R269
Mobile, Wearable Systems and Machine Learning

Prof Cecilia Mascolo, Dr Yang Liu

UNIVERSITY OF CAMBRIDGE

Prof Cecilia Mascolo

- Mobile and Wearable Systems
- On Device Machine Learning
- Mobile and Wearable Health

Dr Yang Liu

- Mobile and Wearable Sensing
- Wearable and Earable Health
- Mobile Interaction Systems

The course

- The course is about anything to do with mobile systems
    - Systems aspects including power, computation
    - Novel Sensing aspects
    - Wearable/Sensor Data learning aspects
    - Modelling and Inference On Device
    - Applications (Health, Sustainability…)

The Schedule

- 5^{th} October(1h) Introduction (TODAY!)
- 12^{nd} October System, Energy and Security
- 19^{th} October Backscatter Comms, Battery Free and Energy Harvesting Devices
- 26^{th} October New Sensing Modalities
- 2^{th} November On Device Machine Learning
- 9^{th} November Machine Learning on Wearable Data
- 16^{th} November Mobile and Wearable Health
- 23^{rd} November (3h) Mobile and Wearable Systems of Sustainability

Assessment

- A total of 7 items of assessment:
  - 1-2 Presentations
  - 5-6 Reports
- Each contributing 1/7
- A class list of attendance will be kept and apologies for absence should be sent to the lecturers prior to 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 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) in Moodle
- Students assigned randomly each week
- Presentations will be assessed for technical content, clarity, engagement, timeliness and question answering

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

Deadlines:

- Assignment 1 due Wednesday 11th October, noon
- Assignment 2 due Wednesday 18th October, noon
- Assignment 3 due Wednesday 25th October, noon
- Assignment 4 due Wednesday 1st November, noon
- Assignment 5 due Wednesday 8th November, noon
- Assignment 6 due Wednesday 15th November, noon
- Assignment 7 due Wednesday 22nd November, 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
- On Device Machine Learning
- Audio for Health Diagnostics
- Machine Learning for Wearable Data

