Learning from Cohort Data

mainly for chronic disorders and their comorbidities

Myra Spiliopoulou

Dec 10, 2020

OTTO VON GUERICKE
UNIVERSITÄT
MAGDEBURG

INF
FACULTY OF
COMPUTER SCIENCE

Medical Mining in the research lab

Learning from Cohort Data
Myra Spiliopoulou
Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Mission: Serve the medical researchers, and through them physicians and patients

Approach: We develop

- ML methods for learning on streams of data for each patient
- ML methods for filling the gaps
- ML methods for the acquisition of labels, features and constraints
- Complete workflows and frameworks

In parallel, we work on experimental data:

- experiments and studies on human behaviour [analysis]
- experiments on crowdworking [our own designs]

Our projects and our inspiring partners

Ongoing Projects:

- ⋆ ImmunLearning (2019 - 2022): EFRE project on a diagnostic test for immunocompetence for elderly people (with U Med OVGU)
- ⋆ CHRODIS+ (2017-2020) EU Joint Action on "Implementing good practices for chronic diseases"
- ⋆ UNITI (2020-2022) EU project on "Unification of Treatments and Interventions for Tinnitus Patients"

Further cooperations in medical research:

- Learning on longitudinal epidemiological data (U Med Greifswald)
- Intelligent wearables for patients with diabetic foot (U Med Magdeburg)
- Phenotyping, patient evolution - clinic & m/eHealth (U Med Regensburg)
- Phenotyping and patient response to treatment (CHARITE)

The KMD Team

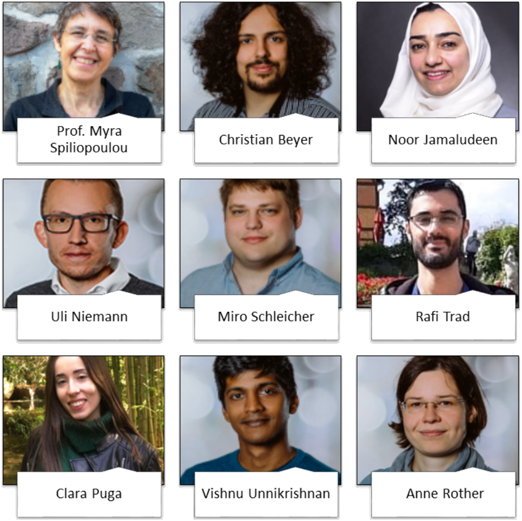
Prof. Myra Spiliopoulou
Christian Beyer
Noor Jamaludeen
Uli Niemann
Miro Schleicher
Rafi Trad
Clara Puga
Vishnu Unnikrishnan
Anne Rother

Agenda

This is Part III of Tutorial 8:

- Structure of clinical data; prediction methods; open challenges
- Structure of mobile, dynamic data; prediction methods; open challenges

RECALL: The CRISP-DM Circle [Chapman et al., 2000]

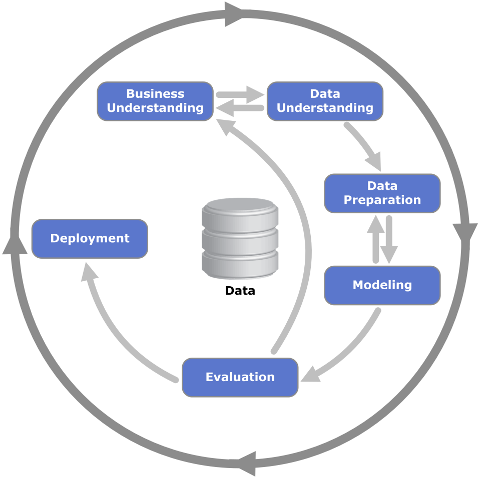

1 2

[1] **License and attribution:** Kenneth Jansen, 1 March 2013, licensed under the Creative Commons Attribution-Share Alike 3.0 Unported license, via Wikimedia Commons

[2] downloaded at Nov 26, 2020

Chronic Disorders / Diseases / Conditions

Example 1: Diabetes

Example 2: Tinnitus

[3] [4]

[3] **Attribution:** Mikael Häggström, 27 Feb 2009, public domain

[4] downloaded at Nov 26, 2020

Tinnitus as subject of research

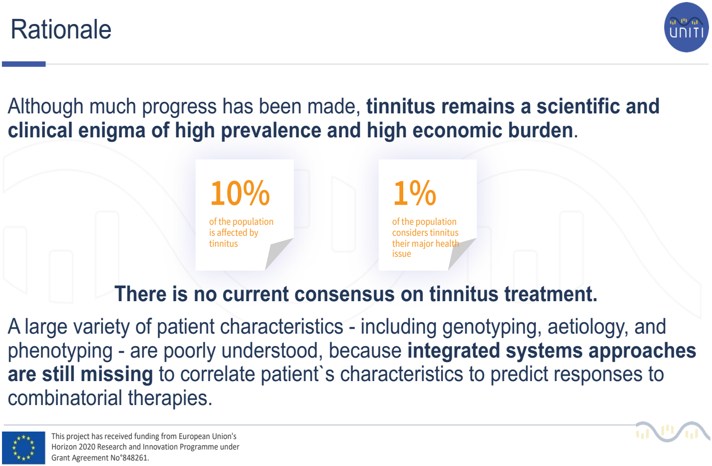

[5] H2020 project UNITI 'Unification of Treatments and Interventions for Tinnitus Patients', since Jan 1, 2020

