Mobile and Sensor Systems

Lecture 5: Sensor Data Inference
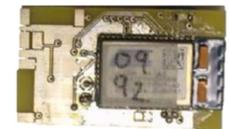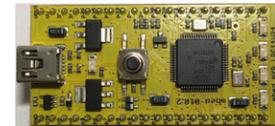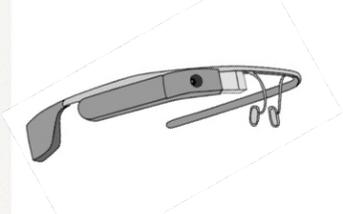
Prof C Mascolo

[ack to Prof T. Plötz for some material]

In this lecture

- We will talk about mobile sensing

- We will describe the challenges in sensor inference

- We will talk about the steps involved in the sensor to inference process

- We will introduce traditional feature based and neural network based modelling

Mobile and Wearable Sensing

Sensors

- Microphone
- Camera
- GPS
- Accelerometer
- Compass
- Gyroscope
- WiFi
- Bluetooth
- Proximity
- Light
- NFC (near field communication)

UNIVERSITY OF CAMBRIDGE

Applications

- **Individual sensing**:
  - Fitness and health applications
  - behaviour intervention applications
- **Group/community sensing**:
  - groups to sense common activities and help achieving group goals
  - examples: assessment of neighbourhood safety, environmental sensing, collective recycling efforts
- **Urban-scale sensing**:
  - large scale sensing, where large number of people have the same application installed
  - examples: tracking speed of disease across a city, congestion and pollution in a city

UNIVERSITY OF CAMBRIDGE

Sensor Based Inference Systems Characteristics

- Offline/Online Inference

- Continuous/Periodic/Isolated Inference

- In all cases, collecting ground truth is key:
 - data needs precise labels…

[Bulling et al 2013]
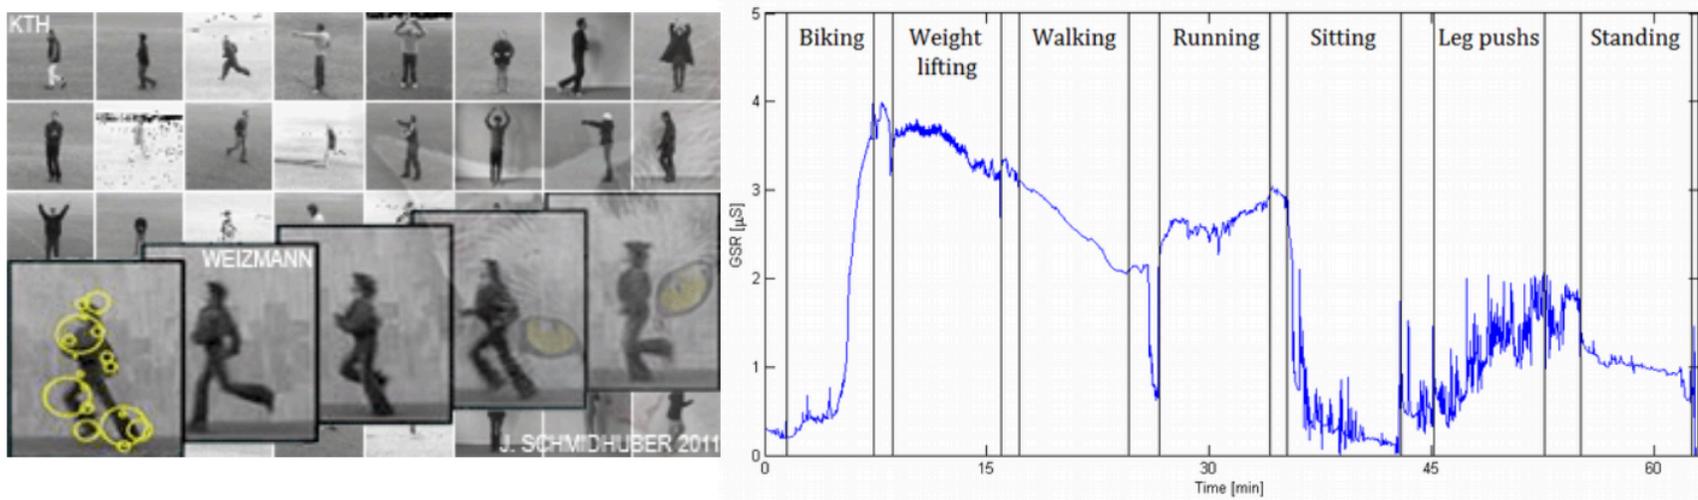
Sensor Data Analysis Challenges: Differences in Users and Device Positions

- Users have differences which influence the readings on the sensors (eg different gait)

# Sensor Data Analysis Challenges: Range of Activities

- Activities (classes) can often be many and sometimes very similar (or similar for some users)

- Collection of balanced data among those classes (for ground truth) can be challenging
  - The nature of the data collection might impose imbalances in the data

- Ground Truth Annotation is hard...
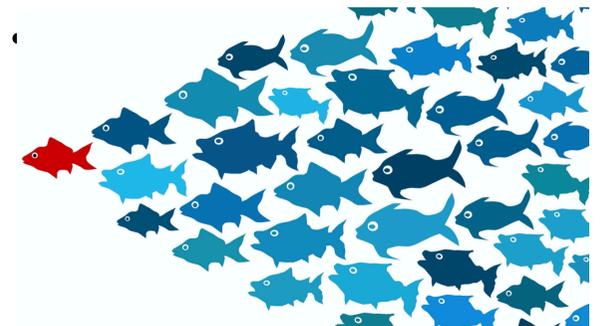
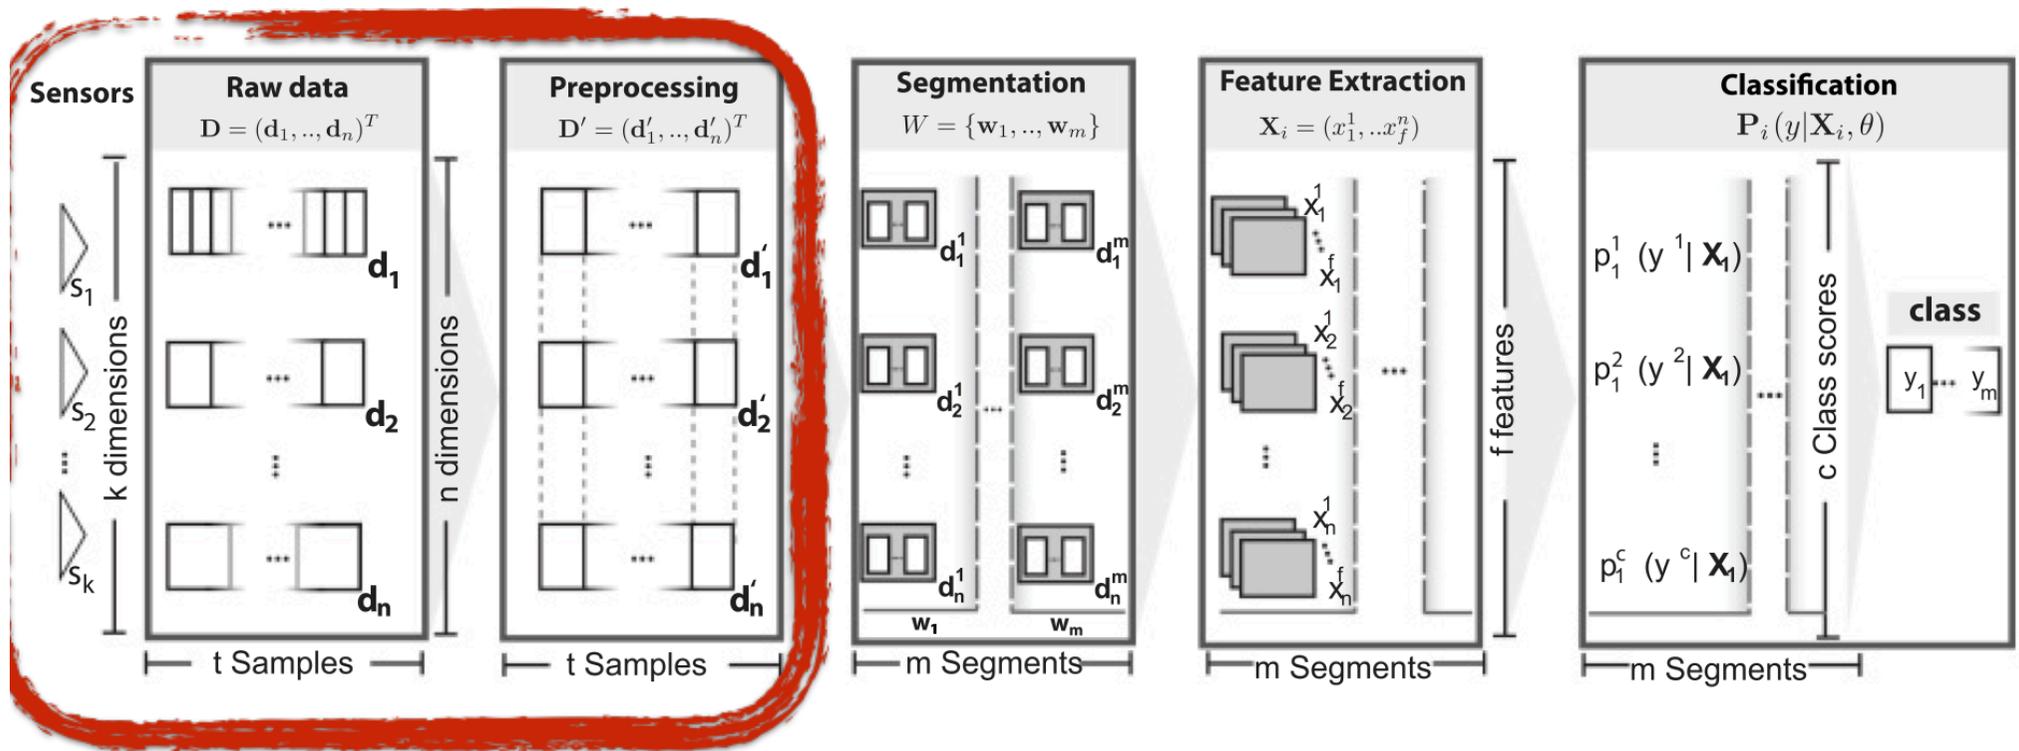# Activity Recognition from Sensor Data

