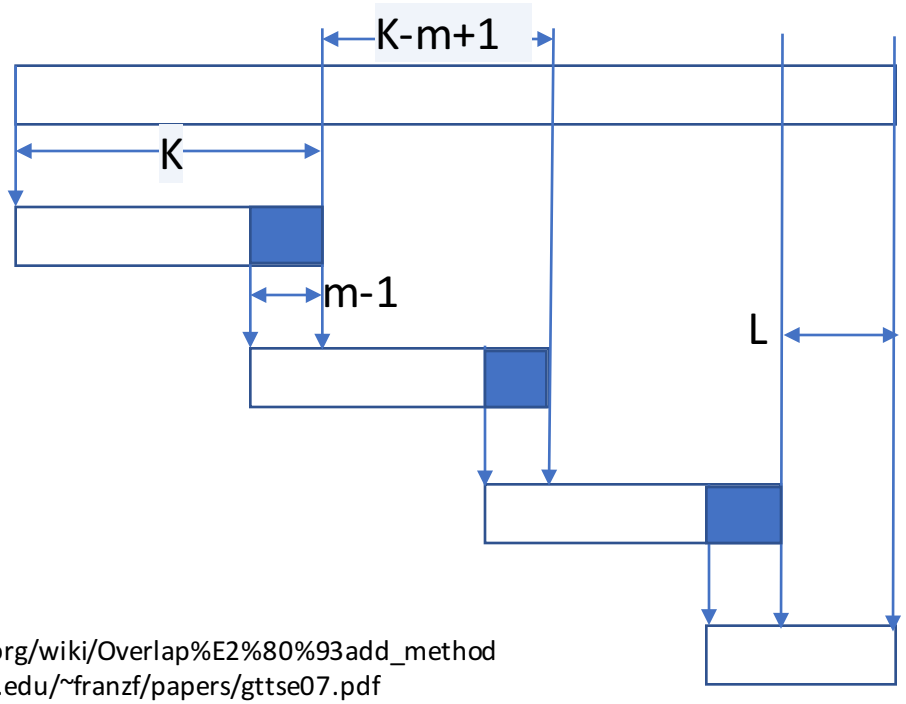

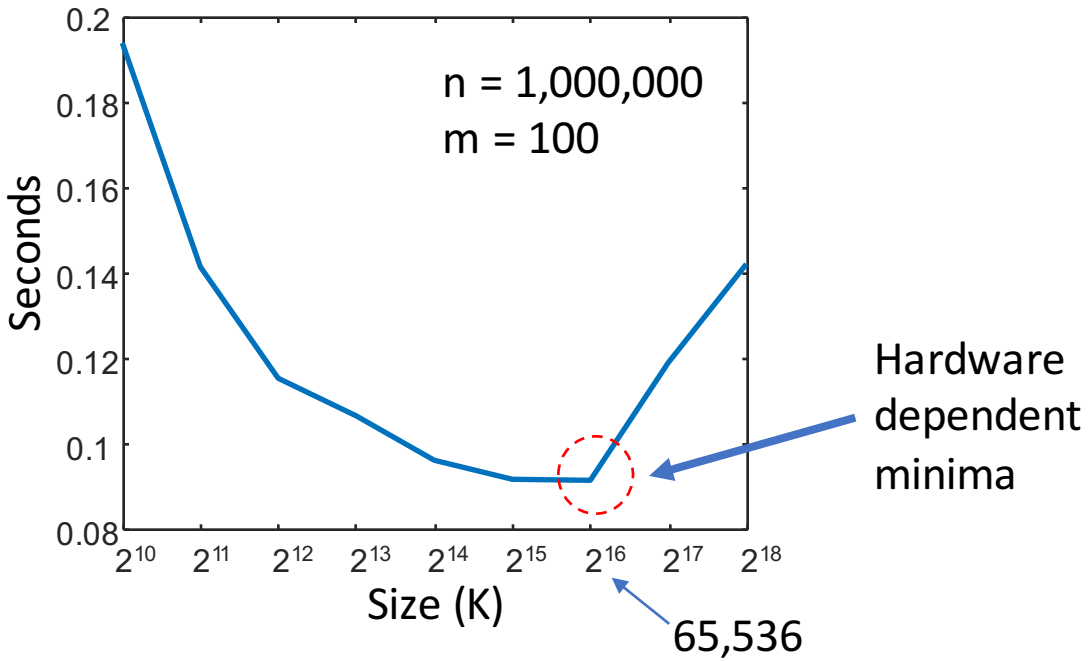

[a] https://en.wikipedia.org/wiki/Overlap%E2%80%93add_method
[b] https://users.ece.cmu.edu/~franzf/papers/gttse07.pdf

Mueen's Algorithm for Similarity Search (MASS) (8 of 9)

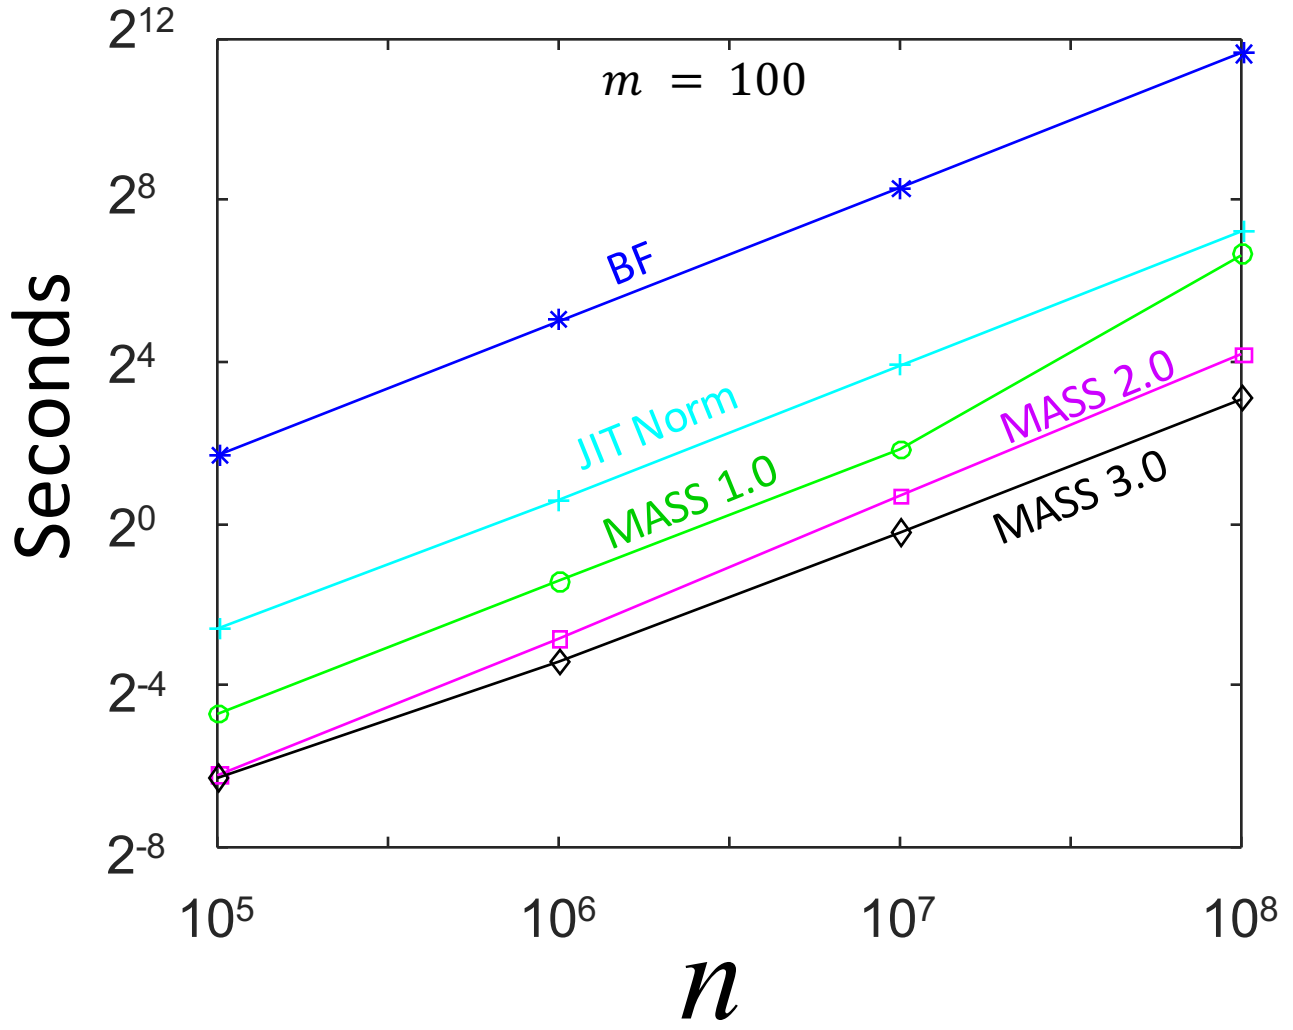

Mueen's Algorithm for Similarity Search (MASS) (9 of 9)

- Design a new application specific processor to do FFT efficiently [a][b].
- Keep the general processor and optimize DRAM to compute FFT efficiently [c].
- Keep a commodity box, add a GPU to do FFT efficiently [d].
- Keep a commodity box, add an FPGA to do FFT efficiently [e].
- Take a commodity box, and use sparse FFT [f].

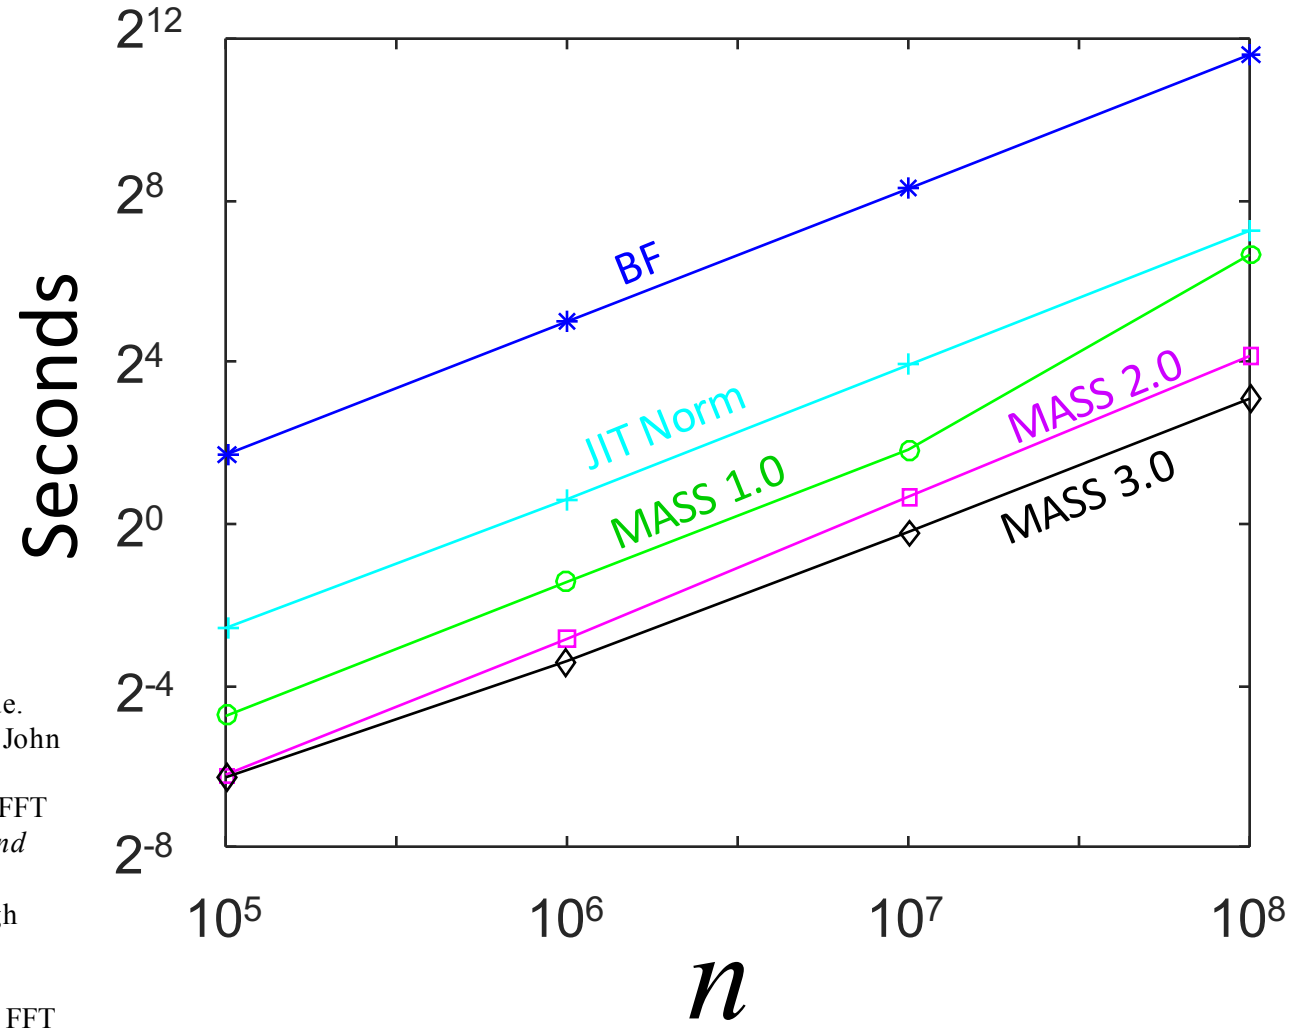

[a] The Fast Fourier Transform in Hardware: A Tutorial Based on an FPGA Implementation, George Slade.
[b] Appendix C: Efficient Hardware Implementations of FFT Engines, Nasserbakht, Mitra (Ed. Bingham, John A. C.) ADSL, VDSL, and Multicarrier Modulation, John Wiley & Sons, Inc. 2001
[c] B. Akın, F. Franchetti and J. C. Hoe, "Understanding the design space of DRAM-optimized hardware FFT accelerators," *2014 IEEE 25th International Conference on Application-Specific Systems, Architectures and Processors*, Zurich, 2014, pp. 248-255.
[d] Naga K. Govindaraju, Brandon Lloyd, Yuri Dotsenko, Burton Smith, and John Manferdelli. 2008. High performance discrete Fourier transforms on graphics processors. In *Proceedings of the 2008 ACM/IEEE conference on Supercomputing* (SC '08).
[e] Ren Chen, Shreyas G. Singapura, and Viktor K. Prasanna. 2017. Optimal dynamic data layouts for 2D FFT on 3D memory integrated FPGA. *J. Supercomput.* 73, 2 (February 2017), 652-663.
[f] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. 2012. Simple and practical algorithm for sparse Fourier transform. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete algorithms* (SODA '12). 1183-1194.