Medical Data

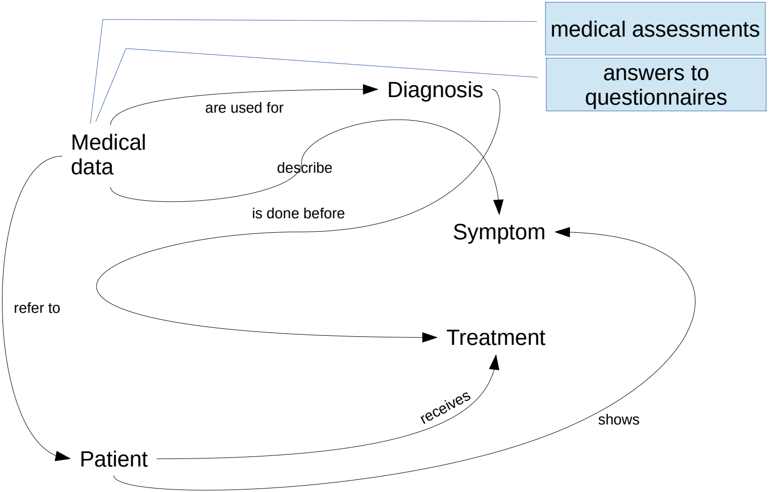

Example: The Tinnitus Case History Questionnaire [6] cf. [Langguth et al., 2007]

<u>"Items list" for tinnitus case history questionnaires</u>

Items are ordered according to their level of significance:
Category "A" (= essential) in bold type.

Background

1. **Age.**
2. **Gender.**
3. Handedness.
4. **Family history of tinnitus (parent, sibling, children).**

Tinnitus history

5. **Initial onset. Time?**
6. **Initial onset. Mode? Gradual or abrupt?**
7. Initial onset. Associated events? Hearing change, Acoustic trauma, Otitis media, Head trauma, Whiplash, Dental Treatment, Stress, Other.
8. **Pattern. Steady? Pulsatile? Other?**
9. **Site. Right ear? Left ear? Both ears? (symmetrical?) Inside head?**
10. Intermittent or constant?
11. fluctuant or non-fluctuant?
12. **Loudness. Scale 1-100. At worst & at best?**
13. Quality. Own words / Give a list of choices.
14. Pure tone or Noise? Uncertain / polyphonic?
15. Pitch. Very high? High? Medium? Low?
16. **Percentage of awake time aware of tinnitus?**
17. Percentage of awake time annoyed by tinnitus?
18. Previous tinnitus treatments (no, some, many)?

Modifying influences

19. **Natural masking? Music, everyday sounds, other sounds?**
20. Aggravated by loud noise?
21. **Altered by head and neck movement or touching of head or upper limbs (specification of the respective movements)?**
22. Daytime nap. Worse? Better? No effect?
23. Effect of nocturnal sleep on daytime tinnitus?
24. Effect of stress?
25. Effect of medications? Which?

Related conditions

26. **Hearing impairment?**
27. **Hearing aids (No, left ear, right ear, both ears; effect on tinnitus)?**
28. **Noise annoyance or intolerance?**
29. Noise induced pain?
30. Headaches?
31. Vertigo/dizziness?
32. Temporomandibular disorder?
33. Neck pain?
34. Other pain syndromes?
35. Under treatment for psychiatric problems?

As an example of how the above items can be expressed for patients to complete see the

TINNITUS SAMPLE CASE HISTORY QUESTIONNAIRE (TSCHQ)

[6] `https://www.tinnitusresearch.net/images/files/migrated/consensusdocuments/en/Items-list.pdf`, downloaded: Nov 26, 2020

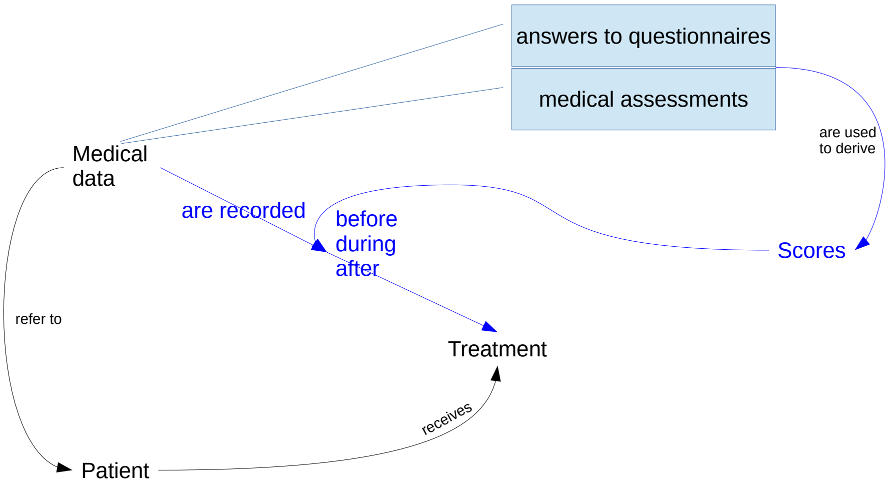
answers to questionnaires
medical assessments
are used to derive
Medical data
are recorded
before
during
after
Scores
refer to
Treatment
receives
Patient

Medical Data for Learning

Learn what?

→ the effects of a treatment
→ the side-effects of a treatment
→ the symptoms
→ the comorbidities
...

on different subpopulations

...and identify predictive features

Medical Data for Learning

Learn what?

- → the effects of a treatment
- → the side-effects of a treatment
- → the symptoms
- → the comorbidities
- . . .

. . . and identify predictive features

on different subpopulations

Learn for whom?

- the physician
- the medical researcher

Medical Data for Learning

Learn what?

→ the effects of a treatment
→ the side-effects of a treatment
→ the symptoms
→ the comorbidities
. . .

on different subpopulations

. . . and identify predictive features

Learn for whom?

- the physician
- the medical researcher

Learn on what?

· on an export of patient records
· on a cohort
· on multiple cohorts

Cohorts [Glenn, 2005]

The term "cohort"

Quoting [Glenn, 2005], page 2: "The term *cohort* originally referred to a group of warriors or soldiers, and the term is now sometimes used in a very general sense to refer to a number of individuals who have some characteristic in common."

Cohorts [Glenn, 2005]

The term "cohort"

Quoting [Glenn, 2005], page 2: "The term *cohort* originally referred to a group of warriors or soldiers, and the term is now sometimes used in a very general sense to refer to a number of individuals who have some characteristic in common."

