杭州师范大学 Hangzhou Normal University 数学学院

陈建功大讲堂

Topological time series analysis: with applications to biomedical and speech signal processing

报告人：朱一飞（南方科技大学）

报告时间：2022年11月10日（周四）14:30

报告地点：勤园21号楼306学术报告厅

报告摘要

We give an overview of topological approaches to analyzing time–dependent data, with an emphasis on detection of periodic phenomena. This methodology enjoys robustness afforded by continuous deformation and change of measures, captures interesting geometric features underlying the data, and requires a reasonable computational cost. We illustrate these by reporting progress on two specific applications: (i) automated and real–time detection of mouse scratching behavior, joint with Fangyi Chen and Zhen Zhang, and (ii) classification of voiced and unvoiced speech signals, joint with Meng Yu.

报告人简介

Yifei Zhu is an Assistant Professor in mathematics of Southern University of Science and Technology. His current research focuses on interactions of algebraic topology with algebraic geometry and number theory, especially moduli spaces from spectral algebraic geometry in the context of the Langlands program, as well as applications of geometry and topology to interdisciplinary research, including condensed–matter physics and material science.

主办方 杭州师范大学数学学院

Topological time series analysis: with applications to biomedical and speech signal processing

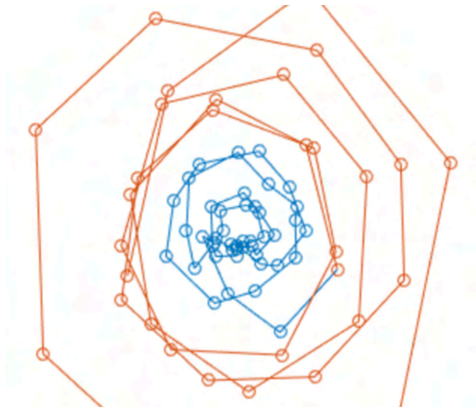

朱一飞 (SUSTech)

杭州师范大学数学学院陈建功大讲堂, 2022.11

Periodic phenomena: a motivating example

Let $\mathbb{T}^2 = (\mathbb{R}/\mathbb{Z})^2$ be the 2D torus. Consider the dynamical system given by

$$\Phi_\sigma : \mathbb{T}^2 \times \mathbb{R} \to \mathbb{T}^2$$
$$\big((a,b),t\big) \mapsto (a+t, b+\sigma t)$$

Periodic phenomena: a motivating example

Let $\mathbb{T}^2 = (\mathbb{R}/\mathbb{Z})^2$ be the 2D torus. Consider the dynamical system given by

$$\Phi_\sigma : \mathbb{T}^2 \times \mathbb{R} \to \mathbb{T}^2$$
$$\big((a,b),t\big) \mapsto (a+t, b+\sigma t)$$

If σ is rational, then every orbit is periodic.

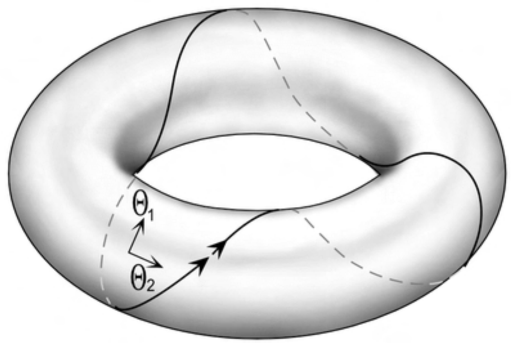

(a)

Periodic phenomena: a motivating example

Let $\mathbb{T}^2 = (\mathbb{R}/\mathbb{Z})^2$ be the 2D torus. Consider the dynamical system given by

$$\begin{aligned}\Phi_\sigma : \mathbb{T}^2 \times \mathbb{R} &\to \mathbb{T}^2 \\ \big((a,b),t\big) &\mapsto (a+t, b+\sigma t)\end{aligned}$$

If σ is rational, then every orbit is periodic. Otherwise every orbit is dense in $\mathbb{T}^2$.

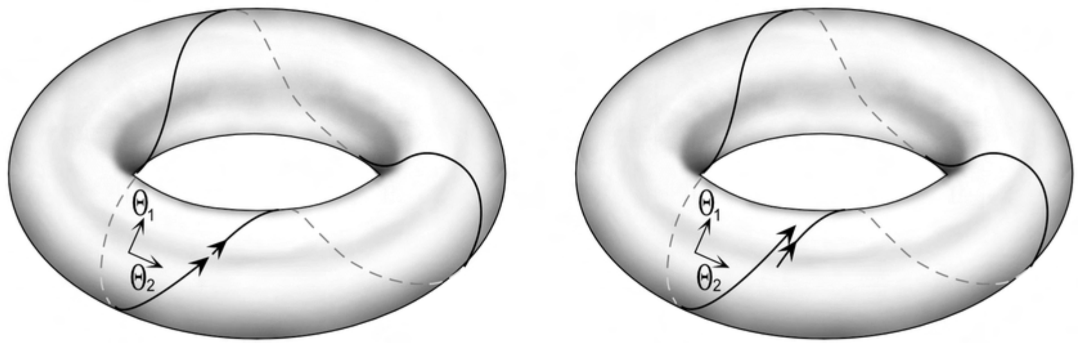

(a) (b)

Periodic phenomena: a motivating example

Let $\mathbb{T}^2 = (\mathbb{R}/\mathbb{Z})^2$ be the 2D torus. Consider the dynamical system given by

$$\Phi_\sigma : \mathbb{T}^2 \times \mathbb{R} \to \mathbb{T}^2$$
$$\big((a, b), t\big) \mapsto (a + t, b + \sigma t)$$

If σ is rational, then every orbit is periodic. Otherwise every orbit is dense in $\mathbb{T}^2$.

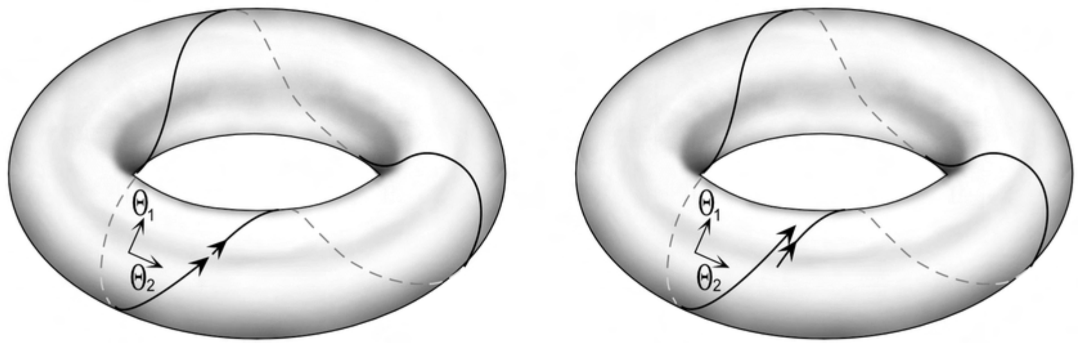

(a) (b)

From time series to topological shapes

Most periodic time series can be realized by a topological circle S^1 embedded in a Euclidean space of higher dimension.

Ideas of topological data analysis (TDA)

The topological type (more precisely, homotopy type) is robust against perturbations.

$y = \sin x$

Ideas of topological data analysis (TDA)

The topological type (more precisely, homotopy type) is robust against perturbations.

$y = \sin x$

$y = \sin x + \epsilon(x)$

Ideas of topological data analysis (TDA)

Features of topological shapes, such as the number of holes, can be captured by algebraic invariants that are computable.

Ideas of topological data analysis (TDA)

Features of topological shapes, such as the number of holes, can be captured by algebraic invariants that are computable.

Comparing these invariants effectively distinguishes the topological types of shapes.

Ideas of topological data analysis (TDA)

Features of topological shapes, such as the number of holes, can be captured by algebraic invariants that are computable.

Comparing these invariants effectively distinguishes the topological types of shapes.

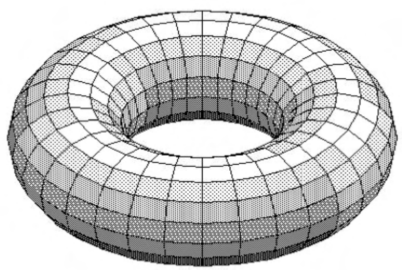

$$H_k(\mathbb{T}^2) = \begin{cases} \mathbb{Z} & k = 0 \\ & \\ & \end{cases}$$

Ideas of topological data analysis (TDA)

Features of topological shapes, such as the number of holes, can be captured by algebraic invariants that are computable.

Comparing these invariants effectively distinguishes the topological types of shapes.

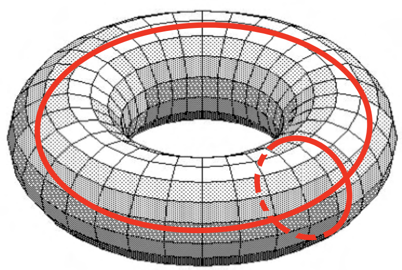

$$H_k(\mathbb{T}^2) = \begin{cases} \mathbb{Z} & k = 0 \\ \mathbb{Z} \oplus \mathbb{Z} & k = 1 \\ & \\ & \end{cases}$$

Ideas of topological data analysis (TDA)

Features of topological shapes, such as the number of holes, can be captured by algebraic invariants that are computable.

Comparing these invariants effectively distinguishes the topological types of shapes.

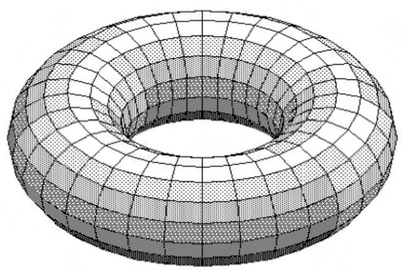

$$H_k(\mathbb{T}^2) = \begin{cases} \mathbb{Z} & k = 0 \\ \mathbb{Z} \oplus \mathbb{Z} & k = 1 \\ \mathbb{Z} & k = 2 \end{cases}$$

Ideas of topological data analysis (TDA)

Features of topological shapes, such as the number of holes, can be captured by algebraic invariants that are computable.

Comparing these invariants effectively distinguishes the topological types of shapes.

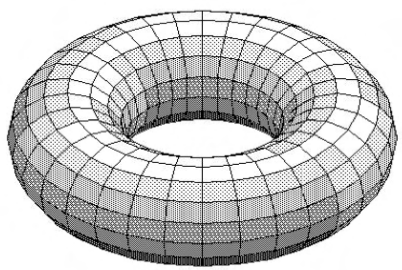

$$H_k(\mathbb{T}^2) = \begin{cases} \mathbb{Z} & k = 0 \\ \mathbb{Z} \oplus \mathbb{Z} & k = 1 \\ \mathbb{Z} & k = 2 \\ 0 & k > 2 \end{cases}$$

Ideas of topological data analysis (TDA)

Features of topological shapes, such as the number of holes, can be captured by algebraic invariants that are computable.

Comparing these invariants effectively distinguishes the topological types of shapes.

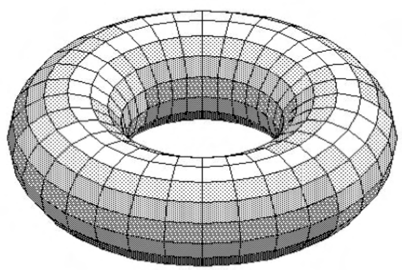

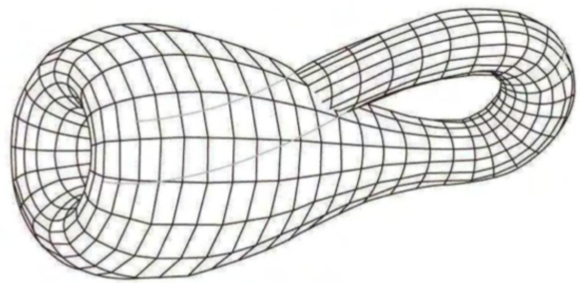

$$H_k(\mathbb{T}^2) = \begin{cases} \mathbb{Z} & k = 0 \\ \mathbb{Z} \oplus \mathbb{Z} & k = 1 \\ \mathbb{Z} & k = 2 \\ 0 & k > 2 \end{cases}$$

$$H_k(\text{Klein bottle}) = \begin{cases} \mathbb{Z} & k = 0 \\ \mathbb{Z} \oplus \mathbb{Z}/2 & k = 1 \\ 0 & k = 2 \\ 0 & k > 2 \end{cases}$$

Topological time series analysis

Let us make the assumption that sampled signals are distributed over a manifold.

Topological time series analysis

Let us make the assumption that sampled signals are distributed over a manifold. To topologically analyze time series, we then proceed as follows:

Step1 Embed the data into a Euclidean space of suitable dimension;

Topological time series analysis

Let us make the assumption that sampled signals are distributed over a manifold. To topologically analyze time series, we then proceed as follows:

<u>Step1</u> Embed the data into a Euclidean space of suitable dimension;

<u>Step 2</u> Compute the algebraic invariants for statistical inference.

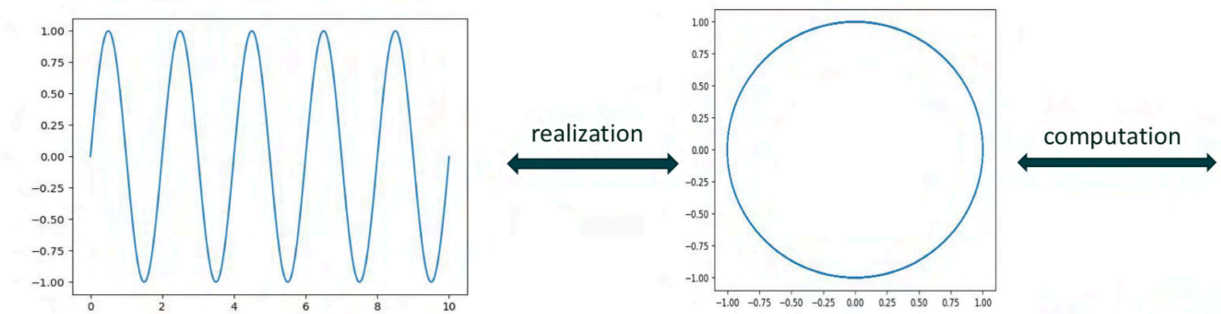

$$H_k(S^1) = \begin{cases} \mathbb{Z} & k = 0 \\ \mathbb{Z} & k = 1 \\ 0 & k > 1 \end{cases}$$

Topological time series analysis

Let us make the assumption that sampled signals are distributed over a manifold. To topologically analyze time series, we then proceed as follows:

<u>Step1</u> Embed the data into a Euclidean space of suitable dimension;

<u>Step 2</u> Compute the algebraic invariants for statistical inference.

