EEP 596: AI and Health Care || Lecture 6

Dr. Karthik Mohan

Univ. of Washington, Seattle

Apr 14, 2022

Logistics

- Ayush Quiz Section: Friday, 6-7 pm pst

New Time

Pytorch for Deep Learning

Structured content
- Recap of concepts
- Walk through Libraries (for DL)
- open-ended discussion

Logistics

- Ayush Quiz Section: Friday, 6-7 pm pst
- Mathew Grading OH: Saturday, 5-6 pm pst → end of week/Sunday

→ Related to Assignment Grading/ Rubrics

Logistics

- Ayush Quiz Section: Friday, 6-7 pm pst
- Mathew Grading OH: Saturday, 5-6 pm pst
- Ayush OH: Sunday 12-1 pm pst

× Karthik OH: on wed after class (today after class)
(wednesday)

Logistics

- Ayush Quiz Section: Friday, 6-7 pm pst
- Mathew Grading OH: Saturday, 5-6 pm pst
- Ayush OH: Sunday 12-1 pm pst
- Pytorch library for DL covered in quiz section this week with examples

Logistics

- Ayush Quiz Section: Friday, 6-7 pm pst
- Mathew Grading OH: Saturday, 5-6 pm pst
- Ayush OH: Sunday 12-1 pm pst
- Pytorch library for DL covered in quiz section this week with examples
- Reference papers/material for this lecture (and future lectures) on last slide of this deck

Logistics

- Ayush Quiz Section: Friday, 6-7 pm pst
- Mathew Grading OH: Saturday, 5-6 pm pst
- Ayush OH: Sunday 12-1 pm pst
- Pytorch library for DL covered in quiz section this week with examples
- Reference papers/material for this lecture (and future lectures) on last slide of this deck
- Anything from your side?

→ If you recently joined → Join Discord for updates

Last Lecture

- **Use cases for anomaly detection**

Last Lecture

- **Use cases for anomaly detection**
- **Anomaly Detection Pipeline for Health Care**

Last Lecture

- **Use cases for anomaly detection**
- **Anomaly Detection Pipeline for Health Care**
- **Imputation Methods**

Last Lecture

- **Use cases for anomaly detection**
- **Anomaly Detection Pipeline for Health Care**
- **Imputation Methods**
- **Anomaly Detection Baselines**

Today

- **Wearables and Data Access - Mathew**

Today

- **Wearables and Data Access - Mathew**
- **Sleep and Relaxation case study**] (Karthik)

Today

- **Wearables and Data Access - Mathew**
- **Sleep and Relaxation case study**
- **Introduction to Deep Learning for Anomaly Detection**

Self Study: HR variations during Accupressure/(Marma)

Apple Watch Data

Premium

Using the Cardiogram app to do continuous HR recording (5 second granularity) and graph.

csv

TS	HR
9:44:03	65
9:44:07	60
⋮	⋮

Apple watch

Self Study: HR variations during Accupressure/Marma

Start of Accupressure (Marma)

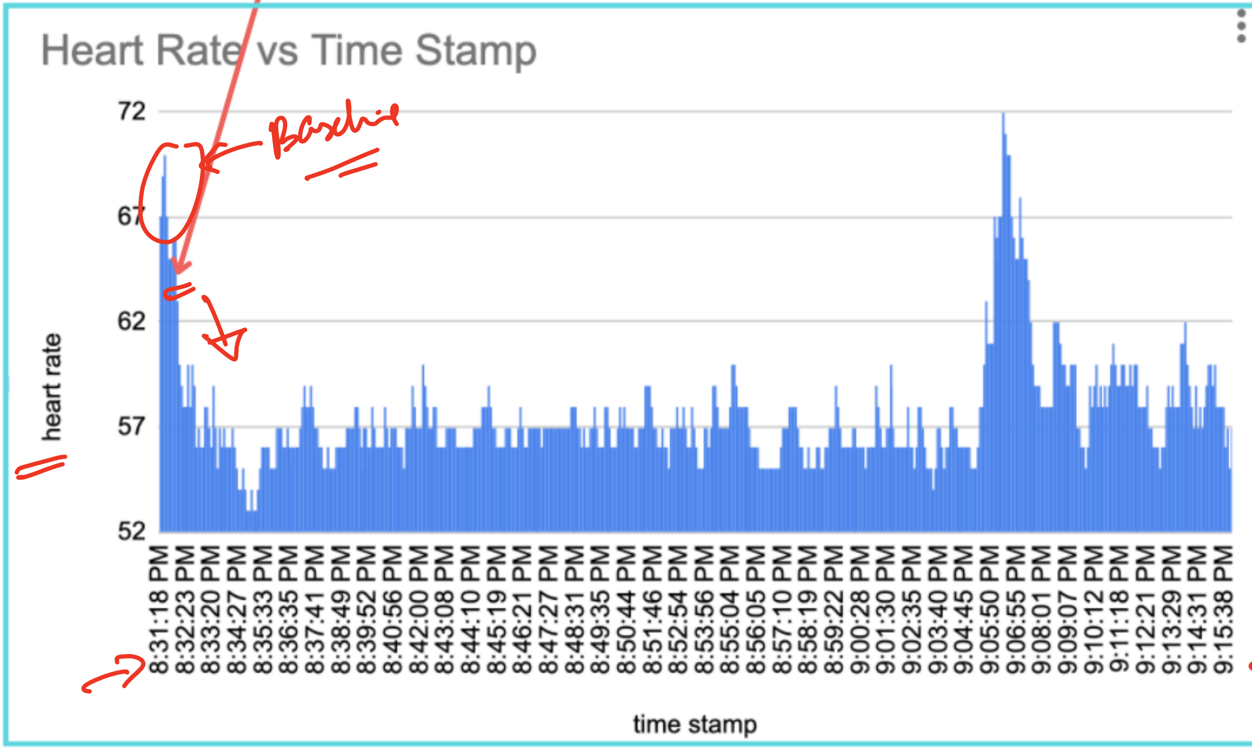

Self Study: HR variations during Accupressure/Marma

(India)

Start of Accupressure (Marma)

