Activity Recognition and Time Series Analysis

Kristof Van Laerhoven - Ubiquitous Computing
kvl@eti.uni-siegen.de

Outline

- Definitions of context-aware computing
- Early examples
- From Context to Activity
- Classifying activities:
  - Capturing real activities
  - Sensor coverage
  - Features
  - Classifiers
  - Evaluation of activity recognition
  - Ground truth
  - Shape matching

Outline

- **Definitions of context-aware computing**
- Early examples
- From Context to Activity
- Classifying activities:
  - Capturing real activities
  - Sensor coverage
  - Features
  - Classifiers
  - Evaluation of activity recognition
  - Ground truth
  - Shape matching

Definitions of Context-Aware Computing

- BN **Schilit**, MM **Theimer**. *Disseminating active map information to mobile hosts*. In IEEE network 8 (5), 22-32, **1994**:
  "Context-aware computing is the ability of a mobile user's applications to discover and react to changes in the environment they are situated in."
  - ➡ information: location, nearby persons and objects, changes
- PJ **Brown**. *The Stick-e Document: a Framework for Creating Context-Aware Applications*. Electronic Publishing, 259-272, **1996**:
  "In general the context part of the note can be a combination of elements of the environment that the user's computer knows about."
  - ➡ information: location, adjacency of other objects, critical states (e.g., temperature > 25ºC), computer states (e.g., changes made in files), imaginary companions, time
- others: J **Pascoe**. In ISWC'98, 92-99, 1998; A **Dey**, G Abowd, A Wood. In Knowledge-Based Systems, 11, 3-13, 1999

Definitions of Context-Aware Computing

- A **Dey**, G **Abowd**. *Towards a Better Understanding of Context and Context-Awareness*. In, **2000**:

  "These definitions are too specific. Context is all about the whole situation relevant to an application and its set of users. We cannot enumerate which aspects of all situations are important, as this will change from situation to situation. In some cases, the physical environment may be important, while in others it may be completely immaterial."

  "Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves."

Context as an information source

- **eases interaction** with computers through **common understanding**
  both user and computer *perceive* the same and
  have a similar background *knowledge*

- supports **implicit interaction** between user and computer
  users do not have to directly give lengthy, explicit, and
  unambiguous commands to make sure the computer has
  all the required information

- **avoids false interpretations** from the computer's side,
  problems with hidden assumptions from both the user's and
  the software developer's perspective occur less frequently

Context as an information source

- **eases interaction** with computers through **common understanding**

“switch all lights on”	user is in the living room (location)
“put this in my calendar”	selected meeting details in e-mail (app)
“when will I get home?”	bicycling, 7km, 9:12 (activity, position, time)

- supports **implicit interaction** between user and computer

phone’s dial key pressed	alarm is going off (app)
user touches smartwatch	user just arrived in a new city (location)
car changes lane	another car is already there! (nearby cars)

- **avoids false interpretations** from the computer’s side

user’s heart rate is 130!	user was jogging for 34 minutes (activity)
user is lying down!	bedtime, bedroom (routine, time, location)
user away from unlocked car!	car key is still in the ignition (objects)

Outline

- Definitions of context-aware computing
- **Early examples**
- From Context to Activity
- Classifying activities:
  - Capturing real activities
  - Sensor coverage
  - Features
  - Classifiers
  - Evaluation of activity recognition
  - Ground truth
  - Shape matching

Early context-aware application examples

- Active Badge system (Olivetti/AT&T, Cambridge, **1997-2001**)
  - ultrasound location of office workers through 720 receivers to cover an area of around 1000m^2 on three floors
  - teleports the user's screen output to nearby computers
  - phone forwarding to the nearest phone

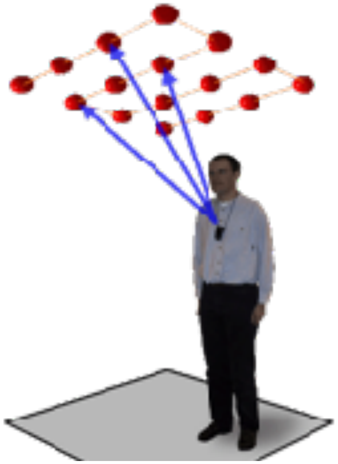

Early context-aware application examples

- Stick-E Notes (Pascoe, **1997**)
  - documents pop up in certain contexts only
  - system contains authoring tool for describing the context

```
<note>
        <required>
        <at> (12.1,42)..(12.3,42.2)
        <facing> 150..207
        <during> December
<body>
        The large floodlit building
        at the bottom of the hill
        is the cathedral.
```

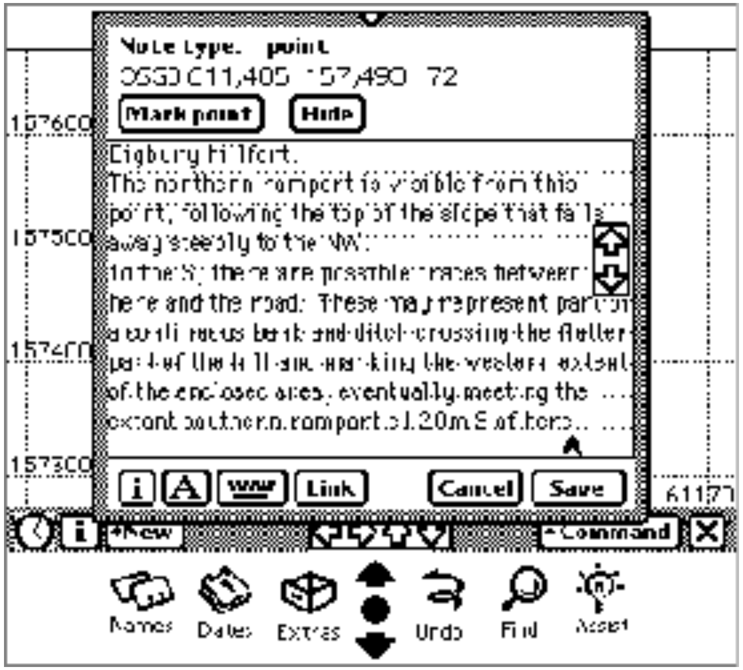

Early context-aware application examples

- Cooltown (Hewlett-Packard, **2000**)
  - people, places, and objects are given a web presence through URLs
  - HP WebBus example:
    - buses equipped with GPS and web server
    - in bus: show location, nearby stops and points of interest
    - waiting for bus: show location and expected arrival time

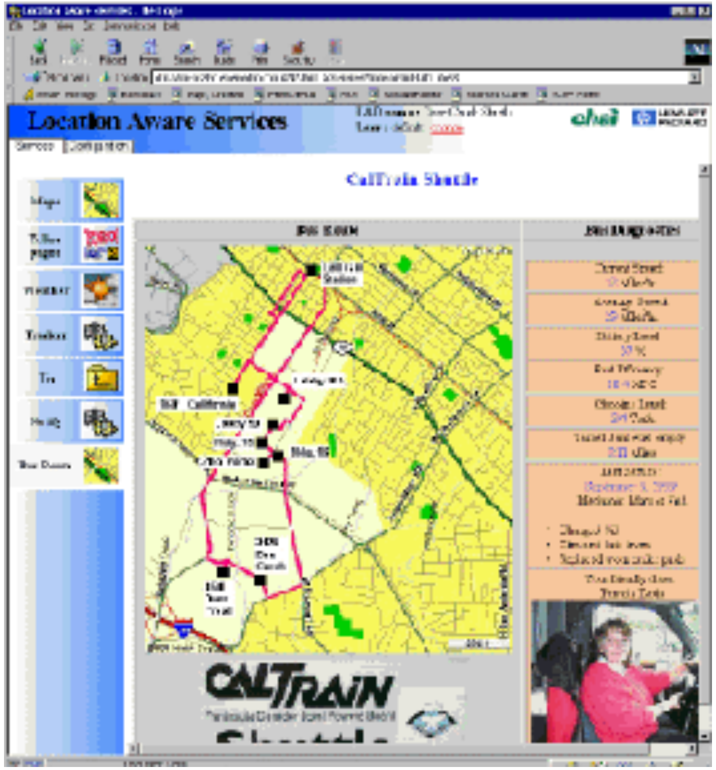

Early context-aware application examples

- Technology for Enabling Awareness (TEA, **1999**)
  - sensors are added to PDAs and phones

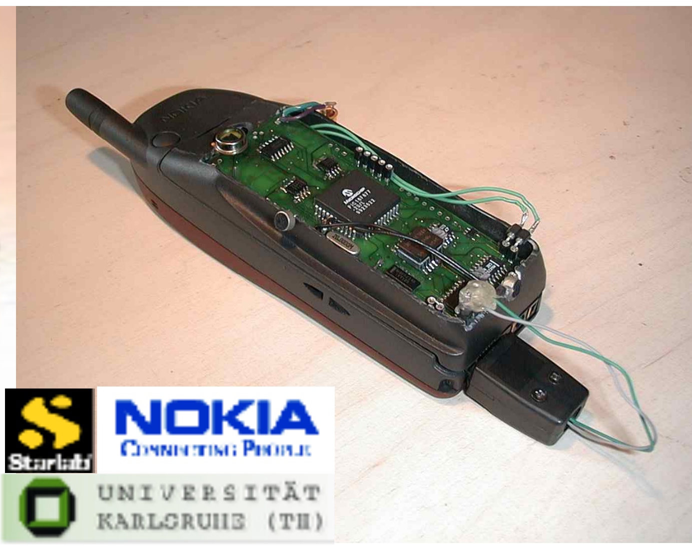

Sensors hidden in battery:

- 2D accelerometer
- 2 photodiodes
- 2 microphones
- capacitive touch
- temperature sensor

→ Serial communication to Nokia 6110 phone → On-phone Automatic Profile Switching

Early context-aware application examples

- Technology for Enabling Awareness (TEA, **1999**)
  - sensors are added to PDAs and phones
  - context aware phone sets profiles automatically
    in user's hand (capacitive touch sensor) -> vibrate when called
    on desk (accelerometer, light sensors) -> low-volume ringing
    in user's pocket (accelerometer, light sensors) -> loud ringing

  - less-critical context awareness:
    user explicitly presses button to change profile,
    list is sorted for most likely profiles according to sensors

A Schmidt and K Van Laerhoven. IEEE Personal Communications, 8(4): 66-71, 08/2001,2001

Outline

- Definitions of context-aware computing
- Early examples
- From Context to Activity
- Classifying activities:
  - Capturing real activities
  - Sensor coverage
  - Features
  - Classifiers
  - Evaluation of activity recognition
  - Ground truth
  - Shape matching

Activity Recognition

- recognize the **actions** and **goals** of one or more agents
- from a series of observations
  - the agents' actions
  - environmental conditions

- Agents: can be people, but also robots or other entities that can perform actions
- "Agents' actions", for instance:
  - physical motions (gestures of hands, arms, legs, etc.)
  - Bio-physical signals (heart rate, muscle activity, etc.)
- "Environmental conditions", for instance:
  - Location as an indication of activity (kitchen, classroom, …)
  - Activities requiring type of environment (swimming ➔ high humidity)

Activity Recognition

- recognize the **actions** and **goals** of one or more agents
- from a series of observations

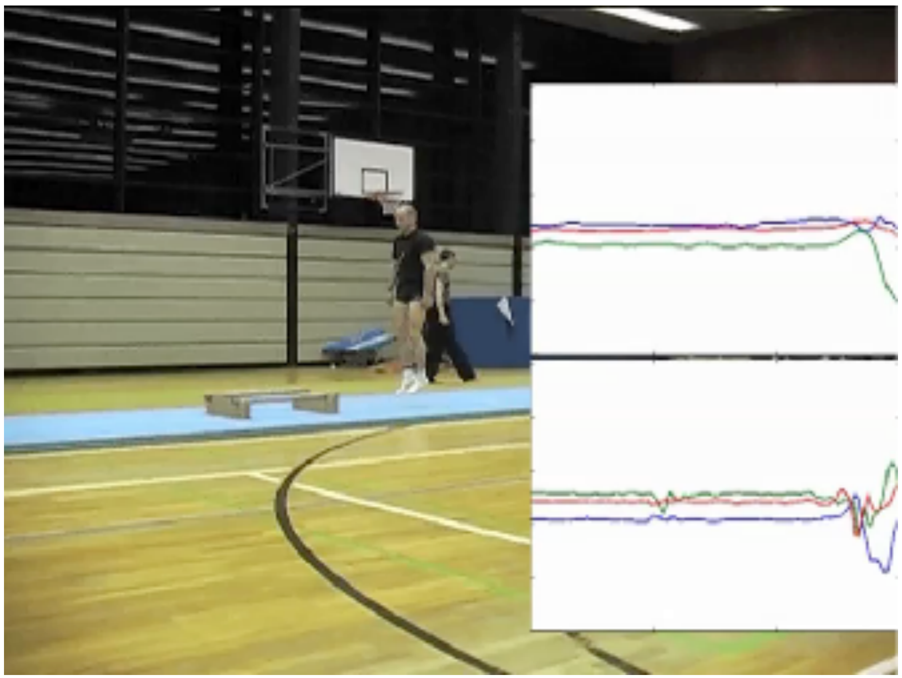


Activity Recognition

- recognize the **actions** and **goals** of one or more agents
- from a series of observations

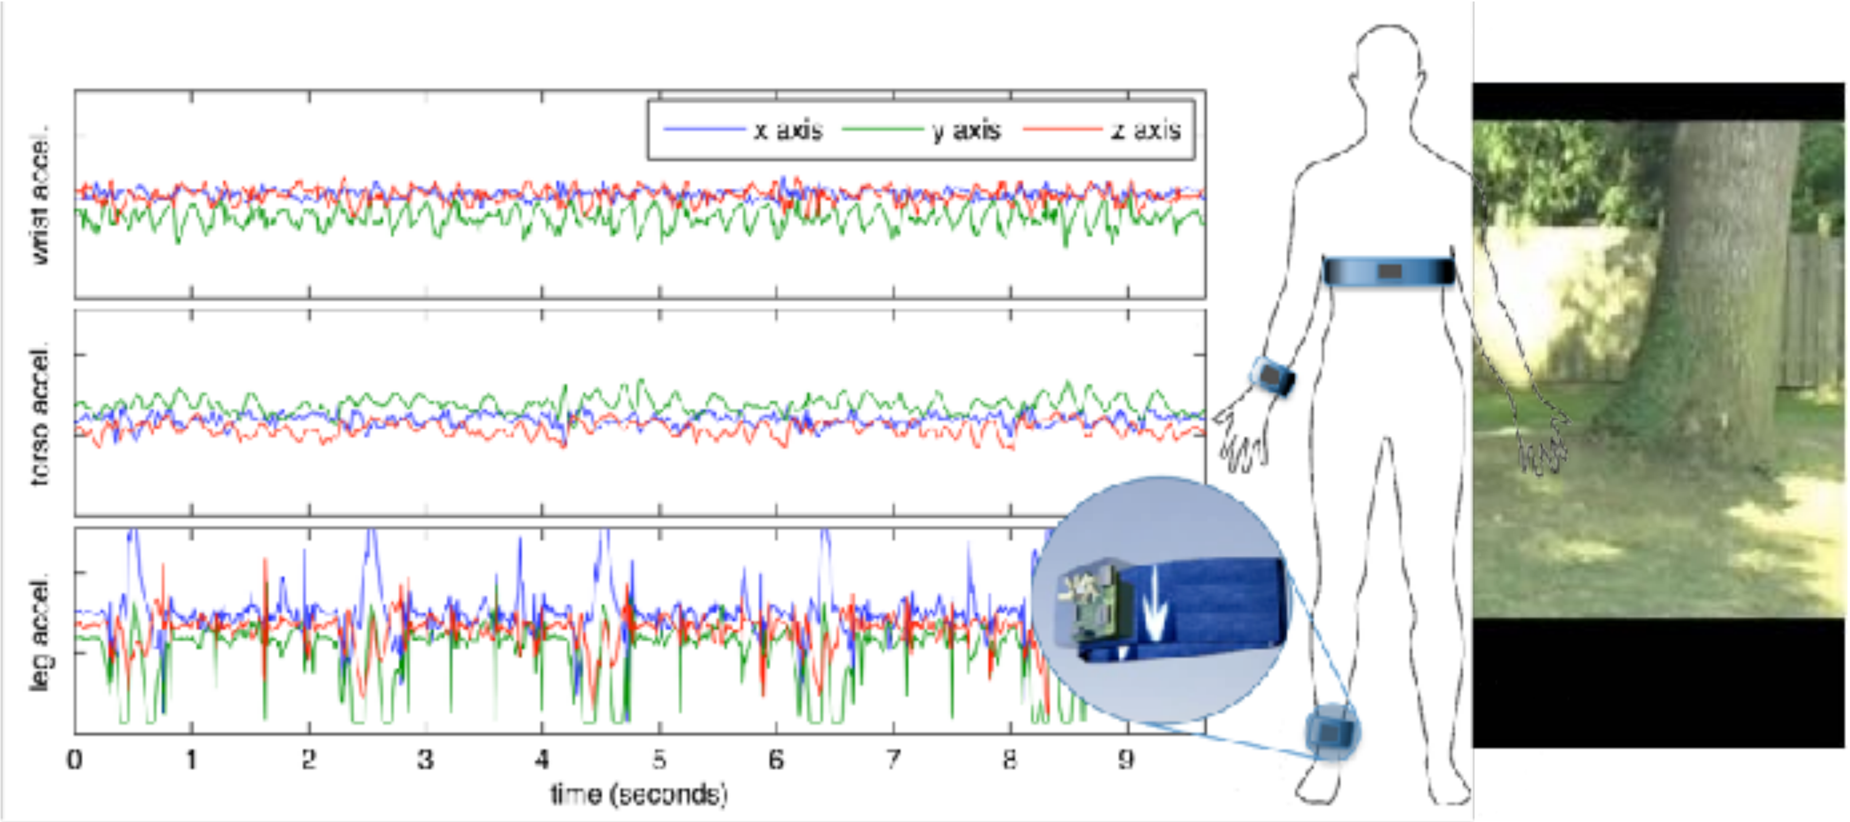

Activity Recognition: instrumented environments

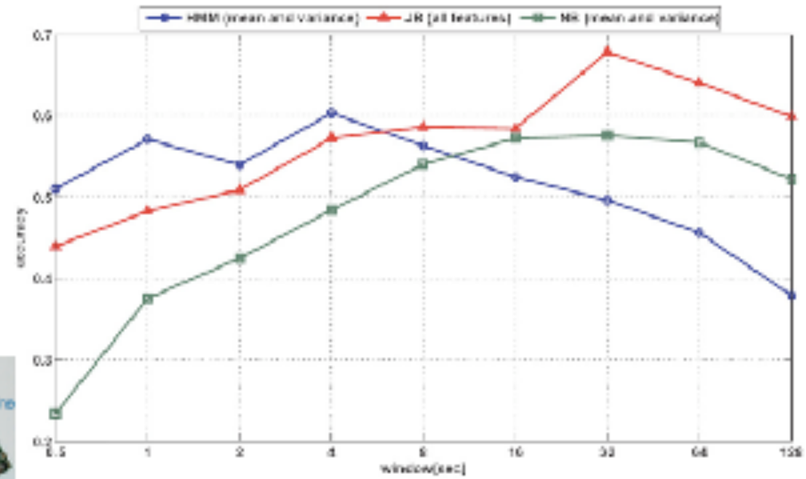

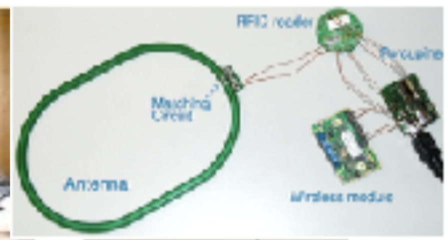

Activity Recognition: instrumented environments

Hogg. **Model-based vision: a program to see a walking person**, 1983

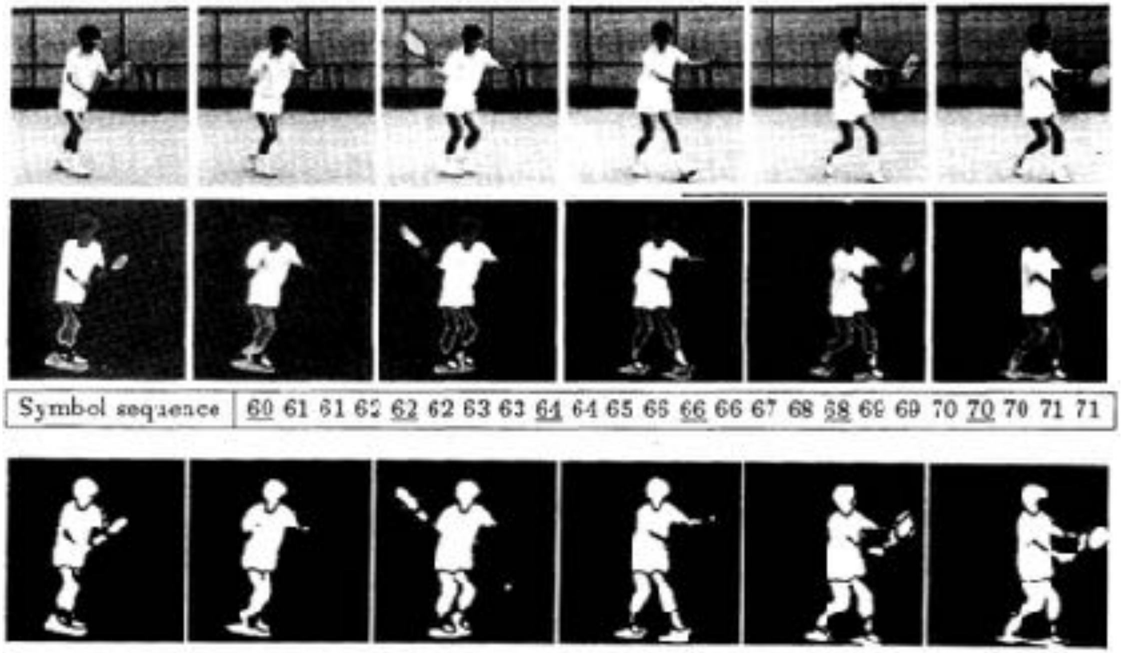

Yamato et al. **Recognizing Human Action in Time-Sequential Images using Hidden Markov Model**, CVPR 1992

