

Mobile and Sensor Systems

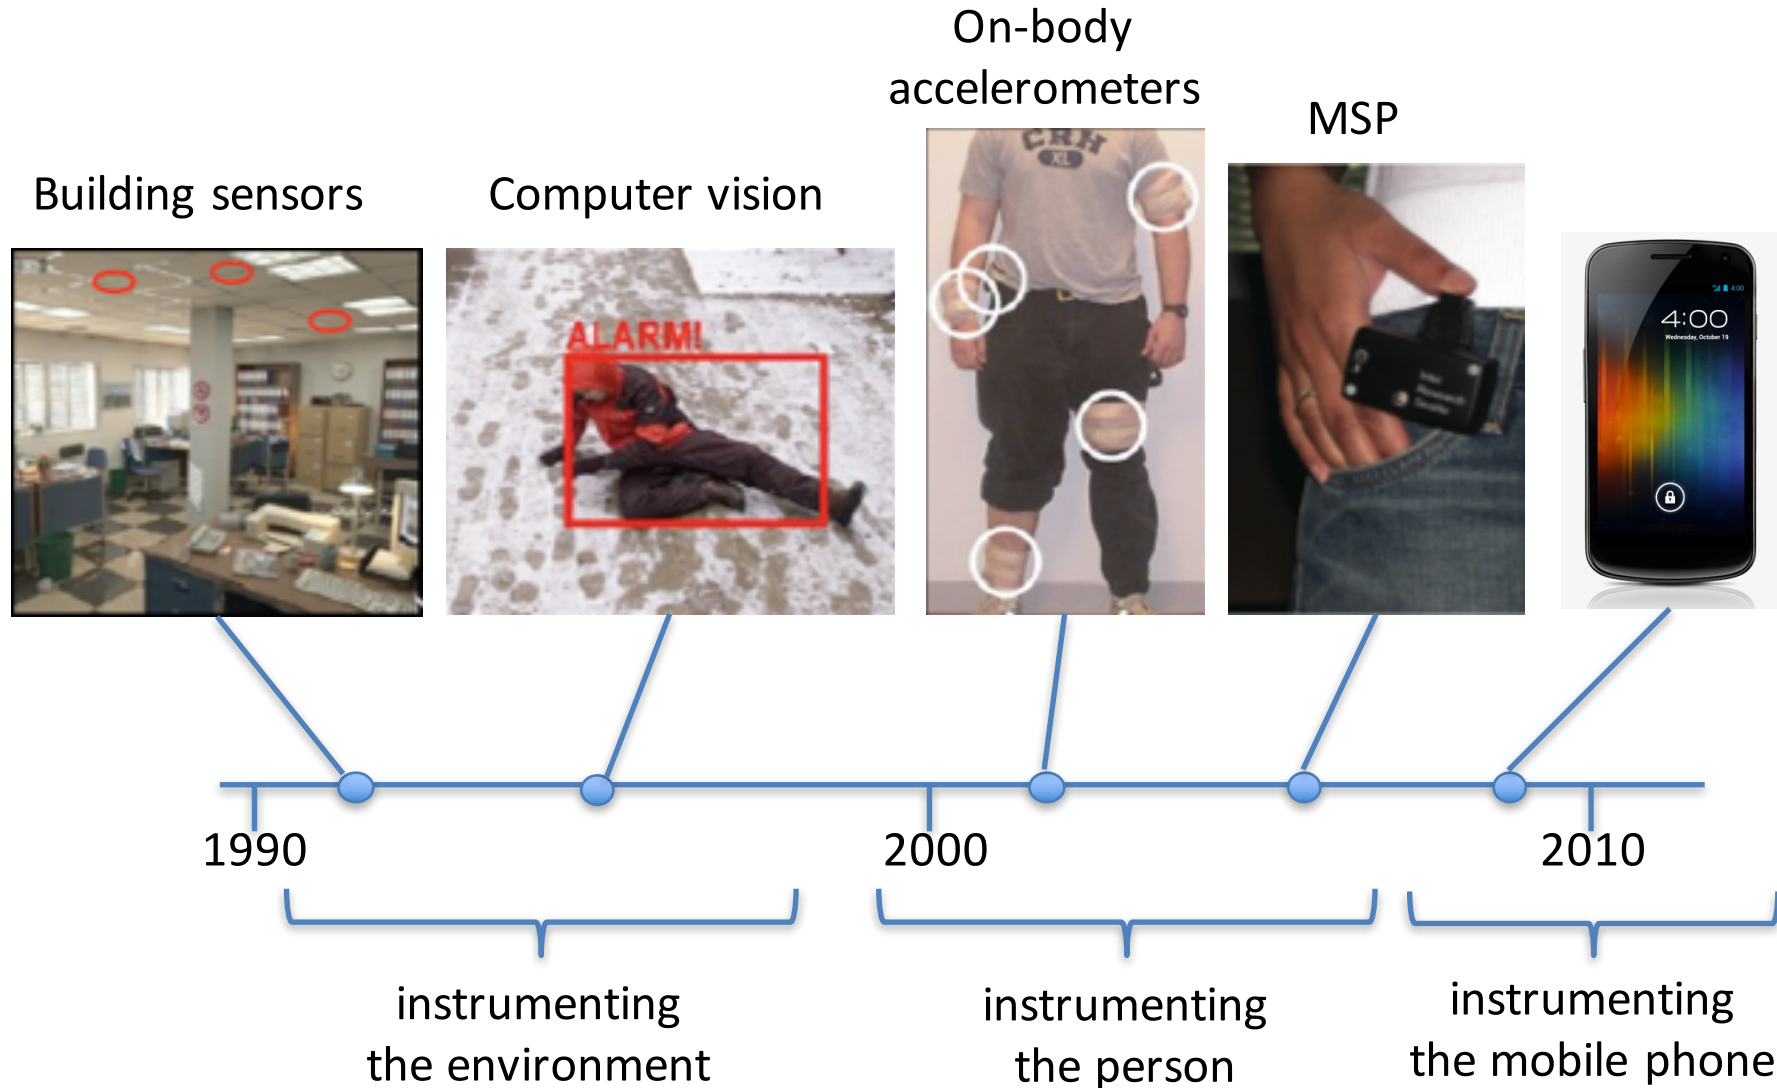
Lecture 5: Mobile Phone Sensing

Dr Nicholas Lane

In this lecture

- We will introduce mobile phone based sensing.
- We will look at the general design pattern of mobile phone sensing applications.
- We will talk about the challenges in terms of energy and activity inference.

History of Sensing Platforms



Mobile Phone Sensing

- Phone manufacturers never intended their devices to act as general purpose s




Specifications

- Sensing components considered as tools for interaction with the world
 - Accelerometer: rotation
 - Gyro: games
 - Microphone: r

Um...the Nokia N95 has an accelerometer inside??

 listen now



Apparently [the Nokia N95 has an accelerometer inside](#). You know, just like those fancy pants iPhone thingies. This is pretty freaky because it means that you're soon going to see a bunch of Symbian type applications which use auto-rotate and stuff like that. If you've got an N95 download [Moving Ball](#) and try it out yourself. You have to ask why Nokia keeps doing this kind of thing. Remember the hidden mic in the N770 and the FM radio that we suddenly found lurking inside the N800 tablet? Weird huh? The demo video is [here](#) and watch out for an

0/1900

20
RAM

Mobile Phone Sensing

- Phone manufacturers never intended their devices to act as general purpose sensing devices
- Sensing components were only considered as tools to facilitate interaction with the phone
 - Accelerometer: Screen rotation
 - Gyro: games
 - Microphone: making calls ☺



Specifications

CPU 332MHz Dual Arm 11
2G Network GSM 850/900/1800/1900
3G Network HSDPA 2100
Display TFT, 16M colours, 240x320
Memory 160MB storage, 64MB RAM
GPS
GPU 3D Graphics HW accelerator
Browser WAP 2.0/xHTML, HTML



Specifications

CPU 332MHz Dual Arm 11
2G Network GSM 850/900/1800/1900
3G Network HSDPA 2100
Display TFT, 16M colors, 240x320
