EEP 596: AI and Health Care || Lecture 5

Dr. Karthik Mohan

Univ. of Washington, Seattle

Apr 11, 2022

Logistics

- **Late days:** 5 late days for the whole quarter - can use it whenever. After you use up the late days - You will incur a penalty on the grade.

Logistics

- **Late days:** 5 late days for the whole quarter - can use it whenever. After you use up the late days - You will incur a penalty on the grade.
- **Mini-Project 1:** Will be based on wearable device and data collected and analyzed from the device. Exact description of mini-project will be released early this week.

Logistics

- **Late days:** 5 late days for the whole quarter - can use it whenever. After you use up the late days - You will incur a penalty on the grade.
- **Mini-Project 1:** Will be based on wearable device and data collected and analyzed from the device. Exact description of mini-project will be released early this week.
- Anything else?

Poll!

Last Lecture

- **Random Forests**

Last Lecture

- **Random Forests**
- **Wearable Devices Overview**

Last Lecture

- **Random Forests**
- **Wearable Devices Overview**
- **Introduction to Anomaly and Change Point Detection**

Today

- **Anomaly and Change Point Detection**

Today

- **Anomaly and Change Point Detection**
- **Case Study: Arrhythmia Data**

New Topic: Wearable Devices

Power of Wearables

Anomaly Detection is key

Deanna Recktenwald, an 18-year-old high school senior, received a ping on her Apple Watch earlier this year that said, “Seek medical attention.” Recktenwald is grateful for taking the message seriously, because emergency room procedures revealed that her kidneys were functioning at a mere 20% and needed immediate intervention. The Apple Watch, with the heart rate tracking function that alerted Recktenwald, is one of many wearable devices in the market that offer patients the ability to track their health, thereby improving their chances of better outcomes.

Case Study: Wearable Devices

"We are working with people with Parkinson's. Our technology can facilitate a meaningful interaction between the wearer and care teams. The wearable can offer an understanding if a new medication is working or not." – Nicholas Constant, **EchoWear LLC**

"Most cases of mild-to-moderately-severe hearing loss can be managed by users themselves, if all the tools work well. The hearing aid advantage will not last more than two years from now." – Alexander Goldin, Founder and CEO, **Alango Technologies**

"We started with sleep. The impact it has, from cognitive function the next day, fasting glucose, hormones, T-cells that fight cancer – these are all linked to our sleep." – Harpreet Singh Rai, CEO**, Ōura Ring**

"Apple launched a revolution in the wearable industry and mainstreamed the device. The "smart" fall detection system that we built is a patented API. It can go into any wearable." – Shea Gregg, CEO, **FallCall Solutions**

Case Study: Wearable Devices

Category	What it is	Examples
Hearables/earbuds	An amplification ear-worn wearable	Apple AirPods, WearandHear, Dime
Smartwatches	Smart watches monitoring activity, health metrics	Apple, Samsung, Fitbit, Garmin, Adapt
Headsets – AR/VR	Internet-connected glasses enable alternate views	MyndVR, Embodied Labs, Rendever
Fitness trackers (no watch)	Step counter, heart rate	Amazon Halo, Vivo, Whoop Strap 3, Fitbit One
Continuous diabetes wearables	Scans detect blood sugar level, patch injects insulin	FreeStyle Libre, Dexcom G5,
Sleep trackers	Wearables noted for sleep tracking	Ōura Ring, Whoop, Fitbit Versa
Wrist-worn Health	Low-sleep indicates risk of dementia	Omron HeartGuide, Amazon Halo, Whoop
Smart jewelry	Ring, Necklace	Trelaware, ADT invisiWear
Dementia zone trackers	Set a range – track movement outside range	MindMe Locate, PocketFinder
Medical Grade wearable, data collection	Blood pressure, mobile EKG, Diabetes patch	Omron HeartGuide, AliveCor, Tidepool
Medical Alert/PERS/Safety	Emergency call, fall detection – in home or out	Medical Guardian, Lively Wearable, UnaliWear