Numerical Errors in MASS

- Three possible sources of numerical errors in MASS
- Convolution
  - Numerical error in Convolution can appear if n is very large, MASS 3.0 reduces such error by dividing the computation in batches
- Division by the standard deviation
  - If a subsequence is constant, the standard deviation is zero, causing divide by zero errors. MASS reduces such error by checking the standard deviations ahead of the division.
- Square root of the distances
  - Theoretically it is not possible to have negative squared distance. However, for exact matches, the squared distance can be a very small negative number (-1.09e-15), resulting imaginary numbers in the output.

```
d(1:n) = 0;
Q = zNorm(query);
Stdv = movstd(T,[0 m-1]);
Q = Q(end:-1:1);
Q(m+1:n) = 0;
dots = ifft( fft(T).* fft(Q) );
dist = 2*(m-(dots(m:n))./Stdv));
dist = sqrt(dist);
```

Outline

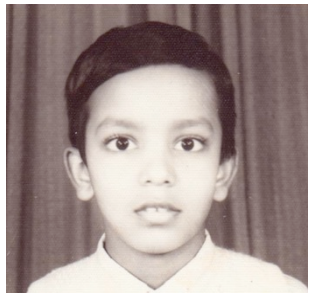

Act 1

- Our Fundamental Assumption
- What is the (MP) Matrix Profile?
- Properties of the MP
- Developing a Visual Intuition for MP
- Basic Algorithms
  - MP Motif Discovery
  - MP Time Series Chains
  - MP Anomaly Discovery
  - MP Joins (self and AB)
  - MP Semantic Segmentation
- From Domain Agnostic to Domain Aware: The Annotation Vector (A simple way to use domain knowledge to adjust your results)
- The "*Matrix Profile and ten lines of code is all you need*" philosophy.
- Break

Act 2

- Background on time series mining
  - Similarity Measures
  - Normalization
- Distance Profile
  - Brute Force Approach
  - Just-in-time Normalization
  - The MASS Algorithm
  - Extensions of MASS ⟵ Rest Zone
- Matrix Profile
  - STAMP
  - STOMP
  - GPU-STOMP
  - SCRIMP
- Open problems to solve

MASS Extensions: Approximate Distance Profile

- Can we trade speed with accuracy in MASS?
- We use piecewise aggregate approximation (PAA) to reduce dimensionality of both T and Q
- `sqrt(w)* MASS(PAA(T,w),PAA(Q,w),m/w)`

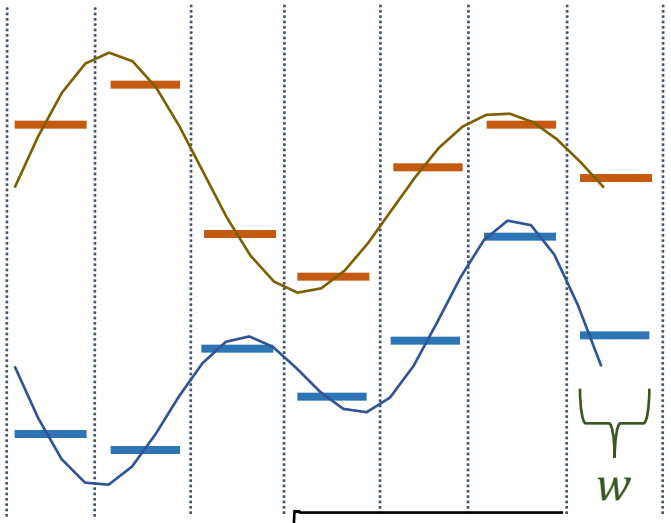

$$d(x,y) = \sqrt{\sum_{i=1}^{m}(x_i - y_i)^2}$$

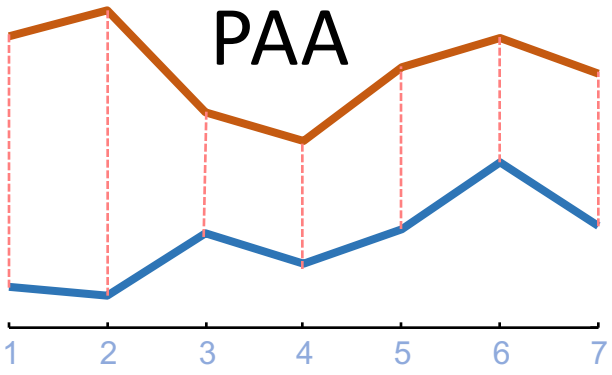

$$LB(x,y) = \sqrt{w\sum_{i=1}^{m/w}(x_i - y_i)^2}$$

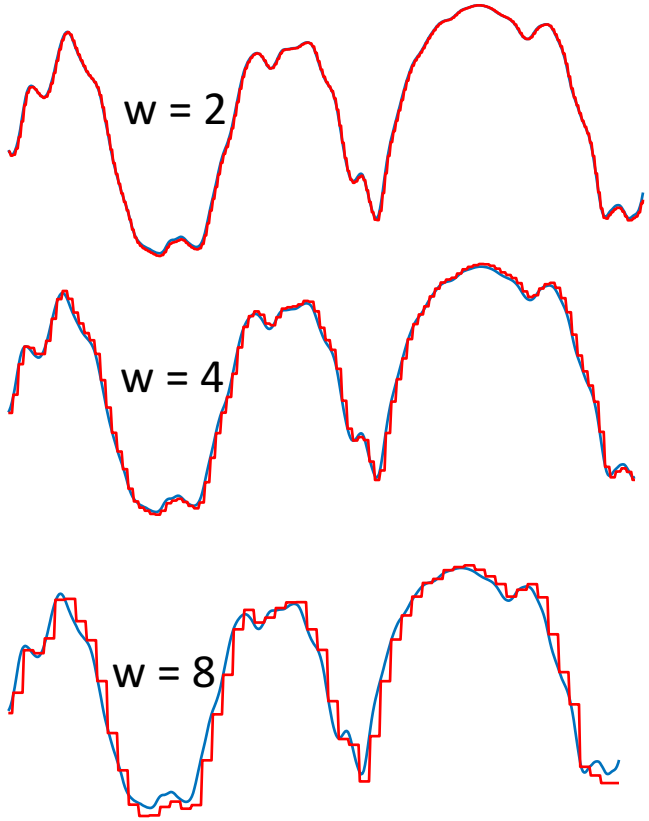

Approximate distance profile in Red

Selina Chu, Eamonn J. Keogh, David M. Hart, Michael J. Pazzani: Iterative Deepening Dynamic Time Warping for Time Series. SDM 2002: 195-212

MASS Extensions: Weighted Query

Given two time series

$$\boldsymbol{x} = x_1 \ldots x_n$$

and

$$\boldsymbol{y} = y_1 \ldots y_n$$

their z-Normalized weighted Euclidean distance is defined as:

$$\hat{x}_i = \frac{x_i - \mu_x}{\sigma_x} \qquad \hat{y}_i = \frac{y_i - \mu_y}{\sigma_y} \qquad d^2(x,y) = \sum_{i=1}^{n} w_i (\hat{x}_i - \hat{y}_i)^2$$

- $|w| = m$
- $\forall\, w_i \geq 0$
- $\sum w_i = 1$

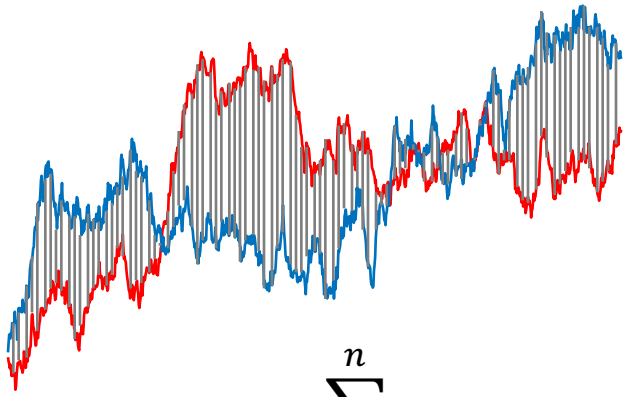

$$d(x,y) = \sum_{i=1}^{n} (\hat{x}_i - \hat{y}_i)^2$$

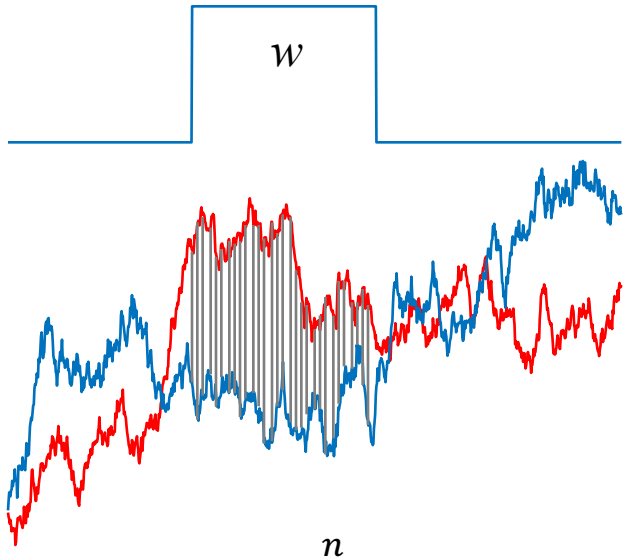

$$d^2(x,y) = \sum_{i=1}^{n} w_i (\hat{x}_i - \hat{y}_i)^2$$

MASS Extensions: Weighted Query

- Working formula that we have been using so far, will not work anymore.
- The new working formula is below.
- We need three sliding dot products instead of *one* in regular MASS

$$d^2(x,y) = \sum_{i=1}^{n} w_i(\hat{x}_i - \hat{y}_i)^2 = \sum_{i=1}^{n} w_i(\frac{x_i - \mu_x}{\sigma_x} - \hat{y}_i)^2$$

$$= \sum_{i=1}^{n} w_i x_i^2 - 2\mu_x w_i x_i - 2w_i x_i \hat{y}_i + \mu_x^2 w_i - 2\mu_x w_i \hat{y}_i + w_i \hat{y}_i{}^2$$

Terms with x need sliding dot products

Terms without x are precomputed

MASS Extensions: Weighted Query Example

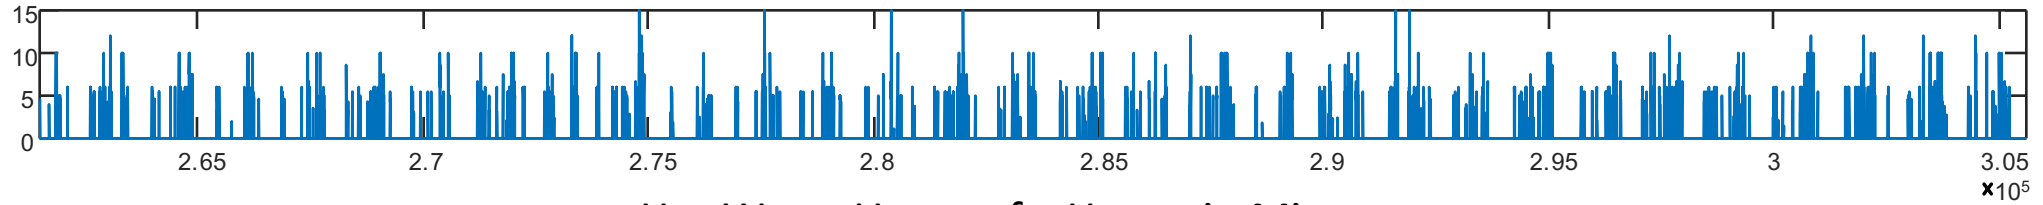

Hot Water Usage of a House in Minutes

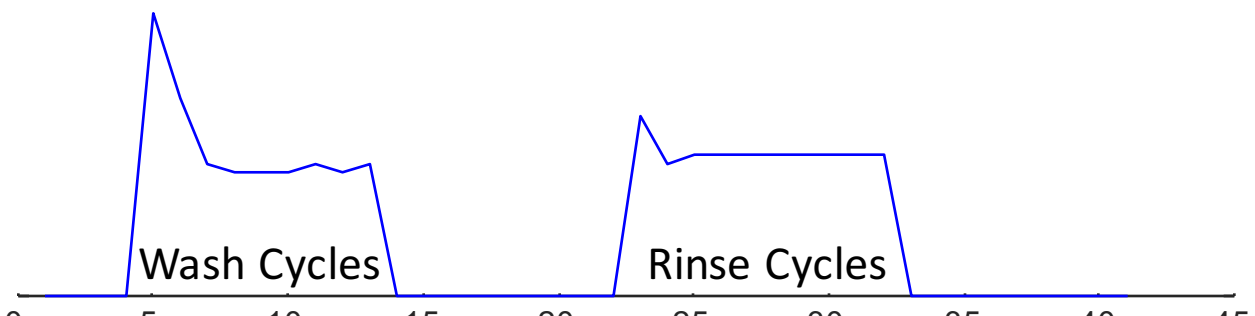

Washer

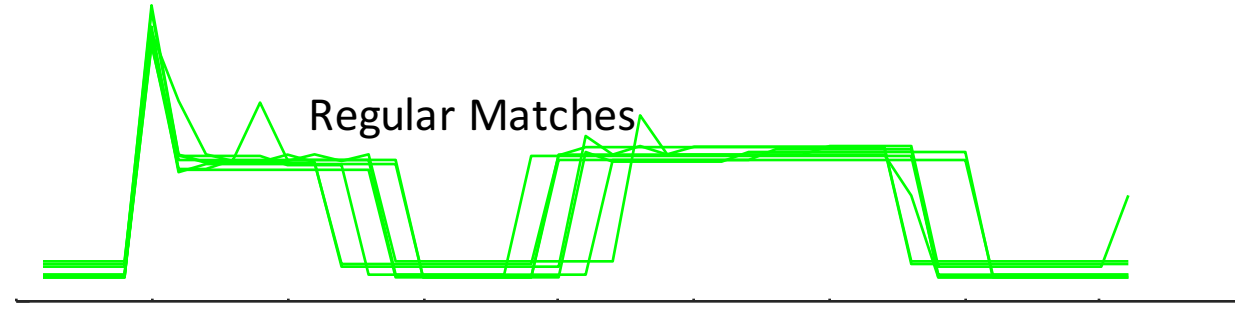

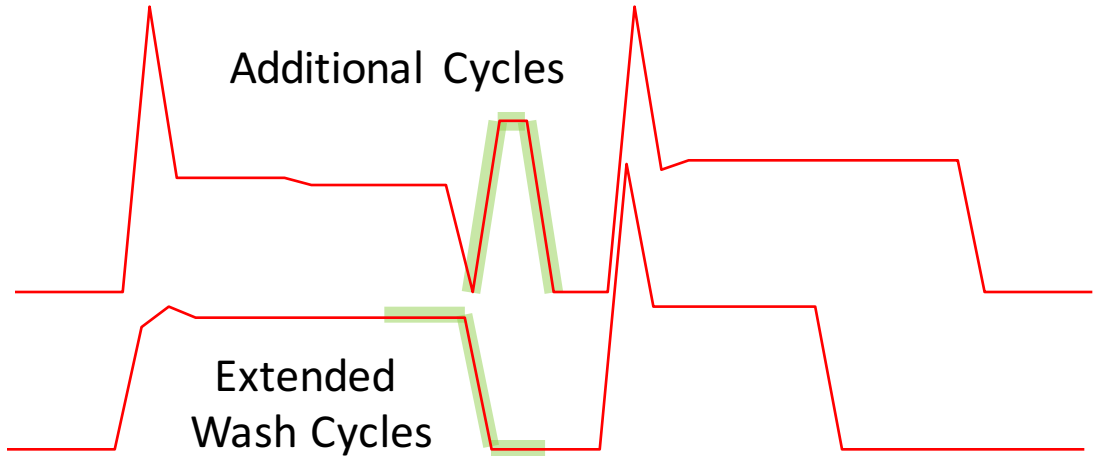

MASS Extensions: Query with Gap(s)

We may wish to consider a more general type of query. Queries with gaps, as in the purple example.

I want the first third to slope down, and the last third to slope up with the plateau at the end. However, I don't care what goes in the middle.

Either of the two red patterns would be acceptable matches.

Note that the relative offset and scaling of the two parts do *not* matter, and we should be invariant to them. In other words, the exact slope does not matter in the above query statement.

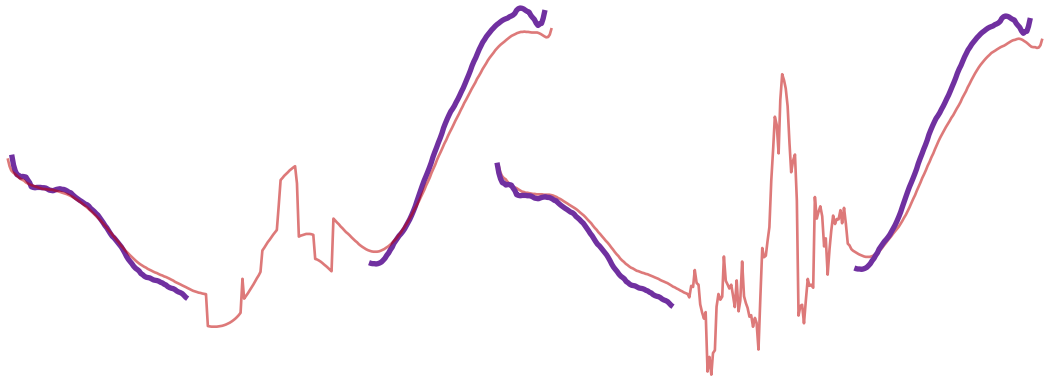

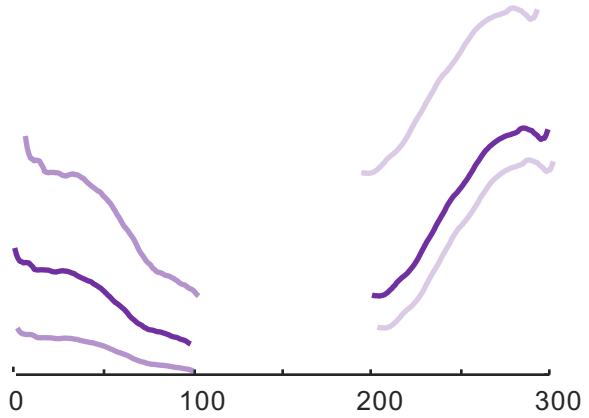

MASS Extensions: Query with Gap(s)

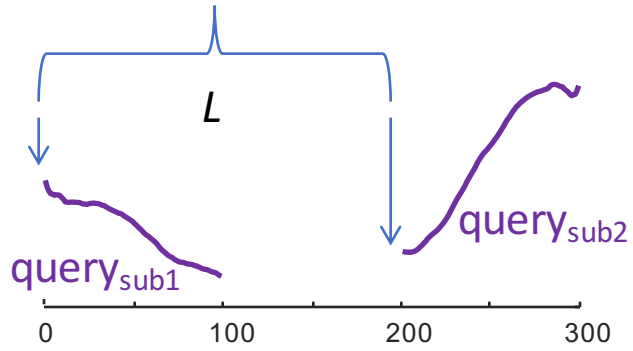

1. Measure the distance between beginning of the two patterns, call it *L*
2. Compute…

```
Dist_profile1  = MASS(T, query_sub1);
Dist_profile2  = MSAS(T, query_sub2);
```

3. Pad the *end* of `Dist_profile1` and the *beginning* of `Dist_profile2` with a vector of *L* infinities to add a lag of *L*

   For example, if *L* = 5, then use {inf, inf, inf, inf, inf}

4. Create `Dist_profile3 = Dist_profile1 + Dist_profile2`
5. We can now use `Dist_profile3` to find the nearest neighbor(s) as before.

