R249
Advanced Topics in Mobile Systems and Mobile Data Machine Learning

Prof Cecilia Mascolo, Dr Jagmohan Chauhan

UNIVERSITY OF CAMBRIDGE

Prof Cecilia Mascolo

- Mobile Systems
- Mobile Data Analysis
- Mobile Health

Dr Jagmohan Chauhan

- Mobile Systems
- Mobile Sensing
- Mobile Health
- Usable Security
- On Device Learning

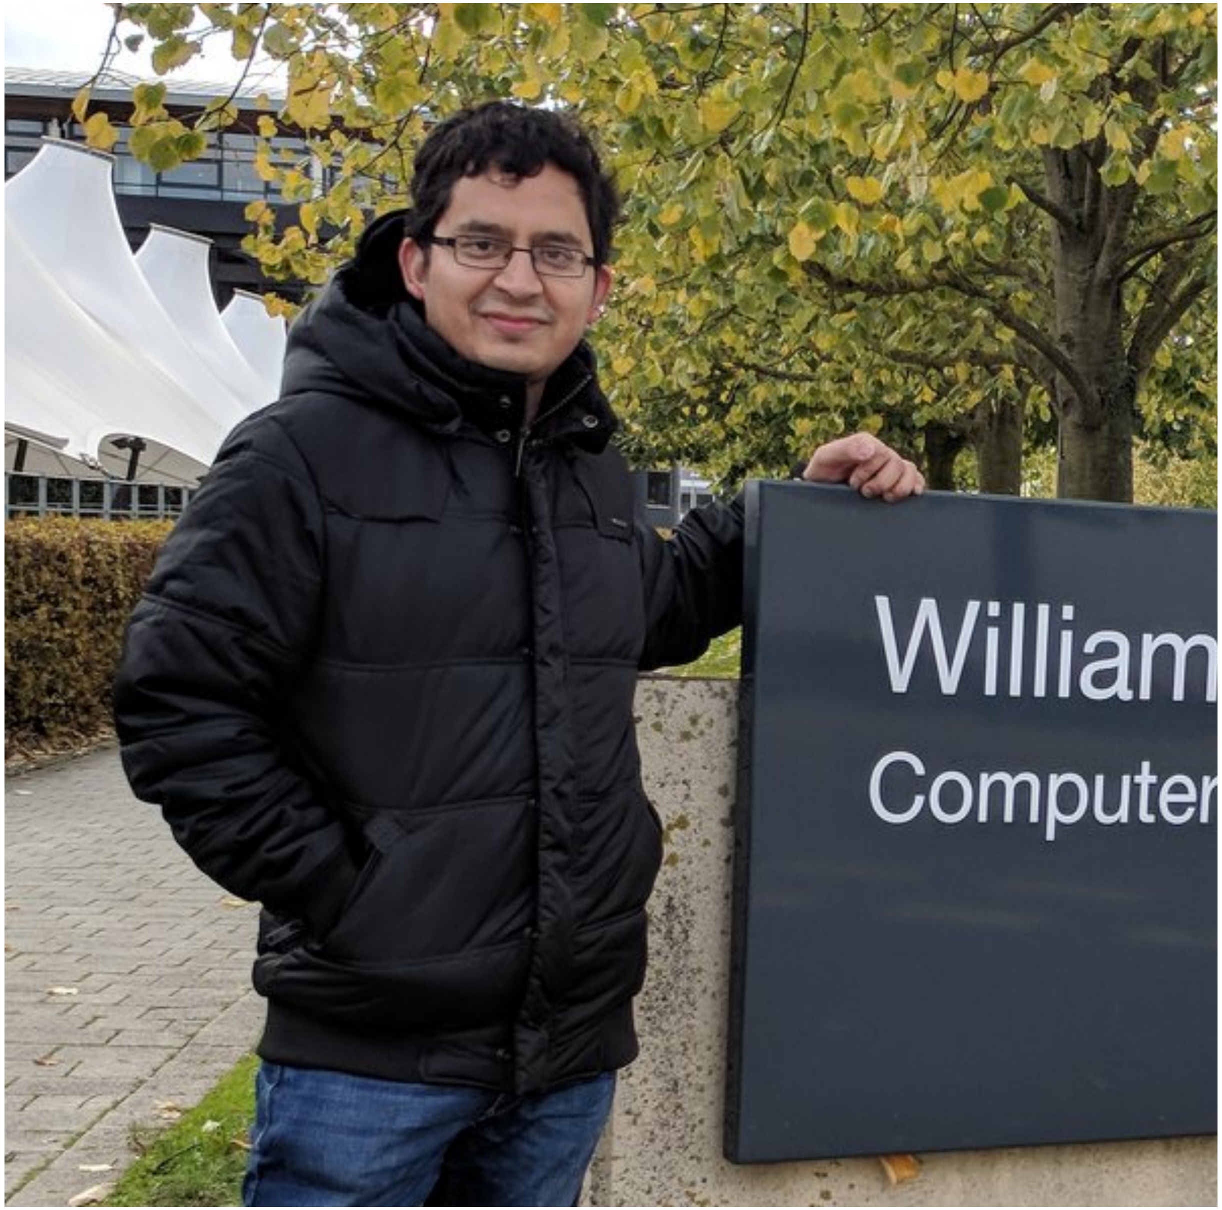

The course

- The course is about anything to do with mobile systems
    - Systems aspects including power, computation
    - Novel Sensing aspects
    - Mobility/Sensor Data learning aspects
    - Modelling and Inference On Device
    - Mobile Health

The Schedule

- 15th October(1h) Introduction (TODAY!)
- 22nd October System, Energy and Security
- 29th October Activity Recognition with Machine Learning and Mobile Sensor Data
- 5th November On Device Machine Learning
- 12th November Backscatter Communication, Battery Free and Energy Harvesting Devices
- 19th November New Sensing Modalities
- 26th November Mobile Systems and Data for COVID-19
- 25th January (3h) Mobile Health

Assessment

- A total of 7 items of assessment:
  - 1-2 Presentations
  - 5-6 Reports
- Each contributing 1/7
- A contribution tick for class attendance and participation
- A class list of attendance will be kept and apologies for absence should be sent to the lecturers prior the lecture.

Written Reports

- **In the weeks when a student is not presenting**
- Student assigned a paper among the ones listed to be presented for the following week.
- Write no more than 1000 words (recommendation would be for ~800 words report).
- A template list of headings online
- All in PDF please in Moodle!
- [Students presenting will submit slides and a video presentation instead of a report]

Form

- Paper Report Summary of the paper (200 words)
- Discussion on novelty of the paper as stated (200 words)
- Positives of this Paper (100 words)
- Negatives of this Paper (100 words)
- Ideas for Future Work, Critical discussion of potential impact and context setting (200 words)

How to Read a Paper

- https://www.cl.cam.ac.uk/teaching/2021/R249/report-guidelines.pdf

- Summary of the paper and key findings: Describe what the paper is about, the key problems it is trying to solve, its motivation (and maybe why it is an important problem) and the key contributions the paper spells out. Note that this is probably not the right place for your subjective views about the contributions.

- Discussion on novelty : Novelty of the contribution wrt to literature. Note that if the paper is not extremely recent, the novelty needs to be put in the context of the time at which the paper is published. You will want to comment on the novelty at the time as well as contextualize with respect to the current literature. Here is your chance to comment on the contribution value with a more subjective angle.

How to Read a Paper

- Positives of this Paper: Things to note, for instance, are if the paper is seminal, in the sense that works seem to have been citing this a lot, if it is very novel, if it has a thorough evaluation. Note that is often hard to be positive about a paper than finding flaws: remember to consider the difficulty of getting the work done and presented when you judge.

- Negatives of this Paper: Here is where you can be critical and highlight the limitations of the work. Is the novelty limited? Is the evaluation constrained or artificial? Is the writing difficult? Note that highlighting more negatives than positives does not mean higher marks for your report. It always depends on what you write and how you justify it.

How to Write a Report

- Ideas for Future Work: Critical discussion of potential impact and context setting. This is the space where you describe what potential the paper has. It might be that you have already set the paper into context in the novelty section so you can link to that and discuss more about the impact achieved and the future potential. If the paper is recent you can speculate on the take up of the research community or industry. This is really the space for your more subjective speculations and views.
- Write concisely and precisely
- Use scientific arguments

Presentation

- Each student will present 1-2 times
- No report when presenting
- Submission of slides (in PDF) and **a video** of 20 mins (no more) in **mp4**
  - **Don’t send too big files please (try to reduce sampling before sending)**
- Students assigned randomly each week
- Presentations will be assessed for technical content, clarity, engagement, timeliness and question answering
- Presentation will be **streamed** during the lesson

- All classes will be online and recorded
  - Recordings only shared with the students in the class and only kept for the duration of the course.

What do I put in the slides?

- https://www.cl.cam.ac.uk/teaching/1920/R249/presentation-guidelines.pdf
- Structure similar to a report in terms of what to cover however remember your audience: some students have not read the paper as carefully as others (because assigned to other papers)!
- Slides Format and Content: Remember that your slides are not your script. Use both channels (your talk and your slides)
- Keep to the time!
- Rehearse! Think of presentations you liked (or not liked!)
- Use silence and pauses…
- Q&A: don’t be defensive. Do right by the authors.

Report and Slides/Video Deadlines

Michaelmas Term Deadlines:

- Assignment 1 due Wednesday 21 October, noon
- Assignment 2 due Wednesday 28 October, noon
- Assignment 3 due Wednesday 4 November, noon
- Assignment 4 due Wednesday 11 November, noon
- Assignment 5 due Wednesday 18 November, noon
- Assignment 6 due Wednesday 25 November, noon

- Lent Term Deadlines:
  - Assignment 7 due Friday 22 January, noon

The Papers!

- http://www.cl.cam.ac.uk/teaching/2021/R249/materials.hml
- http://www.cl.cam.ac.uk/teaching/2021/R249/paper-assignment.txt

About the group’s research…

- Devices for Behaviour Monitoring
- Wearables and Mobile Systems and Data for Health
- Mobile data analysis for Urban Planning
- Audio for Health Diagnostics
- On Device Machine Learning

Devices for Behaviour Monitoring

Wearables for Indoor Interaction Monitoring

Tracking serendipitous interactions: How individual cultures shape the office. C. Brown, C. Efstratiou, I. Leontiadis, D. Quercia, C. Mascolo. In Proceedings of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW 2014). Baltimore, Maryland, USA. February 2014.