Figure 8 Examples of wearable categories and some of the offerings

Anomaly Detection: Arrythmia

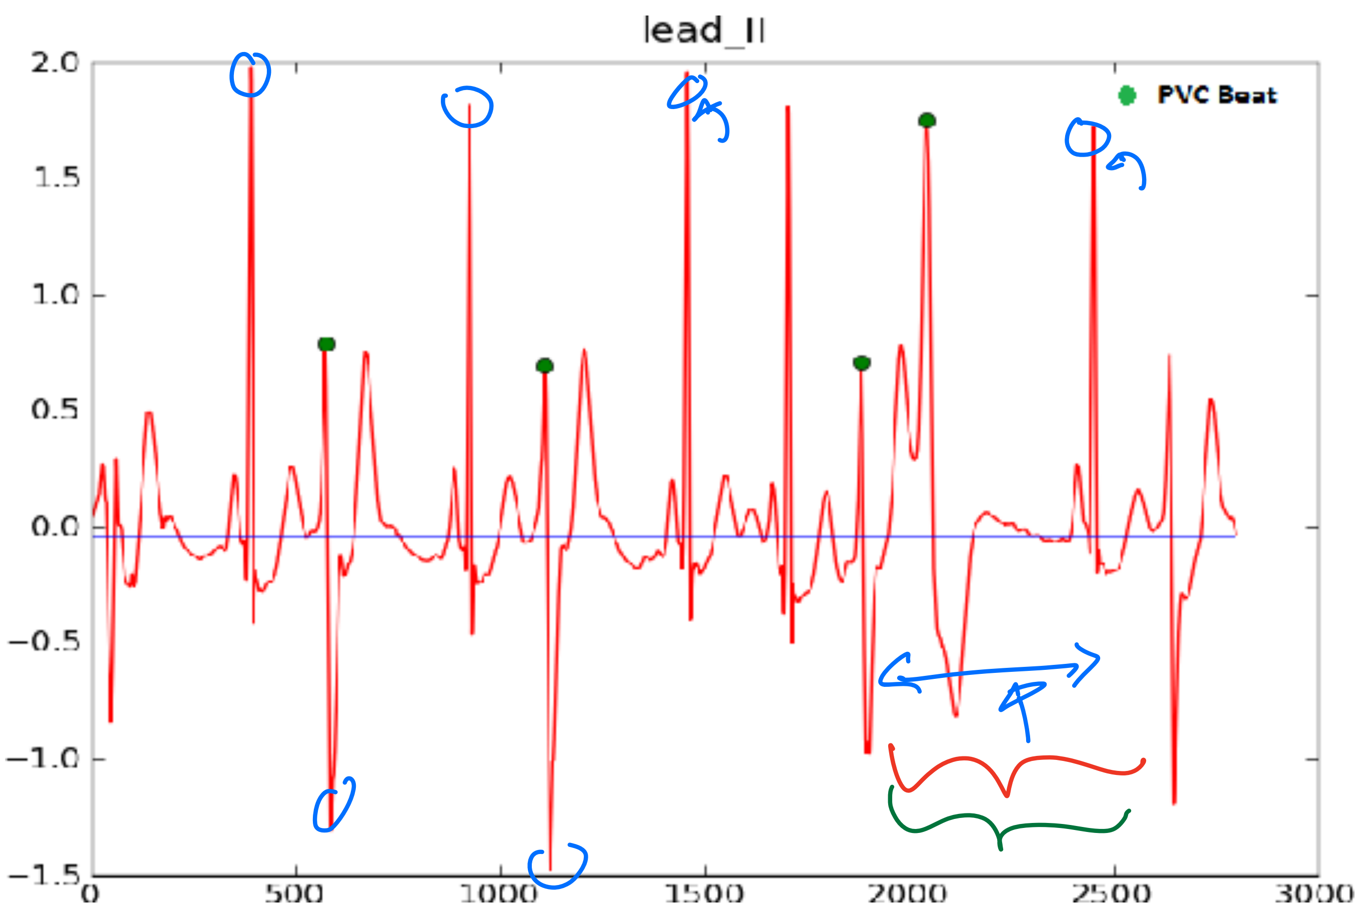

Broad list of methods

Categorization

- Offline anomaly detection
- Real-time anomaly detection

Broad list of methods

Categorization

- Offline anomaly detection
- Real-time anomaly detection

wearables

Categorization

- Time-series data anomaly detection
- Regular anomaly detection

Anomaly Detection Baselines

Baselines method

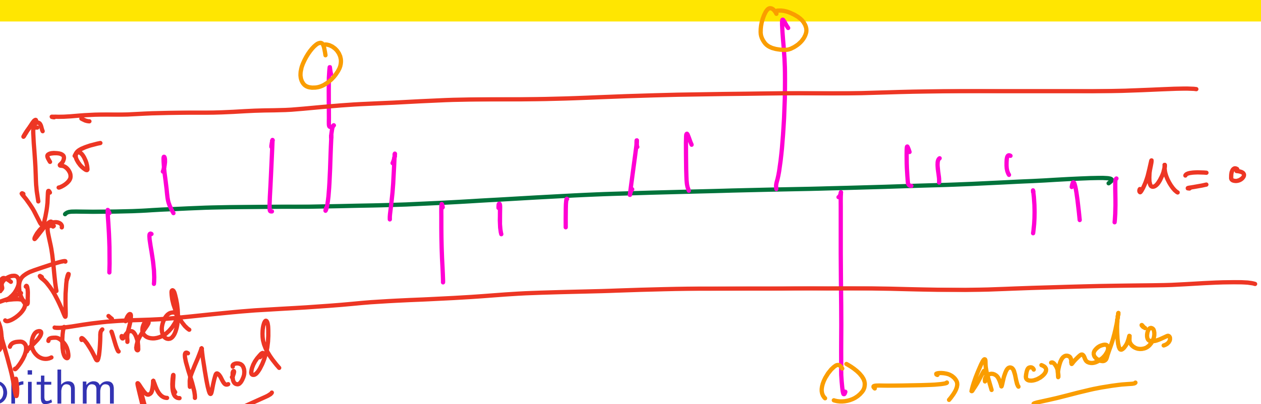

↳ unsupervised method

Baseline Algorithm

Classify a data point as an anomaly if your data point is α standard deviations (σ) away from the mean. Here α is typically greater than 3.

σ $\qquad$ μ

Applications
- BMI/Health Indicator
- BP
- Fever
- Pulse

$X \rightarrow$ Is this an anomaly?

$X > \mu + 3\sigma$ or $X < \mu - 3\sigma$?

Yes → Anomaly!

→ No: Not an anomaly!

Anomaly Detection Baselines

Temperature Example

The mean human body temperature is 98.4 F. Assume now that the thermometer is accurate but normal body temperature fluctuations are expected to be within 0.5 degrees F, then there is cause for concern if temperature deviates beyond $\underbrace{98.4 \pm 1.5}$ for $\underbrace{\alpha = 3}$.

$\sigma = 0.5$

$\alpha = 3$

$\mu = 98.4$

$98.4 + 3\sigma = 98.4 + 1.5$

$= \underbrace{99.9}$

Anomaly Detection Baselines

Temperature Example

The mean human body temperature is 98.4 F. Assume now that the thermometer is accurate but normal body temperature fluctuations are expected to be within 0.5 degrees F, then there is cause for concern if temperature deviates beyond 98.4 ± 1.5 for $\alpha = 3$.