MASS Extensions: Query with Gap(s)

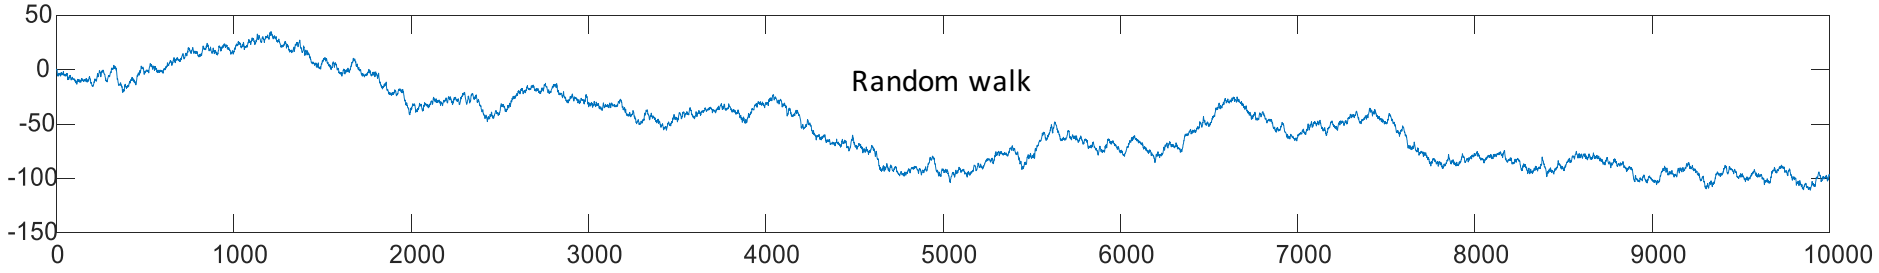

Dist_profile1 = MASS(tag, $query_{sub1}$);

Slide by L slots

Dist_profile2 = MASS(tag, $query_{sub2}$);

Dist_profile3 = Dist_profile1 + Dist_profile2

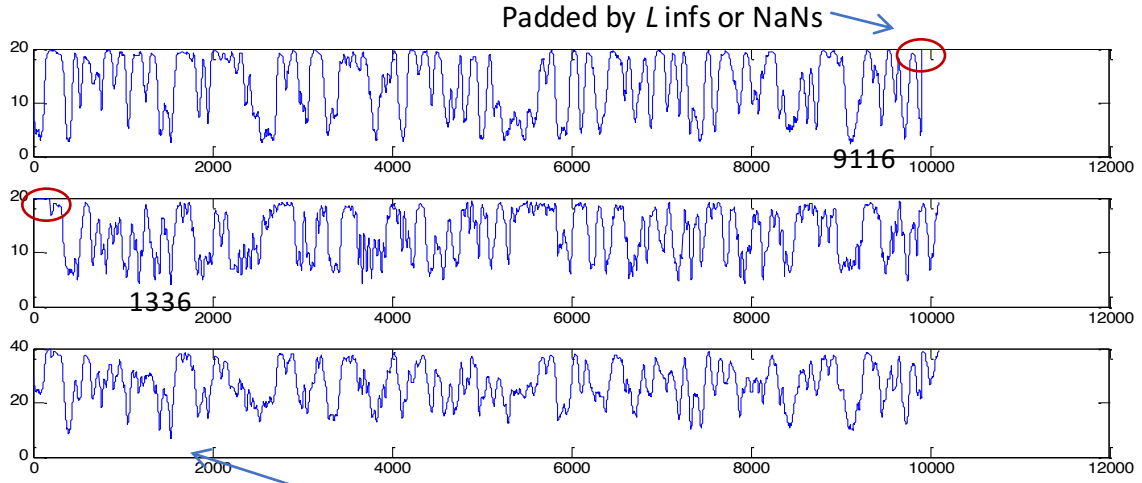

Note that for the best match, each part of the query scaled differently.
For the image to the right, I fixed the left side and scaled the right side (to give context), but both sides needed rescaling somewhat

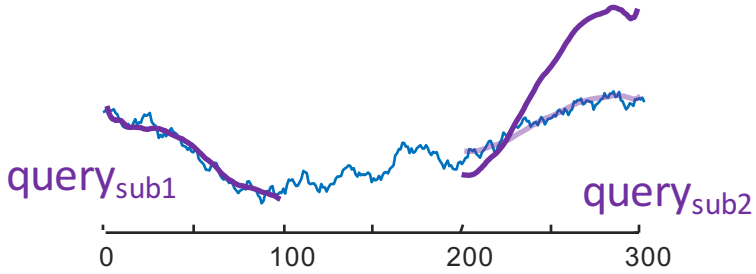

Outline

Act 1

- Our Fundamental Assumption
- What is the (MP) Matrix Profile?
- Properties of the MP
- Developing a Visual Intuition for MP
- Basic Algorithms
  - MP Motif Discovery
  - MP Time Series Chains
  - MP Anomaly Discovery
  - MP Joins (self and AB)
  - MP Semantic Segmentation
- From Domain Agnostic to Domain Aware: The Annotation Vector (A simple way to use domain knowledge to adjust your results)
- The "*Matrix Profile and ten lines of code is all you need*" philosophy.
- Break

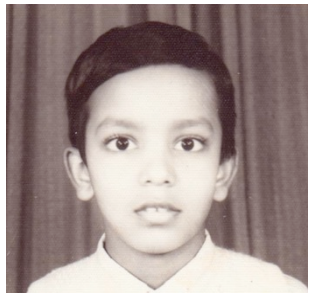

Act 2

- Background on time series mining
  - Similarity Measures
  - Normalization
- Distance Profile
  - Brute Force Approach
  - Just-in-time Normalization
  - The MASS Algorithm
  - Extensions of MASS
- Matrix Profile
  - STAMP
  - STOMP
  - GPU-STOMP
  - SCRIMP
- Open problems to solve

Matrix Profile from Distance Profiles

Query, the 1st subsequence in the time series

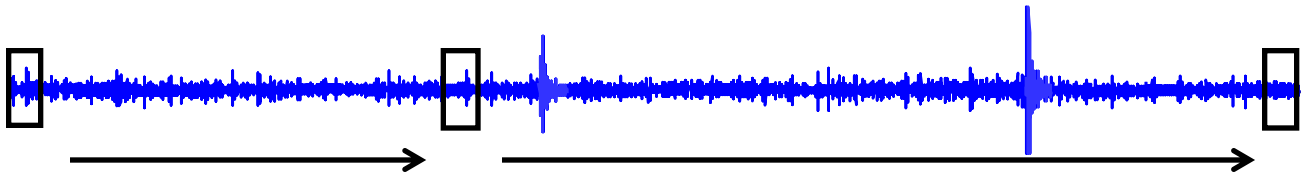

Obtain the z-normalized Euclidean distance between **Query** and each window (subsequence) in the time series. We would obtain a vector like this:

$d_{1,1}$(= 0)	$d_{2,1}$	...	$d_{n-m+1,1}$	⟶	D_1

$d_{i,j}$ is the distance between the i^{th} subsequence and the j^{th} subsequence.

We can obtain D_2, D_3, … D_{n-m+1} similarly.

Matrix Profile from Distance Profiles

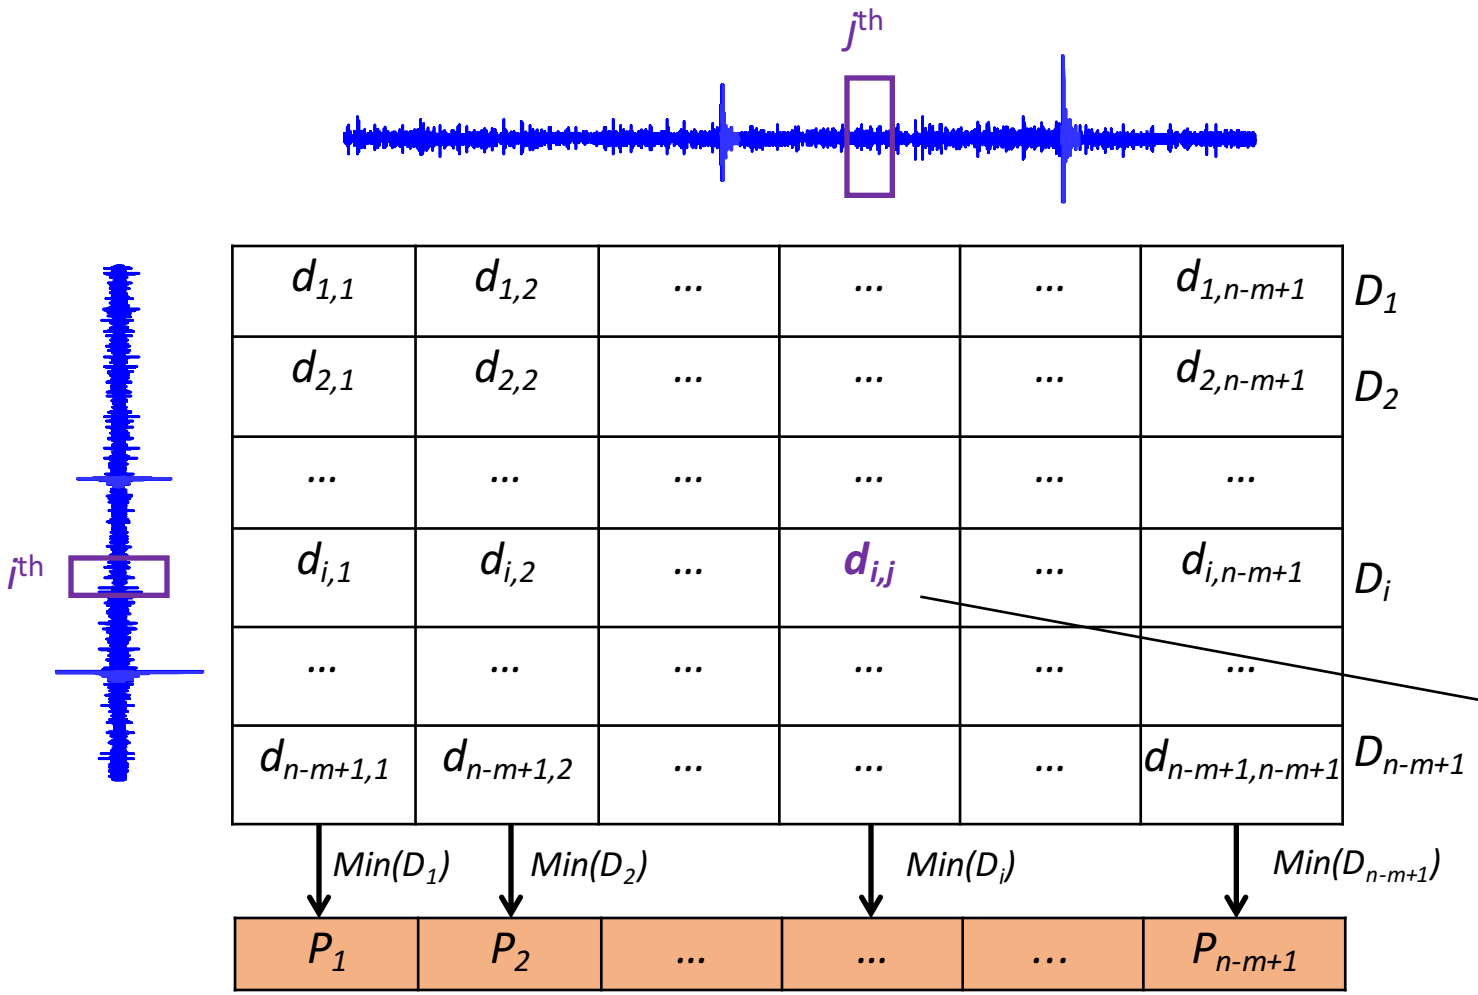

The distance matrix is symmetric
The diagonal is zero
Cells close to the diagonal are very small

$d_{i,j}$ is the distance between the i^{th} window and the j^{th} window of the time series

Matrix Profile: a vector of distance between each subsequence and its nearest neighbor

STMP: Scalable Time Series Matrix Profile Algorithm

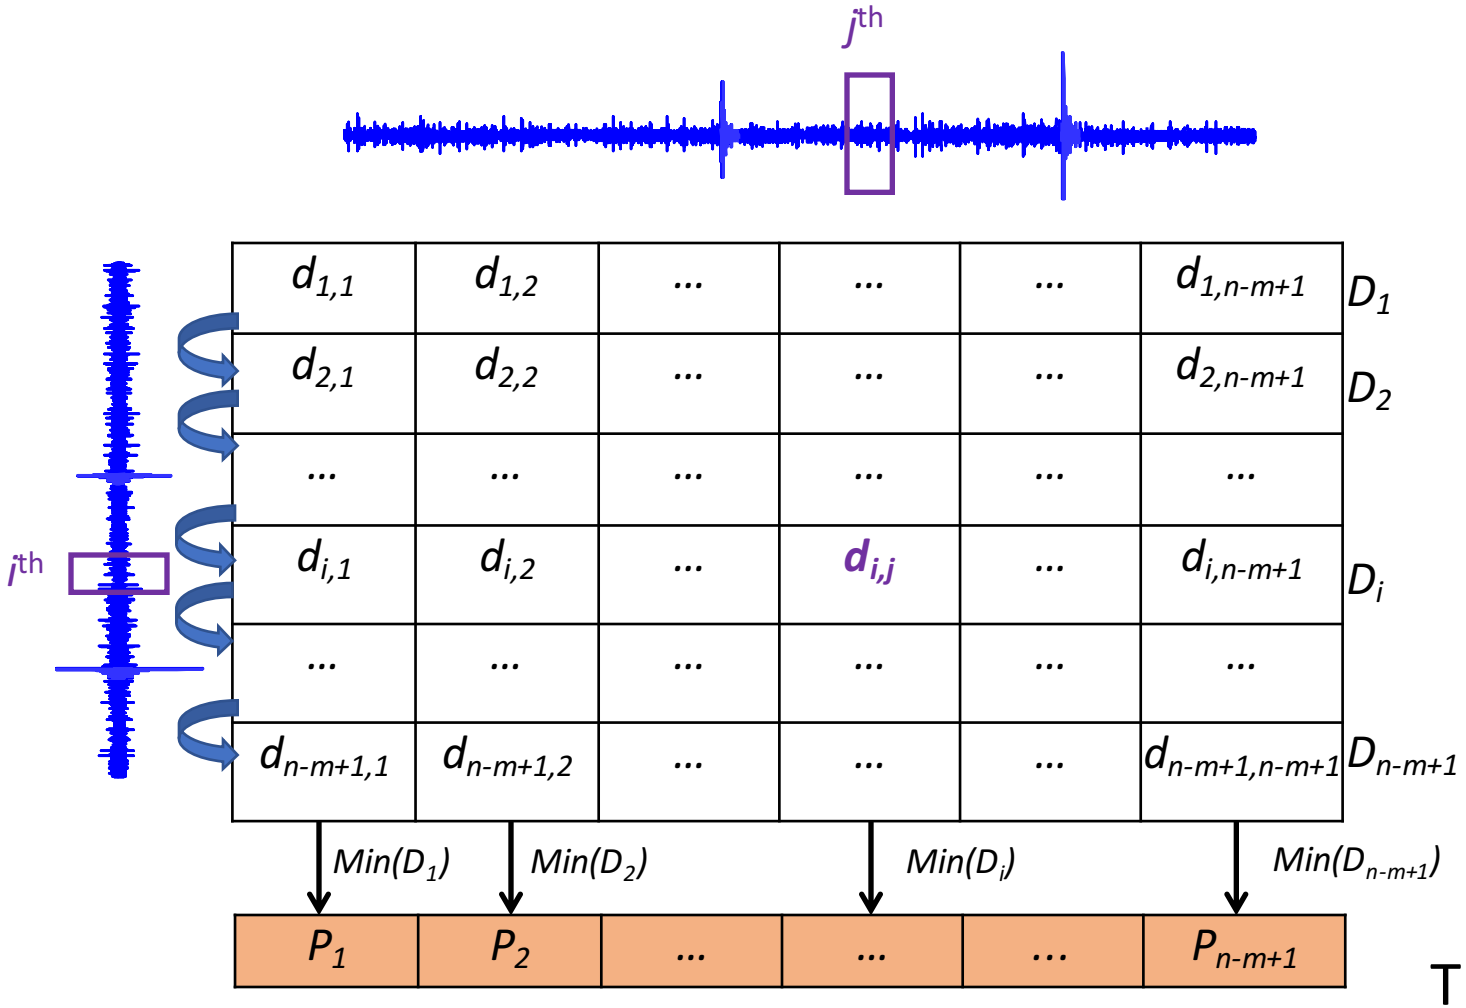

Matrix Profile: a vector of distance between each subsequence and its nearest neighbor

```
MP(1:n-m+1) = inf;
for i = 1:n-m+1
        d = MASS(T,T(i:i+m-1));
        MP = min([MP ; d]);
end
```