- Activity recognition aims to recognize the **actions of an individual** from a **series of observations** on the individual's actions and the **environmental conditions**.

UNIVERSITY OF CAMBRIDGE

Activity Recognition
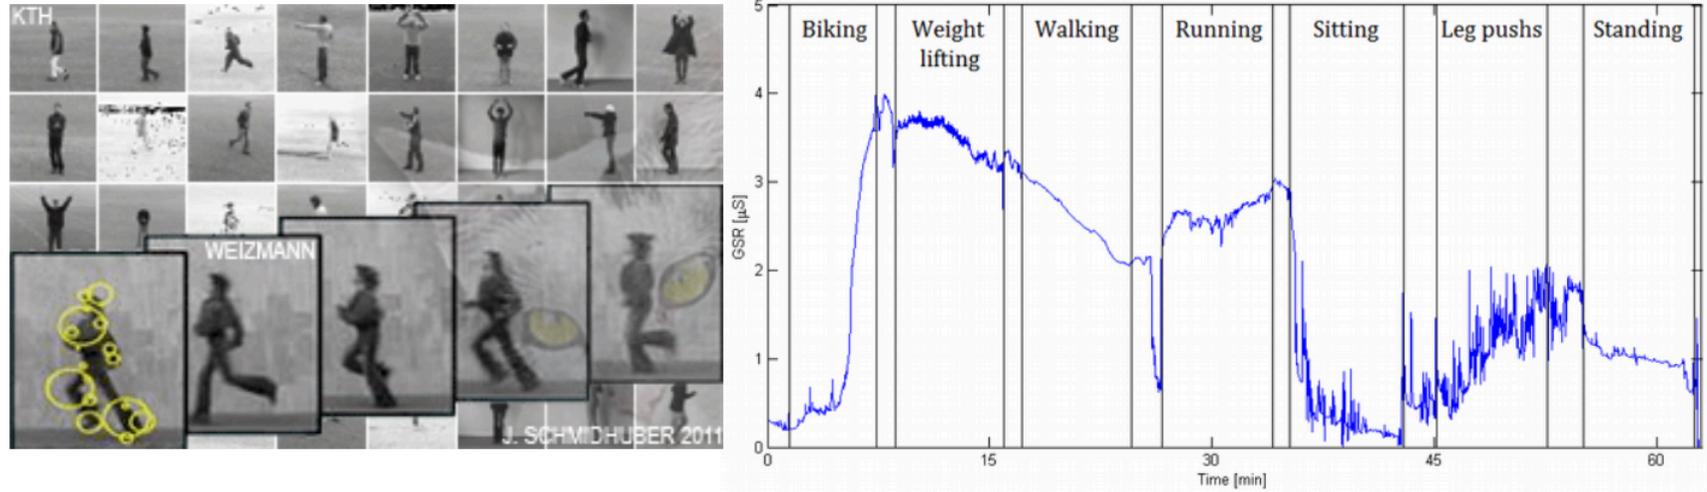
- Wearables and mobiles produce **sequential data**

Sensor Inference Pipeline

Sensor Raw Data and Preprocessing

- Acquisition can happen
  - from different sensors (at different locations on the body or orientation, or from different sensors, acceleration or GPS…)
  - At different sampling rate (eg for energy reasons)
- Sensor data can be corrupted or contain errors
- Preprocessing synchronizes and removes artifacts (calibration, unit conversion, normalization, resampling, synchronization..)

Accelerometer Preprocessing

- Magnitude

$$\forall i : m_i = \sqrt{x^2 + y^2 + z^2}$$
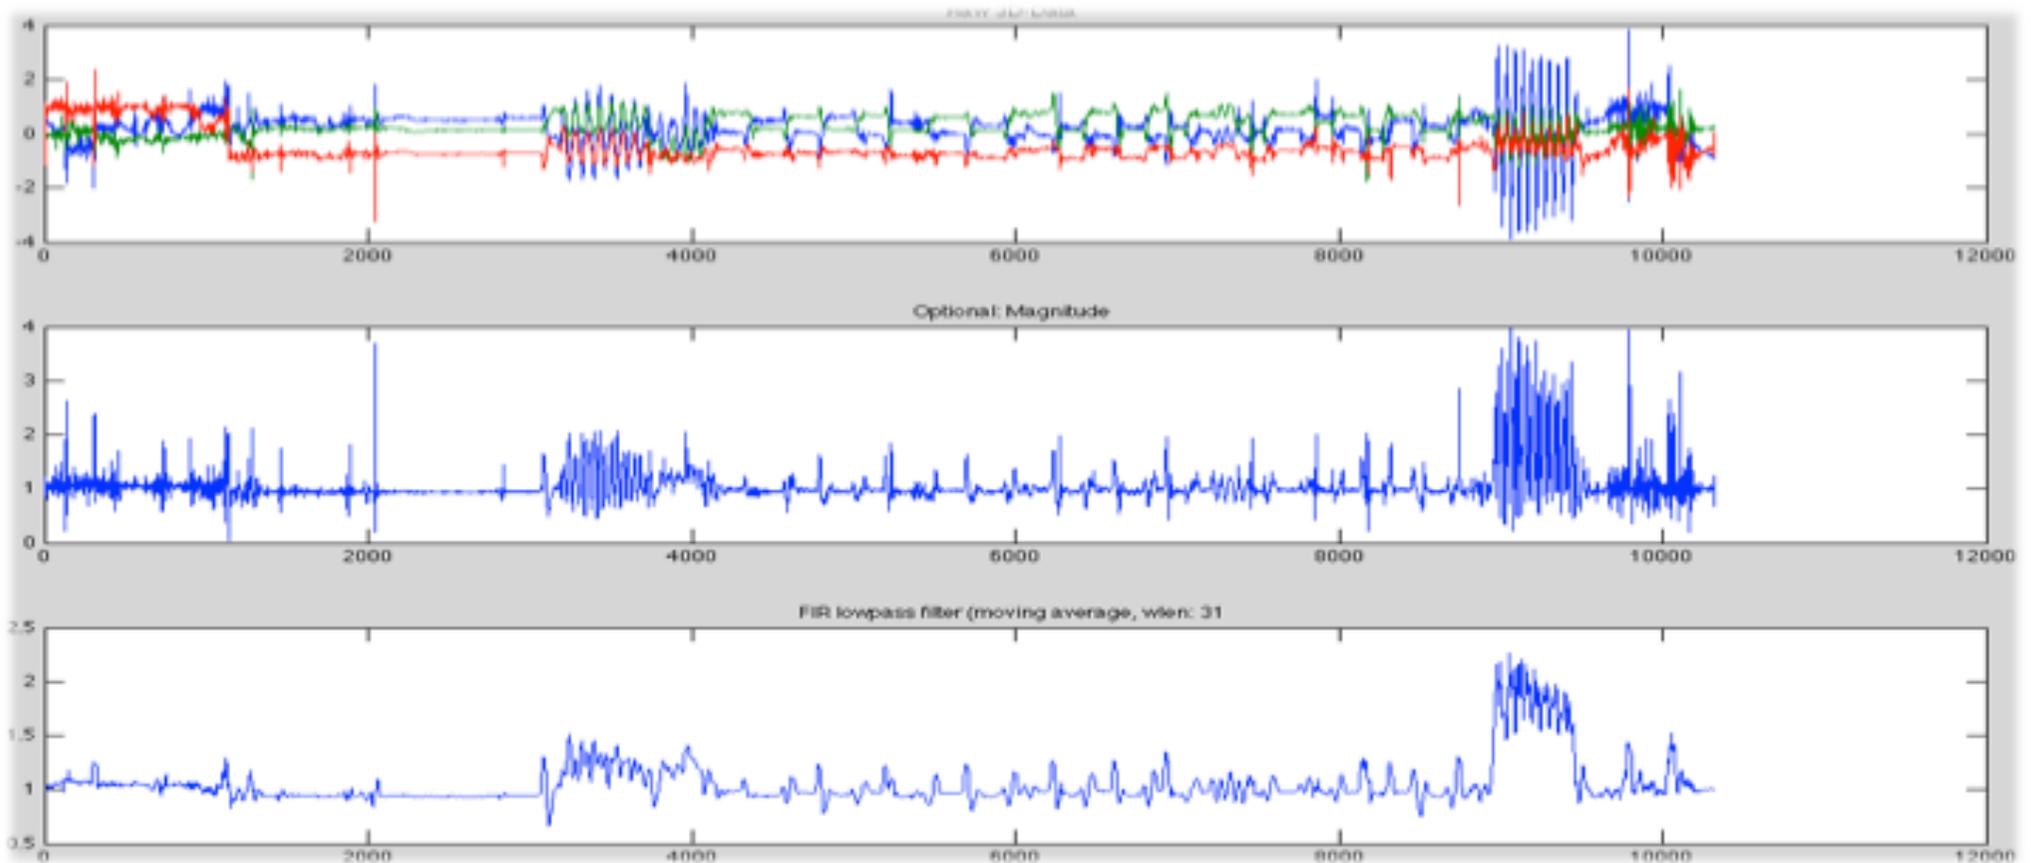
Sampling Rate