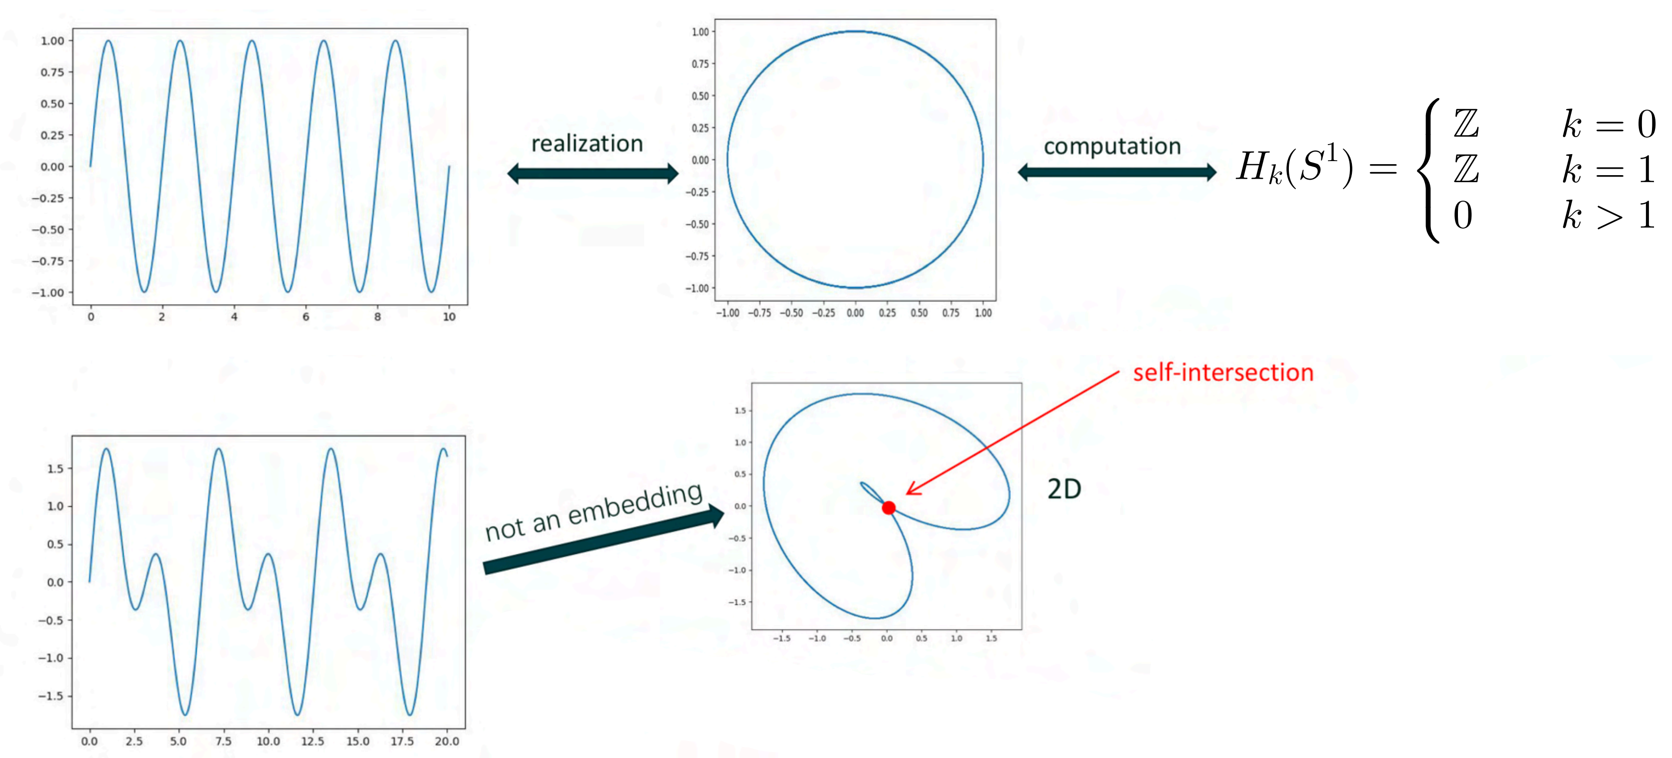

Topological time series analysis

Let us make the assumption that sampled signals are distributed over a manifold. To topologically analyze time series, we then proceed as follows:

<u>Step1</u> Embed the data into a Euclidean space of suitable dimension;

<u>Step 2</u> Compute the algebraic invariants for statistical inference.

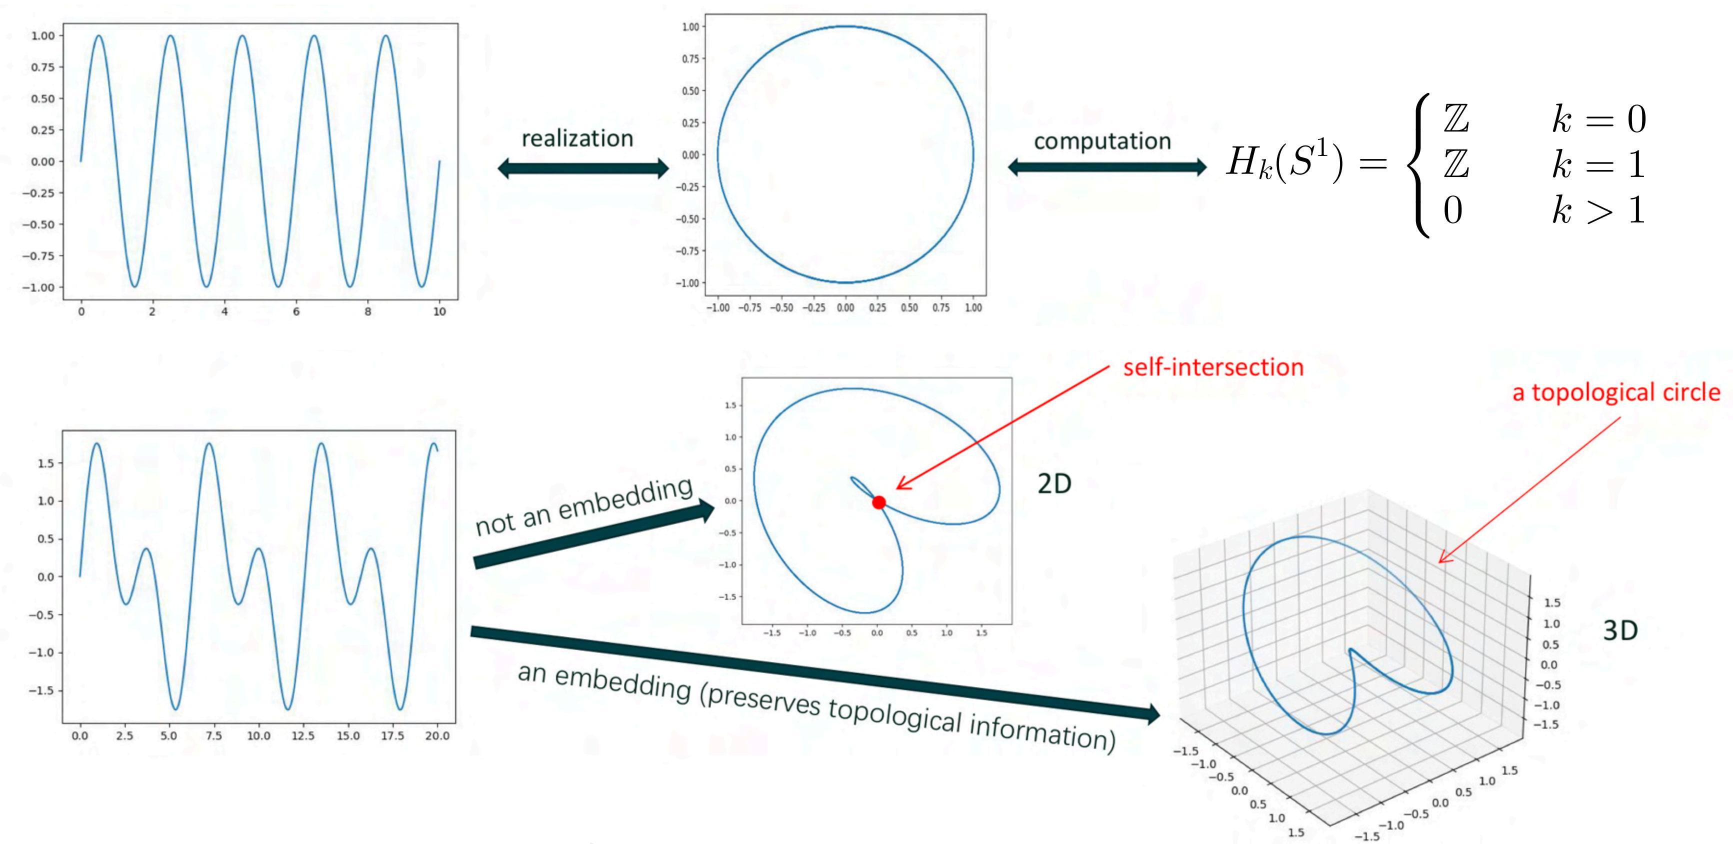

An application: detection of wheeze in medical science (pulmonology)

Wheezes are abnormal lung sounds and usually imply obstructive airway diseases.

An application: detection of wheeze in medical science (pulmonology)

Wheezes are abnormal lung sounds and usually imply obstructive airway diseases.

The most important characteristic of wheeze signals is their periodic patterns.

An application: detection of wheeze in medical science (pulmonology)

Wheezes are abnormal lung sounds and usually imply obstructive airway diseases.

The most important characteristic of wheeze signals is their periodic patterns.

The accuracy of topological periodicity detection is 98.39% (Emrani et al., IEEE Signal Processing Letters, 2014), while in two earlier papers with different methods they are 86.2% and 95.5%.

An application: detection of wheeze in medical science (pulmonology)

Wheezes are abnormal lung sounds and usually imply obstructive airway diseases.

The most important characteristic of wheeze signals is their periodic patterns.

The accuracy of topological periodicity detection is 98.39% (Emrani et al., IEEE Signal Processing Letters, 2014), while in two earlier papers with different methods they are 86.2% and 95.5%.

As a warm-up, our research group has reproduced their results using the original data and open-source TDA programming package.

An application: detection of wheeze in medical science (pulmonology)

wheeze

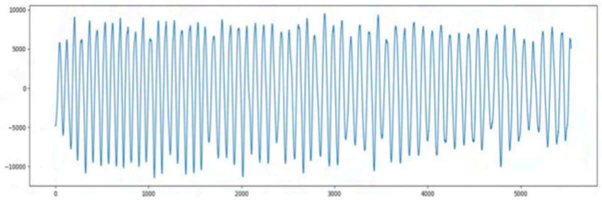

normal

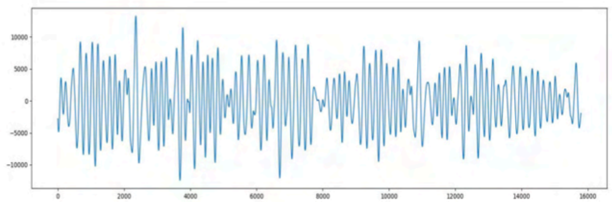

Original sound signals

An application: detection of wheeze in medical science (pulmonology)

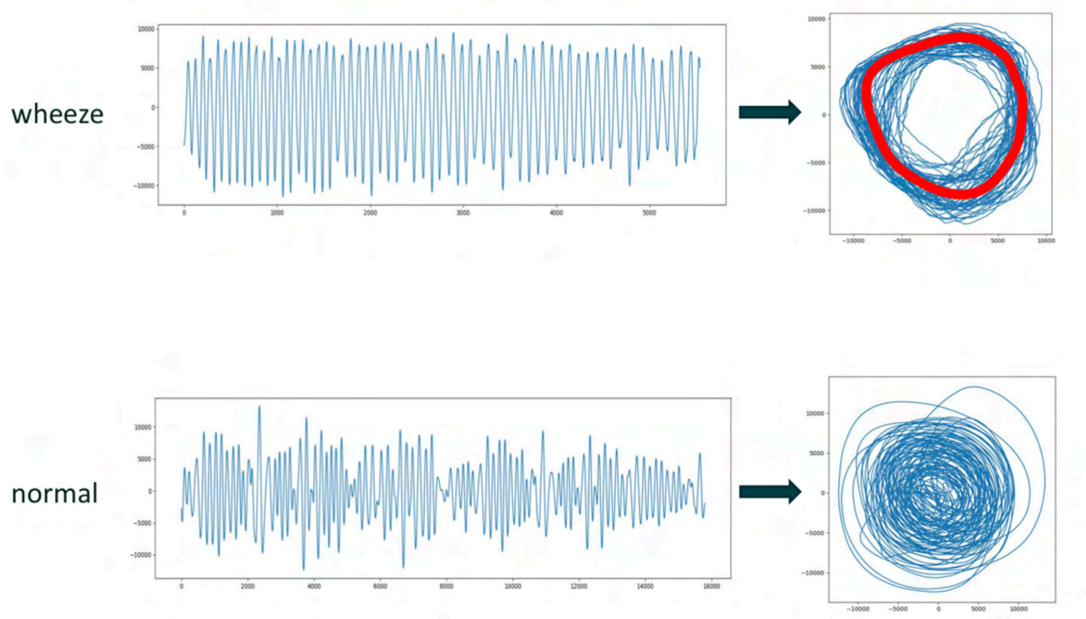

Original sound signals

Realized topological shapes embedded in 2D Euclidean space

An application: detection of wheeze in medical science (pulmonology)