Face to Face Interactions

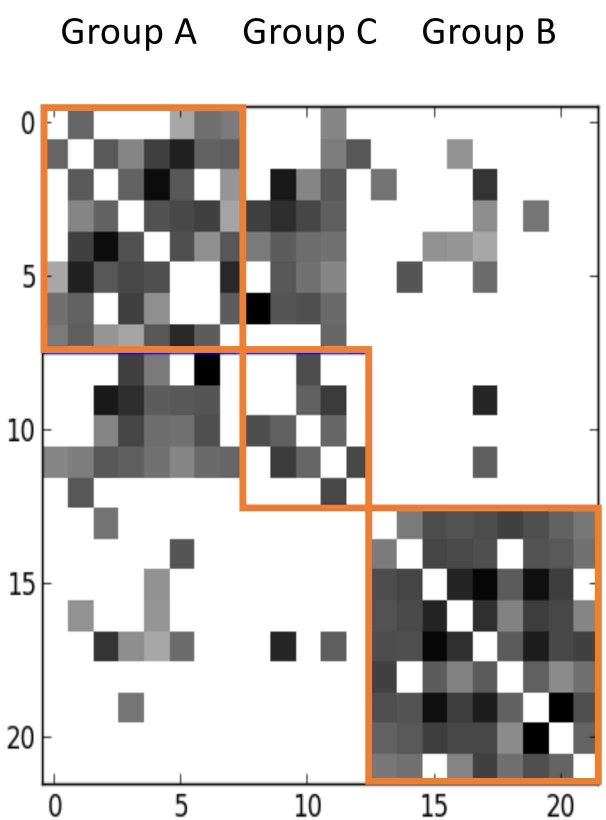

Old building

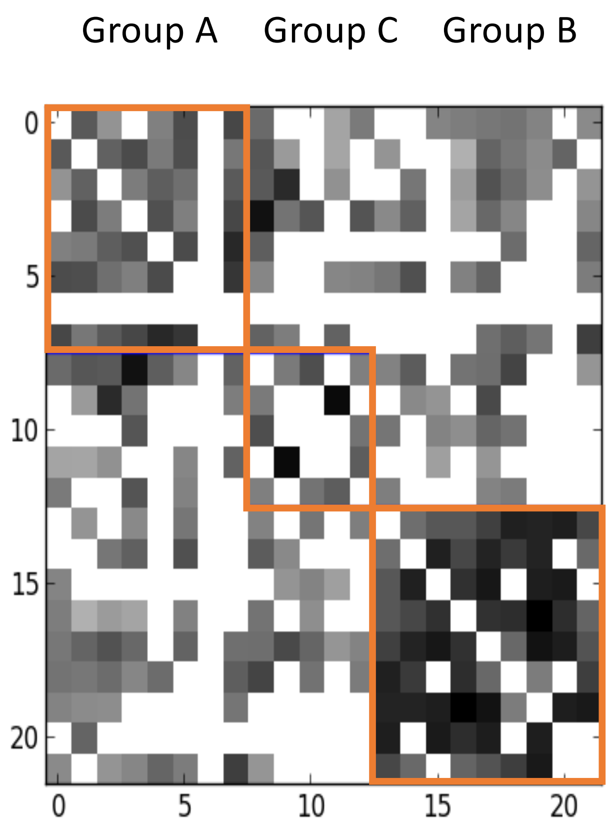

New building

The architecture of innovation: Tracking face-to-face interactions with ubicomp technologies. C. Brown, C. Efstratiou, I. Leontiadis, D. Quercia, C. Mascolo, J. Scott, P. Key. In Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (Ubicomp 2014). Seattle, WA, USA. September 2014.

Bluetooth and interactions

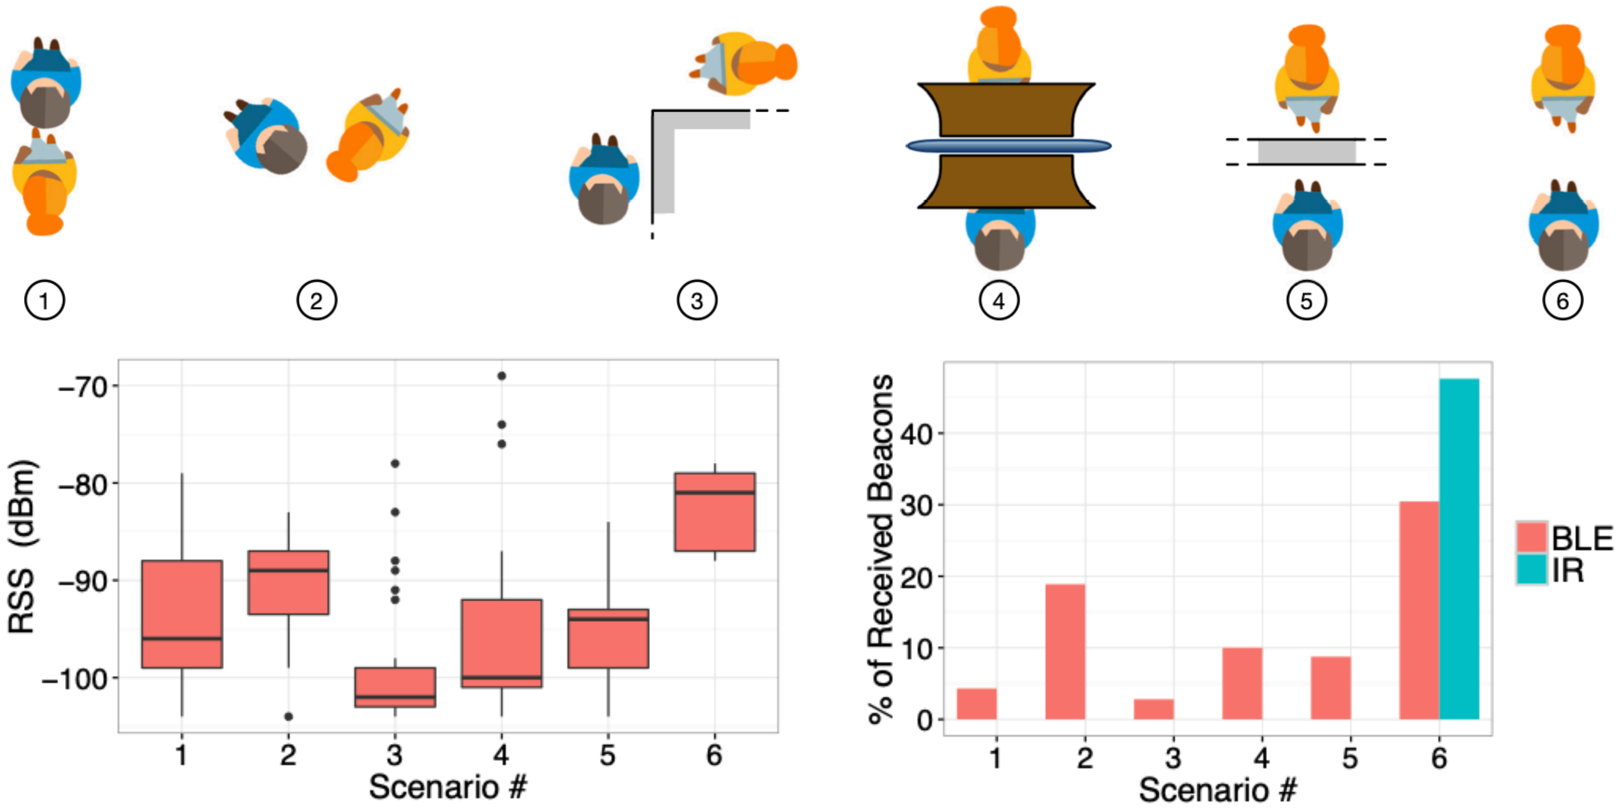

(a) (b)

Measuring Interaction Proxemics with Wearable Light Tags. A. Montanari, Z. Tian, E. Francu, B. Lucas, B. Jones, X. Zhou, C. Mascolo. In Procs of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT). Volume 2(1). 2018