Activity Recognition: instrumented environments

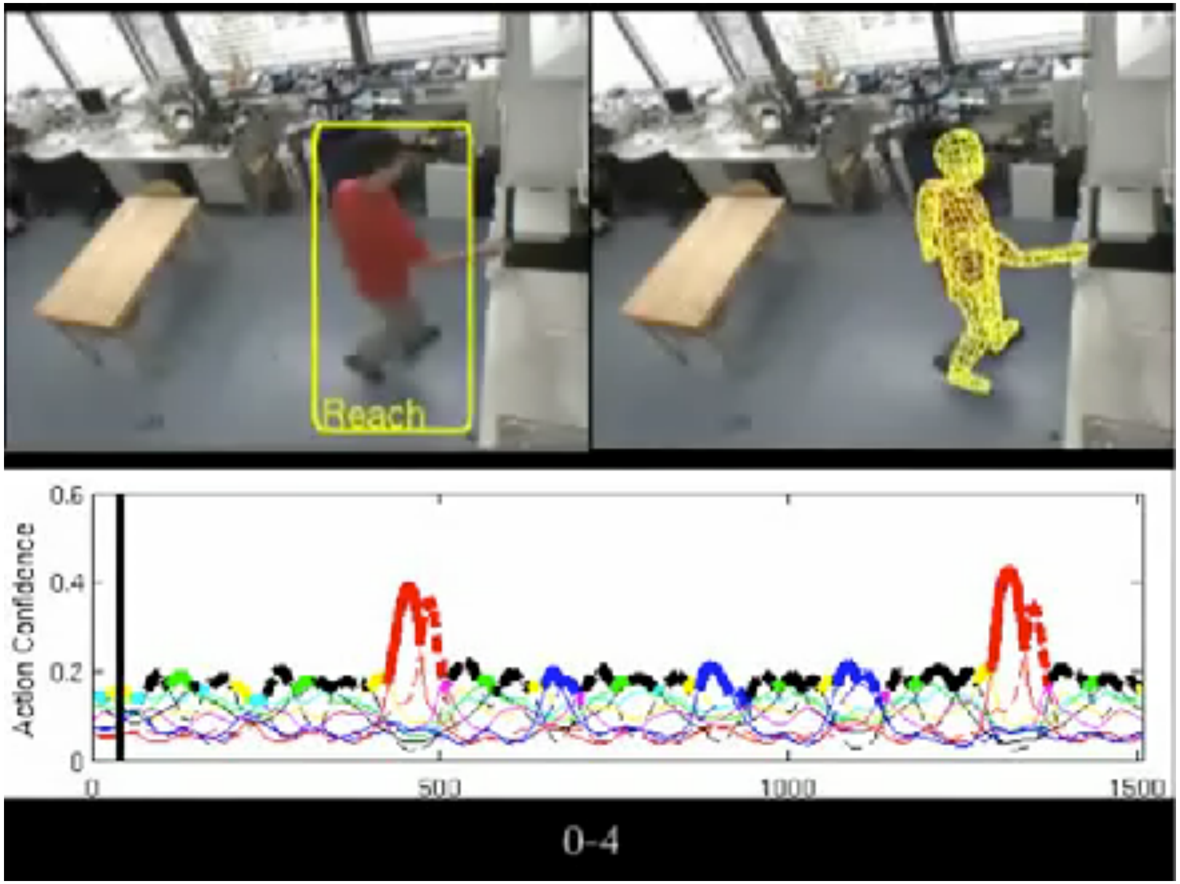

Gall et al. **2D Action Recognition Serves 3D Human Pose Estimation**, ECCV 2010

Ubiquitous Computing

UNIVERSITÄT SIEGEN

Activity Recognition: instrumented apparel

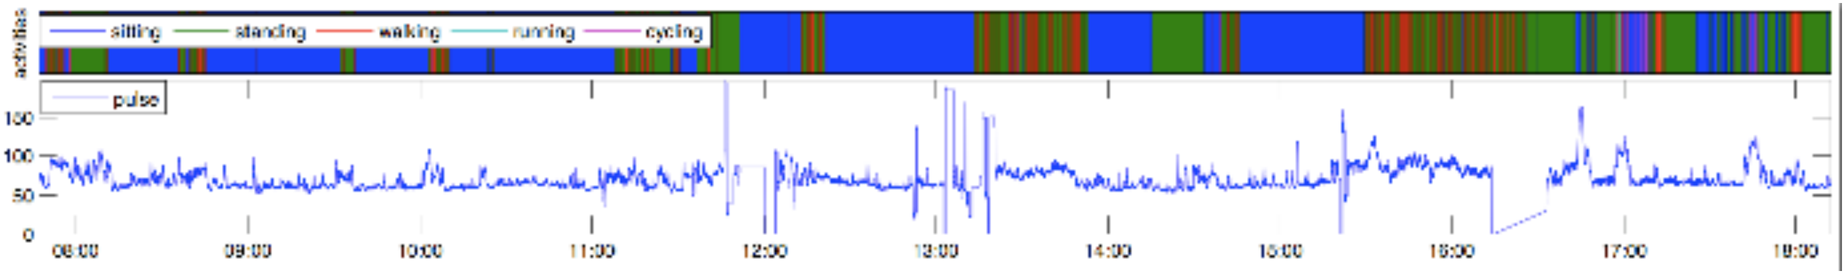

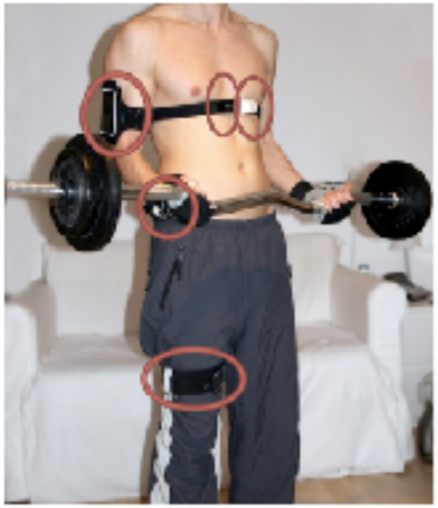

	Exercise	Posture	Type
1	Walking	-	Cardio
2	Running		
3	Cycling		
4	Rowing		
5	Elliptical trainer		
6	Wide grip lat pulldown	Sitting	Back
7	Barbell rear delt row	Standing	
8	Hyperextensions	Standing	
9	Barbell bench press	Lying	Chest
10	Butterfly	Sitting	
11	Front barbell raise	Standing	Shoulders
12	Dumbell lateral raise	Sitting	
13	Barbell curl	Standing	Arms
14	Cable triceps extensions		
15	Barbell squat	Standing	Legs
16	Table top crunch	Lying	Abs

Activity Recognition: instrumented participants

Activity Recognition

- recognize the **actions** and **goals** of one or more agents
- from a series of observations

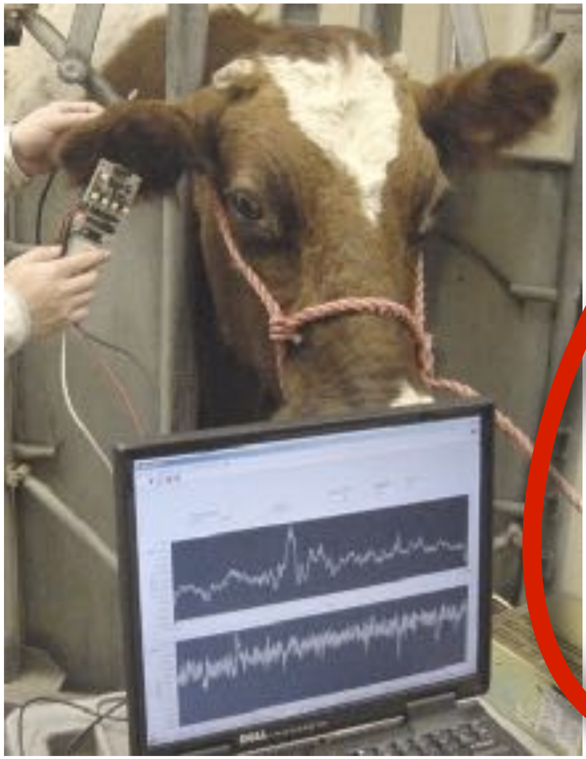

Why Activity Recognition?

- **Elderly Care.** Monitoring Activities of Daily Living (ADLs) to estimate quality of self-care
- **Fitness and Well-Being.** Fighting obesity, coaching toward more active life.
- **Psychiatry.** Correlation of activities with moods, mood swings, manic depressions.
- **Security / Workflow Monitoring.** Tracking maintenance staff, crowds, managing accountability
- **Groupware.** Sharing activity information in groups / over social networks.
- **Memory Support.** Managing diaries, auto-filling journals for later accounting.

Activity Recognition: How well does it work?

"How well does our activity recognition system work?"

- **Measures and benchmarks:**
  - How long does the system last?

    -> battery, power efficiency
  - How big is it? How deployable / wearable is it?

    -> weight, size, minimization of hardware, comfort rating
- **How often is the system right?**
  - Is it better than system X?
  - And if yes, how much?

From Context to Activity

- Location and activities: office, arch, up/downstairs, ping pong room

  Golding and Lesh, Indoor Navigation Using a Diverse Set of Cheap, Wearable Sensors, ISWC 1999.

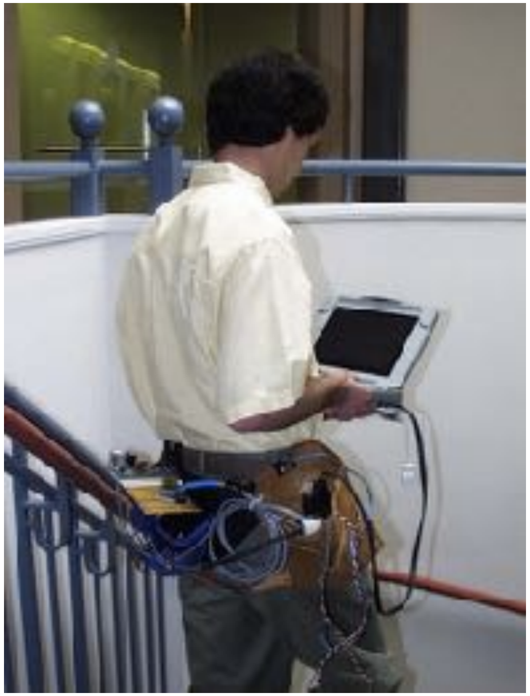

3D Accelerometers, 3D Magneto-meters, Light, Temperature

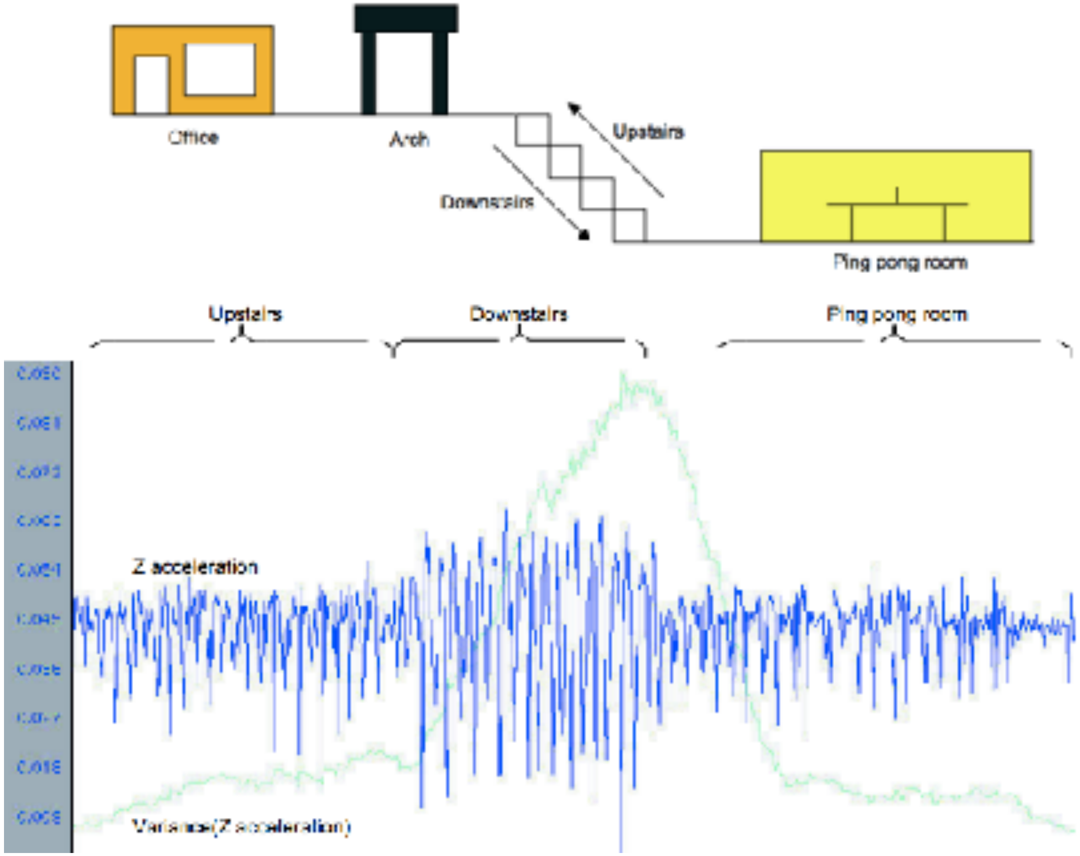

From Context to Activity

- (Basic) activities: sitting, standing, walking, running, bicycling
  Kristof Van Laerhoven & Ozan Cakmakci, “What shall we teach our pants?”, ISWC 2000.

2D Accelerometers

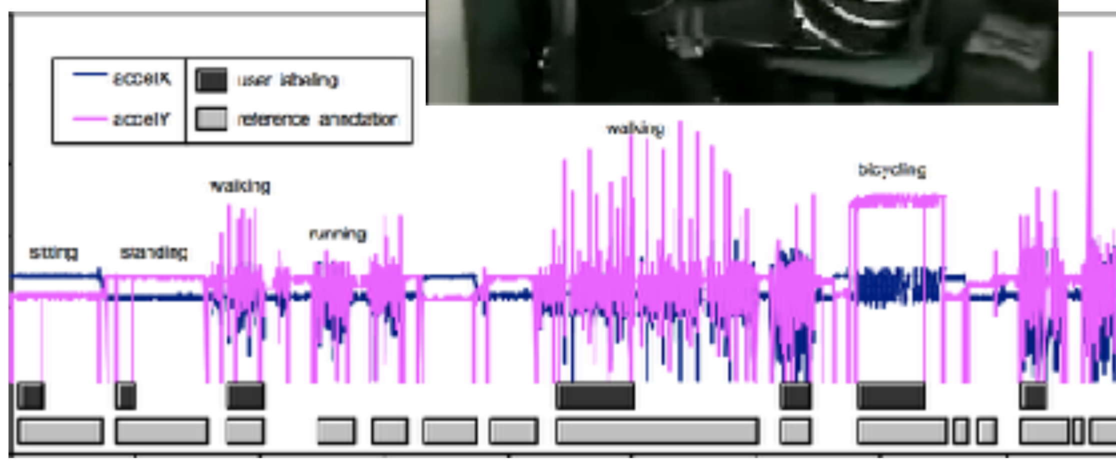

From Context to Activity

- Early research efforts focused on hardware prototypes that integrated and combined new sensors:
  - capacitive sensing around the neck

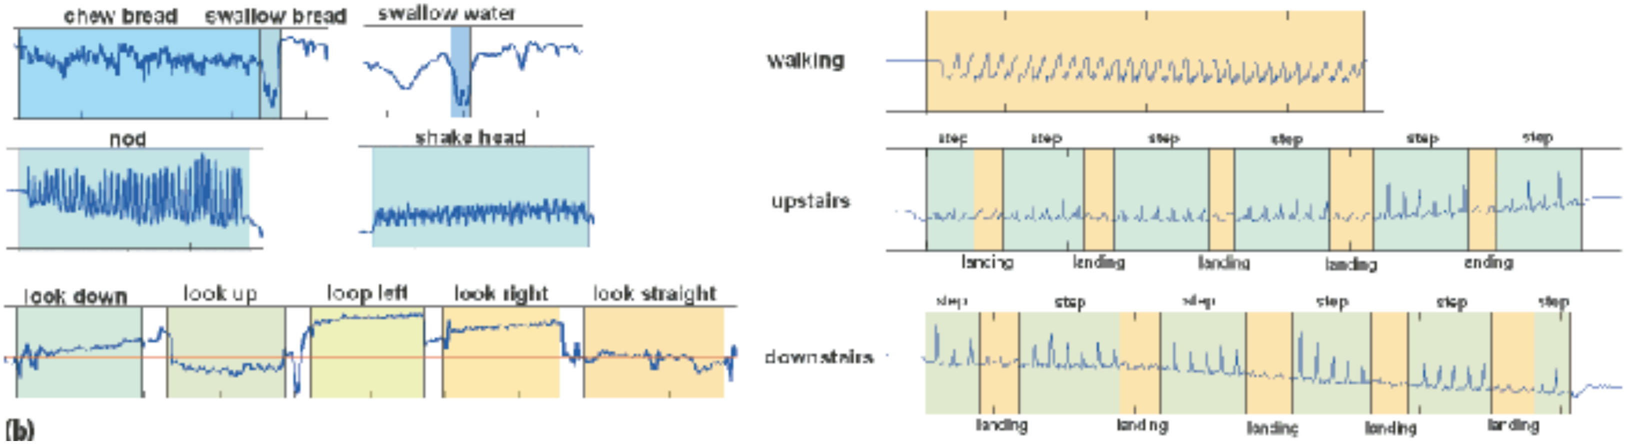

Oliver Amft, Paul Lukowicz, Jingyuan Cheng, Daniel Roggen, "On-Body Sensing: From Gesture-Based Input to Activity-Driven Interaction", Computer, vol. 43, no. , pp. 92-96, October 2010

From Context to Activity

- Early research efforts focused on hardware prototypes that integrated and combined new sensors:
  - capacitive sensing around the neck
  - sensing pressure changes in a house

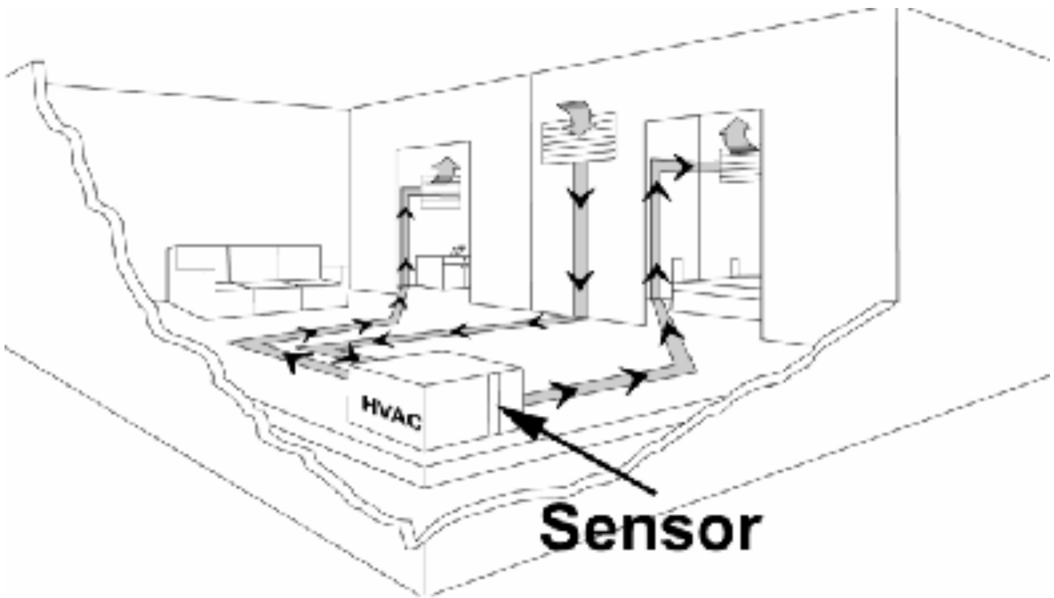

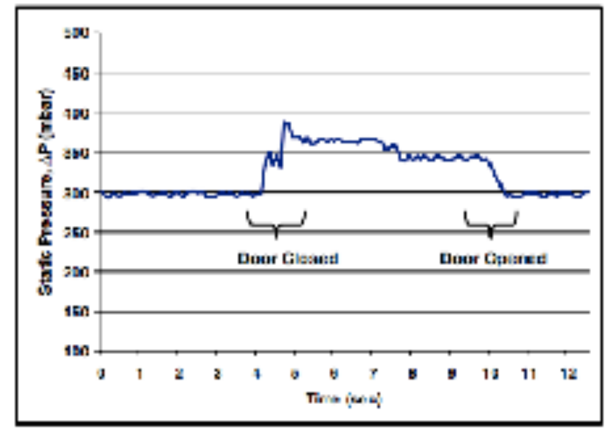

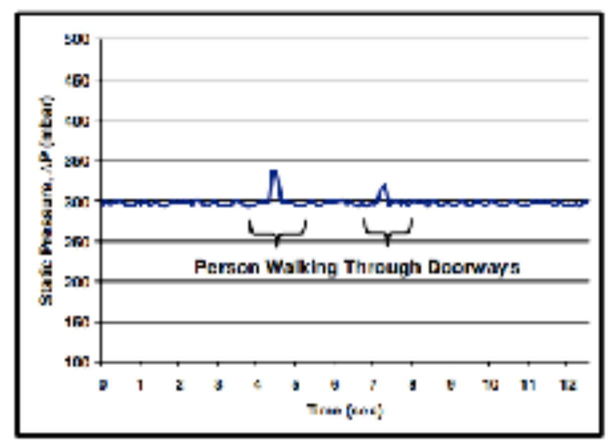

Shwetak Patel, Matthew S. Reynolds, Gregory D. Abowd, 2008

From Context to Activity

- Early research efforts focused on hardware prototypes that integrated and combined new sensors:
  - capacitive sensing around the neck
  - sensing pressure changes in a house
  - reading detection from EOG