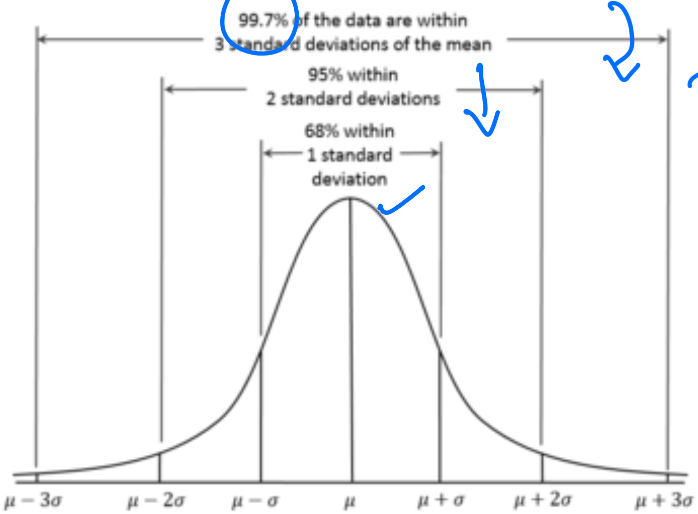

→ Gaussian Distn.

6σ

Anomaly Detection Baselines

$\alpha = 3 \rightarrow$ Good

$\alpha = 6!$

Temperature Example

FP

FN

	α	Outcome
1	2	Lots of false positives and un-necessary trips to urgent care
2	6	Almost no false positives. Might miss early signs of a flu
3	4	Fewer false positives. Get to urgent care at the right time!

False Negative

$98.4 + 6 \times 0.5 = 101.2$

Trade-off bet- FP & FN

Anomaly Detection Trade-off!

[Precision vs Recall]

Anomaly Detection Baselines

Faulty thermometer

Assume you only have a faulty thermometer to measure your body temperature. Sometimes its accurate and sometimes it is not. You get to know that its reading can fluctate up to 3 degrees over or below the true temperature. You measure your temperature and its 102 degrees F. Should you head to the urgent care?

take many measurements

Get new thermometer

Anomaly Detection Baselines

Faulty thermometer

Assume you only have a faulty thermometer to measure your body temperature. Sometimes its accurate and sometimes it is not. You get to know that its reading can fluctate up to 3 degrees over or below the true temperature. You measure your temperature and its 102 degrees F. Should you head to the urgent care?

False positives vs False Negatives

Anomaly Detection methods get caught between controlling false positives and not missing True positives (i.e. having false negatives). Would you rather flag a social media post as inappropriate and capture 95% of

Precision | Recall

Anomaly Detection for Wearables Framework

sensors

MDPI

Review

Anomaly Detection Framework for Wearables Data: A Perspective Review on Data Concepts, Data Analysis Algorithms and Prospects

Jithin S. Sunny [1,†], C. Pawan K. Patro [2,*,†], Khushi Karnani [1], Sandeep C. Pingle [2], Feng Lin [2], Misa Anekoji [2], Lawrence D. Jones [2], Santosh Kesari [3,4] and Shashaanka Ashili [2]

[1] Rhenix Lifesciences, Hyderabad 500038, India; jithin@rhenix.org (J.S.S.); khushi@rhenix.org (K.K.)
[2] CureScience, San Diego, CA 92121, USA; spingle@curescience.org (S.C.P.); flin@curescience.org (F.L.); manekoji@curescience.org (M.A.); ljones@curescience.org (L.D.J.); shashi@curescience.org (S.A.)
[3] Pacific Neuroscience Institute, Providence Saint John's Health Center, Santa Monica, CA 90404, USA; santoshkesari@gmail.com
[4] Department of Translational Neurosciences, Saint John's Cancer Institute at Providence Saint John's Health Center, Santa Monica, CA 90404, USA
* Correspondence: info@curescience.org
† These authors contributed equally to this work.

Citation: Sunny, J.S.; Patro, C.P.K.; Karnani, K.; Pingle, S.C.; Lin, F.; Anekoji, M.; Jones, L.D.; Kesari, S.; Ashili, S. Anomaly Detection Framework for Wearables Data: A Perspective Review on Data Concepts, Data Analysis Algorithms and Prospects. *Sensors* **2022**, *22*, 756. https://doi.org/10.3390/s22030756

Abstract: Wearable devices use sensors to evaluate physiological parameters, such as the heart rate, pulse rate, number of steps taken, body fat and diet. The continuous monitoring of physiological parameters offers a potential solution to assess personal healthcare. Identifying outliers or anomalies in heart rates and other features can help identify patterns that can play a significant role in understanding the underlying cause of disease states. Since anomalies are present within the vast amount of data generated by wearable device sensors, identifying anomalies requires accurate automated techniques. Given the clinical significance of anomalies and their impact on diagnosis and treatment, a wide range of detection methods have been proposed to detect anomalies. Much of what is reported herein is based on previously published literature. Clinical studies employing wearable devices are also increasing. In this article, we review the nature of the wearables-associated data and the downstream processing methods for detecting anomalies. In addition, we also review supervised and un-supervised techniques as well as semi-supervised methods that overcome the challenges of missing and un-annotated healthcare data.