The term "cohort" in "cohort analysis"

Quoting [Glenn, 2005], page 2: "Here and in other literature on cohort analysis, however, the term is used in a more restricted sense to refer to those individuals (human or otherwise) who experienced a particular event during a specified period of time. . . .

Learning on the cohorts of an RCT

Randomised Clinical Trial

- Development of a robust **RCT protocol** for single and combinational therapies
- Harmonization of interventions among clinical centers
- **Recruitment** of tinnitus patients for the RCT
- **Conduction** of the RCT
- Validation of the Decision Support System (DSS)

5 clinical centers

- Klinikum der Universitaet Regensburg (REG)
- Katholieke Universiteit Leuven (KUL)
- Charité – Universitaetsmedizin Berlin (CHA)
- Ethniko Kai Kapodistriako Panepistimo Athinon (UOA)
- Servicio Andaluz de Salud (GRA)

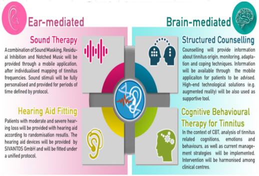

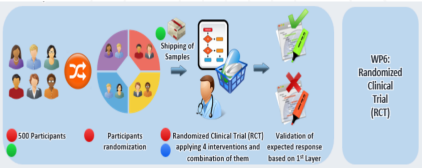

This project has received funding from European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No°848261.

Learning from Cohort Data
Myra Spiliopoulou
Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Narrowing down the subject area: Small datasets, large feature spaces, and two goals: assist the medical researcher, while reducing the data demand.

Learning: static data first

Learning from Cohort Data
Myra Spiliopoulou
Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Learning: static data first

→ To what extend can we predict an outcome after treatment (T1), when using only features recorded before treatment (T0) ?

→ How to build a minimal set of cumulatively predictive features?

Learning from Cohort Data
Myra Spiliopoulou
Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

→ To what extend can we predict whether tinnitus distress will be compensated or not after treatment (T1), when using only features recorded before treatment (T0) ?

→ How to build a minimal set of cumulatively predictive features?

Prediction and Minimization of the Feature Space [Niemann et al., 2020b]

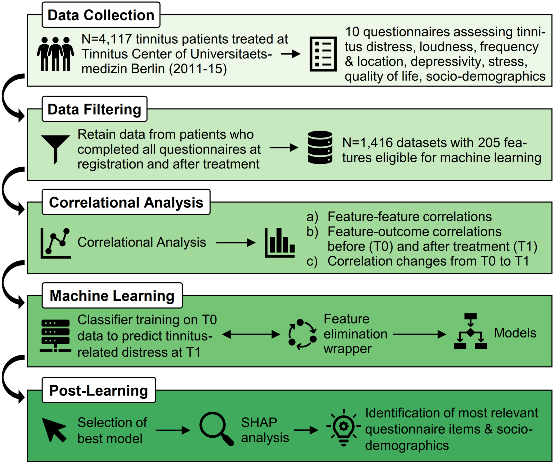

[7]

[7] Figure 1 from [Niemann et al., 2020b]

Model learning on T0 to predict T1 for target TQ_score $\in \{0, 1\}$

11 algorithms: LASSO, RIDGE, Generalized Partial Least Squares (GPLS), SVM, a feedforward NN with one hidden layer, weighted kNN, NB, CART, C5.0, RF, **Gradient Boosted Trees (GBT)**

hyperparam selection: grid search

k=10 folds: 10-fold stratified cross validation

AUC: as performance measure

Feature filtering and scoring

Model reliance [Fisher et al., 2018]: of feature f as "classification error on the original training set [vs] classification error on a modified version of the training set where the values of f are randomly permuted."

SHAP value [Lundberg and Lee, 2017] of feature f for instance x "as change in the expected value of the prediction if for f the feature vector of x is observed instead of [a] random [one]"

Prediction and Minimization of the Feature Space [Niemann et al., 2020b]

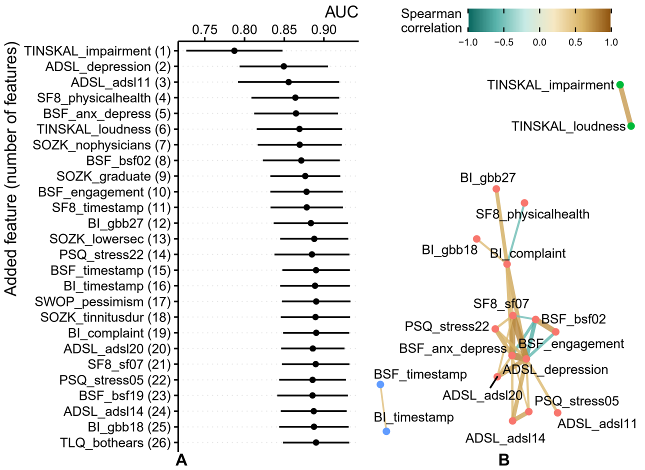

8

---

[8] Figure 6 from [Niemann et al., 2020b]: (A) features sorted on importance, and achieved AUC of a feature and its predecessors; (B) graph of correlations (for absolute values of 0.5 and higher)

Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Learning on static data $\surd$

$\surd$ Workflows for learning on static data and showing the solutions

$\surd$ First attempts to sort out features from a huge feature space

The feature space and the *data drain* problem

Data collection: 4,103 tinnitus patients, older than 18
treated at the Tinnitus Center of Charité Universitätsmedizin Berlin
between Jan. 2011 and Oct. 2015, 2having tinnitus for at least 3 months

Four studies:

- √ Prediction and feature space minimization [Niemann et al., 2020b]: 1,416 patients

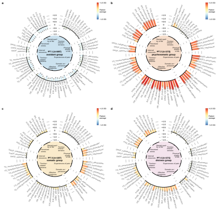

- Phenotypes at T0 and T1 [Niemann et al., 2020d]: 1,228 patients
- Predicting depression severity after treatment [Niemann et al., 2020c]: 1,490 patients
- Identifying gender-specific differences among chronic tinnitus patients [Niemann et al., 2020a]: 1,228 patients (m:609, f:619)

What makes the difference?