4 mins in - HR at 53

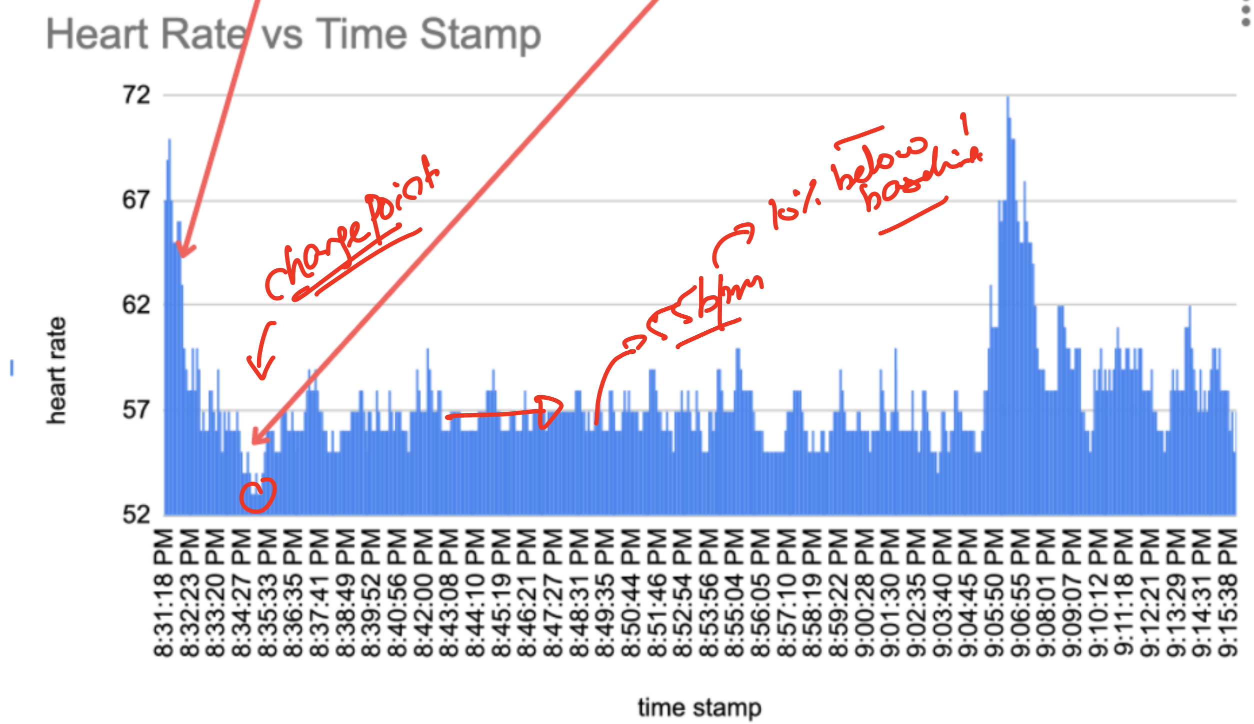

Self Study: HR variations during Accupressure/Marma

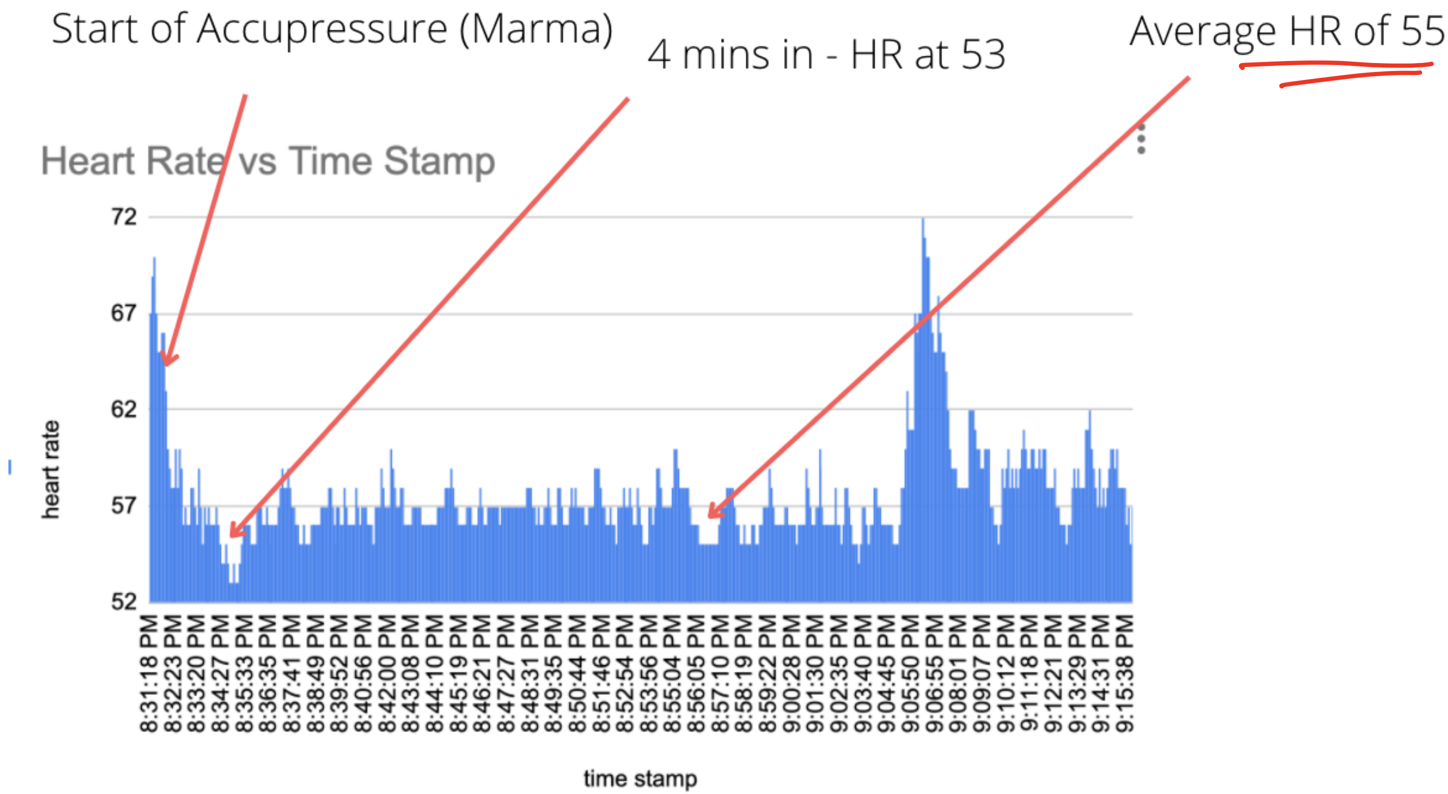

Self Study: HR variations during Accupressure/Marma

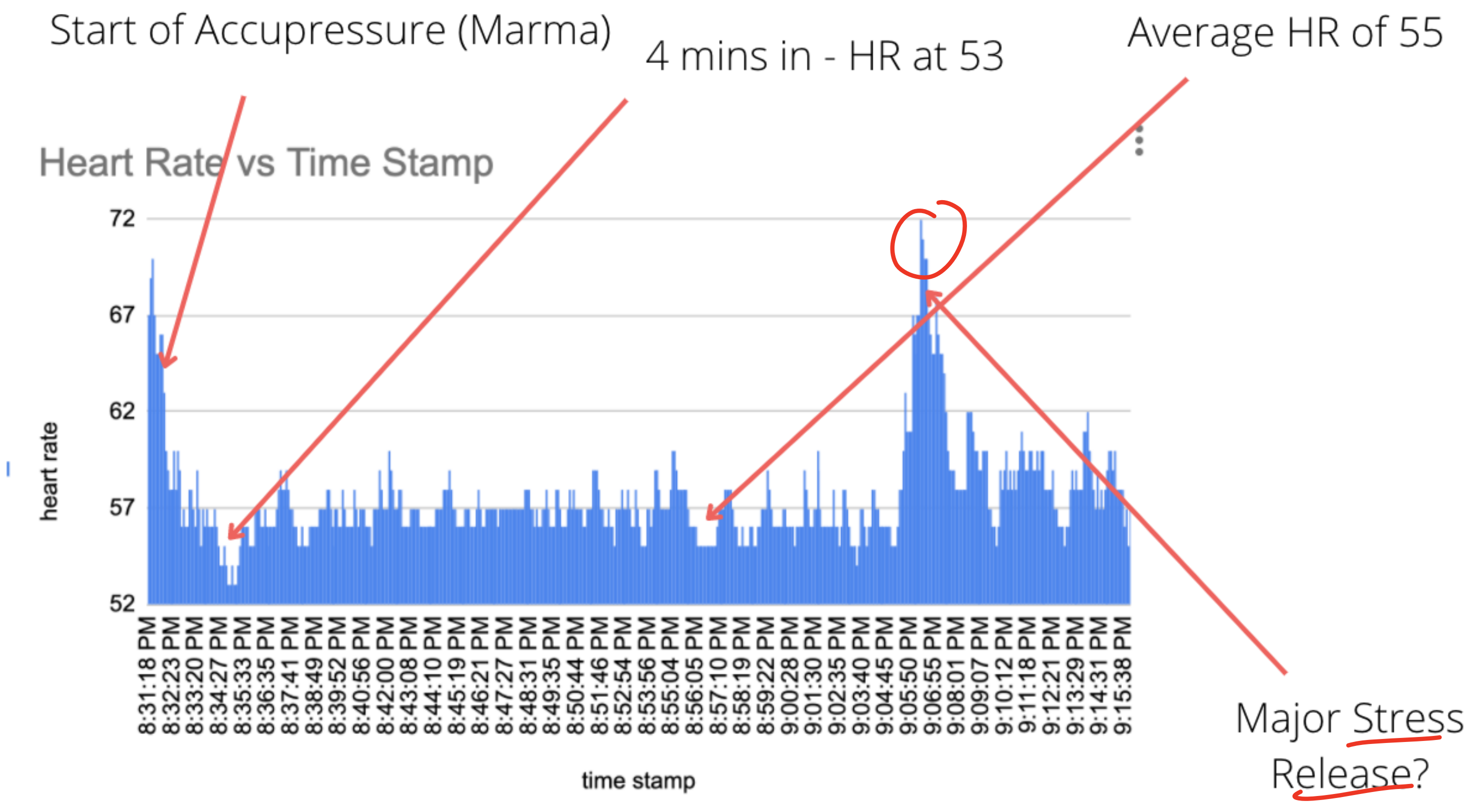

Self Study: HR variations during Accupressure/Marma

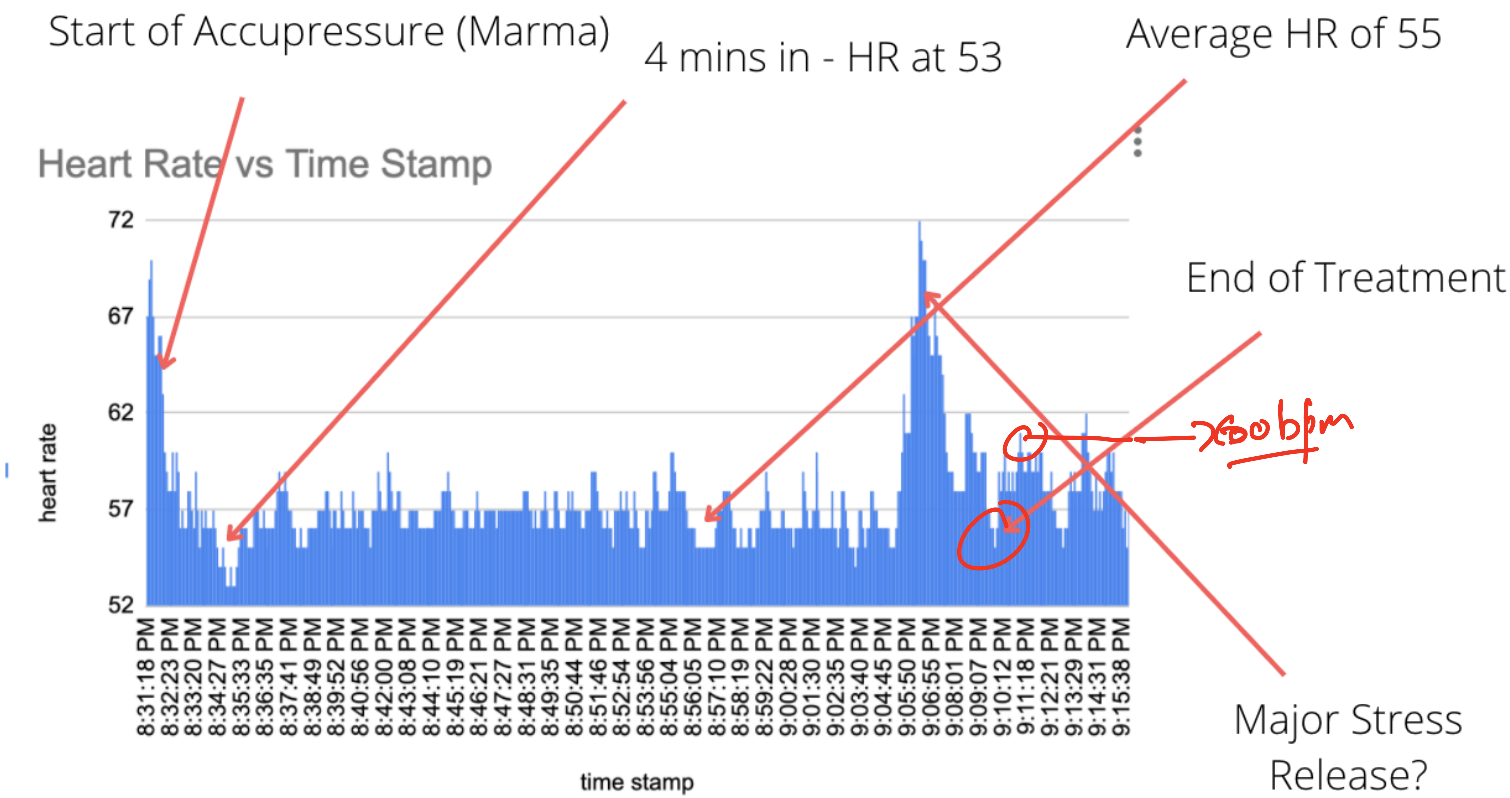

Self Study: HR variations during Accupressure/Marma

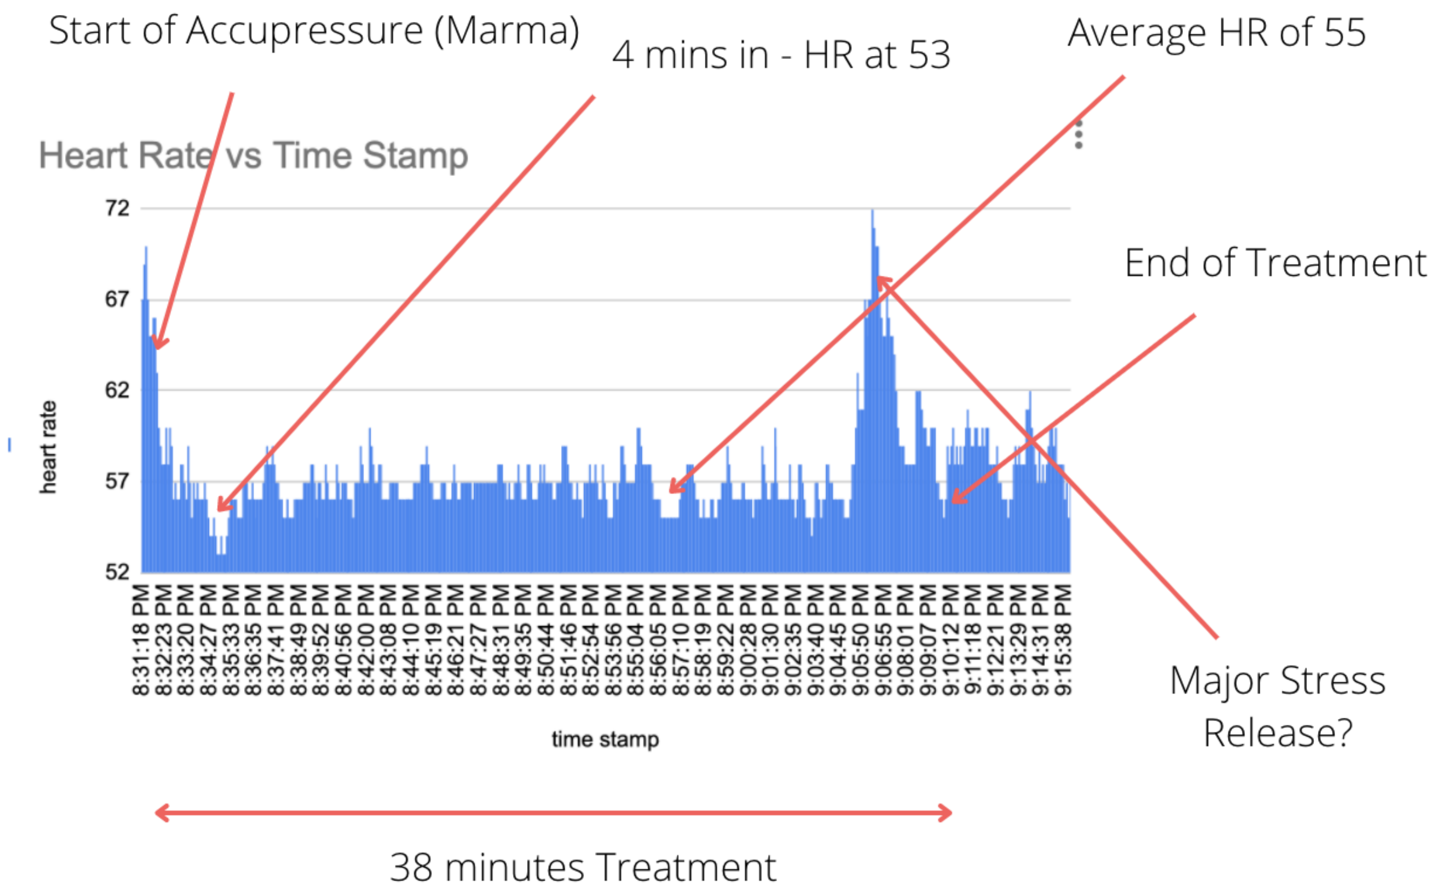

Sleep Phases

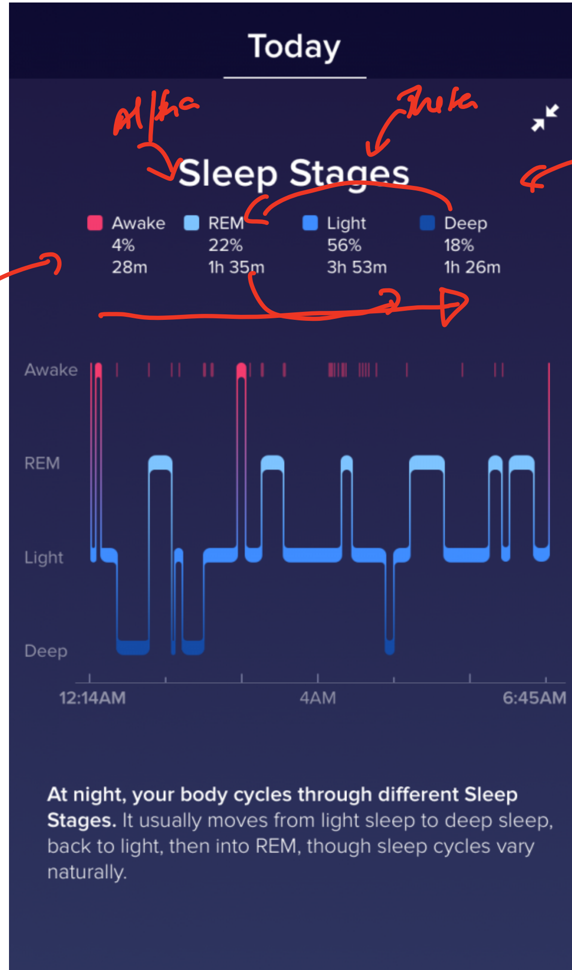

Sleep Phases

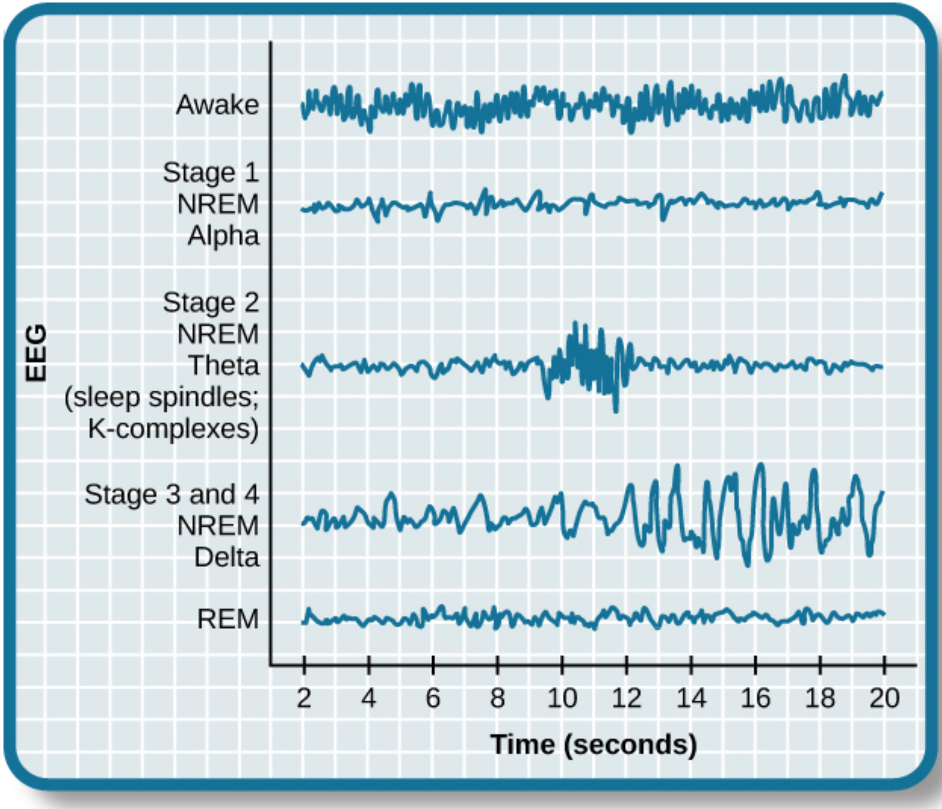

NREM 1

Light sleep 2

Deep sleep

Self Study: Sleep Cycle detection through HR

Me
contemplating
sleep

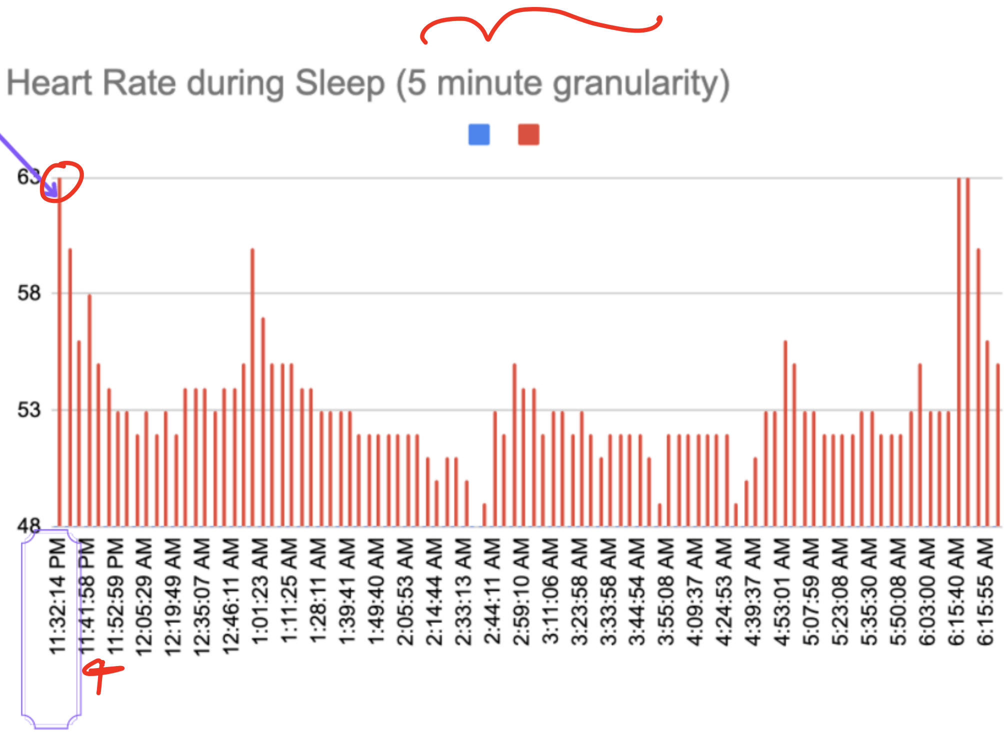