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

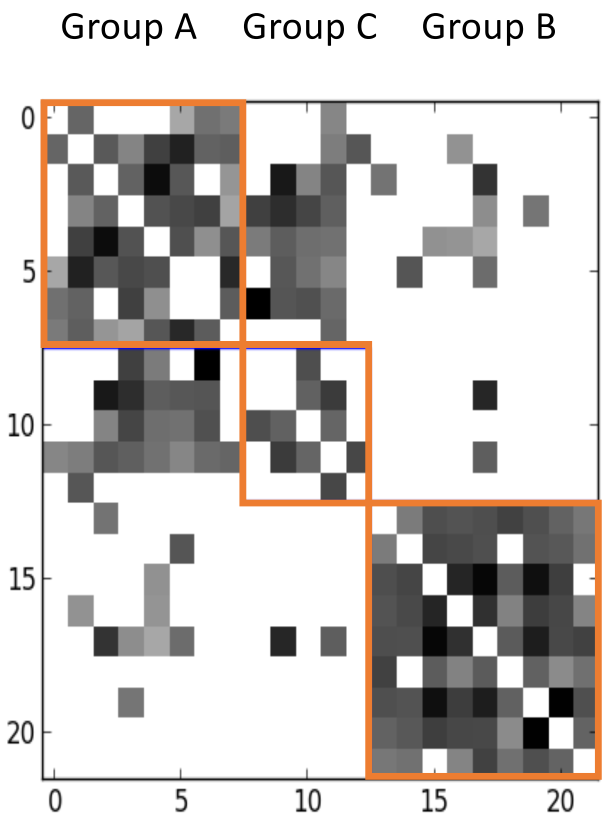

Old building

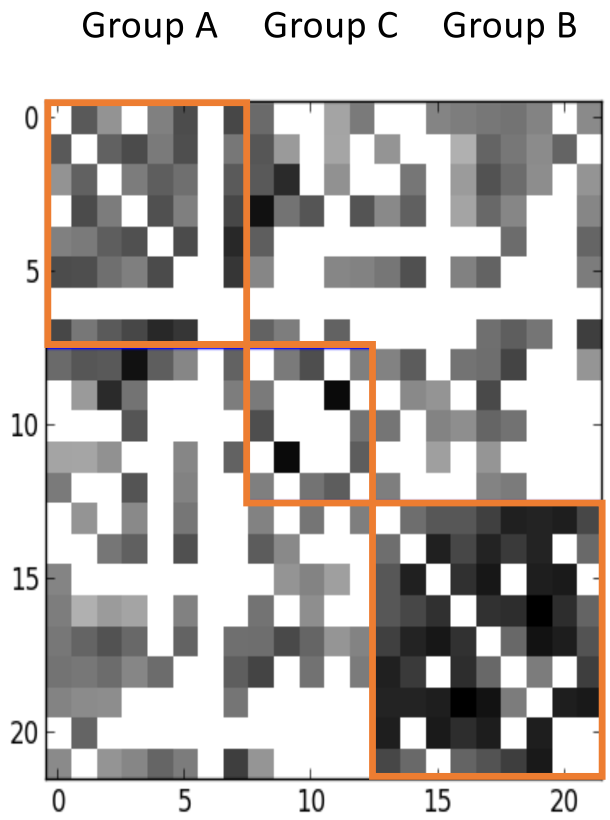

New building

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

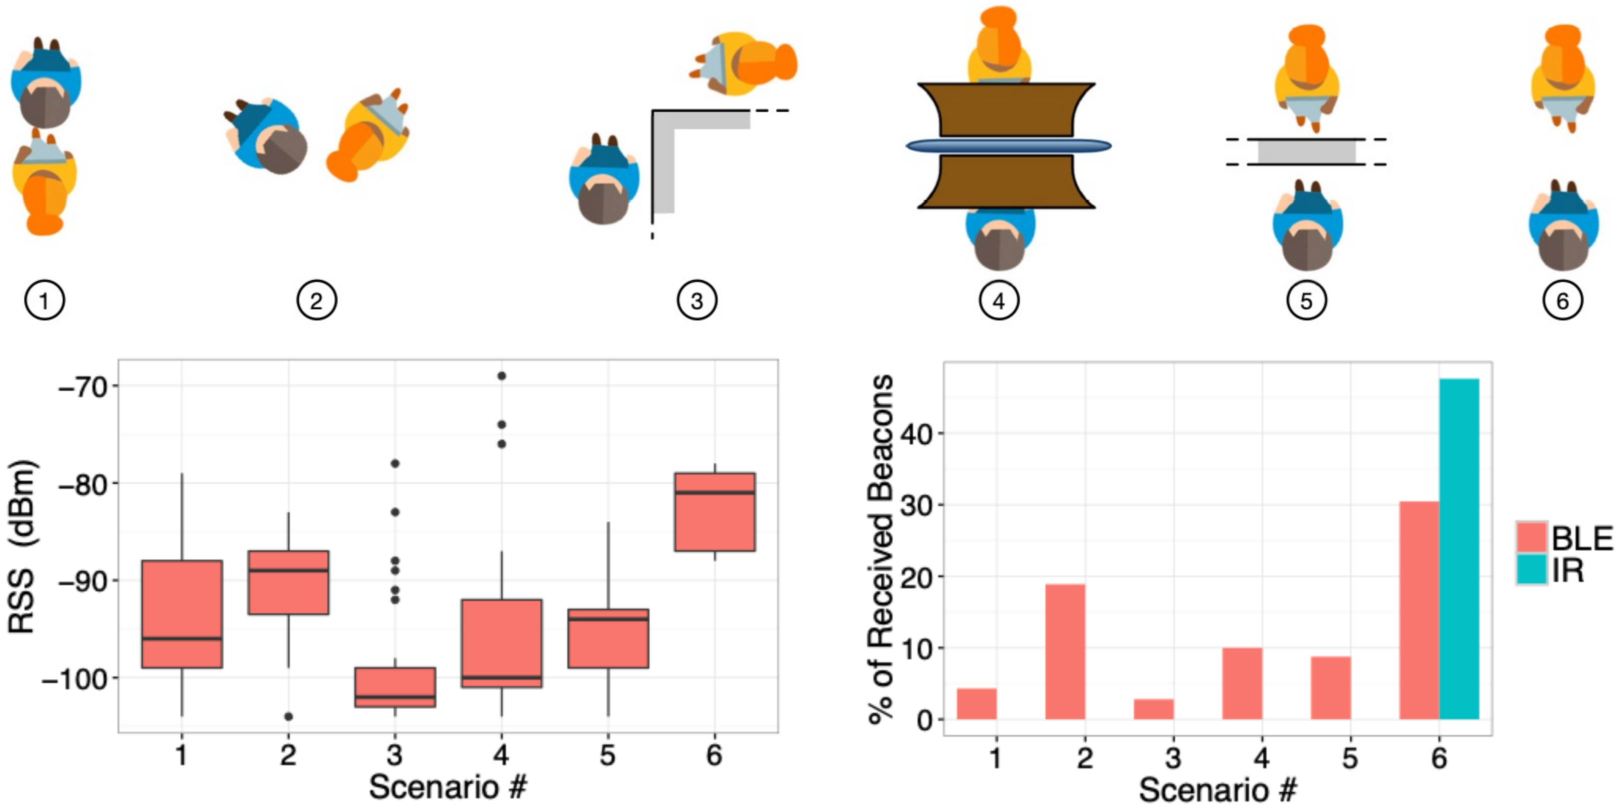

(a) (b)