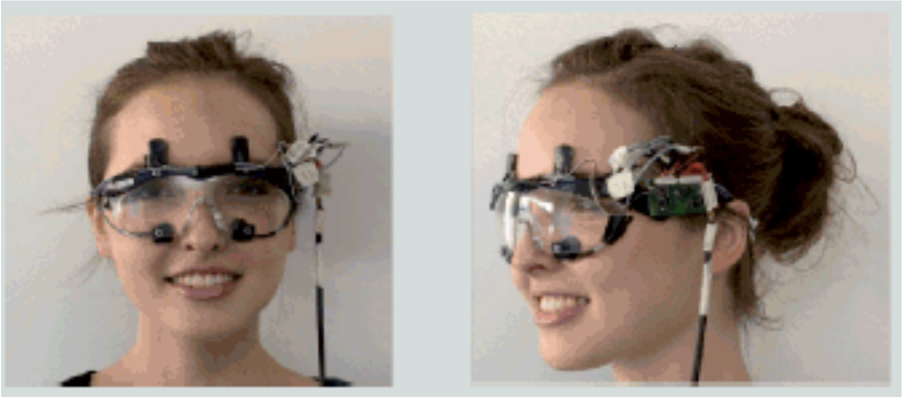

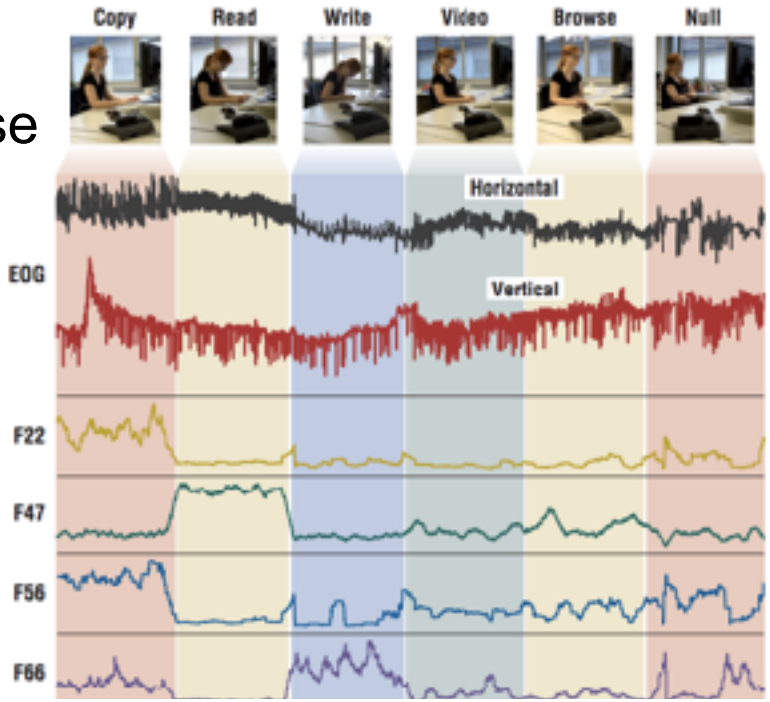

Gerhard Troster, Andreas Bulling, Daniel Roggen, "What's in the Eyes for Context-Awareness?", IEEE Pervasive Computing, vol. 10, no. , pp. 48-57, April-June 2011

From Context to Activity

- Early research efforts focused on hardware prototypes that integrated and combined new sensors:
  - capacitive sensing around the neck
  - sensing pressure changes in a house
  - reading detection from EOG
  - …

  => focus on **feasibility**
  less focus on perfect recognition, usability, durability

From Context to Activity

- Early research efforts focused on hardware prototypes that integrated and combined new sensors
- Followed up by more research in data analysis and classifiers

=> focus on **accuracy**
less focus on perfect usability, durability

From Context to Activity

- Early research efforts focused on hardware prototypes that integrated and combined new sensors
- Followed up by more research in data analysis and classifiers
- Deployments with actual users, over longer stretches of time

=> focus on **usability**, **durability, applicability**

Outline

- Definitions of context-aware computing
- Early examples
- From Context to Activity
- **Classifying activities:**
  - Capturing real activities
  - Sensor coverage
  - Features
  - Classifiers
  - Evaluation of activity recognition
  - Ground truth
  - Shape matching

Classifying Activities: Capturing *real* activities

Sleeping outside the box, Rattenborg et al., Biology, 2008:
“Sloths [...] sleep in captivity **15.85h**, but **9.63h** in the tropical rainforest”

Albanian	përtaci	
Basque	nagikeria	
Belarusian	ленасць	
Bosnian	lijenost	
Bulgarian	ленивец	
Catalan	mandra	
Croatian	lijenost	“lazy one”
Czech	lenochod	
Danish	dovendyr	
Dutch	luiaard	
Estonian	laiskus	
Finnish	laiskiainen	“lazy animal”
French	la paresse	
Galician	preguiza	
German	Faultier	
Greek	νωθρότητα	
Hungarian	lajhár	“lazyness”
Icelandic	letidýr	
Irish	Sloth	“deadbeat”
Italian	bradipo	
Latvian	slinkums	
Lithuanian	tingumas	
Macedonian	мрзеливост	
Maltese	sloth	“lubberliness”
Norwegian	dovendyr	
Polish	lenistwo	
Portuguese	preguiça	
Romanian	lene	“sluggish”
Russian	леность	
Serbian	лењост	
Slovak	lenivosť	
Slovenian	lenivec	“slacker”
Spanish	perezoso	
Swedish	lättja	
Ukrainian	лінощі	
Welsh	sloth	
Yiddish	סלאָט	

Classifying Activities: Capturing *real* activities

White coat hypertension:
10% of patients show at their doctor, *but not in real life*, a high blood pressure and receive unnecessary medication

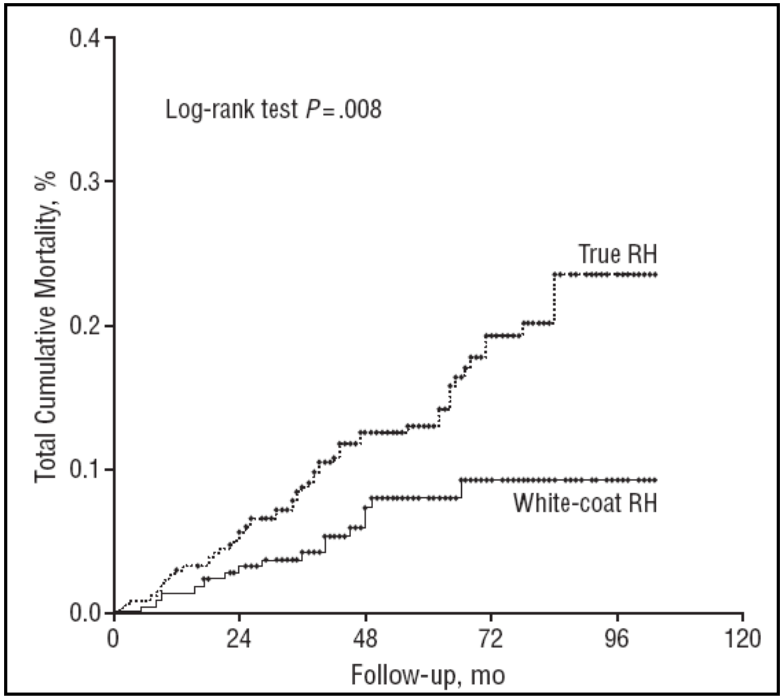

=> A lab study ≠ a real deployment
=> Many short-term studies ≠ a long-term study

Classifying Activities: Capturing *real* activities

Systematic distortions in trials occur frequently:

- **Affective valence effect** (Kihlstrom, 2000)
  Information associated with positive affect is more easily remembered
- **Mood congruent memory effect** (Kihlstrom, 2000)
  Persons with happy moods recall more positive events and fewer negative events than people in sad moods
- **Peak end rule & duration neglect** (Kahneman et al., ‘93)
  People judge an experience largely based on how they felt at its peak
  Judgments of unpleasantness of painful experiences depend very little on the duration of those experiences
- **Digit bias** (Shiffman & Paty, 2005)
  Cigarette smokers tend to “heap” their reports of smoking around particular numbers reflecting the number of cigarettes per pack or fractions thereof (e.g., 10 or 20)
- ...

Classifying Activities: Capturing *real* activities

Common pitfalls:

- **wrong sensors**: sensors do not capture the essence of the activities
- **wrong environment**: activities are not performed in the environment or circumstances where activities usually occur
- **wrong users**: users are different from target users
- **insufficient samples**: not enough data is captured

Classifying Activities: Sensor coverage

"One sensor does not say that much"

- Assume a person uses a <u>perfect</u> orientation sensor on upper leg to detect the activities "sitting" and "standing"
- Let's see how well this works from an information point of view…

15:37
sitting

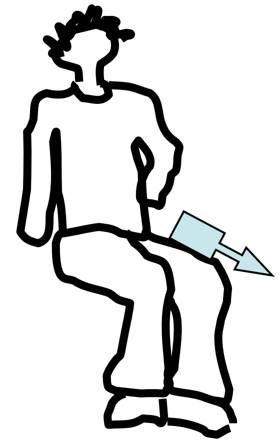

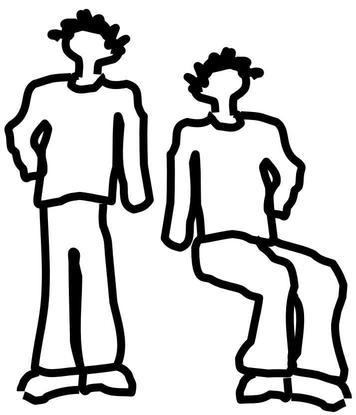

16:02
standing

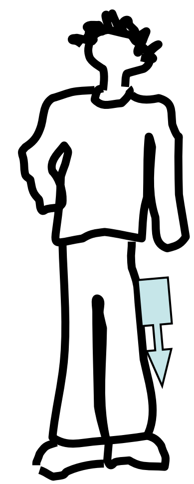

Is it working perfectly?

Classifying Activities: Sensor coverage

- Since only one leg is monitored, bending that leg and standing on the other leg can be falsely classified as sitting

- For the same reason, sitting on a high enough area with one leg dangling downwards can be falsely classified as standing

- Even worse, other activities (not sitting or standing) can be wrongly classified

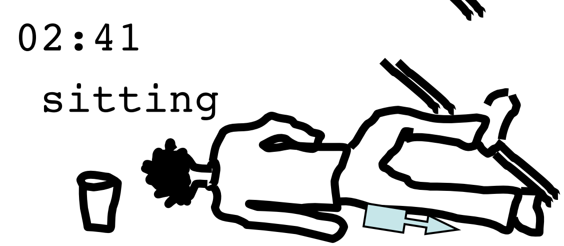

Classifying Activities: Sensor coverage

Solution: Multiple sensors, networked together

- Multiple of the same sensors, elsewhere located
  - Observing the same phenomenon elsewhere

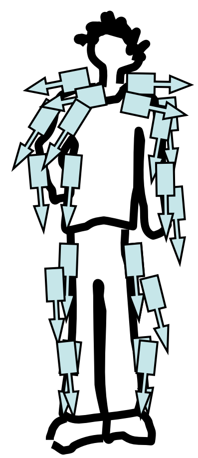

Classifying Activities: Sensor coverage

Solution: Multiple sensors, networked together

- Multiple different sensors
  - Observing a different phenomenon
    - e.g., accelerometer to measure motions and an RFID reader to read handheld objects to detect together object-based activities

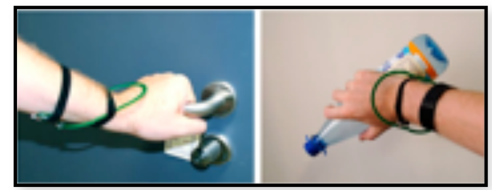

  - Observing the same phenomenon, differently
    - e.g., accelerometers to measure impact of steps with gyroscopes to measure leg's angle to detect together walking, running, …

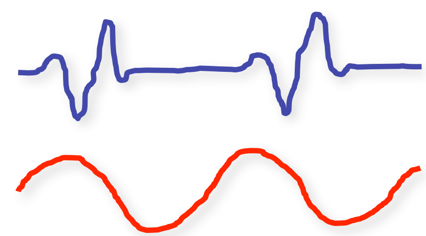

    - e.g., accelerometers together with GPS data to detect walking, running, …

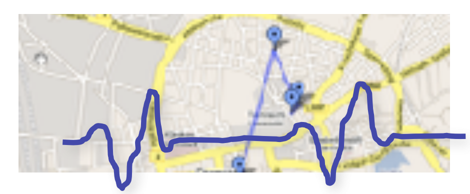

Classifying Activities: From Sensor Signals to Recognition

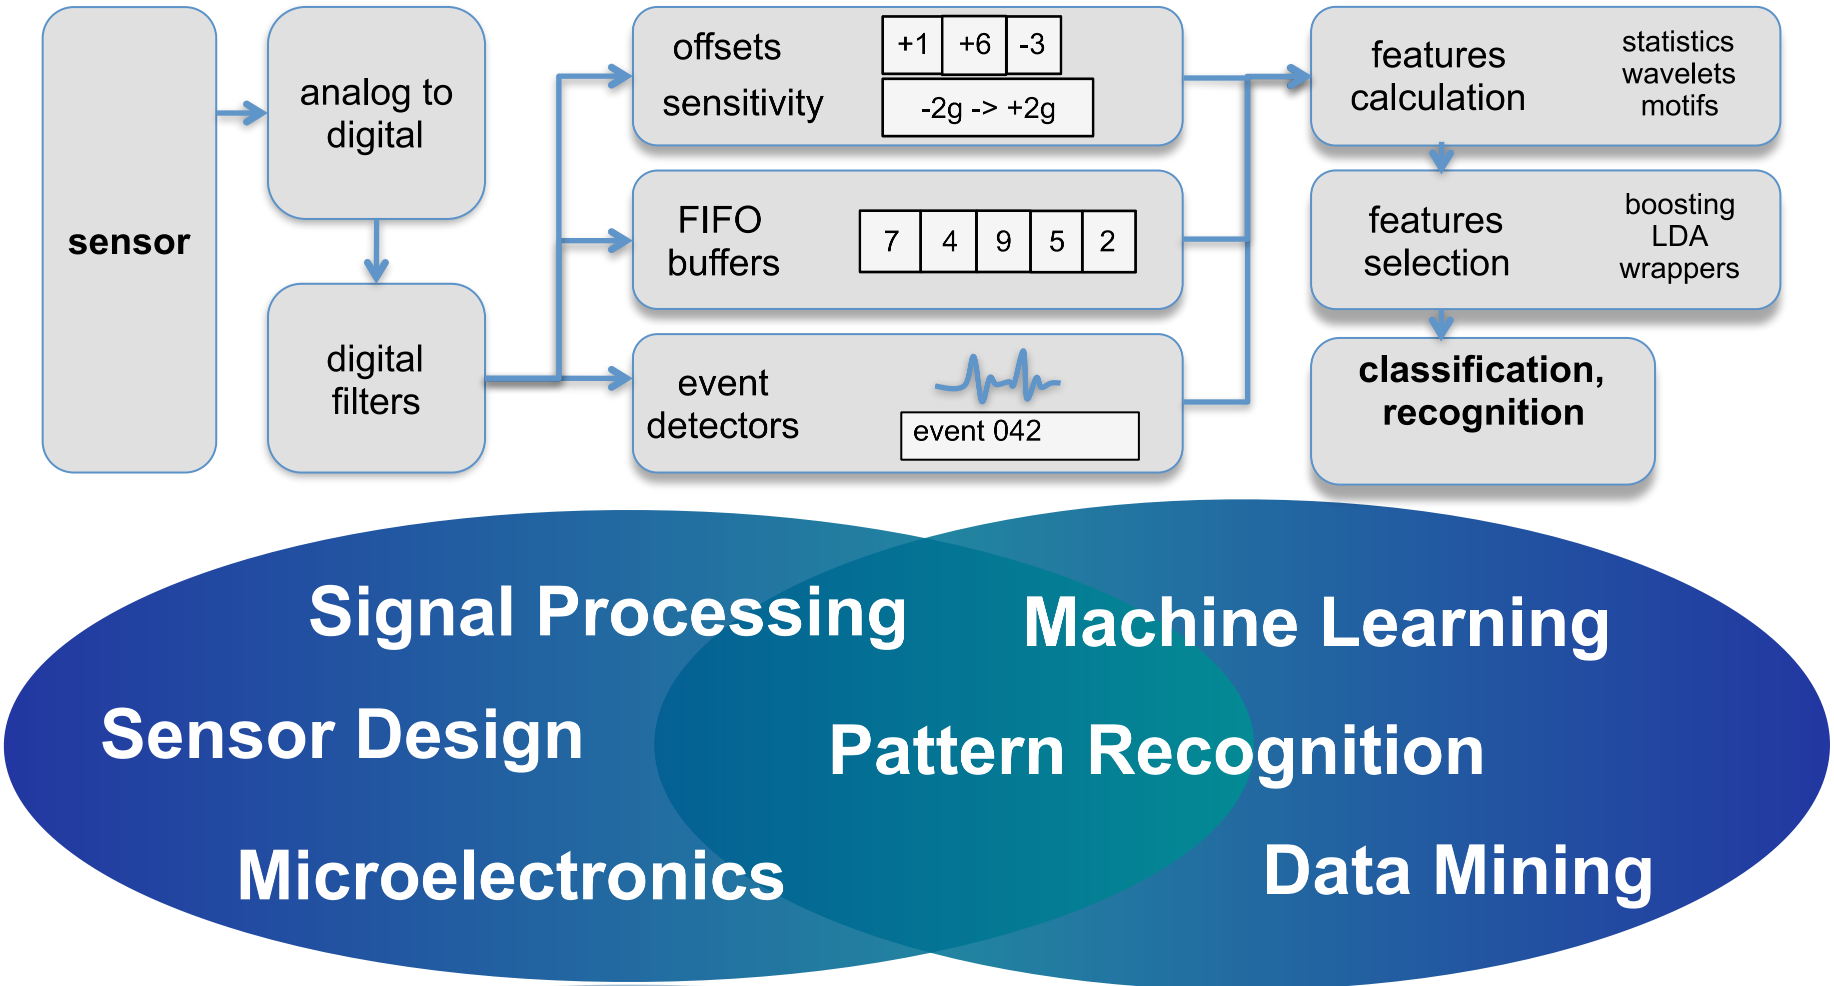

Classifying Activities: From Sensor Signals to Recognition

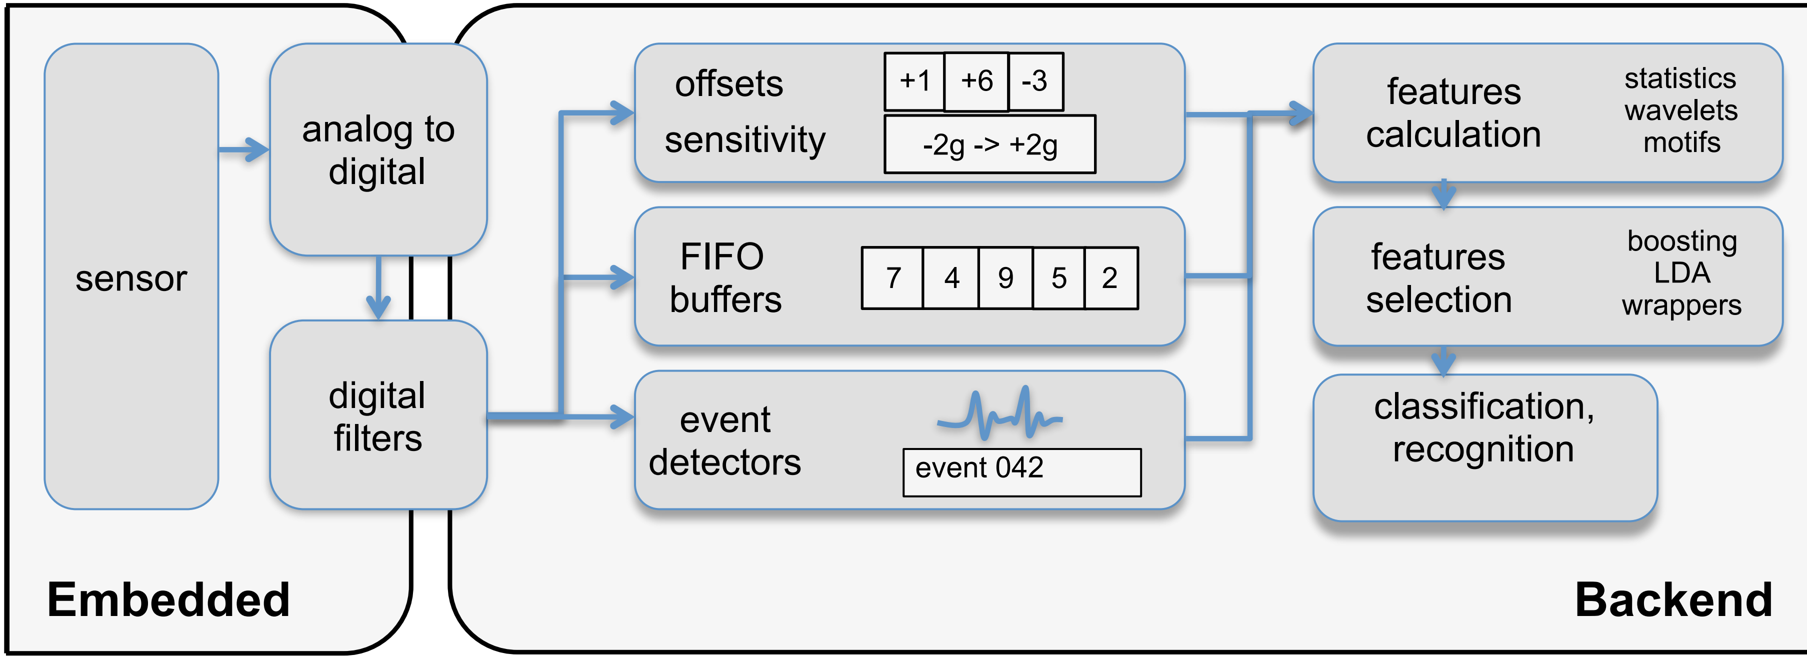

year	name	size (mm)	axes	output	I (uA)	detectors	range (g)
1995	ADXL05	$10 \times 10 \times 4.5$	1	voltage	8k-800	-	1-5
1999	ADXL202	$10 \times 7.4 \times 3$	2	duty cycle/volt.	600	-	2
2003	LIS3L02AQ	$7 \times 7 \times 1.8$	3	voltage	850	-	2,6
2006	ADXL330	$4 \times 4 \times 1.45$	3	voltage	320	-	3
2007	SMB380	$3 \times 3 \times 1$	3	SPI, I^2C, 1 int.	200	freefall, motion,	2,4,8
2007	LIS331DL	$3 \times 3 \times 1$	3	SPI, I^2C, 2 int.	290	freefall, motion, taps	2,8
2009	ADXL345	$3 \times 5 \times 1$	3	SPI, I^2C, 2 int.	145	freefall, motion, taps	2,4,8,16
2010	BMA220	$2 \times 2 \times 1$	3	SPI, I^2C, 1 int.	250	freefall, motion, taps, turn	2,4,8,16