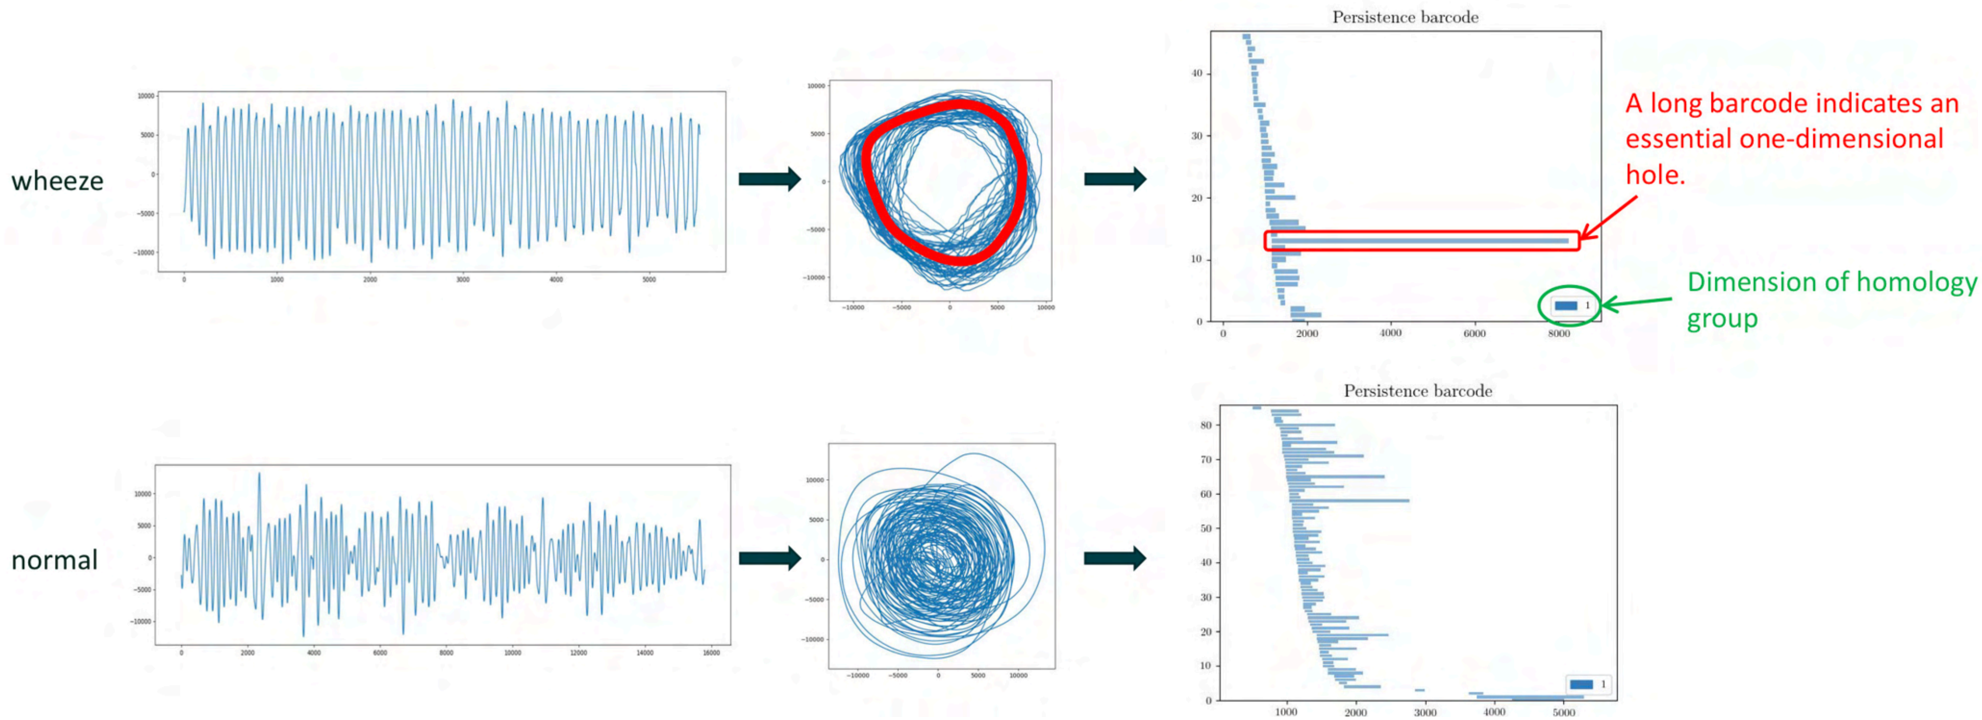

Original sound signals

Realized topological shapes embedded in 2D Euclidean space

“Persistence barcodes” as representations of the algebraic invariant (1D homology group)

A pipeline for topological time series analysis

Time series data: $x_1, x_2, x_3, x_4, \ldots$

A pipeline for topological time series analysis

Time series data: $x_1, x_2, x_3, x_4, \ldots$

Preprocessing → Sliding window embedding

2D Euclidean shape: point cloud $(x_1, x_{1+\tau}), (x_2, x_{2+\tau}), (x_3, x_{3+\tau}), \ldots$

delay

A pipeline for topological time series analysis

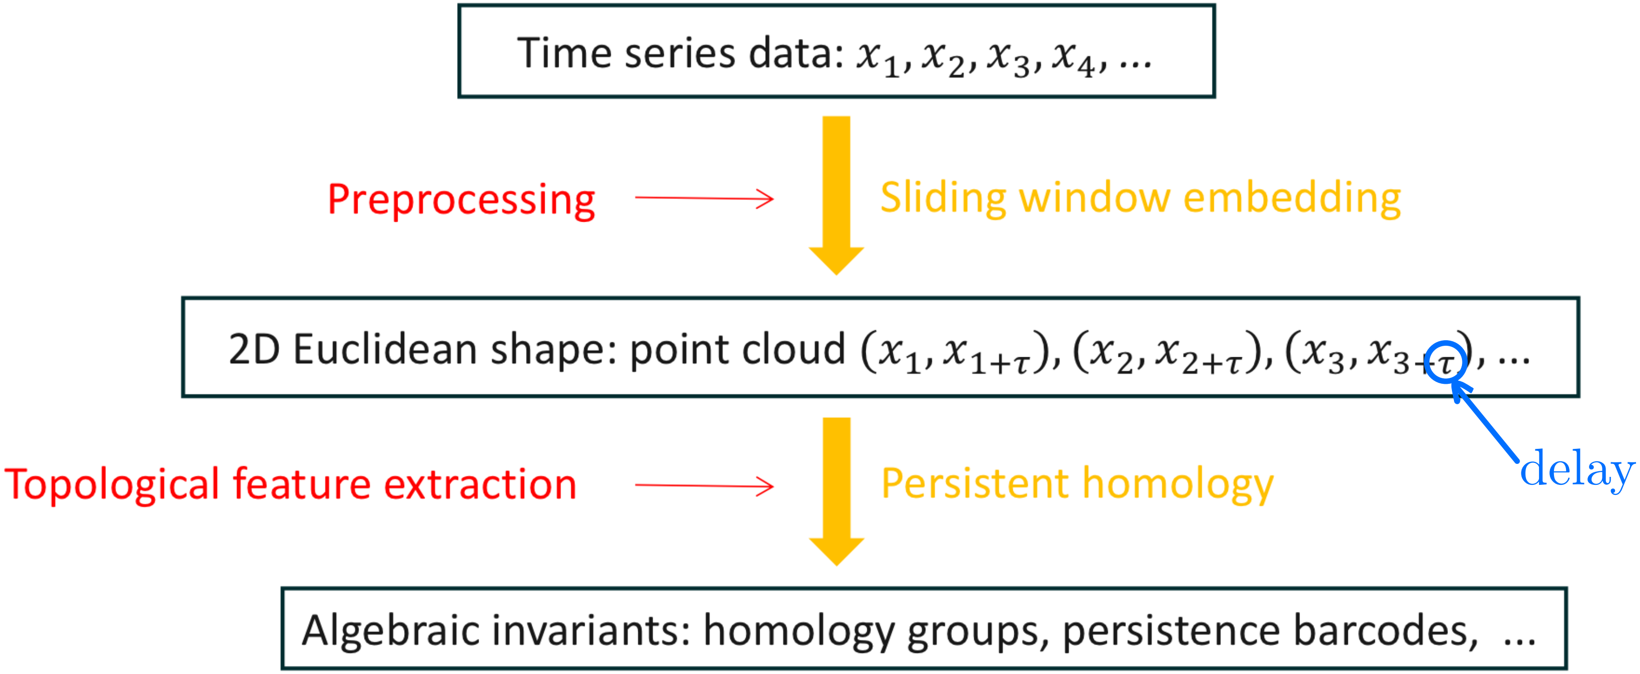

A pipeline for topological time series analysis

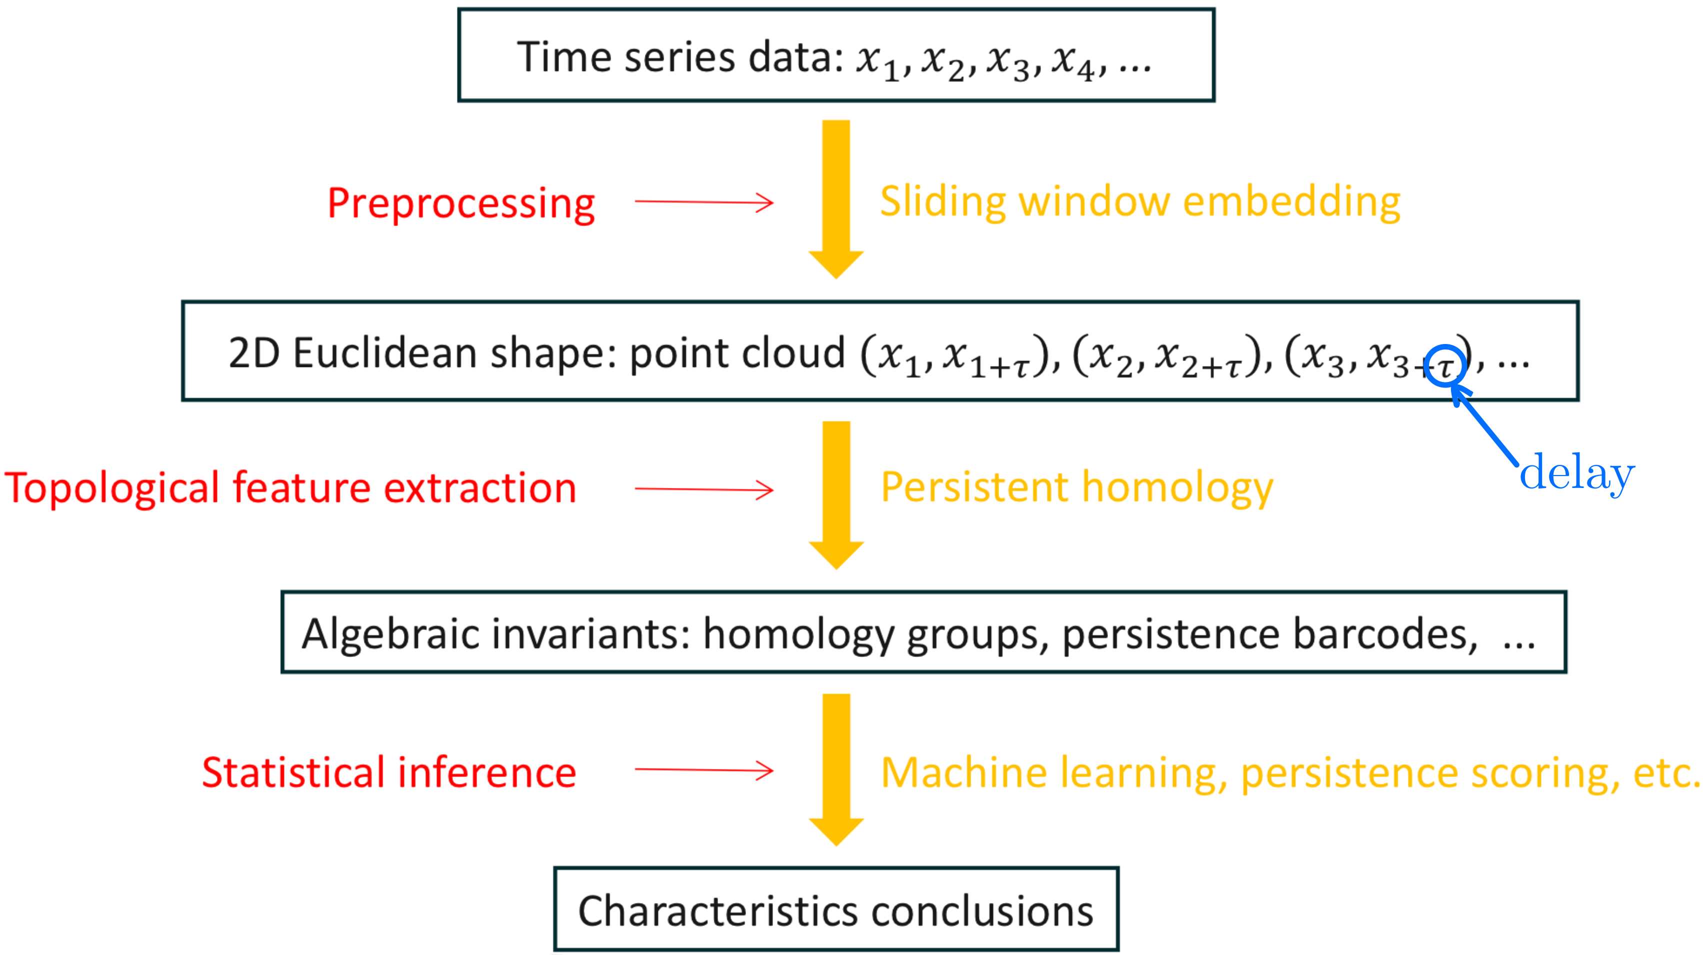

Application I: detection of mouse scratching behavior

Joint with the biomedical engineering group led by Fangyi Chen and the data science group led by Zhen Zhang, both at SUSTech, we applied topological methods to the problem of automated and real-time detection of mouse scratching behavior, with motivations from pharmacology.

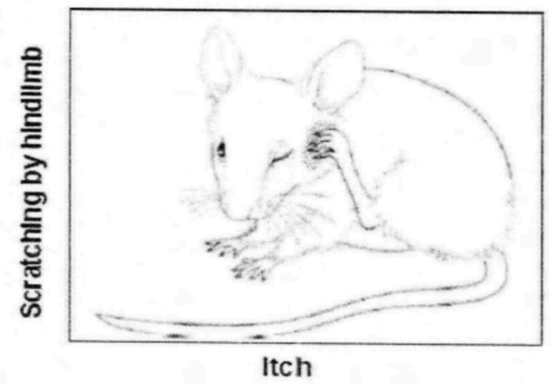

Application I: detection of mouse scratching behavior

Joint with the biomedical engineering group led by Fangyi Chen and the data science group led by Zhen Zhang, both at SUSTech, we applied topological methods to the problem of automated and real-time detection of mouse scratching behavior, with motivations from pharmacology.

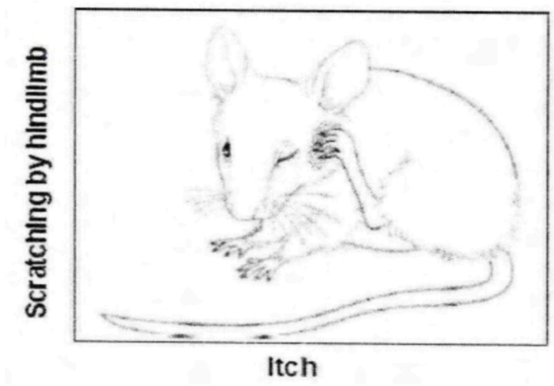

Prior to our group's involvement, machine learning via neural networks was applied with satisfactory accuracy (https://yifeizhu.github.io/scratch.mp4).