Keywords: anomaly detection; heart rate; wearables; missing data; machine learning

Anomaly Detection for Wearables Framework

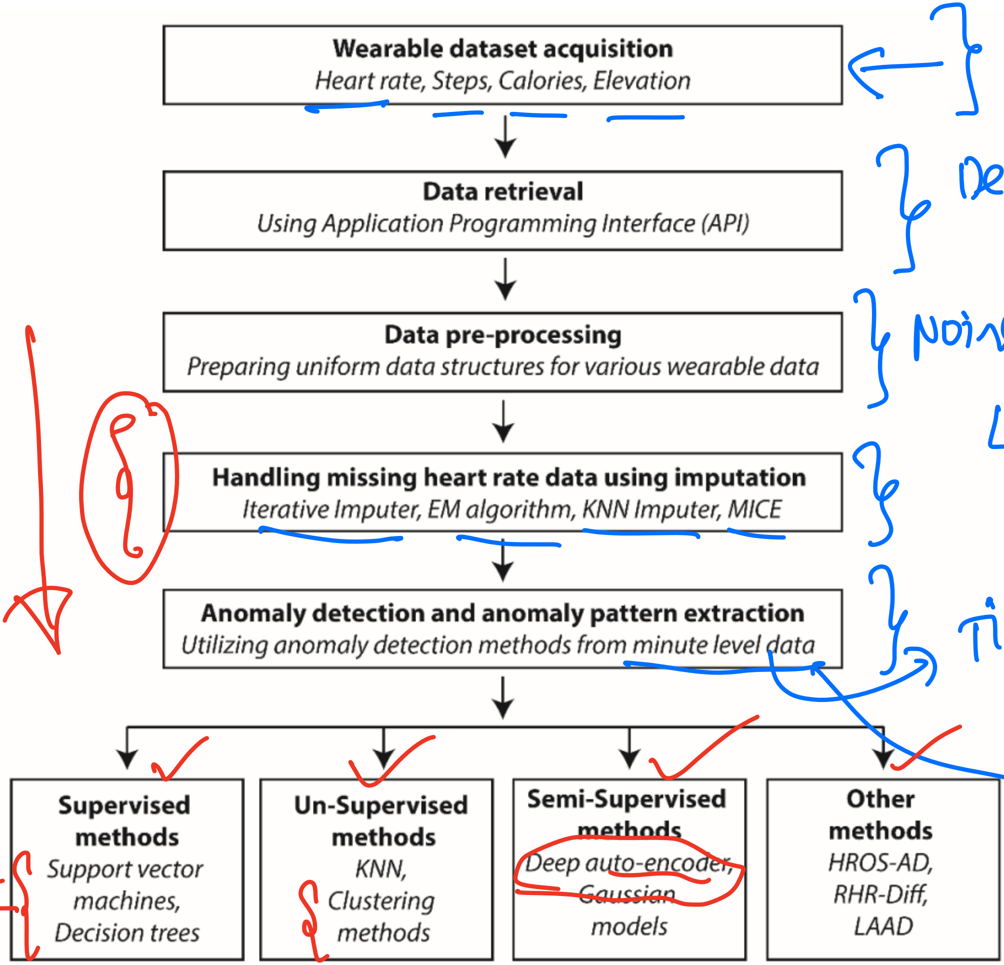

Apple watch

Device → cloud (csv)

Noise in data

Missing Data

Time bucket that makes sense

Granularity

Classification

Imputation Methods for wearable analytics

Methods	Definition	Accuracy	References
Mean value imputation (MVI)	The values are filled using calculating the mean for a missing value	Biased	[42]
Maximum Likelihood (ML)	A likelihood function is evaluated and then sum or integrate over the missing data	Unbiased parameter estimation	[43]
Hot Deck Imputation	A data matrix for all instances created is chosen as a source for missing values	Replication of values may cause bias	[44]
Multiple Imputation (MI)	Starts by introducing random variation and generates several datasets with slightly different imputed values. Statistical analysis on each to find the optimal one	Comparable to ML	[45]
Multivariate Imputation by Chained Equations (MICE)	The method first identifies an imputation model for each column followed by random draws from the observable data	Comparable to ML	[46]
Expectation–Maximization with Bootstrapping (EMB)	Initially the likelihood function is evaluated using model parameters. Next, with the updated parameters, the likelihood function is maximized, and the parameters are updated to return a new distribution	Comparable to ML	[47]

Imputation Methods: Mean Imputation

Mean Imputation

Replace missing value at a (row,col) index in a data table by the mean of the column values that the index belongs to.

ICE #1

Median Imputation

Consider the following data table

x_1	x_2	x_3	x_4	x_5	x_6
1	2	3	1	4	5
1	5	5	4	2	3
1	2	5	3	1	2
2	3		4	3	

What would median imputation output for the missing values in row 4?

a. 3 5
b. 5 3
c. 5 5
d. 3 3

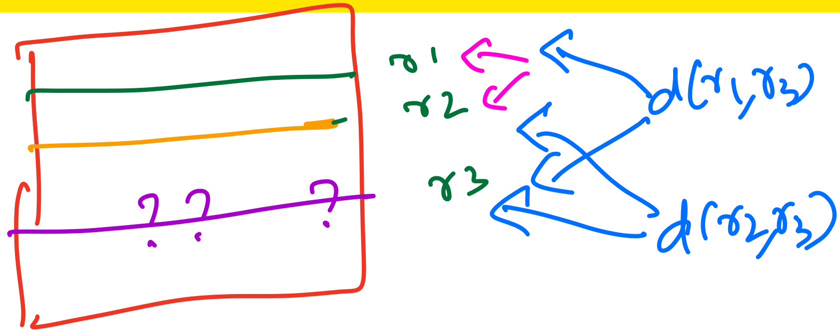

Hot Deck Imputation

For any given row that needs imputation, find a **donor** row that can help with imputation. Donor rows can be chosen based on the nearest distance to the row that needs imputation. The nearest distance is computed on the values that are available in both the donor row and the row needing imputation.