Measuring Animal Behaviour

Tagging Animals
WILDSENSING

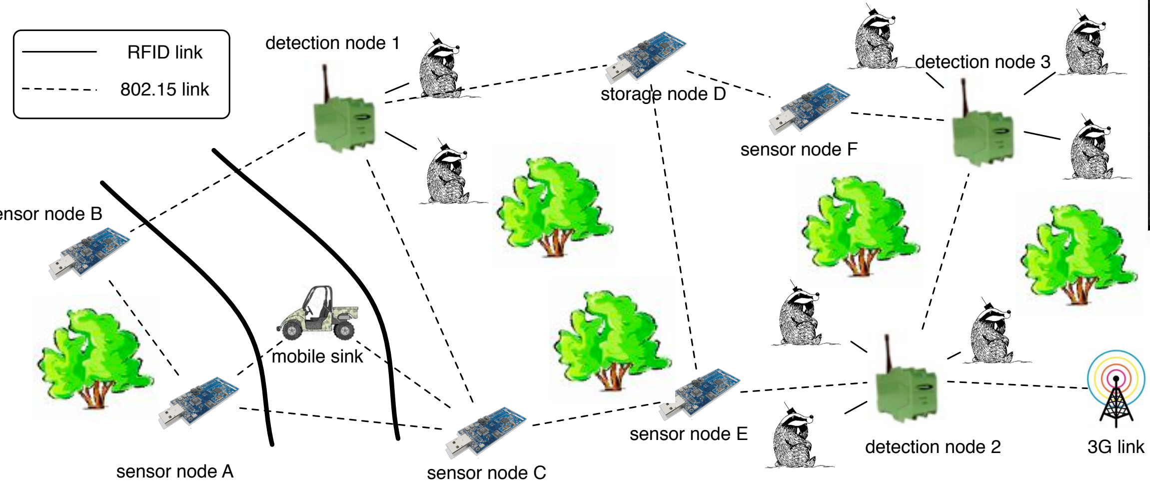

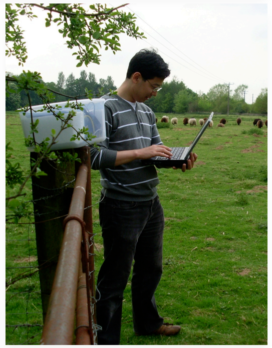

Understanding Animal Movement

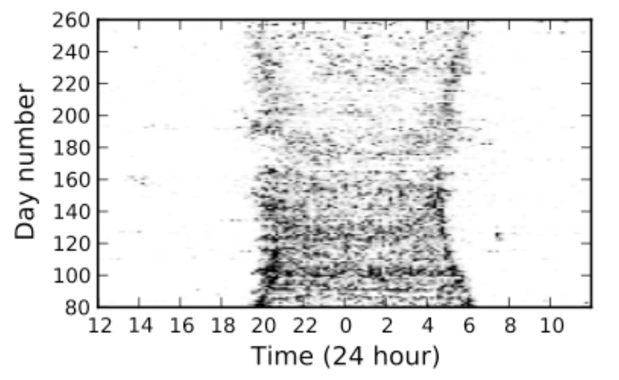

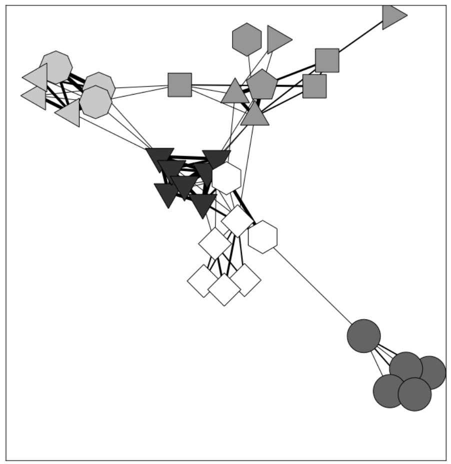

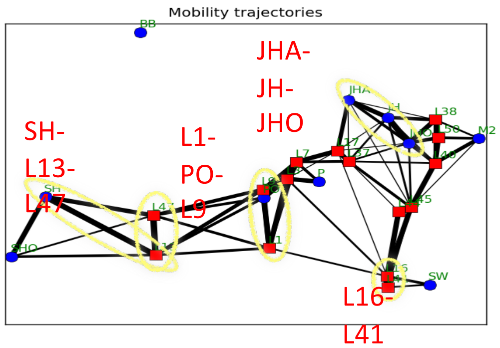

Multi-modal Sensing

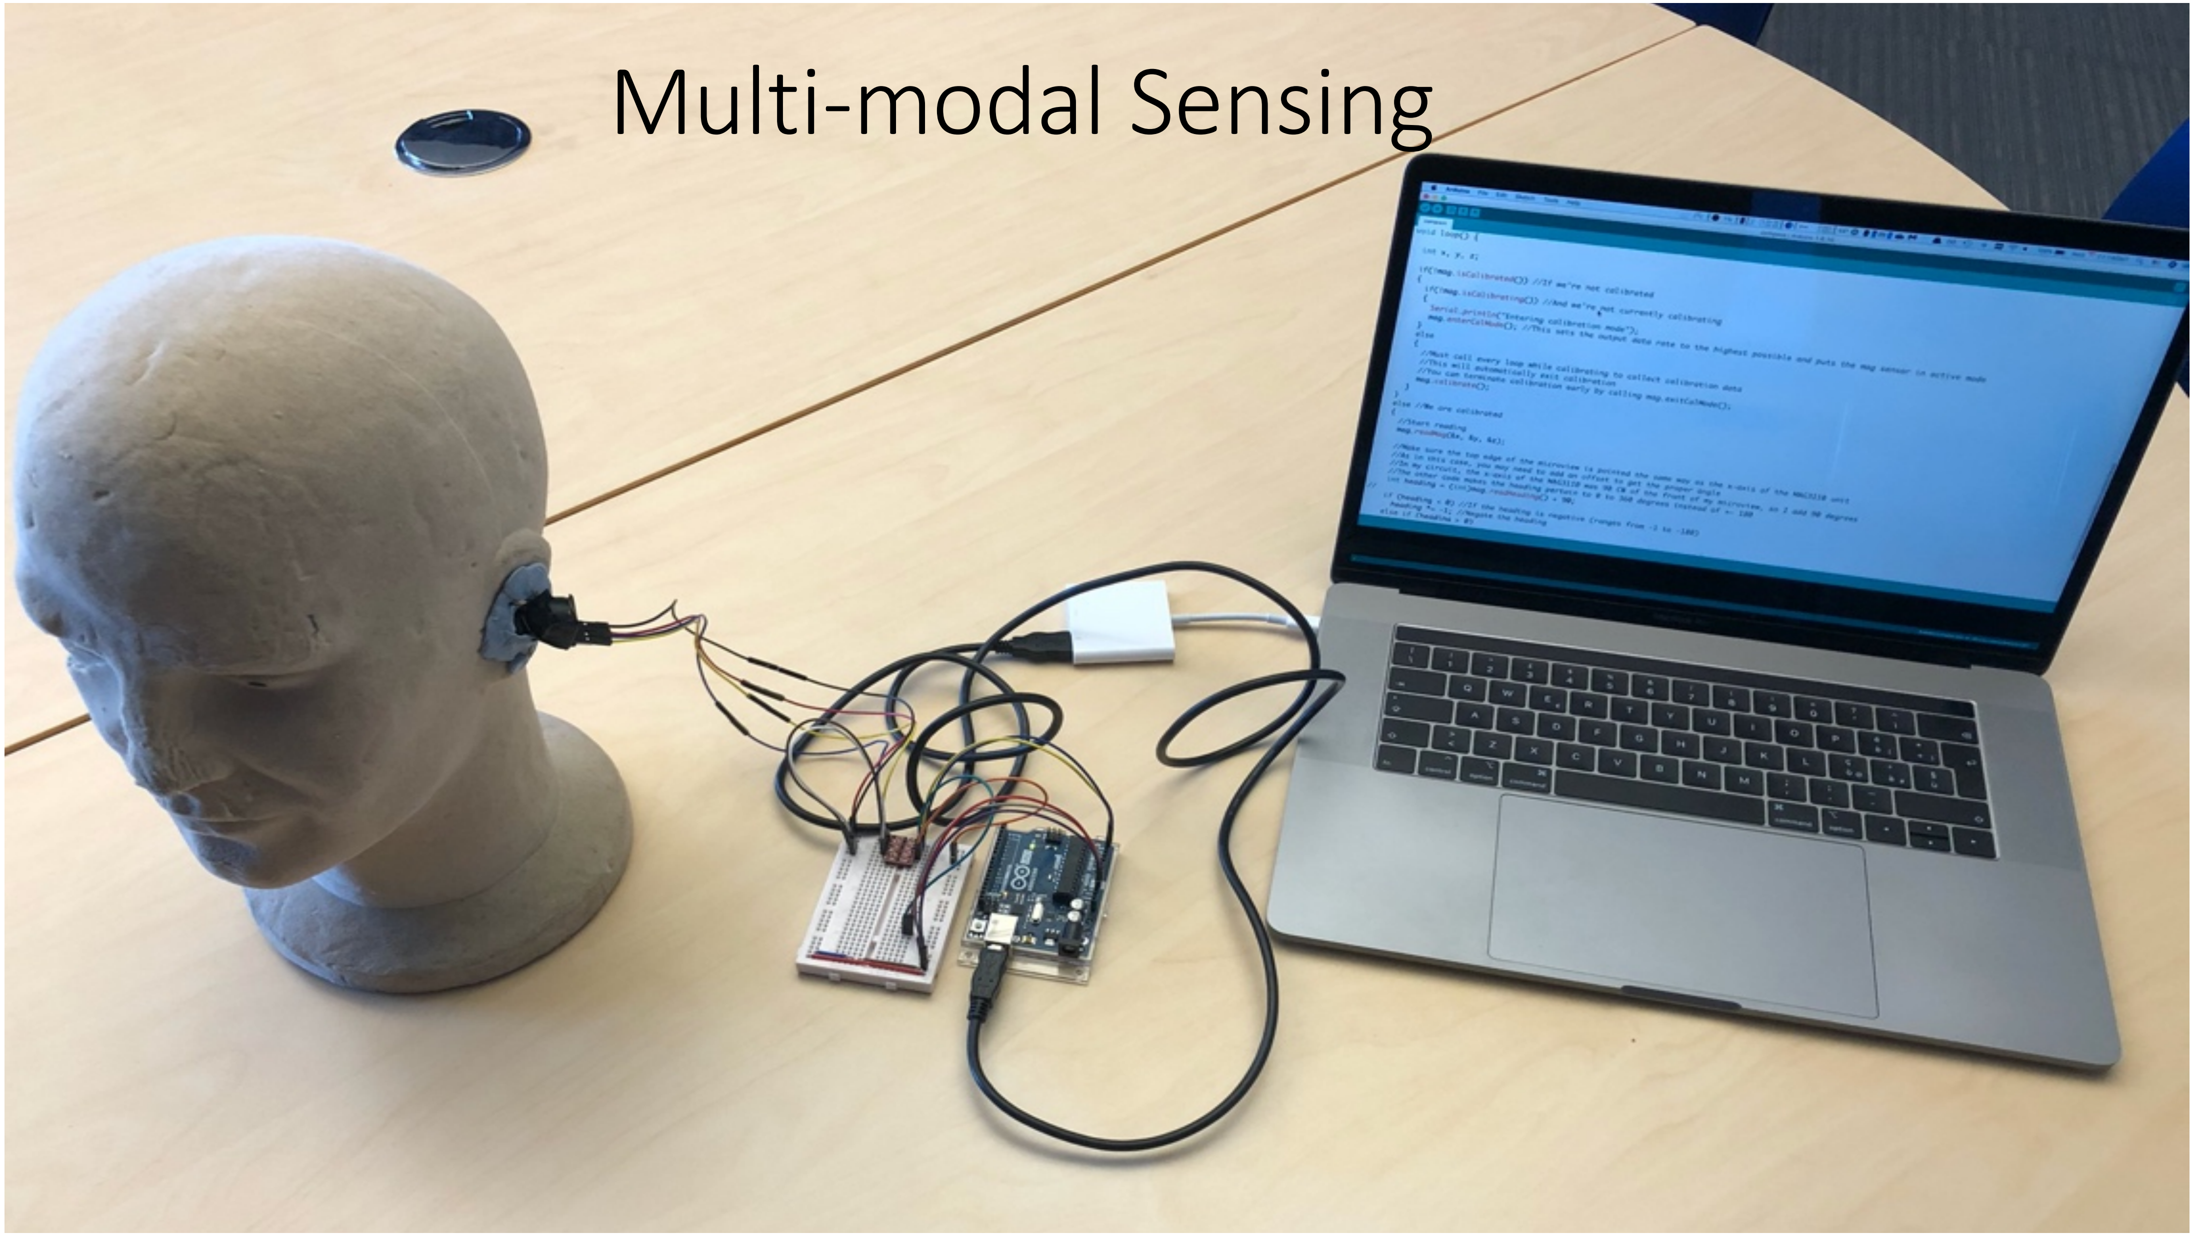

Accelerometer Data and Mood

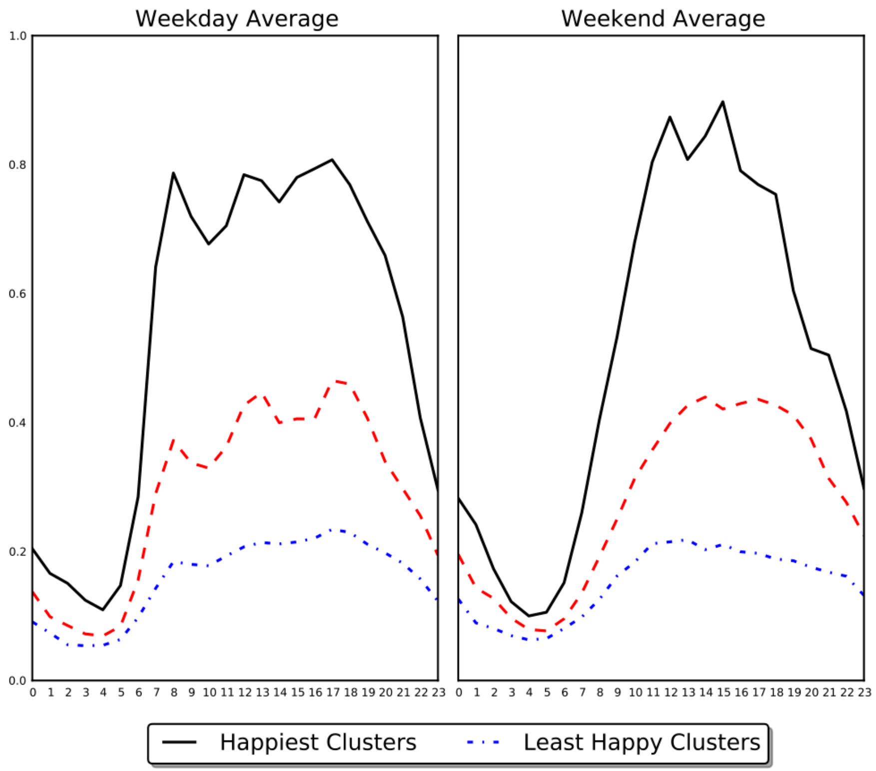

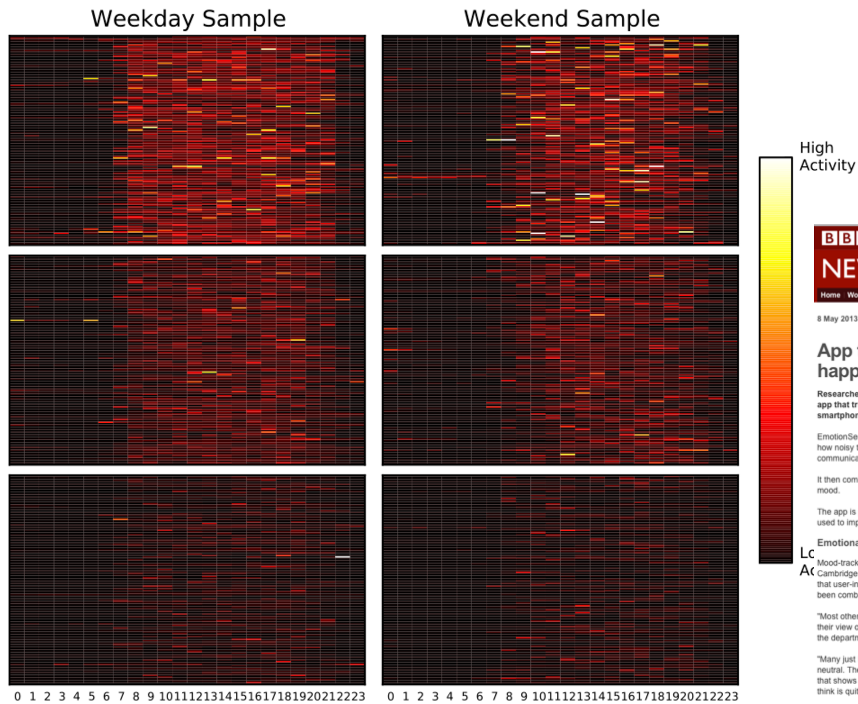

BBC NEWS TECHNOLOGY

8 May 2013 Last updated at 03:51

App taps phone and personal clues to your happiness

Researchers at Cambridge University have developed an app that tries to track happiness by combining smartphone data with users' perception of mood.

EmotionSense collects information about where users are, how noisy the environment is and whom they are communicating with.

It then combines this data with the user's own report about mood.

The app is part of a project to see how mobile phones can be used to improve health and wellbeing.

The app aims to combine phone data with perceived emotions

Emotional state

Mood-tracking apps already exist but the team from the Cambridge Computer Laboratory think this is the first time that user-input data and phone information sources have been combined.