Memory 160MB storage, 64MB RAM
GPS
Accelerometer
GPU 3D Graphics HW accelerator
Browser WAP 2.0/xHTML, HTML

Mobile Phone Sensing

- The mobile phone sensing domain is filled with “hacks”, and imaginative techniques that were used to circumvent the limitations of a platform that was **designed for a different purpose.**
- However, manufacturers have started to change direction
 - In the near future we expect the release of
 - New hardware platforms that facilitate back-ground sensing
 - New OS frameworks that incorporate a general purpose sensing middleware

Phone Sensing vs Sensor Networks

Sensor Networks


- Well suited for sensing the environment
- Specialized hardware designed to accurately monitor specific phenomena
- All resources dedicated to sensing
- High cost of deployment and maintenance (regular recharging thousands of sensor nodes)

Phone Sensing

- Well suited for sensing human activities
- General purpose hardware, often not well suited for accurate sensing of the target phenomena
- Multi-tasking OS. Main purpose of the device is to support other applications
- Low cost of deployment and maintenance (millions of potential users where each user charges their own phone)

But not sure if users will keep you app on their device!

Sensors



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

Applications

- **Individual sensing:**
 - fitness 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

Physical Activity

- Example Inferences:
 - {walking, running, up/down stairs}
- Sensors used: accelerometer, gyroscope, compass
- Applications:
 - Health/behaviour intervention
 - “Presence sharing”



Miluzzo et al. Sensing Meets Mobile Social Networks: The Design, Implementation and Evaluation of the CenceMe Application. Proceedings of SenSys'08.

Consolvo et al. Flowers or a Robot Army? Encouraging Awareness & Activity with Personal, Mobile Displays. Proceedings of UbiComp'08.