ICE #2

Hot Deck Imputation

Consider the following data table

x_1	x_2	x_3	x_4	x_5	x_6
1	2	3	1	4	5
1	5	5	4	2	3
1	2	5	3	1	2
2	3		4	3	

Find the donor row index for the row with missing values based on ℓ_1 distance.

$$\|x_1 - x_2\|_1 = \sum_i |x_{1i} - x_{2i}|$$

- (a) Row 1
- (b) Row 2
- (c) Row 3
- (d) None of the above

Non-negative Matrix Factorization } multiple imputation

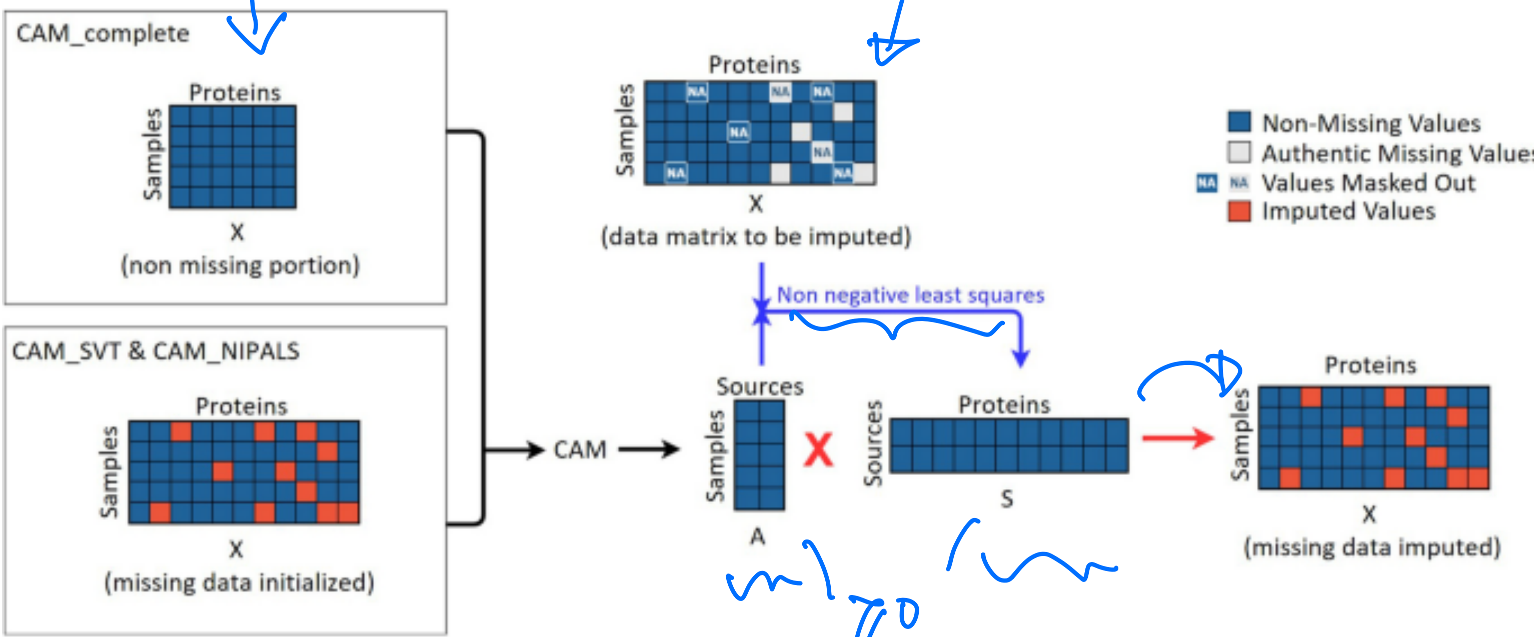

$X = AS$

X

$$\min_{\substack{A,S\\ A\geq 0\\ S\geq 0}} \sum_{(i,j)\in \underline{\underline{\Omega}}} \left(X_{ij} - A_{i,\cdot}^{T} S_{\cdot j}\right)^2$$

Ω — Indices for which data is observed → Blue cells above!

→ Gradient Descent

Non-negative Least Squares

Non-negative MF (NMF)

NMF vs SVD

Singular value De-composition

Why not SVD?

Why not use SVD for data imputations?

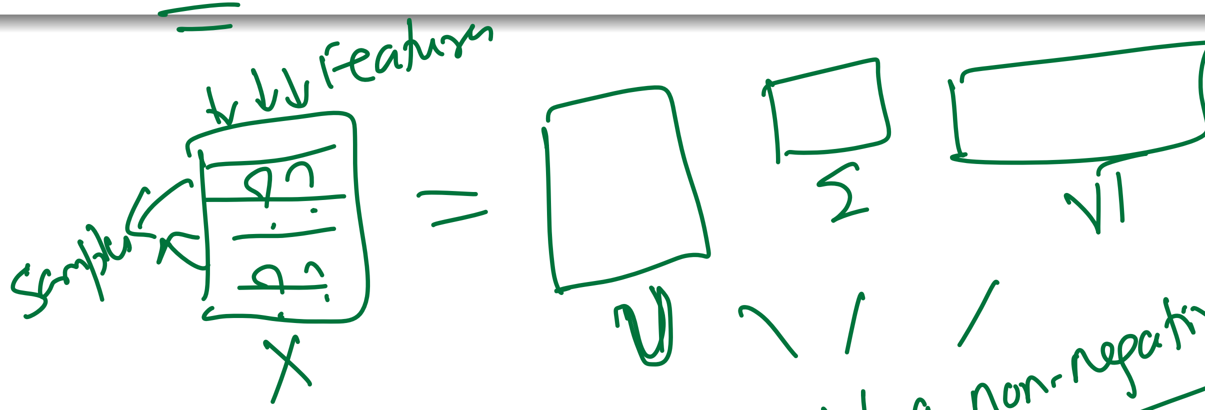

Not a non-negative MF!