Time complexity of STMP is $O(n^2 \log n)$
Space complexity of STMP is $O(n)$

Matrix Profile Index

- Matrix Profile Index (MPI) can be maintained as we compute the profile
- Vectorized `min` functions can be used to efficiently maintain MPI
- Overhead is negligible

```
MP(1:n-m+1) = inf;
MPI(1:n-m+1) = -1;
for i = 1:n-m+1 in a random order
    d = MASS(T,T(i:i+m-1));
    [MP, ind] = min([MP ; d]);
    MPI(ind==2) = i;
end
```

Assumes `MP` and `d` as column vectors

Trivial Match I

- We need to understand the idea of a “trivial match”. It shows up for definitions of discords, motifs, similarity search etc.
- Suppose we search for the **query**, of length 20, in the time series…
- …we find its best match is at location 50…

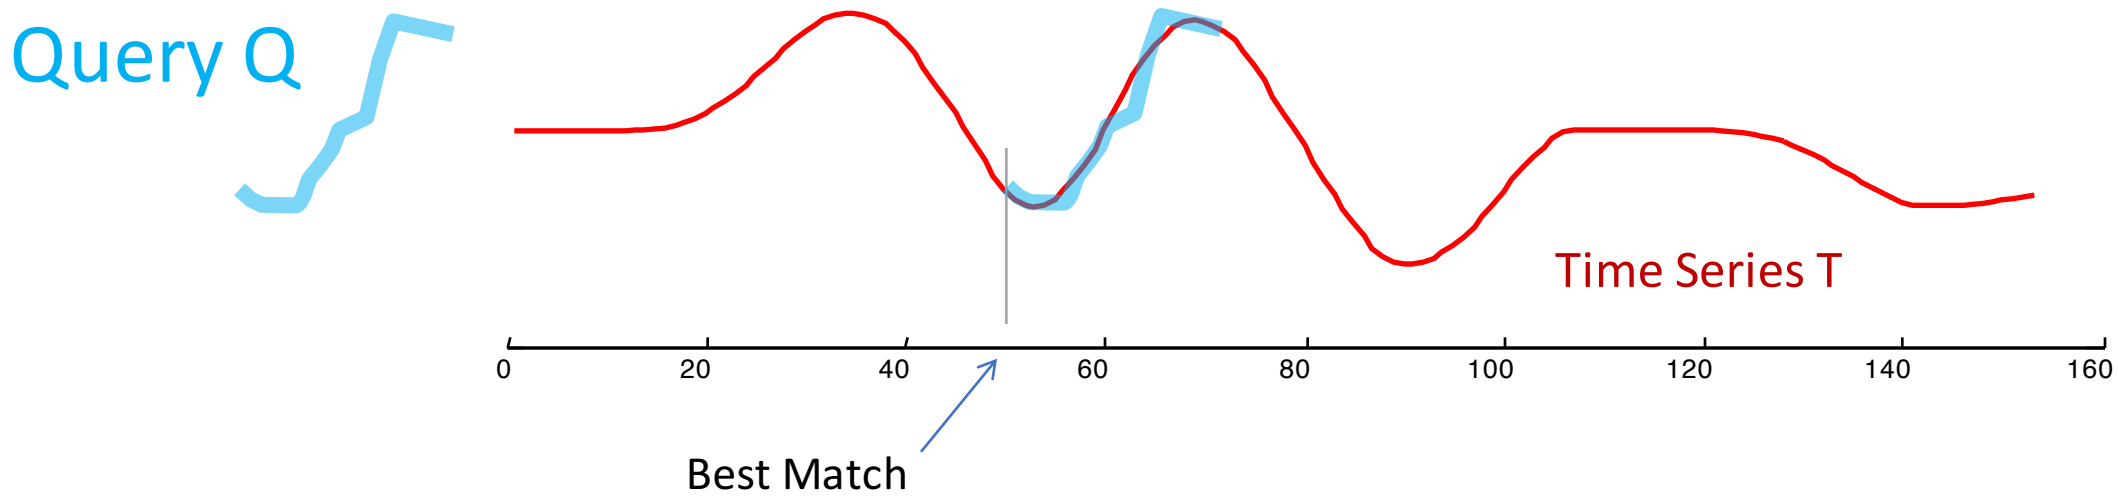

Trivial Match II

- Where is the second best match? It is probably going to be at 49 or 51, but that is *trivial*, it is just a minor "variant" of the original match.
- (try togging backwards and forwards between this and the last slide)

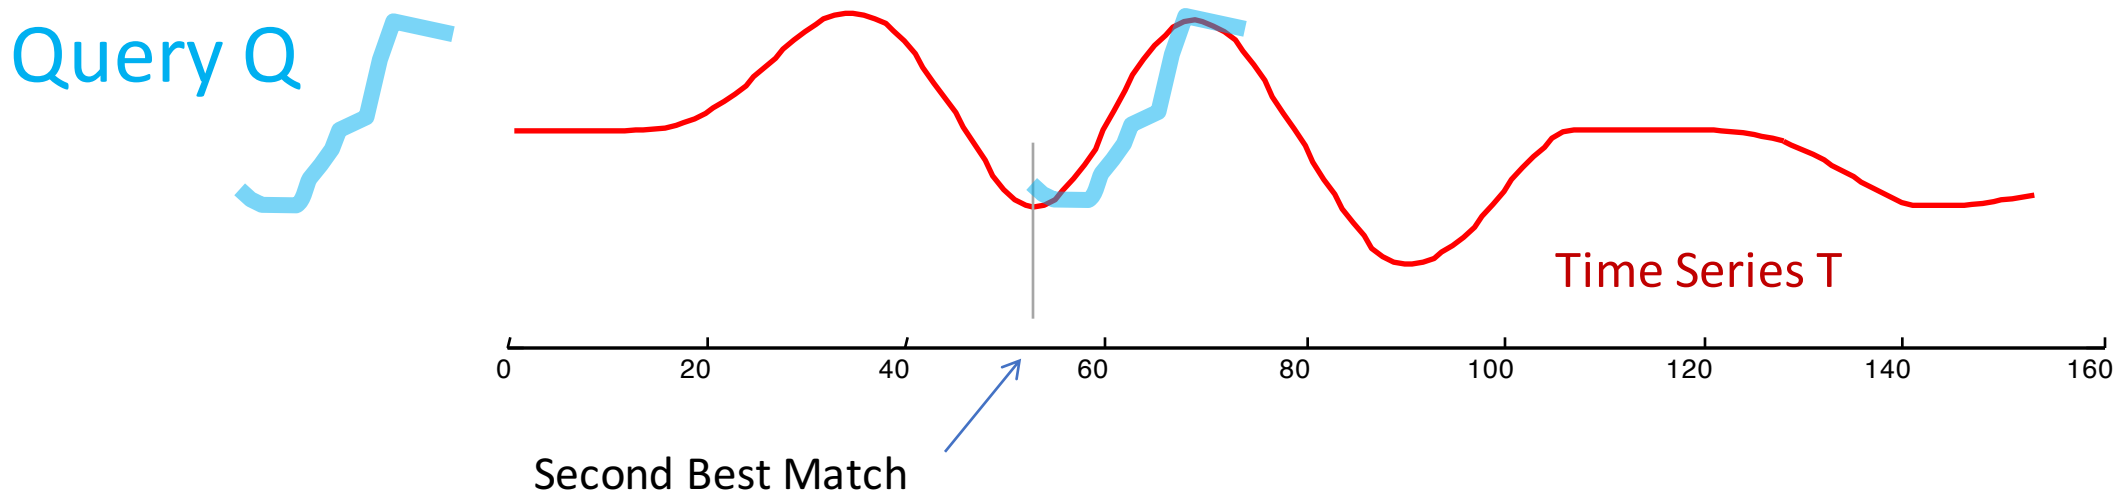

Trivial Match III

- To avoid trivial matches, we find the first match, then we set an *exclusion zone* around the best match, then find the *second* best match, etc.
- The size of the *exclusion zone* is not critical, ½ *m* works well.

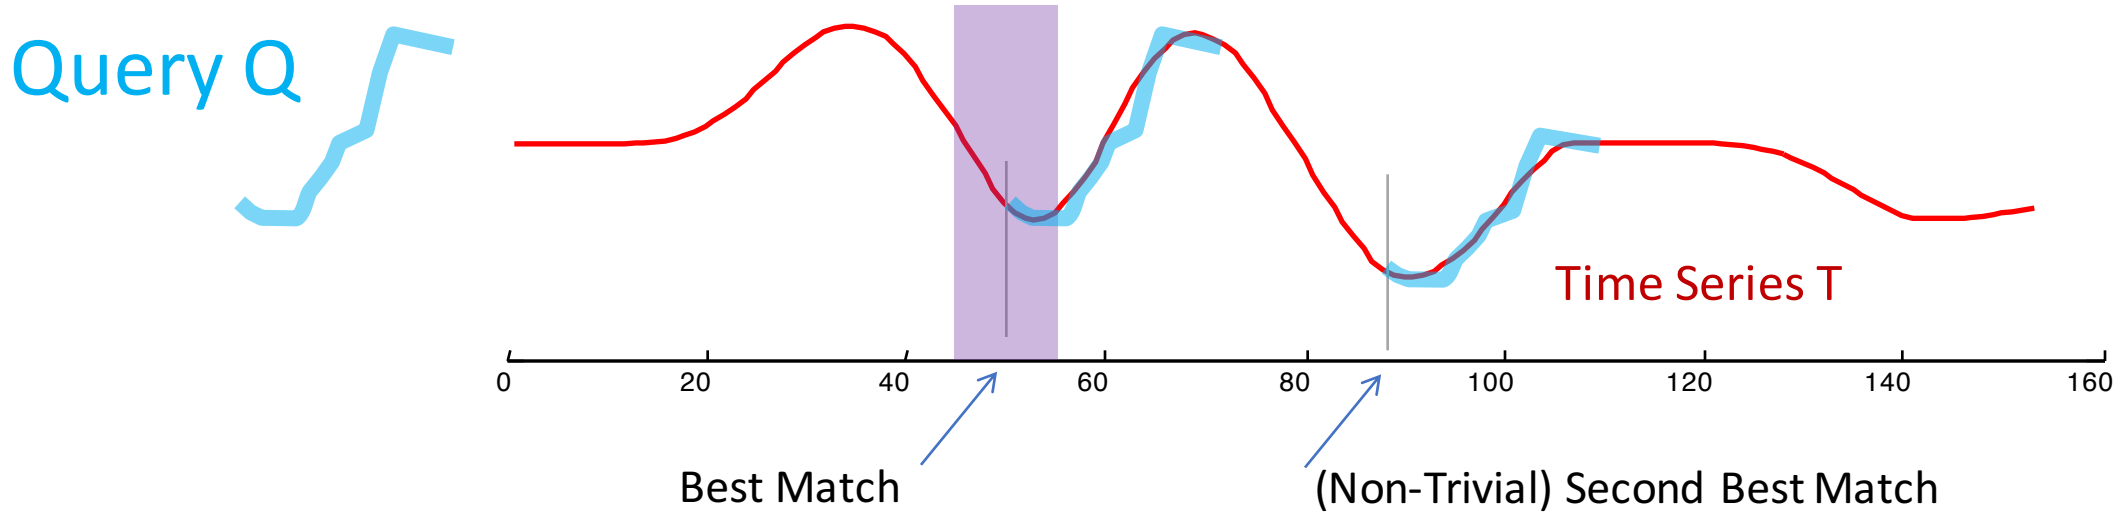

Trivial Match IIII Special Case

- When computing the MP, we will be extracting the subsequences from the time series *itself*.
- Clearly such queries will just find "themselves" as their own nearest neighbor!
- The distance from the query to any part of the subsequence that overlaps the *exclusion zone* is kept undefined (NaN).

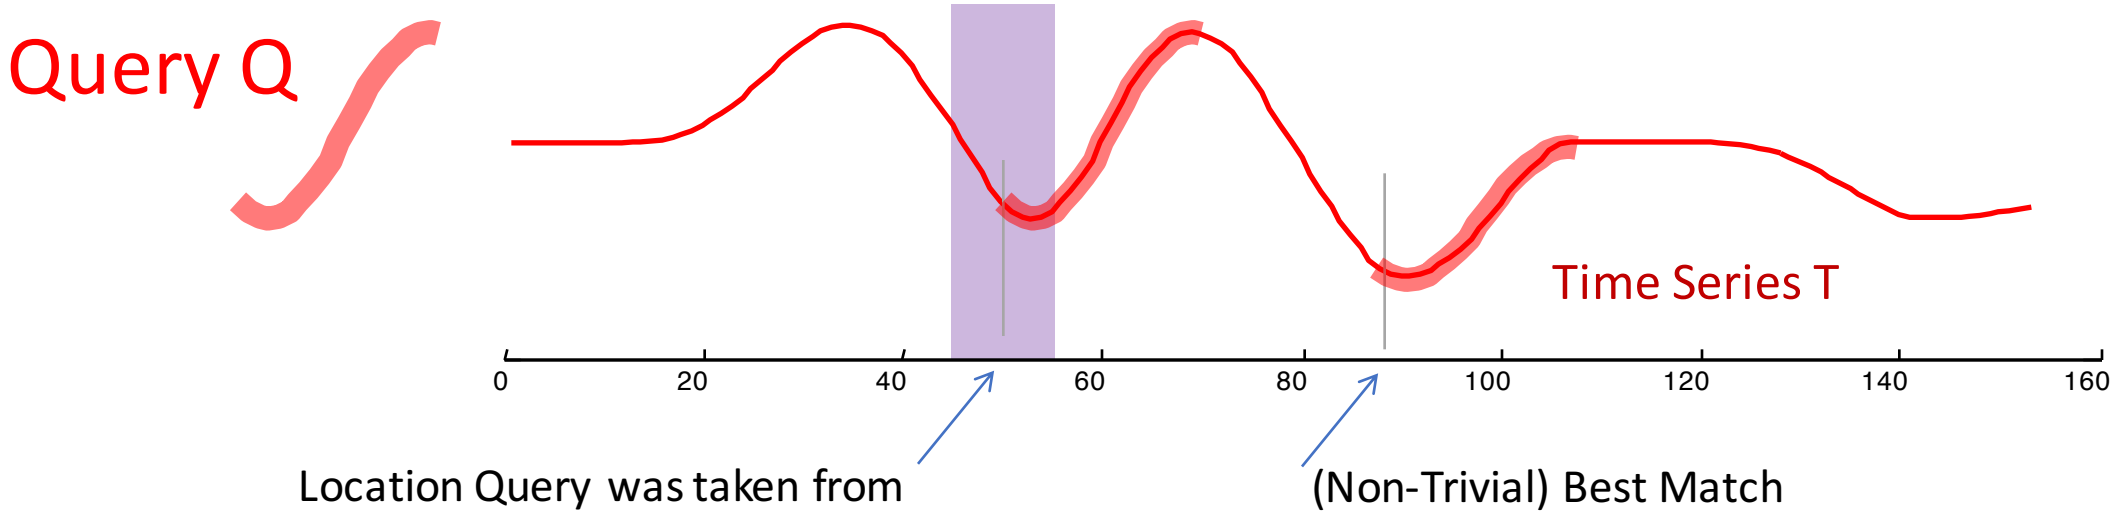

STMP: Scalable Time Series Matrix Profile Algorithm

```
MP(1:n-m+1) = inf;
MPI(1:n-m+1) = -1;
for i = 1:n-m+1
        d = MASS(T,T(i:i+m-1));
        d( max(i-m/4,1) : min(i+m/4-1,n-m+1))= NaN;
        [MP, ind] = min([MP ; d]);
        MPI(ind==2) = i;
end
```

Size of Exclusion Zone depends on the smoothness of the time series. A default value of m/2 works well. Data adaptive heuristic could be to set the length of the exclusion zone equal to m or the period (1/f) of the highest energy frequency, whichever is smaller.