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

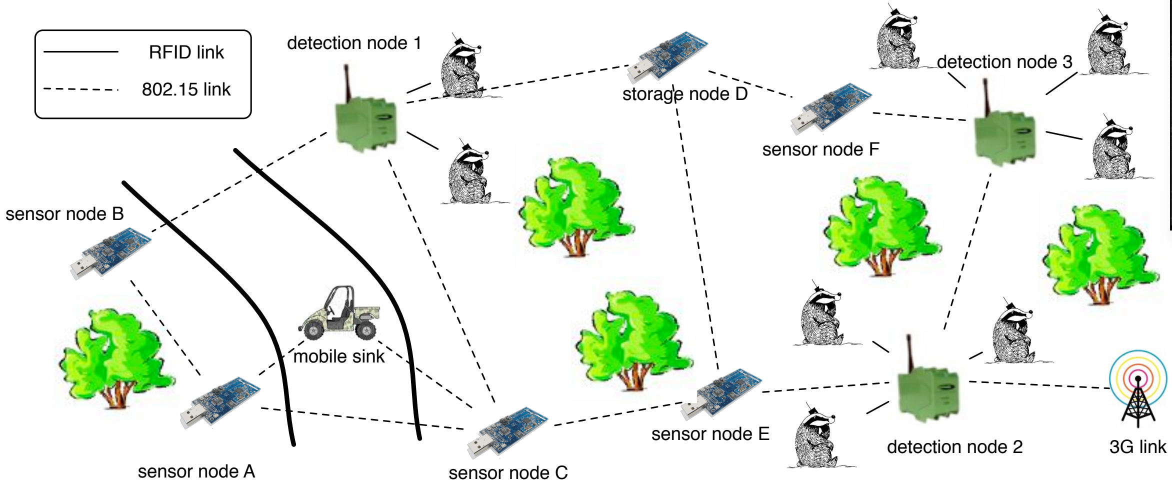

Understanding Animal Movement

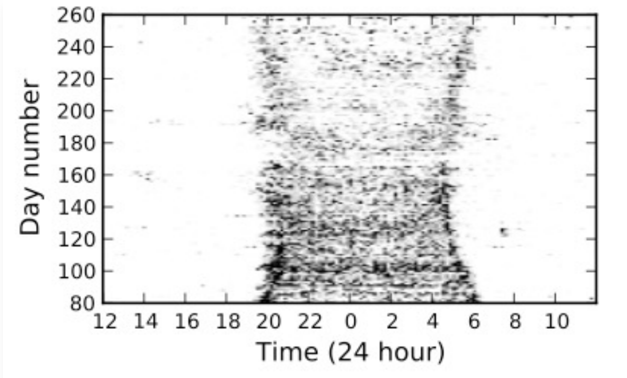

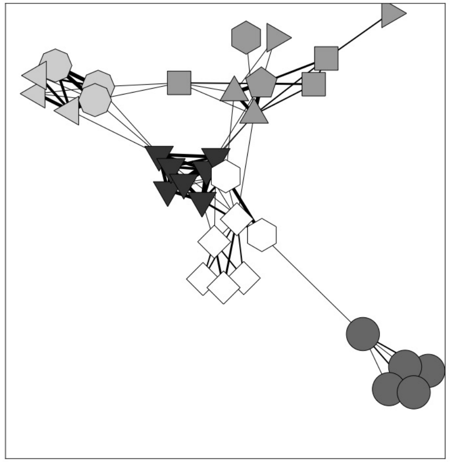

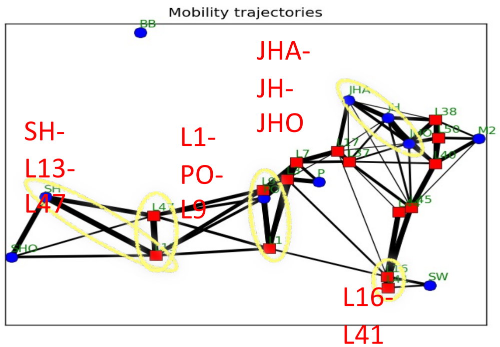

Accelerometer Data and Mood

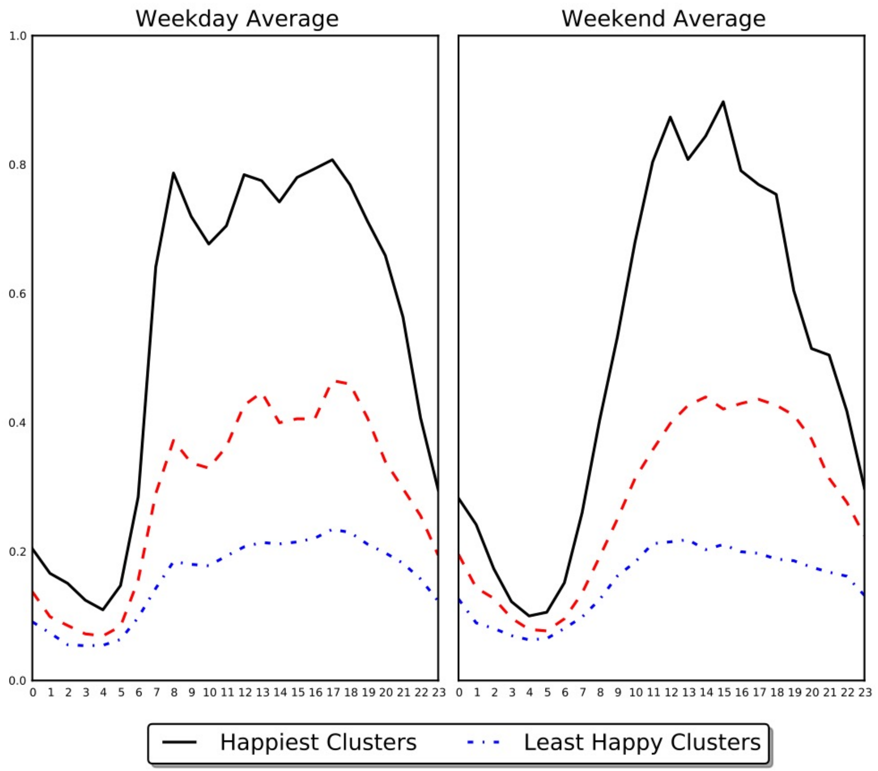

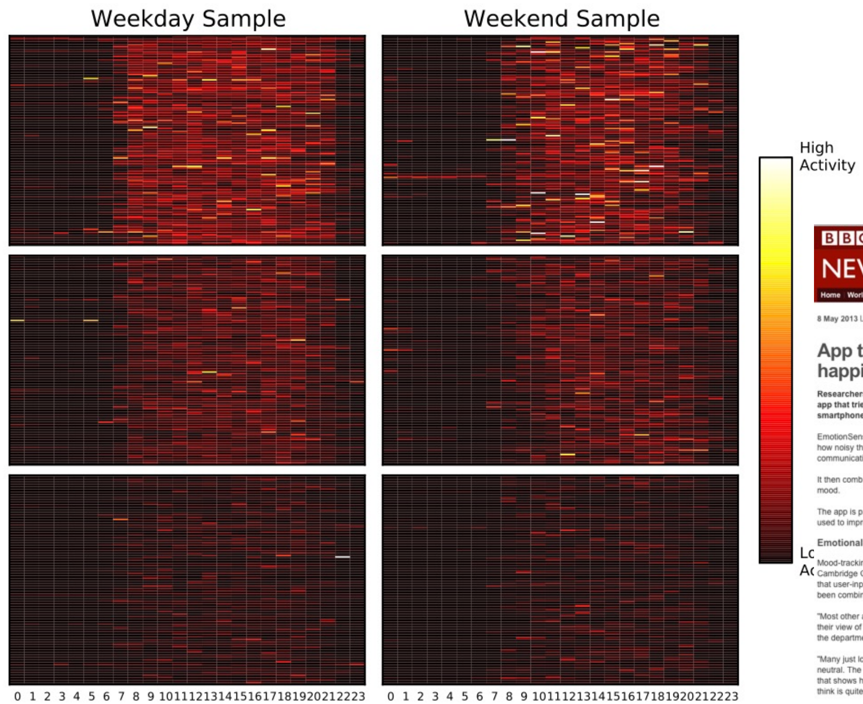

BBC Sign in News Sport Weather iPlayer TV

NEWS TECHNOLOGY

Home World UK England N. Ireland Scotland Wales Business Politics Health Education

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

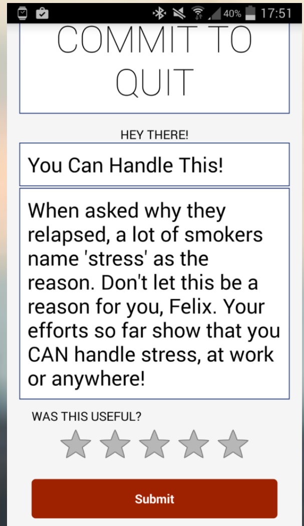

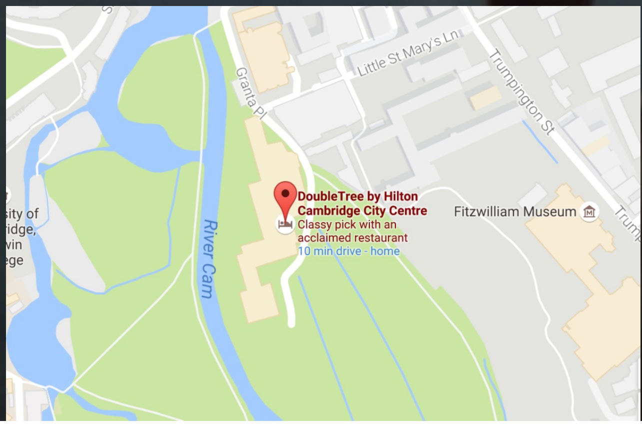

The feasibility of a context sensing smoking cessation smartphone application (Q Sense): a mixed methods study. Felix Naughton, Sarah Hopewell, Neal Lathia, Rik Schalbroeck, Chloe Brown, Cecilia Mascolo, Stephen Sutton. JMIR mHealth uHealth. September 2016