Self Study: Sleep Cycle detection through HR

Me contemplating sleep

Hitting the bed

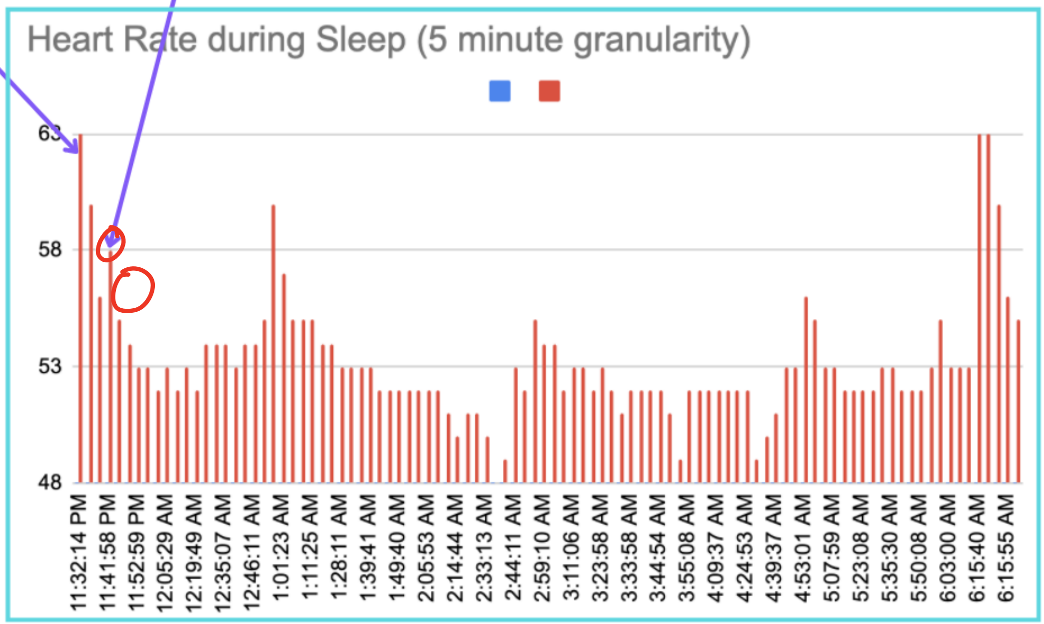

Self Study: Sleep Cycle detection through HR

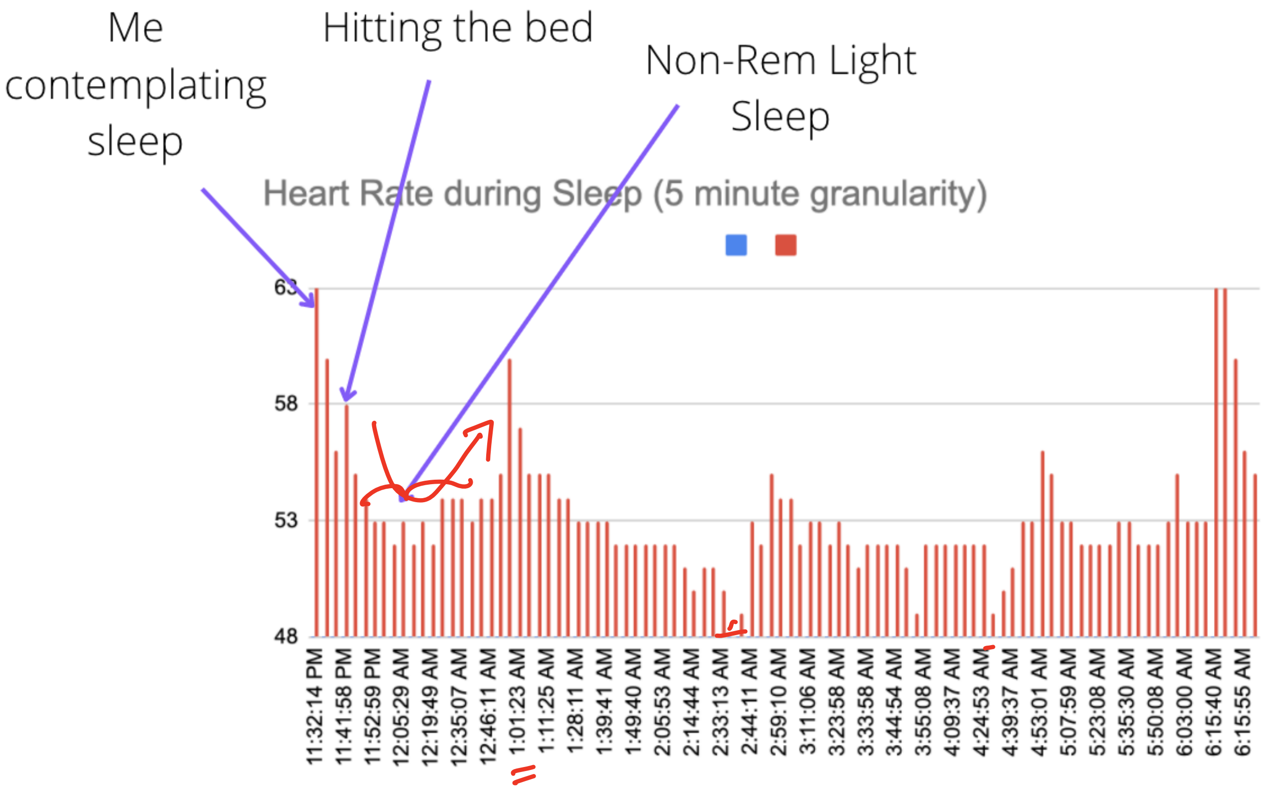

Self Study: Sleep Cycle detection through HR

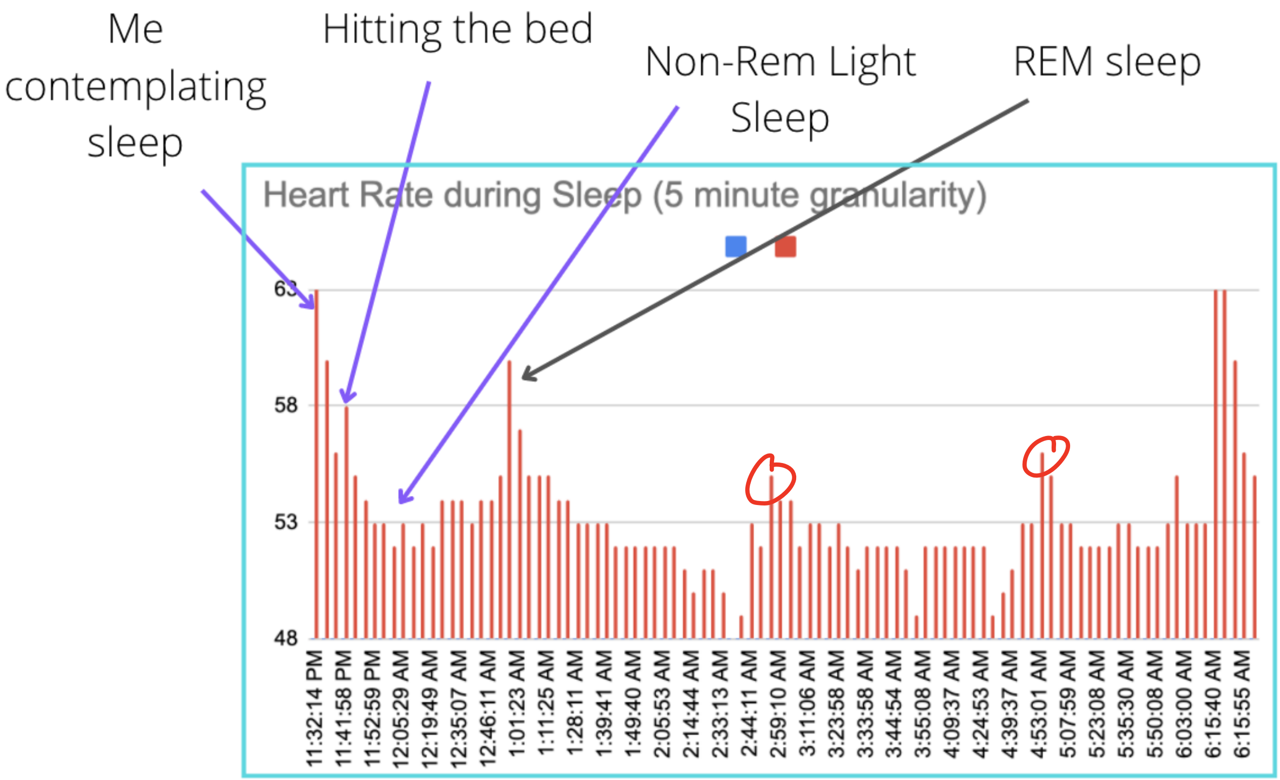

Self Study: Sleep Cycle detection through HR

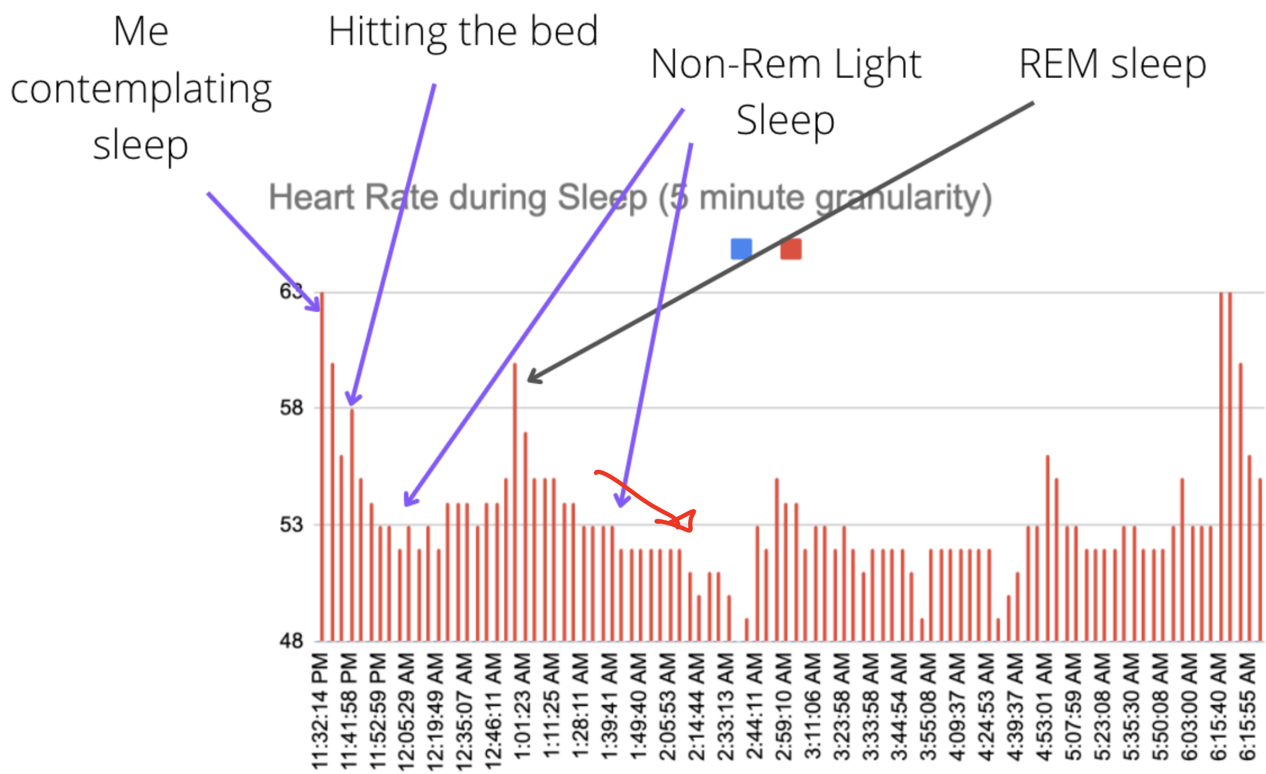

Self Study: Sleep Cycle detection through HR

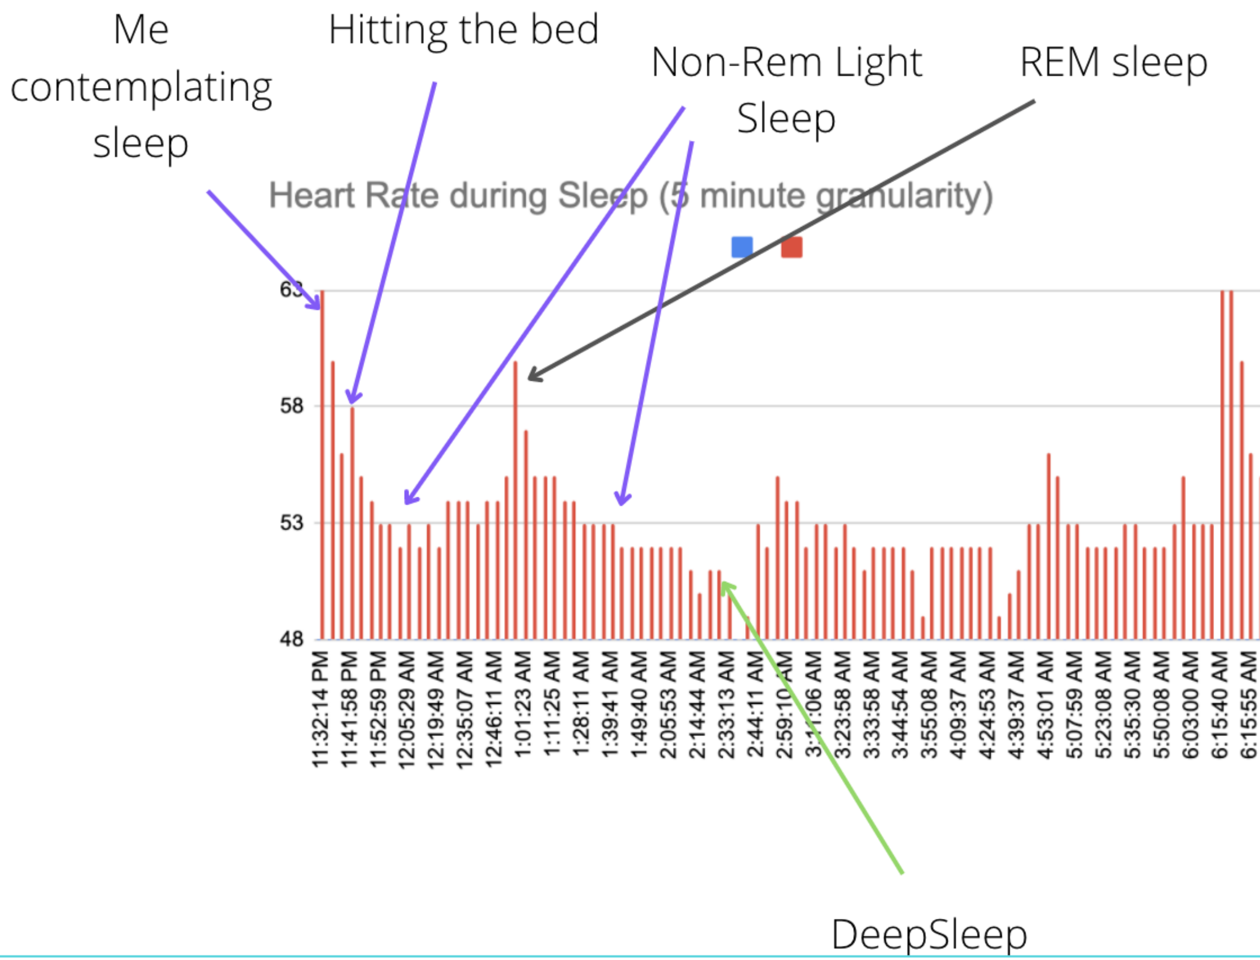

Self Study: Sleep Cycle detection through HR

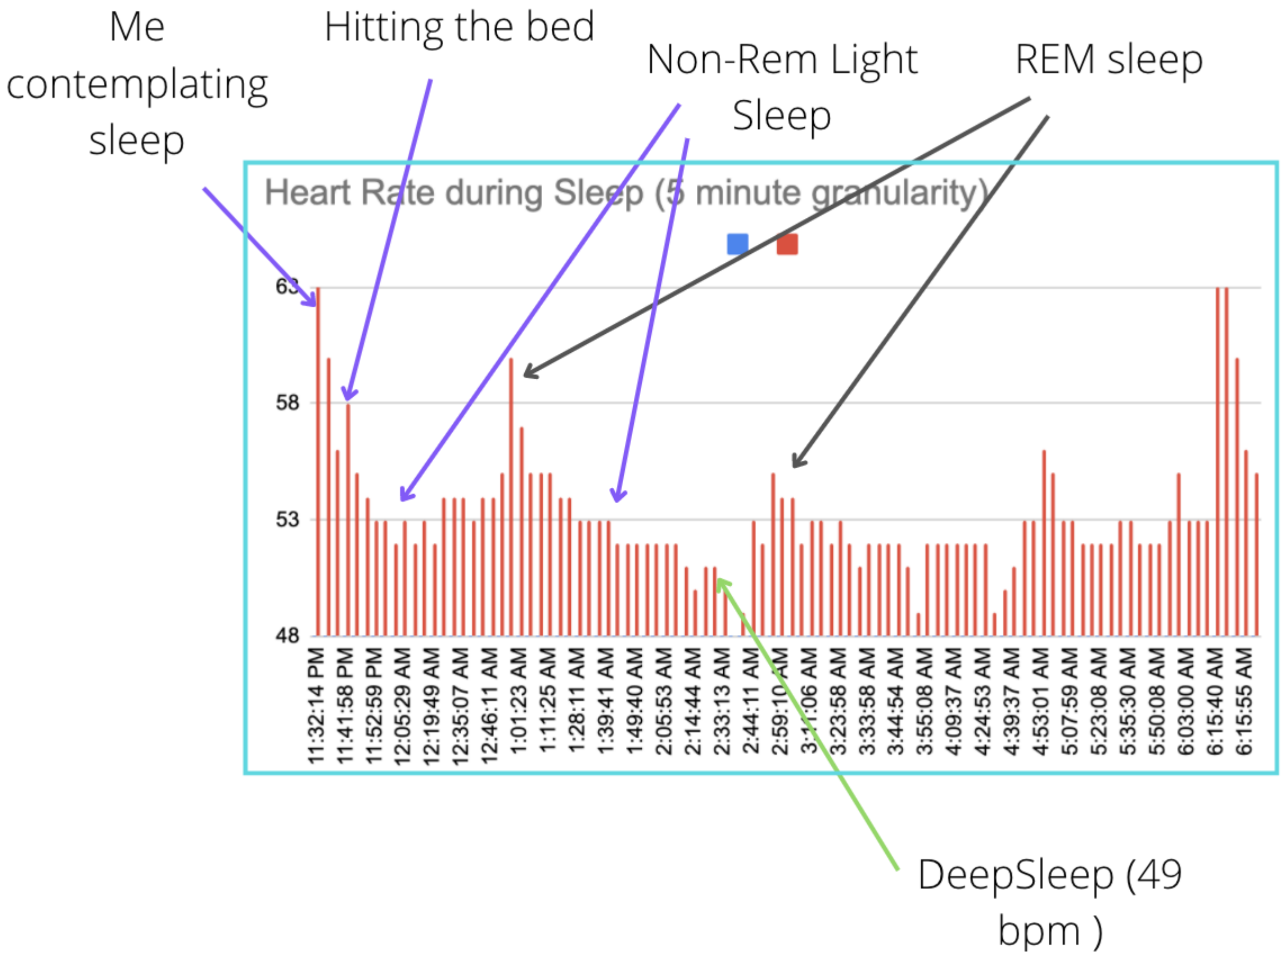

Self Study: Sleep Cycle detection through HR

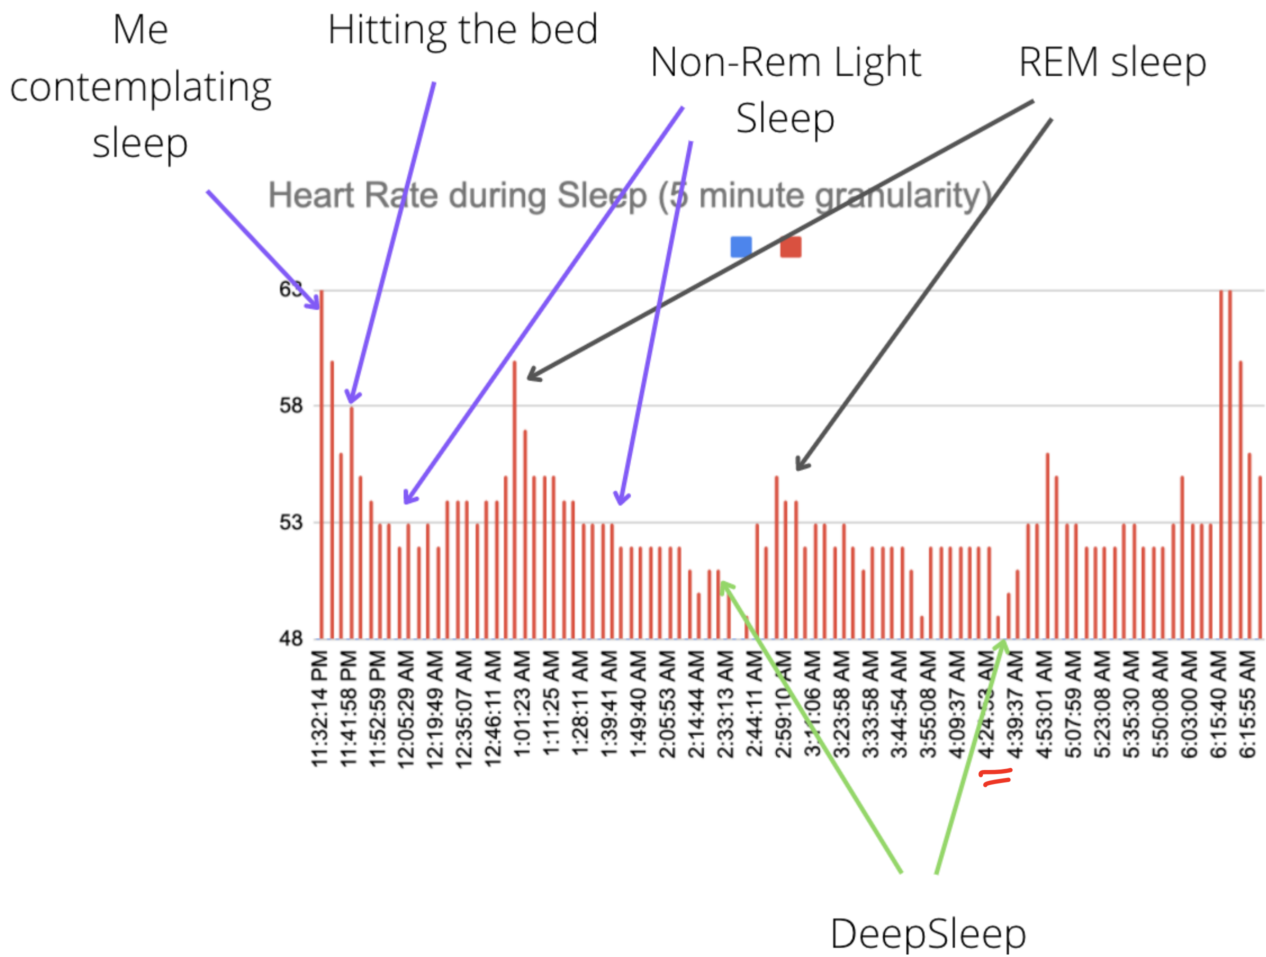

Self Study: Sleep Cycle detection through HR