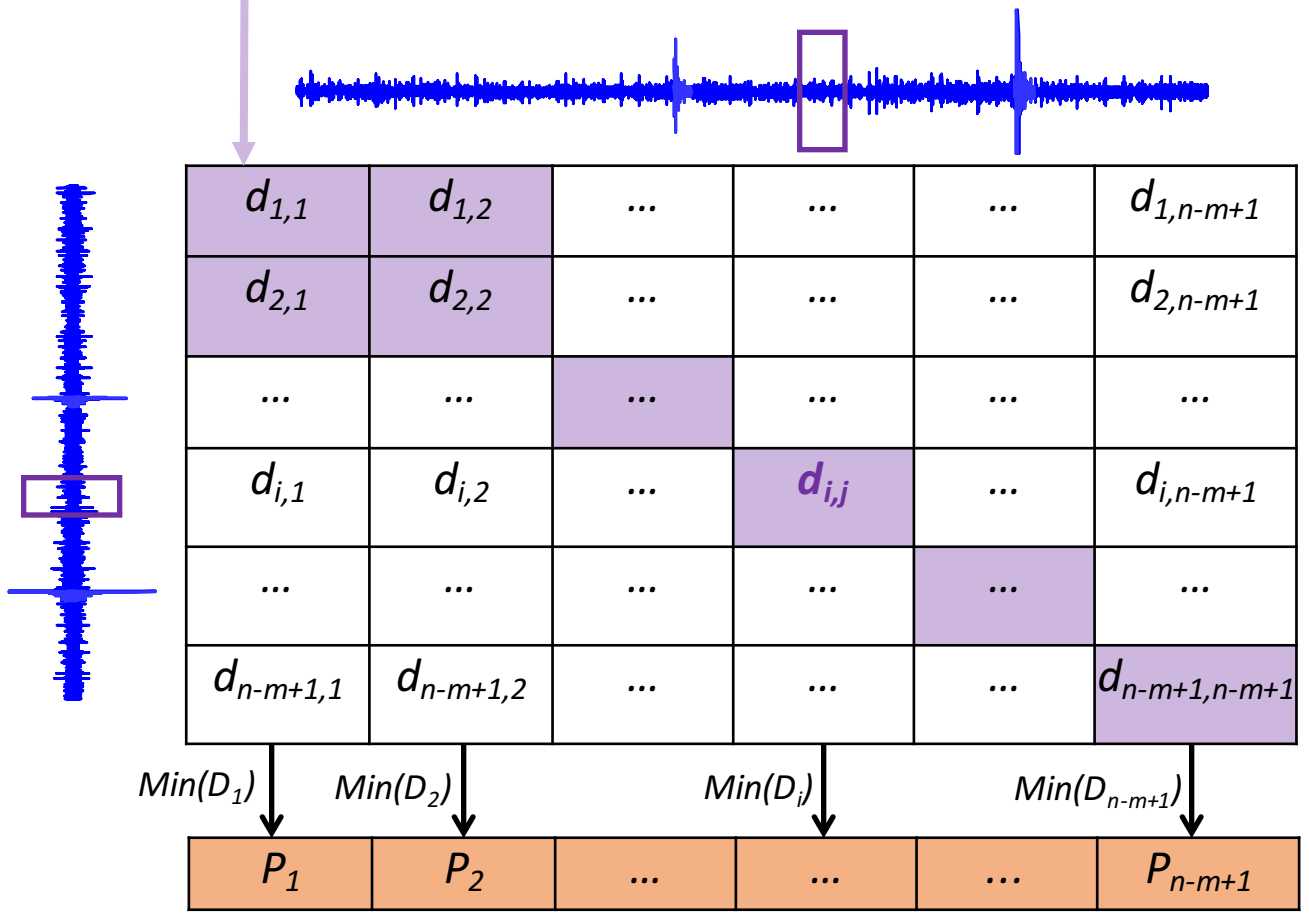

STAMP: Scalable Time Series Anytime Matrix Profile

- Each distance profile is independent of other distance profiles, the order in which we compute them can be random
- A random ordering averages the performance over properties of the time series, such as smoothness, frequency, stationarity, etc.
- The random ordering provides diminishing return, which allows *interrupt-resume* operations *anytime*.

```
MP(1:n-m+1) = inf;
MPI(1:n-m+1) = -1;
for i = 1:n-m+1 in a random order
       d = MASS(T,T(i:i+m-1));
       d( max(i-m/4,1) : min(i+m/4-1,n-m+1))= NaN;
       [MP, ind] = min([MP ; d]);
       MPI(ind==2) = i;
end
```

Exploring the Anytime Property (1 of 3)

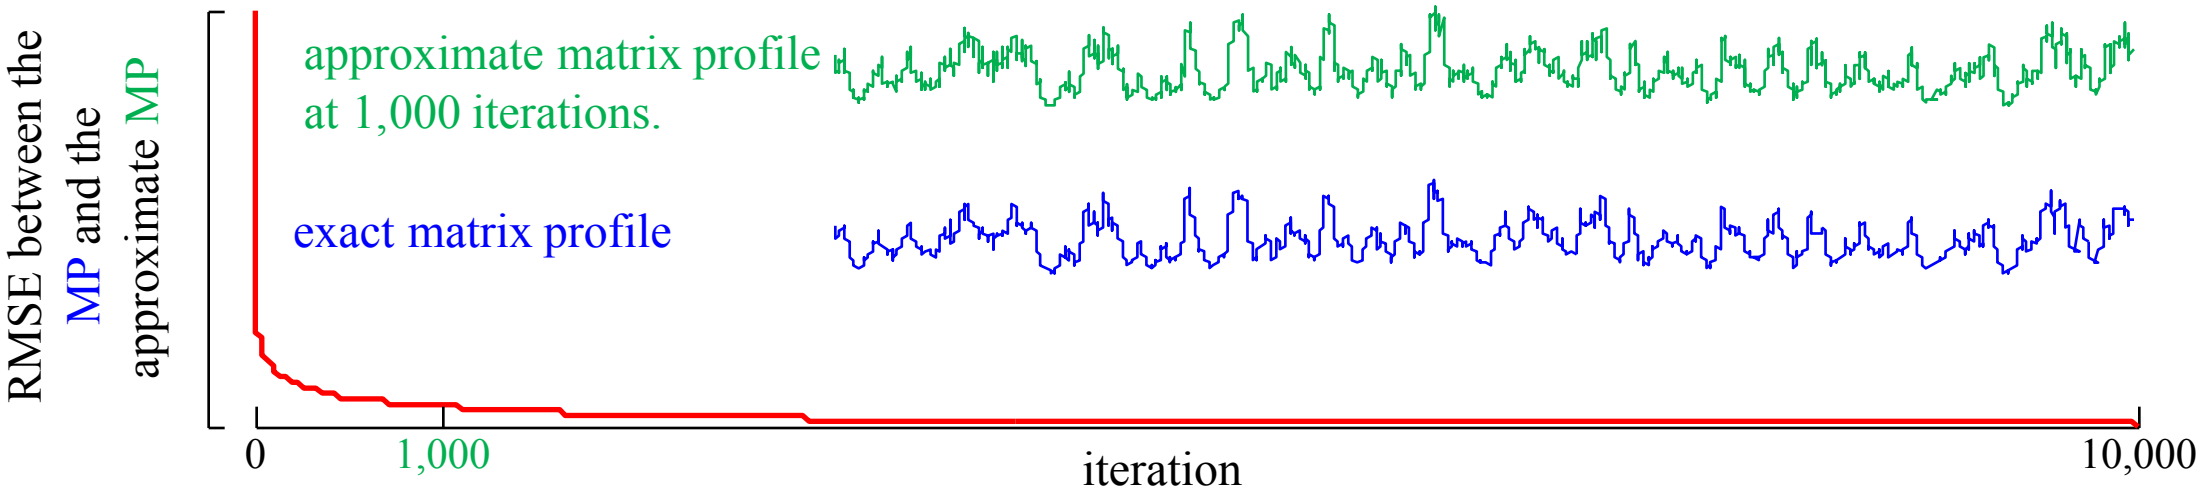

- The approximate matrix profile at 1,000 iteration is extremely similar to the exact solution.
- The convergence rate increases, for larger datasets.

Exploring the Anytime Property (2 of 3)

Play Video

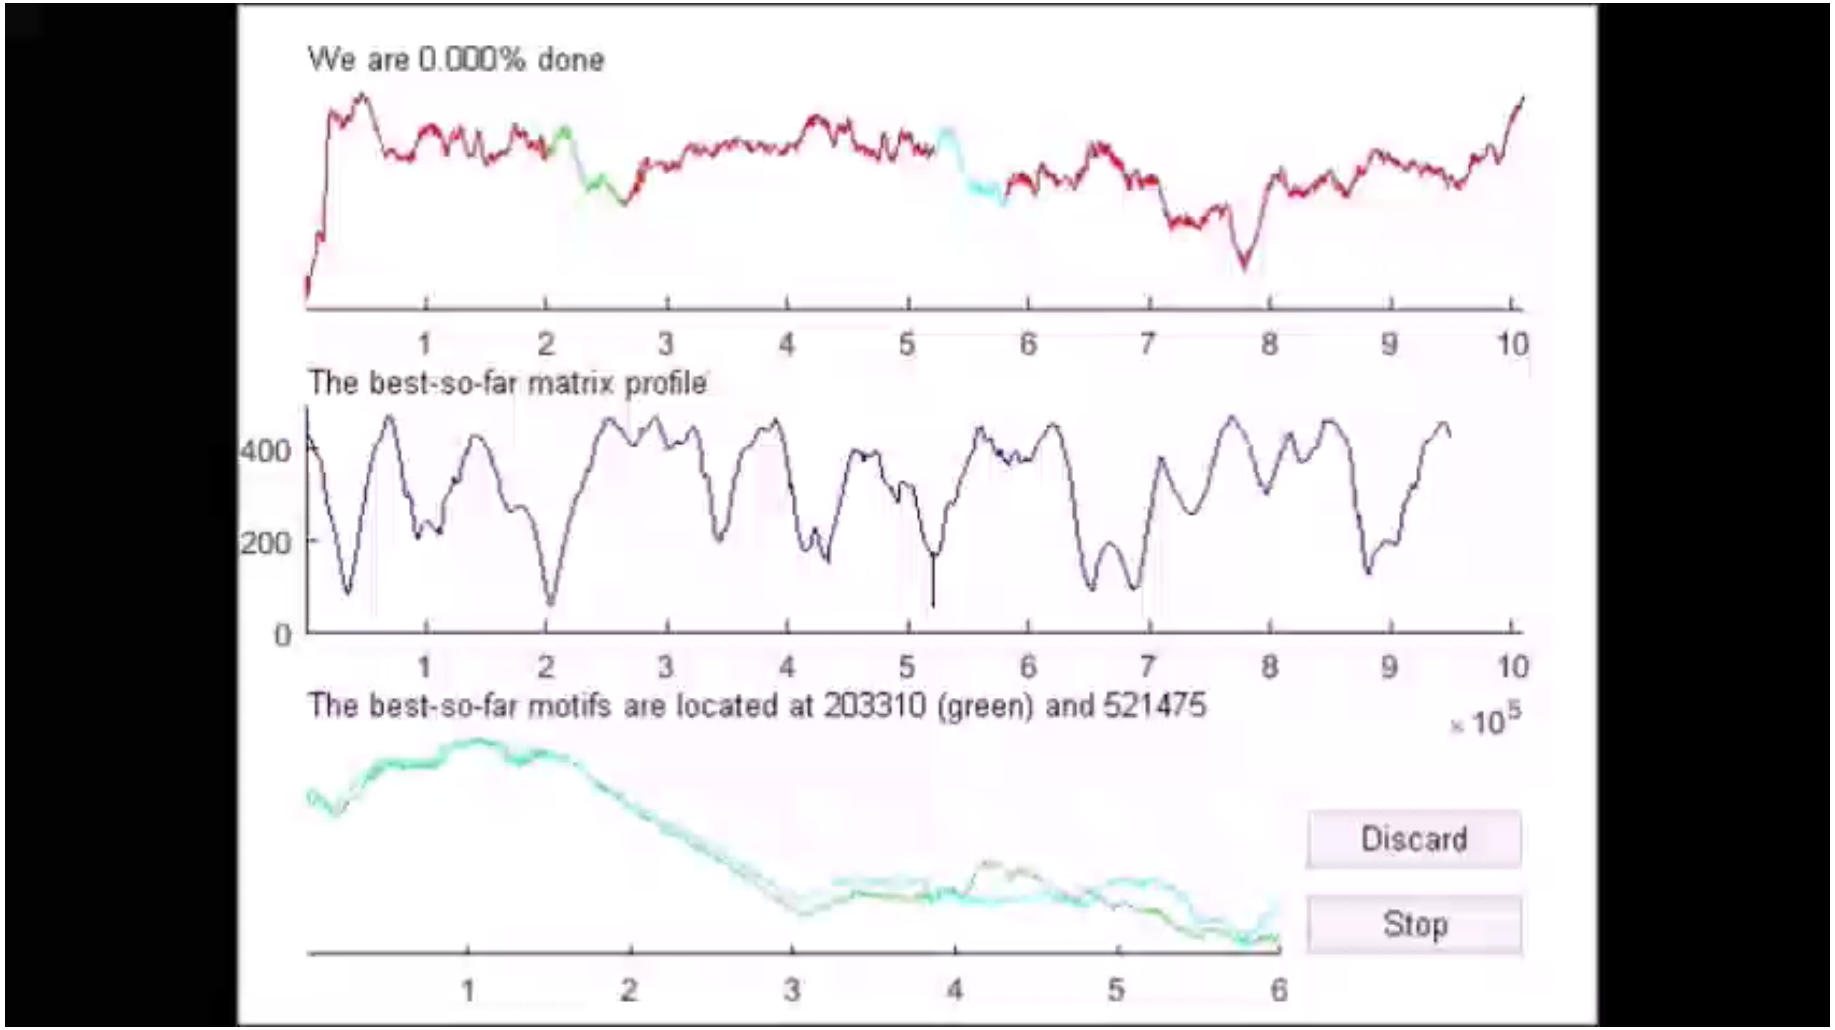

Exploring the Anytime Property (3 of 3)

After doing only 1/500th of the computations, the basic shape of the MP has converged, and we have found the final correct motif.

As an aside, note that the dimensionality of the subsequences here is 60,000.

