

14744 Mobile and Embedded Software Design


## Modeling Time Series

David Varodayan  
varodayan@cmu.edu

# What is a time series?

A sequence of measurements taken at regular time intervals


<https://map.purpleair.com>



# Why model time series?

## Identifying phases


<https://covid19.sccgov.org/dashboard-cases-and-deaths>



# Why model time series?


Identifying trends

<https://covid19.sccgov.org/dashboard-cases-and-deaths>




# Why model time series?

Identifying  
anomalies



# Why model time series?

Identifying  
sensor failures



# Modeling Time Series

- Asking a question
- Collecting data
- Understanding the data
- Exploring the data
- Constructing a model
- Evaluating and tuning the model
- Iterating
- Interpreting errors

# Asking a question

Can you use a wrist-mounted acceleration sensor to determine whether you are brushing your teeth long enough?





# Collecting data



## Dataset for ADL Recognition with Wrist-worn Accelerometer Data Set

Download: [Data Folder](#), [Data Set Description](#)

**Abstract:** Recordings of 16 volunteers performing 14 Activities of Daily Living (ADL) while carrying a single wrist-worn tri-axial accelerometer.

<b>Data Set Characteristics:</b>	Multivariate, Time-Series	<b>Number of Instances:</b>	N/A	<b>Area:</b>	Computer
<b>Attribute Characteristics:</b>	N/A	<b>Number of Attributes:</b>	3	<b>Date Donated</b>	2014-02-11
<b>Associated Tasks:</b>	Classification, Clustering	<b>Missing Values?</b>	N/A	<b>Number of Web Hits:</b>	88049

### Source:

Barbara Bruno, Fulvio Mastrogiovanni, Antonio Sgorbissa  
Laborium - Laboratory for Ambient Intelligence and Mobile Robotics  
DIBRIS, University of Genova,  
via Opera Pia 13, 16145, Genova, Italia (IT)

<https://archive.ics.uci.edu/ml/datasets/Dataset+for+ADL+Recognition+with+Wrist-worn+Accelerometer>

# Understanding the data (MANUAL.txt)

## Human motion primitives

The dataset provides labelled recorded executions of a number of simple human activities, which are defined as Human Motion Primitives (HMP):

1. brush\_teeth: to brush one's teeth with a toothbrush  
(complete gesture)
2. climb\_stairs: to climb a number of steps of a staircase
3. comb\_hair: to comb one's hair with a brush  
(complete gesture)
4. descend\_stairs: to descend a number of steps of a staircase
5. drink\_glass: to pick a glass from a table, drink and put it back on the table
6. eat\_meat: to eat something using fork and knife  
(complete gesture)
7. eat\_soup: to eat something using a spoon  
(complete gesture)
8. getup\_bed: to get up from a lying position on a bed
9. liedown\_bed: to lie down from a standing position on a bed
10. pour\_water: to pick a bottle from a table, pour its content in a glass on the table and put it back on the table
11. sitdown\_chair: to sit down on a chair
12. standup\_chair: to stand up from a chair
13. use\_telephone: to place a telephone call using a fixed telephone  
(complete gesture)
14. walk: to take a number of steps

# Understanding the data (MANUAL.txt)

## Accelerometer specifications

Type: tri-axial accelerometer  
Measurement range: [- 1.5g; + 1.5g]  
Sensitivity: 6 bits per axis  
Output data rate: 32 Hz  
Location: attached to the right wrist of the user with:  
- x axis: pointing toward the hand  
- y axis: pointing toward the left  
- z axis: perpendicular to the plane of the hand

## File naming conventions

Each file in the dataset follows the following naming convention:

Accelerometer-[START\_TIME]-[HMP]-[VOLUNTEER]

where:













- [START\_TIME]: timestamp of the starting moment of the recording in the format [YYYY-MM-DD-HH-MM-SS]
- [HMP]: name of the HMP performed in the recorded trial, following the naming convention specified in Section 2 of this manual
- [VOLUNTEER]: identification code of the volunteer performing the recorded motion in the format [gN] where:
  - "g" indicates the gender of the volunteer (m -> male, f -> female)
  - "N" indicates the progressive number associated to the volunteer

# Exploring the data


12 traces collected  
from 3 volunteers

One trace looks like

22 49 35  
22 49 35  
22 52 35  
22 52 35  
21 52 34  
22 51 34  
20 50 35  
22 52 34  
22 50 34  
22 51 35  
21 51 33  
⋮