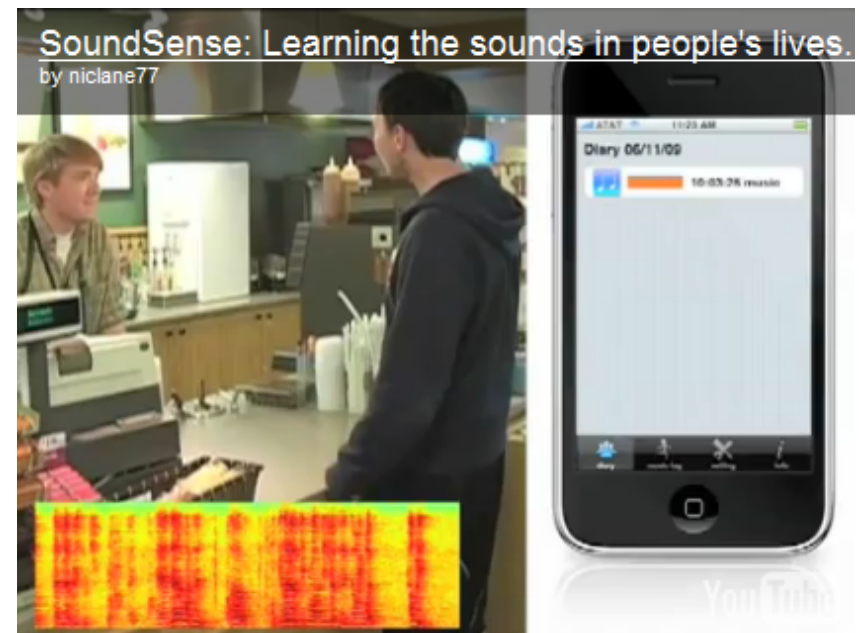
Transportation Mode

- Example Inferences:
 - {bike, bus, car}
- Sensors Used: accelerometer, GPS, WiFi, (location technologies)
- Applications:
 - Intelligent Transportation
 - Smart Commuting



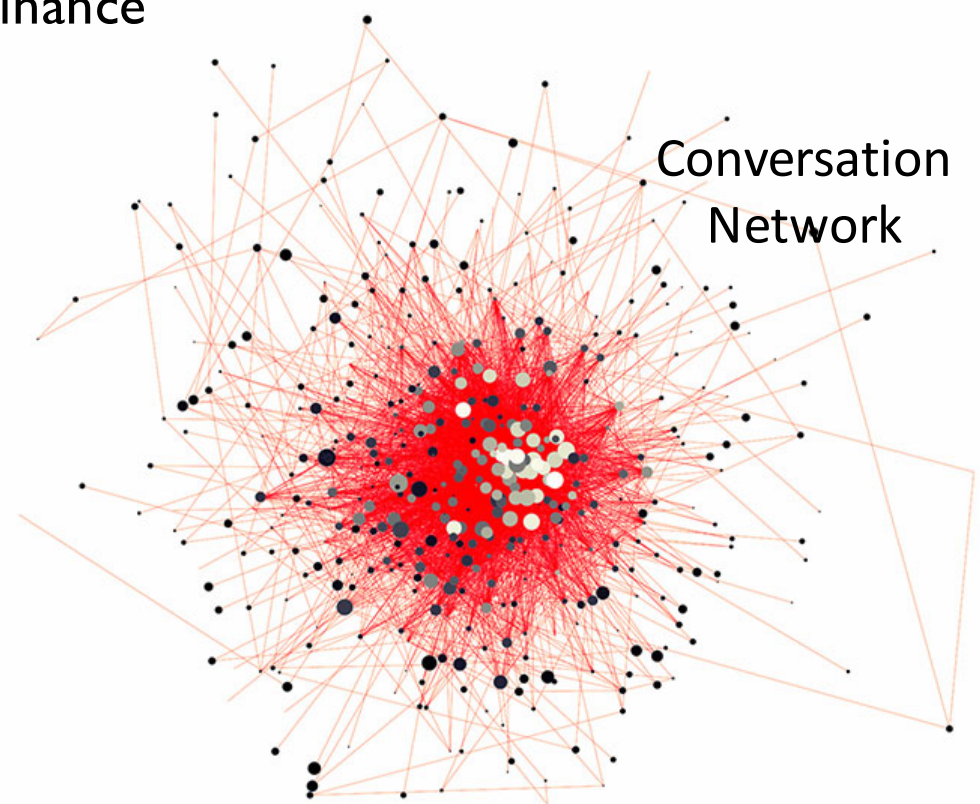
Context and Environment

- Examples:
 - {conversation, music, party, activity-related sounds}
- Sensors: microphone, camera
- Applications:
 - Automated Diary
 - Health & Wellness



Human Voice and Conversations

- Example Analysis:
 - Turn-taking, Stress, Speaker Dominance
- Sensors Used: microphone
- Applications:
 - Social network analysis
 - Stress



Detecting Emotions

- Example inference:
 - Emotional state, location and co-location with others
- Sensors used:
 - Microphone, bluetooth, GPS
 - Map speaking features to emotional state
- Application:
 - Behaviour intervention
 - Computational social science
 - Using mobile sensing for quantifying theories in social science



Mobile Systems for Computational Social Science

SOCIAL SCIENCE

Computational Social Science

David Lazer,¹ Alex Pentland,² Lada Adamic,³ Sinan Aral,^{2,4} Albert-László Barabási,⁵ Devon Brewer,⁶ Nicholas Christakis,¹ Noshir Contractor,⁷ James Fowler,⁸ Myron Gutmann,³ Tony Jebara,⁹ Gary King,¹ Michael Macy,¹⁰ Deb Roy,² Marshall Van Alstyne^{2,11}

A field is emerging that leverages the capacity to collect and analyze data at a scale that may reveal patterns of individual and group behaviors.


We live life in the network. We check our e-mails regularly, make mobile phone calls from almost any location, swipe transit cards to use public transportation, and make purchases with credit cards. Our movements in public places may be captured by video cameras, and our medical records stored as digital files. We may post blog entries accessible to anyone, or maintain friendships through online social networks. Each of these transactions leaves digital traces that can be compiled into comprehensive pictures of both individual and group behavior, with the potential to transform our understanding of our lives, organizations, and societies.