Early Alzheimer’s Disease Diagnostics

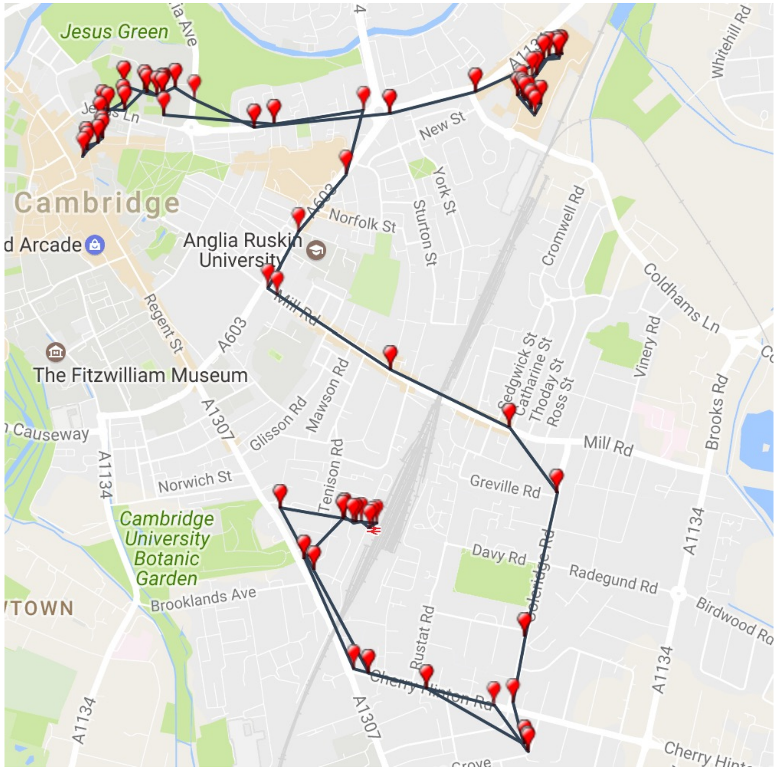

Differentiating features

- Patients travel the same segments more often
- Patients explore less – entropy of spatial distribution is low