| Dataset	#Classes	Q	q (S=0.99)	Δ
Skoda	11	96Hz	12Hz	-87.5%
PAMAP2-Hand			32Hz	-68%
PAMAP2-Chest	13	100Hz	33Hz	-67%
PAMAP2-Ankle			42Hz	-58%
USC-HAD	12	100Hz	17Hz	-83%
PHealth	10	100Hz	15Hz	-85%
Walk8	4	250Hz	18Hz	-92.8%

UNIVERSITY OF
CAMBRIDGE

[Khan et a 2016]

Data Segmentation

- Localize temporal patterns of interest
- But you do not know what/where these are…
- Sliding window approach
 - Works well with quasi-periodic activities
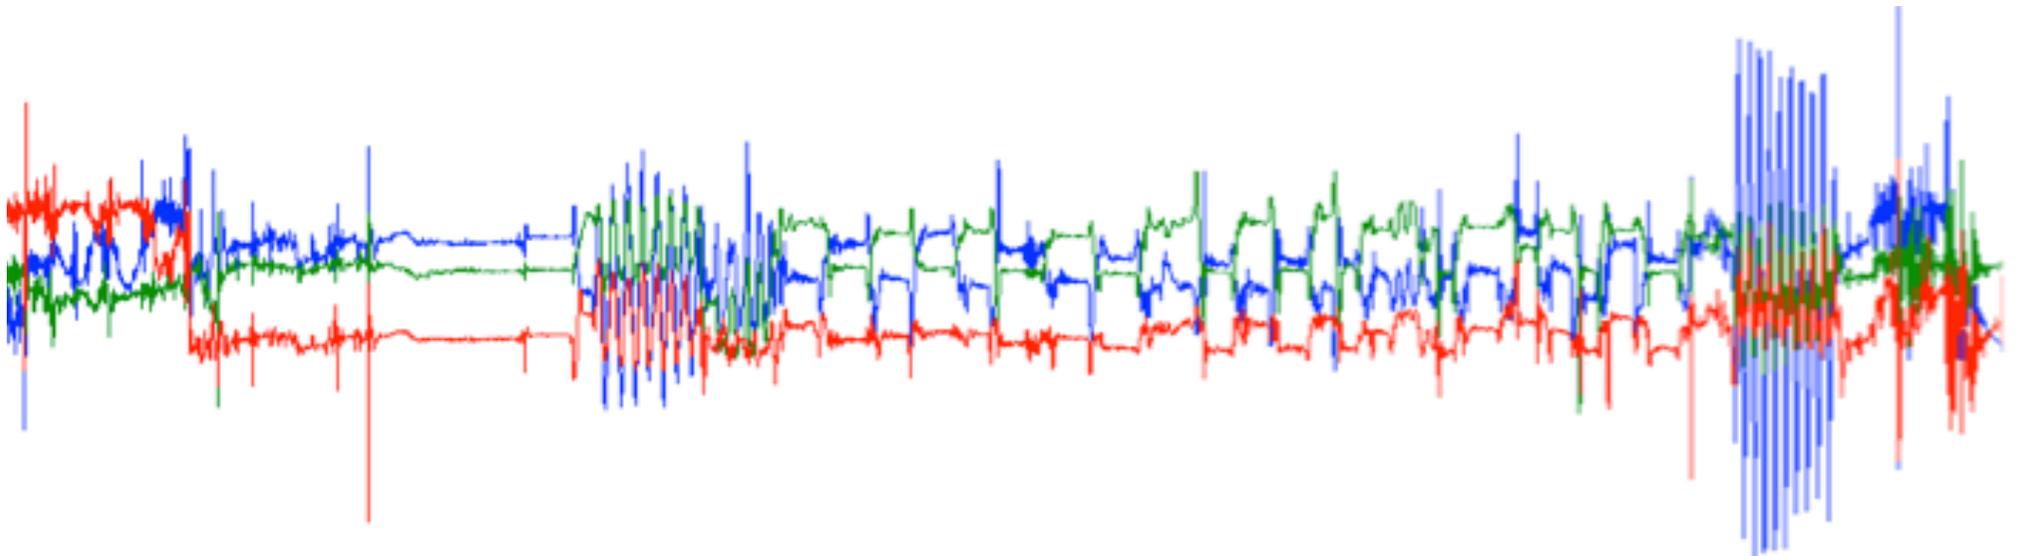
Feature Extraction
Example: Activity Recognition
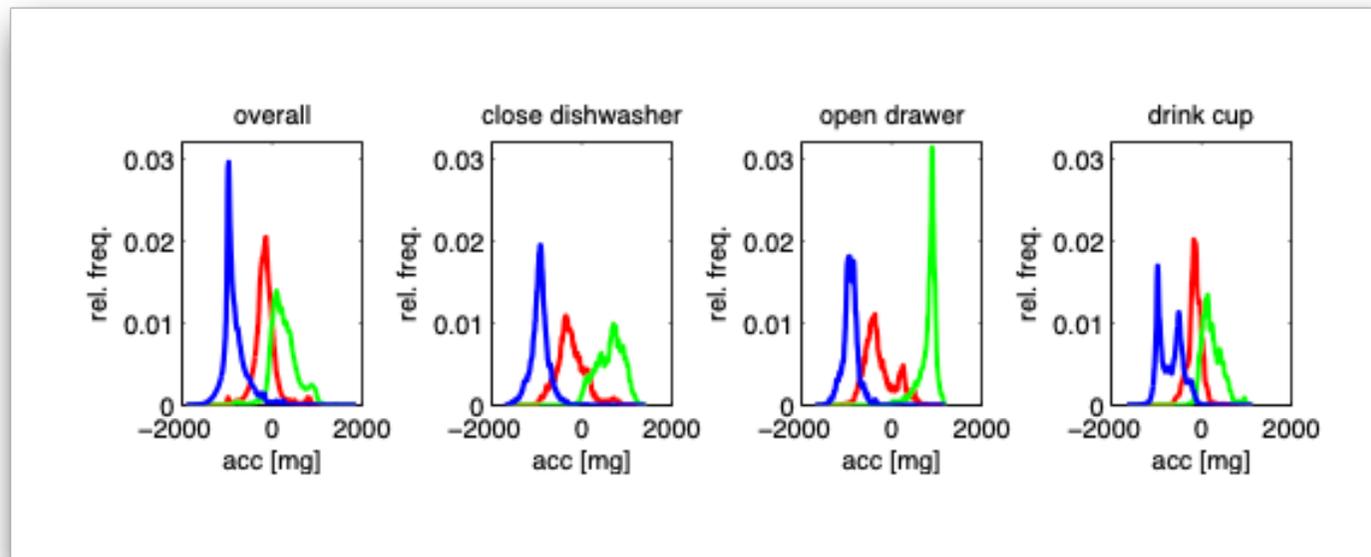
Physical Activity using Accelerometer

- Sensor: accelerometer

- Activities: sitting, standing, waking, running

- Features:
  - Mean (can help distinguish between standing and sitting).
  - Standard deviation
  - Number of peaks (can help distinguish between waking and running).