The capacity to collect and analyze massive amounts of data has transformed such fields as biology and physics. But the emergence of a data-driven “computational social science” has been much slower. Leading journals in economics, sociology, and political science show little evidence of this field. But computational social science is occurring—in Internet companies such as Google and Yahoo, and in govern-

ment agencies such as the U.S. National Security Agency. Computational social science could become the exclusive domain of private companies and government agencies. Alternatively, there might emerge a privileged set of academic researchers presiding over private data from which they produce papers that cannot be

critiqued or replicated. Neither scenario will serve the long-term public interest of accumulating, verifying, and disseminating knowledge.

What value might a computational social science—based in an open academic environment—offer society, by enhancing understanding of individuals and collectives? What are the




¹Harvard University, Cambridge, MA, USA. ²Massachusetts Institute of Technology, Cambridge, MA, USA. ³University of Michigan, Ann Arbor, MI, USA. ⁴New York University,

Challenges to Accurate Sensing

- Complex Natural Environments
- Diverse User Populations
- Data Noise
- Differences Sensor Position
- Breadth of Activities



Differences in
Sensor Position



Noisy Data



Diverse User
Populations


16

Mobile Phone Sensing Design


- Typical mobile phone sensing applications follow a common design pattern:
 - Collect raw data using the sensors of the mobile phone.
 - Infer a particular activity of interest using the sensor values:
 - physical activity: is the user running?
 - context detection: is the user in a place full of other people?
 - Expose the high-level result to the user or use that result to adapt the behaviour of the application.

Development Design Patterns

- Collect data (labelled or unlabelled)



- Inference pipeline



- Mobile Sensing App
 - Extras: storage, networking, sharing, privacy



Sensing

- Sensing is resource intensive



BATTERY



CPU



MEMORY




STORAGE

- The mobile phone's purpose is to support multiple applications.
- A mobile phone sensing application needs to maintain a balance between:
 - The amount of resources needed to operate.
 - The accuracy of the detection that is achieved.