"Most other attempts at software like this are coarse-grained in terms of their view of what a feeling is," said Dr Jason Rentfrow, a senior lecturer in the department of psychology at Cambridge University.

"Many just look at emotions in terms of feeling happy, sad, angry or neutral. The aim here is to use a more flexible approach, to collect data that shows how moods vary between people. That is something which we think is quite unique to the system we have designed," he said.

When the app is opened for the first time, a sensor that tells the researchers what time of day it is is unlocked. The app spends roughly a

Related Stories

Plan to crowdsource a happy city

What are the top five happiest parts of the UK?

Is it a good idea to measure stress?

Mobile sensing at the service of mental well-being: a large-scale longitudinal study. S Servia, K. Rachuri, C. Mascolo, P. Rentfrow, N. Lathia, G. Sandstrom.
In Proceedings of 26th International World Wide Web Conference (WWW 2017).

Behaviour Intervention

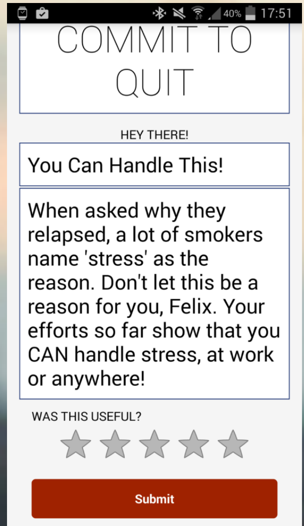

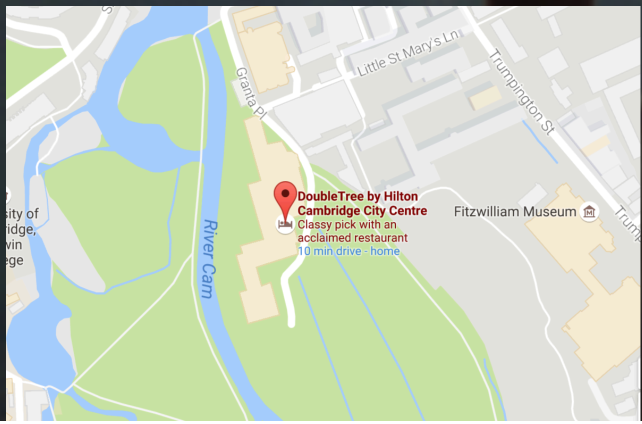

The feasibility of a context sensing smoking cessation smartphone application (Q Sense): a mixed methods study. Felix Naughton, Sarah Hopewell, Neal Lathia, Rik Schalbroeck, Chloe Brown, Cecilia Mascolo, Stephen Sutton. JMIR mHealth uHealth. September 2016