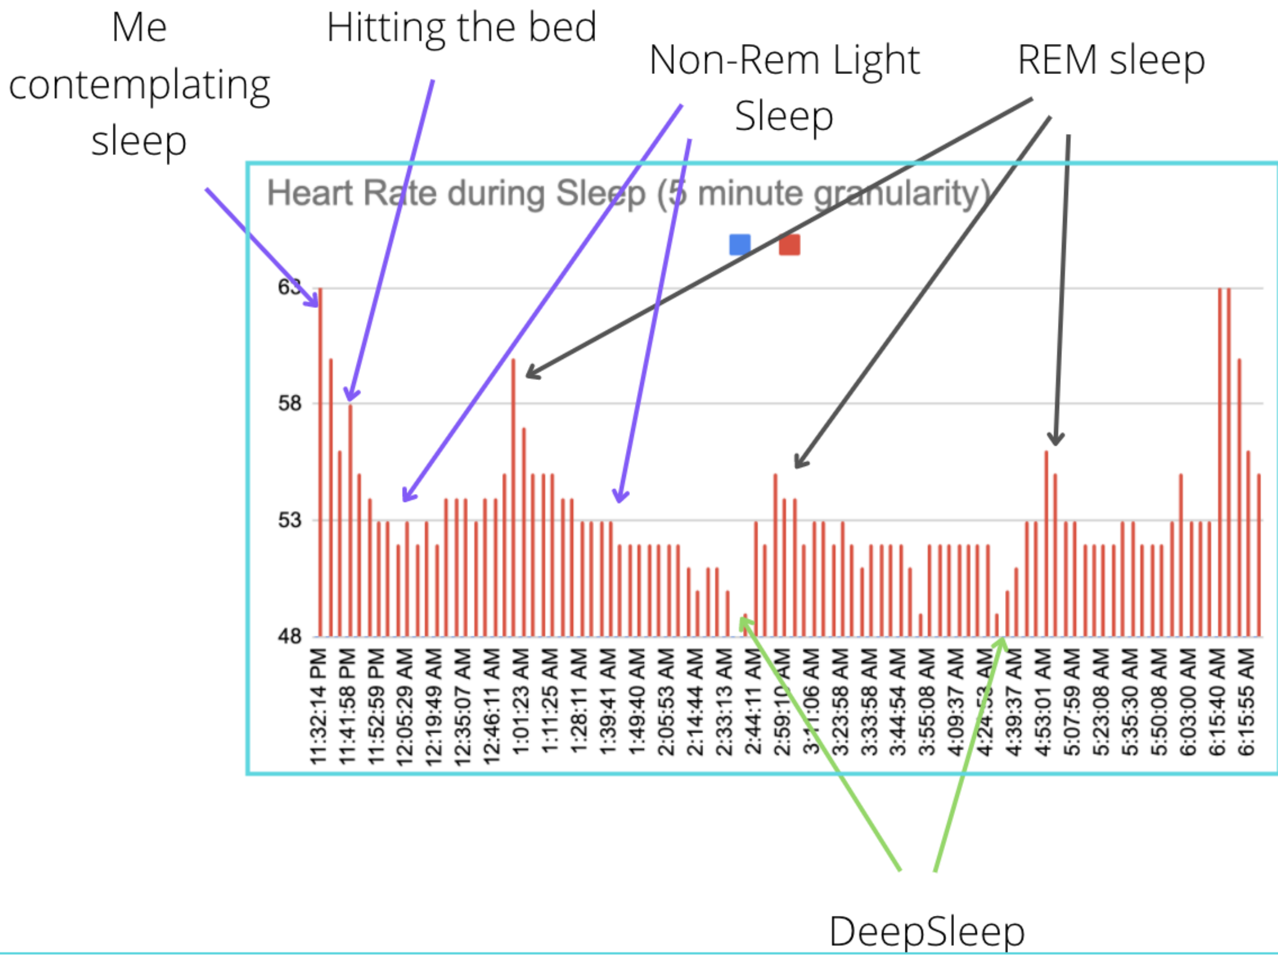

Self Study: Sleep Cycle detection through HR

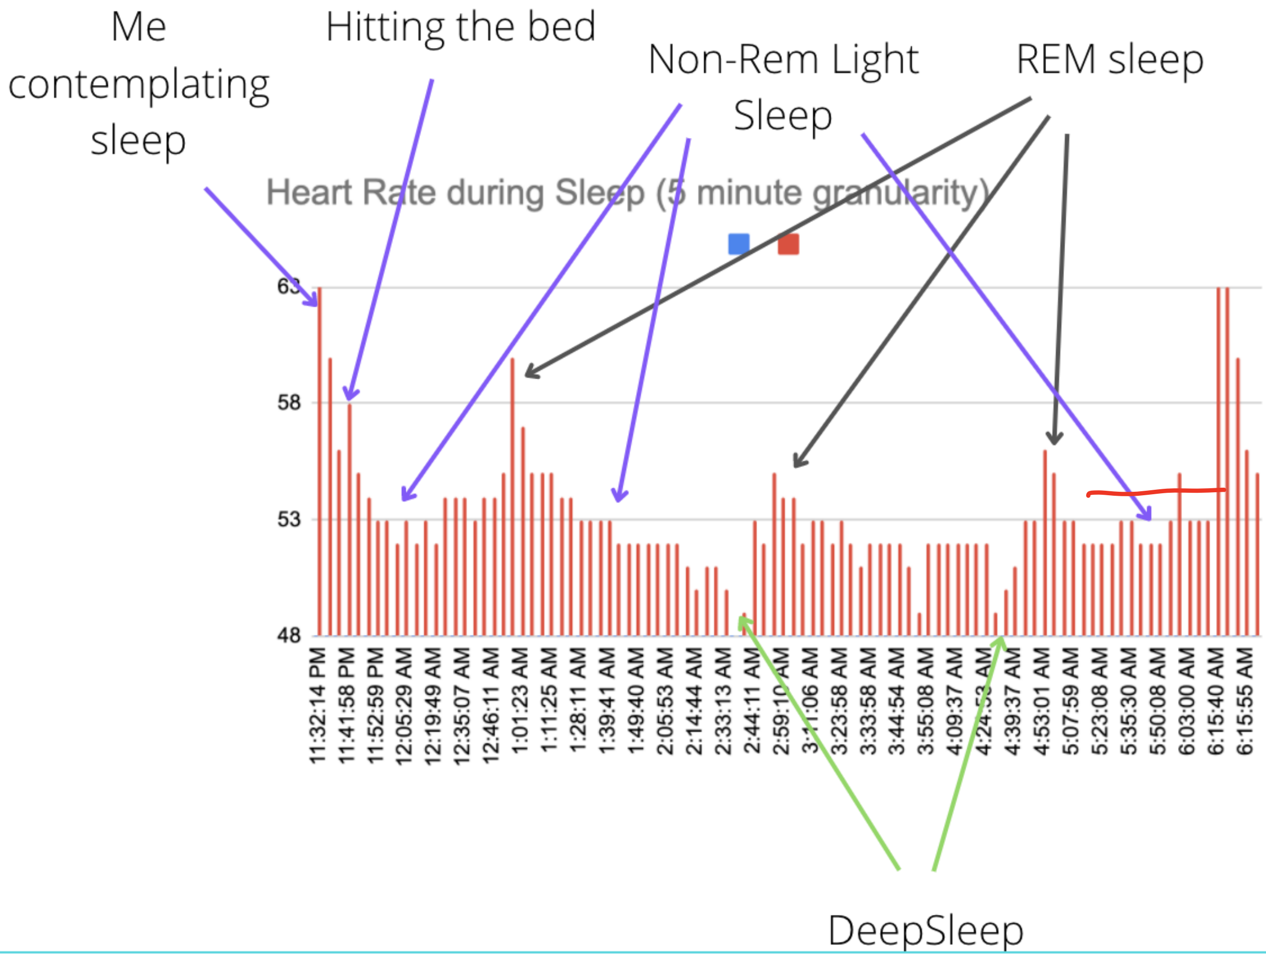

Self Study: Sleep Cycle detection through HR

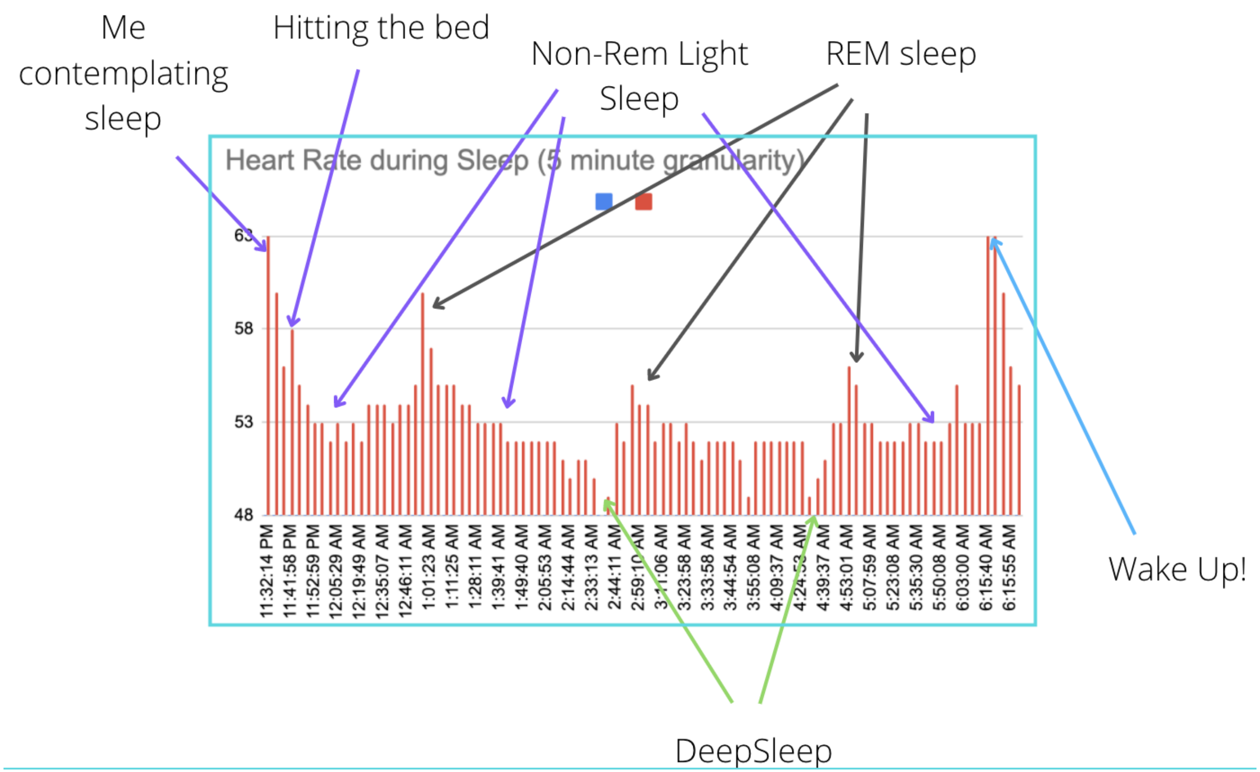

Self Study: Sleep Cycle detection through HR

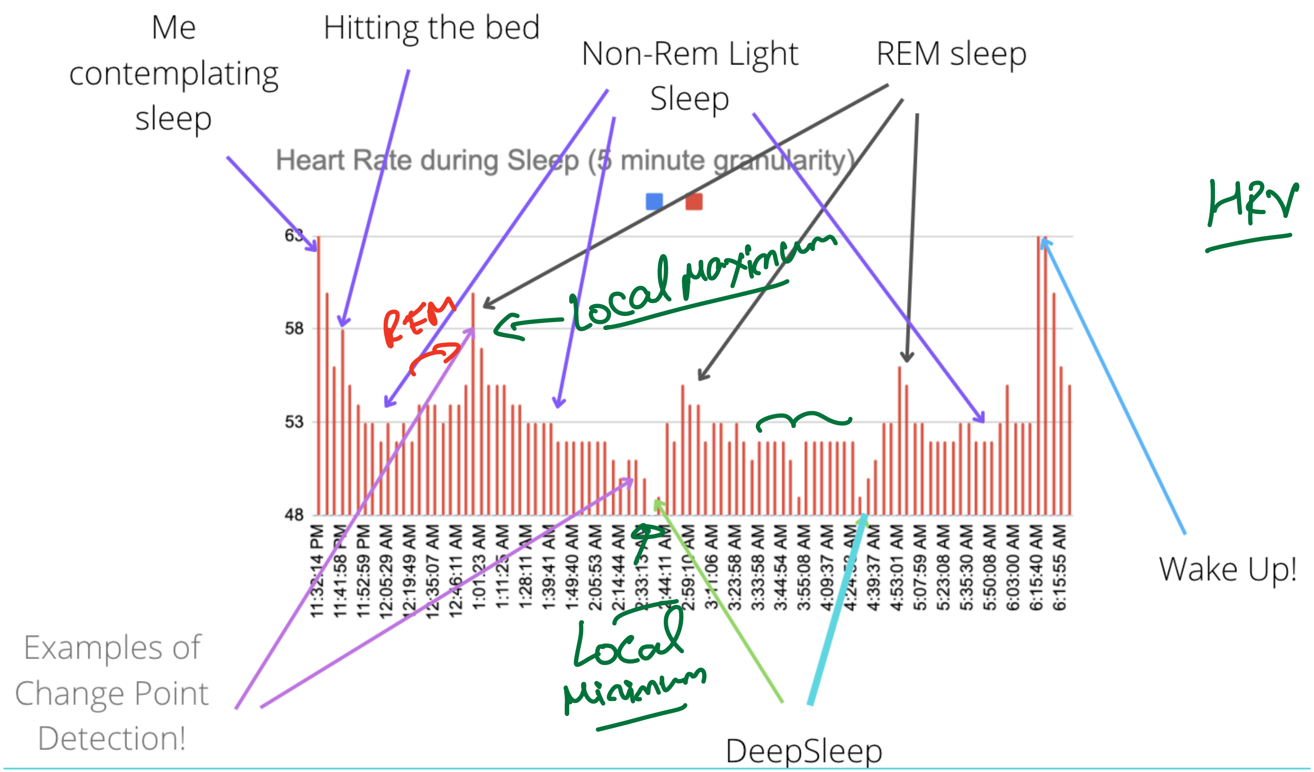

Anomaly Detection in IoT context

Properties of a good Anomaly Detector for IoT data streams

1. **Speed:** Ability to handle data coming in rapidly. In the case of heart rate monitoring and Arrhythmia Detection, data coming in every 5 seconds (e.g. Cardiogram App for Apple Watch).

Anomaly Detection in IoT context

Properties of a good Anomaly Detector for IoT data streams