 Accelerometer-2011-04-11-13-28-18-brush\_teeth-f1  
 Accelerometer-2011-04-11-13-29-54-brush\_teeth-f1  
 Accelerometer-2011-05-30-08-35-11-brush\_teeth-f1  
 Accelerometer-2011-05-30-09-36-50-brush\_teeth-f1  
 Accelerometer-2011-05-30-10-34-16-brush\_teeth-m1  
 Accelerometer-2011-05-30-21-10-57-brush\_teeth-f1  
 Accelerometer-2011-05-30-21-55-04-brush\_teeth-m2  
 Accelerometer-2011-05-31-15-16-47-brush\_teeth-f1  
 Accelerometer-2011-06-02-10-42-22-brush\_teeth-f1  
 Accelerometer-2011-06-02-10-45-50-brush\_teeth-f1  
 Accelerometer-2011-06-06-10-45-27-brush\_teeth-f1  
 Accelerometer-2011-06-06-10-48-05-brush\_teeth-f1

# Exploring the data (m2)



Accelerometer-2011-05-30-21-55-04-brush\_teeth-m2.txt

# Exploring the data (f1)



Accelerometer-2011-04-11-13-28-18-brush\_teeth-f1.txt


# Constructing a model

- Differencing
- Data fusion
- Smoothing
- Thresholding

[https://colab.research.google.com/drive/1wFFN-fXXnsFkd3giMFsplV\\_fL6zRSBcA?usp=sharing](https://colab.research.google.com/drive/1wFFN-fXXnsFkd3giMFsplV_fL6zRSBcA?usp=sharing)




# Differencing (f1)



$$\text{jerk}(t) = \text{acceleration}(t) - \text{acceleration}(t - 1)$$

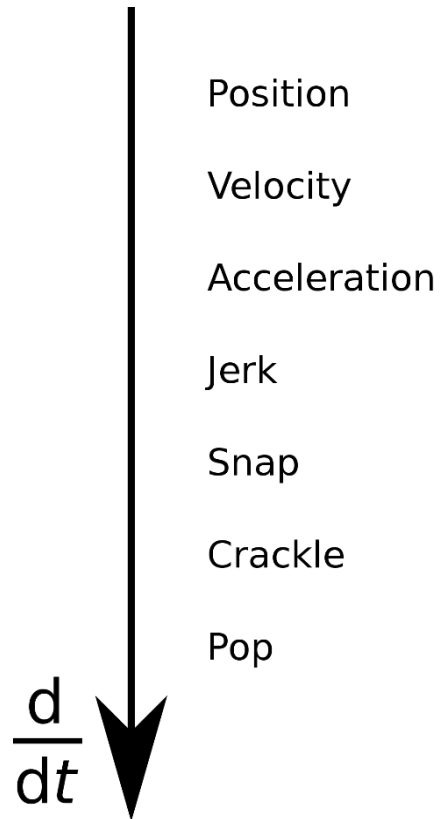


# More differencing (f1)



$$\text{snap}(t) = \text{jerk}(t) - \text{jerk}(t - 1)$$


# Jerk, snap?



[https://en.wikipedia.org/wiki/Fourth, fifth, and sixth derivatives of position](https://en.wikipedia.org/wiki/Fourth,_fifth,_and_sixth_derivatives_of_position)

[https://en.wikipedia.org/wiki/Snap, Crackle and Pop](https://en.wikipedia.org/wiki/Snap,_Crackle_and_Pop)

# Data fusion (f1)



$$\|jerk(t)\| = \sqrt{(jerk_x(t))^2 + (jerk_y(t))^2 + (jerk_z(t))^2}$$

# Smoothing


- Moving average

$$\text{smoothed}(t) = \frac{1}{N} \sum_{s=t-N+1}^t \|\text{jerk}(s)\|$$

- Exponential smoothing

$$\text{smoothed}(t) = \alpha \cdot \|\text{jerk}(t)\| + (1 - \alpha) \cdot \text{smoothed}(t - 1)$$


# Exponential smoothing (f1)



$$\text{smoothed}(t) = \alpha \cdot \|\text{jerk}(t)\| + (1 - \alpha) \cdot \text{smoothed}(t - 1), \text{ where } \alpha = 0.5, 0.2, 0.1$$


# Thresholding (f1)

After smoothing with  $\alpha = 0.2$



$\text{indicator}(t) = (\text{smoothed}(t) > \text{threshold}), \text{ where } \text{threshold} = 7$


# Comparing the indicator to raw data (f1)



Accelerometer-2011-04-11-13-28-18-brush\_teeth-f1.txt



# Comparing the indicator to raw data (m2)



Accelerometer-2011-05-30-21-55-04-brush\_teeth-m2.txt



# Have we constructed a good model?

So far we have:

- constructed a model by looking at one accelerometer trace
- “tuned” the model parameters ( $\alpha = 0.2$ , threshold = 7) on the same trace
- checked the result on one other trace

Can we be more principled about evaluating and tuning our model?