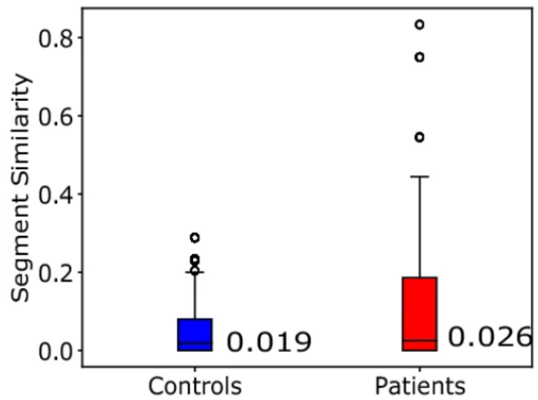

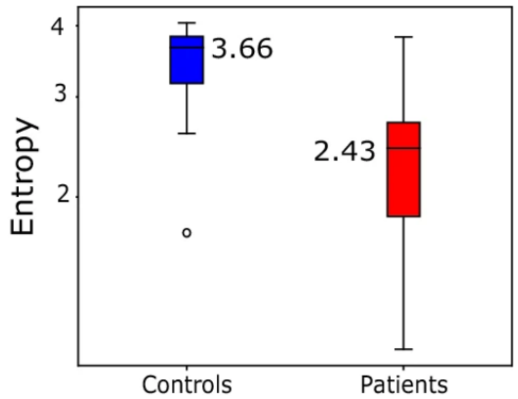

Classifier Sensitivity: 0.71, Specificity: 0.83

VO2max = Fitness: finding a better proxy

Audio Based Health Diagnostics

Emotionsense
Capturing Emotions from Microphone in the Wild

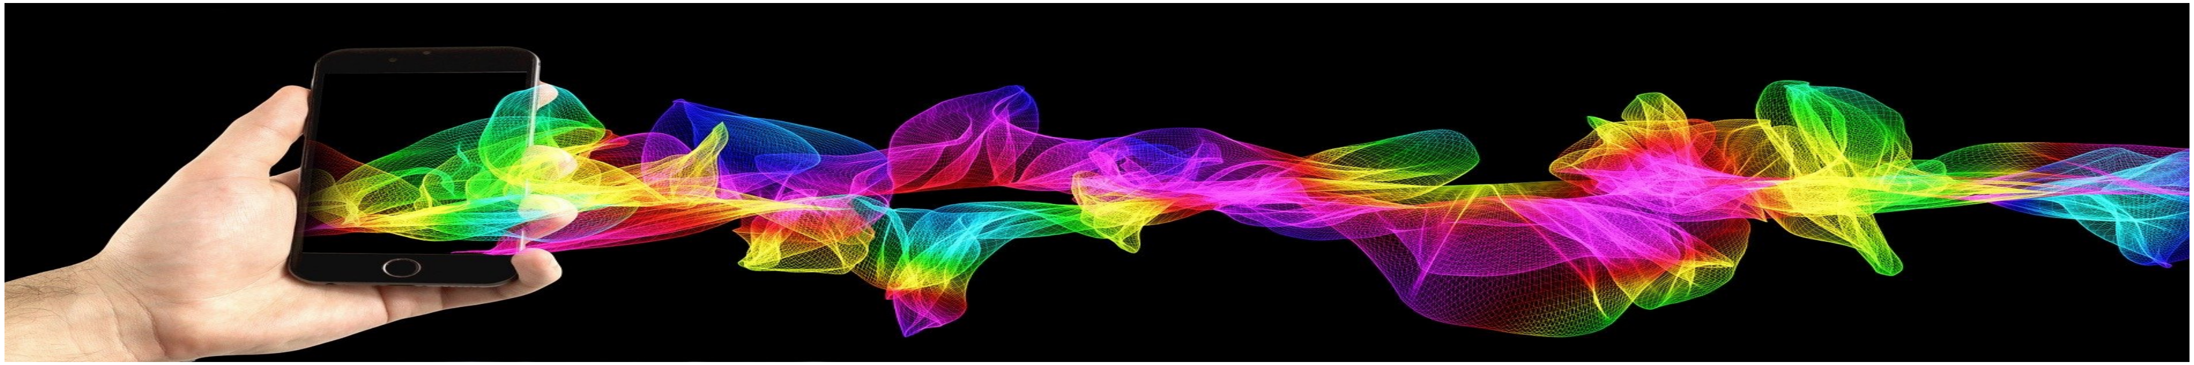

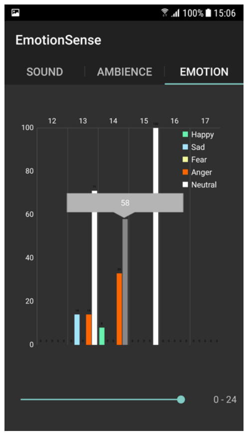

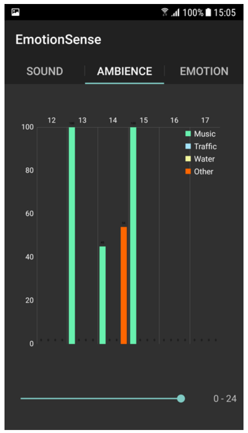

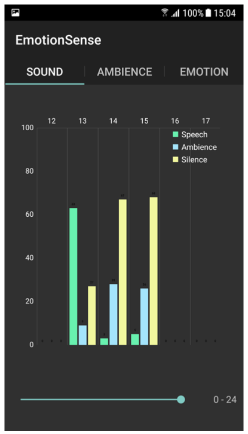

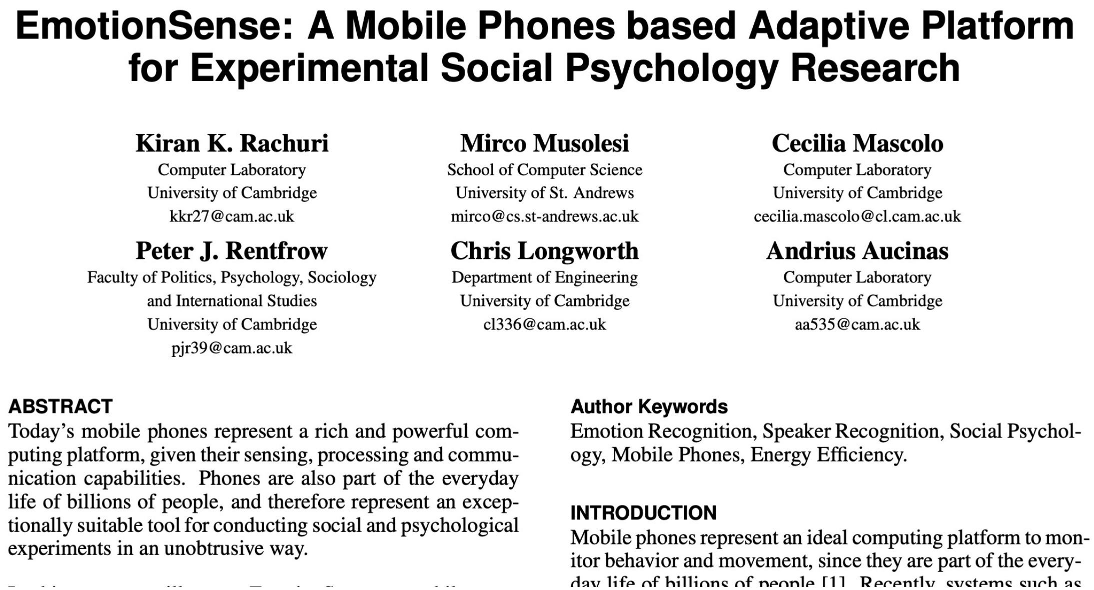

EmotionSense: A Mobile Phones based Adaptive Platform for Experimental Social Psychology Research

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

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.

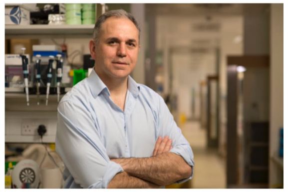

Andres Floto

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

Can AI diagnose patients with severe respiratory tract infections?

Researchers using AI to identify patients at risk of developing severe respiratory tract infections have just won support from UK Research and Innovation.

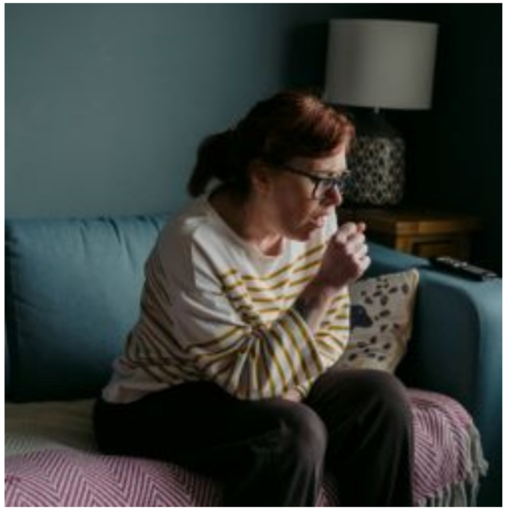

The RELOAD study - REspiratory disease progression through LOngitudinal Audio Data machine learning - is being funded as part of a new government mission to support AI innovation to speed up health research.

The project will use AI to analyse the sounds of patients' breathing and speech and diagnose those who are at risk of becoming severely ill.

It is one of 22 new projects that Michelle Donelan, Secretary of State for Science and Technology, revealed last week would receive funding to explore how to develop and use AI in health.

> *Identifying patients at higher risk could reduce hospital admissions, cases of severe illness and the number who die.*
>
> *— Cecilia Mascolo*

The Principal Investigator on the RELOAD project is **Cecilia Mascolo**, Professor of Mobile Systems here. The collaborative project also includes **Professor Pietro Cicuta** in the Cambridge University Department of Physics; **Professor Nick Francis**, head of the Medical School at Southampton University; and **Professor Anna Barney**, Professor of Biomedical Acoustic Engineering at Southampton University.

On-Device In-Ear Microphone Sensing Systems

D. Ma, A. Ferlini, C. Mascolo. OESense: Employing Occlusion Effect for In-ear Human Motion Sensing. ACM MobiSys. June 2021.