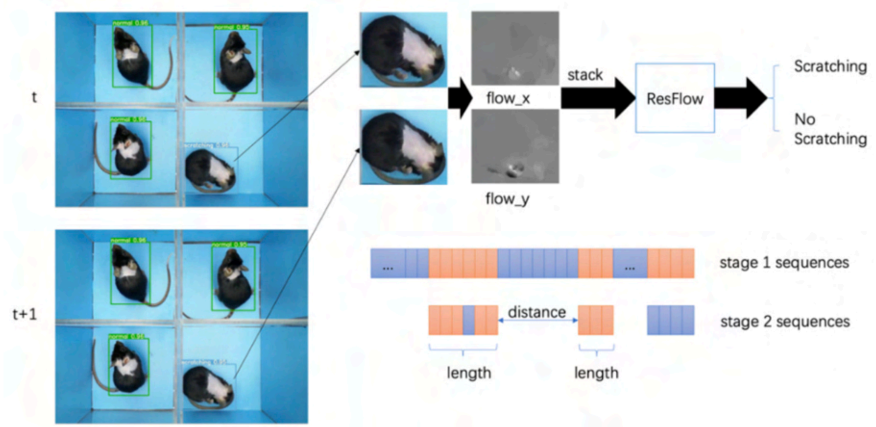

Application I: detection of mouse scratching behavior

Joint with the biomedical engineering group led by Fangyi Chen and the data science group led by Zhen Zhang, both at SUSTech, we applied topological methods to the problem of automated and real-time detection of mouse scratching behavior, with motivations from pharmacology.

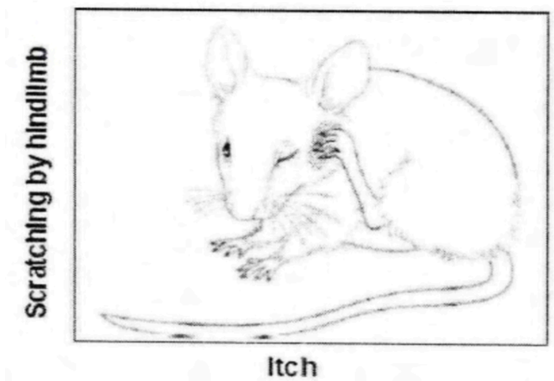

Prior to our group's involvement, machine learning via neural networks was applied with satisfactory accuracy (https://yifeizhu.github.io/scratch.mp4).

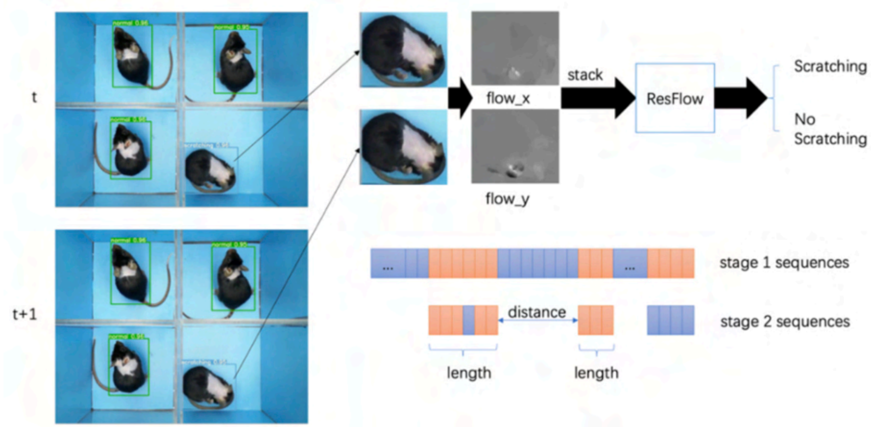

However, the learning process was time consuming, which is impractical for time-sensitive purposes and lab efficiency.

Application I: detection of mouse scratching behavior

We observed that the scratching behavior exhibits periodicity.

Application I: detection of mouse scratching behavior

We observed that the scratching behavior exhibits periodicity. In the meantime, however, global movements of a mouse may significantly reduce the pattern.

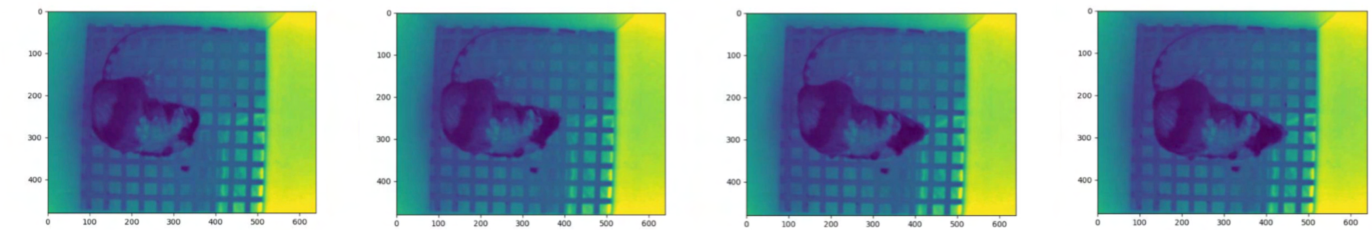

Application I: detection of mouse scratching behavior

We observed that the scratching behavior exhibits periodicity. In the meantime, however, global movements of a mouse may significantly reduce the pattern.

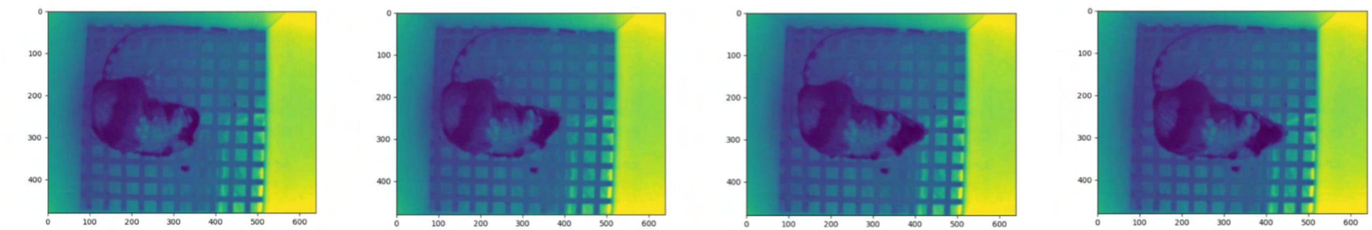

To resolve this issue, we adopted the following approaches:

Approach 1 Sum up all 460 x 640 pixels to extract a series of 1D data which ignores differences caused by global movements. Too coarse?

Application I: detection of mouse scratching behavior

We observed that the scratching behavior exhibits periodicity. In the meantime, however, global movements of a mouse may significantly reduce the pattern.

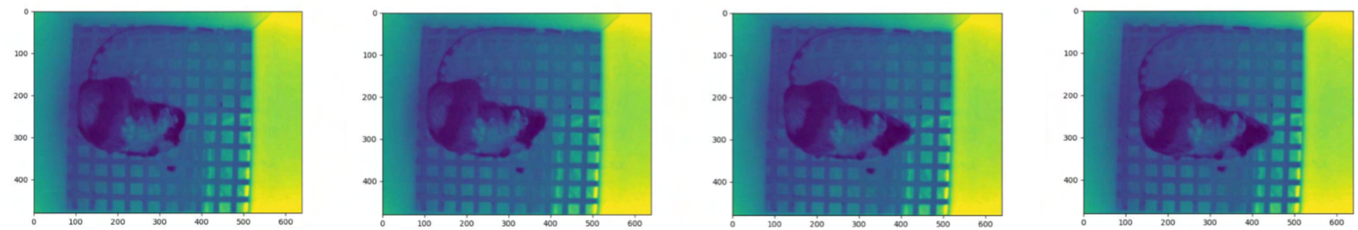

To resolve this issue, we adopted the following approaches:

Approach 1 Sum up all 460 x 640 pixels to extract a series of 1D data which ignores differences caused by global movements. Too coarse?

Approach 2 Blur the images by pooling, and feed the topological pipeline with reduced 100-dimensional data. Still too refined?

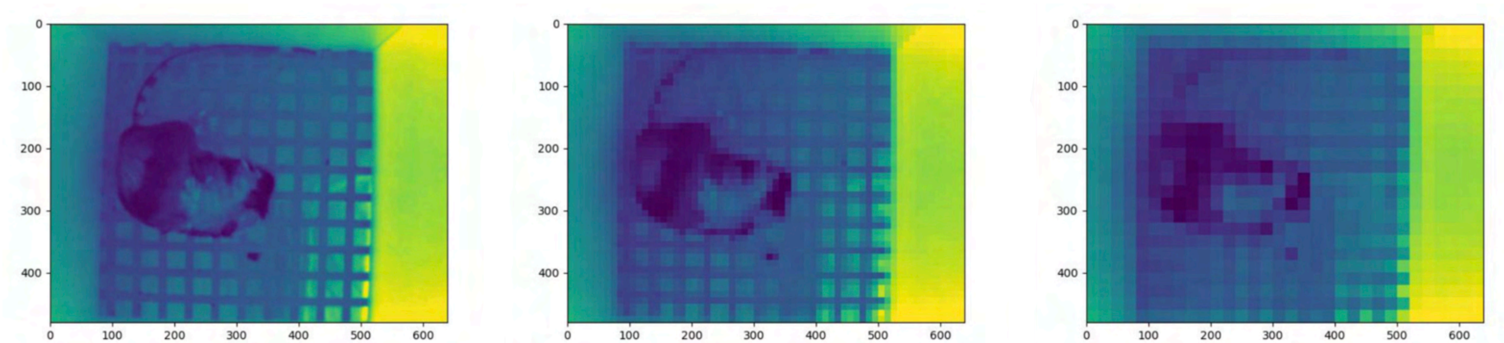

Application I: detection of mouse scratching behavior

Approach 1 (1D data), combined with carefully designed filtration for wave signals + suitably chosen geometric statistics, yielded a close-to-real-time, decently accurate detection performance.

Application I: detection of mouse scratching behavior

Approach 1 (1D data), combined with carefully designed filtration for wave signals + suitably chosen geometric statistics, yielded a close-to-real-time, decently accurate detection performance.

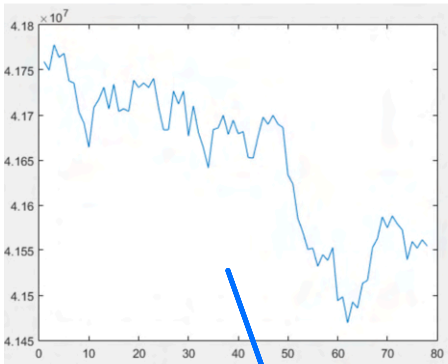

2 filtrations:

$$f' = f * K_3$$
$$f'' = f' \square f' * K_7$$

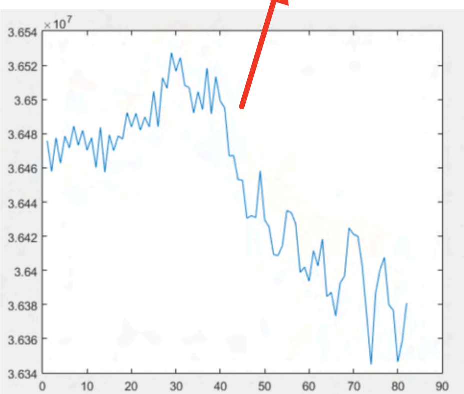

Application I: detection of mouse scratching behavior

Approach 1 (1D data), combined with carefully designed filtration for wave signals + suitably chosen geometric statistics, yielded a close-to-real-time, decently accurate detection performance.

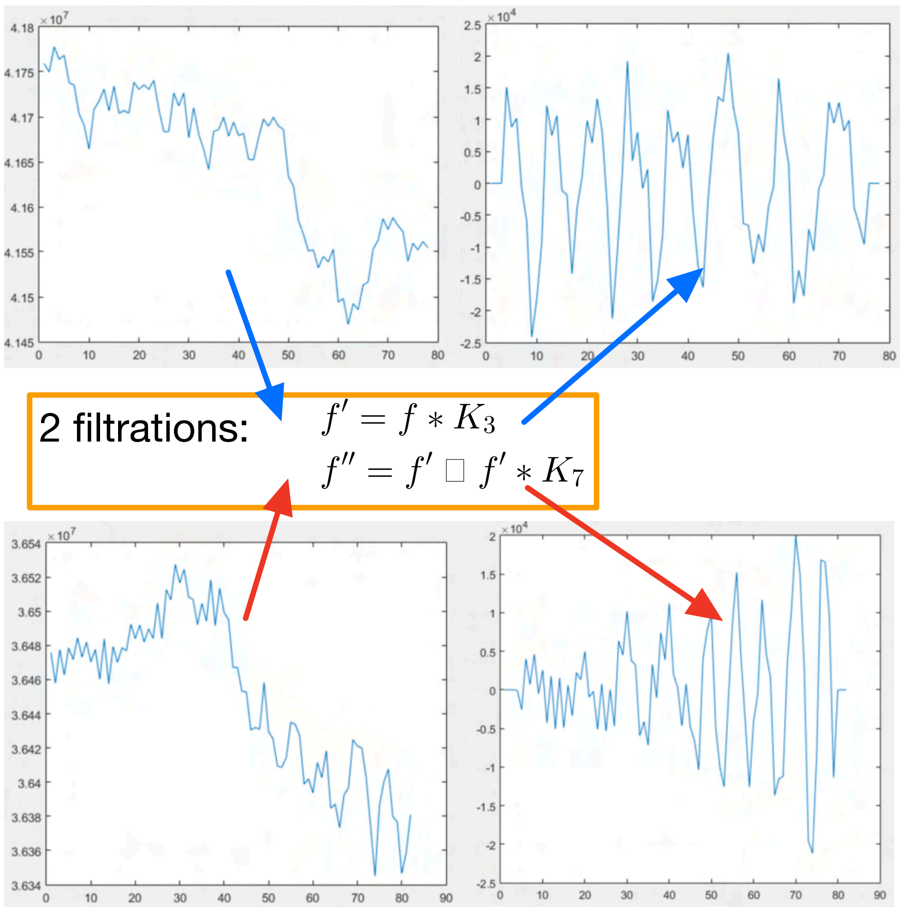

Application I: detection of mouse scratching behavior

Approach 1 (1D data), combined with carefully designed filtration for wave signals + suitably chosen geometric statistics, yielded a close-to-real-time, decently accurate detection performance.