1. **Speed:** Ability to handle data coming in rapidly. In the case of heart rate monitoring and Arrhythmia Detection, data coming in every 5 seconds (e.g. Cardiogram App for Apple Watch).
2. **Memory:** Ability to handle massive amounts of data with limited memory. One data point a second is 86400 data points a day. With multiple sources of data, this can go to a million data points a day - However, anomaly detectors may only be able to use a small window size around the current timestamp to analyze and detect anomalies.

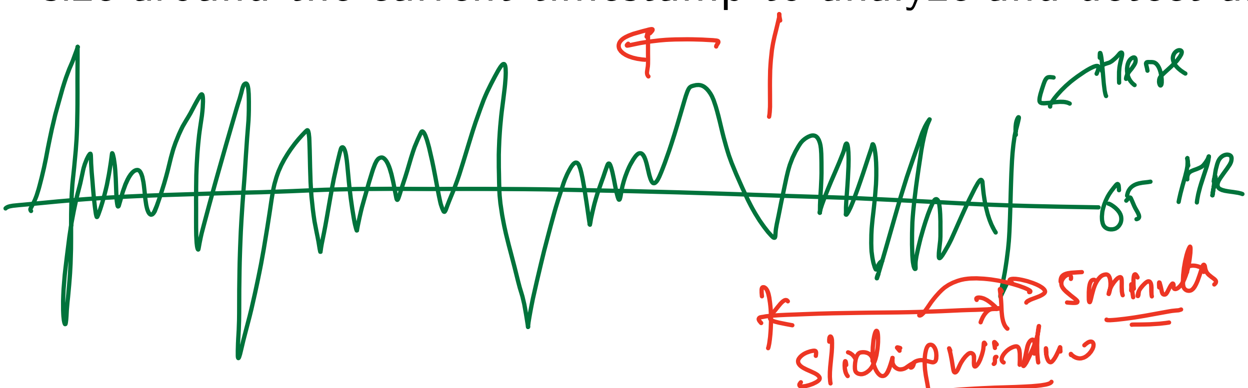

Anomaly Detection in IoT context

Properties of a good Anomaly Detector for IoT data streams

1. **Speed:** Ability to handle data coming in rapidly. In the case of heart rate monitoring and Arrhythmia Detection, data coming in every 5 seconds (e.g. Cardiogram App for Apple Watch).
2. **Memory:** Ability to handle massive amounts of data with limited memory. One data point a second is 86400 data points a day. With multiple sources of data, this can go to a million data points a day - However, anomaly detectors may only be able to use a small window size around the current timestamp to analyze and detect anomalies.
3. **High dimensionality:** Heart rate is single dimension. Combine this with other sensors, and many more dimensions emerge and make the data stream more complex. Health care applications might be good on this.