UNIVERSITY OF CAMBRIDGE
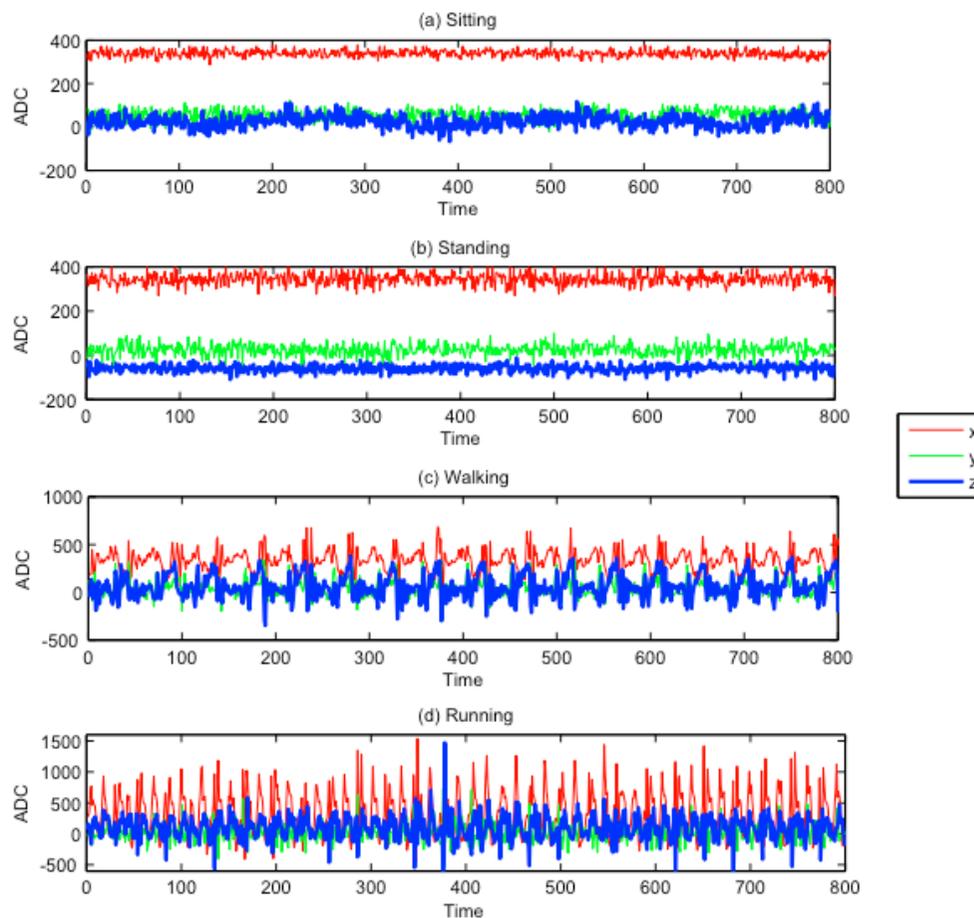
Feature Extraction: Conversation Detection

- FFT (Fast Fourier Transform) of audio (from microphone)

Human voice

Noisy environment
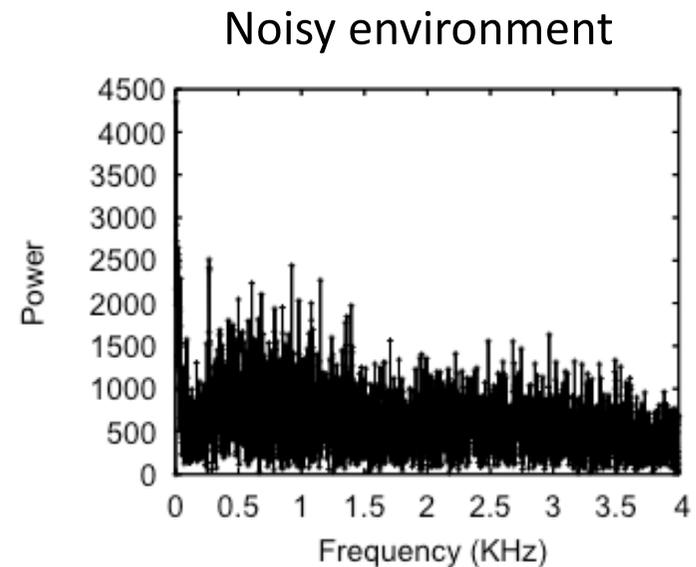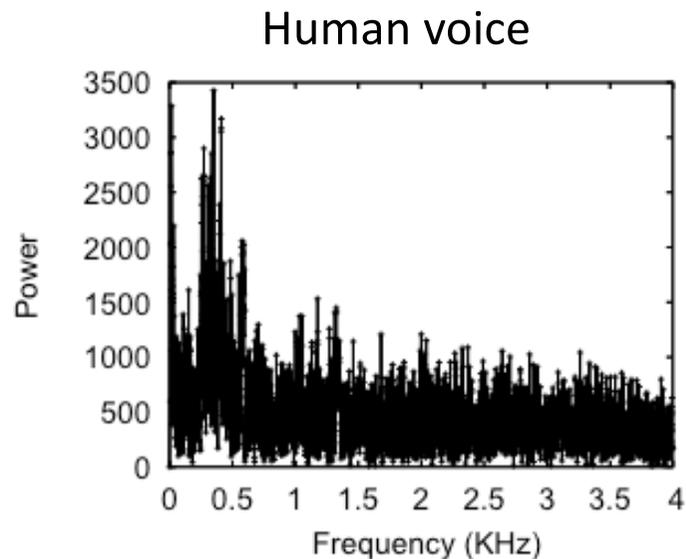
- Sound samples of human voice present most of their energy within the 0-4 KHz spectrum.

Feature Extraction: Conversation Detection

- Selecting as Features the mean and standard deviation of the FFT power
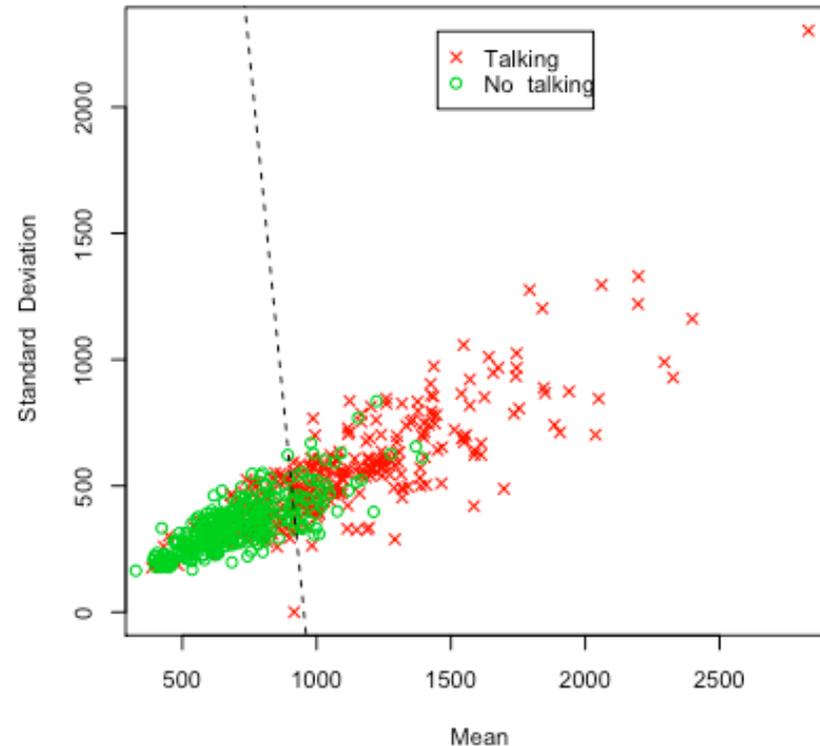
- Using a simple threshold line, could give a relatively accurate detection (with a high number of false positives, however)

Inference

- The process of mapping raw sensor data to meaningful high-level events. Inference Pipeline:

```
Sensing → Feature Extraction → Classification → Detected Event
```

- Designing an Inference Engine:
  – Collecting raw sensor data, typically labelled with ground truth information.
  – Data set should also cover states we are not trying to detect but look similar (e.g. detect *walking* : we need data also for *running* and *standing*).
  – Train the inference engine with the collected data.
  – Applying the inference engine to the target application.