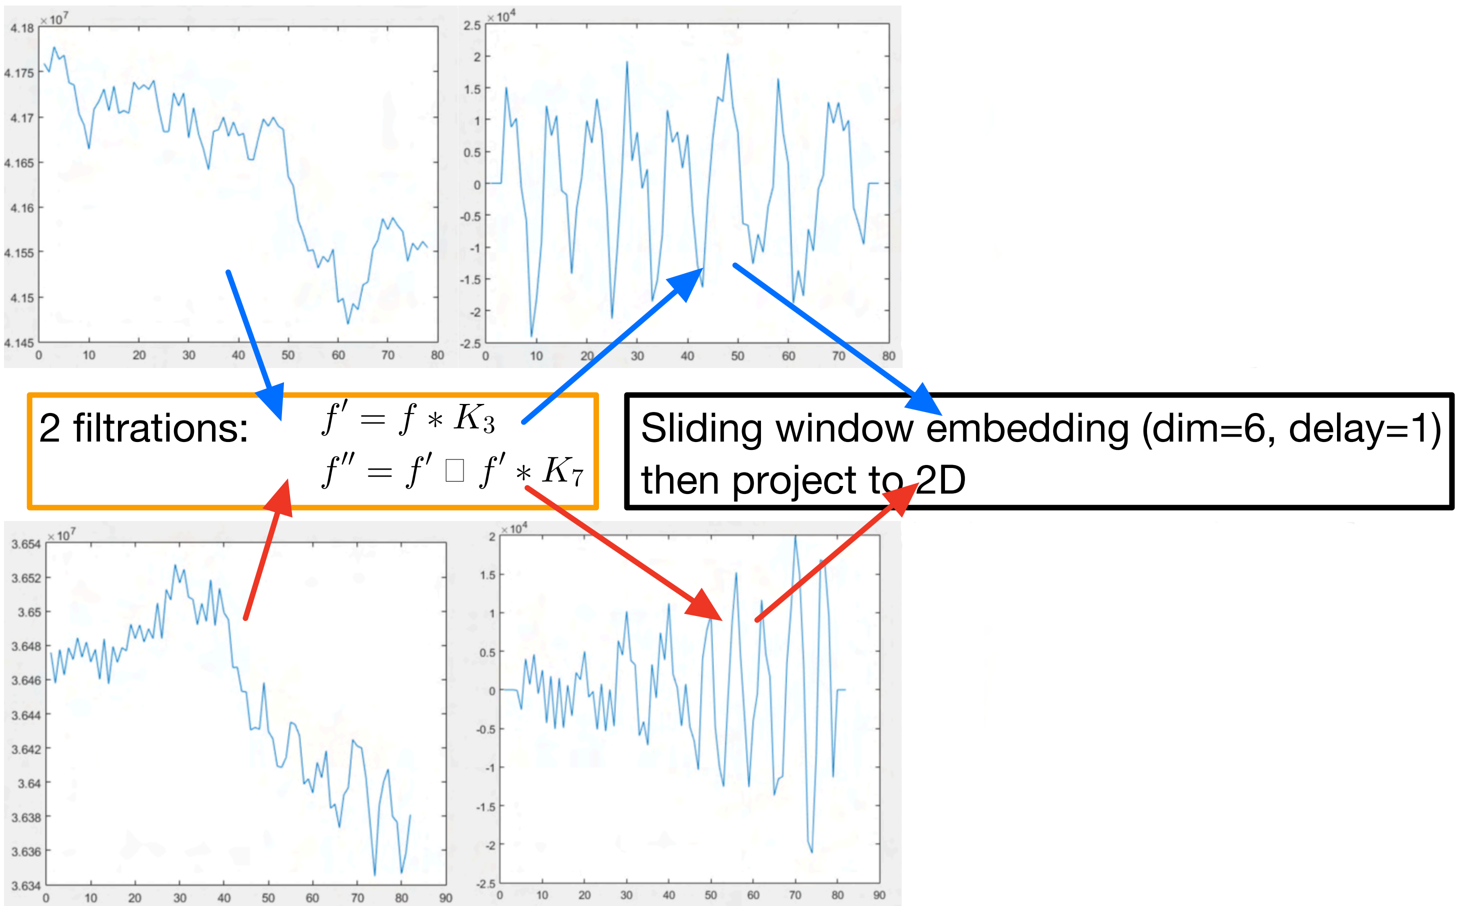

Application I: detection of mouse scratching behavior

Approach 1 (1D data), combined with carefully designed filtration for wave signals + suitably chosen geometric statistics, yielded a close-to-real-time, decently accurate detection performance.

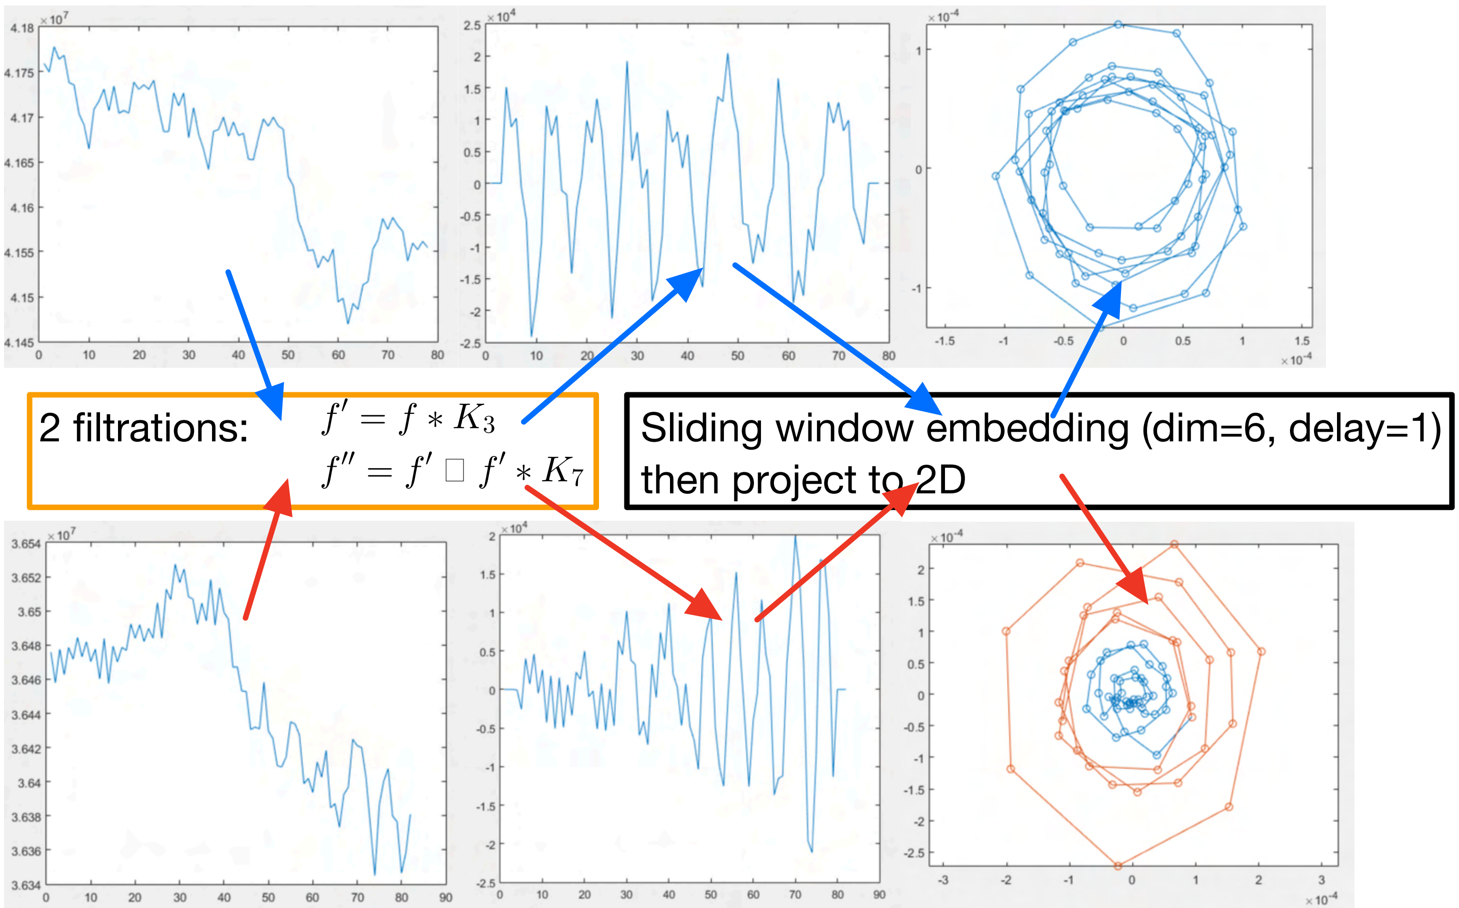

Application I: detection of mouse scratching behavior

Approach 1 (1D data), combined with carefully designed filtration for wave signals + suitably chosen geometric statistics, yielded a close-to-real-time, decently accurate detection performance.

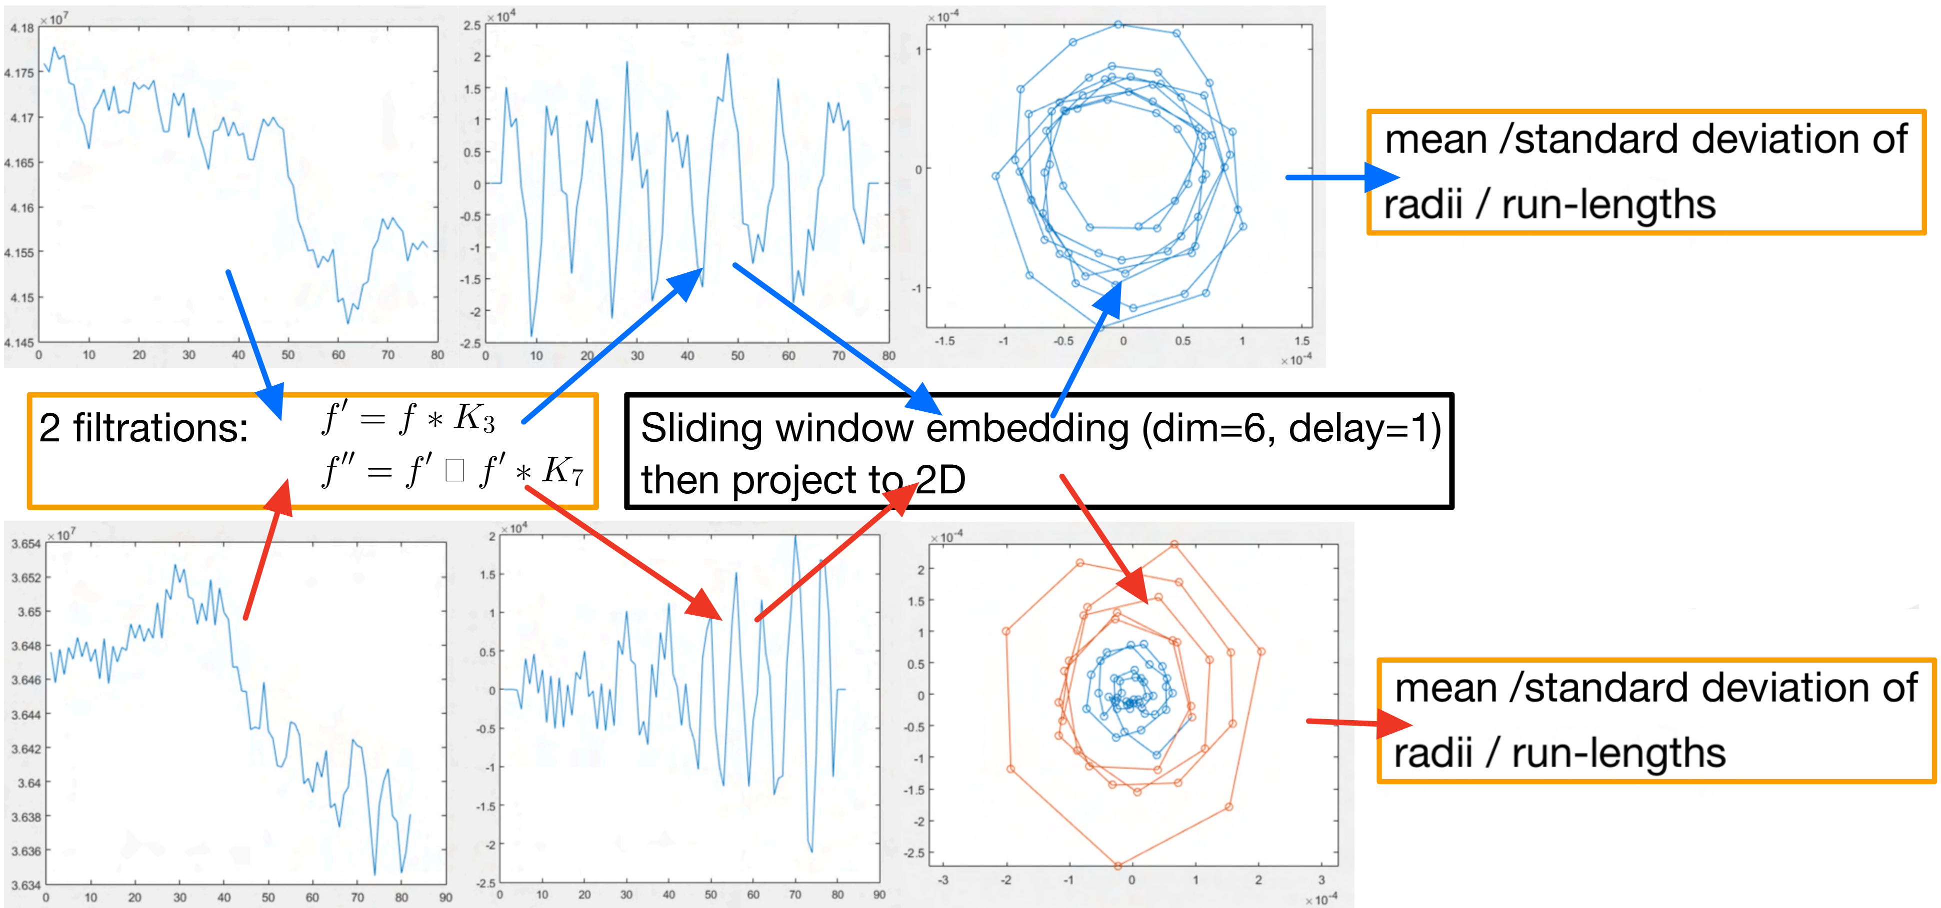

Application I: detection of mouse scratching behavior

Approach 1 (1D data), combined with carefully designed filtration for wave signals + suitably chosen geometric statistics, yielded a close-to-real-time, decently accurate detection performance.