Occlusion Effect

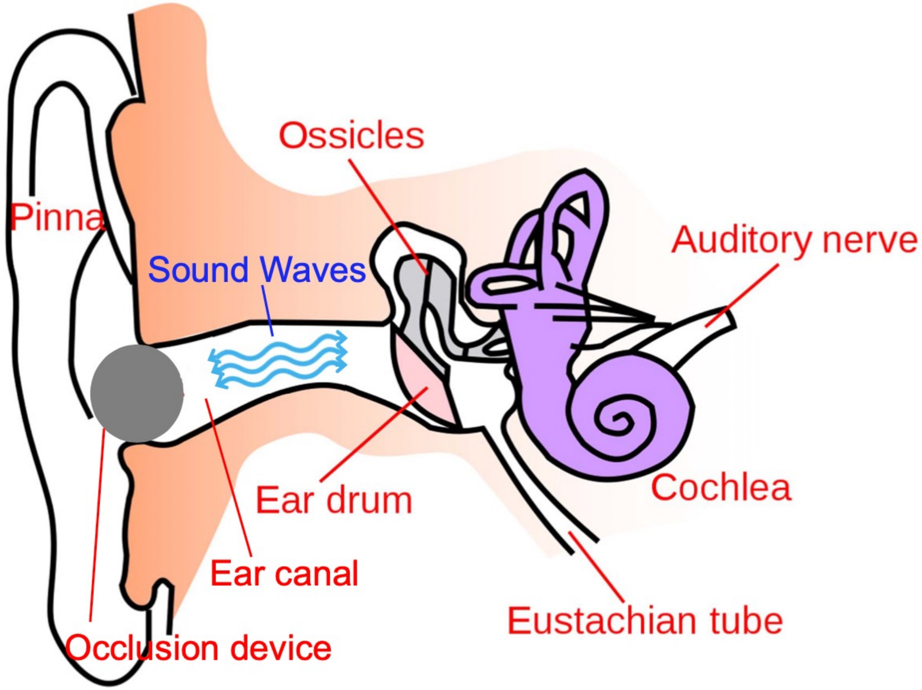

Left Earbud

Right Earbud

In-air Mic

Walk

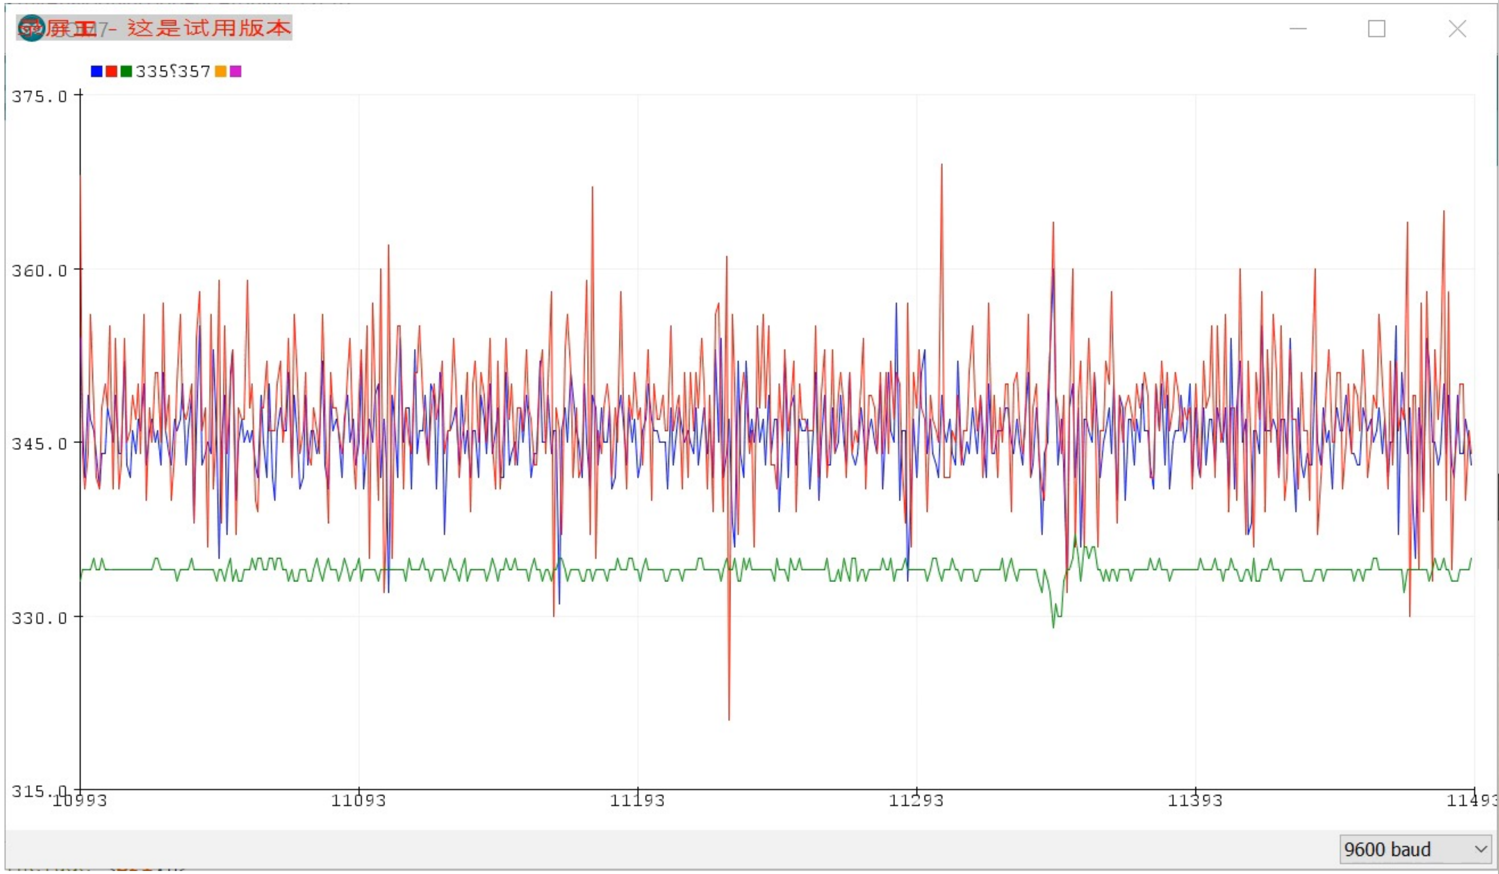

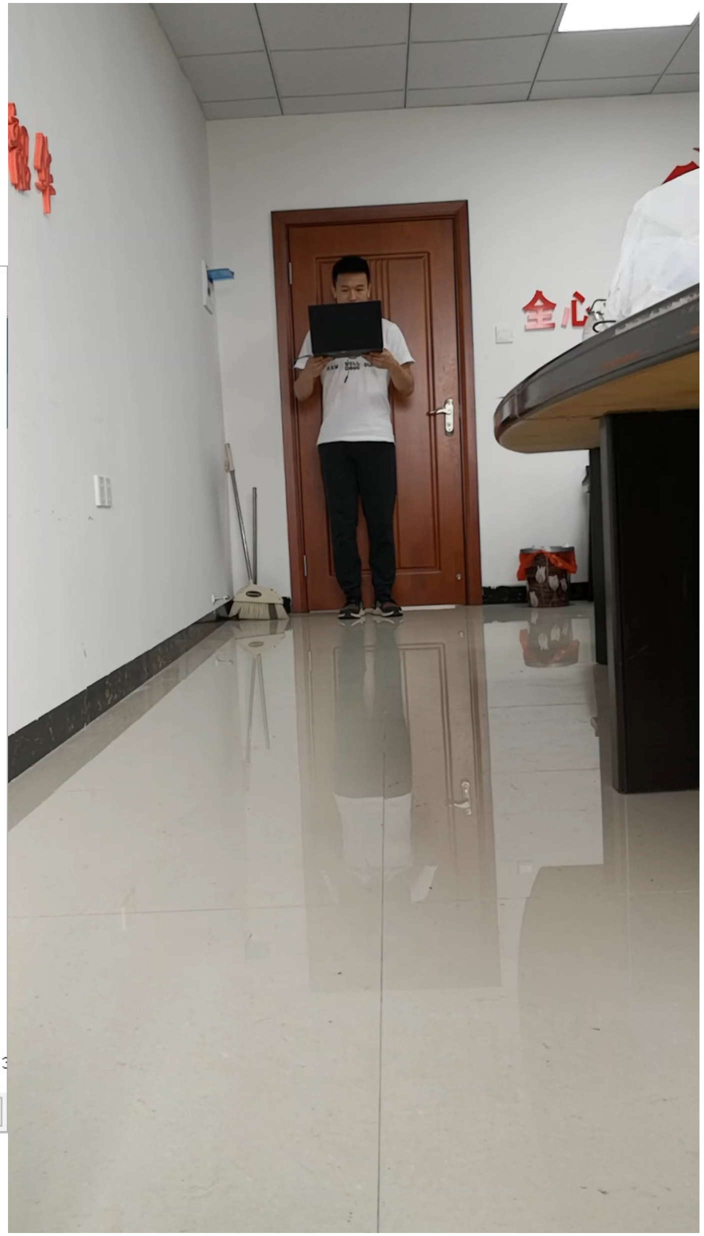

Gait Recognition from In-Ear Microphone