Classifying Activities: From Sensor Signals to Recognition

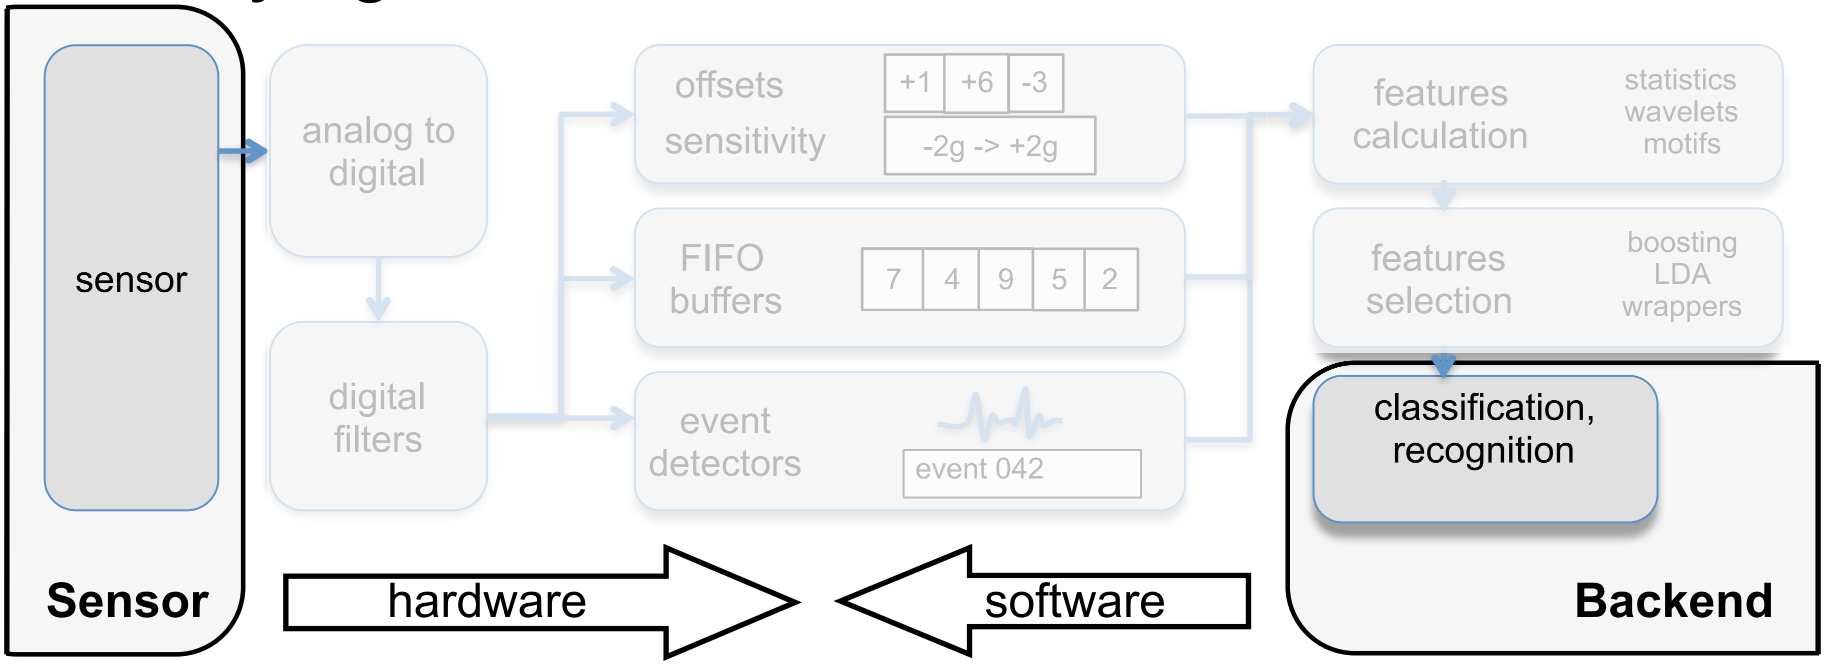

- **Design:** Hardware and software choices affect each other
- **Constraints**
  - real-world phenomena → Recognition Accuracy
  - battery-operated units → Runtime
  - long-term deployment → Reliability

Outline

- Definitions of context-aware computing
- Early examples
- From Context to Activity
- Classifying activities:
    - Capturing real activities
    - Sensor coverage
    - **Features**
    - Classifiers
    - Evaluation of activity recognition
    - Ground truth

Ubiquitous Computing

UNIVERSITÄT SIEGEN

Classifying Activities: Features

- Example: Thresholds

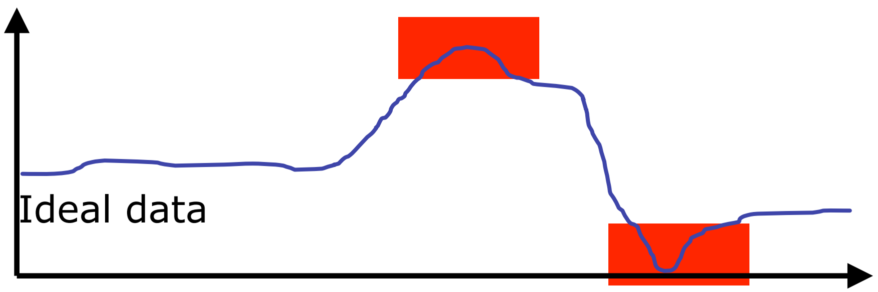

not always the best:

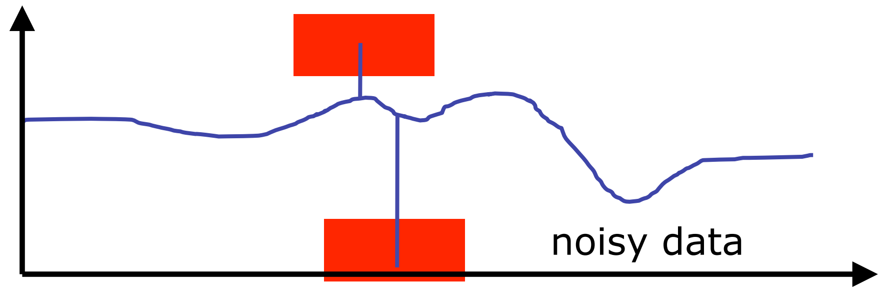

Embedded Features in MEMS sensors

Case study: ADXL345, double-tap feature example

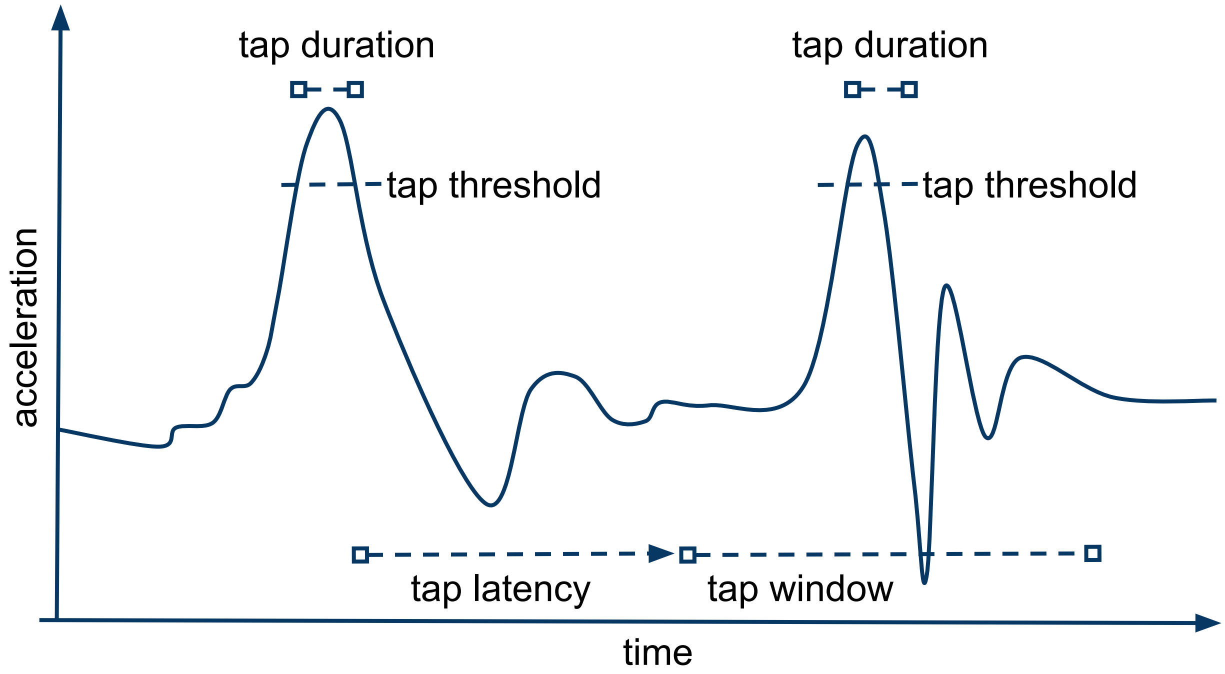

Classifying Activities: Features

- Time series of sensor data:
  - "sitting": 
  - "walking": 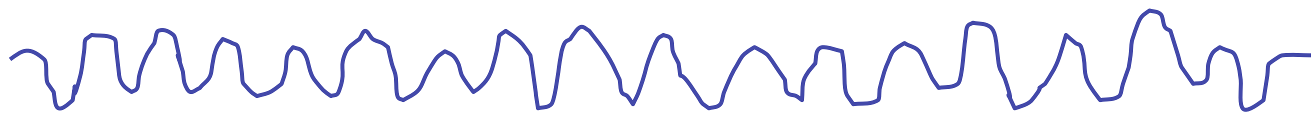

- By looking at the time series, what features would work well to distinguish the activities "sitting" and "walking"?

-> very simple solution: (maximum – minimum) over sliding window:

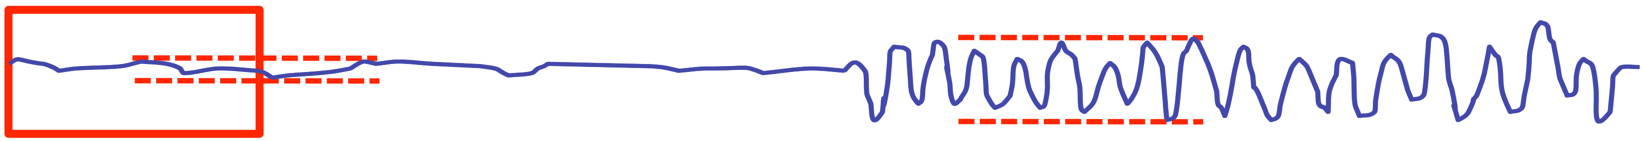

Classifying Activities: Features

- Time series of sensor data:
  - "sitting": 
  - "walking": 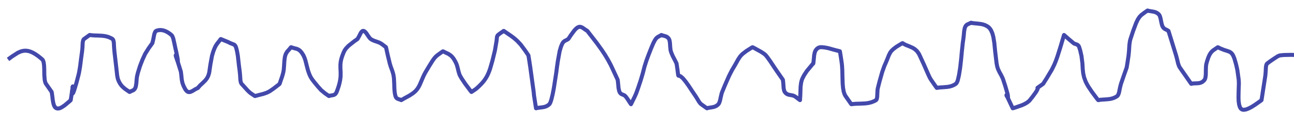
  - "running": 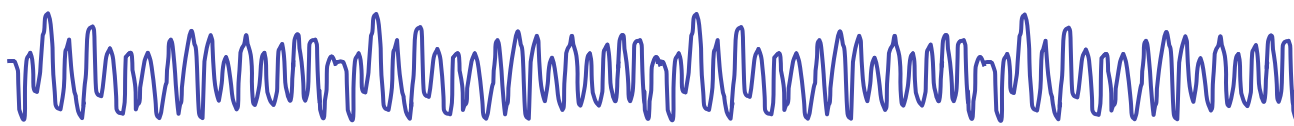
- By looking at the time series, what features would work well to distinguish the activities "sitting" and "walking"?

-> (max – min) over sliding window does not work that well:

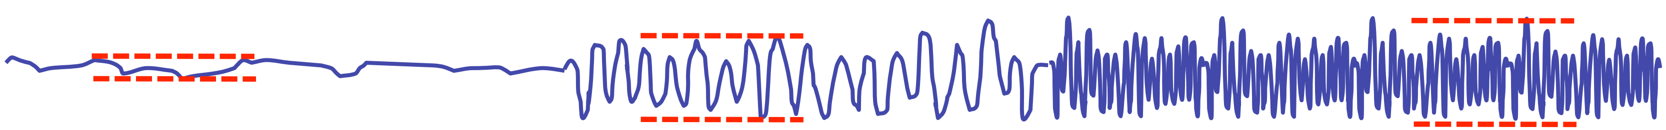

Classifying Activities: Features

- Basic statistics over buffer often work well, and are fast to calculate:
  - **Minimum** and **maximum** over **sliding window** (e.g., last 9 values)
  - **Mean** (or average) over sliding window
  - **Variance** over sliding window

t_{-8} … t_0

122	131	117	102	98	114	122	134	125

Min = 98

Max = 134

Mean = (122 + 131 + 117 + 102 + 98 + 114 + 122 + 134 + 125) / 9

= 118.3

Variance = $((122-118.3)^2 + (131-118.3)^2 + (117-118.3)^2 + \ldots + (125-118.3)^2) / 9$

= 130.9

Classifying Activities: Features

- For large (sliding) windows, such features are **much** easier to process:

  (118.3, 130.9) instead of (122,131,117,102,98,114,122,134,125)

- But: features are selected (most distinctive for separating classes) with domain knowledge (information on class, sensor signal, etc.)

- And: window size becomes a defining factor

- Alternatives: feature selection, e.g.:
  - wrapper method: incrementally try out subsets and select the best
  - boosting: provides a ranking of best performing features…

Classifying Activities: Classifiers

- So far we have converted the raw sensor data in features (e.g., mean and variance)
- Feature space:

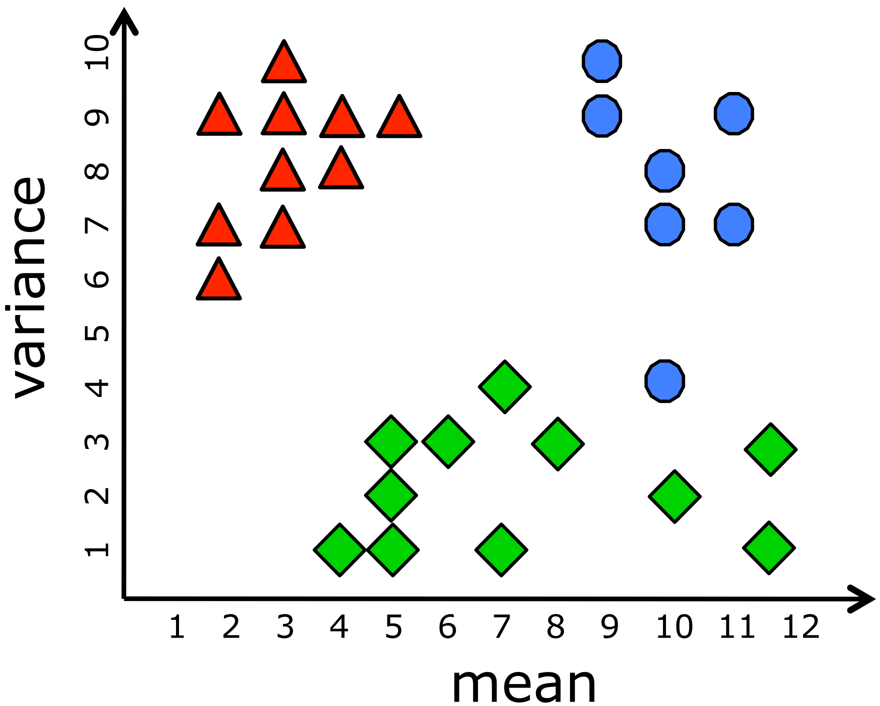

- Now we want an algorithm to learn how to classify new data points, given the feature space of data for which we know the class => **training data**
- Later we can test the algorithm on different data => **test data**
- Classes:

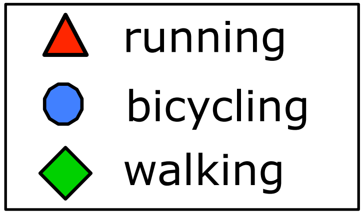

Outline

- Definitions of context-aware computing
- Early examples
- From Context to Activity
- Classifying activities:
    - Capturing real activities
    - Sensor coverage
    - Features
    - Classifiers
    - **Evaluation of activity recognition**
    - Ground truth

Classifying Activities: Evaluation of Classifiers

Typical: Split up the data so that each part will become (part of) training data and test data, without using the same data for both

n-Fold Cross Validation

e.g., 3-fold cross validation:

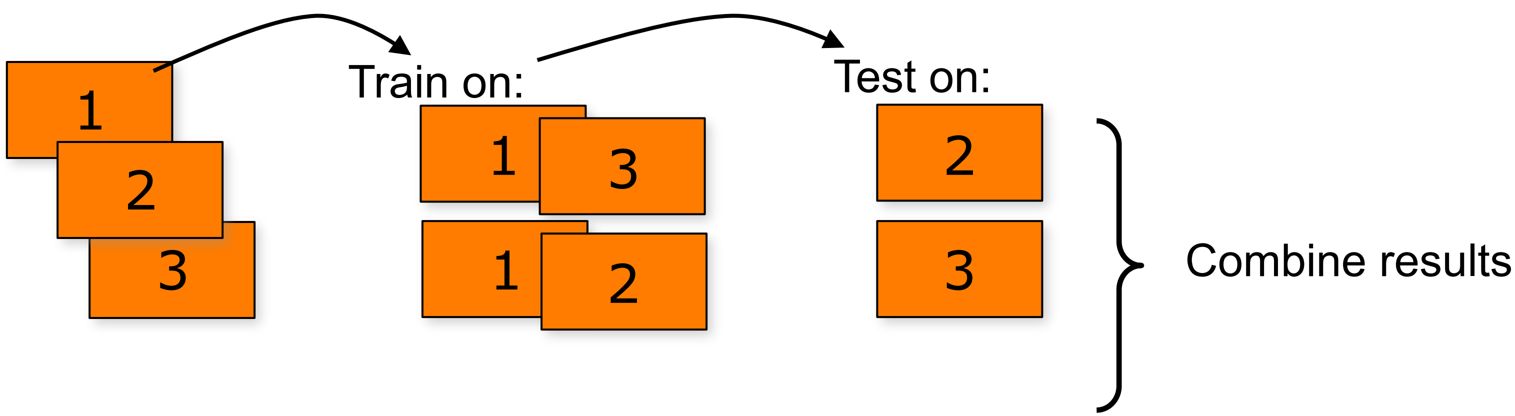

Classifying Activities: Evaluation of Classifiers

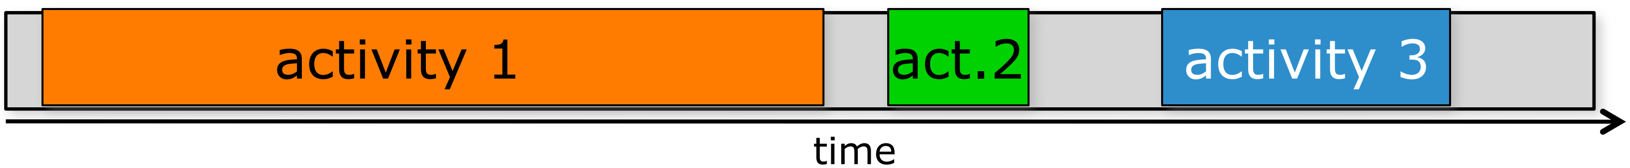

In real life,

1. activity classes never cover the whole data set
2. activities never last equally long

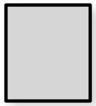
: void class / background class

Classifying Activities: Ground truth

- We need to know the positives and negatives in a data set
- Thus: Data needs to be annotated for training and testing

- Ideally everything is observed and annotated while data streams in, e.g.:
  - using video footage (synchronized with sensor data)
  - using experiment observers
    (annotating the sensor data stream directly)
  - using secondary sensors
    (e.g., augmented cigarette lighter for annotating smoking instances)

Classifying Activities: Ground truth

- Data needs to be annotated for **training** and **evaluation**
- Ideally everything is observed and annotated while data streams in
- What about long-term experiments?