UNIVERSITY OF CAMBRIDGE

Classification

- Feature extraction produces a feature vector.

- The classification matches the feature vector to a pre-defined set of high-level classes.

- The classification engine is typically based on machine-learning techniques and is trained using labelled training data.

- Common classification algorithms include:
  - K Nearest Neighbour.
  - Naive Bayes classifier.
  - Decision Trees.
  - Hidden Markov Models.

Activity Recognition Classification
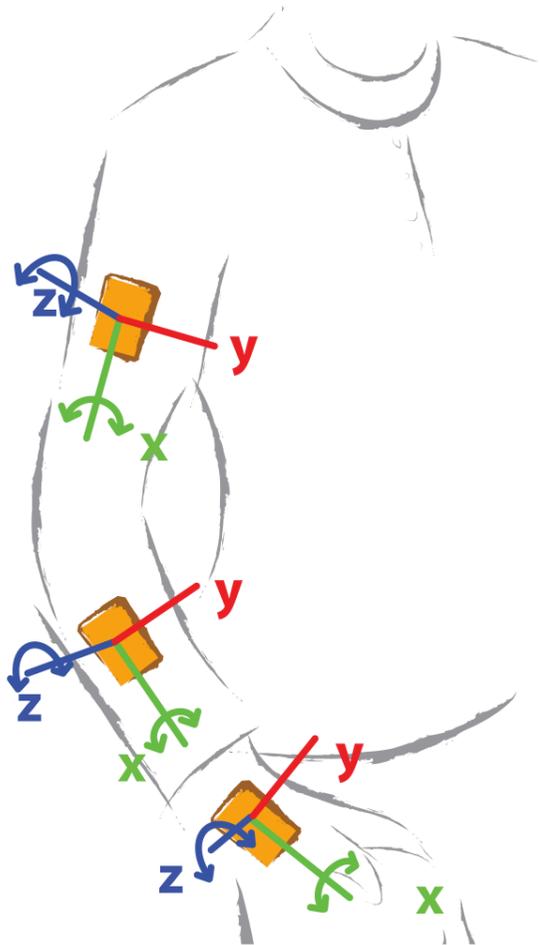
Activities

opening a window
closing a window
watering a plant
turning book pages

drinking from a bottle
cutting with a knife
chopping with a knife
stirring in a bowl

forehand
backhand
and smash

Data Acquisition	Segmentation	Feature Calculation	Modeling & Inference	Classification
ACC_{Wrist}	Fixed-Size (1s) Sliding Window	Mean, Variance	Nearest Class Centroid (NCC)	$\arg\max\limits_{y \in \mathcal{Y}} \ p_i(y\|\mathbf{X}_i, \theta)$
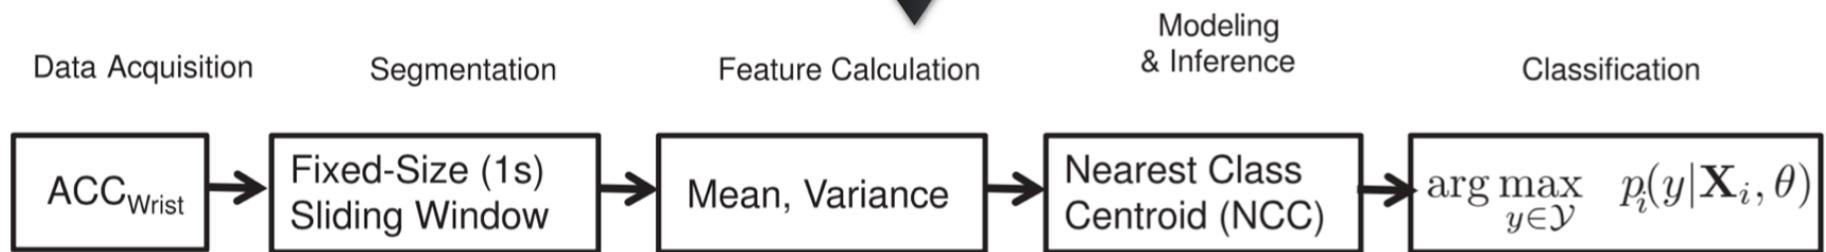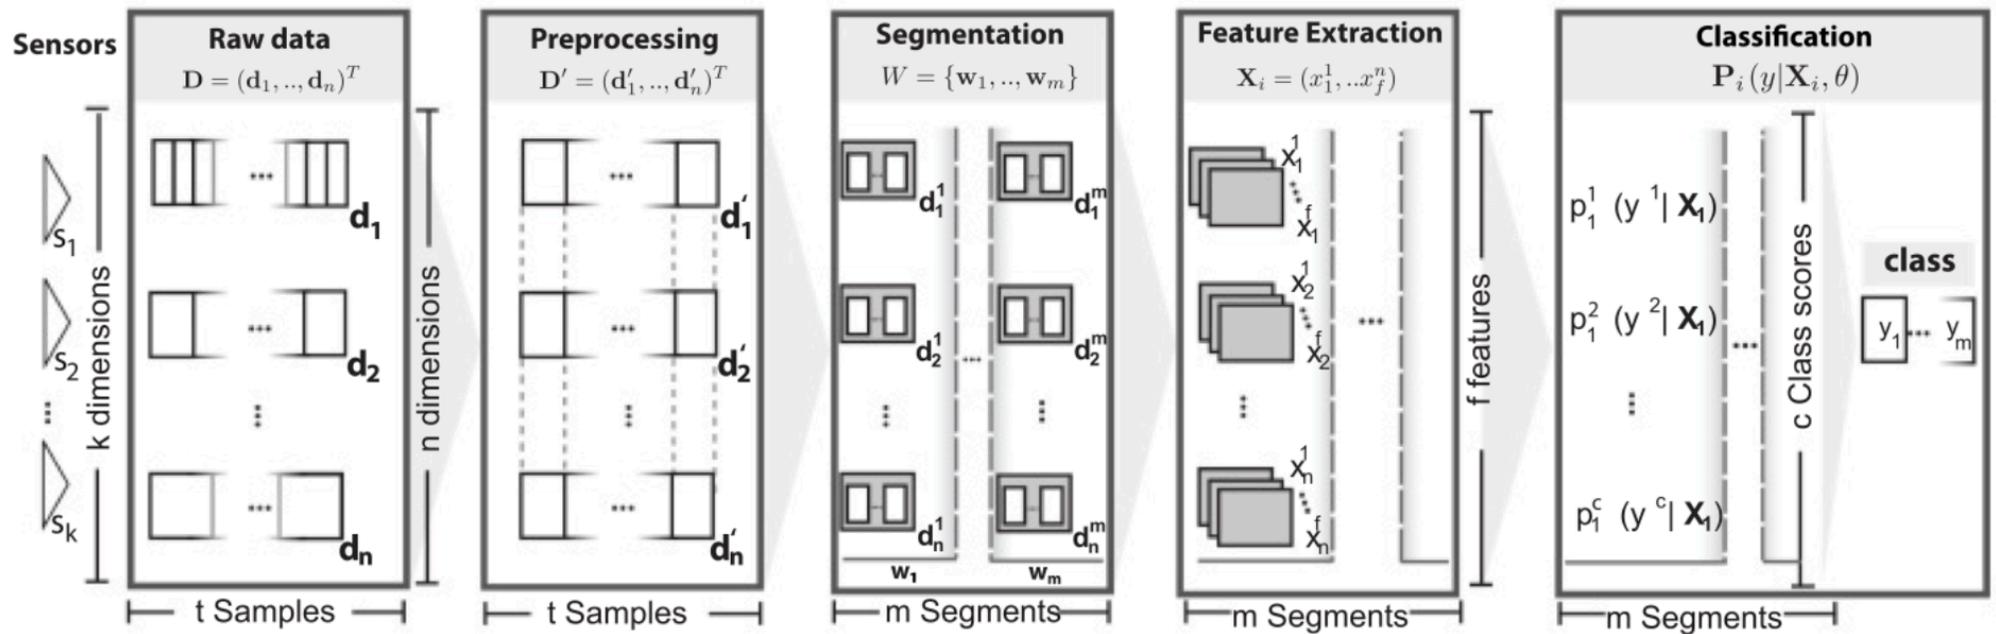
Classification Results:
Person Dependence and Multiple Sensors
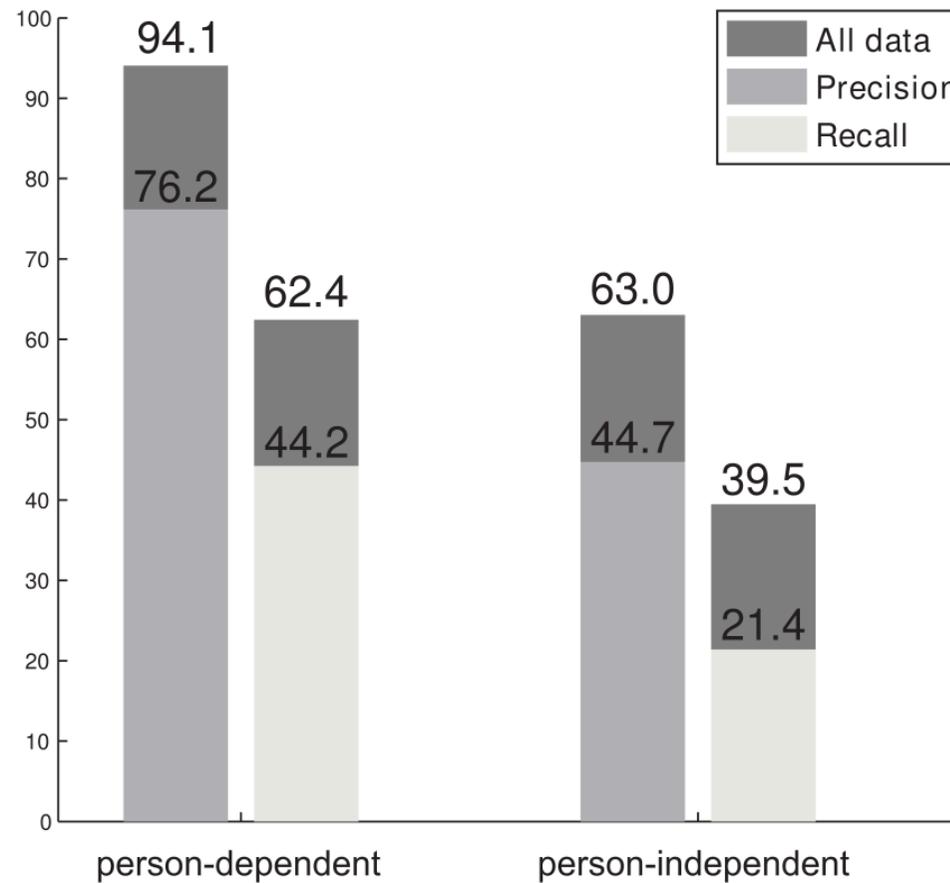
Confusion Matrix on Activities