Early Alzheimer’s Disease Diagnostics

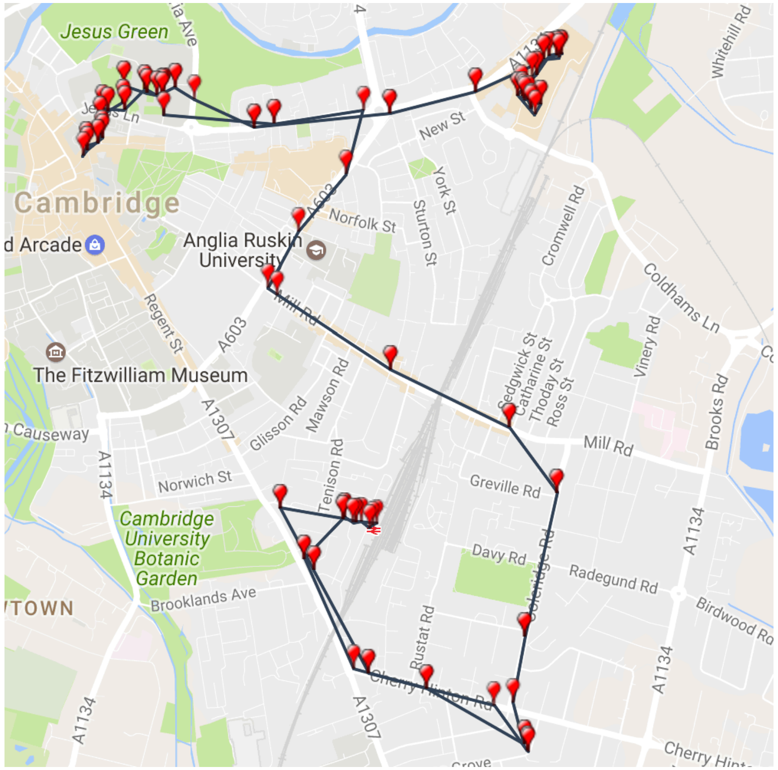

VO2max = Fitness: finding a better proxy

Mobile Data Analysis for Urban Planning

Fo

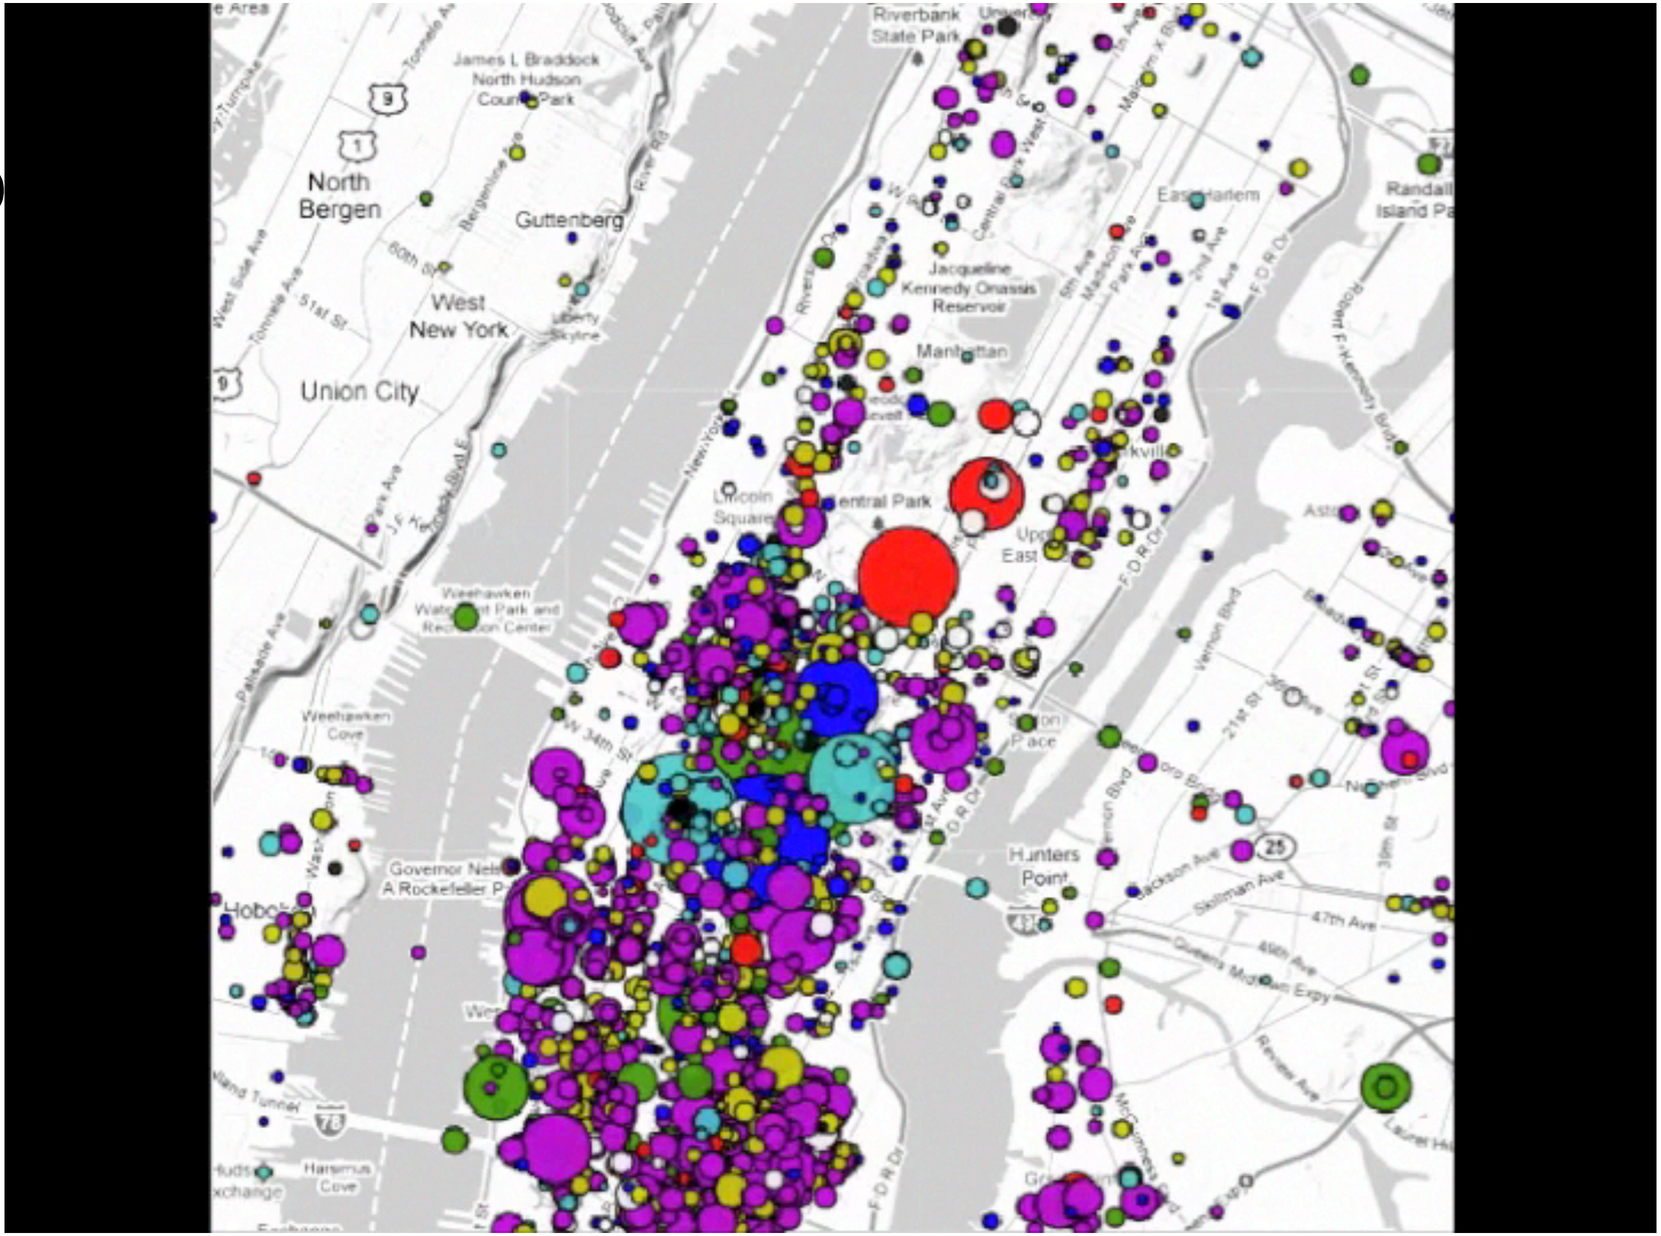

Riverbank State Park
James L Braddock North Hudson County Park
North Bergen
Guttenberg
West New York
Union City
Jacqueline Kennedy Onassis Reservoir
Manhattan
East Harlem
Randall Island Pa
Central Park
Lincoln Square
Weehawken Waterfront Park and Recreation Center
Weehawken Cove
Governor Nelson A Rockefeller P
Hunters Point
Harsimus Cove
47th Ave
Queens Midtown Expy

Predicting Popularity of Venues

Predicting the temporal activity patterns of new venues.
Krittika D'Silva, Anastasios Noulas, Mirco Musolesi Cecilia Mascolo, and Max Sklar.
In EPJ Data Science. 2018.