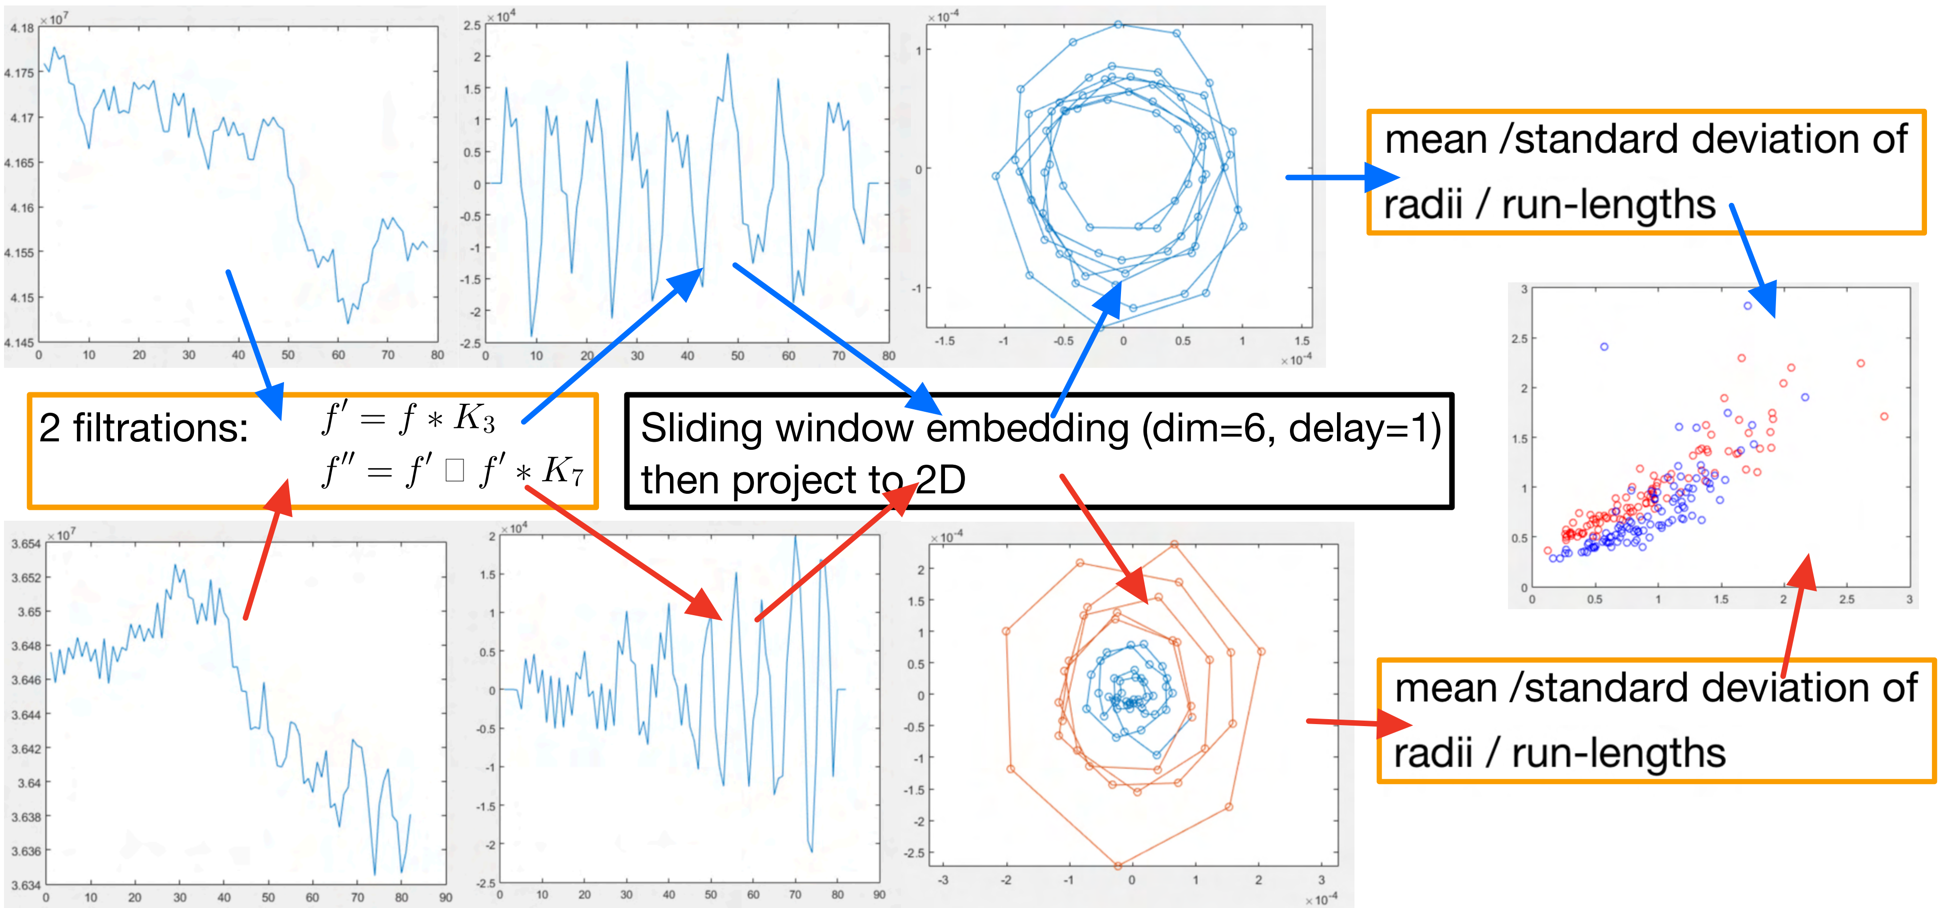

Application I: detection of mouse scratching behavior

Approach 2 (multi-dimensional data), combined with persistent homology and its representations, yielded recognizable characteristics but required considerable computational time.

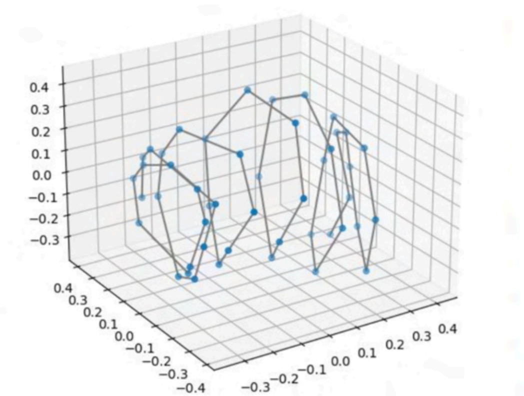

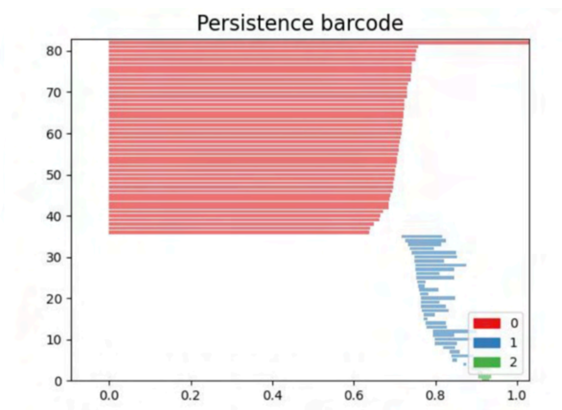

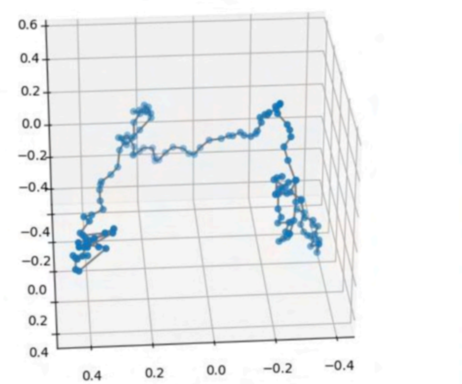

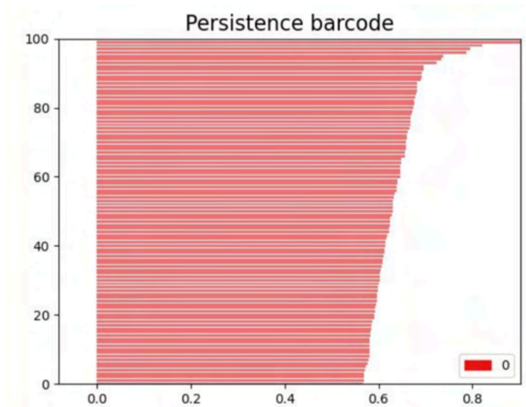

Application II: classification of voiced and unvoiced speech signals

Joint with Meng Yu of Tencent AI Lab, we applied topological methods to classify voiced/unvoiced and vowel/consonant speech data, with motivations from industrial applications.

Application II: classification of voiced and unvoiced speech signals

Joint with Meng Yu of Tencent AI Lab, we applied topological methods to classify voiced/unvoiced and vowel/consonant speech data, with motivations from industrial applications.

We were inspired by Carlsson et al.'s discovery of the Klein-bottle distribution of local natural images, as well as their subsequent recent work of topological convolutional neural networks learning video data.

Application II: classification of voiced and unvoiced speech signals

Joint with Meng Yu of Tencent AI Lab, we applied topological methods to classify voiced/unvoiced and vowel/consonant speech data, with motivations from industrial applications.

We were inspired by Carlsson et al.'s discovery of the Klein-bottle distribution of local natural images, as well as their subsequent recent work of topological convolutional neural networks learning video data. We would like to understand an analogous "moduli space" for speech data and how its input may enable smarter learning.

Application II: classification of voiced and unvoiced speech signals

Joint with Meng Yu of Tencent AI Lab, we applied topological methods to classify voiced/unvoiced and vowel/consonant speech data, with motivations from industrial applications.

We were inspired by Carlsson et al.'s discovery of the Klein-bottle distribution of local natural images, as well as their subsequent recent work of topological convolutional neural networks learning video data. We would like to understand an analogous "moduli space" for speech data and how its input may enable smarter learning.

Display of speech signals

There are speech signal processing softwares for professional use.

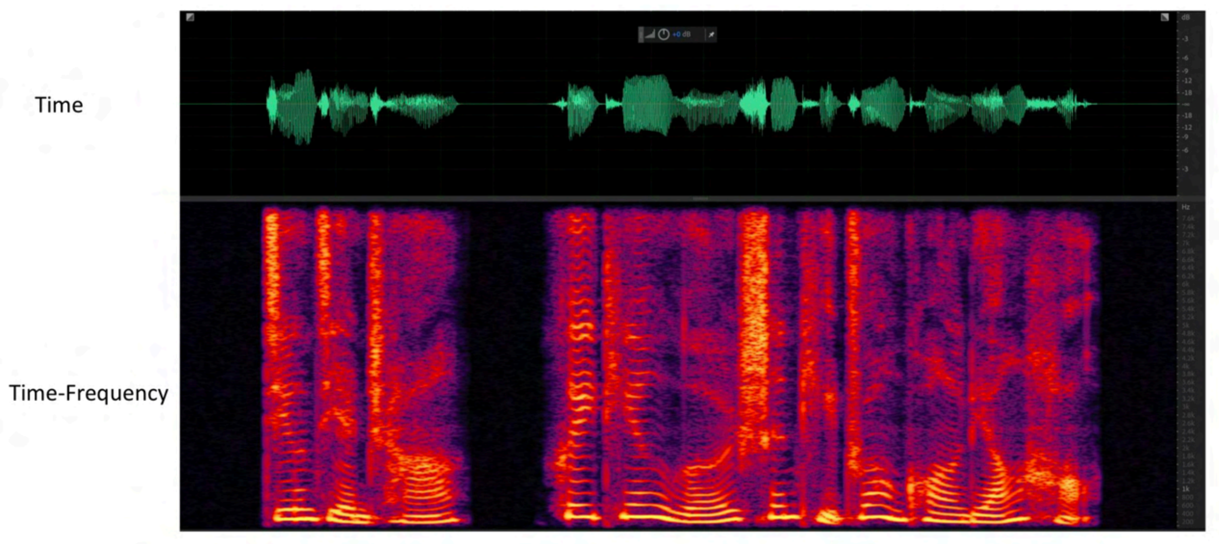

Application II: classification of voiced and unvoiced speech signals

Joint with Meng Yu of Tencent AI Lab, we applied topological methods to classify voiced/unvoiced and vowel/consonant speech data, with motivations from industrial applications.

We were inspired by Carlsson et al.'s discovery of the Klein-bottle distribution of local natural images, as well as their subsequent recent work of topological convolutional neural networks learning video data. We would like to understand an analogous "moduli space" for speech data and how its input may enable smarter learning.

Display of speech signals

There are speech signal processing softwares for professional use.