groundtruth \ classification	NULL	Open window	Drink	Water plant	Close window	Cut	Chop	Stir	Book	Forehand	Backhand	Smash	recall
NULL	24267	216	444	3228	48	24	60	75	45		3		85.42
Open window	3849	1938	453	291	48	12	9		24				29.26
Drink	3984	927	3780	321	3	9							41.89
Water plant	3984	726	774	3735	21	57	15						40.11
Close window	3891	381	1173	945	1533								19.35
Cut	2940		264	450		6585	456		3				61.55
Chop	2895	168	435	153		909	5742		126				55.06
Stir	4947	39	135	42	21	474	561	4392	207				40.60
Book	4560	27	144	951		354	1725	60	6687				46.09
Forehand	3195	330		144	609	9	66		3	969	6	3	18.17
Backhand	3003	207	21		21	3	6	24	33		1302		28.18
Smash	1860	57		78	185		42	45		1567	137	230	5.47
precision	38.29	38.64	49.59	36.13	61.59	78.06	66.14	95.56	93.81	38.21	89.92	98.71	

UNIVERSITY OF CAMBRIDGE

Beyond Features: Deep Learning

- Movement away from hand-crafted experimentally driven features towards models that combine feature and classification phases
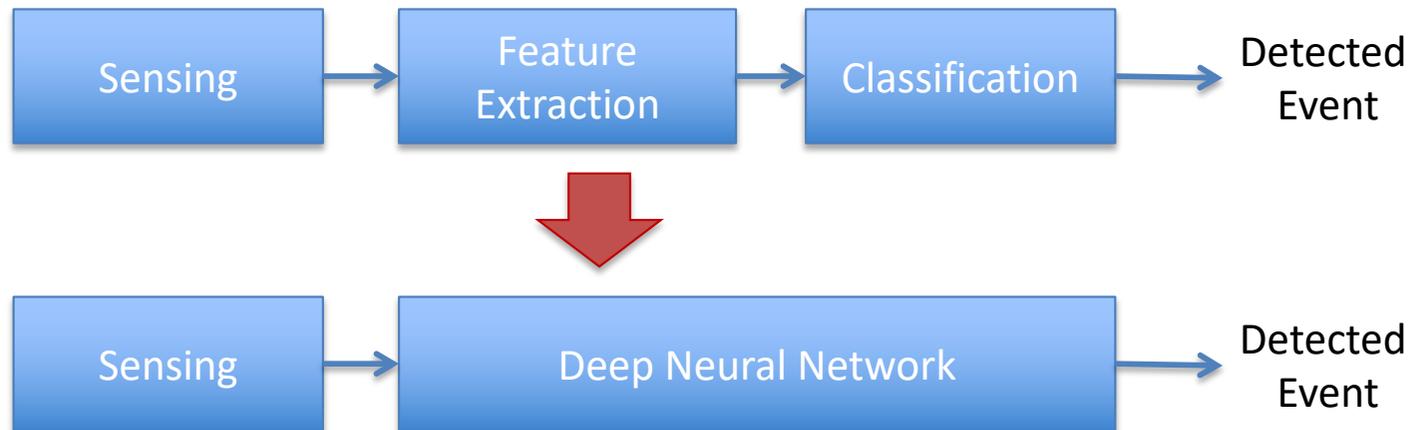
- Paradigm of learning discriminative representations ("feature representation learning") directly from large amounts of relatively raw data ("end-to-end learning")
- Modelling techniques (e.g., training algorithms, network architecture) are less tied to specific domains and tasks.
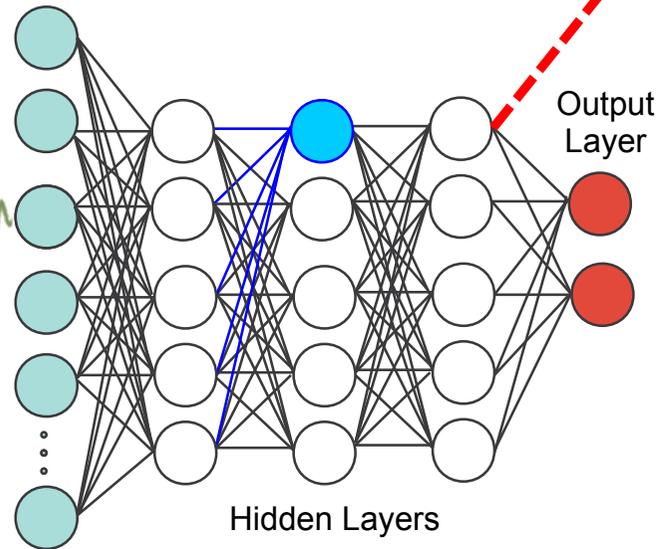
Representative Deep Neural Network

Sensor Data

Input Layer

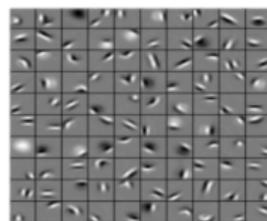$$x_i^{L+1} = \sum_j w_{i,j}^{L+1} x_j^L$$
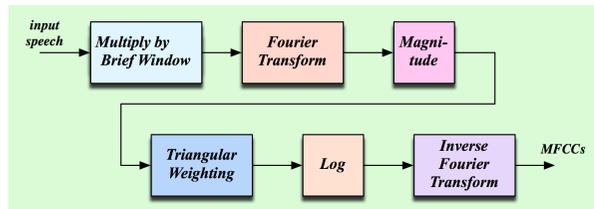
$$\mathcal{A}_i^{L+1} = \frac{1}{1 + \exp(-x_i^{L+1})}$$

Output Layer

"Sad"

Inference

Hidden Layers

- Just one of dozens of types of deep learning that exist (CNNs, RNNs, etc.)