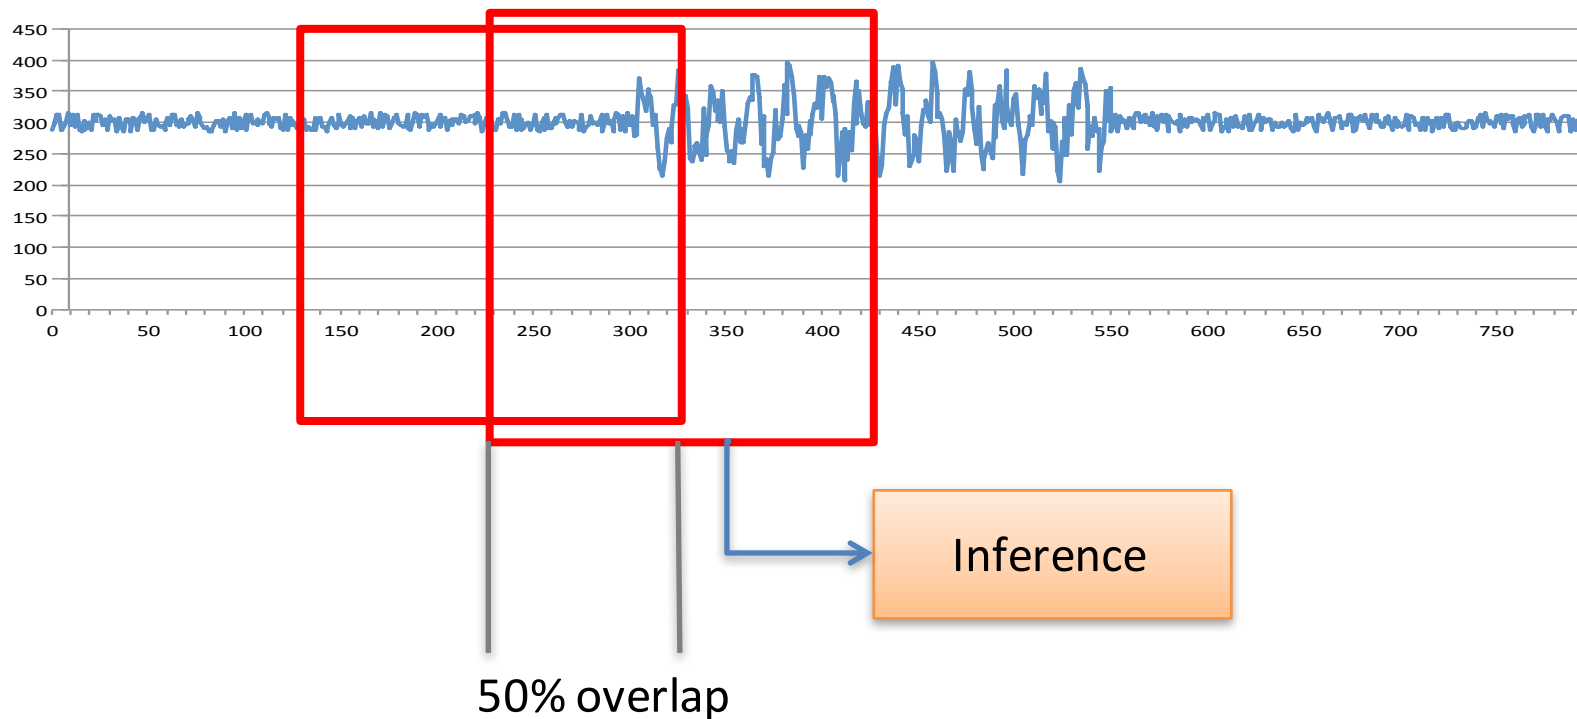
Sensing Modes

Continuous Sensing




Sensing Modes

Continuous Sensing




Sensing Modes

Continuous Sensing




Sensing Modes

Continuous Sensing




- Highly accurate data.
- Very costly in terms of battery and CPU usage:
 - Continuous sensing on multiple sensors can reduce phone stand-by to 6 hours.
 - May be used on “cheap” sensors e.g. accelerometer.


Sensing Mode Duty Cycling



Sensing Mode Duty Cycling




Sensing Mode Duty Cycling



- Lower impact on battery.
- Less accurate, interesting events may take place during sleeping.

Adaptive Duty Cycling

- Adjust the duration of sleeping periods according the rate of events that are detected.
- If no events are detected sleep for longer.
- When new events are detected reduce the sleeping time.



New event

New event

No event!



reduce sleep time reduce sleep time increase sleep time

Adaptive Duty Cycling

- Typical approaches (depending on the type of events)
 - Exponential increase – linear decrease.
 - Linear increase – exponential decrease.
- Reduces the energy cost (compared to continuous sampling).
- Maintains high accuracy (compared to duty cycling).

But

- Requires a good understanding of the application domain:
 - in conversation detection a new voice events may be followed immediately by more such events, so faster sleep-time decrease may be necessary.
 - a location change event may not be followed by an immediate new event, so slower decrease may be applicable.

Inference

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





- 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.

Feature Extraction

- Identifying features in a data set that can be used to infer a particular type of activity.
- The set of selected features depends on the type of sensor and the type of activity that is detected.
- The design process typically involves off-line analysis of training data to identify the right features for the particular inference engine.
 - Usually an iterative process where different features are tested.
- Examples:
 - Conversation detection.
 - Physical activity detection.

Feature Extraction: Conversation Detection


- Applying FFT on the audio samples, and comparing training data that are labelled as “conversation” and “non-conversation noise”.



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

Feature Extraction: Conversation Detection

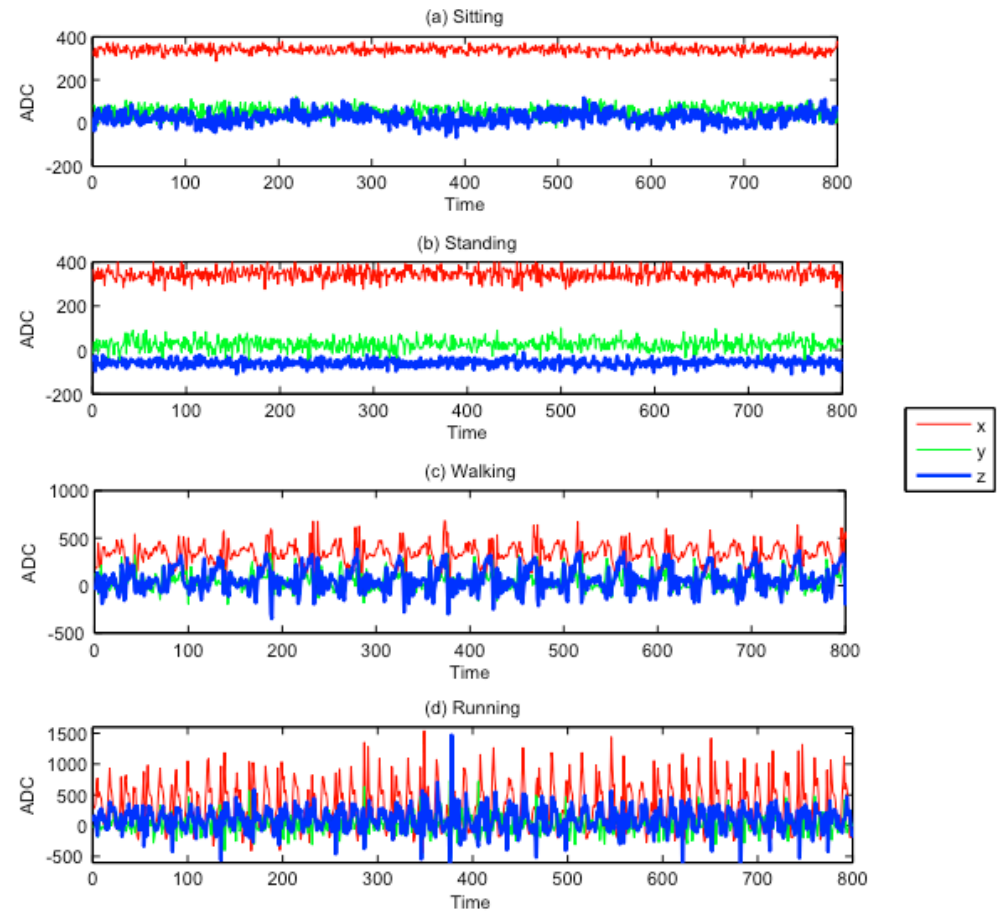
- Selecting as Features the mean and standard deviation of the FFT power



- Using a simple threshold line, could give a relatively accurate detection (with a high number of false positives, however)

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).