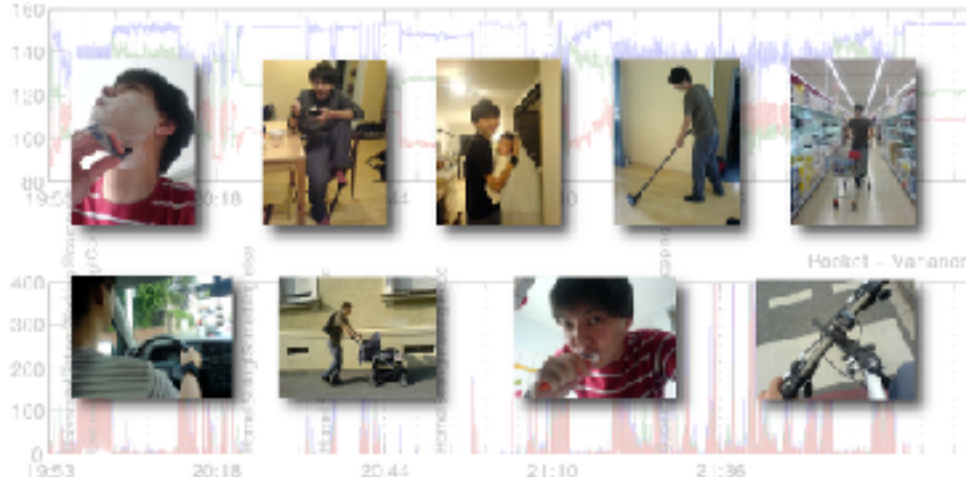

Tam Huynh, Human Activity Recognition with Wearable Sensors

Classifying Activities: Ground truth by **Time Diary**

- Typical for early experiments
- Sensor data is recorded with time stamps
- Time stamps are synchronized with a diary's timed activities

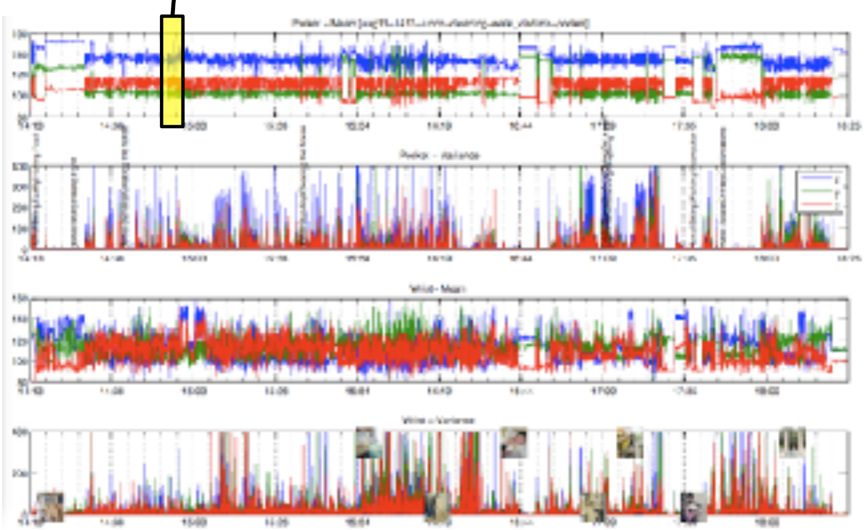

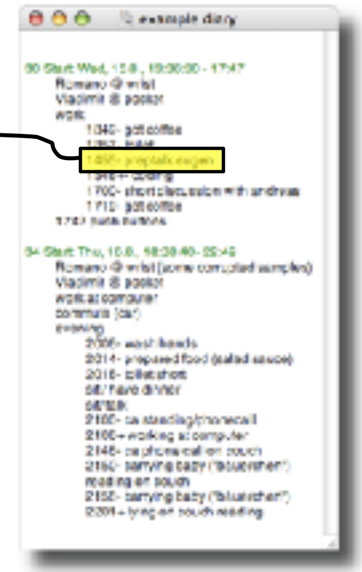

- Time diary can be kept by an observer or the test subject

☹ Large effort required from the observer/test subject

Classifying Activities: Ground truth by **Experience Sampling**

- Designed for more situated and flexible experiments where test subjects walk around freely
- A phone is carried by the user with annotation software
- The phone gives out an alarm every 15-30 minutes, and asks for the current activity

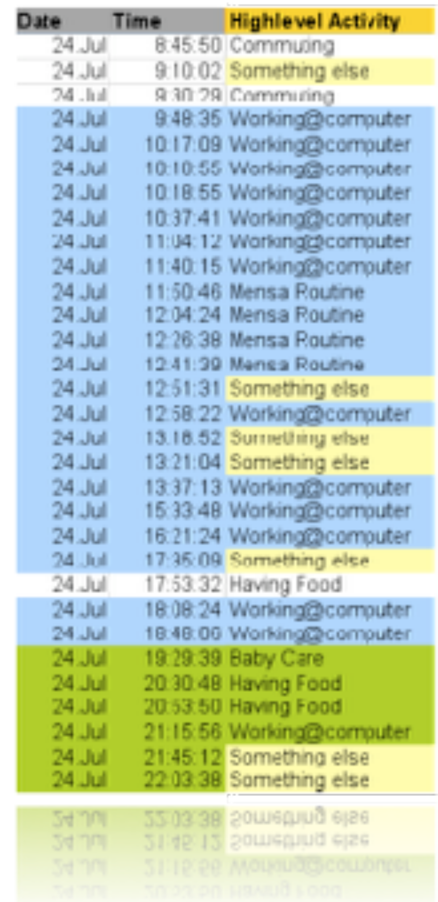

Date	Time	Highlevel Activity
24 Jul	8:45:50	Commuting
24 Jul	9:10:02	Something else
24 Jul	9:30:29	Commuting
24 Jul	9:48:35	Working@computer
24 Jul	10:17:09	Working@computer
24 Jul	10:10:55	Working@computer
24 Jul	10:18:55	Working@computer
24 Jul	10:37:41	Working@computer
24 Jul	11:04:12	Working@computer
24 Jul	11:40:15	Working@computer
24 Jul	11:50:46	Mensa Routine
24 Jul	12:04:24	Mensa Routine
24 Jul	12:26:38	Mensa Routine
24 Jul	12:41:39	Mensa Routine
24 Jul	12:51:31	Something else
24 Jul	12:58:22	Working@computer
24 Jul	13:18:52	Something else
24 Jul	13:21:04	Something else
24 Jul	13:37:13	Working@computer
24 Jul	15:33:48	Working@computer
24 Jul	16:21:24	Working@computer
24 Jul	17:35:09	Something else
24 Jul	17:53:32	Having Food
24 Jul	18:08:24	Working@computer
24 Jul	18:48:00	Working@computer
24 Jul	19:29:39	Baby Care
24 Jul	20:30:48	Having Food
24 Jul	20:53:50	Having Food
24 Jul	21:15:56	Working@computer
24 Jul	21:45:12	Something else
24 Jul	22:03:38	Something else

- ☹ Interrupts test subject
- ☹ Does usually not cover short activities
- ☹ Redundant queries for long activities

Classifying Activities: Ground truth by **Self recall**

- Test subjects annotate themselves, using:
  - Their memory of the recent past
  - Sensor data visualizations
- Test subjects wear a sensor, and annotate after the logging is done (e.g., every evening)
- No interruptions

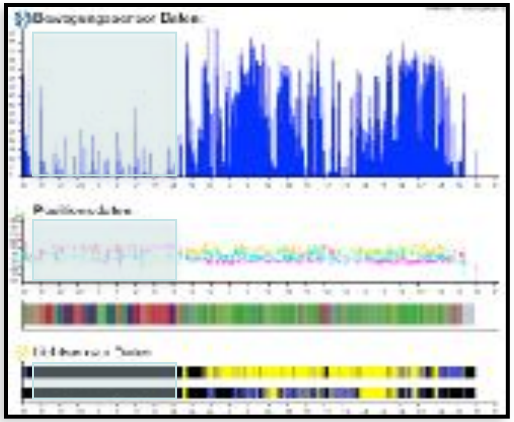

☹ Not as accurate (depends on memory)

☹ Still requires daily a bit of time

Outline

- Definitions of context-aware computing
- Early examples
- From Context to Activity
- Classifying activities:
  - Capturing real activities
  - Sensor coverage
  - Features
  - Classifiers
  - Evaluation of activity recognition
  - Ground truth
  - **Shape matching**

Classifying Activities: Dealing with large time series

Example of a very specific feature: Shape matching

- Step 1: approximate the time series in segments
- Step 2: store the set of segments from a known pattern
- Step 3: calculate distance between known pattern and new data

Two algorithms needed:

1. Approximation algorithm

   Compresses the data and reduces it to its essentials

- Matching algorithm

   Compares pieces of approximated data and decides whether they are similar enough

Shape matching

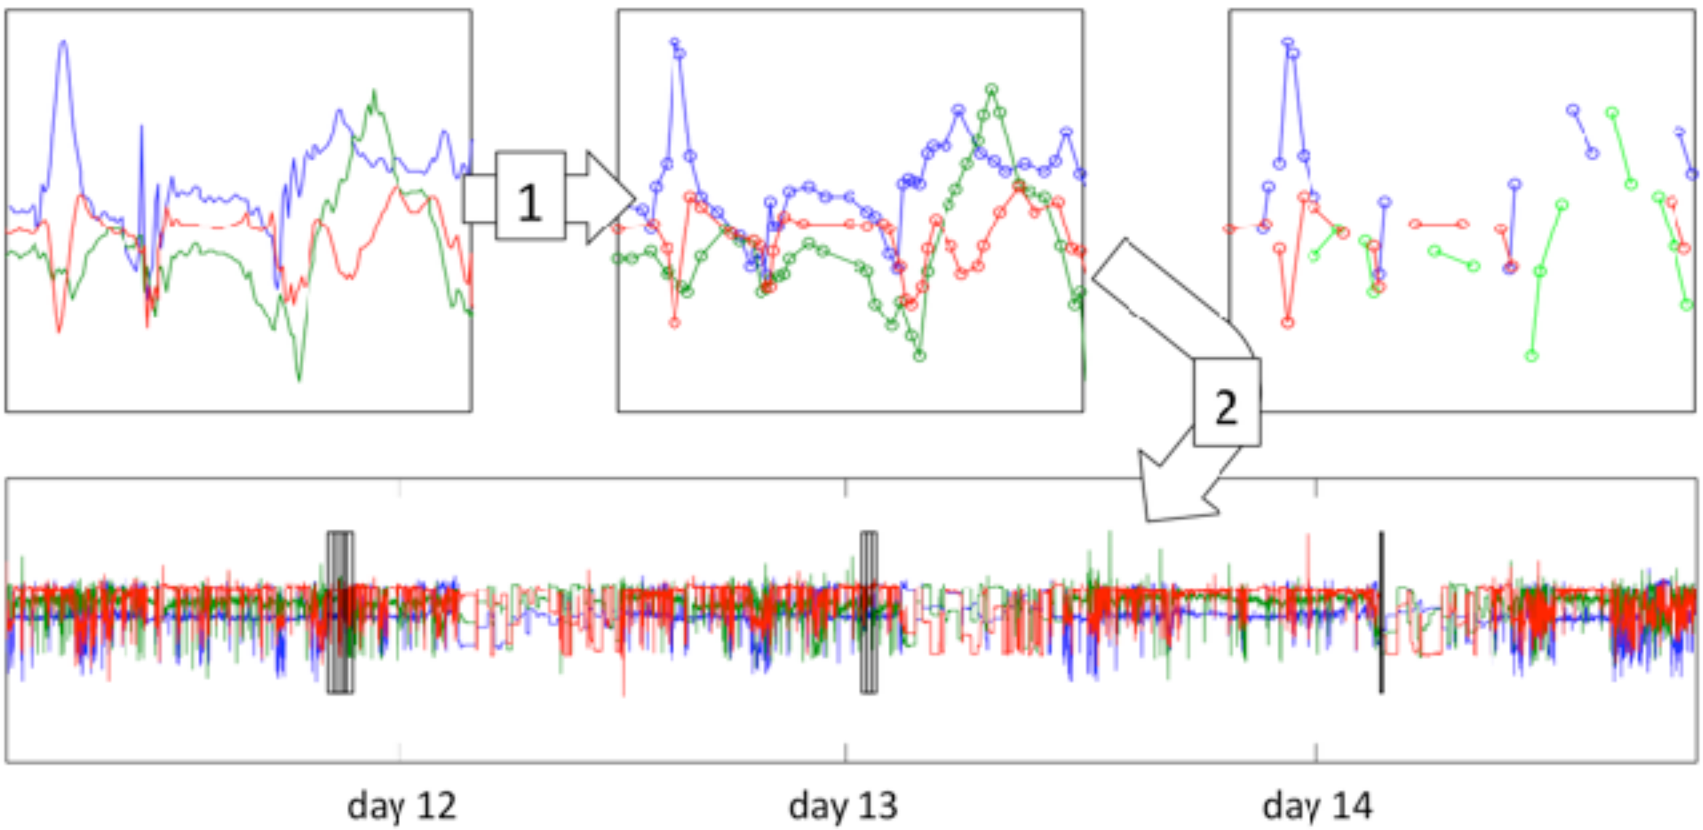

Shape matching

Time Series

t1 = [123 127 125 129 139 143 128 122 117 102 120]

Piecewise Approximation

- Sliding Window
- SWAB
- Piecewise Aggreg. Approximation
- Adaptive Piecewise Constant Approx.
- …

Basic Statistics

- minimum
- maximum
- moments
  mean
  variance

Discrete Fourier Transform

- first coefficients
- exponential bands
- spectral energy

Discrete Wavelet Transform

- Haar
- Daubechies
- Symlet
- Coiflet

Symbolic

- PERSIST
- SAX
- iSAX
- SDL

Singular Value Decomposition

Sliding window

Segments are grown from a previous position to a new raw data point, until the approximation cost is too high

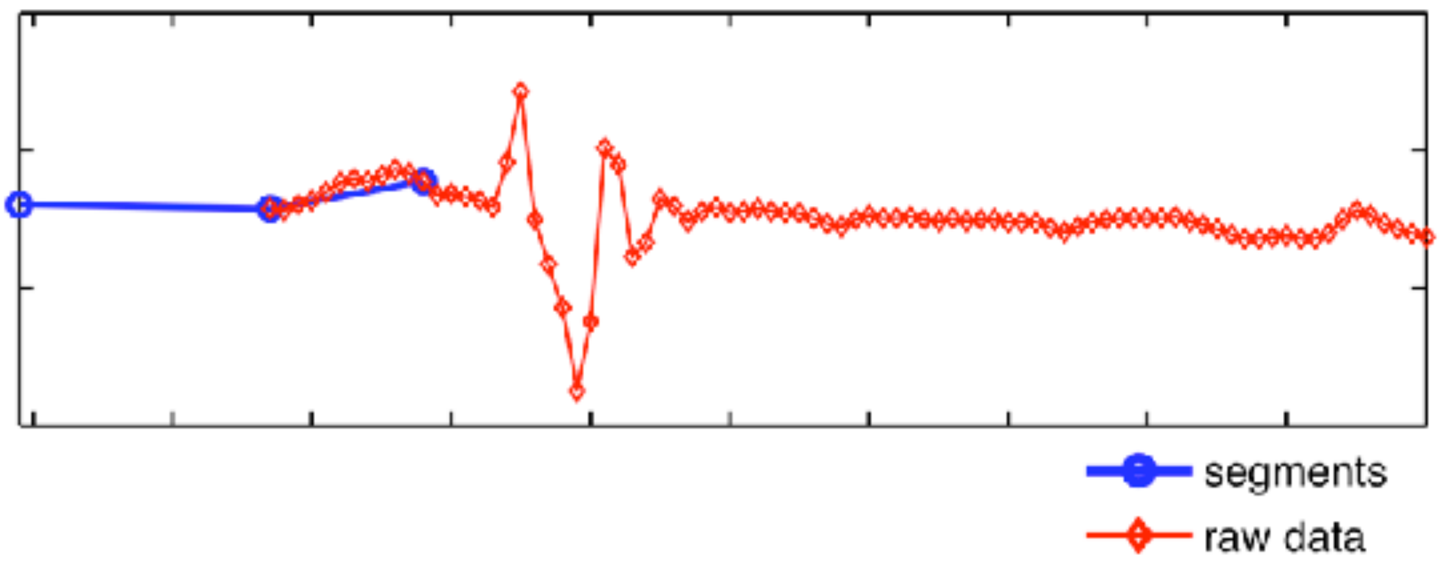

Sliding window

Approximation cost: difference between all interpolated points on the segment and the raw data points

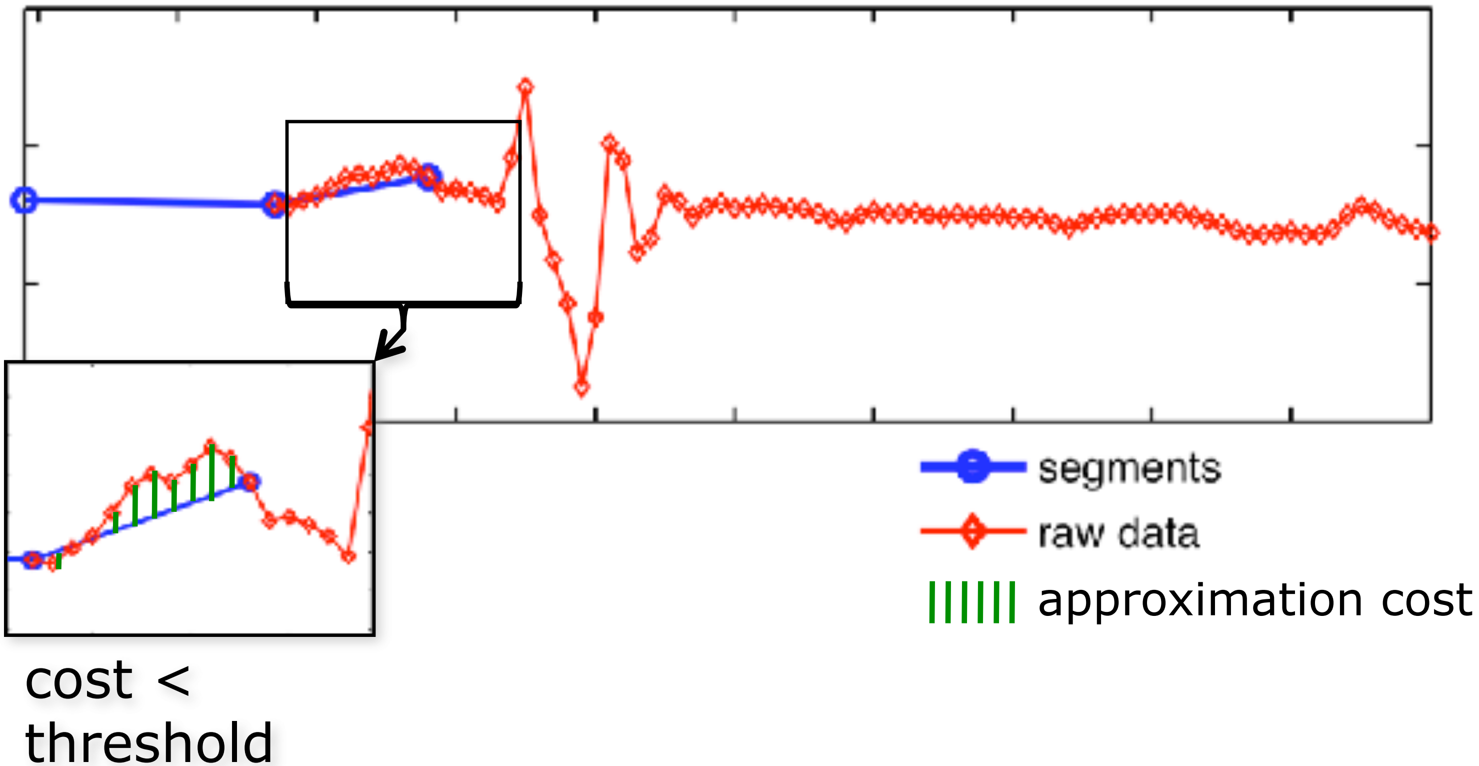

Sliding window

When the cost is too high, the previous raw data point is taken as the new position from which we start growing the next segment