UNIVERSITY OF CAMBRIDGE

Example of
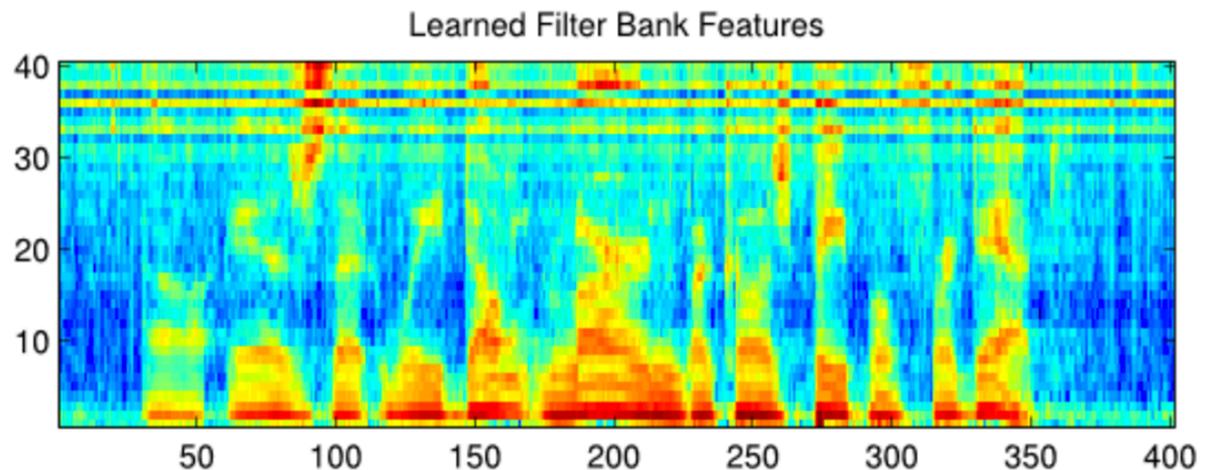# Feature Representation Learning

(1) Mel frequency cepstral coefficients

- Result of decades of research into audio
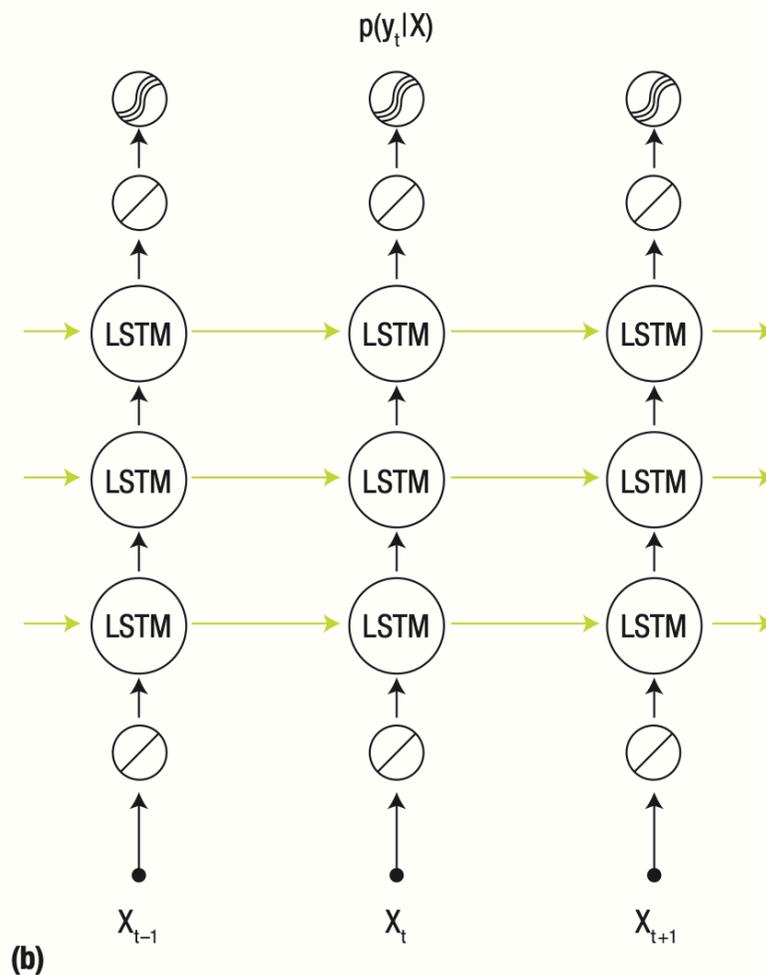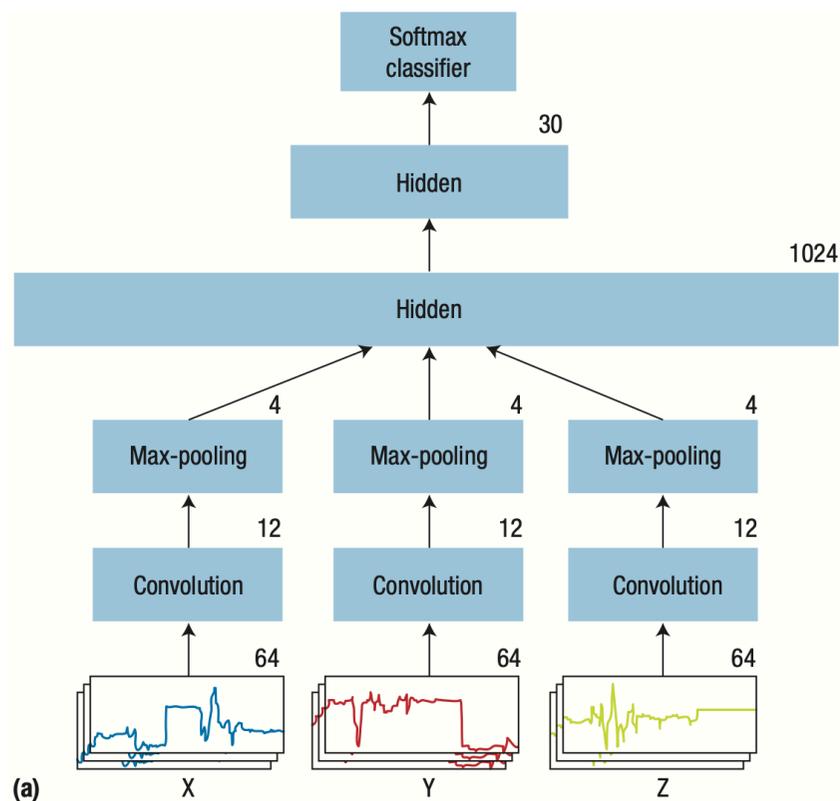- Dominant general purpose audio representation
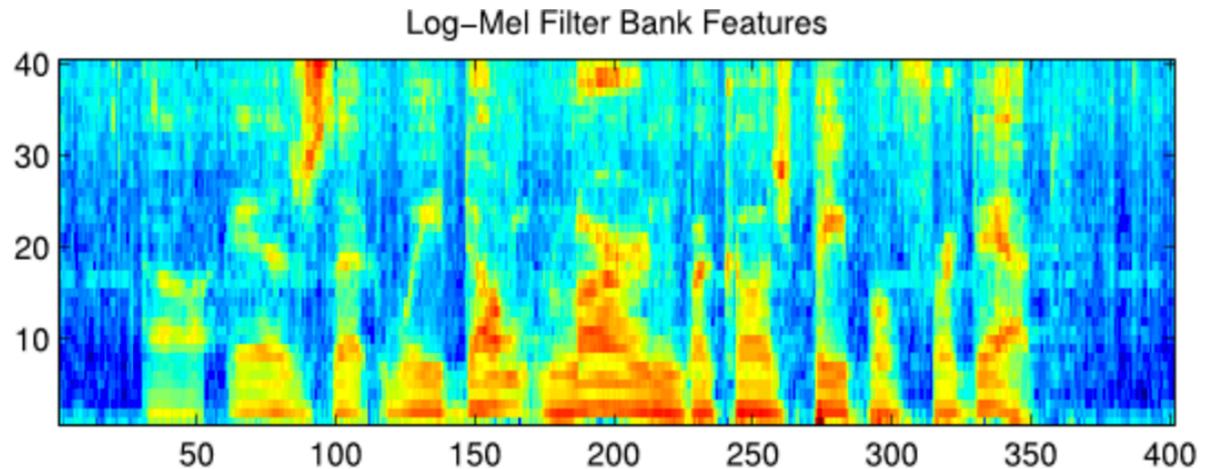
(2) Representation learned directly from data