The feature space and the *data drain* problem

Data collection: 4,103 tinnitus patients, older than 18
treated at the Tinnitus Center of Charité Universitätsmedizin Berlin
between Jan. 2011 and Oct. 2015, 2having tinnitus for at least 3 months

Four studies:

- √ Prediction and feature space minimization [Niemann et al., 2020b]: 1,416 patients

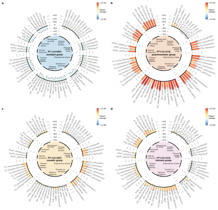

- Phenotypes at T0 and T1 [Niemann et al., 2020d]: 1,228 patients
- Predicting depression severity after treatment [Niemann et al., 2020c]: 1,490 patients
- Identifying gender-specific differences among chronic tinnitus patients [Niemann et al., 2020a]: 1,228 patients (m:609, f:619)

What makes the difference? The medical objective

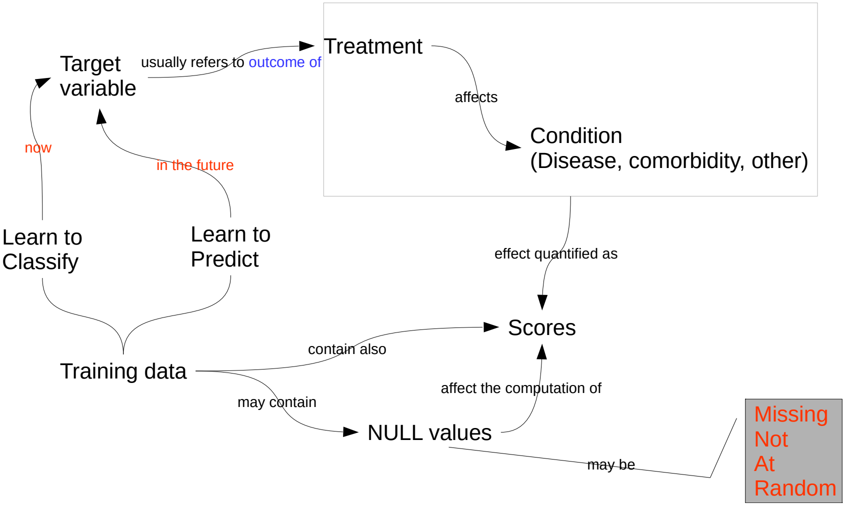
Target variable
usually refers to outcome of
Treatment
affects
Condition (Disease, comorbidity, other)
now
in the future
Learn to Classify
Learn to Predict
effect quantified as
Scores
contain also
Training data
affect the computation of
may contain
NULL values
may be
Missing Not At Random

Learning on small static data √

Open issues:

→ how to reduce the feature space

··· and thus reduce the data demand [working on that ...]

→ how to exploit time and deal with gaps [comes next]

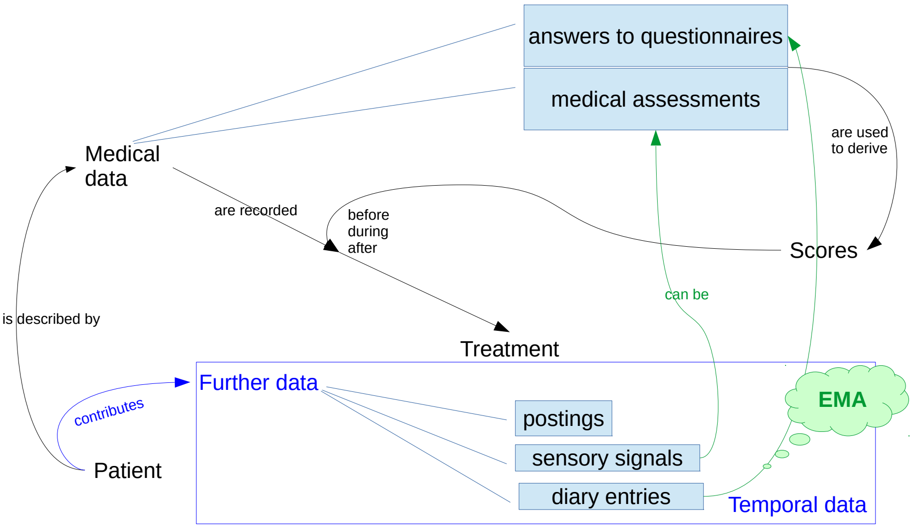

answers to questionnaires
medical assessments
are used
to derive
Medical
data
are recorded
before
during
after
Scores
can be
is described by
Treatment
Further data
contributes
postings
sensory signals
diary entries
EMA
Patient
Temporal data

Learning patterns of movement for patient wellbeing

Example: Wandering

"Wandering is among the most frequent, problematic, and dangerous behaviors for elders with dementia. Frequent wanderers likely suffer falls and fractures, which affect the safety and quality of their lives." [Lin et al., 2012]

- "Travel behavior of nursing home residents perceived as wanderers and non-wanderers" [Martino-Saltzman et al., 1991]
- "Detecting wandering behavior based on GPS traces for elders with dementia" [Lin et al., 2012]
- "Location prediction using GPS trackers: Can machine learning help locate the missing people with dementia?" [Wojtusiak and Nia, 2019]

. . . on mitigating the side-effects of smartphone battery-saving

"Inferring Mobility Measures from GPS Traces with Missing Data" [Barnett and Onnela, 2020]

9

[9] Thanks to Priyanka Mohan for the literature collection and to Dr. Nadine Diersch – *Aging & Cognition Research Group, German Center for Neurodegenerative Diseases (DZNE), Magdeburg* – working on real-world navigation behavior