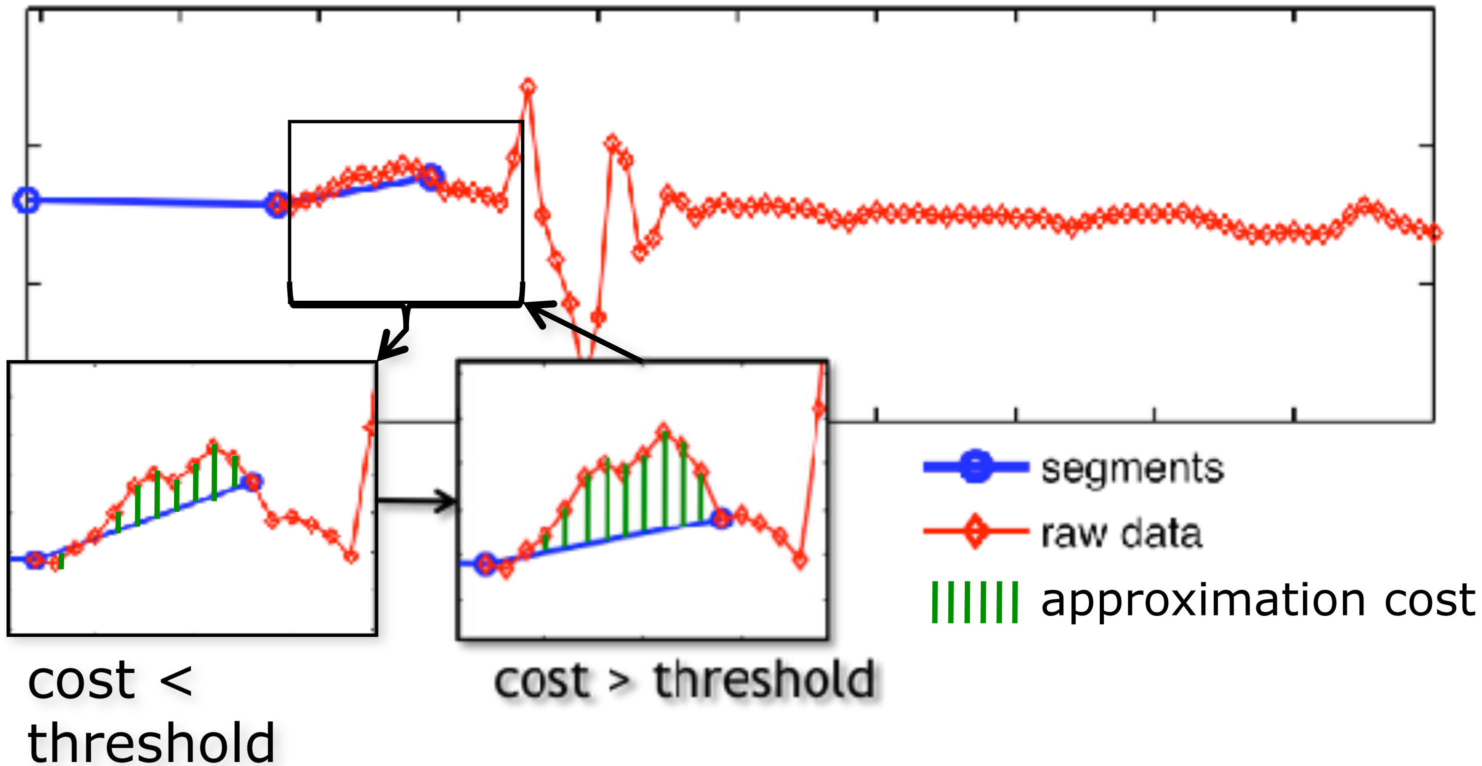

Sliding window

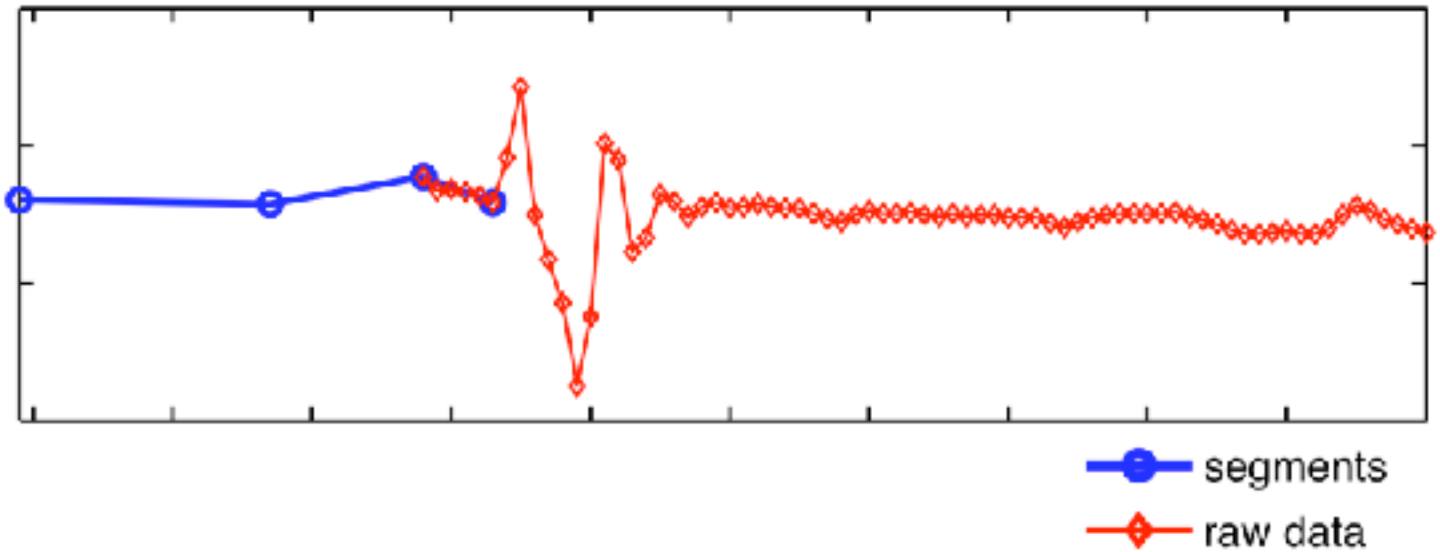

Sliding window

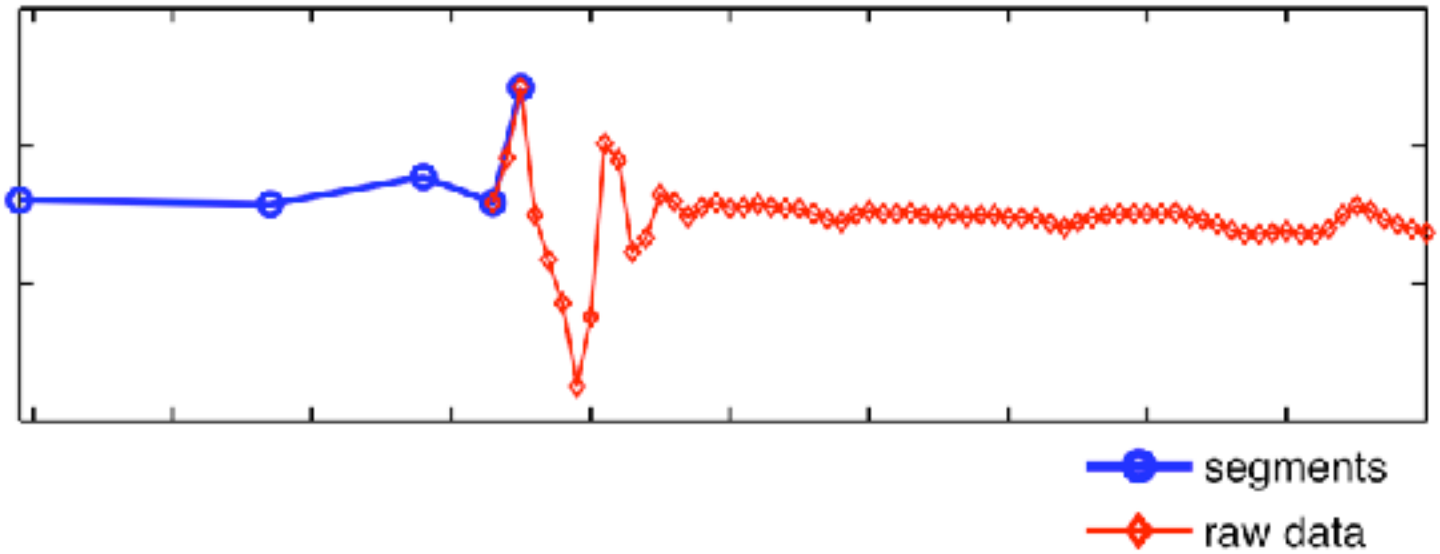

Sliding window

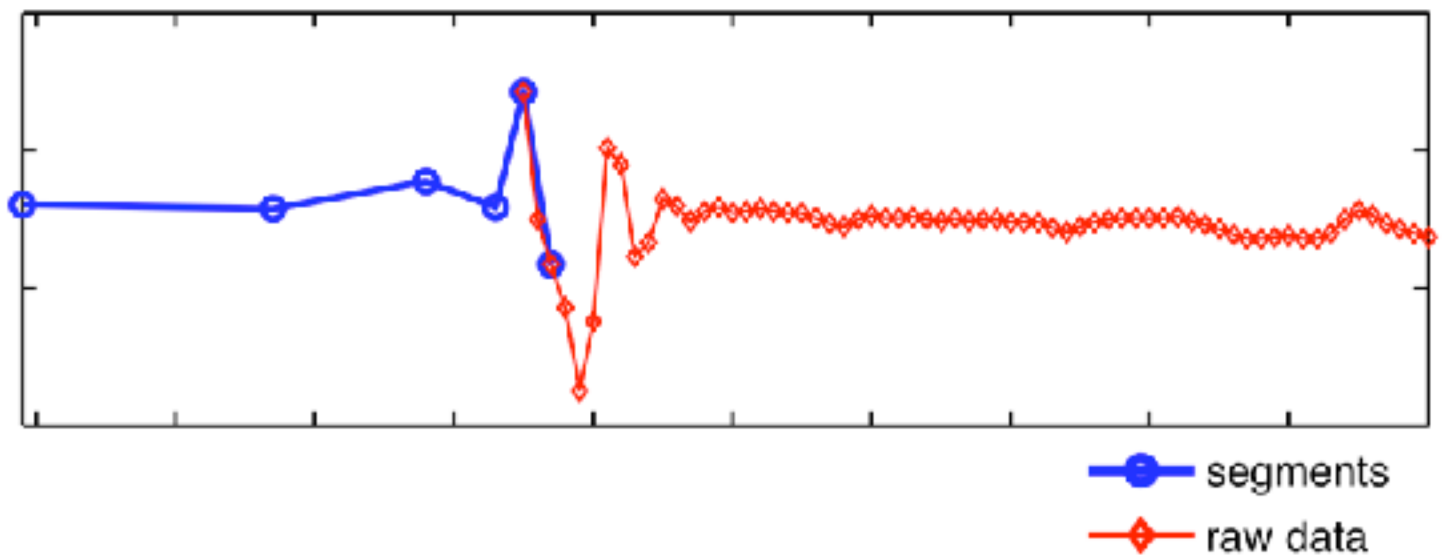

Sliding window

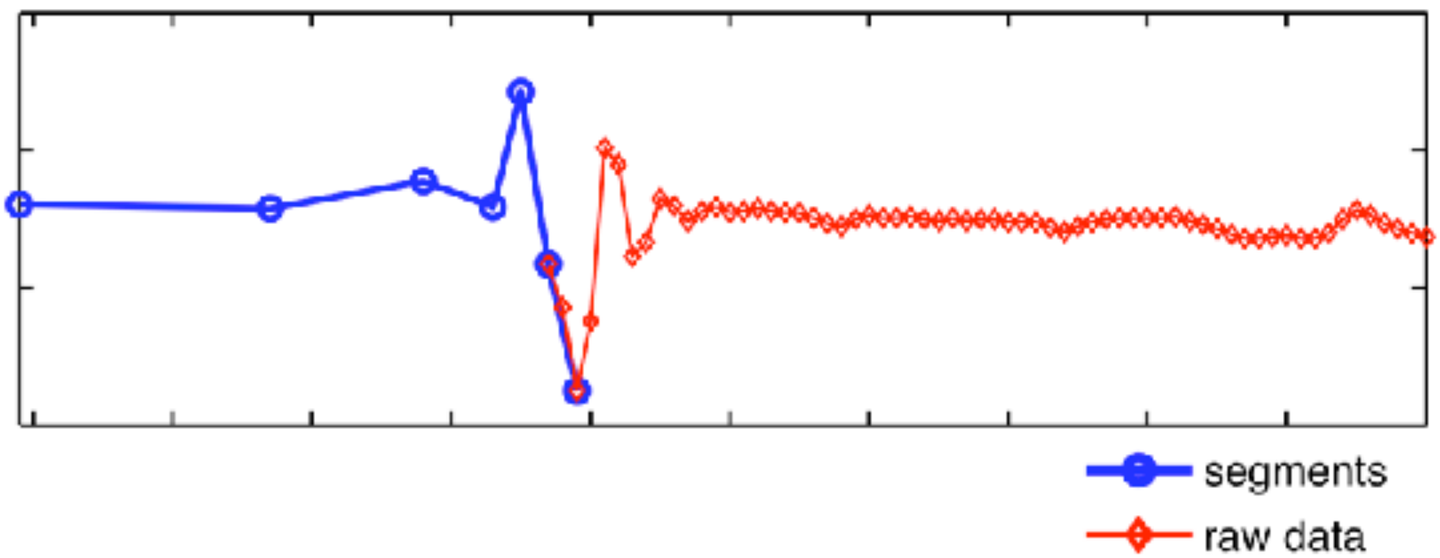

Sliding window

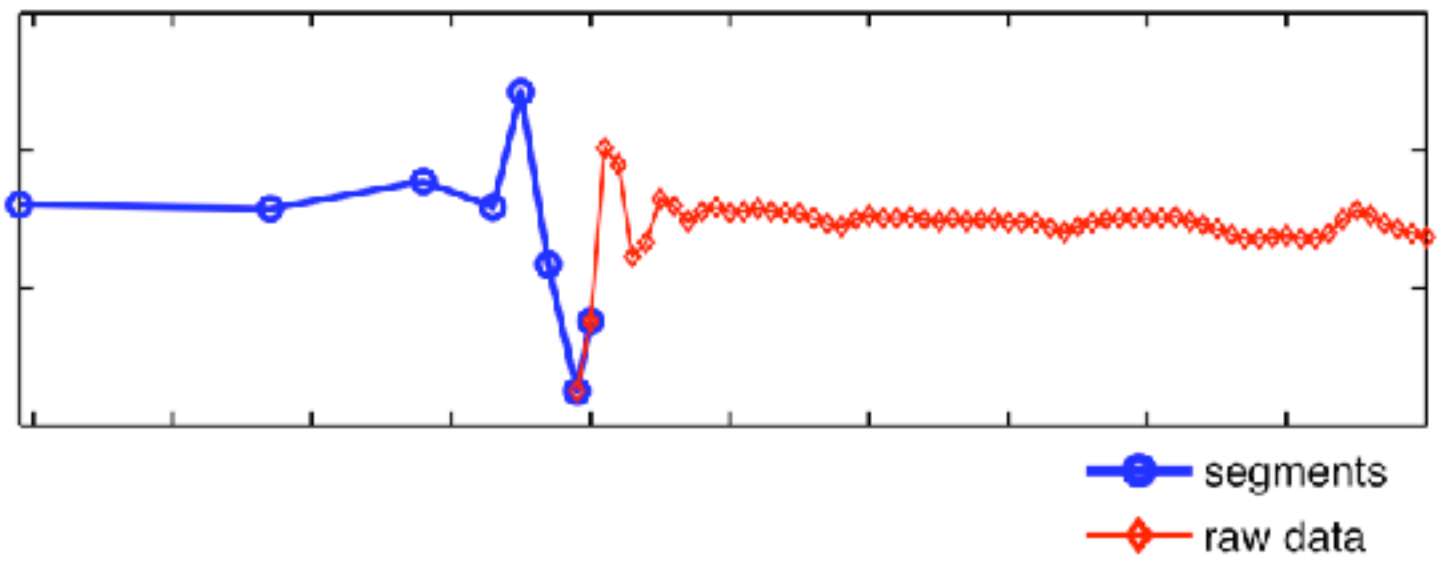

Sliding window

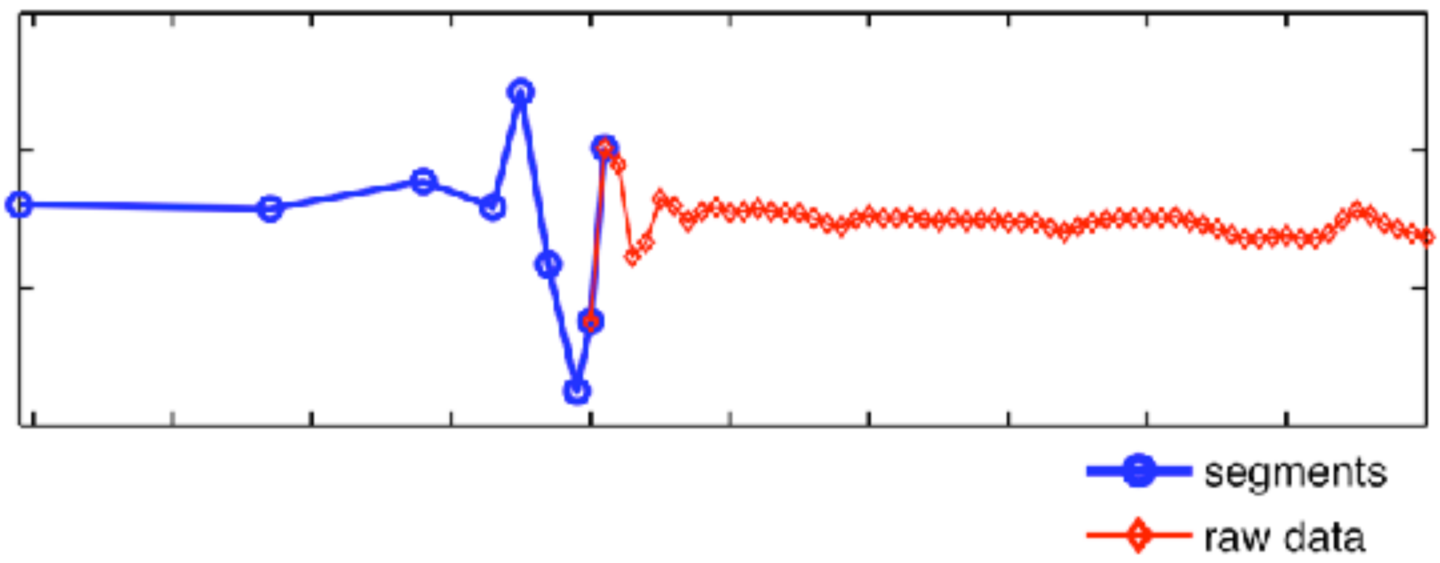

Sliding window

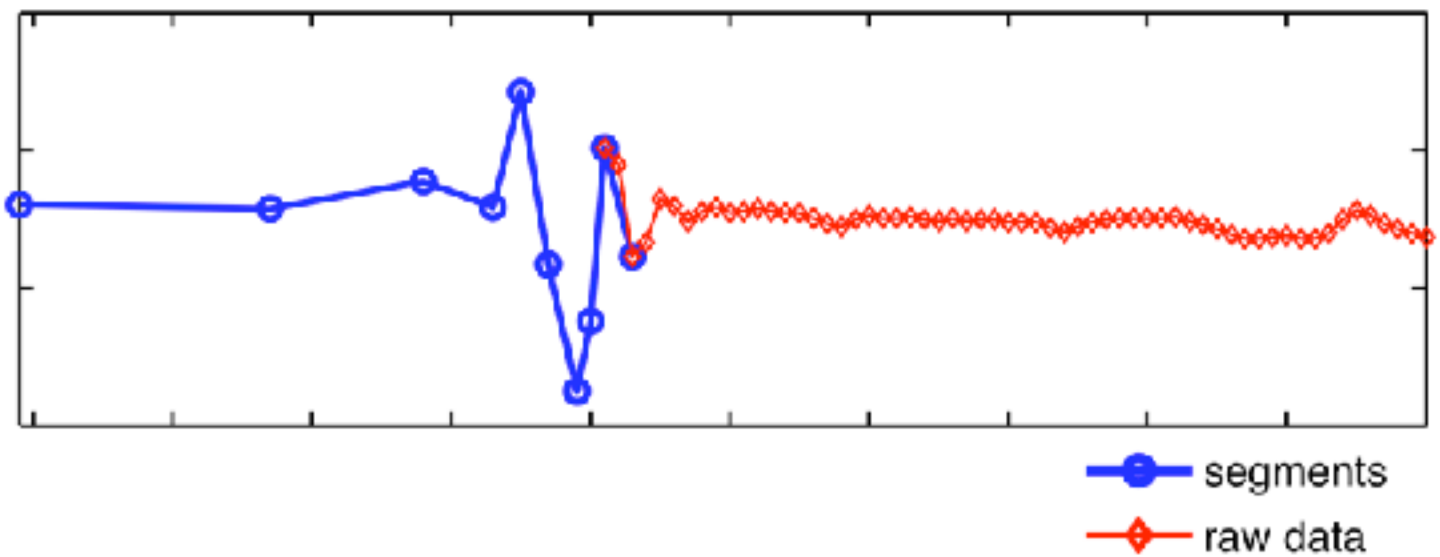

Sliding window

9 segment coordinates (18 bytes)
instead of 19 raw samples (19 bytes)