Predicting Venue Closure

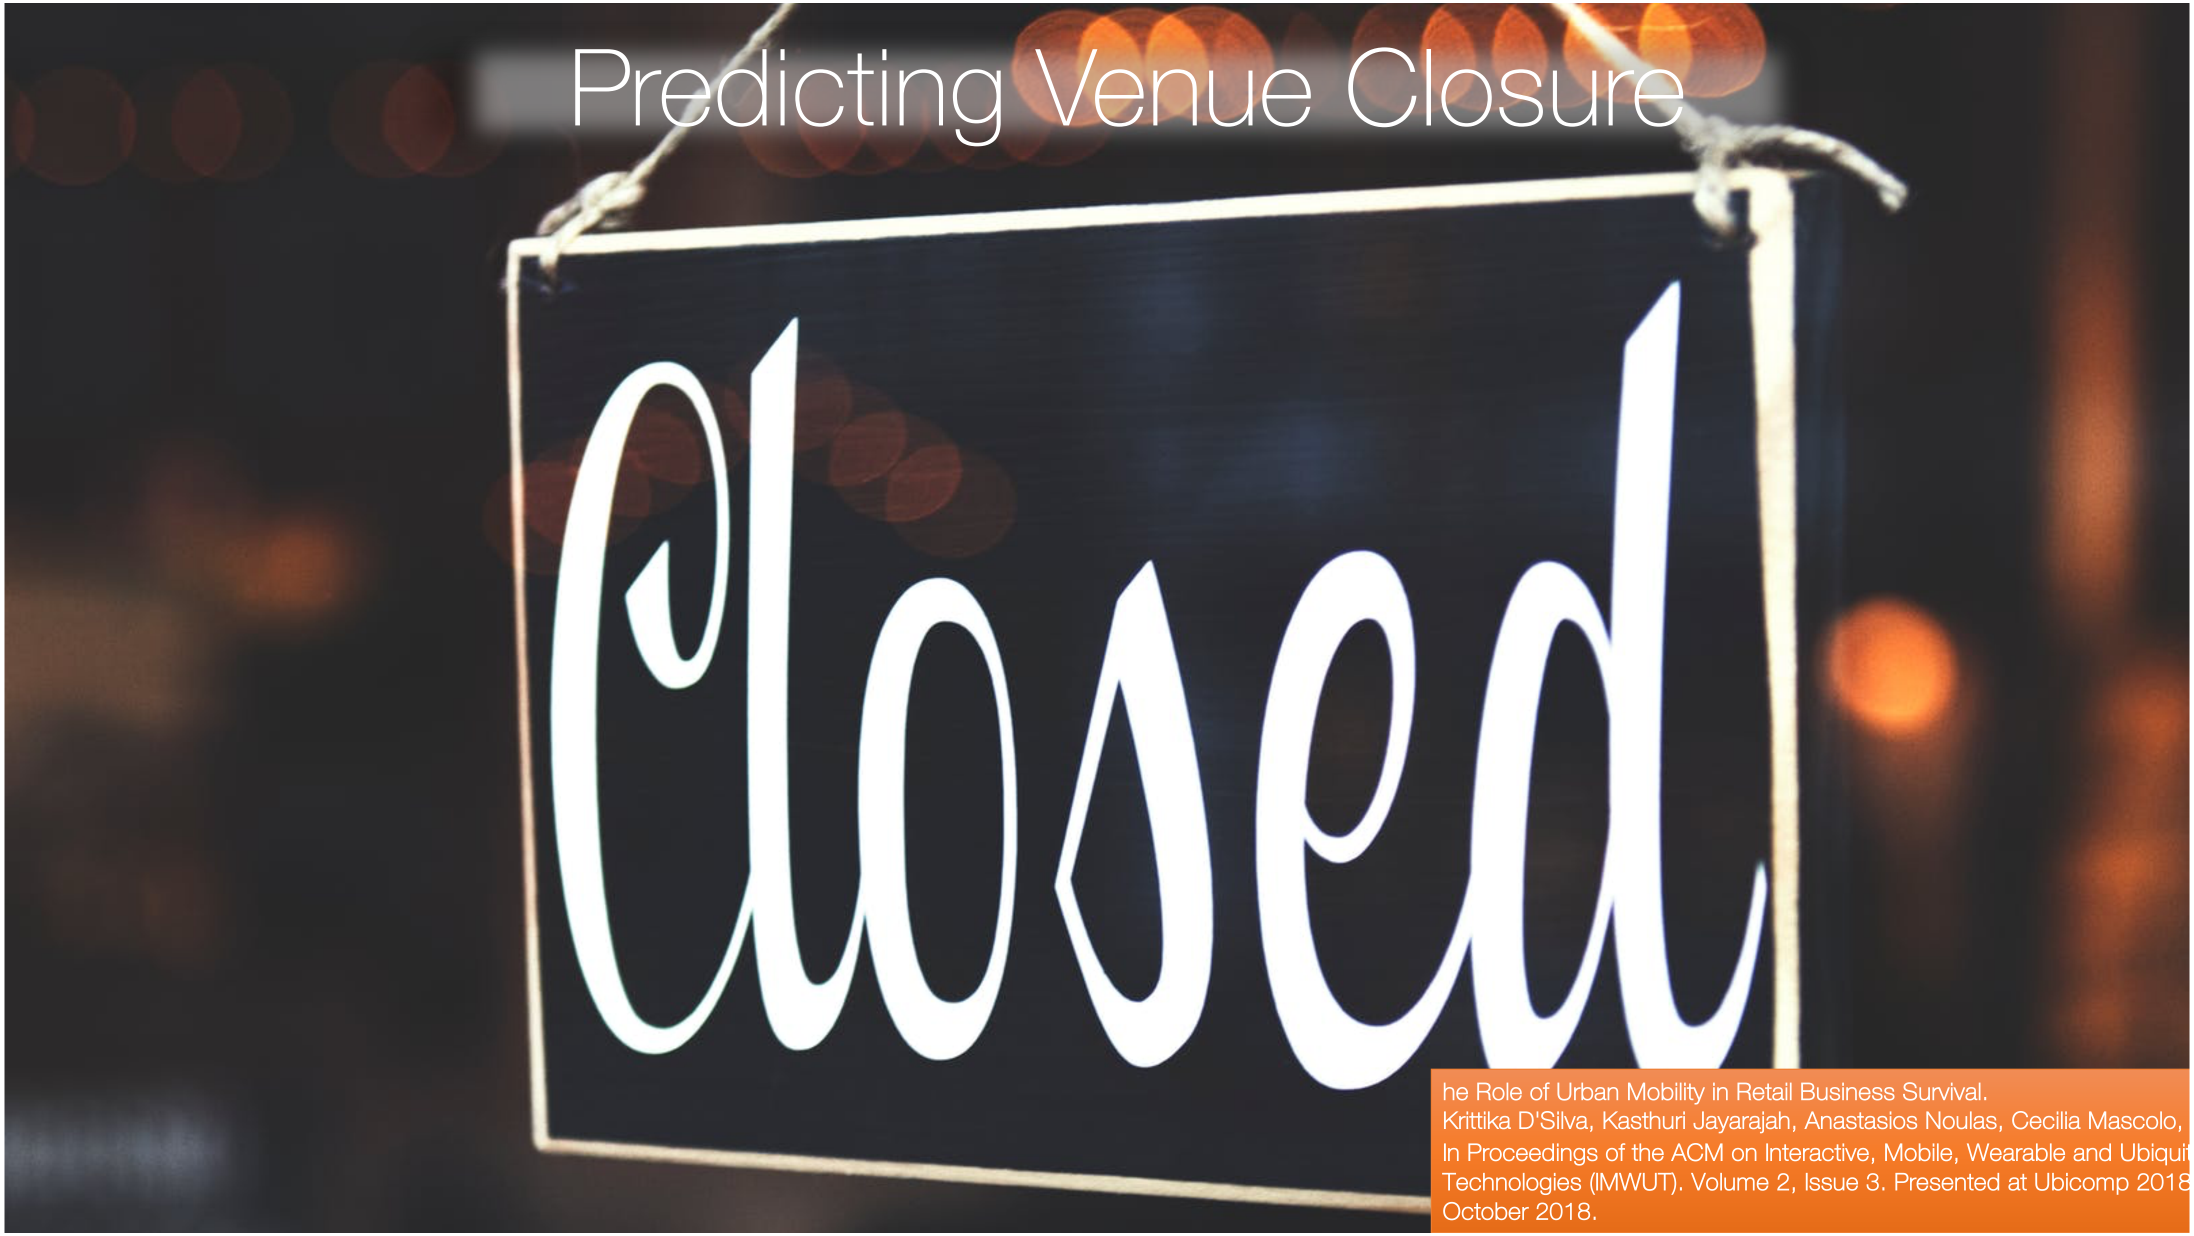

he Role of Urban Mobility in Retail Business Survival.
Krittika D'Silva, Kasthuri Jayarajah, Anastasios Noulas, Cecilia Mascolo,
In Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquit
Technologies (IMWUT). Volume 2, Issue 3. Presented at Ubicomp 2018
October 2018.

Audio Based Health Diagnostics

Emotionsense
Capturing Emotions from Microphone in the Wild

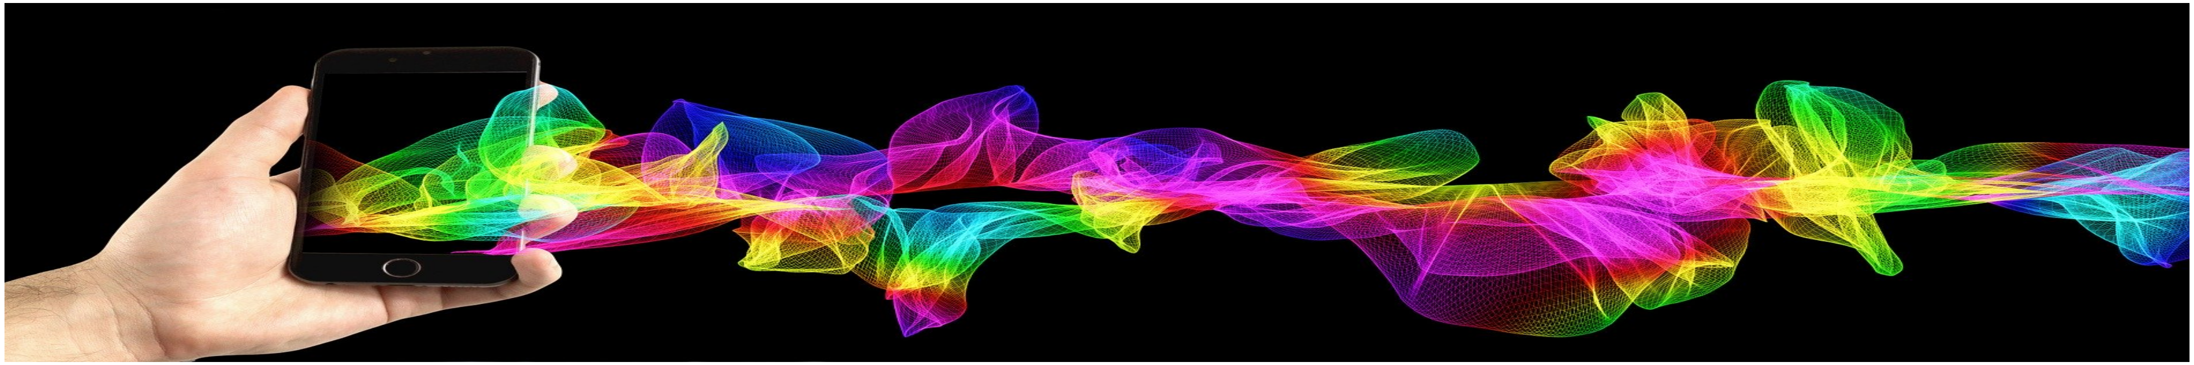

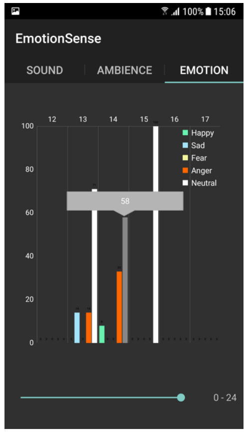

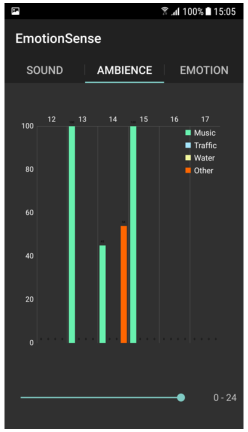

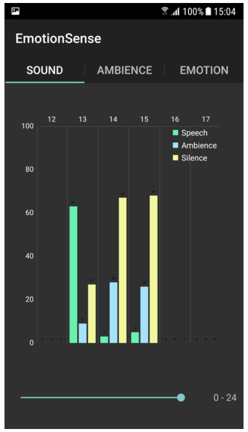

EmotionSense: A Mobile Phones based Adaptive Platform for Experimental Social Psychology Research

Kiran K. Rachuri
Computer Laboratory
University of Cambridge
kkr27@cam.ac.uk

Mirco Musolesi
School of Computer Science
University of St. Andrews
mirco@cs.st-andrews.ac.uk

Cecilia Mascolo
Computer Laboratory
University of Cambridge
cecilia.mascolo@cl.cam.ac.uk

Peter J. Rentfrow
Faculty of Politics, Psychology, Sociology and International Studies
University of Cambridge
pjr39@cam.ac.uk

Chris Longworth
Department of Engineering
University of Cambridge
cl336@cam.ac.uk

Andrius Aucinas
Computer Laboratory
University of Cambridge
aa535@cam.ac.uk

ABSTRACT

Today's mobile phones represent a rich and powerful computing platform, given their sensing, processing and communication capabilities. Phones are also part of the everyday life of billions of people, and therefore represent an exceptionally suitable tool for conducting social and psychological experiments in an unobtrusive way.

Author Keywords

Emotion Recognition, Speaker Recognition, Social Psychology, Mobile Phones, Energy Efficiency.