The arrowline of the mHealth potential [Kumar et al., 2013]

Measurement → Diagnostic → Treatment/prevention → Global

where **Measurement** encompasses:

→ on-person or embedded sensor sampling in real time
- Global Positioning System
- Ecological Momentary Assessment

↓

Conversion of the
"raw sensor data into meaningful information related to behaviors, thoughts, emotions . . . and clinical states and disorders." [Mohr et al., 2017]

Converting sensor data into information [Mohr et al., 2017]

Learning from Cohort Data
Myra Spiliopoulou

Medical Mining in the KMD Lab

Business Understanding
- Termini
- Medical data for learning

Learning on Static Data
- Prediction with few features
- Data drain

Learning on Dynamic Data
- **Learning on sensors**
- Learning on user inputs
  - → Learning from the 'long' users
  - → Prediction for (almost) all users
  - → Interpreting the gaps

Closing

The "sensemaking" approach of [Mohr et al., 2017]:

- ⋆ Clinical state
- ↑ "Behavioral marker: behaviors, thoughts, feelings, traits, or states . . . "
- ↑ "[Low level] feature: a measureable property of a phenomenon, which is proximal to, and constructed from, sensor data"
- ↑ Raw sensor data

Converting sensor data into information [Mohr et al., 2017]

The "sensemaking" approach of [Mohr et al., 2017]:

- $\star$ Clinical state
- $\uparrow$ "Behavioral marker: behaviors, thoughts, feelings, traits, or states . . . "
- $\uparrow$ "[Low level] feature: a measureable property of a phenomenon, which is proximal to, and constructed from, sensor data"
- $\uparrow$ Raw sensor data

Example [Mohr et al., 2017] (Sections 2.3, 3.1)

- Behavioral marker: Sleep disruption
- Low-level features: Bedtime/waketime, phone usage, movement intensity, ambient noise
- Clinical targets: depression, bipolar disorder, schizophrenia

The arrowline of the mHealth potential [Kumar et al., 2013]

Measurement $\rightarrow$ Diagnostic $\rightarrow$ Treatment/prevention $\rightarrow$ Global

where **Measurement** encompasses:

- on-person or embedded sensor sampling in real time
- Global Positioning System

$\rightarrow$ Ecological Momentary Assessment

The arrowline of the mHealth potential [Kumar et al., 2013]

Measurement → Diagnostic → Treatment/prevention → Global

where **Measurement** encompasses:

- on-person or embedded sensor sampling in real time
- Global Positioning System

→ Ecological Momentary Assessment

"Ambulatory assessment" [Fahrenberg et al., 2007]

Quoting from https://psycnet.apa.org/record/2007-18155-002:

> *Ambulatory assessment refers to the use of computer-assisted methodology for self-reports, behavior records, or physiological measurements, while the participant undergoes normal daily activities. . . .*

CHRODIS+ "Implementing good practices for chronic diseases" [10]

Joint Action (09/17- 11/20) funded by the EU and participating organisations

GOAL: Contribute to reducing the burden of chronic diseases

by promoting the implementation of policies and practices that have been demonstrated to be successful. The development and sharing of these tested policies and projects across EU countries is the core idea behind this action.

WP7: Fostering the quality of care for people with chronic diseases:

through the implementation of a set of quality criteria and recommendations defined in the previous JA CHRODIS.

Task 3 on mHealth tools: three pilot studies on self-empowerment of patients with help of mobile technology:

- Cantabrian Health Service – CSC in Spain
- National Center of Public Health and Analyses – NCPHA in Bulgaria
- University Hospital Regensburg – UHREG in Germany

[10] `http://chrodis.eu/about-us/`

EMA for different CHRODIS+ T7.3 pilot studies

T7.3: NPCHA pilot on diabetes ⟶ [Unnikrishnan et al., 2020a]

T7.3: UHREG pilot TinnitusTips on tinnitus [Unnikrishnan et al., 2020b]

Questionnaire – many times a day

ID	Item description and valuerange	
s01	Did you perceive the tinnitus right now?	Y/N
s02	How loud is your tinnitus right now?	0...100
s03	How distressed are you by your tinnitus right now?	0...100
s04	How well do you hear right now?	0...100
s05	How much are you limited by your hearing right now?	0...100
s06	How stressed do you feel right now?	0...100
s07	How exhausted do you feel right now?	0...100
s08	Are you wearing a hearing aid right now?	Y/N

Groups of participants: Y, A, B

Groups Y and A received non-personalized tips from the beginning; group B received the tips only from the middle of the study onwards.

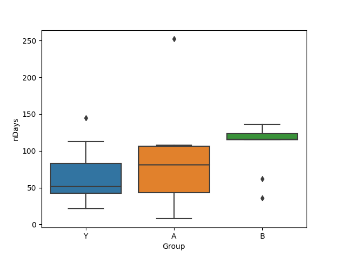

1. Learning from the ‘long’ users

Learning from Cohort Data
Myra Spiliopoulou
Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ **Learning from the ‘long’ users**
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Analyzing EMA on TrackYourTinnitus [Probst et al., 2017]

Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

- Total assessments: 25,863 Retained: 17,209, after excluding assessments with missing values in any of the target variables and days with less than three assessments.
- 350 participants (253m/94f) with average age 45.4 (over 333, SD=12.1) and median since tinnitus onset 5.4Y (from 0 to 61.8Y)
- Median days per participant 11 (from 1 to 415) with median number of assessments per day 4 (from 3 to 18)

Analyzing EMA on TrackYourTinnitus [Probst et al., 2017]

Learning from Cohort Data
Myra Spiliopoulou

Medical Mining in the KMD Lab

Business Understanding
- Termini
- Medical data for learning

Learning on Static Data
- Prediction with few features
- Data drain

Learning on Dynamic Data
- Learning on sensors
- Learning on user inputs
  - → Learning from the 'long' users
  - → Prediction for (almost) all users
  - → Interpreting the gaps

Closing

- Total assessments: 25,863 Retained: 17,209, after excluding assessments with missing values in any of the target variables and days with less than three assessments.
- 350 participants (253m/94f) with average age 45.4 (over 333, SD=12.1) and median since tinnitus onset 5.4Y (from 0 to 61.8Y)
- Median days per participant 11 (from 1 to 415) with median number of assessments per day 4 (from 3 to 18)

Selection of findings

- "tinnitus was significantly louder in the late evening compared to the afternoon and early evening."
- "stress-level increased from morning to afternoon, decreased from afternoon to evening, and did not differ compared to the night"
- "the effects of time-of-day on tinnitus loudness and tinnitus distress were still significant (i.e., after controlling for the effects of stress)."