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:
 - Naive Bayes classifier.
 - Decision Trees.
 - Hidden Markov Models.

Naive Bayes Classifier

- Given a set of features F_1, \dots, F_n and a classifier C estimate the probability

$$p(C|F_1, \dots, F_n)$$

- This can be approximated as

$$p(C|F_1, \dots, F_n) = \frac{1}{Z} p(C) \prod_{i=1}^n p(F_i|C)$$

Where Z is a constant (scaling factor) and can be ignored in comparisons

- Using the training dataset we estimate the distributions $p(F_i|C)$
- During runtime, given a set of values for the features f_1, \dots, f_n we select a classifier that maximizes

$$p(C = c) \prod_{i=1}^n p(F_i = f_i|C = c)$$

Classification Example

- Trying to detect walking and running activities using accelerometer
- We collected 8 data sets labelled with the right class
- We select as features:
 - F1: mean acceleration
 - F2: standard deviation
- We need to calculate the distributions

$$p(F_i | C = c_j)$$

for each feature and class

Training data set

F1 Mean	F2 StdDev	Class
384.68	52.31	walking
410.24	114.39	running
392.21	71.26	walking
383.04	61.11	walking
375.32	91.01	running
399.52	109.32	running
377.36	83.01	walking
395.01	78.34	running

Classification Example

- We assume Gaussian distributions and therefore we can characterise the distributions using the mean and variance for all combinations.

	Mean F1	Var F1	Mean F2	Var F2
walking	384.32	28.12	66.92	131.23
running	395.02	160.00	98.27	207.97

- With these calculations, given a new set of values for F1 and F2 we can estimate the probability that the user is walking or running
- Under the Gaussian distribution assumption this is given by:

c_w :walking
 μ_w :mean
 σ_w^2 :variance

$$P(F_1 = x | c_w) = \frac{1}{\sqrt{2\pi\sigma_w^2}} e^{-\frac{(x-\mu_w)^2}{2\sigma_w^2}}$$

Classification Example

- The classifier is ready and we can run it in our application:
- A new sensor sample is analysed and features are extracted.
- Assume a new input with features $F_1 = 391.2$ and $F_2 = 58.5$
- The classifier calculates:

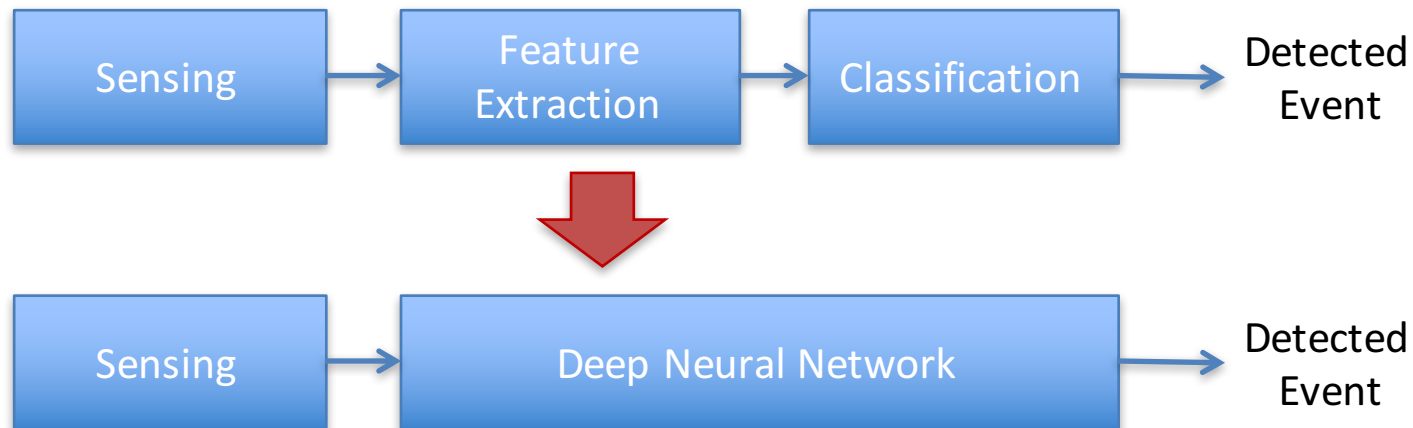
$$p(C = \textit>walking} | f_1, f_2) = 1.21e - 03$$

$$p(C = \textit>running} | f_1, f_2) = 2.71e - 05$$

and selects the class with the highest probability: walking

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



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


“Sad”