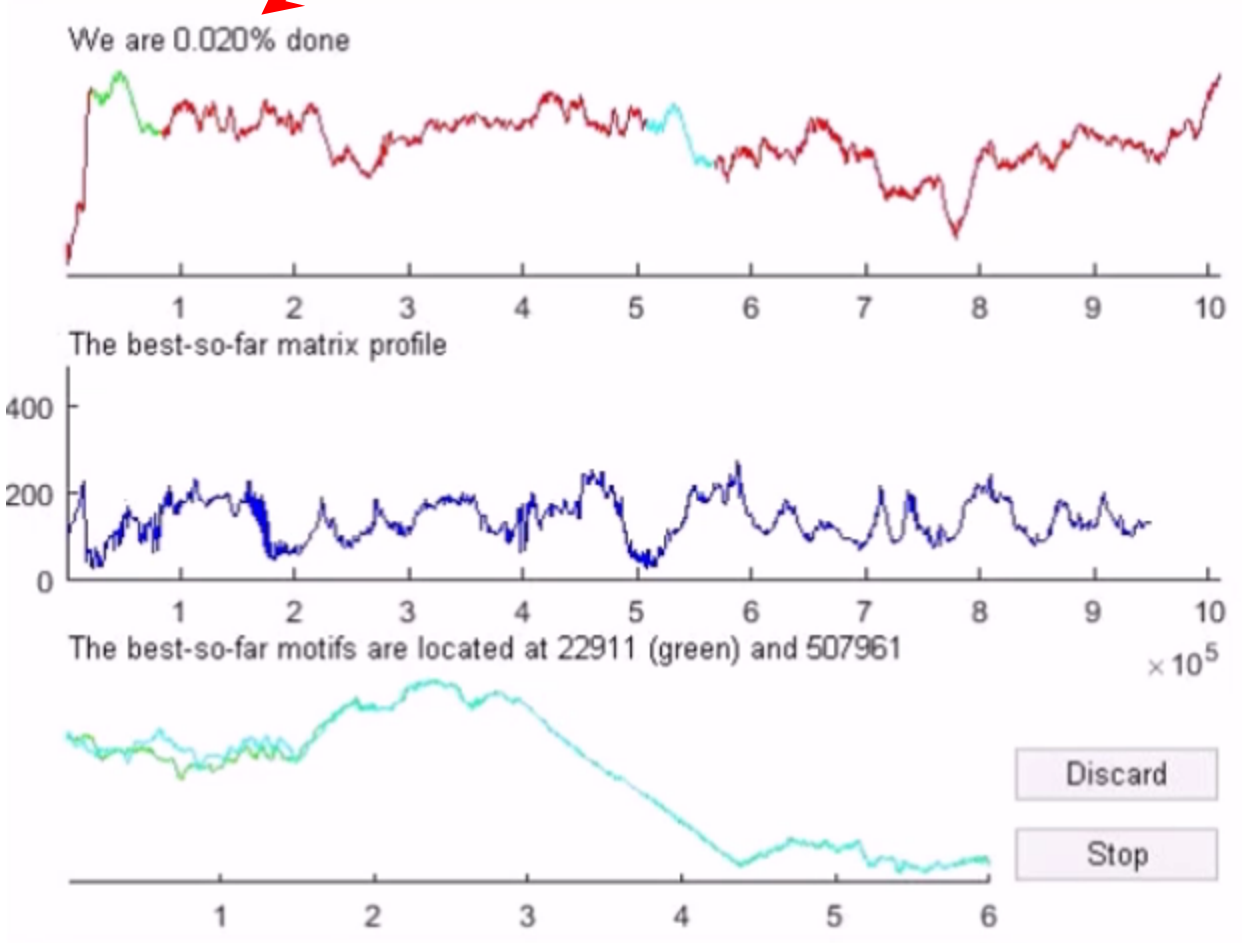

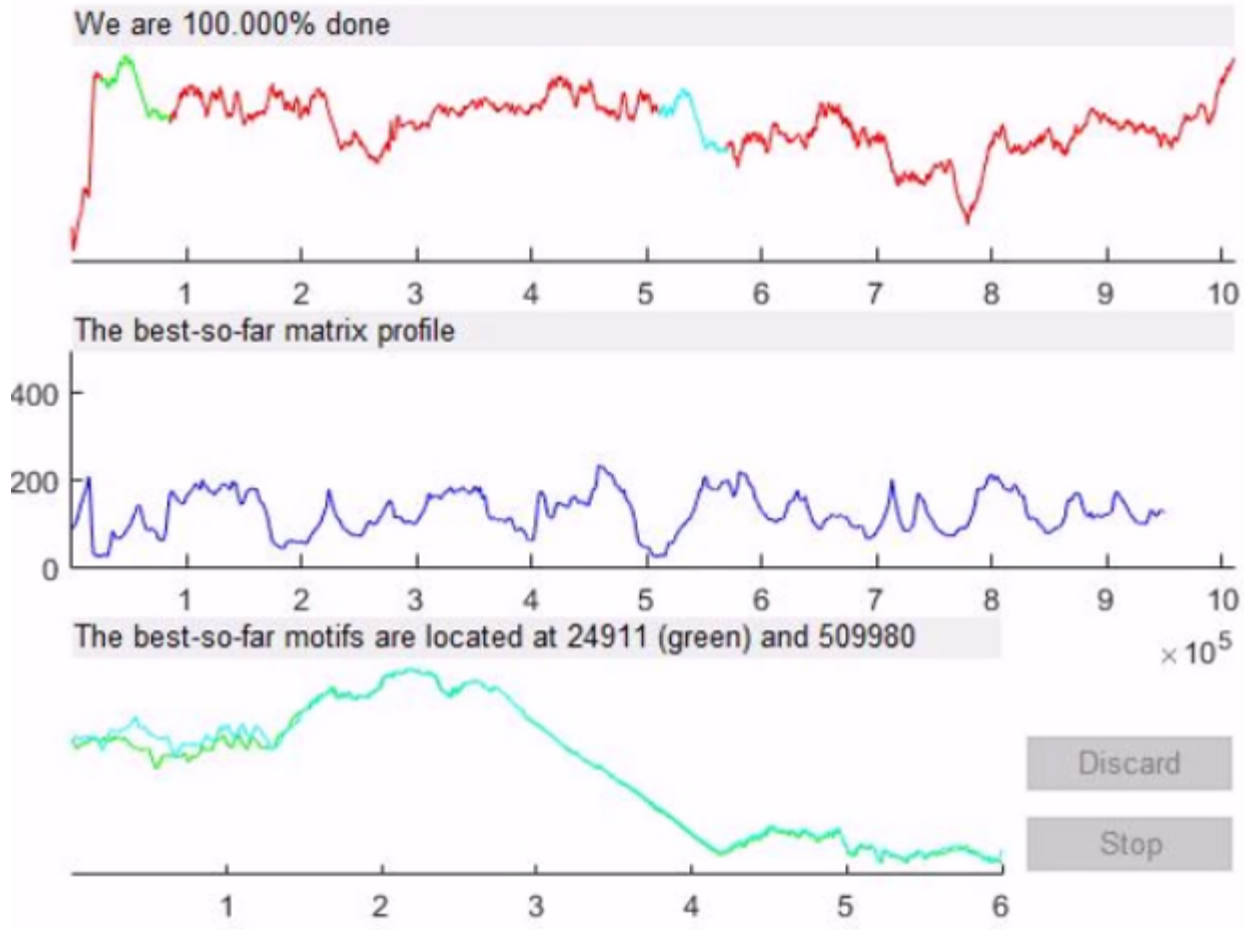

STOMP: Scalable Time series Ordered Matrix Profile

Can we improve STAMP?

STAMP computes distance profiles for rows in random order. Each distance profile is computed independently. However, the successive rows are profiles of two queries that overlap in $m-1$ observations.

We can exploit the overlap between successive queries while computing their distance profiles to build an $O(n^2)$ time, $O(n)$ space algorithm.

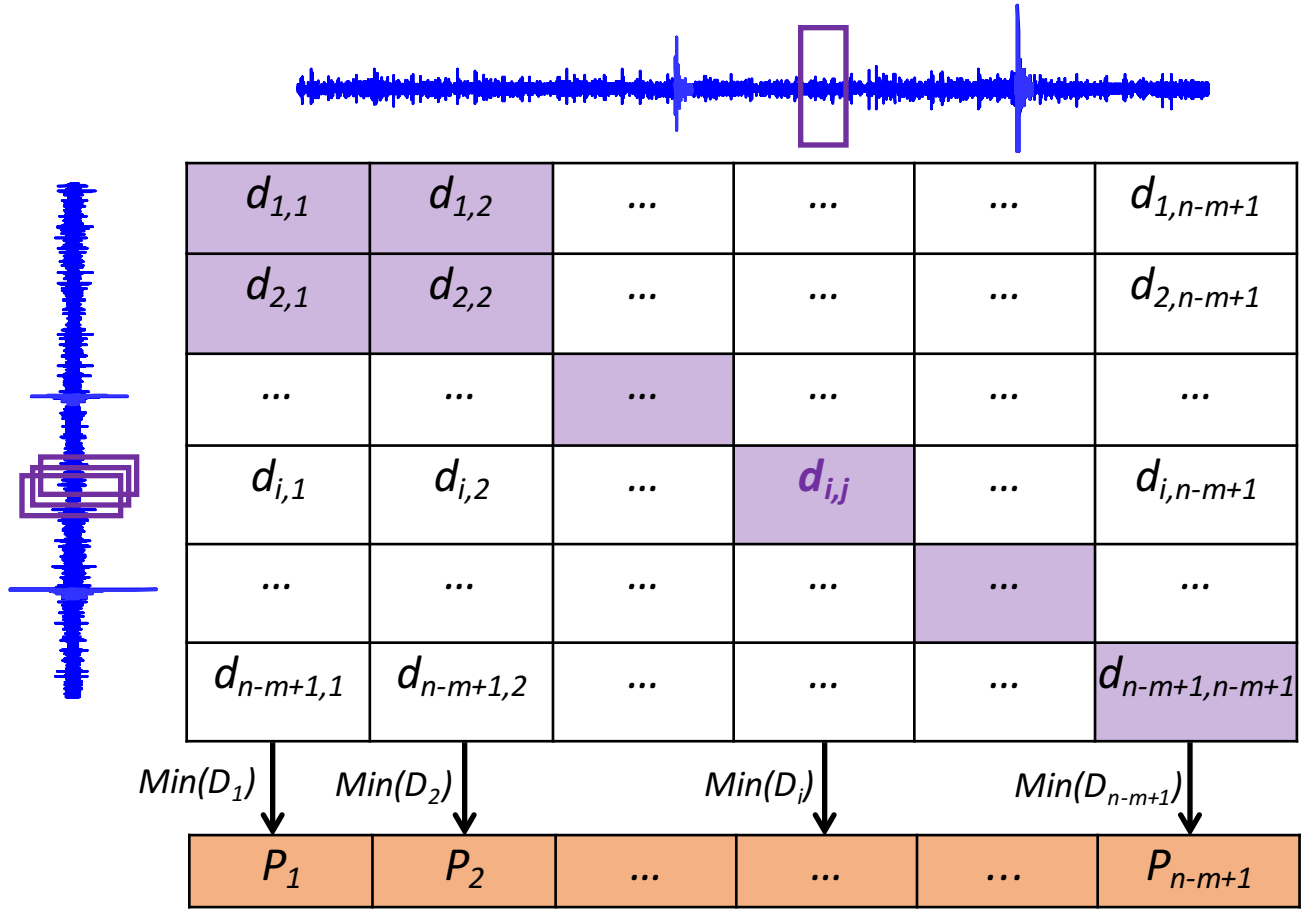

STOMP: Scalable Time series Ordered Matrix Profile

We have an ***O(n²)*** time, ***O(n)*** space algorithm called STOMP to evaluate it.

Recall our working formula:

$$T_i T_j = \sum_{k=0}^{m-1} t_{i+k} t_{j+k}$$

Dot product of the *i*th window and the *j*th window.

$$d_{i,j} = \sqrt{2m\left(1 - \frac{T_i T_j - m\mu_i\mu_j}{m\sigma_i\sigma_j}\right)}$$

- We precompute and store the means and standard deviations in *O(n)* space and time using the `movmean` and `movstd` functions.
- Once we know $T_i T_j$, it takes *O(1)* time to compute $d_{i,j}$

The relationship between T_iT_j and $T_{i+1}T_{j+1}$

$$T_iT_j = \sum_{k=0}^{m-1} t_{i+k}t_{j+k}$$

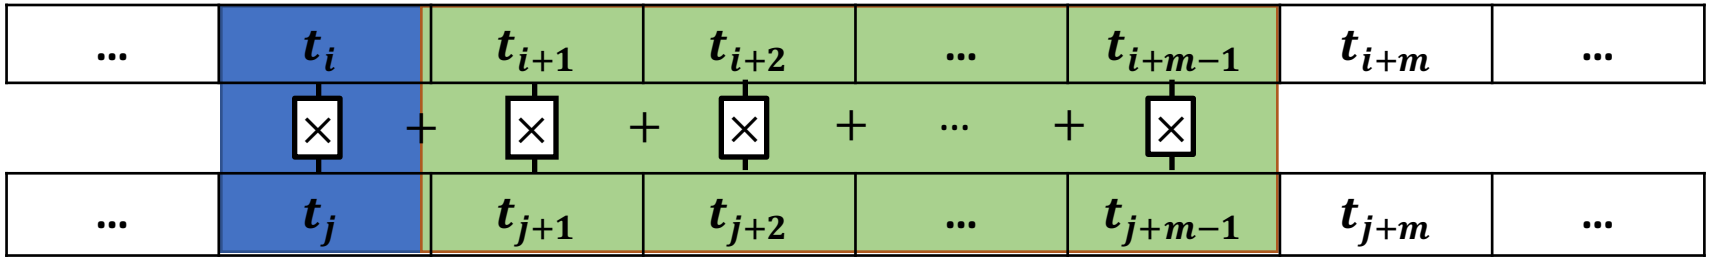

$$T_{i+1}T_{j+1} =$$

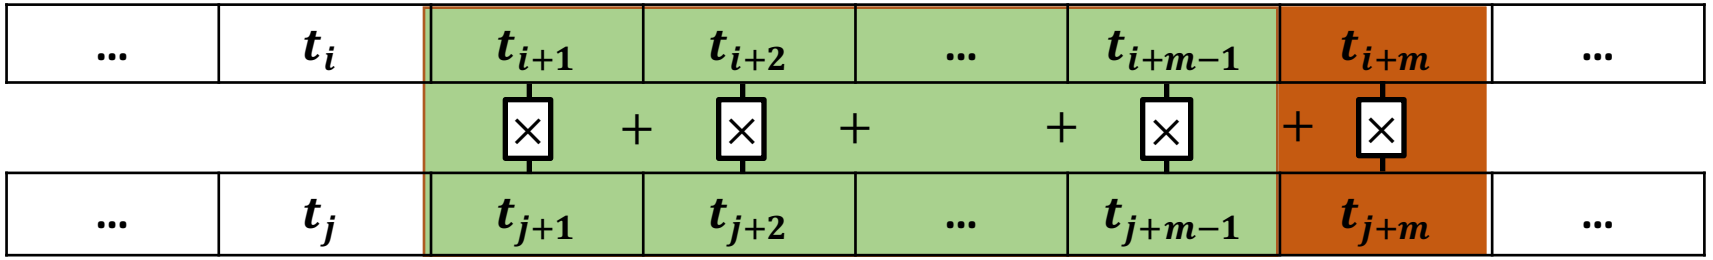

$$T_{i+1}T_{j+1} = T_iT_j - t_it_j + t_{i+m}t_{j+m}$$

$O(1)$ time complexity

STOMP Algorithm: Computing the i^{th} Row

All means and standard deviations are precomputed. This costs linear time and space.

The first column and row of the matrix are identical and pre-computed by MASS.

We iterate through the rows. The previous row is maintained as a local array to feed dot products.

The dot products are converted to distance values and compared against the current best in the profile.