Non-negative Matrix Factorization

Gene-Expression Datasets

Psycho-Neuro Immunology

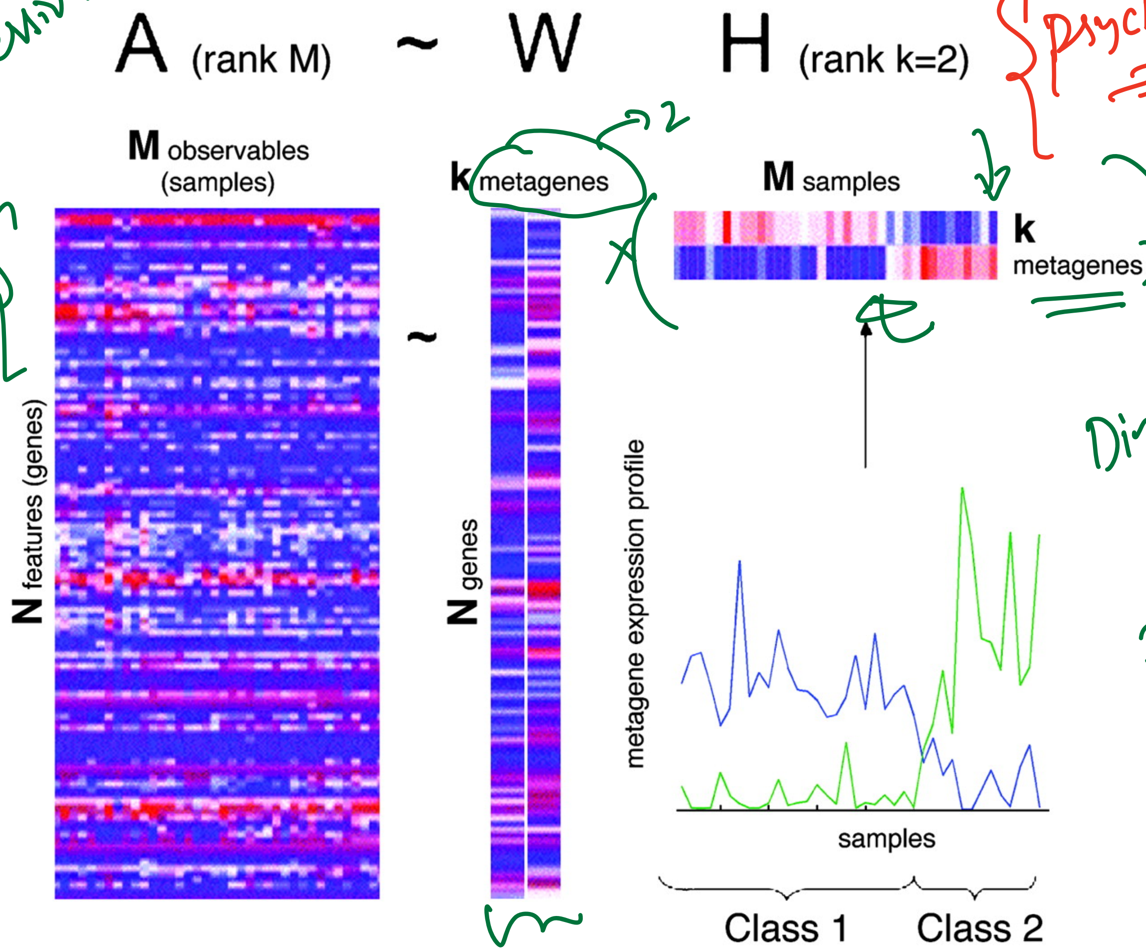

2

20k genes

Dimensionality Reduction & Better Interpretation!