.....→ Inference




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

Example of Feature Representation Learning

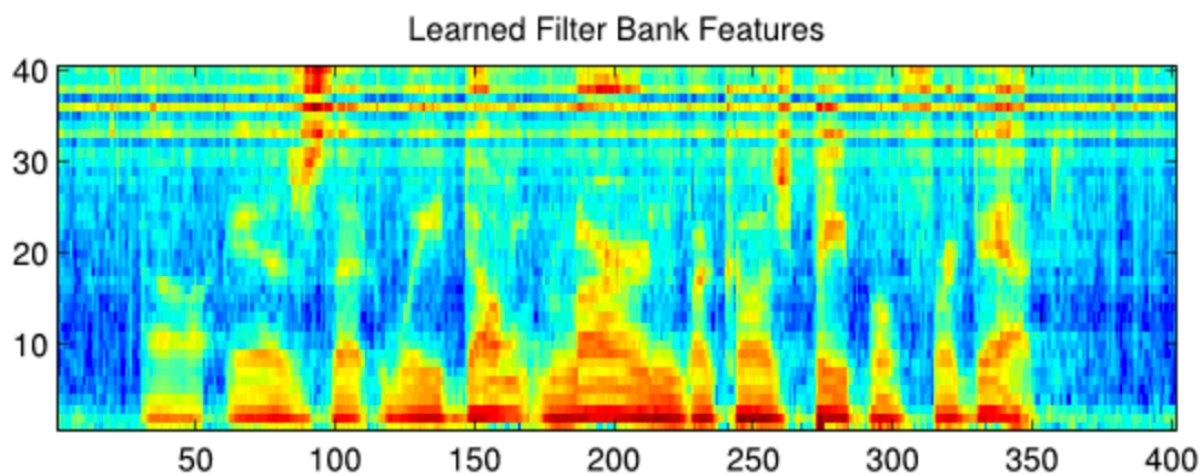


(1) Mel frequency cepstral coefficients

- Result of decades of research into audio
- Dominant general purpose audio representation




(2) Representation learned directly from data



Inference Optimizations

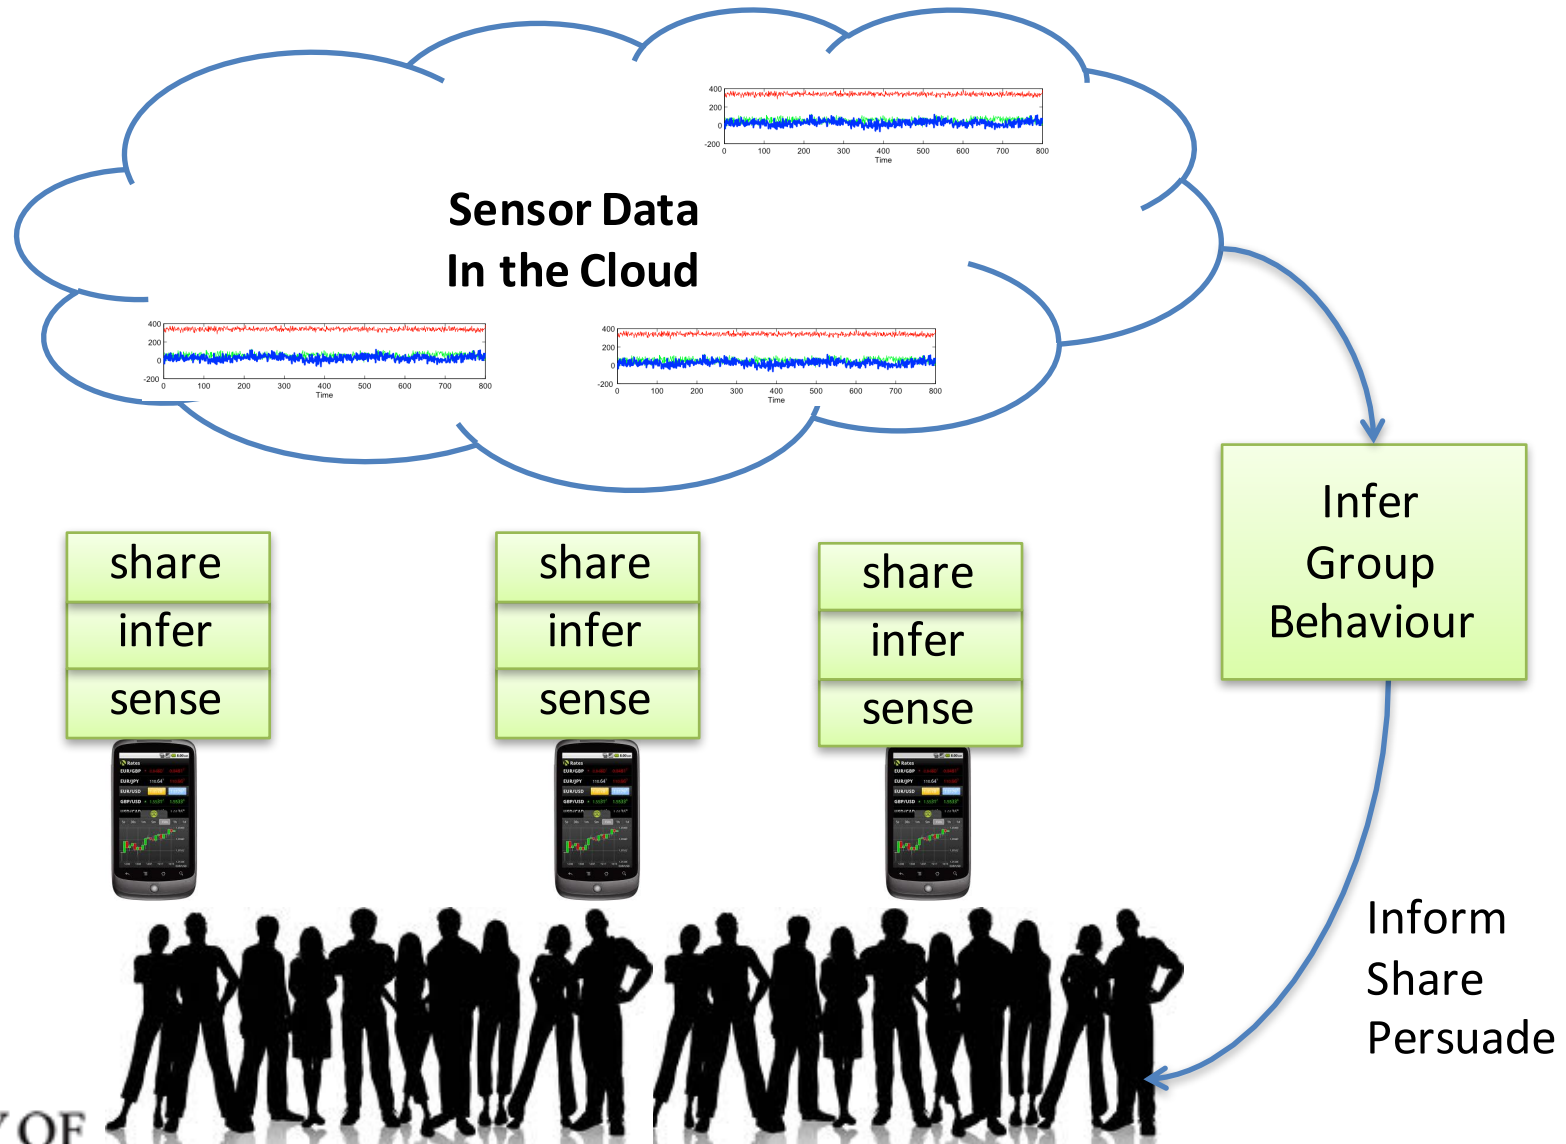
- Adaptive sampling can bring down the energy cost but inference can also be costly.
 - Example: running a speech recognition engine on the phone can have significant impact on the phone's battery life.
- Another aspect is speed of computation.
- Offloading parts of the energy cost to the cloud.