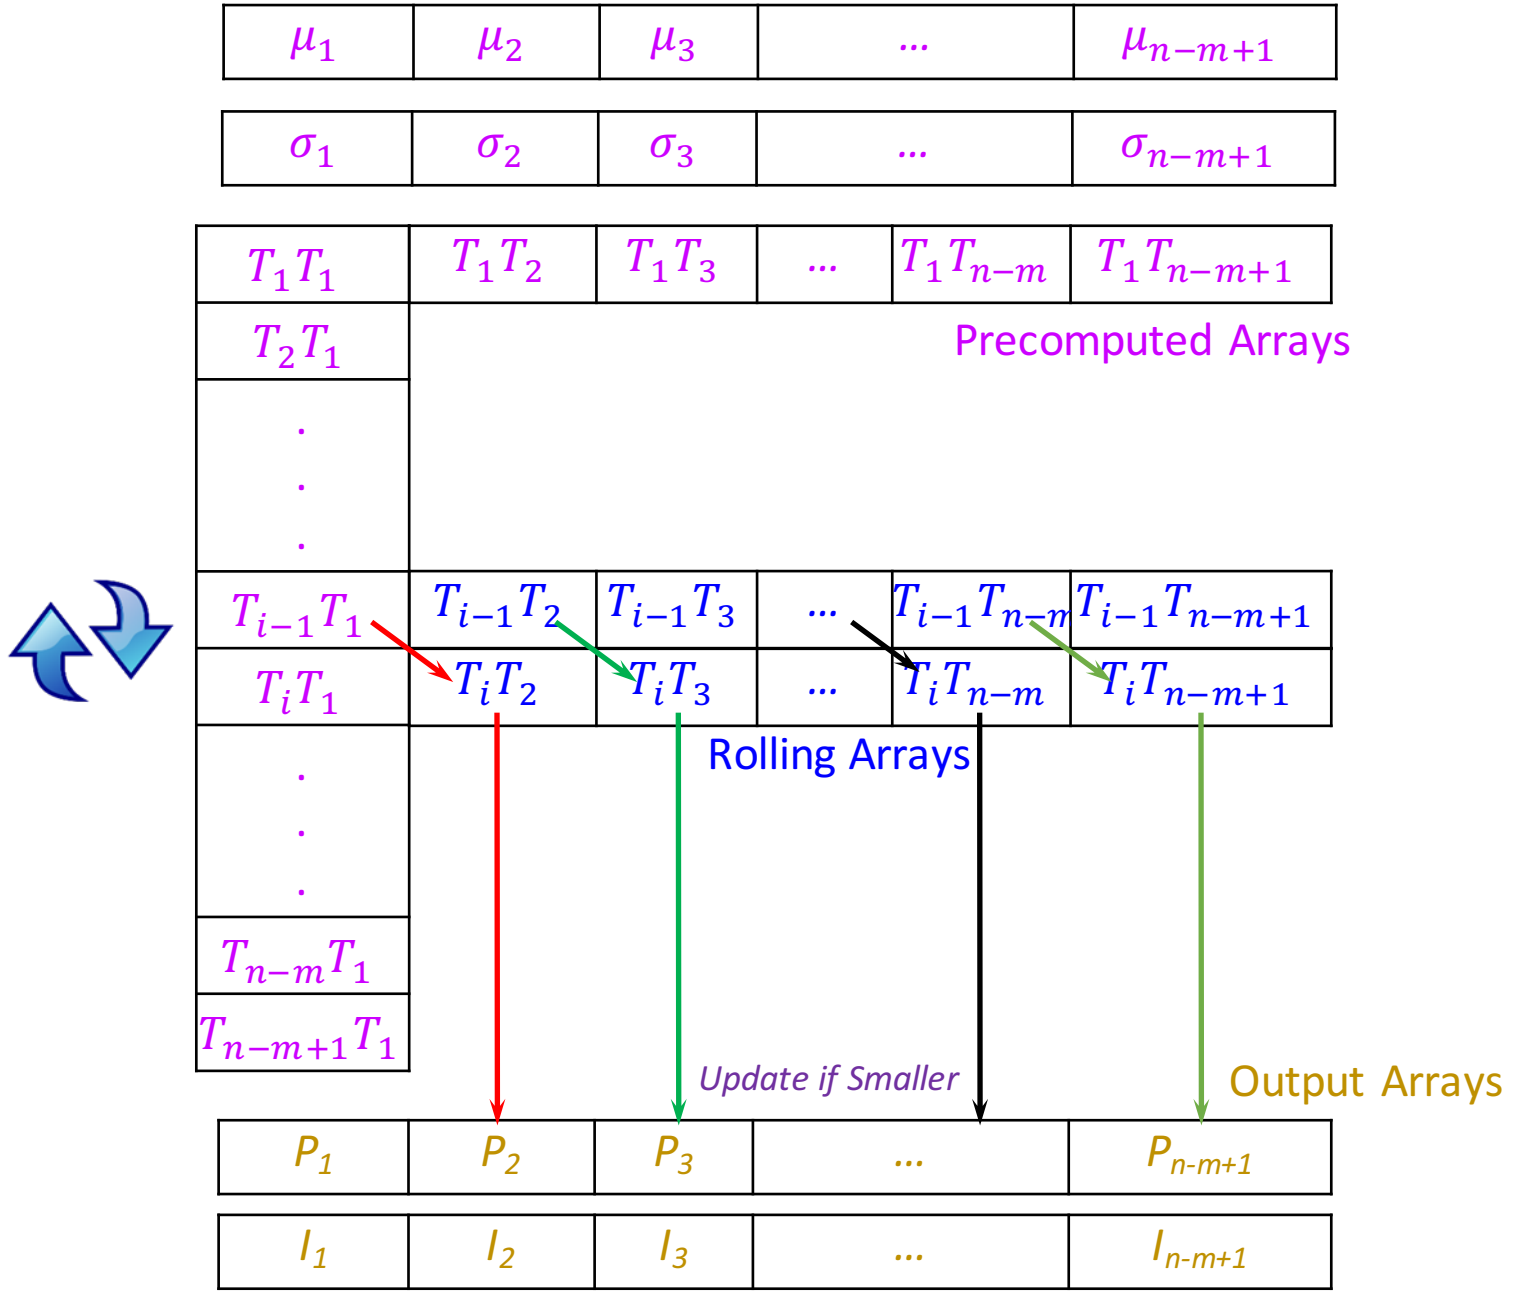

SCRIMP: Scalable Column Independent Matrix Profile

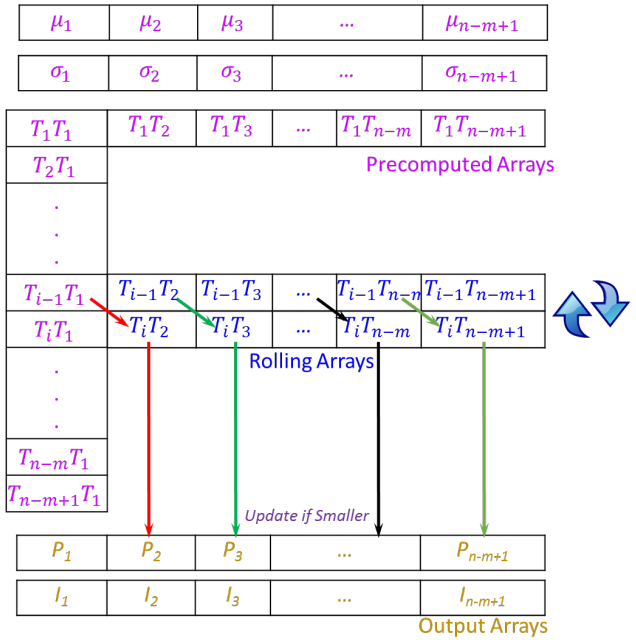

STOMP iterates through rows. Each row depends on the previous row. Therefore random ordering is not suitable.

However, the diagonals are independent of each other. We exploit this property to randomize the computation so we achieve an anytime $O(n^2)$ algorithm.

We avoid computing the upper triangle of the matrix because the distance matrix is symmetric.

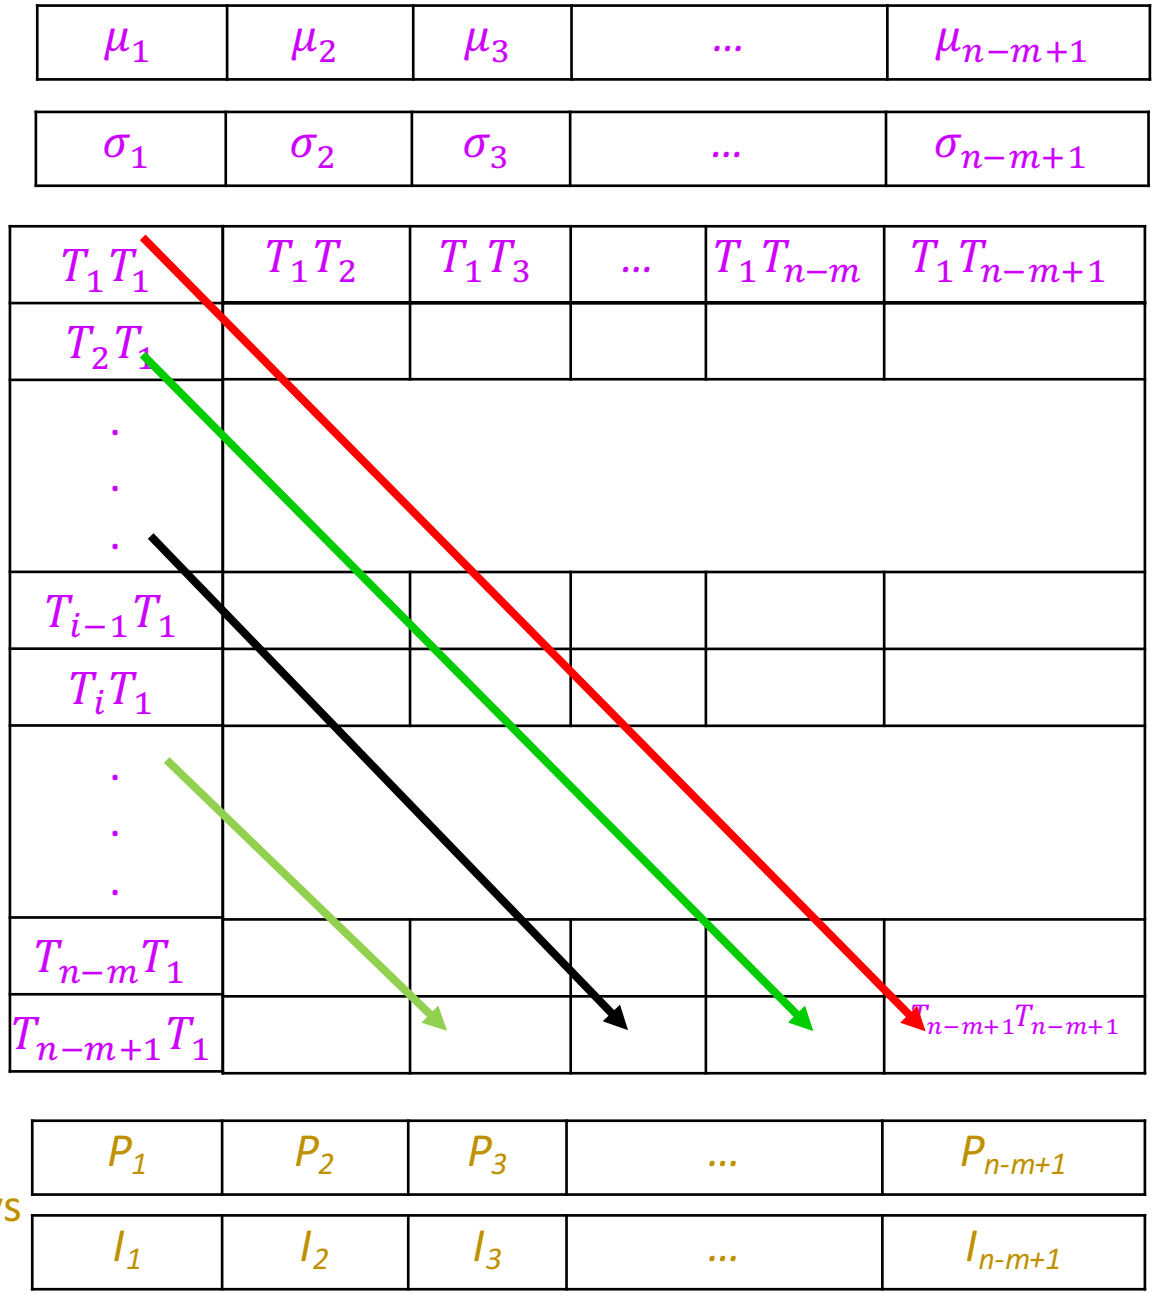

Porting STOMP to GPU

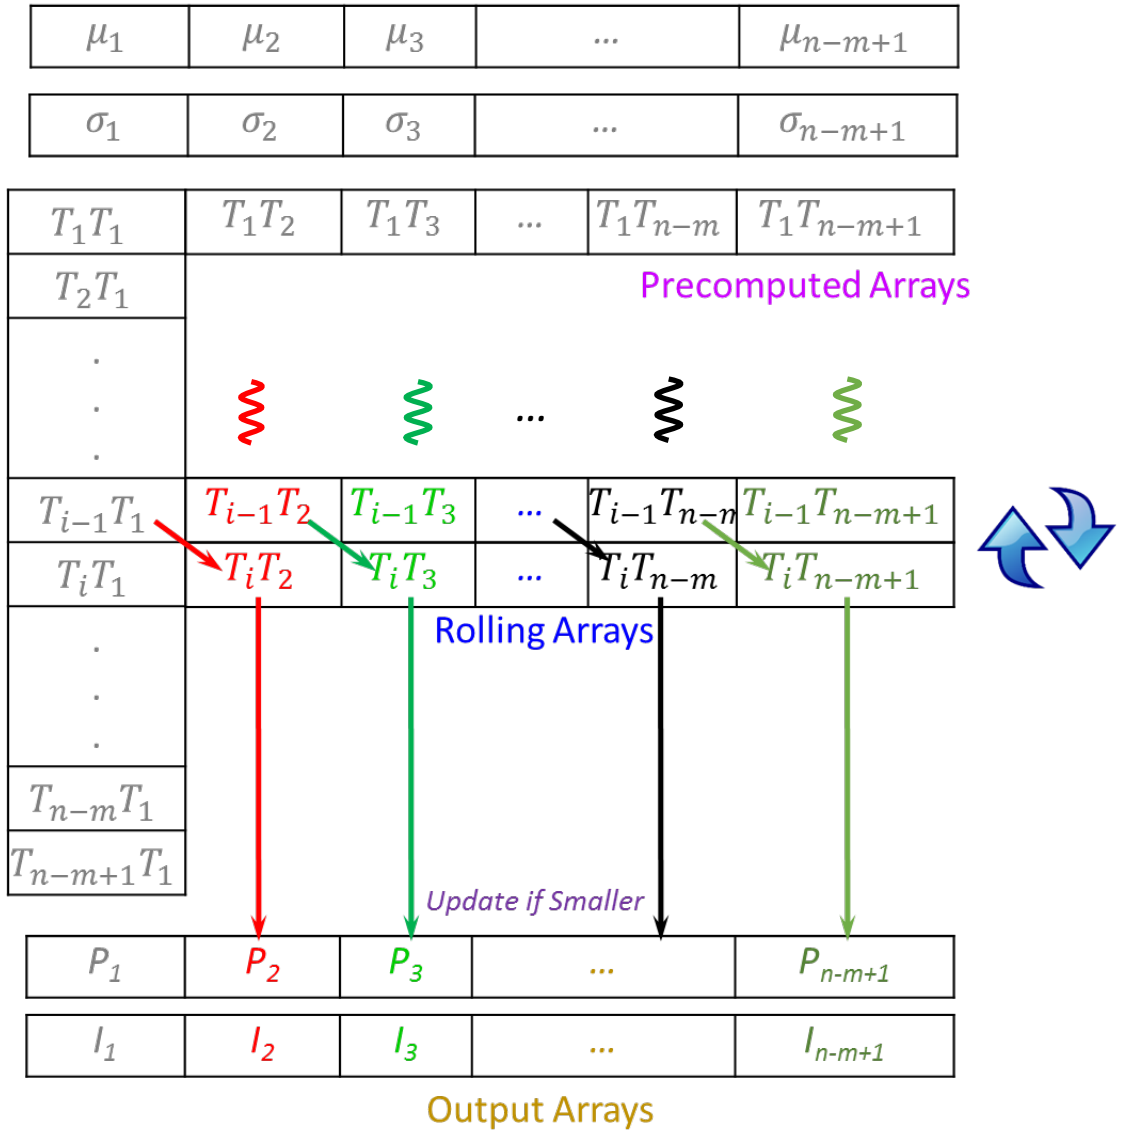

Each thread is assigned to calculate one entry in the matrix profile in every iteration.

All the precomputed arrays are in global shared memory that is accessible to all threads.

Threads are synced after a row is processed to avoid race.

We can further optimize to compute only the lower triangle of the matrix, please see the paper.

Comparison of STAMP, STOMP and GPU-STOMP

For a fix subsequence length m=256: time

Algorithm \ *n*	2^{17}	2^{18}	2^{19}	2^{20}
STAMP	15.1 min	1.17 hours	5.4 hours	24.4 hours
STOMP	4.21 min	0.3 hours	1.26 hours	5.22 hours
GPU-STOMP	10 sec	18 sec	46 sec	2.5 min

For large data, and **for the very first time in the literature, 100,000,000**

Algorithm \ *m* | *n*	**2000 | 17,279,800**	**400 | 100,000,000**
STAMP (*estimated*)	*36.5 weeks*	*25.5 years*
STOMP (*estimated*)	*8.4 weeks*	*5.4 years*
GPU-STOMP	9.27 hours	12.13 days

Comparing the speed of STOMP with existing algorithms

For a time series of length 2^{18}: CPU time(memory usage)

Algorithm \ m	512	1,024	2,048	4,096
STOMP	501s (14MB)	506s (14MB)	490s (14MB)	490s (14MB)
Quick-Motif	27s (65MB)	151s (90MB)	630s (295MB)	695s (101MB)
MK	2040s (1.1GB)	N/A (>2GB)	N/A (>2GB)	N/A (>2GB)

Note: the time and space cost of STOMP is completely independent of any properties (noise, trends, stationarity etc.) of data.
Quick-Motif and MK are pruning based techniques with non-monotonic space need.
STOMP produces more information (i.e. Matrix Profile) while the others find the motifs only.

The Progress Bar Question: How long will it take for STOMP, SCRIMP to finish?

1. Given the inputs, the time is completely *deterministic*.
2. The time is completely independent of, *m* the length of the query. So long as *m* << *n.*
3. The time is completely independent of the structure/noise-level of data itself.

As it happens, these properties are *very* rare in data mining algorithms.

Can we use these properties to estimate time-to-finish?

The Progress Bar Question: How long will it take for STOMP, STAMP or SCRIMP to finish?