Analyzing EMA on TrackYourTinnitus [Probst et al., 2017]

Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

- Total assessments: 25,863 Retained: 17,209, after excluding assessments with missing values in any of the target variables and days with less than three assessments.
- 350 participants (253m/94f) with average age 45.4 (over 333, SD=12.1) and median since tinnitus onset 5.4Y (from 0 to 61.8Y)
- Median days per participant 11 (from 1 to 415) with median number of assessments per day 4 (from 3 to 18)

Selection of findings

- "tinnitus was significantly louder in the late evening compared to the afternoon and early evening."
- "stress-level increased from morning to afternoon, decreased from afternoon to evening, and did not differ compared to the night"
- "the effects of time-of-day on tinnitus loudness and tinnitus distress were still significant (i.e., after controlling for the effects of stress)."

What about users with missing EMA values and less EMA per day?

TrackYourTinnitus data distribution

[Schleicher et al., 2020]

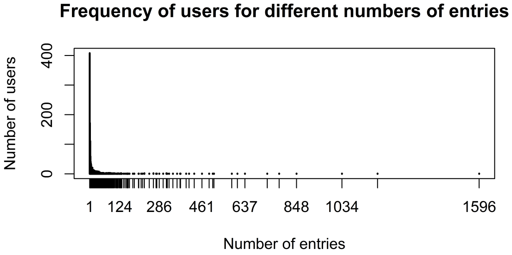

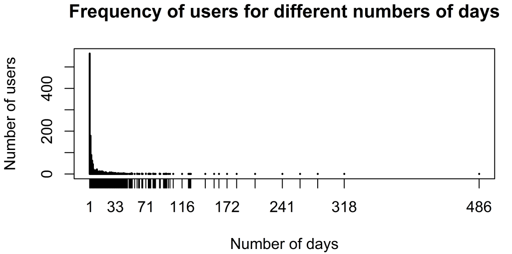

2. Prediction for (almost) all users

Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Data and model augmentation for distress prediction

[Unnikrishnan et al., 2019]

Learning from Cohort Data
Myra Spiliopoulou

Medical Mining in the KMD Lab

Business Understanding
- Termini
- Medical data for learning

Learning on Static Data
- Prediction with few features
- Data drain

Learning on Dynamic Data
- Learning on sensors
- Learning on user inputs
- → Learning from the 'long' users
- → **Prediction for (almost) all users**
- → Interpreting the gaps

Closing

Modeling the prediction problem

Given is a set of participants, for the EMA of whom we have the `distress` values - but only for the first m EMA. For each patient x and EMA $o_{x,j}$ with $j > m$, predict the distress value of $o_{x,j}$.

Data and model augmentation for distress prediction

[Unnikrishnan et al., 2019]

Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Modeling the prediction problem

Given is a set of participants, for the EMA of whom we have the `distress` values - but only for the first m EMA. For each patient x and EMA $o_{x,j}$ with $j > m$, predict the distress value of $o_{x,j}$.

Data cleaning

Removal of:

- participants with less than 5 EMA $\longrightarrow$ 516 participants

Histogramm: #days per participant

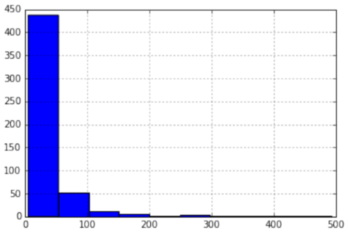

Data and model augmentation in two feature spaces

[Unnikrishnan et al., 2019]

3. kNN-based predictors in D_{EMA}

- Predictor 1 (model augmentation):
  for each y in the neighbourhood of x (including x), learn a linear regression model m_y;
  average the parameters $m_{y,slope}$ and $m_{y,intercept}$ into a final model m_x.
- Predictor 2 (data augmentation):
  place all EMA of x and of its neighbours into a pool;
  learn a linear regression model for m_x over the pool

Data and model augmentation in two feature spaces

[Unnikrishnan et al., 2019]

3. kNN-based predictors in D_{EMA}

- Predictor 1 (model augmentation):
  for each y in the neighbourhood of x (including x), learn a linear regression model m_y;
  average the parameters $m_{y,slope}$ and $m_{y,intercept}$ into a final model m_x.
- Predictor 2 (data augmentation):
  place all EMA of x and of its neighbours into a pool;
  learn a linear regression model for m_x over the pool

2. Exploiting the EMA dataset D_{EMA} and the registration data D_R for learning

- The kNN neighbourhood of each participant x is built on D_R
- Time series of participants are aligned across the time axis – no gaps are filled in D_{EMA}

Augmentation with phenotype information

[Unnikrishnan et al., 2019]

1. Combining tinnitus loudness and distress levels on D_R

Three clusters on loudness and two clusters on distress levels:

Tinnitus Loudness	*Tinnitus Distress*					
	Low TD			*High TD*		
	#P	Avg TD		#P	Avg TD	
Low TL	81	7.4	Avg TL: 28.1	52	15.7	Avg TL: 26.8
Mod TL	83	8.1	Avg TL: 53.0	168	17.0	Avg TL: 54.4
High TL	35	9.2	Avg TL: 77.1	97	18.3	Avg TL: 82.2