Vec for Anomaly Detection

Anomaly Detection in IoT context

Properties of a good Anomaly Detector for IoT data streams

1. **Speed:** Ability to handle data coming in rapidly. In the case of heart rate monitoring and Arrhythmia Detection, data coming in every 5 seconds (e.g. Cardiogram App for Apple Watch).
2. **Memory:** Ability to handle massive amounts of data with limited memory. One data point a second is 86400 data points a day. With multiple sources of data, this can go to a million data points a day - However, anomaly detectors may only be able to use a small window size around the current timestamp to analyze and detect anomalies.
3. **High dimensionality:** Heart rate is single dimension. Combine this with other sensors, and many more dimensions emerge and make the data stream more complex. Health care applications might be good on this.

change point detection + ?

4. **Data Drift:** Ability to handle changing data streams, changing baselines in HR or O2, understanding contexts.

anomaly detection.

Types of Outliers/Anomalies

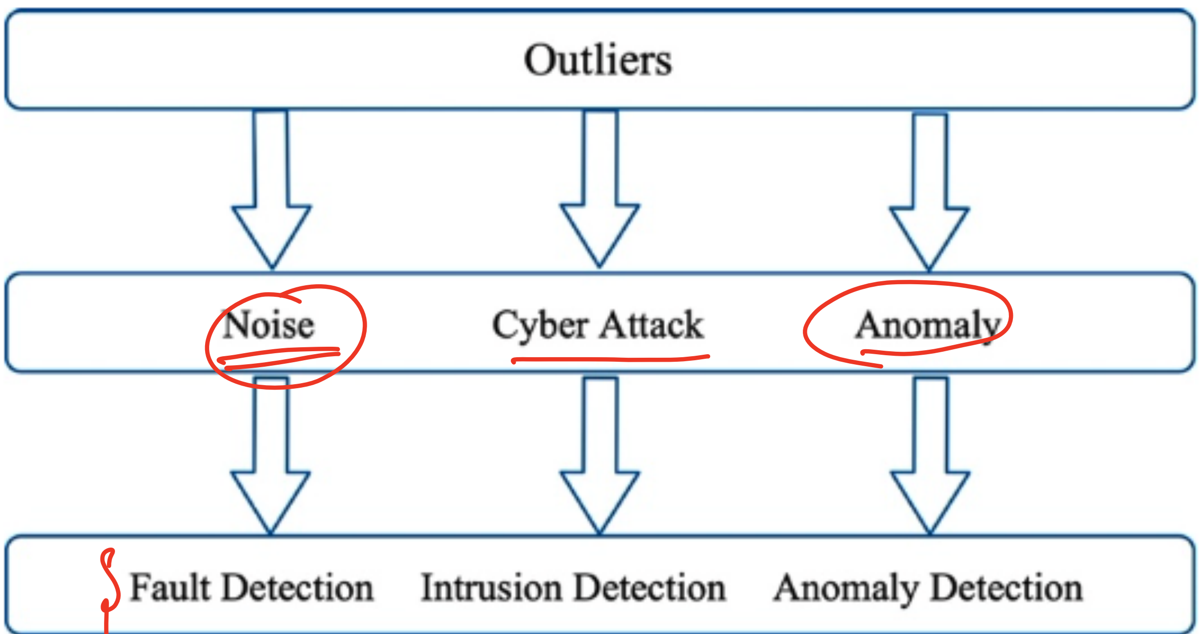

(DOS, DDOS)