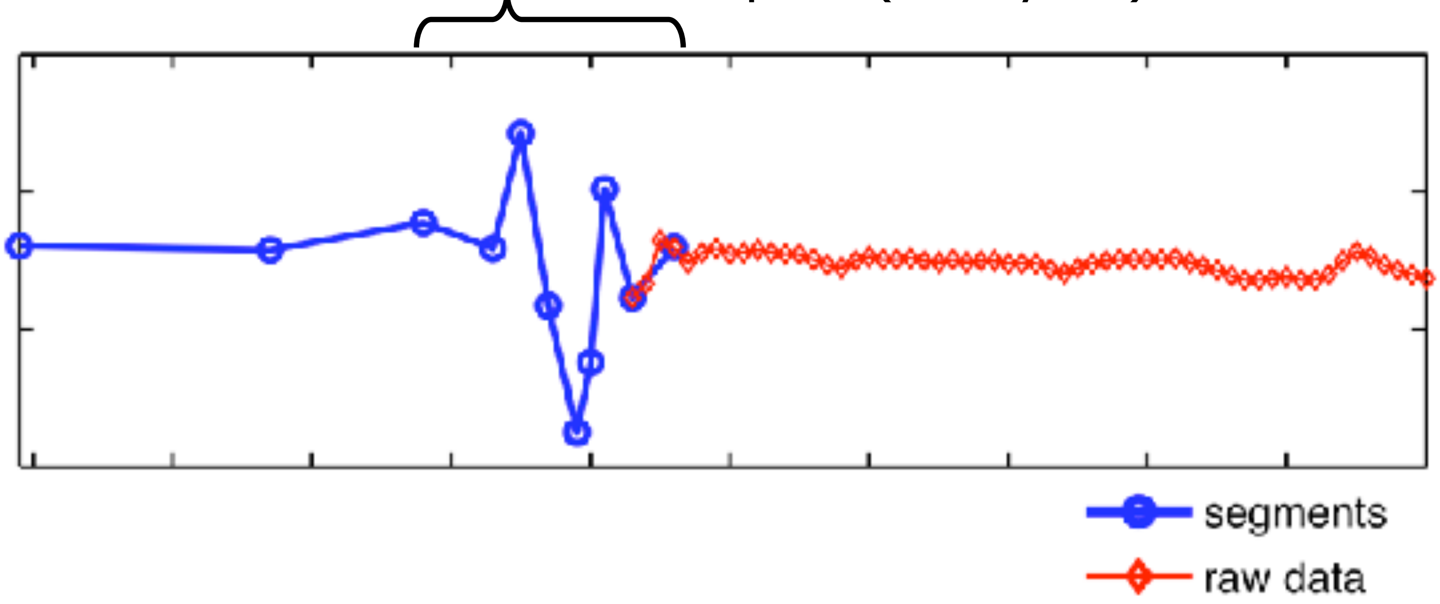

Sliding window

2 segment coordinates (4 bytes)
instead of 33 raw samples (33 bytes)

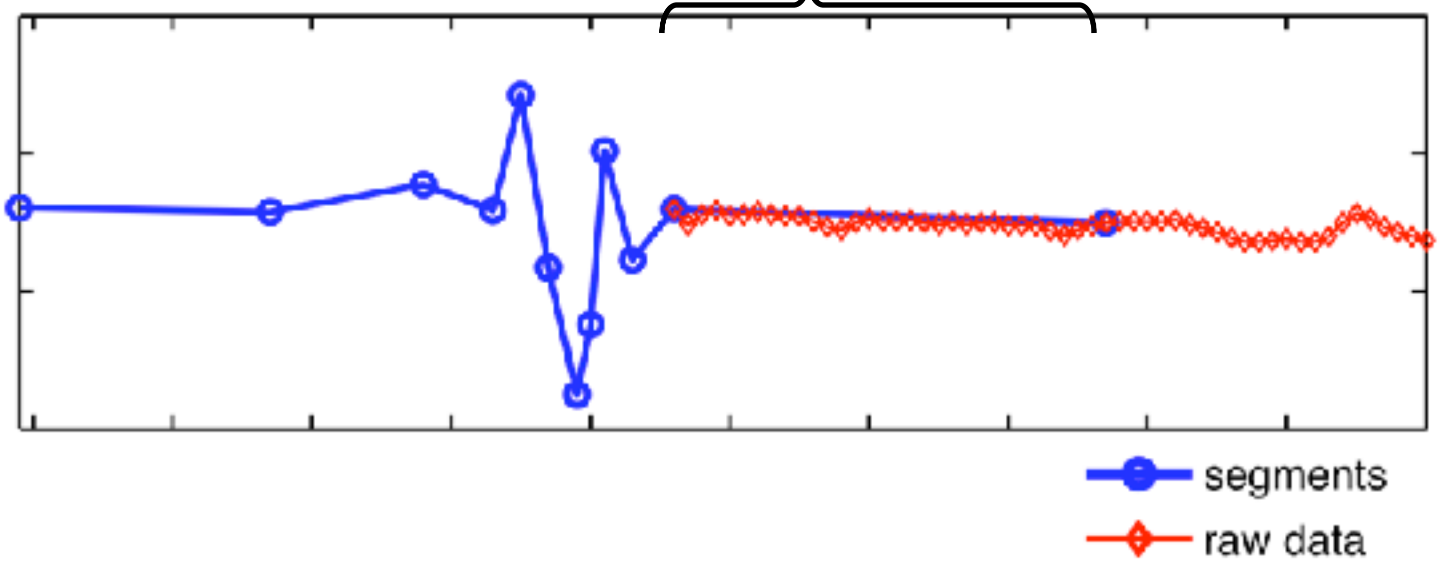

Sliding window

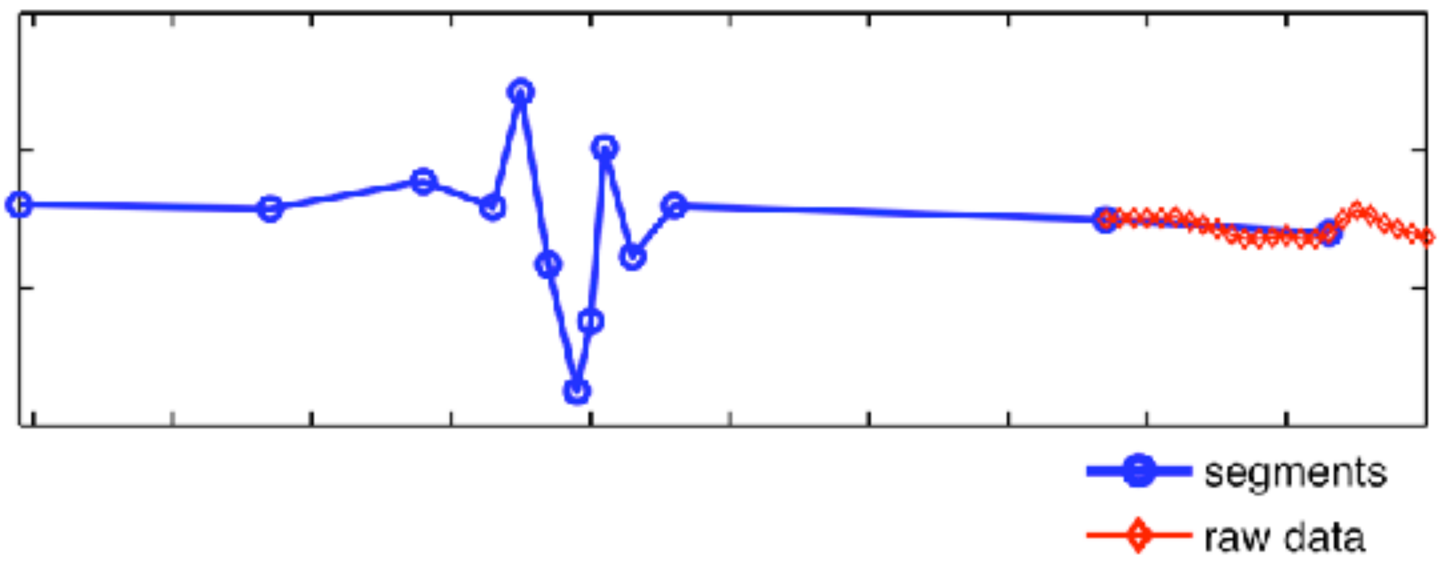

SWAB (Keogh '01)

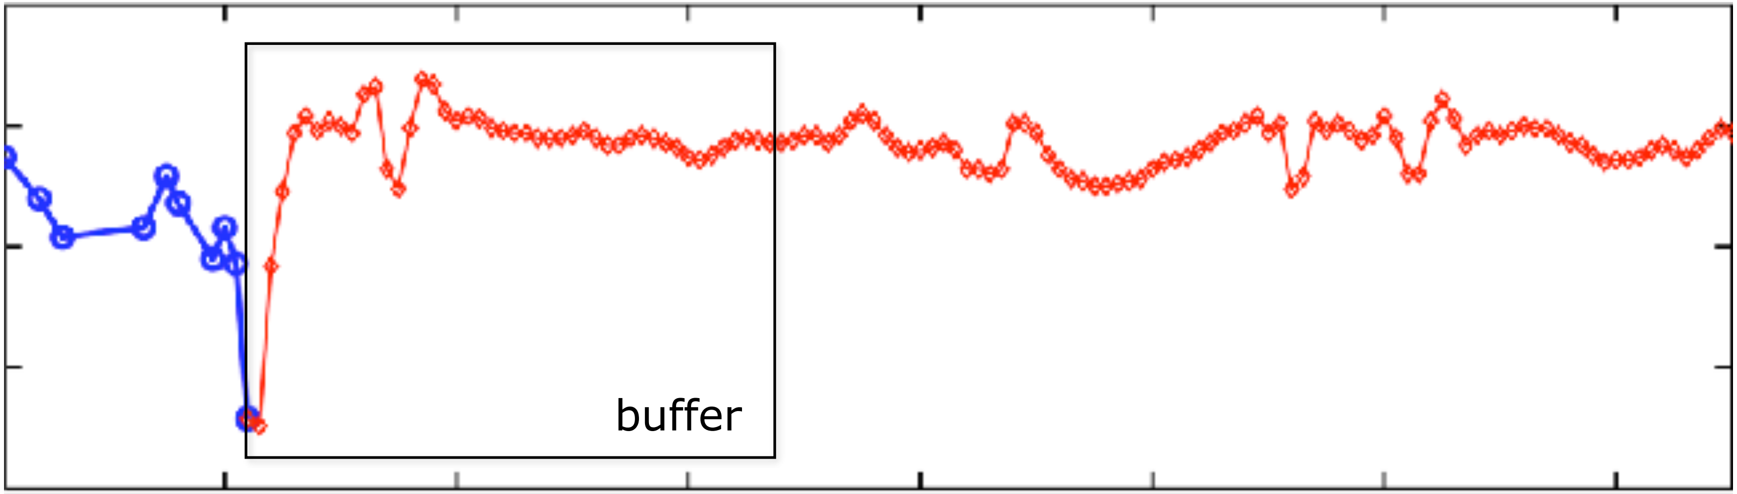

SWAB (Keogh '01)

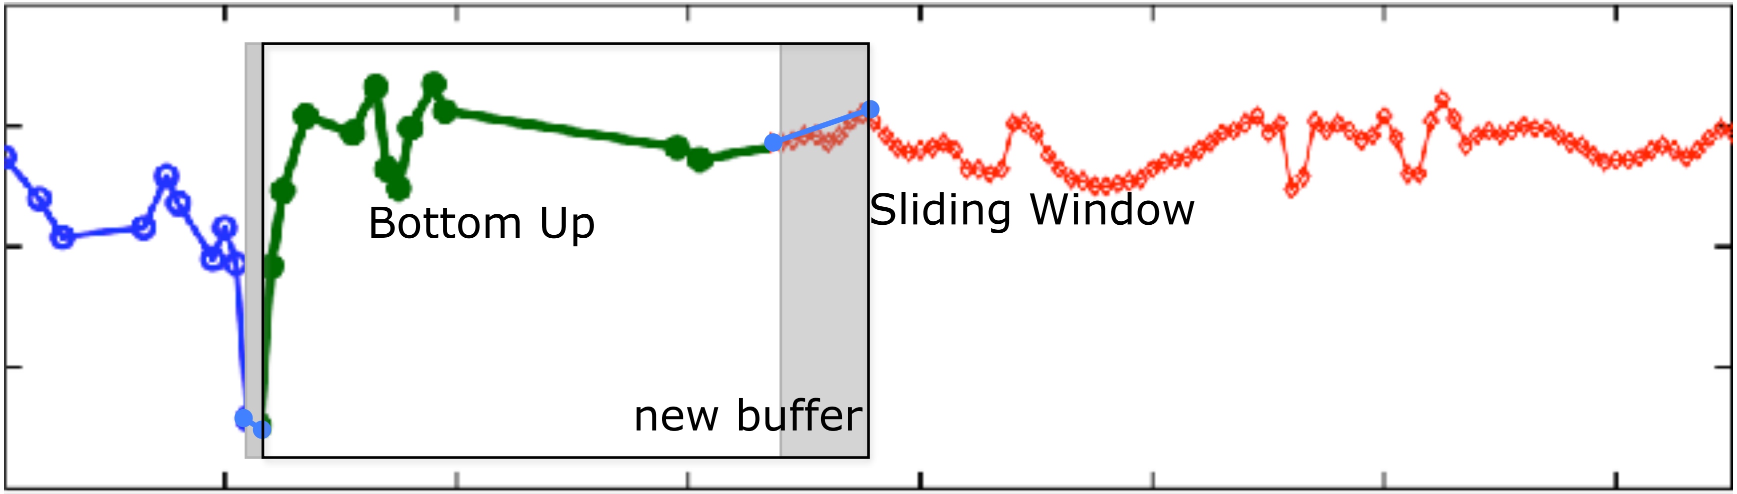

mSWAB (Berlin ’11)

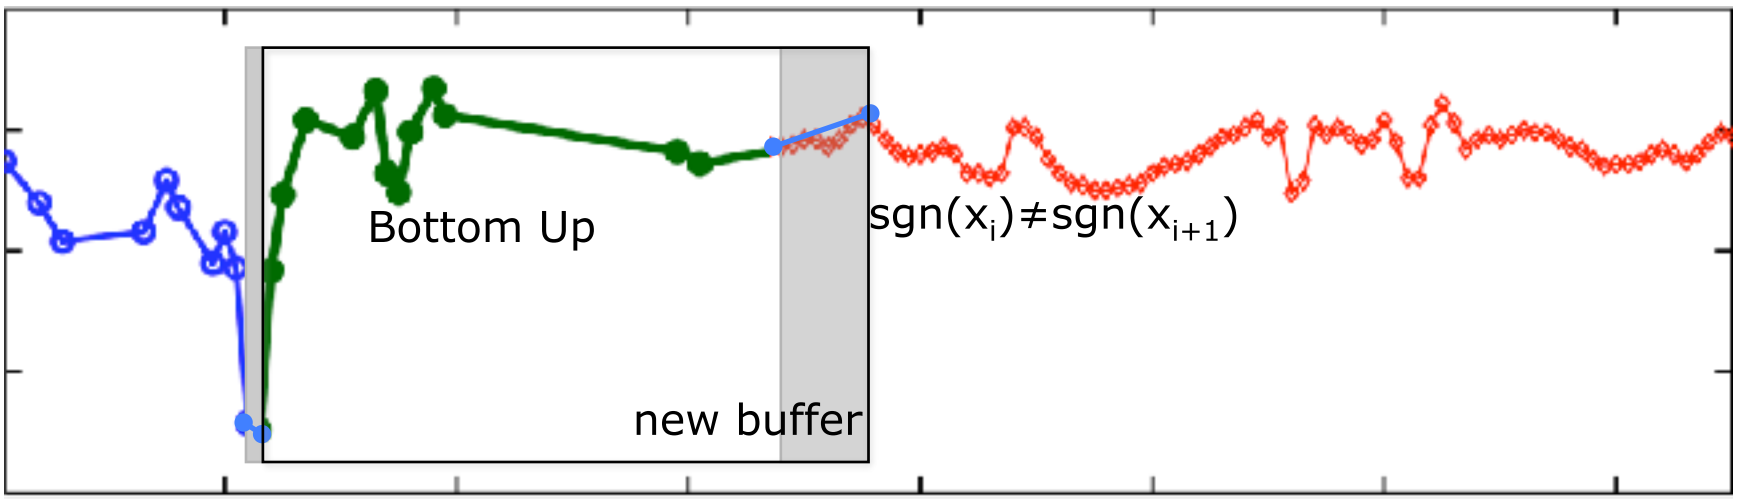

mSWAB (Berlin ’11)

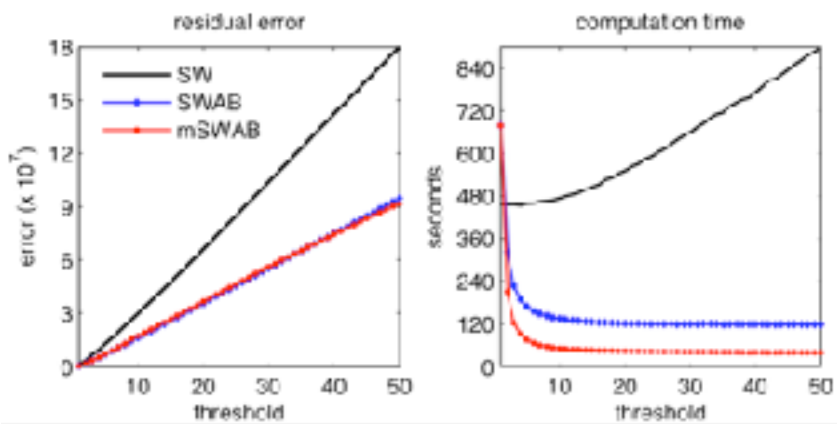

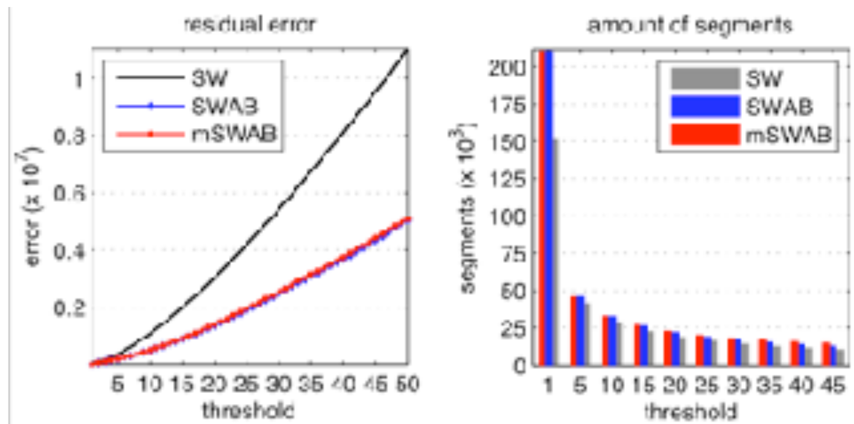

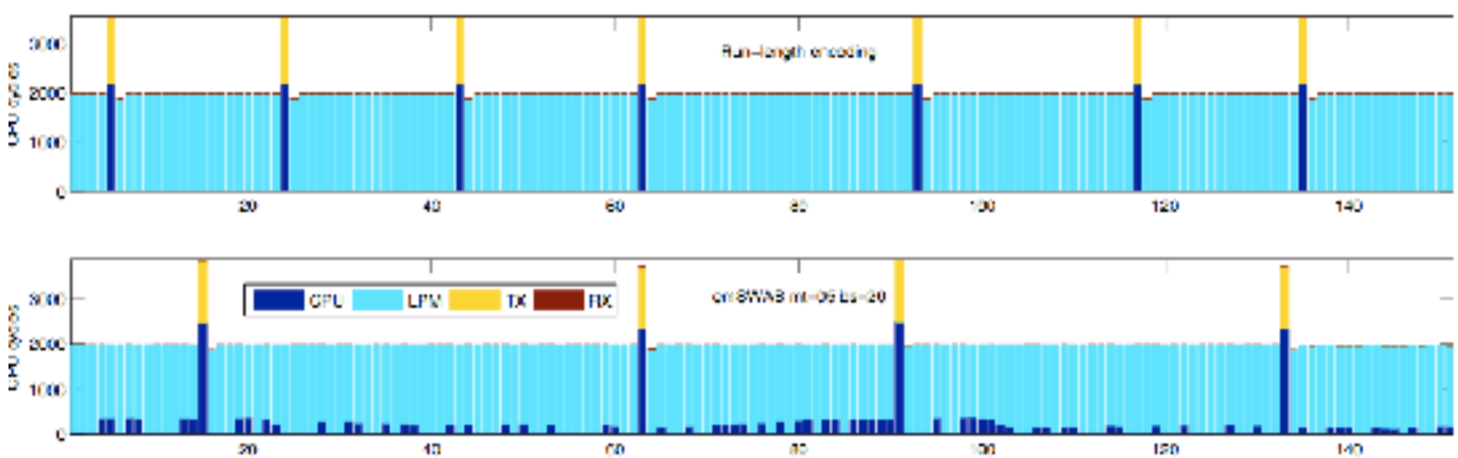

Classifying Activities: Dealing with large time series

…aabab**aabddc**daddaadbccc…

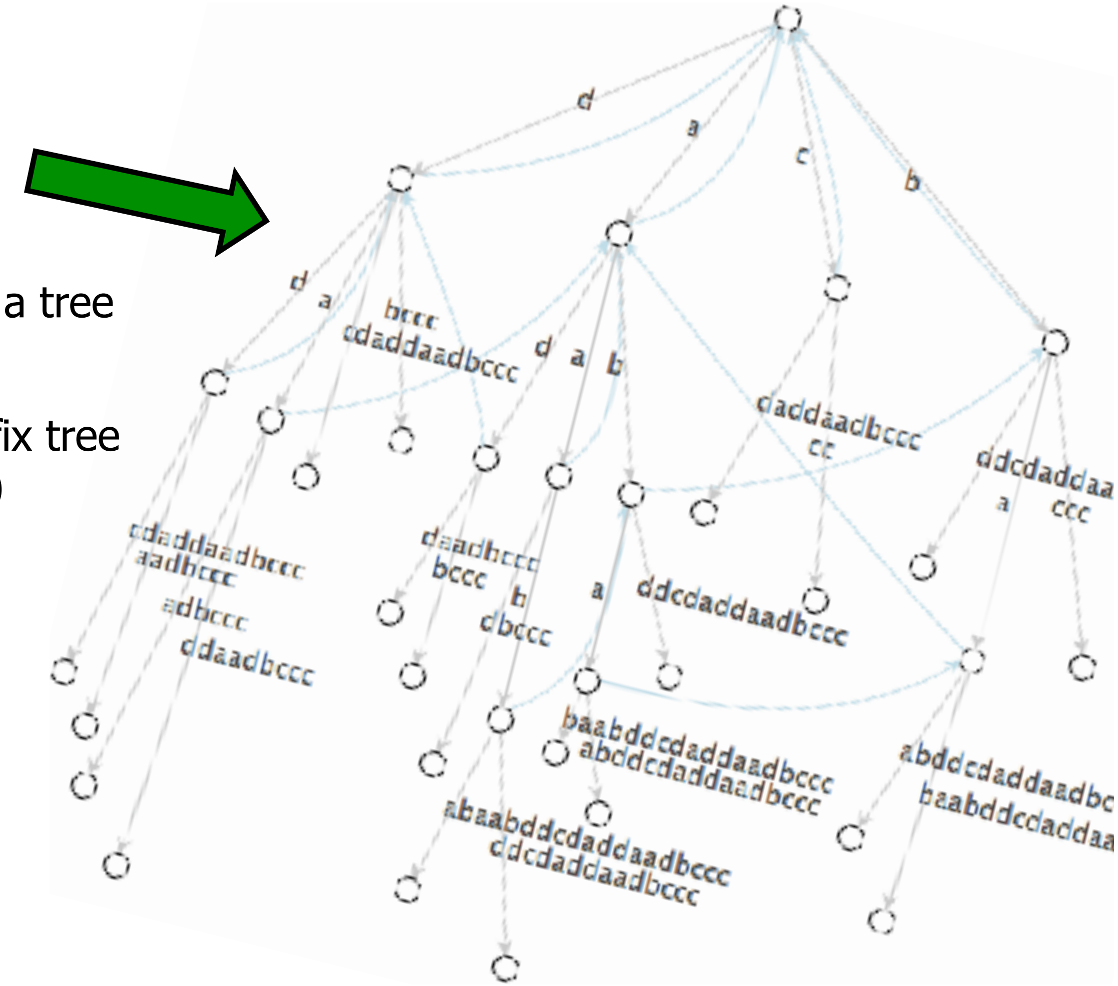

Suffix Tree: represent a string as a tree of suffixes (root -> leaf)