The time only depends on the length of the time series, *n*, and the hardware settings. To compute the time needed to finish, you just need to do *one* calibration run on any particular hardware configuration.

$$TimeNeeded = \frac{Time_{calibration}}{n_{calibration}^{\ 2}} * n_{new}^{\ 2}$$

How well does this work?

We *measured* the time needed for T = 512, 1024, 2048,… ,131072 (we used an old cheap laptop)

We then used the time measured for the 131,072 run, to predict the time needed for all the other runs. We plotted the two curves below. Note that the last point agrees by definition.

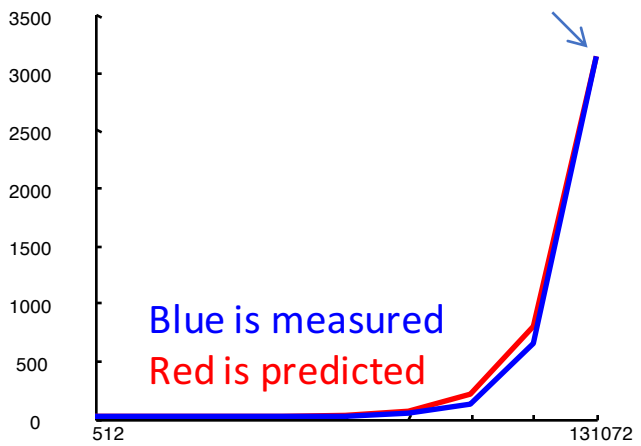

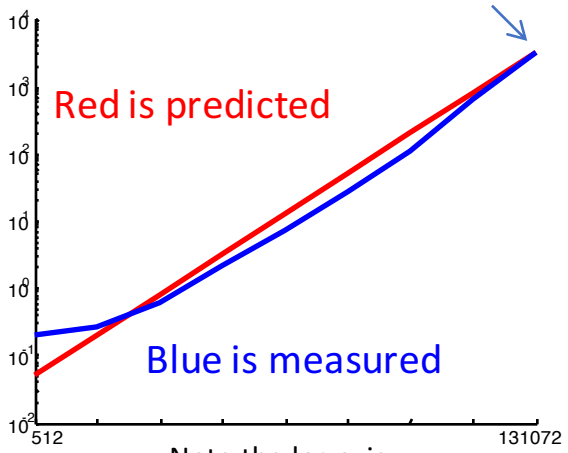

Note the log axis

Outline

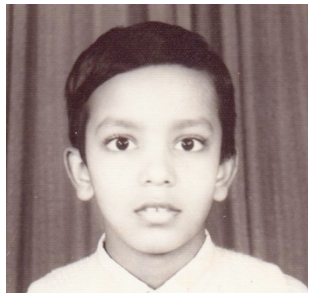

Act 1

- Our Fundamental Assumption
- What is the (MP) Matrix Profile?
- Properties of the MP
- Developing a Visual Intuition for MP
- Basic Algorithms
  - MP Motif Discovery
  - MP Time Series Chains
  - MP Anomaly Discovery
  - MP Joins (self and AB)
  - MP Semantic Segmentation
- From Domain Agnostic to Domain Aware: The Annotation Vector (A simple way to use domain knowledge to adjust your results)
- The "*Matrix Profile and ten lines of code is all you need*" philosophy.
- Break

Act 2

- Background on time series mining
  - Similarity Measures
  - Normalization
- Distance Profile
  - Brute Force Approach
  - Just-in-time Normalization
  - The MASS Algorithm
  - Extensions of MASS
- Matrix Profile
  - STAMP
  - STOMP
  - GPU-STOMP
  - SCRIMP
- Open problems to solve

Open Problems

- Find Lower Bounds
  - An $O(n)$ lower bound to distance profile
  - An $O(n \log n)$ lower bound to matrix profile
- Adapt to L1 distance
  - $O(n \log n)$ algorithm for distance profile
  - $O(n^2)$ algorithm for Matrix Profile
- Adapt to DTW distance
  - Can we create a Matrix Profile of 100M time series under warping?
- Variable Length Matrix Profile
  - Can we rank or sort neighbors of different lengths based on degree of fidelity?
- Scalability
  - Distributed Matrix Profile algorithm for horizontal scalability
  - Can we exploit hardware acceleration techniques for scalability?

Open Problems

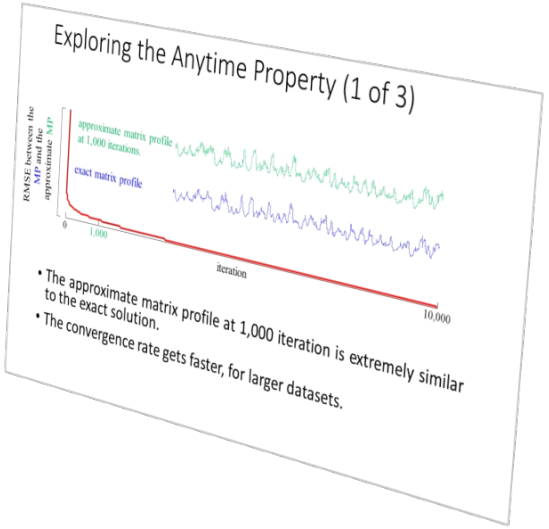

- Domain Appropriate Matrix Profile
  - Recall that both STAMP and SCRIMP converge quickly with *random* ordering. However, could a *data-adaptive* ordering converge even faster?
  - We discussed the Annotation Vector (AV). While we can typically specify an AV with a few lines of code, can we *learn* a domain appropriate AV from user interaction?
- Visualization
  - At some scale, user-interfaces / user-experience become very important, we have largely glossed over this. Can we develop interactive visualization techniques for active exploration over MP?

Questions?

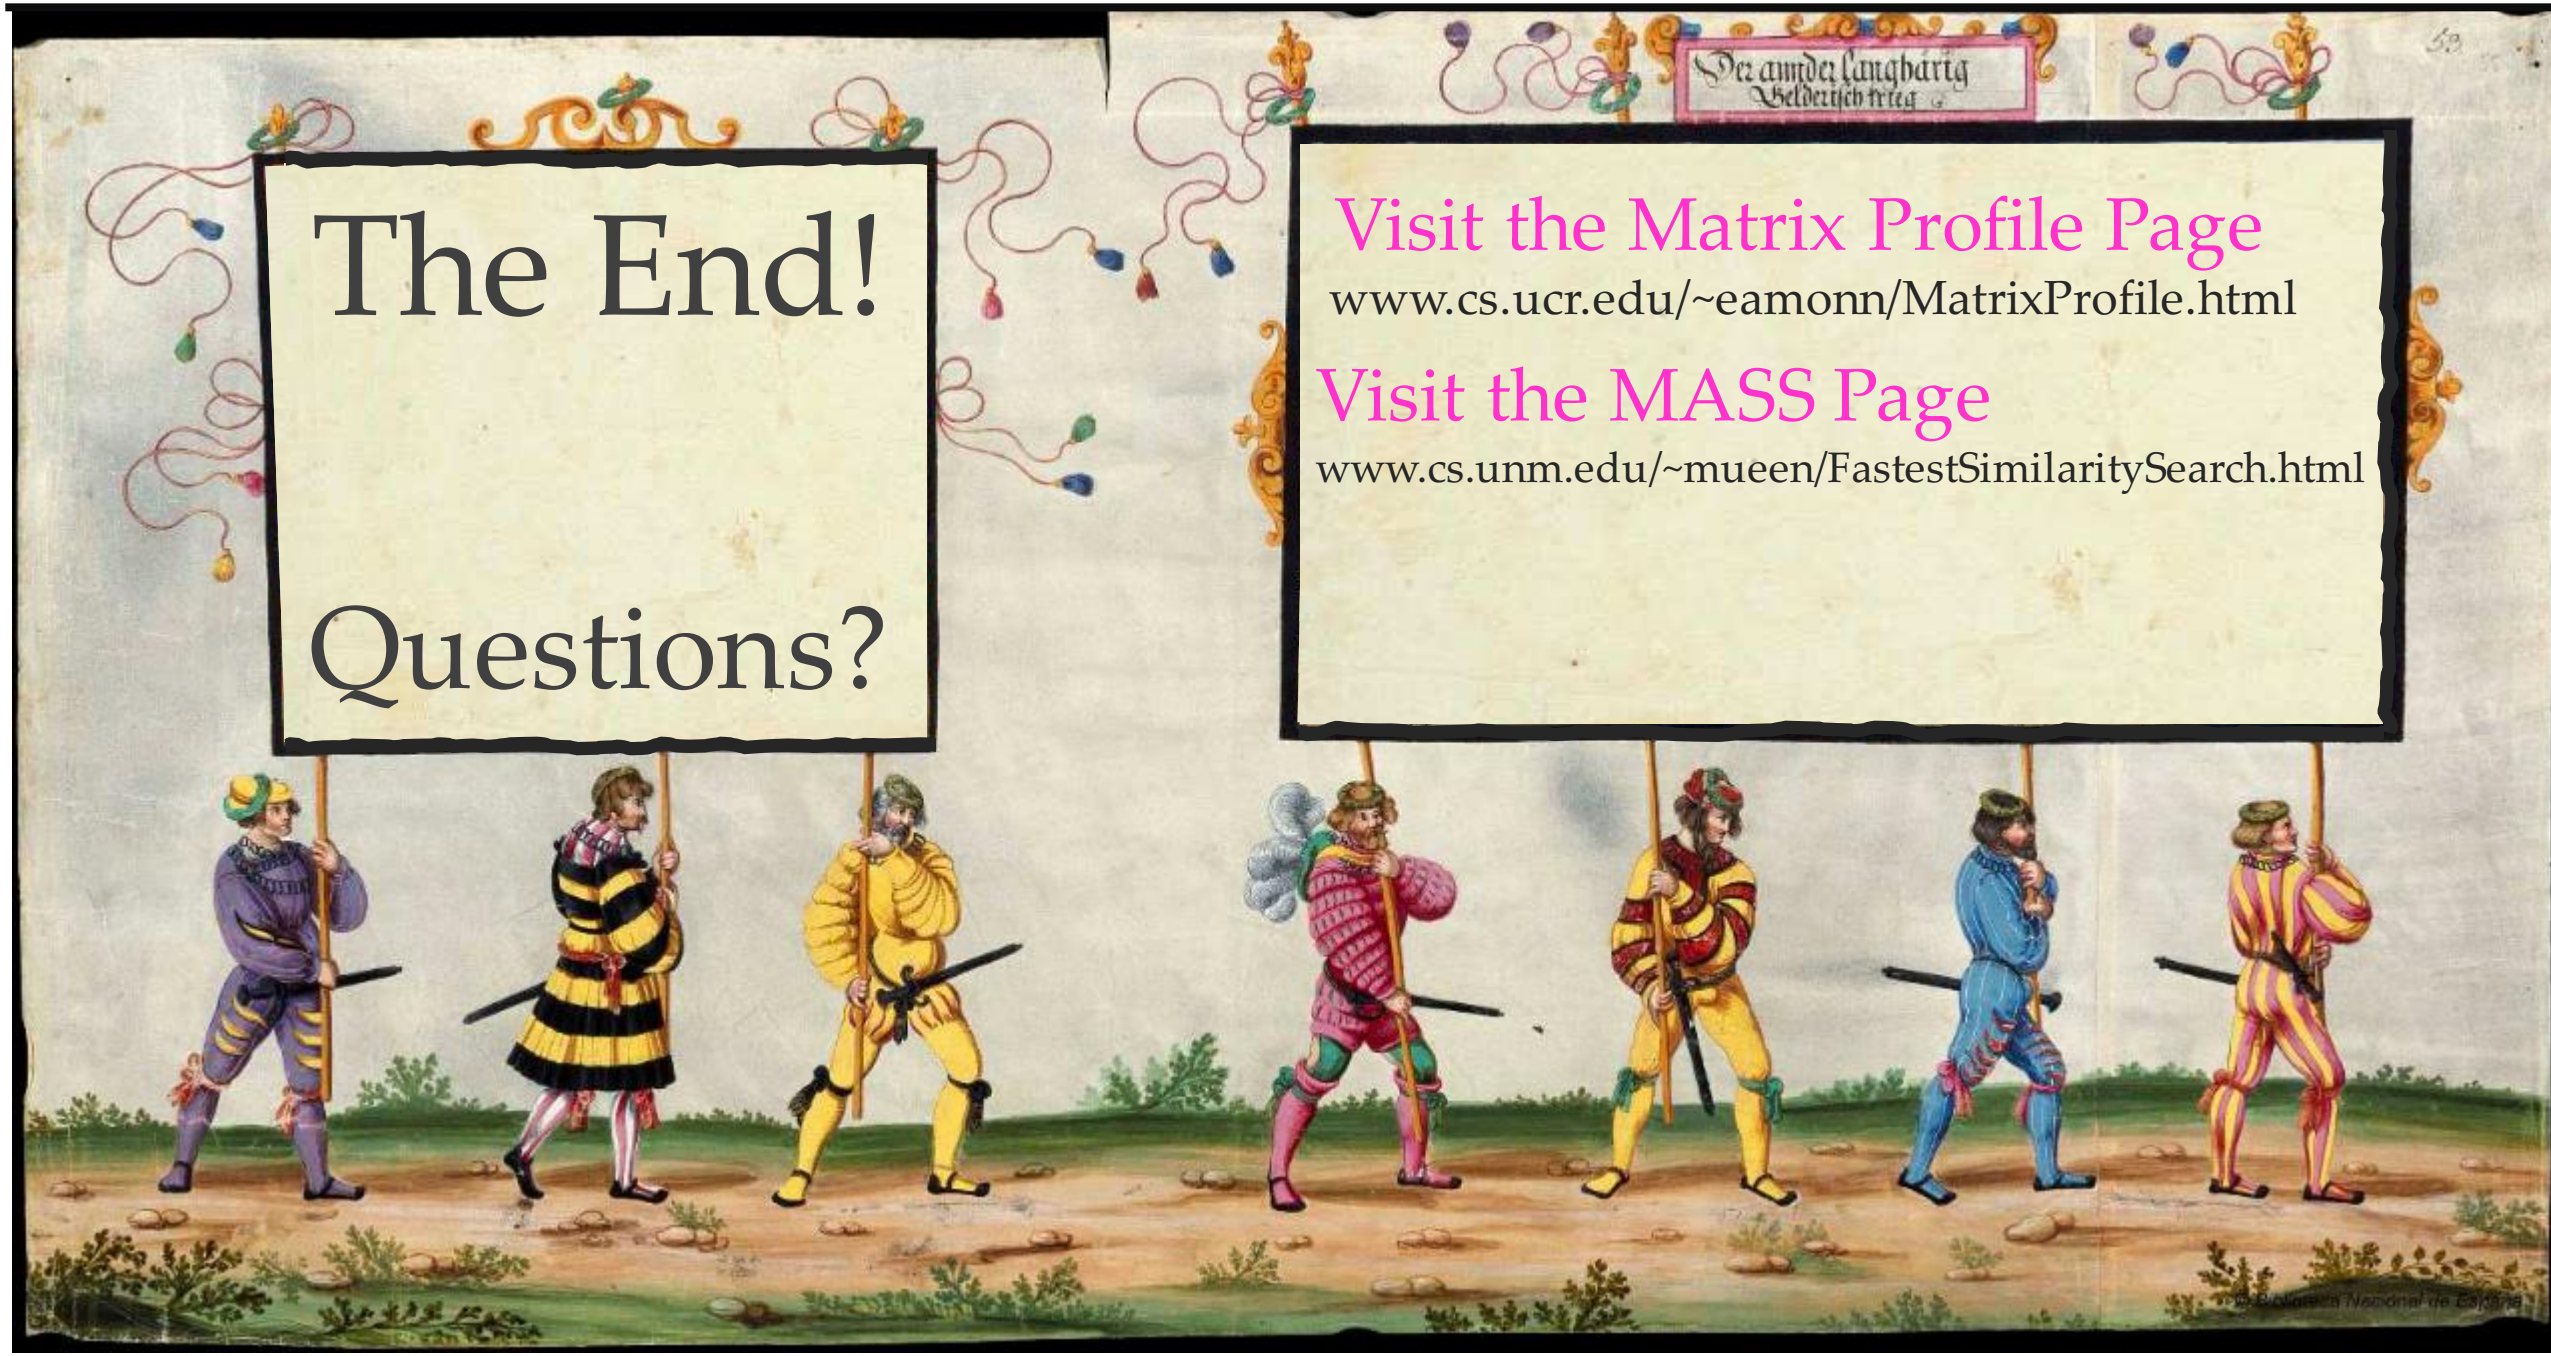
The End!
Questions?
Visit the Matrix Profile Page
www.cs.ucr.edu/~eamonn/MatrixProfile.html
Visit the MASS Page
www.cs.unm.edu/~mueen/FastestSimilaritySearch.html