Log-Mel Filter Bank Features

Learned Filter Bank Features

UNIVERSITY OF
CAMBRIDGE

DNN in Activity Recognition

	PAMAP2	DG	OPP	
Performance	F_m	F_1	F_m	F_w
DNN	0.904	0.633	0.575	0.888
CNN	**0.937**	0.684	0.591	0.894
LSTM-F	0.929	0.673	0.672	0.908
LSTM-S	0.882	**0.760**	0.698	0.912
b-LSTM-S	0.868	0.741	**0.745**	**0.927**
CNN	[Yang *et al.*, 2015]		—	0.851
CNN	[Ordóñez and Roggen, 2016]		0.535	0.883
DeepConvLSTM	[Ordóñez and Roggen, 2016]		0.704	0.917
Delta from median	ΔF_m	ΔF_1	ΔF_m	mean
DNN	0.129	0.149	0.357	0.221
CNN	**0.071**	**0.122**	0.120	**0.104**
LSTM-F	0.10	0.281	0.085	0.156
LSTM-S	0.128	0.297	**0.079**	0.168
b-LSTM-S	0.087	0.221	0.205	0.172

UNIVERSITY OF CAMBRIDGE

[Hammerla et al 2016]

CNN vs LSTM for HAR

CNN+ LSTM for HAR

Leyer 1	Leyer 2	Leyer 3	Leyer 4	Leyer 5	Leyer 6	Leyer 7	Leyer 8
Input leyer	Convolutional leyer	Maxpool leyer	Convolutional leyer	Maxpool leyer	Convolutional leyer	LSTM leyer	Softmax leyer

(c)
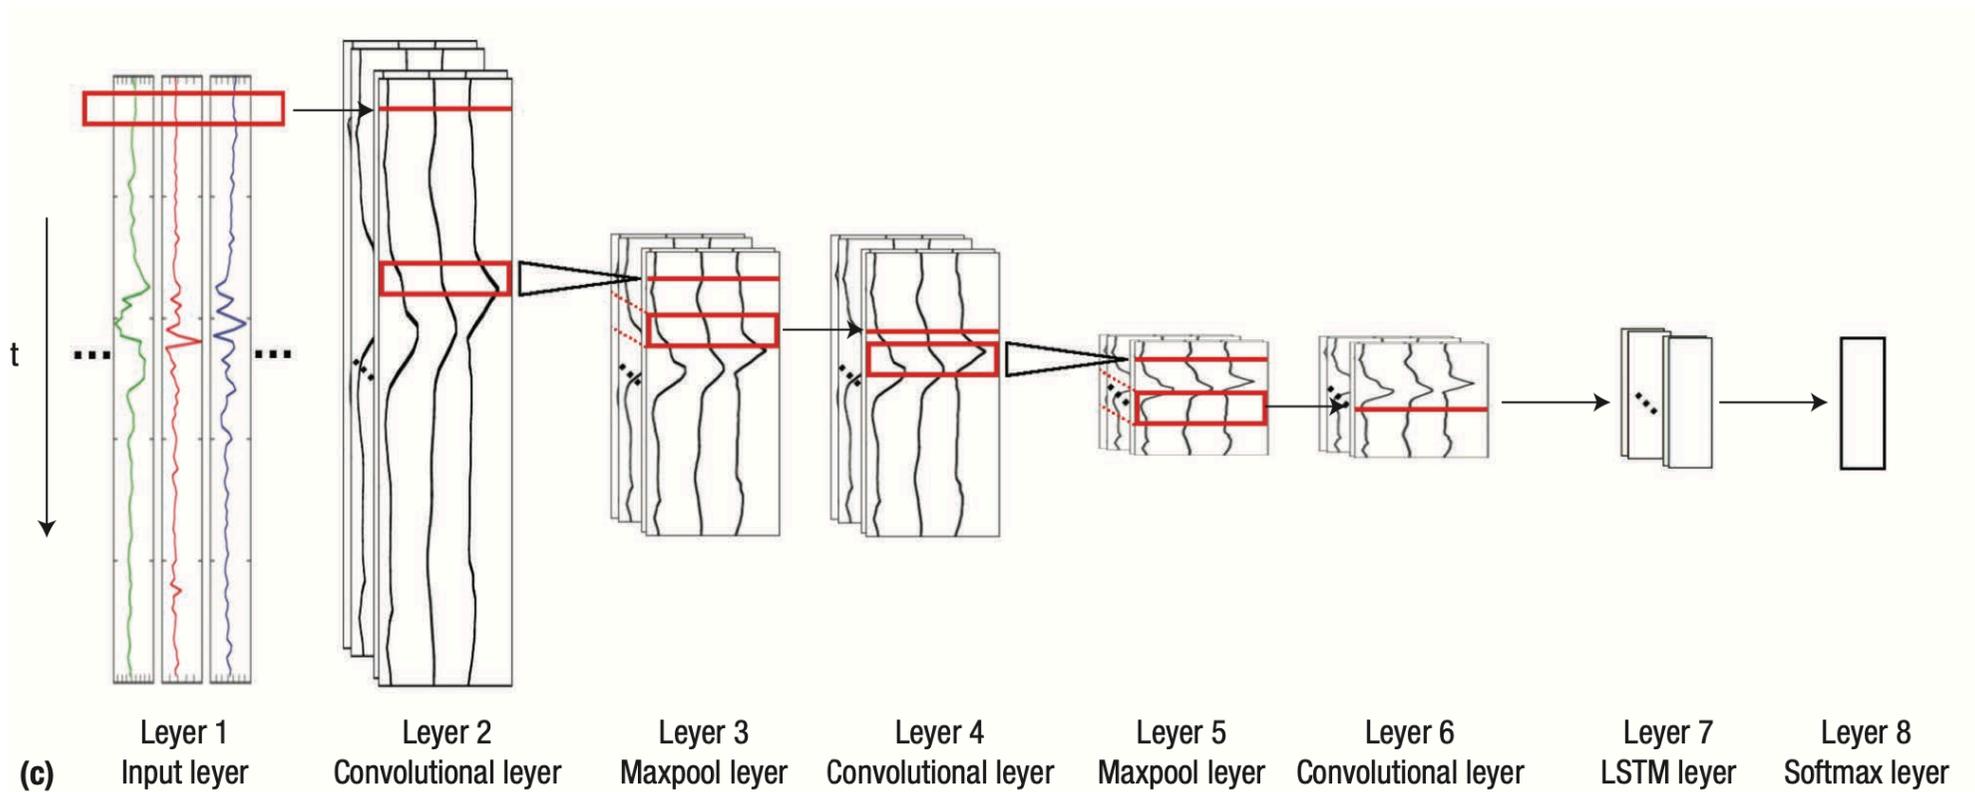
Small Datasets!

- Unlike other domains sensor data lacks large scale labelled datasets
  - Difficult to collect large scale ground truth
- Can lead to overfitting! (DNN with many parameters will memorize small data)
- Solutions:
  - Transfer learning
  - Classifier ensembles

UNIVERSITY OF CAMBRIDGE

References

- N.D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, A. Campbell. A survey of mobile phone sensing. IEEE Computer Magazine. Vol. 48. No 9. September 2010.

- A. Bulling, U. Blanke, and B. Schiele, "A Tutorial on Human Activity Recognition Using Body-worn Inertial Sensors," ACM Computing Surveys, vol. 46, no. 3, Jan. 2013.

- A. Khan, N. Hammerla, S. Mellor, and T. Ploetz, "Optimising sampling rates for accelerometer-based human activity recognition," Pattern Recognition Letters, 2016.

- Nils Y. Hammerla, Shane Halloran, and Thomas Plötz. 2016. Deep, convolutional, and recurrent models for human activity recognition using wearables. In Procs of Int. Joint Conference on Artificial Intelligence (IJCAI'16.) AAAI Press.

- Deep learning for human activity recognition in mobile computing. T Plötz, Y Guan - Computer, 2018

UNIVERSITY OF CAMBRIDGE