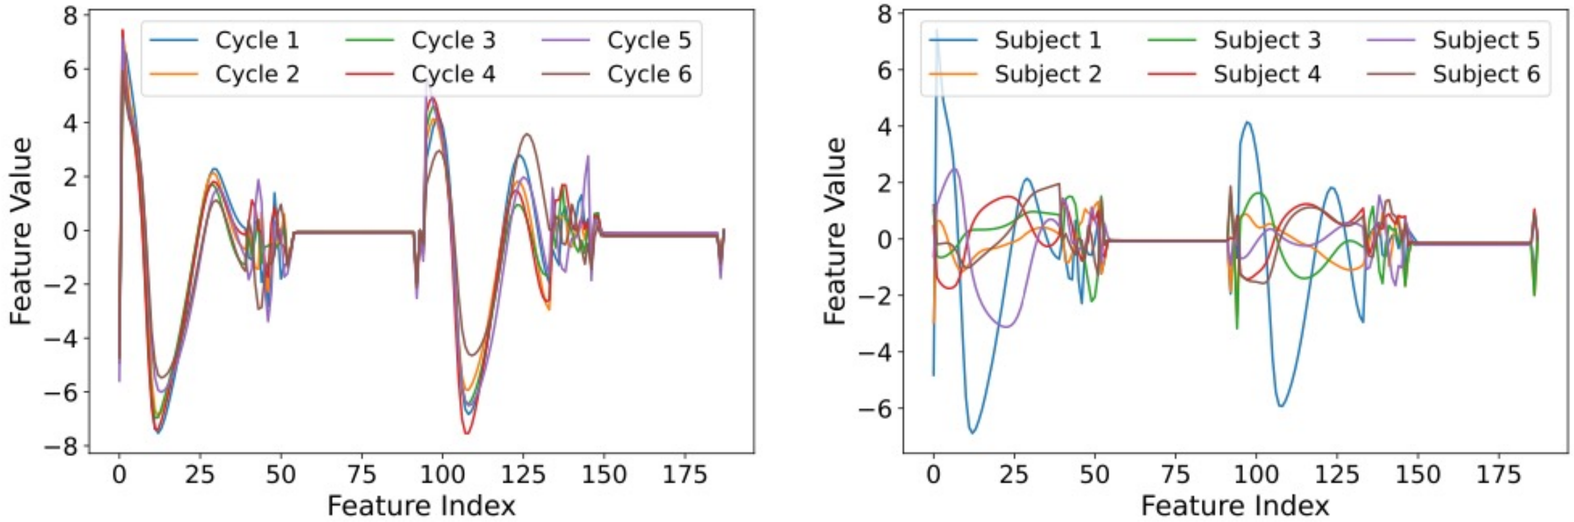

(a) (b)

Heartrate from In-Ear Microphone

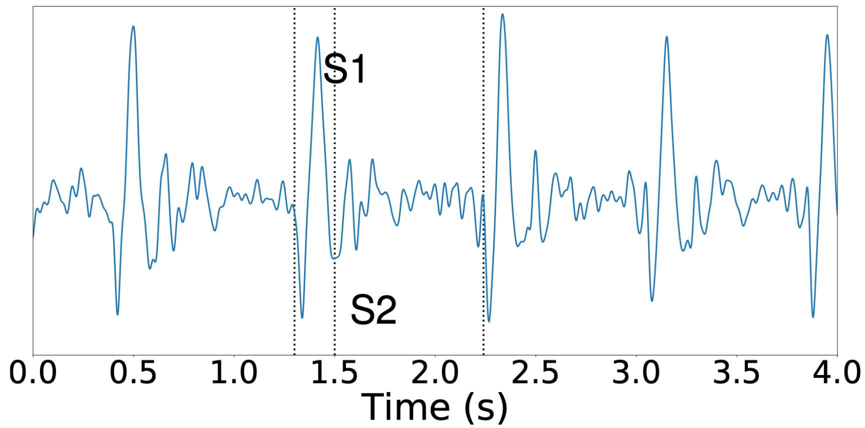

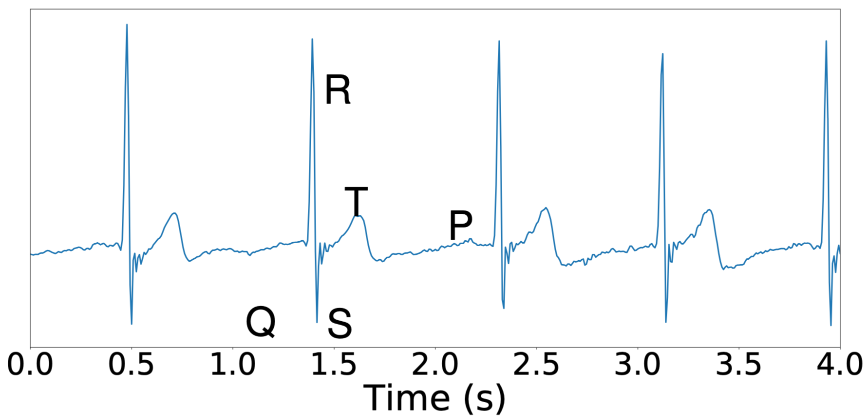

User Identification Performance

Table 2: Power consumption and latency measurement of EarGate.