Non-negative Matrix Factorization

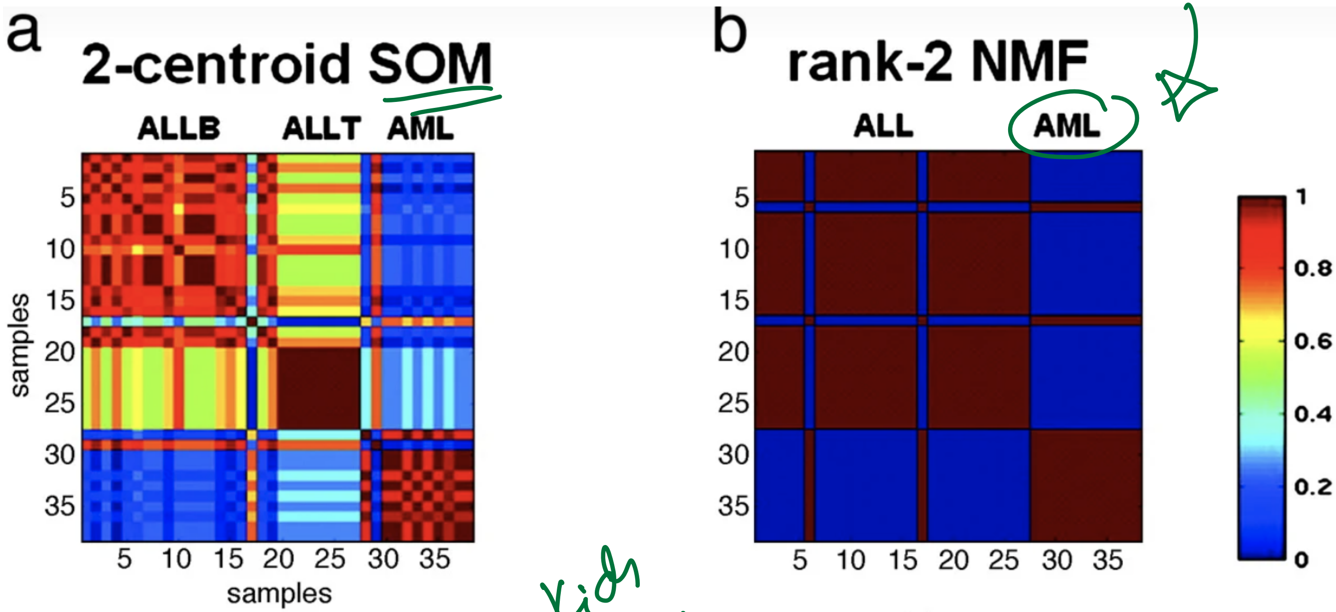

Leukaemia — ALL → Kids
AML → Adults

Anomaly Detection for Wearables Framework

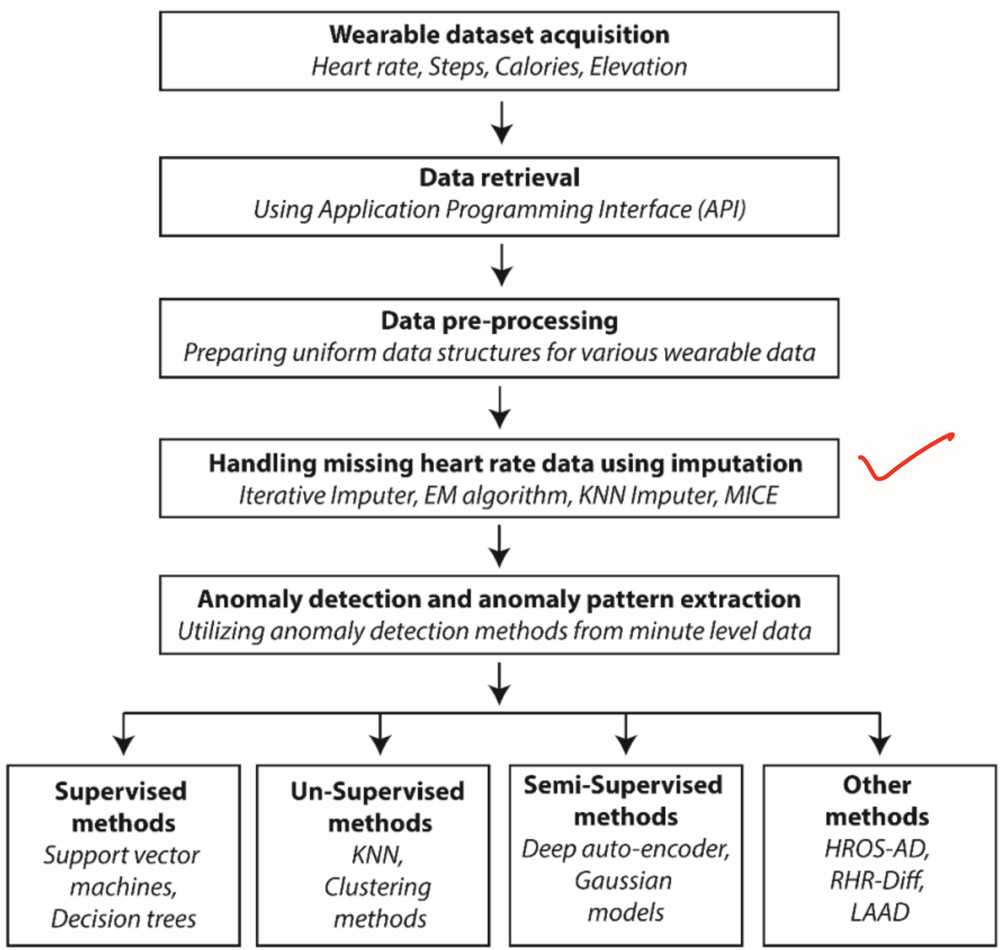

Anomaly Detection: Types of Anomalies

1. **Point Anomaly:** Deviation from a set of data points.

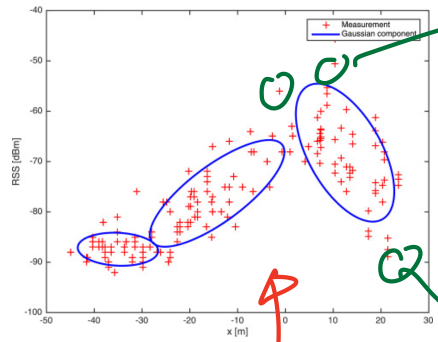

Pt. Anomaly

K=3 Cluster

k-means

Point Anomaly

Anomaly Detection: Types of Anomalies

1. **Point Anomaly:** Deviation from a set of data points.

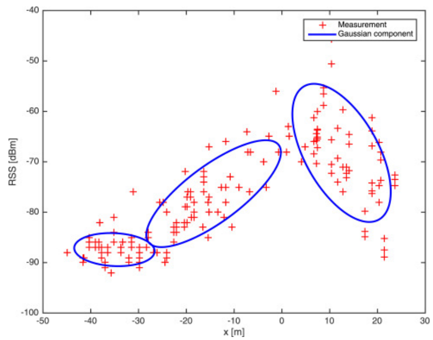

2. **Contextual Anomaly:** Depending on the context, a data point could be an anomaly or not. For instance 35 degrees is not an anomalous temperature for Seattle winter but it is for Seattle summer. Same is true for anomalies in a time-series data e.g. if a person has an average o2 level of 90. Both 100 and 80 look like anomalies, but only lower o2 is a cause for concern (Directional anomalies).

Direction as a context

Tyeps of Anomalies

3. **Collective Anomalies:** No one data point is anomalous but a collection of them become anomalous. E.g. the Arrhythmia time series.