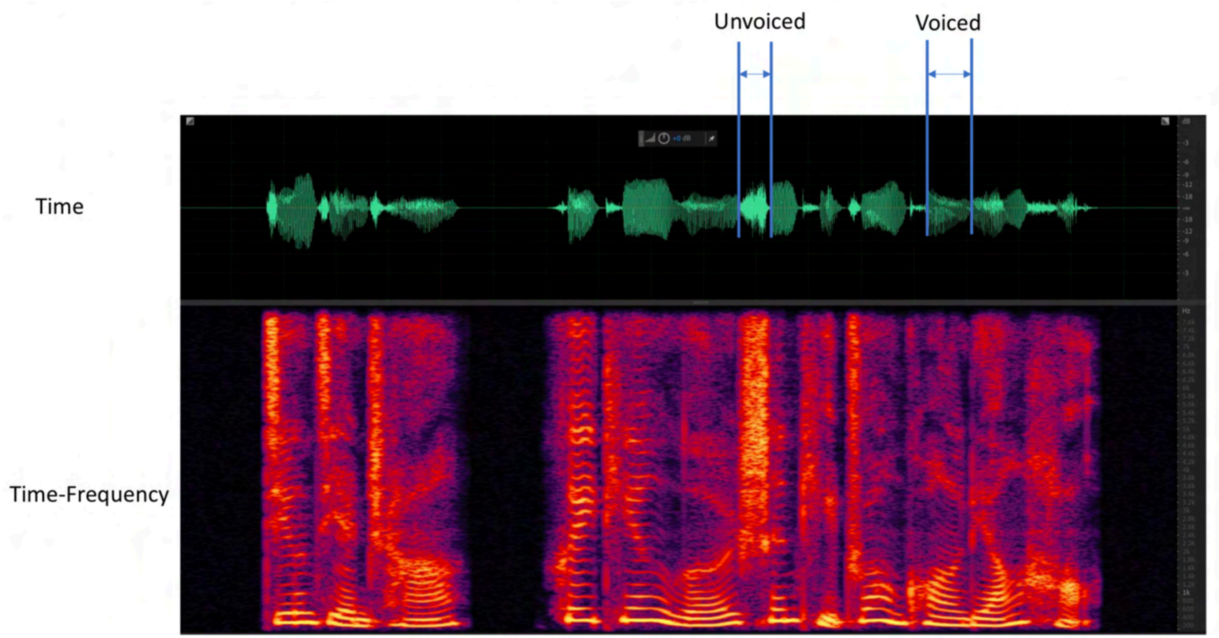

Application II: classification of voiced and unvoiced speech signals

Voiced

Sinusoid in time domain

Harmonics in frequency domain

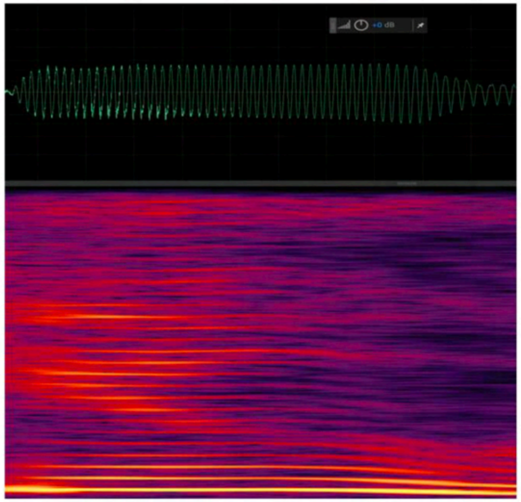

Time and Time-Frequency domain

Frequency response

Application II: classification of voiced and unvoiced speech signals

Voiced

Sinusoid in time domain

Harmonics in frequency domain

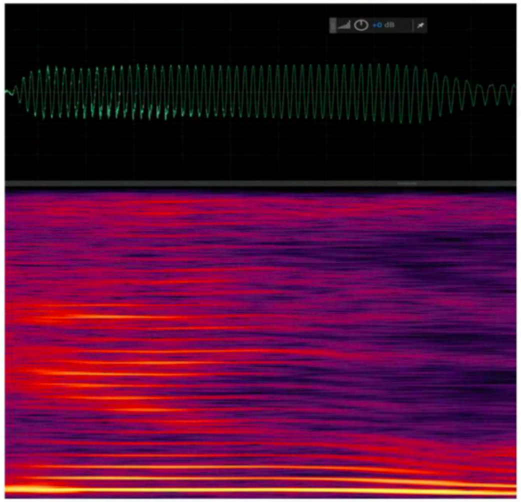

Time and Time-Frequency domain

Frequency response

Unvoiced

Like a white noise

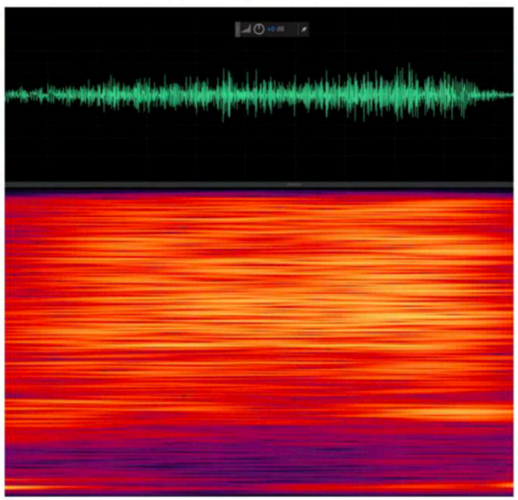

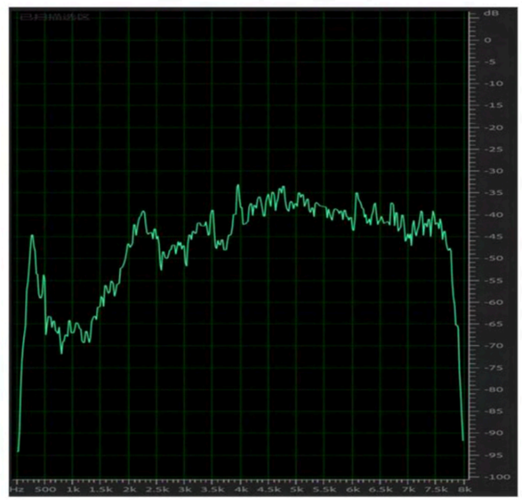

Application II: classification of voiced and unvoiced speech signals

Here is a flowchart for our topological approach:

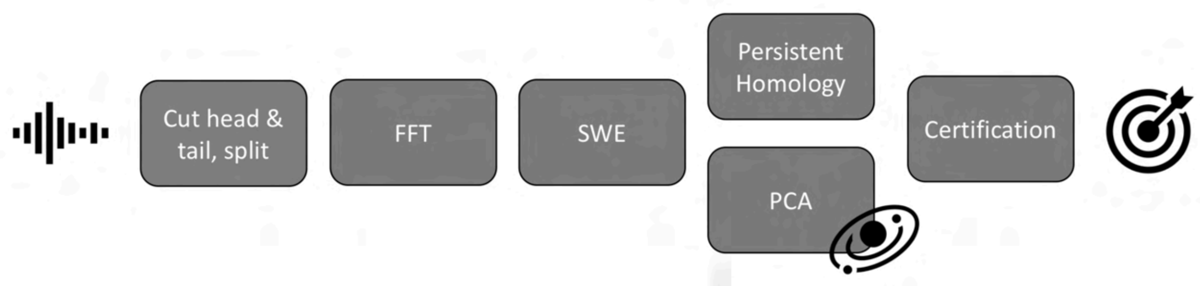

Application II: classification of voiced and unvoiced speech signals

Here is a flowchart for our topological approach:

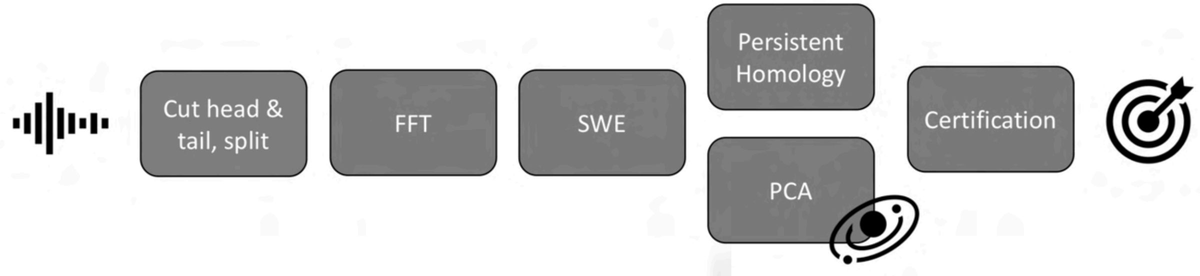

Topological profiles for vowels and consonants

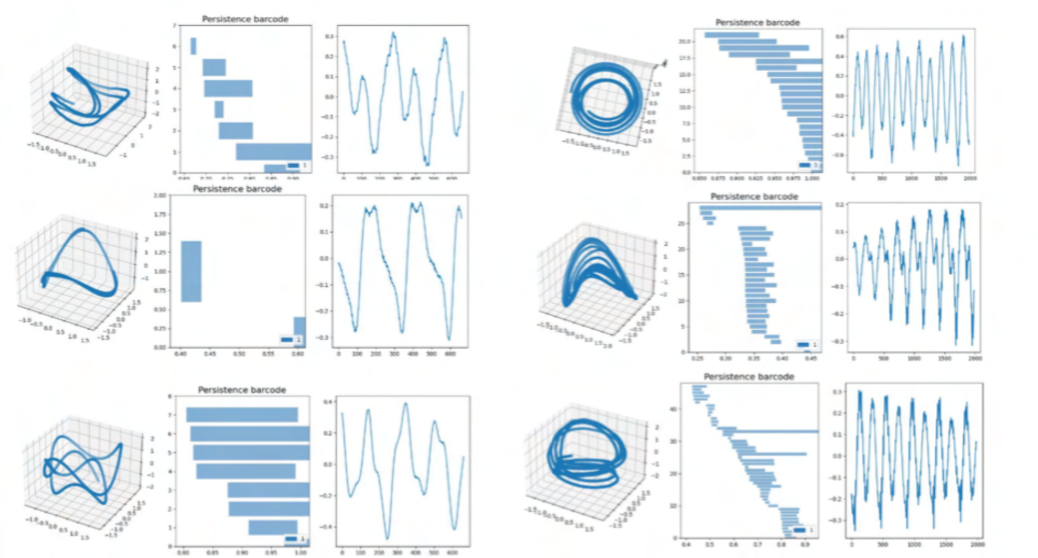

Features for vowels

Left: frame size: 15ms, frame shift: 5ms; Right: frame size: 45ms, frame shift: 22.5ms

Application II: classification of voiced and unvoiced speech signals

Here is a flowchart for our topological approach:

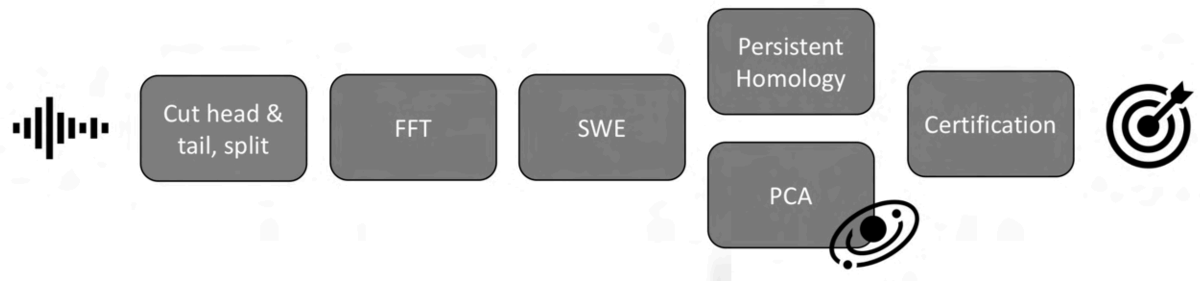

Topological profiles for vowels and consonants

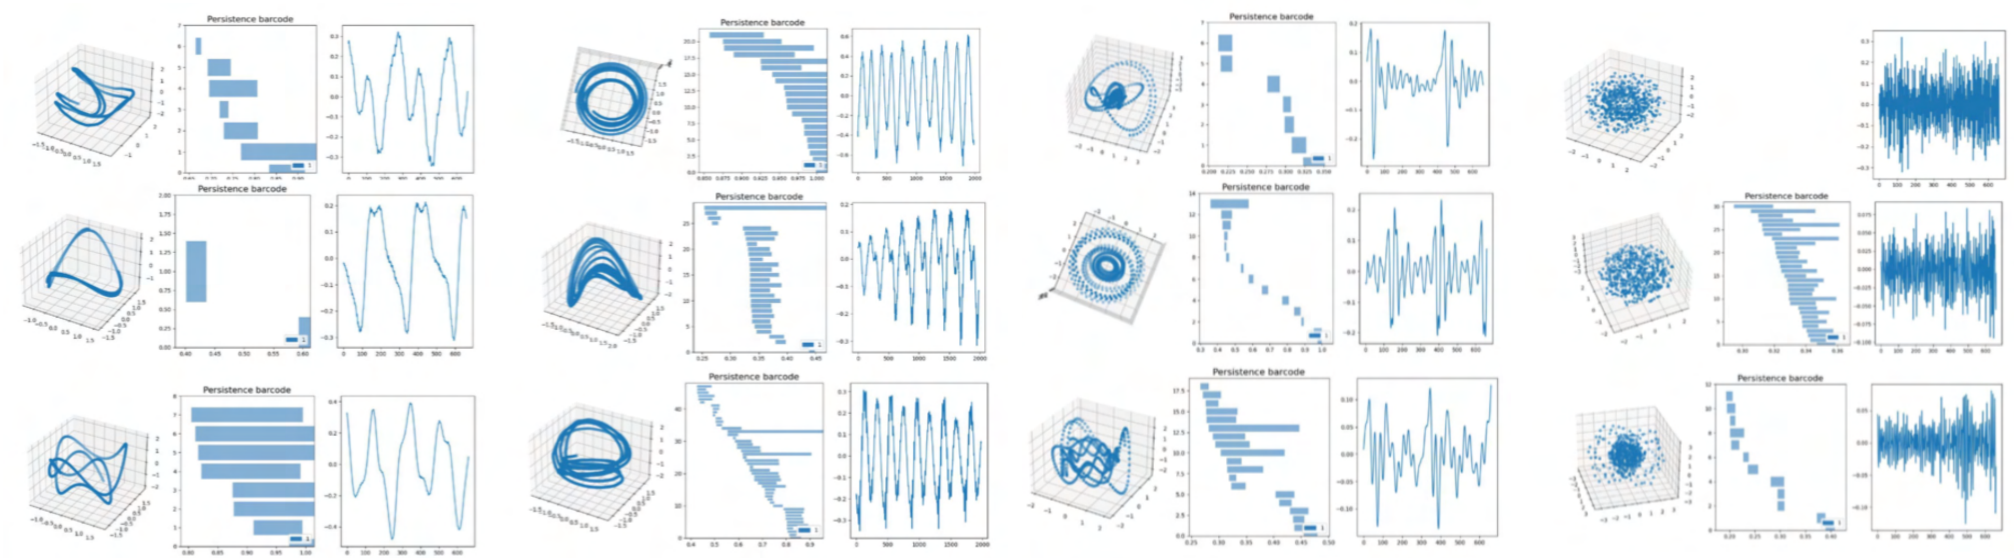

Features for vowels

Features for consonants

Left: frame size: 15ms, frame shift: 5ms; Right: frame size: 45ms, frame shift: 22.5ms

Left: pulmonic consonant; Right: non-pulmonic consonant

Application II: classification of voiced and unvoiced speech signals

Using real-world speech data from the MFA aligner, we further fed the topological features for machine learning, and obtained positive preliminary results for classification.

Application II: classification of voiced and unvoiced speech signals

Using real-world speech data from the MFA aligner, we further fed the topological features for machine learning, and obtained positive preliminary results for classification.