INTRODUCTION

Mobile phones represent an ideal computing platform to monitor behavior and movement, since they are part of the everyday life of billions of people [1]. Recently, systems such as

Acoustic based Sensing for Diagnostics

Use of commodity devices or cheap built devices as sound recorders

Bodily sounds

Additional sensor inputs

Challenges: on device inference, power, robustness to noise, machine learning models, interpretation

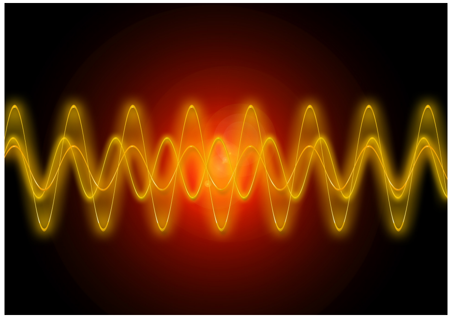

Fine-grained Continuous Diagnosis

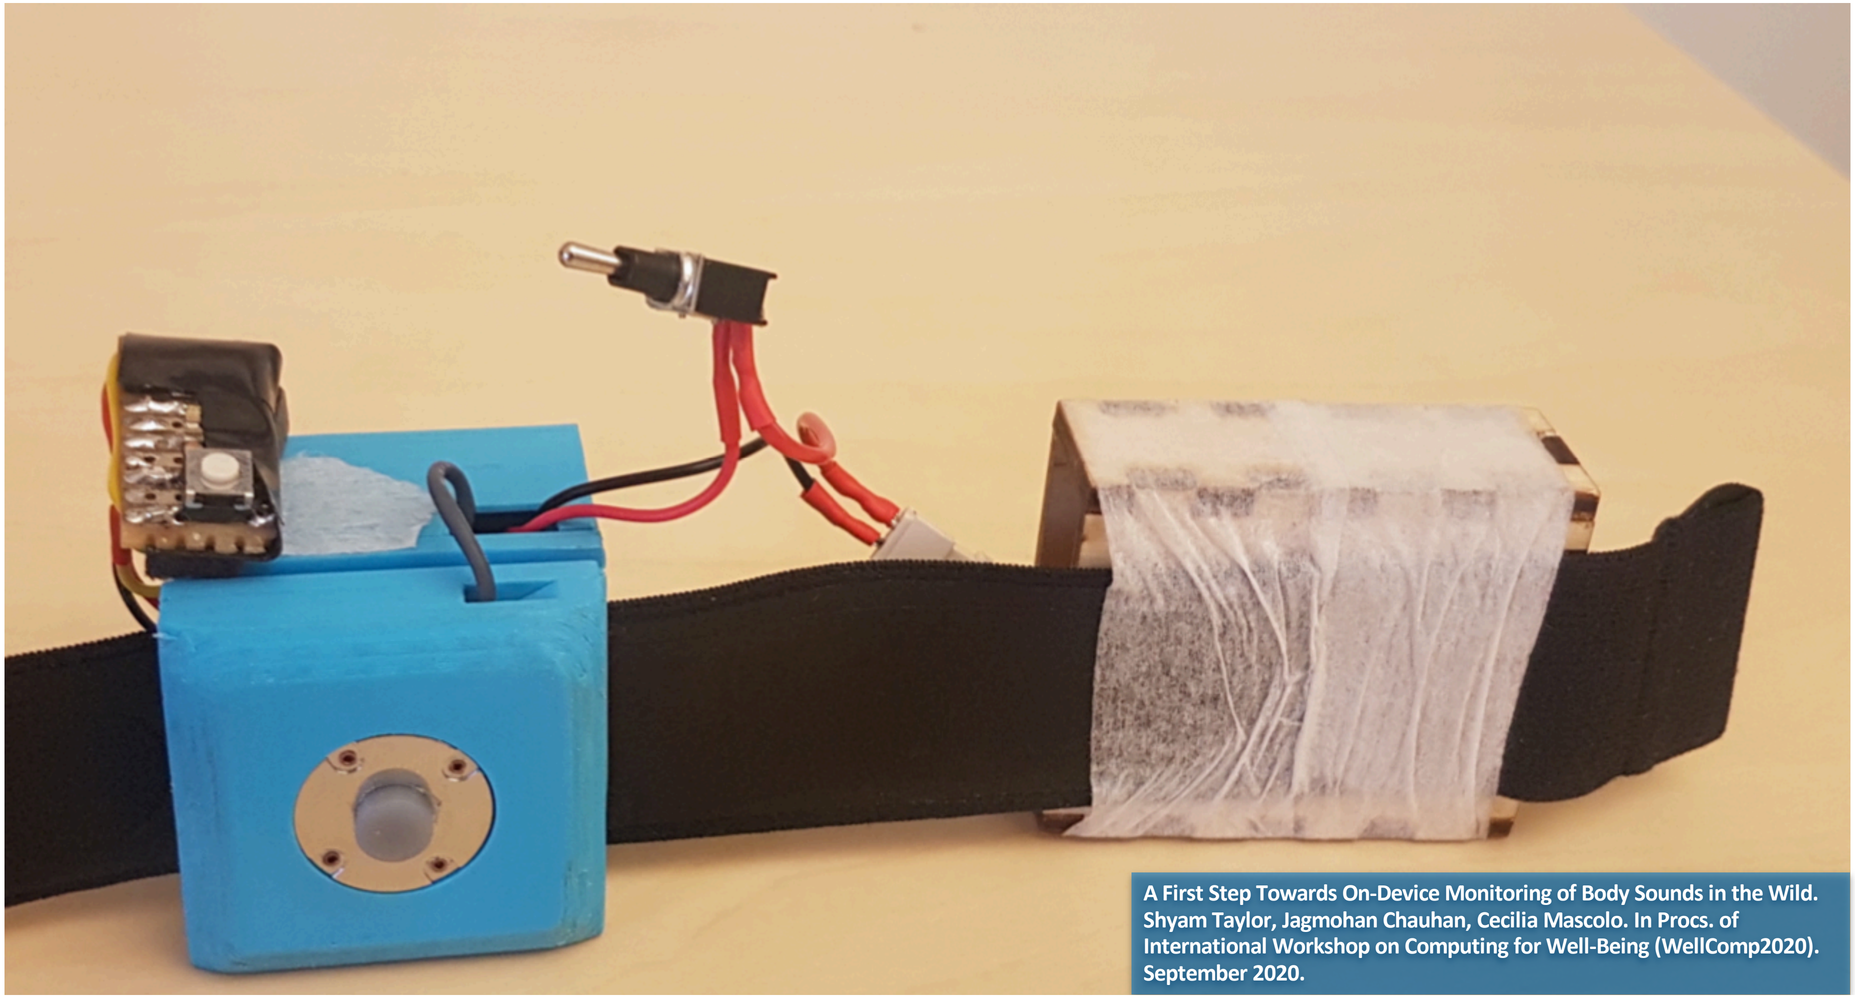

A First Step Towards On-Device Monitoring of Body Sounds in the Wild. Shyam Taylor, Jagmohan Chauhan, Cecilia Mascolo. In Procs. of International Workshop on Computing for Well-Being (WellComp2020). September 2020.

covid-19-sounds.org

UNIVERSITY OF CAMBRIDGE

COVID-19 Sounds App

Upload short recordings of cough and breathing and report symptoms to help researchers from the University of Cambridge detect if a person is suffering from COVID-19. Healthy and *non-healthy* participants welcome.

GET IT ON Google Play

Download on the App Store

or use the **online form**

Cecilia Mascolo

Cecilia is Professor of Mobile Systems. She is an expert in mobile health and mobile data analysis.

Pietro Cicuta

Pietro is Professor of Biological Physics at the Cavendish Laboratory, Cambridge.

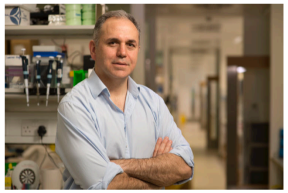

Andres Floto

Andres is Professor of Respiratory Biology and Research Director of the Cambridge Centre for Lung Infection at Papworth Hospital.

On-device and Incremental Machine Learning

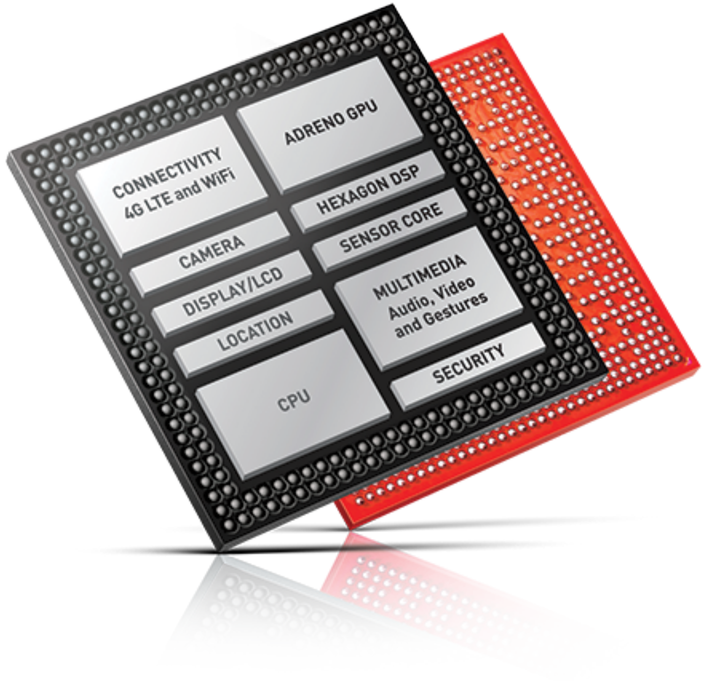
ADRENO GPU
CONNECTIVITY
4G LTE and WiFi
HEXAGON DSP
SENSOR CORE
CAMERA
DISPLAY/LCD
MULTIMEDIA
Audio, Video
and Gestures
LOCATION
SECURITY
CPU

Optimized GPU is Efficient

Optimized GPU with batching outperforms cloud energy-wise

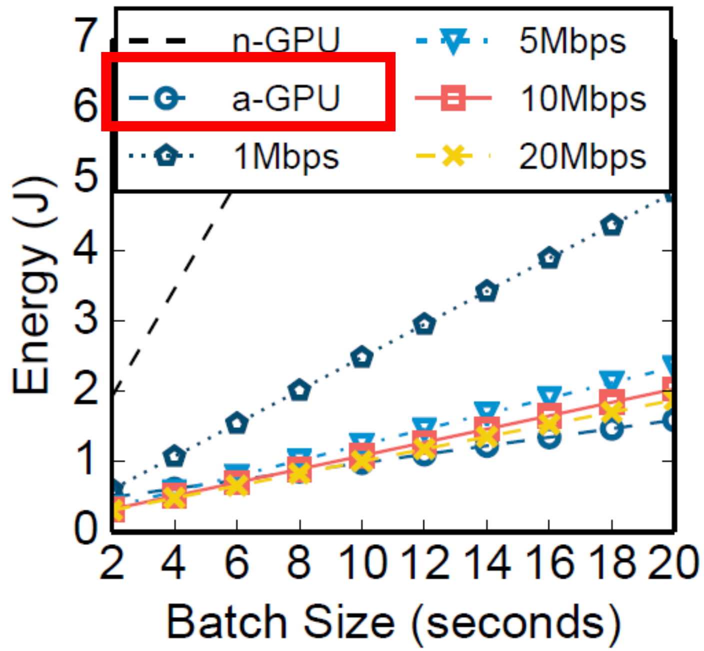

Learning Continuously

Improving models through continual learning

Resource trade offs… what about microcontrollers?

Where is this useful?