Prediction for the six phenotypes [Unnikrishnan et al., 2019]

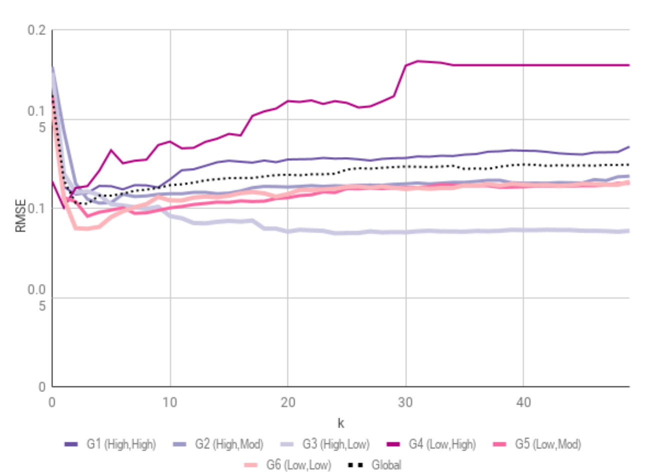

3. Interpreting the gaps

Medical Mining in the KMD Lab

Business Understanding

- Termini
- Medical data for learning

Learning on Static Data

- Prediction with few features
- Data drain

Learning on Dynamic Data

- Learning on sensors
- Learning on user inputs
  - → Learning from the 'long' users
  - → Prediction for (almost) all users
  - → **Interpreting the gaps**

Closing

User Adherence in TrackYourTinnitus [Schleicher et al., 2020]

Adherence as combination of

- ⋆ Interaction duration
  - · Days until first break
- ⋆ Interaction intensity as:
  - · Return after first break (Y/N) within the observation horizon of 30 days

User Adherence in TrackYourTinnitus [Schleicher et al., 2020]

Adherence as combination of

- ⋆ Interaction duration
- ⋆ Interaction intensity as:
  - Days until first break
  - Return after first break (Y/N) within the observation horizon of 30 days

Materials: EMA of 1252 users

- For time series classification: 852 users, after removing 440 users that had only one EMA
- For interpretation of interaction intensity: 816 users, after removing users with incomplete time series and/or missing values

User Adherence in TrackYourTinnitus [Schleicher et al., 2020]

Adherence as combination of

- ⋆ Interaction duration
- ⋆ Interaction intensity as:
  - · Days until first break
  - · Return after first break (Y/N) within the observation horizon of 30 days

Materials: EMA of 1252 users

- For time series classification: 852 users, after removing 440 users that had only one EMA
- For interpretation of interaction intensity: 816 users, after removing users with incomplete time series and/or missing values

Methods

For time series classification on each EMA item separately:

7 Time series classification algorithms with various distance functions

6 conventional classification algorithms

For interaction intensity interpretation: classification rule induction

User Adherence in TrackYourTinnitus [Schleicher et al., 2020]

Adherence as combination of

- ⋆ Interaction duration
- ⋆ Interaction intensity as:
  - Days until first break
  - Return after first break (Y/N) within the observation horizon of 30 days

Materials: EMA of 1252 users

- For time series classification: 852 users, after removing 440 users that had only one EMA
- For interpretation of interaction intensity: 816 users, after removing users with incomplete time series and/or missing values

Methods

For time series classification on each EMA item separately:

7 Time series classification algorithms with various distance functions

6 conventional classification algorithms

For interaction intensity interpretation: classification rule induction

Results

...

Closing: Learning on small mHealth data

Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Huge potential for

- ⋆ patient empowerment
- ⋆ assistance towards caregivers
- ⋆ better diagnostics
- ⋆ deeper insights to chronic diseases

Closing: Learning on small mHealth data

Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Huge potential for

- ⋆ patient empowerment
- ⋆ assistance towards caregivers
- ⋆ better diagnostics
- ⋆ deeper insights to chronic diseases

The dilemma of time series length:

- long time series: useful for phenotyping and prediction
- short time series: little to learn on, very little to predict

Some solutions:

- entity-centric methods that prevent building a global model mainly on the data of very few users
- methods for learning and predicting despite the gaps
- methods for the identification of influence factors

Closing: Learning on small dynamic data

Learning from Cohort Data
Myra Spiliopoulou
Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

Open issues:

- → How to exploit ML-methods designed for large dense datasets
- → Gaps are missing data not at random – how to deal with this?
- → Exploration/Exploitation problem – how to learn with even fewer data?

THANKS TO:

Medical Mining in the KMD Lab
Business Understanding
Termini
Medical data for learning
Learning on Static Data
Prediction with few features
Data drain
Learning on Dynamic Data
Learning on sensors
Learning on user inputs
→ Learning from the 'long' users
→ Prediction for (almost) all users
→ Interpreting the gaps
Closing

- This work was partially supported by the CHRODIS+ Joint Action on "Implementing good practices for chronic diseases".
- This work is partially supported by the European Union's Horizon 2020 Research and Innovation Programme, Grant Agreement Number 848261.

THANK YOU !

Bibliography I

[Barnett and Onnela, 2020] Barnett, I. and Onnela, J.-P. (2020).
Inferring mobility measures from gps traces with missing data.
Biostatistics, 21(2):e98–e112.

[Chapman et al., 2000] Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., Wirth, R., et al. (2000).
Crisp-dm 1.0: Step-by-step data mining guide.
SPSS inc, 9:13.

[Fahrenberg et al., 2007] Fahrenberg, J., Myrtek, M., Pawlik, K., and Perrez, M. (2007).
Ambulatory assessment–monitoring behavior in daily life settings: A behavioral-scientific challenge for psychology.
European Journal of Psychological Assessment, 23(4):206.

[Fisher et al., 2018] Fisher, A., Rudin, C., and Dominici, F. (2018).
All models are wrong but many are useful: Variable importance for black-box, proprietary, or misspecified prediction models, using model class reliance.
arXiv preprint arXiv:1801.01489, pages 237–246.