# Evaluating and tuning the model

We need to:

- collect **ground truth** for each accelerometer trace
- set up an **error cost function** to compare model output to the ground truth
- train the model parameters on a **training set** of traces
- evaluate the model on a separate **testing set** of traces


An alternative to separate training and testing sets is **cross-validation**

# Ground truth (m2)



Accelerometer-2011-05-30-21-55-04-brush\_teeth-m2.txt

# Error cost function (m2)



$$\text{cost} = \frac{1}{L} \sum_{t=0}^{L-1} |\text{indicator}(t) - \text{ground\_truth}(t)|$$

# Training and testing sets













- Split the data set into ~80% training traces and ~20% testing traces
- Find the parameters that minimize the **average training error cost** on the training set only (e.g. using grid search)
- Fix the parameters and evaluate the **average testing error cost** on the testing set only
- Why? Because we want to know the performance of our model on as-yet-unseen data

# The problems with training/testing splits

When we have only a small number of traces, we don't want to “waste” data on testing

If the data is unbalanced in some way, then it unclear how to split

Cross-validation is a solution to these problems

-  Accelerometer-2011-04-11-13-28-18-brush\_teeth-f1
-  Accelerometer-2011-04-11-13-29-54-brush\_teeth-f1
-  Accelerometer-2011-05-30-08-35-11-brush\_teeth-f1
-  Accelerometer-2011-05-30-09-36-50-brush\_teeth-f1
-  Accelerometer-2011-05-30-10-34-16-brush\_teeth-m1
-  Accelerometer-2011-05-30-21-10-57-brush\_teeth-f1
-  Accelerometer-2011-05-30-21-55-04-brush\_teeth-m2
-  Accelerometer-2011-05-31-15-16-47-brush\_teeth-f1
-  Accelerometer-2011-06-02-10-42-22-brush\_teeth-f1
-  Accelerometer-2011-06-02-10-45-50-brush\_teeth-f1
-  Accelerometer-2011-06-06-10-45-27-brush\_teeth-f1
-  Accelerometer-2011-06-06-10-48-05-brush\_teeth-f1

# Leave-one-out cross-validation

- If there are  $K$  traces, we repeat the training/testing procedure  $K$  times
- For  $i \in [1, K]$ 
  - Put the  $i^{\text{th}}$  trace in the testing set and all other traces in the training set
  - Train: Find the parameters that minimize the average training error cost
  - Test: Fix the parameters and evaluate the testing error cost on the  $i^{\text{th}}$  trace only
- Average the testing error cost over all  $K$  iterations
- Finally train the model parameters using all  $K$  traces

# Cross-validation results


- When the model was “tuned” visually based on one trace (f1)
  - Model parameters:  $\alpha = 0.2$ , threshold = 7
  - Average testing error cost = 2.13% (based on the other 11 traces)
- When the model was tuned using leave-one-out cross-validation
  - In each iteration, we optimize using grid search over combinations of
    - $\alpha \in \{0.1, 0.15, 0.2\}$
    - threshold  $\in \{7, 8, 9\}$
  - Average testing error cost = 1.40%
  - Model parameters:  $\alpha = ??$ , threshold = ??



# Iterating


- We have found better parameters for the model with the following stages
  - Differencing
  - Data fusion
  - Smoothing
  - Thresholding
- But can we do even better with a different model?
- Yes, most likely

# The anomaly in the visually “tuned” model (f1)



Using a single threshold makes the indicator sensitive to spikes in the smoothed signal

# How about a model with two thresholds? (f1)



The smoothed signal must cross both thresholds for the indicator to transition


# Cross-validation results

- For the single-threshold model
  - In each iteration, we optimize using grid search over combinations of
    - $\alpha \in \{0.1, 0.15, 0.2\}$
    - $\text{threshold} \in \{7, 8, 9\}$
  - Average testing error cost = 1.40%
- For the double-threshold model
  - In each iteration, we optimize using grid search over combinations of
    - $\alpha \in \{0.1, 0.15, 0.2\}$
    - $\text{threshold\_lo}, \text{threshold\_hi} \in \{7, 8, 9\}$
  - Average testing error cost = 1.13%

# How complex should we make the model?

- If the model is too simple, it cannot extract structure from the trace
  - This is known as **underfitting**
  - The generalization error to as-yet-unseen traces is dominated by **variance**
- If the model is too complex, it becomes overly tuned to the data set
  - This is known as **overfitting**
  - The generalization error to as-yet-unseen traces is dominated by **bias**
- The model of the right complexity trades-off between bias and variance

# Bias and variance visualized



# What if you can't get low bias and variance?

- If the variance is too high, increase the complexity of the model
- If the bias is too high
  - collect more data
  - collect data that is more representative of the as-yet-unseen data