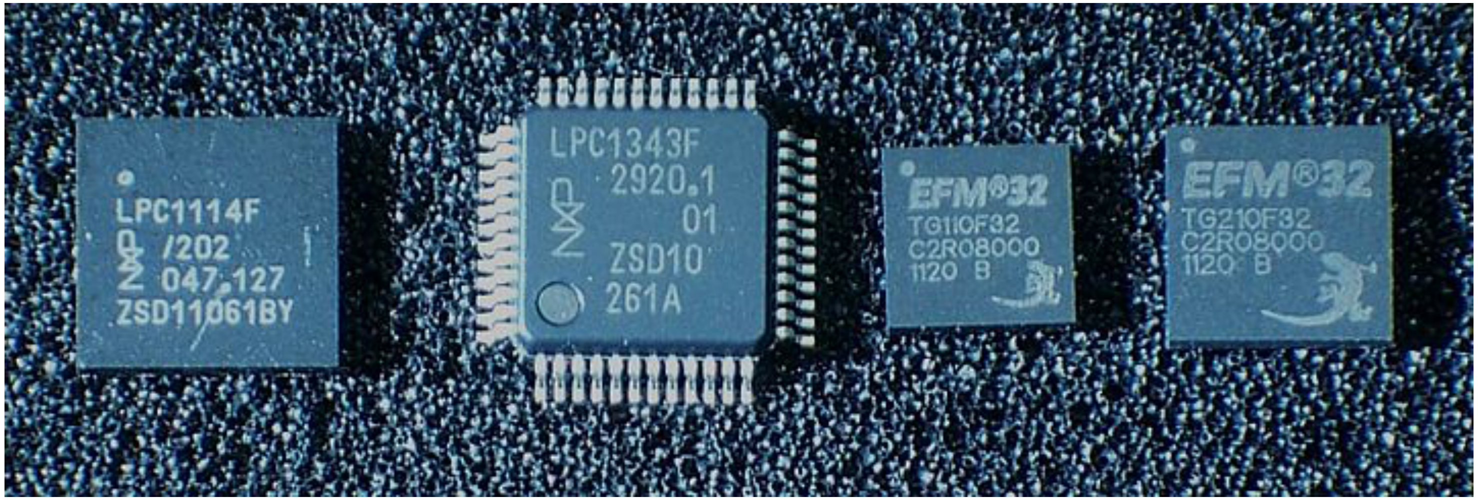

Google

Issues

Big issue with current AI models:

Can’t learn on the fly. Suffers from Catastrophic forgetting (CF).

Mostly designed for efficient inference.

Mobile Sensing Context

- Human Activity Recognition.

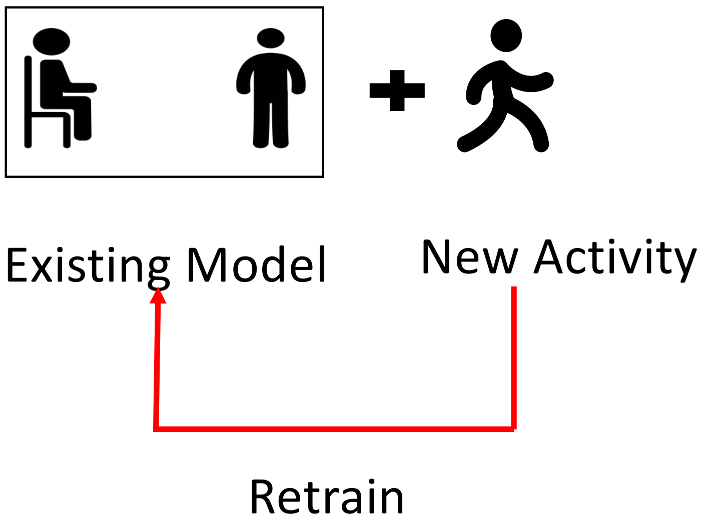

- Lacks Adaptability.
- Retrain – Waste of resources.

Usable Security Context

- User authentication

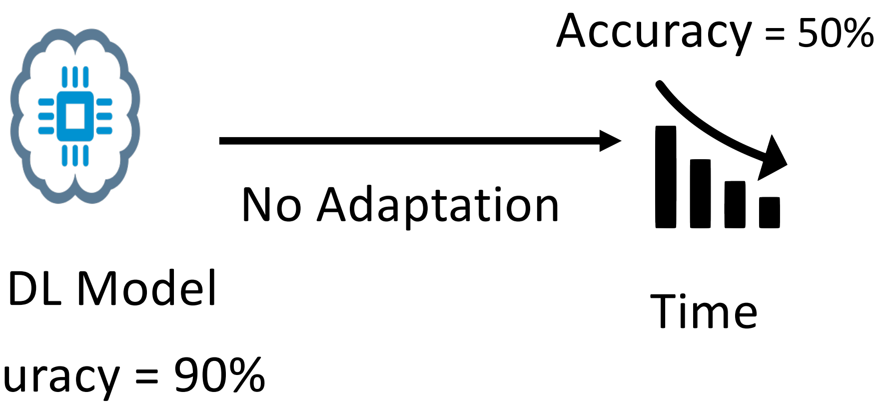

- Affects usability

Models needs to be adaptive and resilient to changes.

Advocate Learning on device: Protect privacy, avoid network costs.

Continual Learning (Sensing)

Accommodate new classes on the fly: Difficult as catastrophic forgetting (CF) happens.

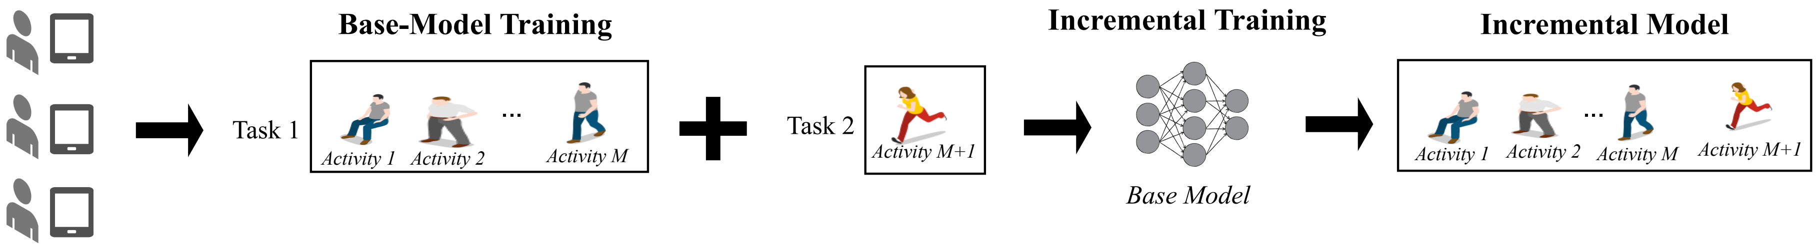

Apply Incremental Learning to learn new classes.

Continual Learning (Sensing)

- Explore the performance of SOTA IL algorithms focusing on accuracy, storage and latency on device – Jetson TX2, a smartphone.
- Various mobile sensing tasks: Human Activity, Emotion, Gesture Recognition.
- Tasks: Add classes (activity, gesture or emotion) on the fly.
- Major findings:

✓ *CF can be largely mitigated with some of the IL algorithms*.

✓ End to end IL = DL *Training (slow)* + IL (fast) is feasible on smartphone CPU w/o burning the device! Important as training is extremely compute intensive.

Continual Learning (User Authentication)

- ContAuth: Solve degrading accuracy. Combination of online and IL models.
- Online: Adapt DNN model to changing user behaviours and IL: add new users for user authentication.

Continual Learning (User Authentication)

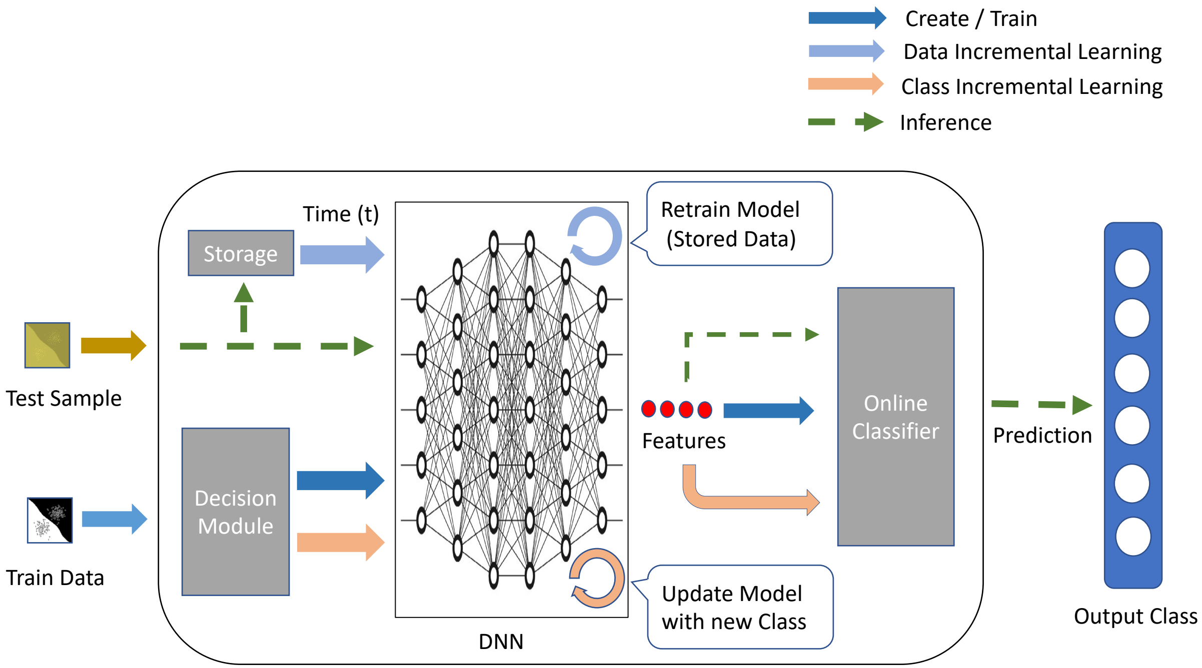

Continual Learning (User Authentication)

- Analyze various modalities: breathing, gait.

- Major findings:

✓ ContAuth can help accuracy to stay > 85% over time.

✓ 35% improvement over non adaptable DNN model.

✓ Can be done completely on device – Jetson TX2.

On Device Learning

- Creating a generic framework to enable training on the device: phones and wearables.
- Efficient learning on Micro Controller units.
- Devise software-based optimizations to do efficient training, continual and multi-task learning.


URBAN
PLANNING


Health Care
Doctor
Hospital
Pharmacist
Nurse
Dentist
First Aid
Surgeon
Emergency

Impact

Cecilia Mascolo cm542

Jagmohan Chauhan jc2161

Dept of Computer Science and Technology,
University of Cambridge

UNIVERSITY OF
CAMBRIDGE