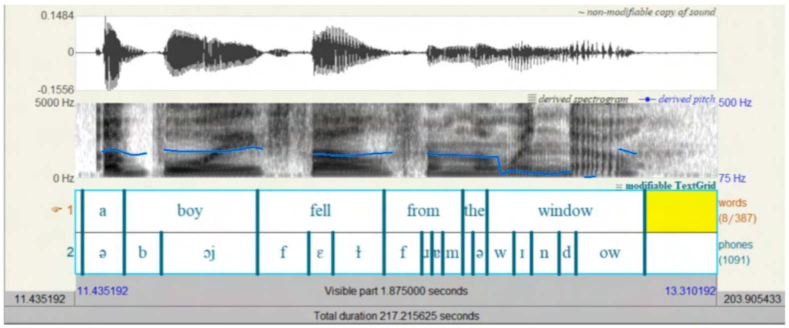


```
vowel_phones=['ɔj','ɛ','ə','ɪ','aj','
ɑ','æ','i','o','ʊ','aw','e','u','a']
consonant_phones=['b','f','m','ɹ','ð'
,'w','h','p','t','z','n','g','dʒ','s'
,'ʃ','v','l','ŋ','k','θ','j','tʃ','ʒ'
,'d']
```

Application II: classification of voiced and unvoiced speech signals

Using real-world speech data from the MFA aligner, we further fed the topological features for machine learning, and obtained positive preliminary results for classification.

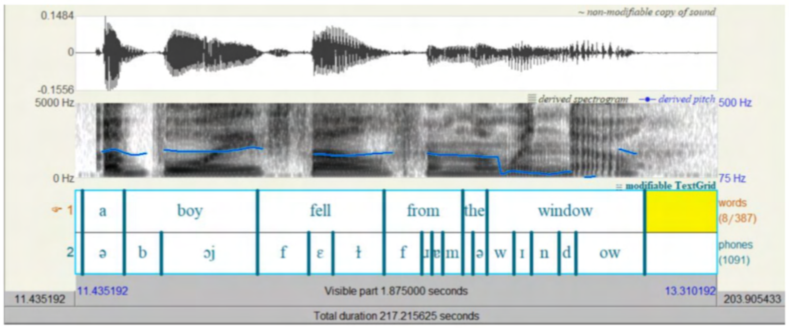


```
vowel_phones=['ɔj','ɛ','ə','ɪ','aj','ɑ','æ','i','o','ʊ','aw','e','u','a']
consonant_phones=['b','f','m','ɹ','ð','w','h','p','t','z','n','g','dʒ','s','ʃ','v','l','ŋ','k','θ','j','tʃ','ʒ','d']
```

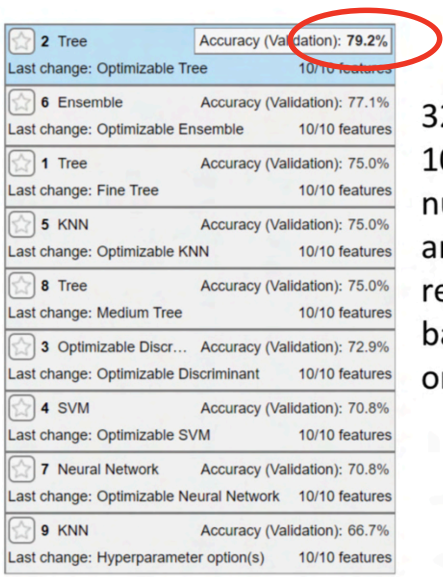

		Accuracy (Validation)	Last change	
2	Tree	79.2%	Optimizable Tree	10/10 features
6	Ensemble	77.1%	Optimizable Ensemble	10/10 features
1	Tree	75.0%	Fine Tree	10/10 features
5	KNN	75.0%	Optimizable KNN	10/10 features
8	Tree	75.0%	Medium Tree	10/10 features
3	Optimizable Discr...	72.9%	Optimizable Discriminant	10/10 features
4	SVM	70.8%	Optimizable SVM	10/10 features
7	Neural Network	70.8%	Optimizable Neural Network	10/10 features
9	KNN	66.7%	Hyperparameter option(s)	10/10 features

32 vowels, 16 consonants.
10 features: 5 are barcodes number of 5 diag, other 5 are number of barcodes that reaches inf(both consider barcode of 1 dimension for only)

Application II: classification of voiced and unvoiced speech signals

Using real-world speech data from the MFA aligner, we further fed the topological features for machine learning, and obtained positive preliminary results for classification.

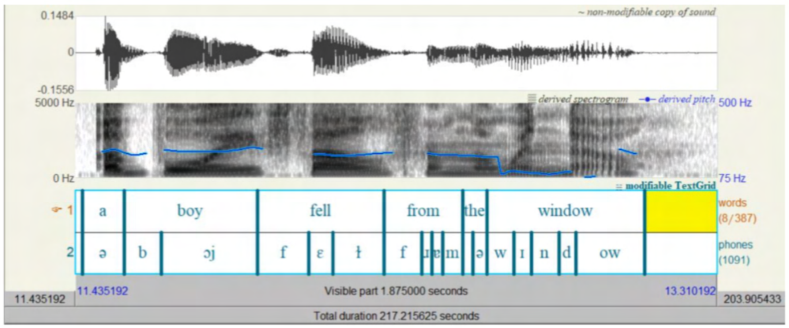


```
vowel_phones=['ɔj','ɛ','ə','ɪ','aj','
ɑ','æ','i','o','ʊ','aw','e','u','a']
consonant_phones=['b','f','m','ɹ','ð'
,'w','h','p','t','z','n','g','dʒ','s'
,'ʃ','v','l','ŋ','k','θ','j','tʃ','ʒ'
,'d']
```

Model	Accuracy (Validation)	Last change	Features
2 Tree	79.2%	Optimizable Tree	10/10 features
6 Ensemble	77.1%	Optimizable Ensemble	10/10 features
1 Tree	75.0%	Fine Tree	10/10 features
5 KNN	75.0%	Optimizable KNN	10/10 features
8 Tree	75.0%	Medium Tree	10/10 features
3 Optimizable Discr...	72.9%	Optimizable Discriminant	10/10 features
4 SVM	70.8%	Optimizable SVM	10/10 features
7 Neural Network	70.8%	Optimizable Neural Network	10/10 features
9 KNN	66.7%	Hyperparameter option(s)	10/10 features

32 vowels, 16 consonants. 10 features: 5 are barcodes number of 5 diag, other 5 are number of barcodes that reaches inf(both consider barcode of 1 dimension for only)

Model	Accuracy (Validation)	Last change	Features
1 Tree	81.5%	Fine Tree	4/4 features
2 Tree	81.5%	Optimizable Tree	4/4 features
7 Tree	81.5%	Medium Tree	4/4 features
4 Tree	78.5%	Coarse Tree	4/4 features
3 KNN	69.2%	Optimizable KNN	4/4 features
5 Neural Network	46.2%	Hyperparameter option(s)	4/4 features
6 Neural Network	46.2%	Narrow Neural Network	4/4 features

32 vowels, 33 consonants. 4 features: bottleneck distance between neighborhood barcode(currently the best result)

Persistent homology

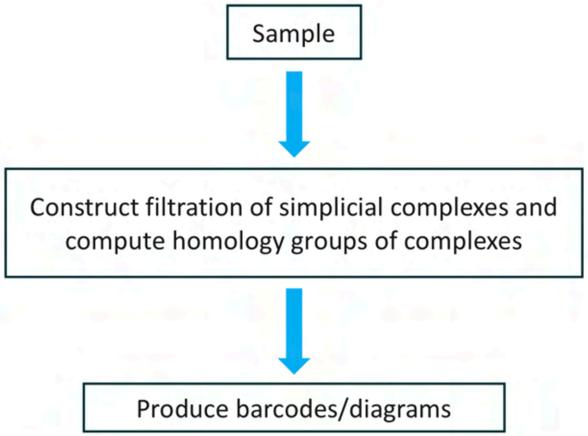

How filtration through varying distance measure reveals essential topological features

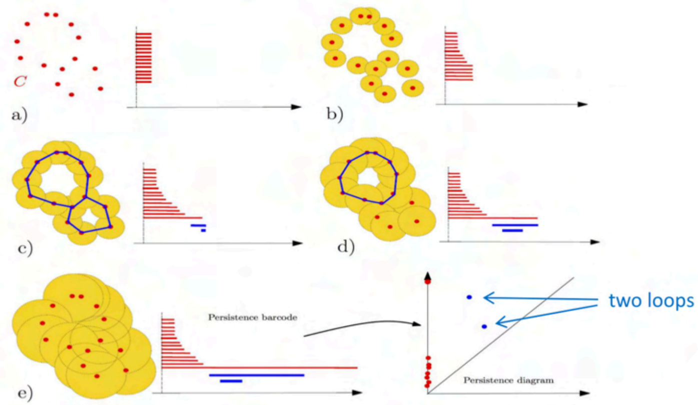

Persistent homology

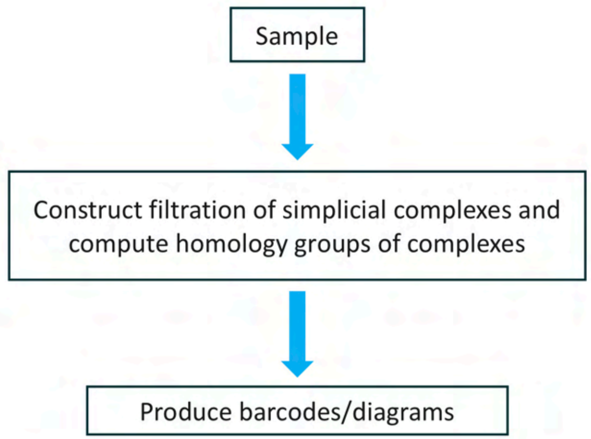

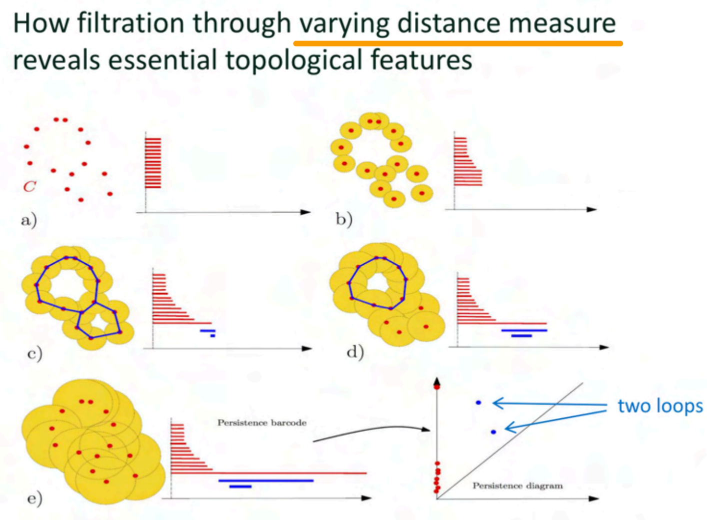

Sliding window embedding

Euclidean embedding of time series data dates back to Takens's work on fluid turbulence in the 1980s.

Persistent homology

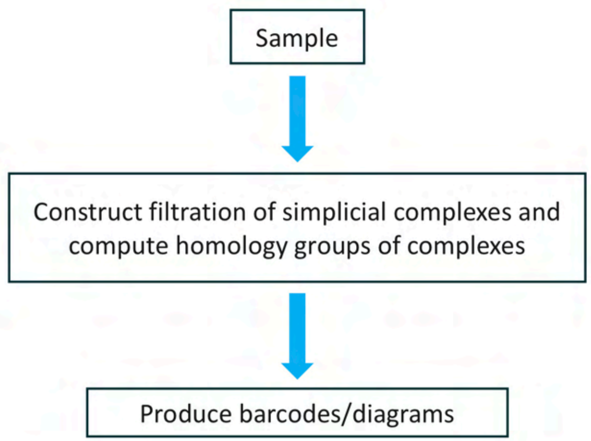

How filtration through varying distance measure reveals essential topological features

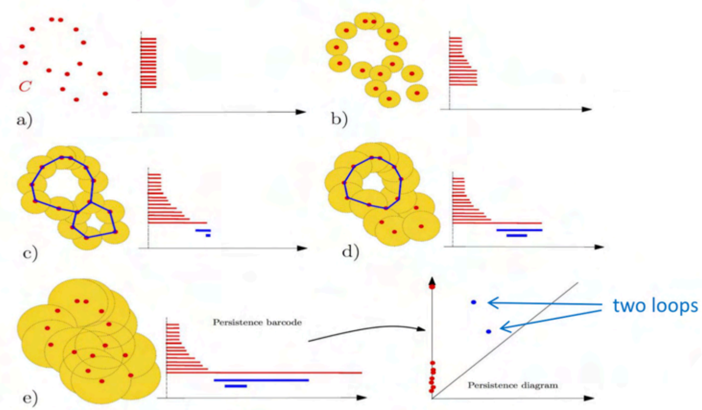

Sliding window embedding

Euclidean embedding of time series data dates back to Takens's work on fluid turbulence in the 1980s.

Theorem (Takens 1981). Let M be a compact manifold of dimension n. Given pairs (ϕ, y) with $\Phi : M \to M$ a smooth diffeomorphism and $y : M \to \mathbb{R}$ a smooth function, it is a generic property that the map $\Phi_{(\phi,\, y)} : M \to \mathbb{R}^{2n+1}$ defined by

$$\Phi_{(\varphi,y)}(x) = \Big(y\,(x)\,,\, y \circ \varphi(x)),\, \ldots\,,\, y \circ \varphi^{2n}(x))\Big)$$

is an embedding.

Thank you.

Credits and references

- Plots of orbits on a torus from
  Jaume Masoliver and Ana Ros Camacho, *Integrability and chaos: the classical uncertainty*, European Journal of Physics, 2011
- TDA diagrams and flowchart by Siheng Yi
- Wheeze picture from https://londonchestspecialist.co.uk/wheeze-treatment-online-appointments-consultation/
- Mouse picture, video, and experimental settings from Min Chen
- Processed mouse scratching pictures by Siheng Yi
- Application I approach 1 designed and realized by Qingrui Qu
- Application I approach 2 designed and realized by Siheng Yi
- Speech signal time-frequency charts from Meng Yu
- Application II designed and realized by Pingyao Feng, with Siheng Yi and Qingrui Qu
- Persistent homology charts from Siheng Yi

- Introductory texts to TDA:
  - Gunnar Carlsson and Mikael Vejdemo-Johansson, *Topological data analysis with applications*, Cambridge University Press, 2021
  - Herbert Edelsbrunner and John L. Harer, *Computational topology: an introduction*, American Mathematical Society, 2010

- More information: https://sustech-topology.github.io/acts/