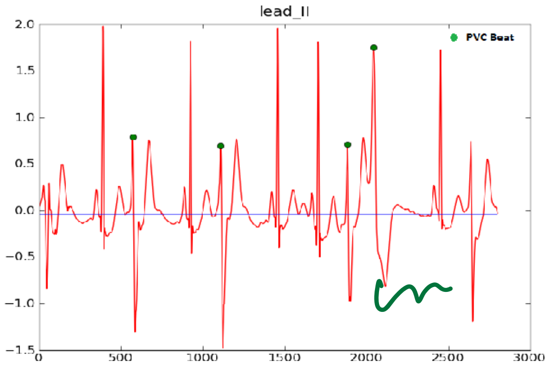

ICE #3

O2

Anomaly Type

You are tasked with detecting if spo2 is trending downwards. So far its been trending downwards. You need to identify the trend by end of day. This is an example of:

a. Point anomaly detection
b. Contextual anomaly detection
c. Collective anomaly detection
d. Not an anomaly detection problem

ICE #4

Anomaly Type

You are tasked with identify which genes are the mutant ones in gene-expression data of patients that are ~~both~~ either healthy ~~and~~ or have cancer. Identifying these mutant genes is an example of:

- a. Point anomaly detection
- b. Contextual anomaly detection
- c. Collective anomaly detection
- d. Not an anomaly detection problem

Time-series Anomaly Detection

Local window anomalies

Fit a linear model that captures local trends and compute a probability for a new data point being an anomaly w.r.t local model.

Time-series Anomaly Detection

Un-supervised Learning

If we don't have enough labels for anomalies (or positive class), we have no choice but to resort to **un-supervised learning**.

Time-series Anomaly Detection

Un-supervised Learning

If we don't have enough labels for anomalies (or positive class), we have no choice but to resort to **un-supervised learning**.

Un-supervised Learning

However, un-supervised learning for anomaly detection is fraught with issues. What are they?

{ Precision | Recall }

Too many FP!

Semi-supervised Learning

Learn good features from un-supervised learning and use a simple classifier - such as logistic regression model to fine tune the probability computations! **Example features:** Deviation from a local linear regression model fit on a local window. Deviation from median of a local window.

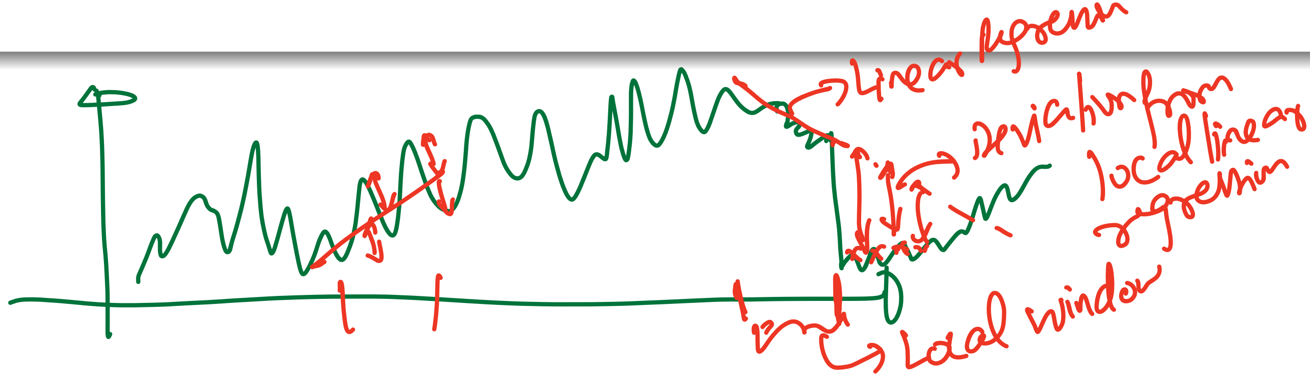

Time-series Anomaly Detection

ICE #5

What are the hyper-parameters for the semi-supervised logistic regression based anomaly detection approach we just described?

a. The weights for the different features learned from un-supervised learning that are then combined to get a probability prediction from logistic regression
b. The number of (unsupervised learning) features used in the logistic regression
c. The size of the local window used to compute these features
d. The probability of a data point being an anomaly

Deep Learning for Anomaly Detection

Next Lecture

Deep Learning based methods for EEG analysis, Arrythmia and detection of other conditions.

Your choice of Wearable - Personalized Analytics Mid-Term Presentation

- Pick a wearable from the market that does tracking of HR, glucose or BP. Ensure you can get end of day csv files or equivalent data from wearable

Your choice of Wearable - Personalized Analytics Mid-Term Presentation

- Pick a wearable from the market that does tracking of HR, glucose or BP. Ensure you can get end of day csv files or equivalent data from wearable
- Do change point detection and anomaly detection on your own data

Your choice of Wearable - Personalized Analytics Mid-Term Presentation

- Pick a wearable from the market that does tracking of HR, glucose or BP. Ensure you can get end of day csv files or equivalent data from wearable
- Do change point detection and anomaly detection on your own data
- Some frequency domain analysis as well and record your insights

Your choice of Wearable - Personalized Analytics Mid-Term Presentation

- Pick a wearable from the market that does tracking of HR, glucose or BP. Ensure you can get end of day csv files or equivalent data from wearable
- Do change point detection and anomaly detection on your own data
- Some frequency domain analysis as well and record your insights
- Present in teams of 2 maybe end of April. How does that sound?

References for this lecture

1. Anomaly Detection Framework for Wearables Data: A Perspective Review on Data Concepts, Data Analysis, Algorithms and Prospects
2. Healthcare and Anomaly Detection: Using Machine Learning to predict Anomalies in Heart Rate Data
3. Review of Hot Deck Imputation