Scheme	Operation	Power (mW)	Latency (ms)	Energy (mJ)
On-device identification	MicRecd	120	1000	168.59 (All)
	LowPassFilt	635	1.83	136.82 (MFCC)
	FeatExtr (All/MFCC)	655/651	71.98/23.62	
	Inference	644	0.44	
Raw Data Offloading	MicRecd	120	1000	123.17 (WiFi)
	TX [OS+Air] (WiFi)	334	9.49+12.8	190.94 (BT)
	TX [OS+Air] (BT)	478	148.41+128	
Feature Offloading	MicRecd	120	1000	168.91 (WiFi, All)
	LowPassFilt	635	1.83	172.27 (BT, All)
	FeatExtr (All/MFCC)	655/651	71.98/23.62	136.67 (WiFi, MFCC)
	TX [OS+Air] (WiFi)(All/MFCC)	332	1.81+0.59/0.39+0.26	139.26 (BT, MFCC)
	TX [OS+Air] (BT)(All/MFCC)	457	8.66+5.94/5.74+2.56	

Raw data 16KB
Tests on Raspberry PI

<100 ms for on device or offloading
<200 mJ for identification

LifeLearner: Hardware-Aware Meta Continual Learning System for Embedded Computing Platforms

Anonymous

ABSTRACT

Continual Learning (CL) allows applications such as user personalization and household robots to learn on the fly and adapt to context. This is an important feature when context, actions, and users change. However, enabling CL on resource-constrained embedded systems is challenging due to the limited labeled data, memory, and computing capacity.

In this paper, we propose LifeLearner, a hardware-aware meta continual learning system that drastically optimizes system resources (lower memory, latency, energy consumption) while ensuring high accuracy. Specifically, we (1) exploit meta-learning and rehearsal strategies to explicitly cope with data scarcity issues and ensure high accuracy, (2) effectively combine lossless and lossy compression to significantly reduce the resource requirements of CL and rehearsal samples, and (3) developed hardware-aware system on embedded and IoT platforms considering the hardware characteristics.

As a result, LifeLearner achieves near-optimal CL performance, falling short by only 2.8% on accuracy compared to an Oracle baseline. With respect to the state-of-the-art (SOTA) Meta CL method, LifeLearner drastically reduces the memory footprint (by 178.7×), end-to-end latency by 80.8-94.2%, and energy consumption by 80.9-94.2%. In addition, we successfully deployed LifeLearner on two edge devices and a microcontroller unit, thereby enabling efficient CL on resource-constrained platforms where it would be impractical to run SOTA methods and the far-reaching deployment

TinyTrain: Deep Neural Network Training at the Extreme Edge

Anonymous Author(s)
Affiliation
Address
`email`

Abstract

On-device training is essential for user personalisation and privacy. With the pervasiveness of IoT devices and microcontroller units (MCU), this task becomes more challenging due to the constrained memory and compute resources, and the limited availability of labelled user data. Nonetheless, prior works neglect the data scarcity issue, require excessively long training time (*e.g.* a few hours), or induce substantial accuracy loss ($\geq$10%). We propose *TinyTrain*, an on-device training approach that drastically reduces training time by selectively updating parts of the model and explicitly coping with data scarcity. *TinyTrain* introduces a task-adaptive sparse-update method that dynamically selects the layer/channel based on a multi-objective criterion that jointly captures user data, the memory, and the compute capabilities of the target device, leading to high accuracy on unseen tasks with reduced computation and memory footprint. *TinyTrain* outperforms vanilla fine-tuning of the entire network by 3.6-5.0% in accuracy, while reducing the backward-pass memory and computation cost by up to 2,286× and 7.68×, respectively. Targeting broadly used real-world edge devices, *TinyTrain* achieves 9.5× faster and 3.5× more energy-efficient training over status-quo approaches, and 2.8× smaller memory footprint than SOTA approaches, while remaining within

Thanks to:

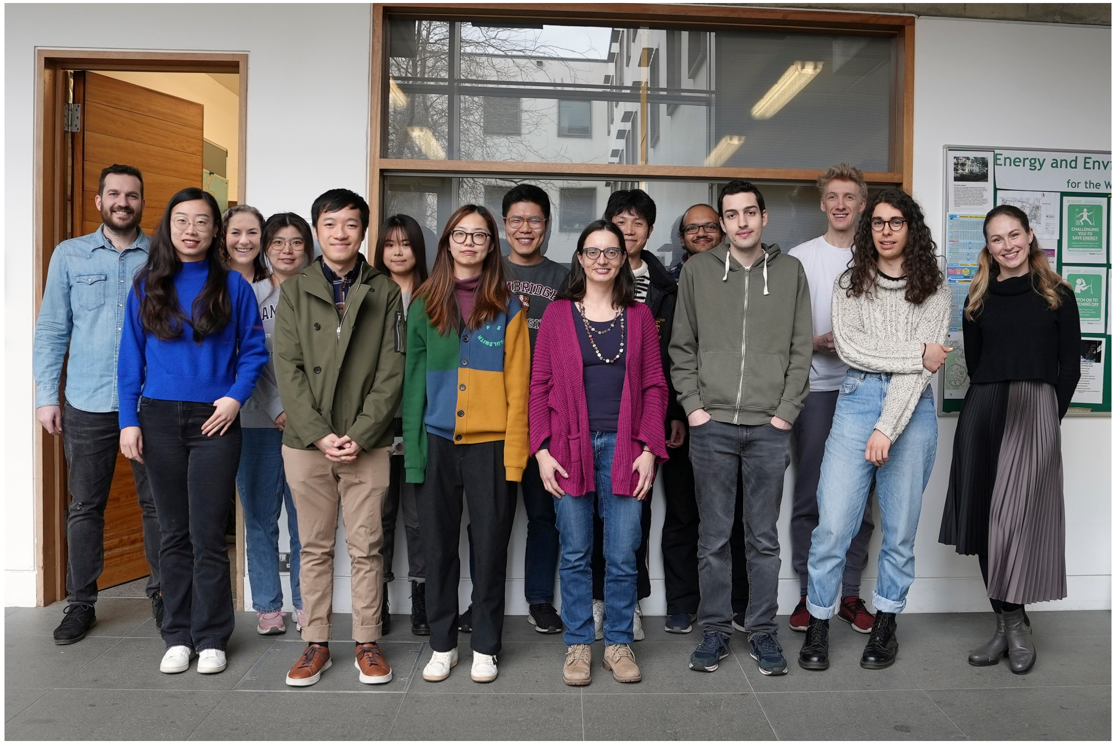

Cecilia Mascolo

cm542@cam.ac.uk

www.cl.cam.ac.uk/~cm542

Yang Liu

yl868@cam.ac.uk

https://yangliu-cs.github.io/YangLiu-CS/