Point Anomaly

→ muscle twitches as an example

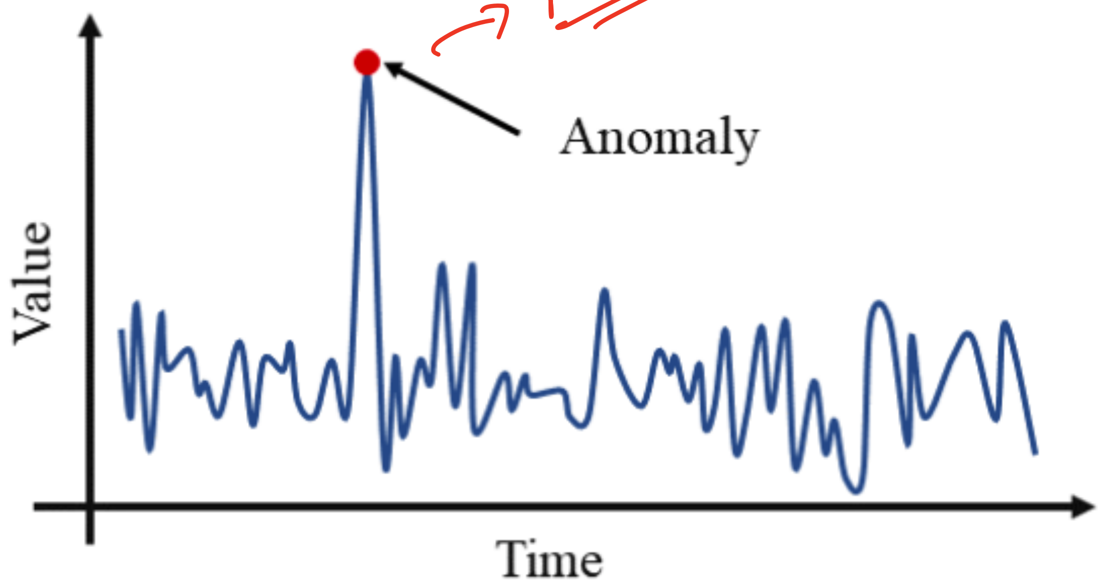

Contextual Anomaly

Sliding windows

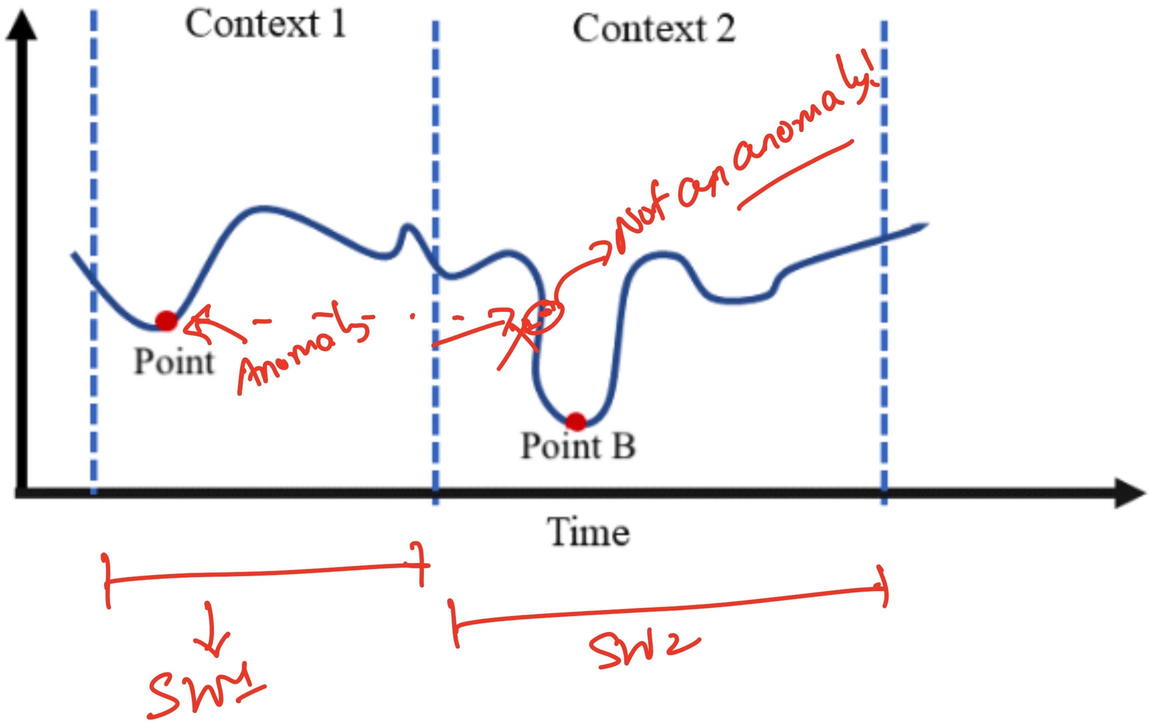

Collective Anomaly

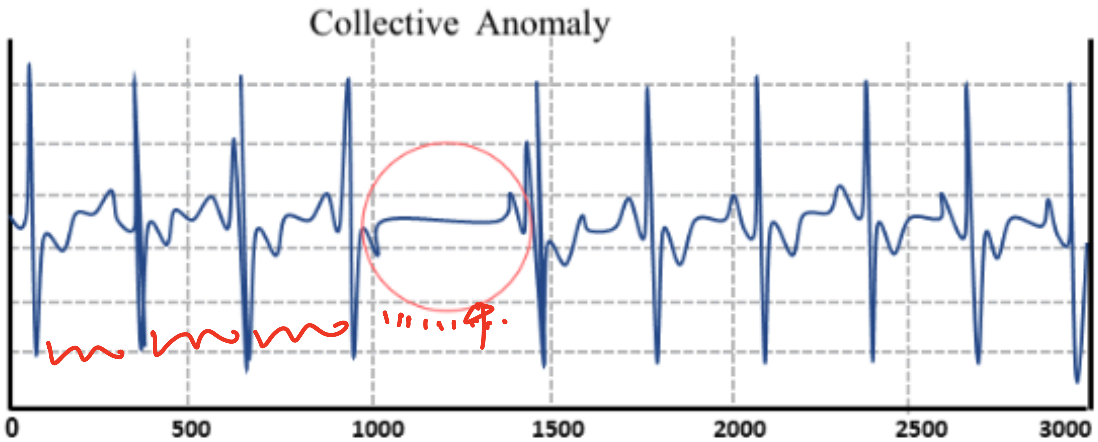

Taxanomy of Anomaly Detection Landscape

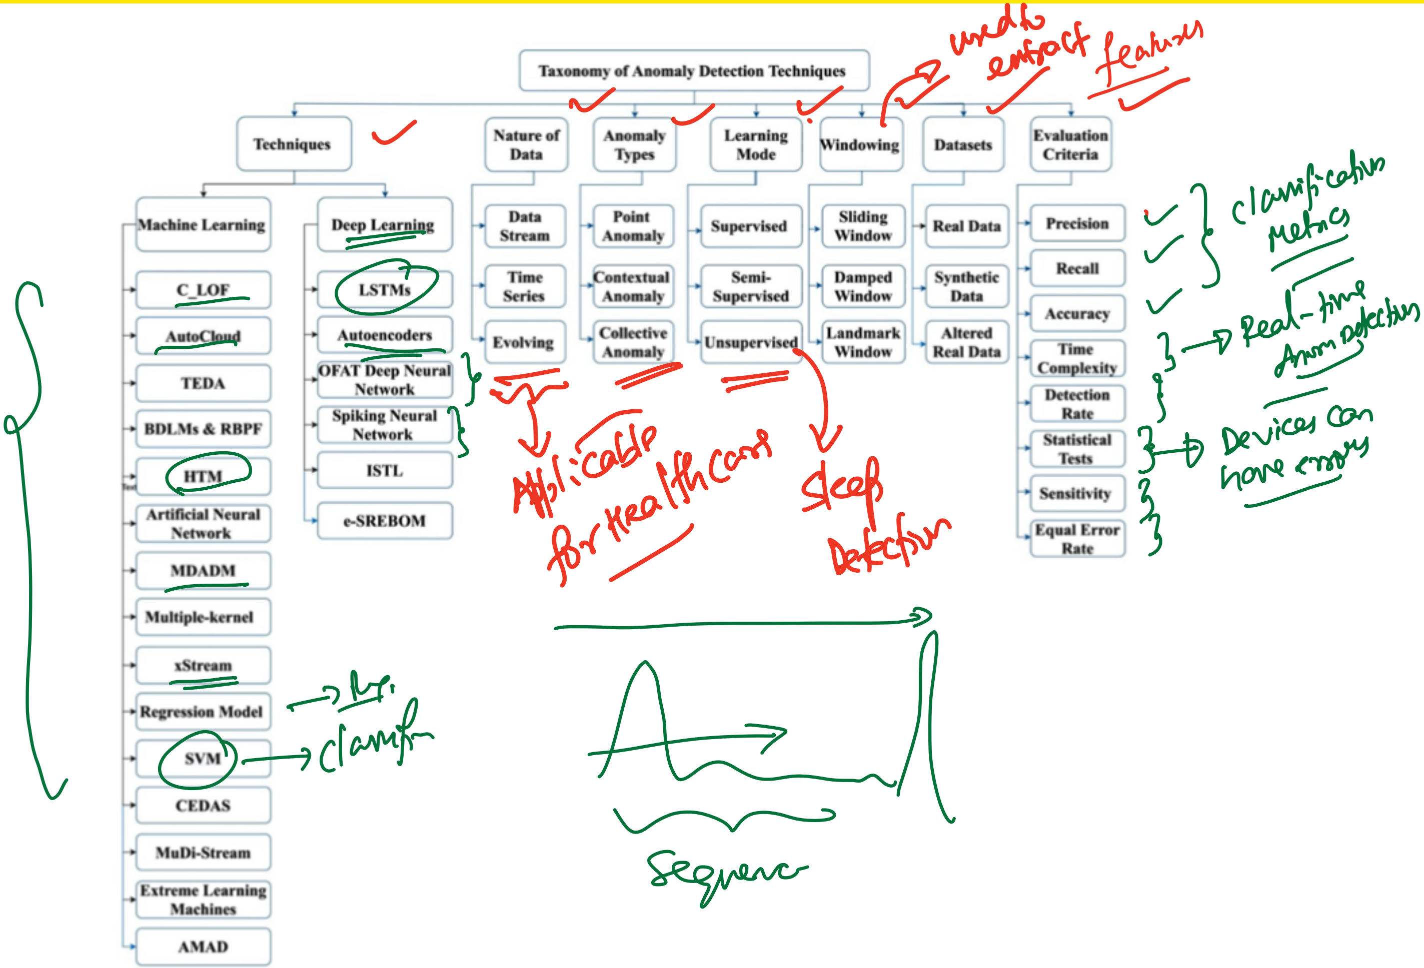

Anomaly Detection Methods

Table 1. Summary of machine learning techniques for data stream anomaly detection.

Techniques	**Nature of the Data**	**Types of Anomaly**	**Anomaly Detection Types**	**Windowing**	**Dataset**	**Evaluation Criteria**
C_LOF [14]	Data Stream (evolving)	Point anomaly	Unsupervised learning using density	Sliding window	synthetic and real-life datasets.	Precision, Recall, and Accuracy
AutoCloud [24]	Data Stream (evolving)	Point anomaly	Unsupervised learning using clustering	Sliding window	Artificial and real dataset	N/A
TEDA Clustering [25]	Data Stream (evolving)	Point anomaly	Unsupervised learning using clustering	Sliding window	Own synthetic data sets	Accuracy, Time complexity
Combination of (BDLMs) & (RBPF) [26]	Data Stream (evolving)	Point anomaly	Unsupervised learning using density	Sliding window	Artificial dataset	Accuracy, the Detection rate
HTM [27]	Data Stream	Point anomaly	Unsupervised learning based on HTM	N/A	Dataset of space imager data stream	Accuracy
Artificial Neural Network [28]	Continuous and image data	Point anomaly	Unsupervised learning on patterns of WSN nodes	Sliding window	The experimental tests that have been conducted and cover more than 27	Accuracy
MDADM [29]	Continuous data	Point anomaly	Supervised learning	N/A	Own dataset	Accuracy

DL Methods

Collective Anomaly

Techniques	Nature of the Data	Types of Anomaly	Anomaly Detection Types	Windowing	Dataset	Evaluation Criteria
LSTMs [40]	Time-Series	Point anomaly	Supervised learning using deep learning	Sliding window	Yahoo Webscope	Confusion matrix.
Autoencoder [41]	Data Stream (evolving)	Point anomaly	Unsupervised learning based on Ensembles neural networks	Sliding window	HTTP, SMTP, SMTP+HTTP, COVERTYPE, SHUTTLE, Weather	AUC
(OFAT) Deep neural network [42]	Time series	Point anomaly	Supervised learning	Window-based	Web traffic dataset, Avocado dataset, Temperature dataset	Statistical tests (average Rank), Mean Average Score (MAS)
Evolving spiking neural network [43]	Data Stream (evolving)	Point anomaly	Unsupervised learning	Sliding window	3 Benchmark dataset	Accuracy
ISTL [44]	Data Stream (evolving)	Point anomaly	Unsupervised learning based on deep learning	Sliding Window	UCSD Pedestrian datasets, Ped 1 and Ped 2) and CUHK Avenue dataset	Accuracy (ACU), Equal Error Rate (EER),
(e-SREBOM) [43]	Data Stream (evolving)	Point anomaly	Unsupervised learning using Spiking Neural Networks (eSNN)	Window-based	Water_tower_dataset, gas_dataset, electric_dataset	Accuracy, Speed, Time to learn

Feed Forward Neural Based Anomaly Detection

Input dimension = 60!

memory of 5 mins

1x60 vector

$W^{(1)}$

$W^{(2)}$

$W^{(3)}$

Anomaly or Not?

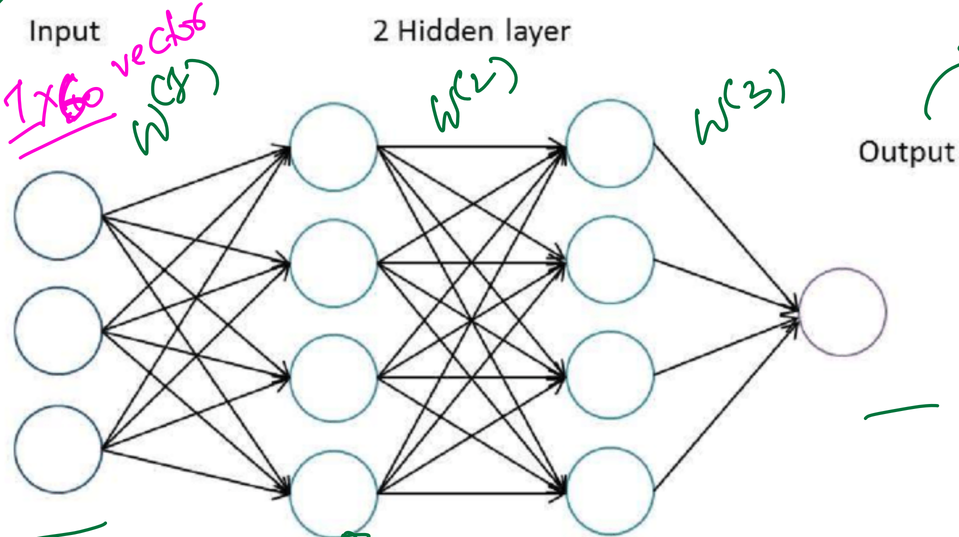

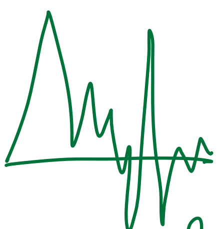

HR

9:05 pm

9:10 pm

60 datapoints at 5 second granularity for window of 5 mins!

non-linear activation fn.

Sliding Windows in Anomaly Detection

Less weight

(Exponential moving Average

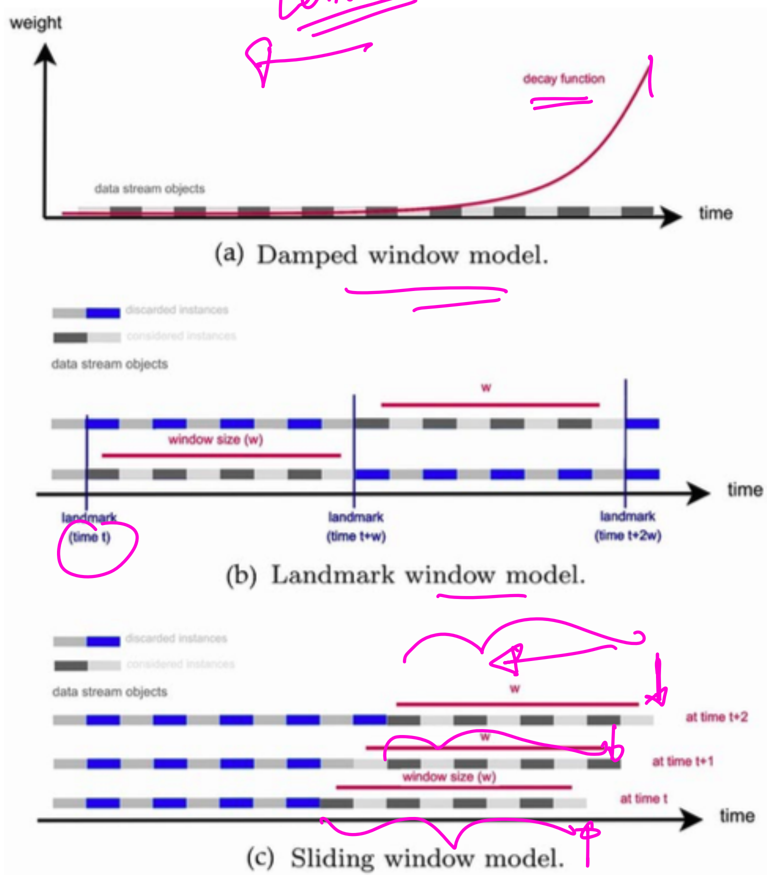

All of data but not practical

→ Typical ✓

Comparison of Methods on different dimensions

Techniques/Methods	Projection	Handling Noisy Data	Limited Time	Limited Memory	Handling Evolving Data	Handling High Dimensional Data	Evolving Features	Scalability
C_LOF [14]				√		√		√
AutoCloud [24]				√	√	√		√
TEDA Clustering [25]				√	√	√		√
Combination of (BDLMs) & (RBPF) [26]		√						√
HTM [27]		√		√				
Artificial Neural Network [28]				√				√
MDADM [29]		√		√	√			
Multi-kernel [30]	√	√			√	√	√	√
xStream [31]	√		√					

Comparison of Methods on different dimensions

Generalizability Scale Data Drift

Techniques/Methods	Projection	Handling Noisy Data	Limited Time	Limited Memory	Handling Evolving Data	Handling High Dimensional Data	Evolving Features	Scalability
Regression Model [32]			√					
Super Vector Machine [33]		√		√				√
HTM [34]	√	√		√	√	√		
CEDAS [36]		√		√				
HTM [35]		√	√	√	√			
MuDi-Stream [37]				√		√		
Extreme Learning Machines [38]	√	√		√	√	√		√
AMAD [39]	√	√	√	√	√	√		√
LSTMs [40]		√		√				
Autoencoder [41]		√			√			√
(OFAT) Deep neural network [42]		√			√	√	√	√
Evolving spiking neural network [43]		√			√			
ISTL [44]			√	√	√			
(e-SREBOM) [43]		√	√	√	√			√

next lecture

High

Inference!

References

1. A review of ML and DL techniques for Anomaly Detection in IoT Data