Computation Distribution

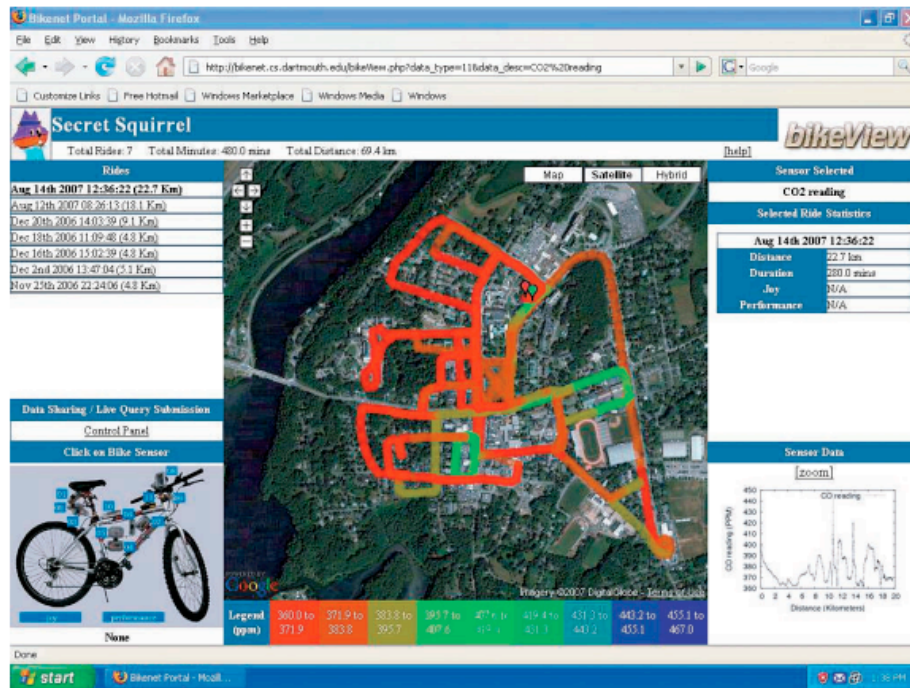
- Challenges:
 - Balance computation energy cost versus network traffic cost.
 - Balance local delay versus remote delay.
- Traffic:
 - Sending raw sensor data may cost more in network energy than what is saved.
- Solution: Split computation.
 - Perform feature extraction on the phone.
 - Perform classification in the cloud.
- Adaptive computation distribution:
 - Decide best place to do computation dynamically.
 - Estimate the cost of off-loading on the fly.

Participatory / Opportunistic Sensing



Participatory Sensing Applications

BikeNet




Mappiness



when are we happy?

The hedonimeters on the right display mappiness users' happiness in real-time, compared against the all-time average. Below, happiness levels are charted hour by hour over the past week.



Click and drag to zoom. Double-click to zoom back out.

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.
- E. Miluzzo, N.D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S.B. Eisenman, X. Zheng, A. Campbell. Sensing meets mobile social networks: the design, implementation and evaluation of the CenceMe application. Proceedings of SenSys'08.
- K.K. Rachuri, M. Musolesi, C. Mascolo, P.J. Rentfrow, C. Longworth, A. Aucinas. EmotionSense: A Mobile Phones based Adaptive Platform for Experimental Social Psychology Research. Ubicomp'10. September 2010.