[Glenn, 2005] Glenn, N. D. (2005).
Cohort Analysis.
Quantitative Applications in the Social Sciences. SAGE, 2nd edition.

[Kumar et al., 2013] Kumar, S., Nilsen, W. J., Abernethy, A., Atienza, A., Patrick, K., Pavel, M., Riley, W. T., Shar, A., Spring, B., Spruijt-Metz, D., et al. (2013).
Mobile health technology evaluation: the mHealth evidence workshop.
American journal of preventive medicine, 45(2):228–236.

Bibliography II

[Langguth et al., 2007] Langguth, B., Goodey, R., Azevedo, A., Bjorne, A., Cacace, A., Crocetti, A., Del Bo, L., De Ridder, D., Diges, I., Elbert, T., et al. (2007).
Consensus for tinnitus patient assessment and treatment outcome measurement: Tinnitus research initiative meeting, regensburg, july 2006.
Progress in brain research, 166:525–536.

[Lin et al., 2012] Lin, Q., Zhang, D., Huang, X., Ni, H., and Zhou, X. (2012).
Detecting wandering behavior based on gps traces for elders with dementia.
In *2012 12th International Conference on Control Automation Robotics & Vision (ICARCV)*, pages 672–677. IEEE.

[Lundberg and Lee, 2017] Lundberg, S. M. and Lee, S.-I. (2017).
A unified approach to interpreting model predictions.
In *Advances in neural information processing systems*, pages 4765–4774.

[Martino-Saltzman et al., 1991] Martino-Saltzman, D., Blasch, B. B., Morris, R. D., and McNeal, L. W. (1991).
Travel behavior of nursing home residents perceived as wanderers and nonwanderers.
The Gerontologist, 31(5):666–672.

[Mohr et al., 2017] Mohr, D. C., Zhang, M., and Schueller, S. M. (2017).
Personal sensing: understanding mental health using ubiquitous sensors and machine learning.
Annual review of clinical psychology, 13:23–47.

[Niemann et al., 2020a] Niemann, U., Boecking, B., Brueggemann, P., Mazurek, B., and Spiliopoulou, M. (2020a).
Gender-specific differences in patients with chronic tinnitus—baseline characteristics and treatment effects.
Frontiers in Neuroscience, 14:487.

Bibliography III

Learning from Cohort Data

Myra Spiliopoulou

Medical Mining in the KMD Lab

Business Understanding

Termini

Medical data for learning

Learning on Static Data

Prediction with few features

Data drain

Learning on Dynamic Data

Learning on sensors

Learning on user inputs

→ Learning from the 'long' users

→ Prediction for (almost) all users

→ Interpreting the gaps

Closing

[Niemann et al., 2020b] Niemann, U., Boecking, B., Brueggemann, P., Mebus, W., Mazurek, B., and Spiliopoulou, M. (2020b).
Tinnitus-related distress after multimodal treatment can be characterized using a key subset of baseline variables.
PLOS ONE, 15(1):1–18.

[Niemann et al., 2020c] Niemann, U., Brueggemann, P., Boecking, B., Mazurek, B., and Spiliopoulou, M. (2020c).
Development and internal validation of a depression severity prediction model for tinnitus patients based on questionnaire responses and socio-demographics.
Scientific Reports, 10(1):4664.

[Niemann et al., 2020d] Niemann, U., Brueggemann, P., Boecking, B., Rose, M., Spiliopoulou, M., and Mazurek, B. (2020d).
Phenotyping chronic tinnitus patients using self-report questionnaire data: cluster analysis and visual comparison.
Scientific Reports, 10(1):16411.

[Probst et al., 2017] Probst, T., Pryss, R. C., Langguth, B., Rauschecker, J. P., Schobel, J., Reichert, M., Spiliopoulou, M., Schlee, W., and Zimmermann, J. (2017).
Does tinnitus depend on time-of-day? an ecological momentary assessment study with the “trackyourtinnitus” application.
Frontiers in Aging Neuroscience, 9:253.

Bibliography IV

[Schleicher et al., 2020] Schleicher, M., Unnikrishnan, V., Neff, P., Simoes, J., Probst, T., Pryss, R., Schlee, W., and Spiliopoulou, M. (2020).
Understanding adherence to the recording of ecological momentary assessments on the example of tinnitus monitoring.
Scientific Reports.
accepted, Dec. 9, 2020.

[Unnikrishnan et al., 2019] Unnikrishnan, V., Beyer, C., Matuszyk, P., Niemann, U., Pryss, R., Schlee, W., Ntoutsi, E., and Spiliopoulou, M. (2019).
Entity-level stream classification: exploiting entity similarity to label the future observations referring to an entity.
International Journal of Data Science and Analytics, pages 1–15.

[Unnikrishnan et al., 2020a] Unnikrishnan, V., Schleicher, M., Shah, Y., Jamaludeen, N., Pryss, R., Schlee, W., and Spiliopoulou, M. (2020a).
Predicting the health condition of mhealth app users with large differences in the amount of recorded observations – where to learn from?
In *International Conference on Discovery Science*. Springer.

[Unnikrishnan et al., 2020b] Unnikrishnan, V., Schleicher, M., Shah, Y., Jamaludeen, N., Pryss, R., Schobel, J., Kraft, R., Schlee, W., and Spiliopoulou, M. (2020b).
The effect of non-personalised tips on the continued use of self-monitoring mhealth applications.
Brain Sciences, 10(12).

Bibliography V

[Wojtusiak and Nia, 2019] Wojtusiak, J. and Nia, R. M. (2019).
Location prediction using gps trackers: Can machine learning help locate the missing people with dementia?
Internet of Things, page 100035.

[Zhang et al., 2014] Zhang, Z., Gotz, D., and Perer, A. (2014).
Iterative cohort analysis and exploration.
Information Visualization (Info Vis), pages 1–19.