- It is possible to construct a suffix tree in linear time, online (Ukkonen)
- Searching substrings, quickly
- Finding out how many times they occur, quickly

Classifying Activities: Dealing with large time series

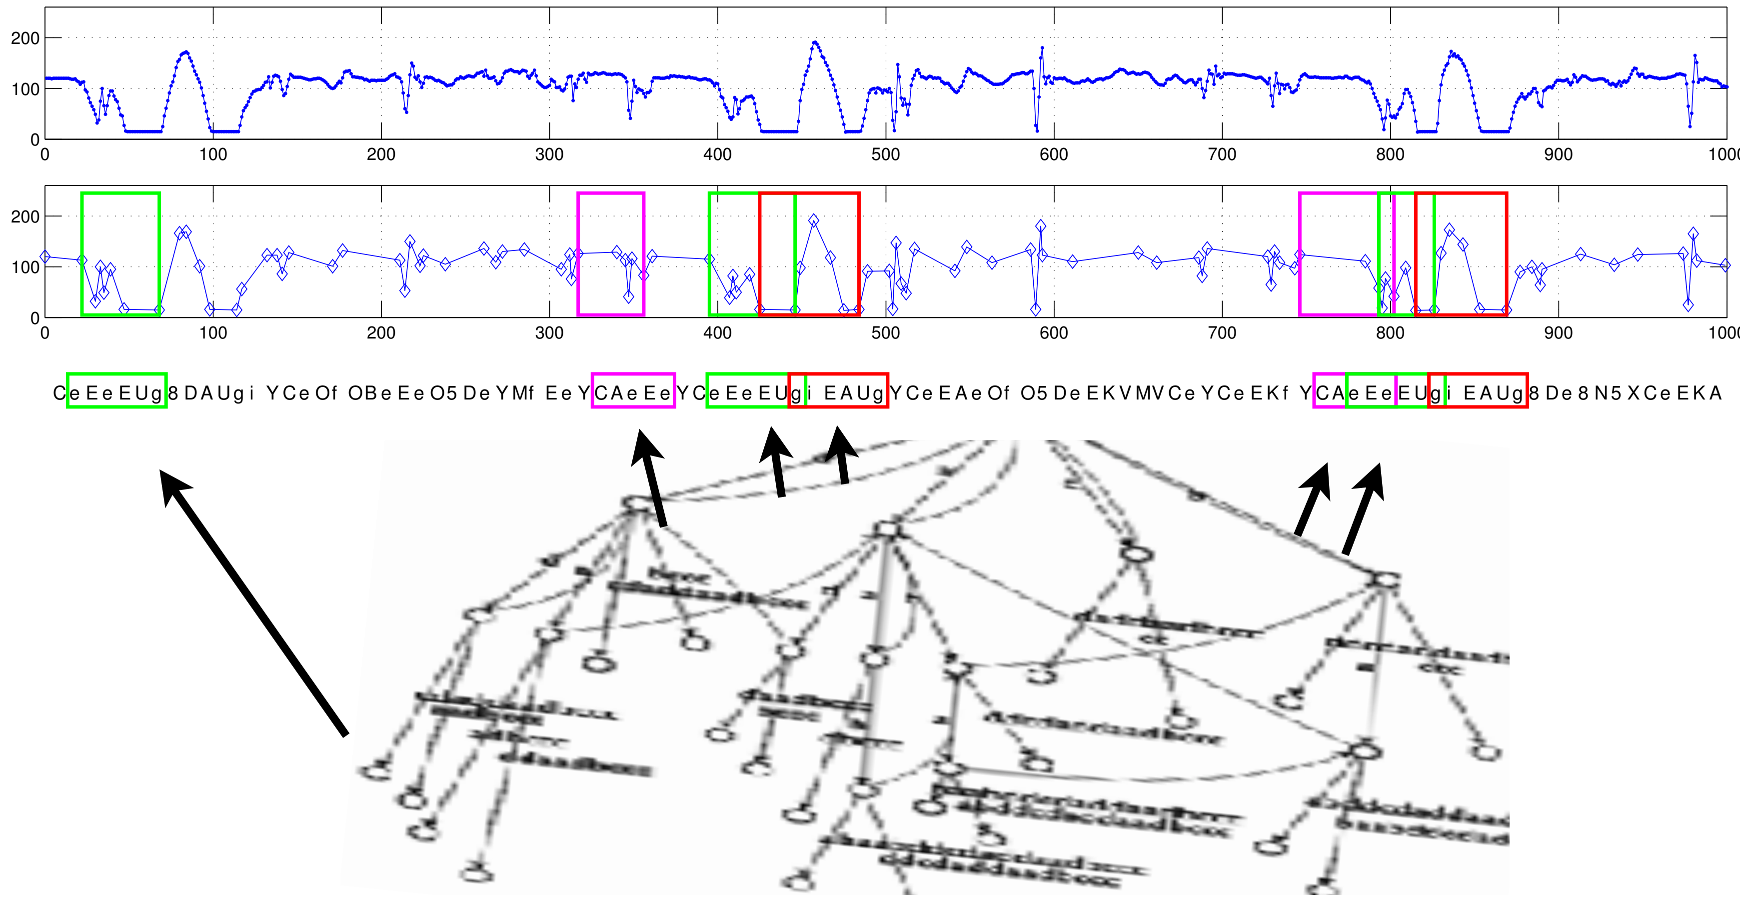

Berlin & Van Laerhoven, Detecting Leisure Activities with Dense Motif Discovery, UbiComp 2012

Classifying Activities: Dealing with large time series

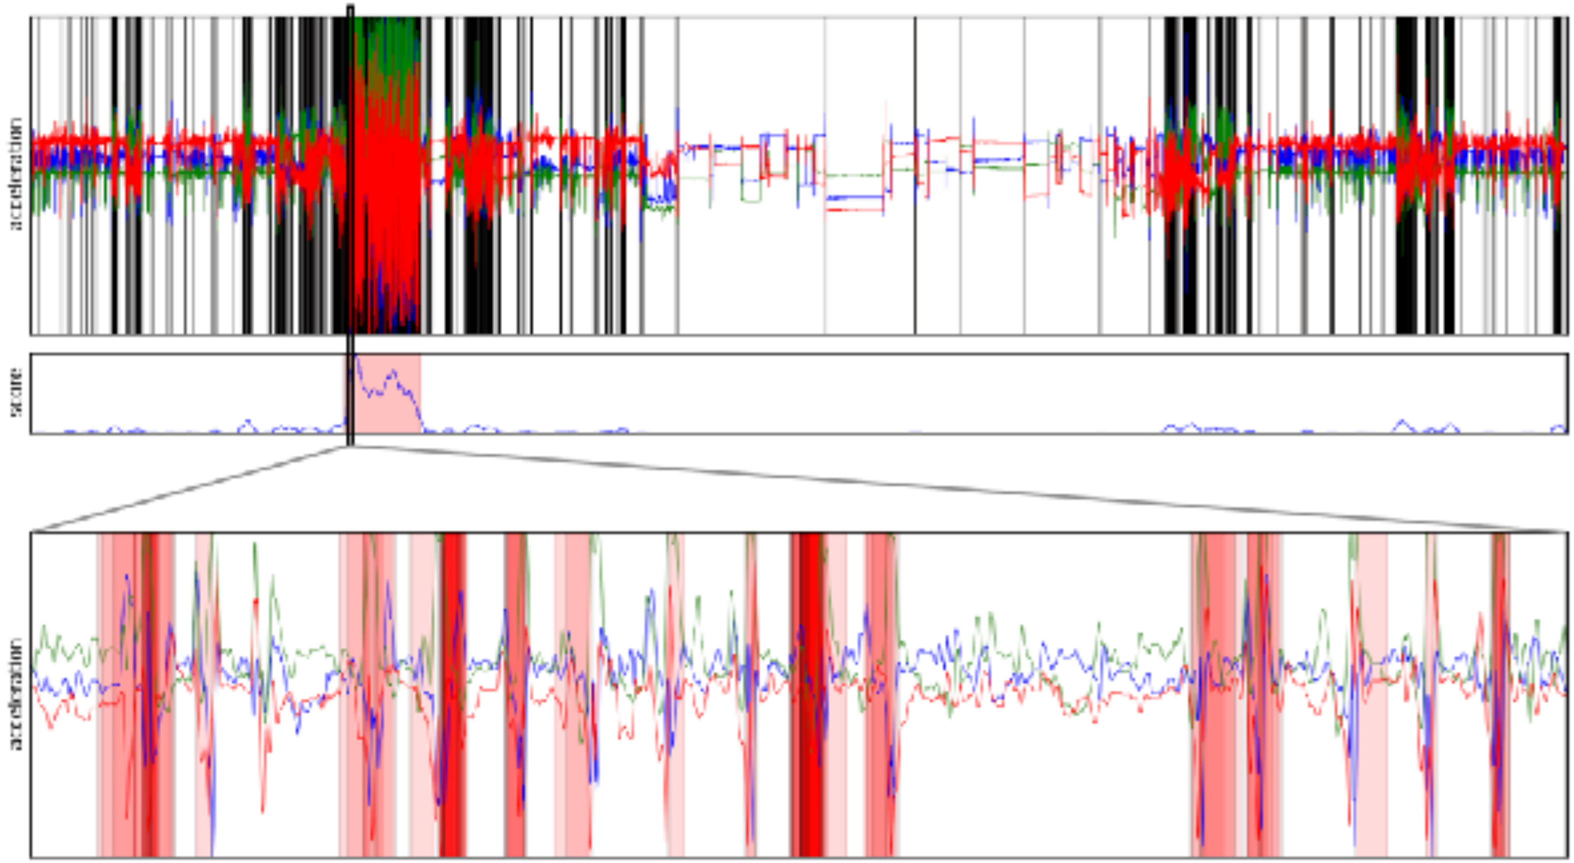

Berlin & Van Laerhoven, Detecting Leisure Activities with Dense Motif Discovery, UbiComp 2012

Classifying Activities: Dealing with large time series

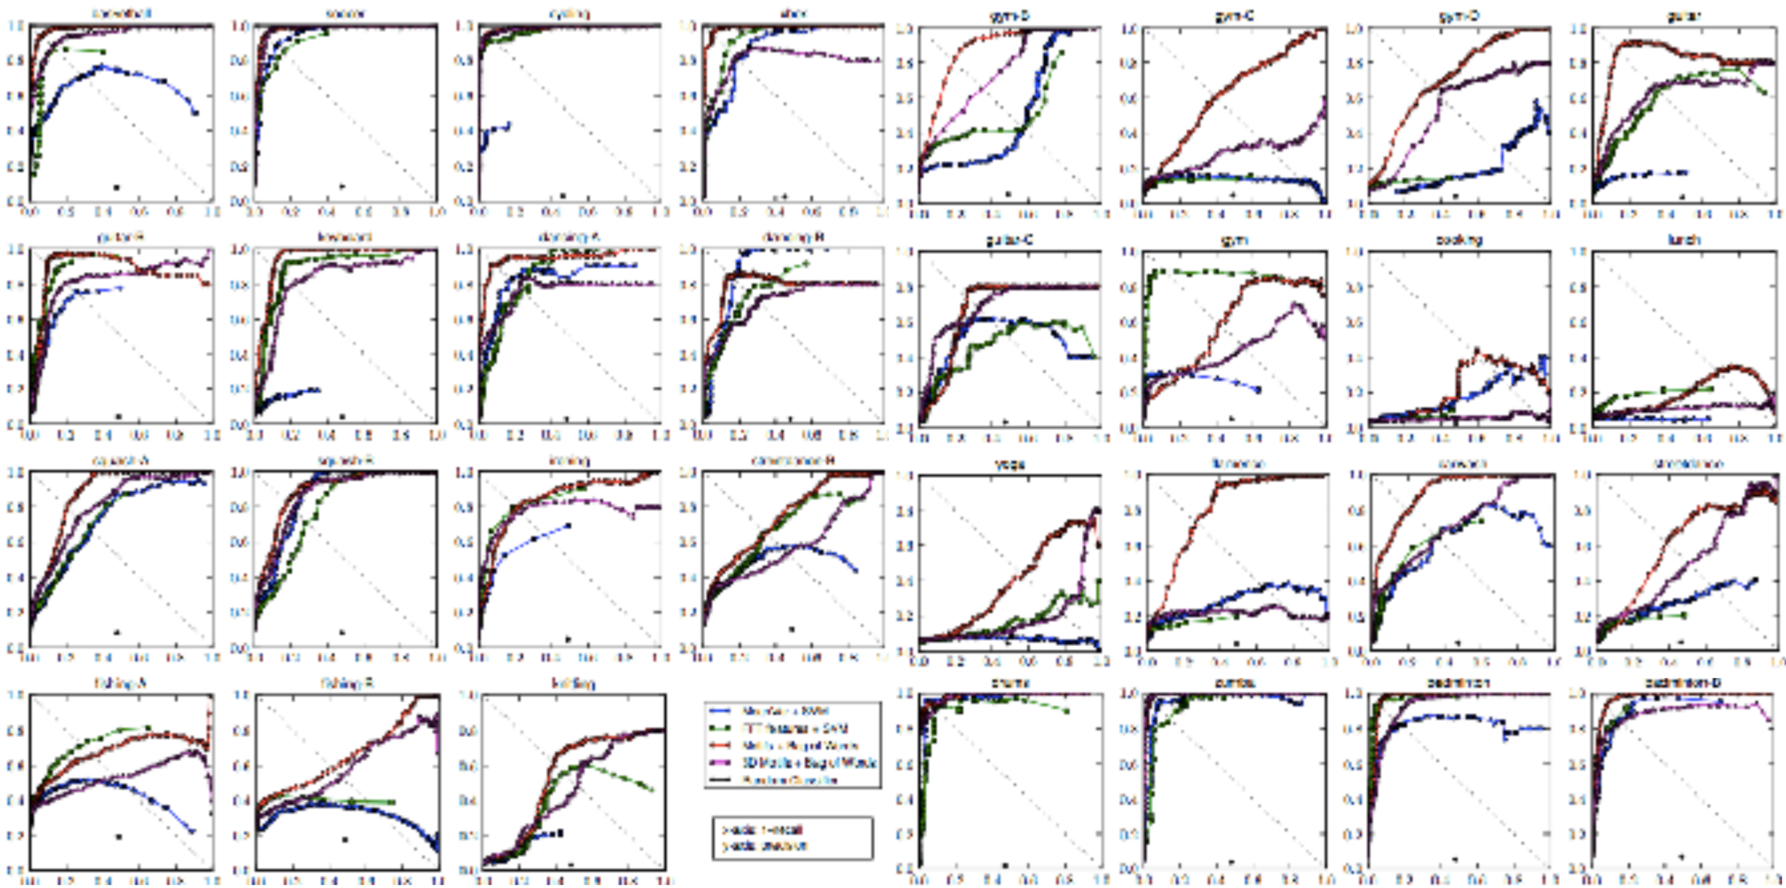

Berlin & Van Laerhoven, Detecting Leisure Activities with Dense Motif Discovery, UbiComp 2012

Classifying Activities: Dealing with large time series

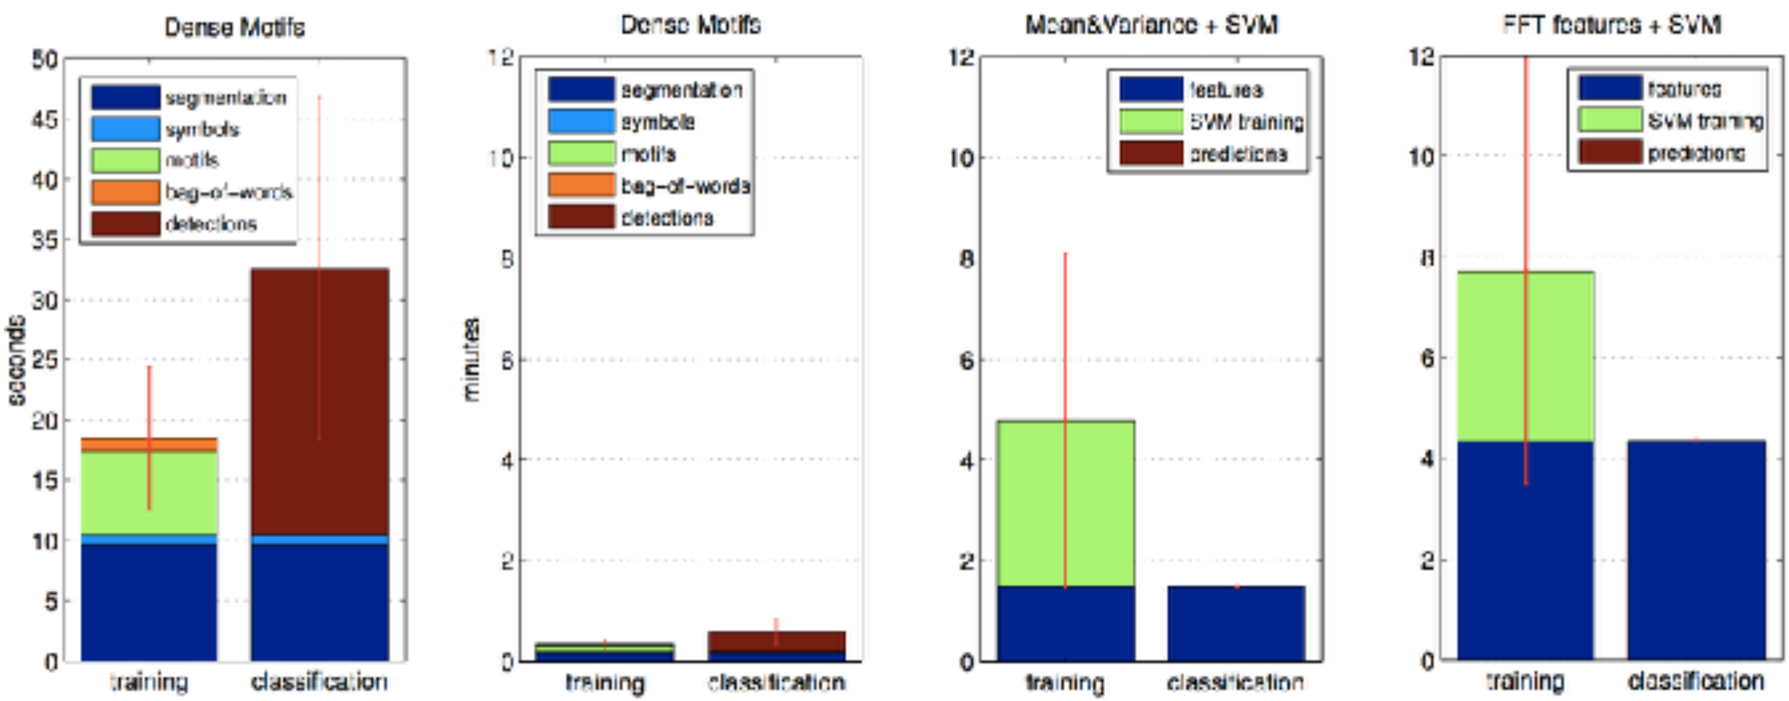

Berlin & Van Laerhoven, Detecting Leisure Activities with Dense Motif Discovery, UbiComp 2012

Outline

- Definitions of context-aware computing
- Early examples
- From Context to Activity
- Classifying activities:
  - Capturing real activities
  - Sensor coverage
  - Features
  - Classifiers
  - Evaluation of activity recognition
  - Ground truth
  - Shape matching