Karlsruhe Institute of Technology

Technology for Pervasive Computing

Recognition of Group Activities using Wearable Sensors

8th International Conference on Mobile and Ubiquitous Systems (MobiQuitous'11)

Dawud Gordon, Jan-Hendrik Hanne, Martin Berchtold, Takashi Miyaki and Michael Beigl

Karlsruhe Institute of Technology (KIT), TecO; TU Braunschweig, AGT Germany

Overview

In-network GAR using Wearable Sensors

- What is GAR?
  - Why is it important?
  - How can it be done?
  - What is the correct approach?
- System for GAR
  - Sensor nodes
  - Mobile phones
  - In-network processing
- Experiment in GAR
  - Different modes evaluated
  - Context abstraction levels
  - Evaluated in terms of power consumption and recognition
- Results
  - Features optimal abstraction level
  - Using HAR as input for GAR creates problems
  - Clustering promising

GAR using Mobile P2P Devices

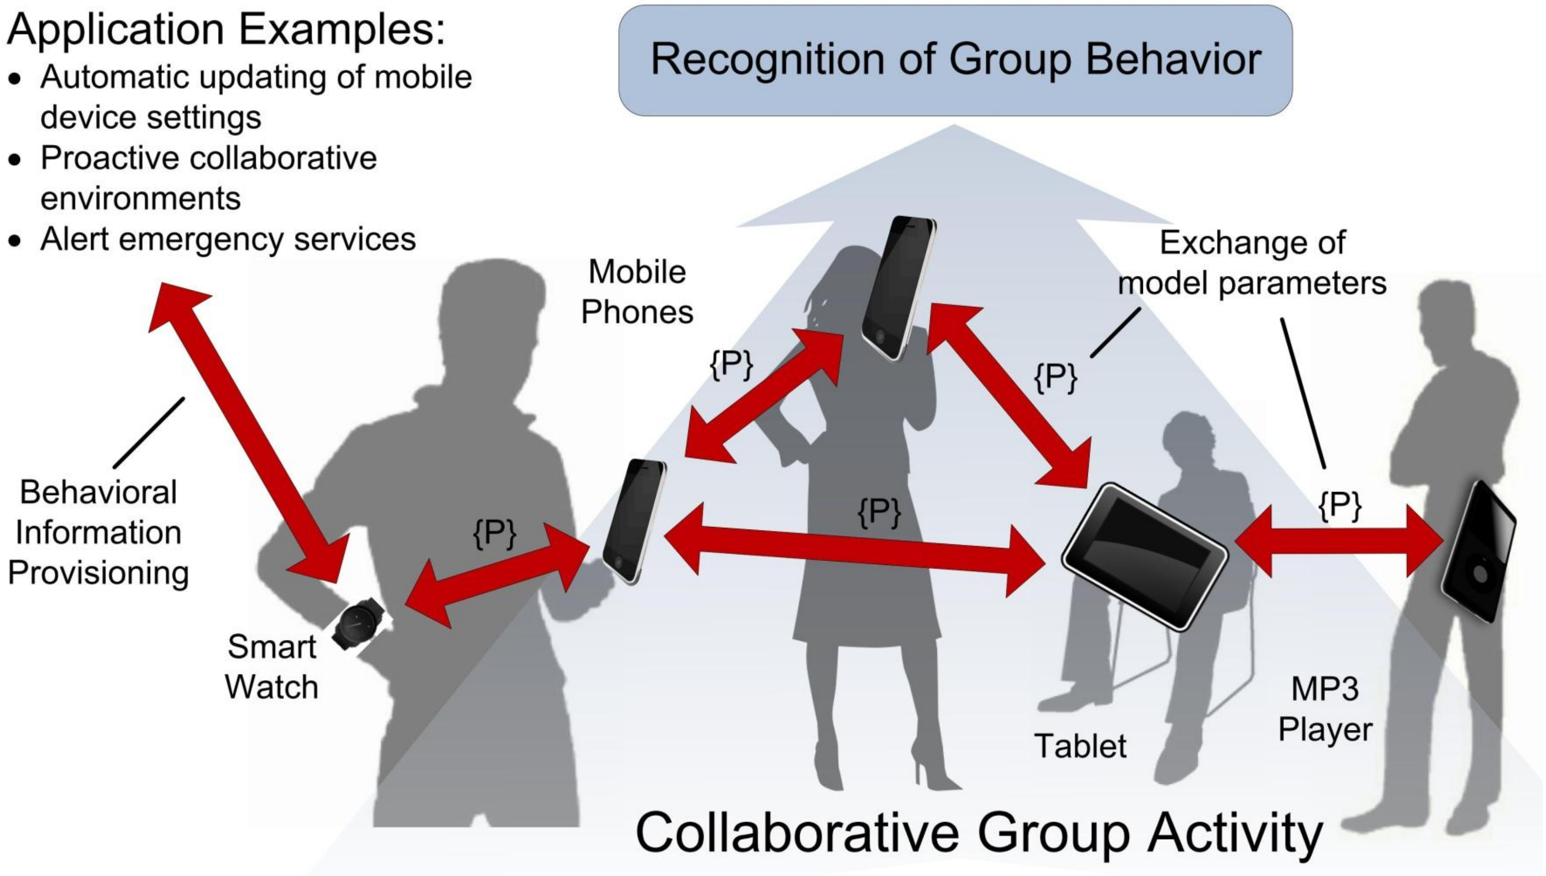

- Devices collaborate to recognize group activity using embedded sensors

How to Approach GAR?

- Group (swarm) behavior studied in the natural kingdom: ants, fish, birds, bees, etc.
- Swarm behavior is emergent behavior resulting from behavior of individuals and interactions between them [Reynolds 1987]
- HAR shown effective for recognizing user activities, interactions
- GAR therefore based on HAR methods

What is Group Activity Recognition?

Karlsruhe Institute of Technology

Bao & Intille 2004

- Observing key points on the body allows activities of the person as a whole to be inferred (HAR)
- In the same way, observing behavior of individuals allows us to infer activities of the group
- The group can be observed as an entity in and of itself. (GAR)

flickr: bade_md

Human Activity Recognition (HAR) using Machine Learning

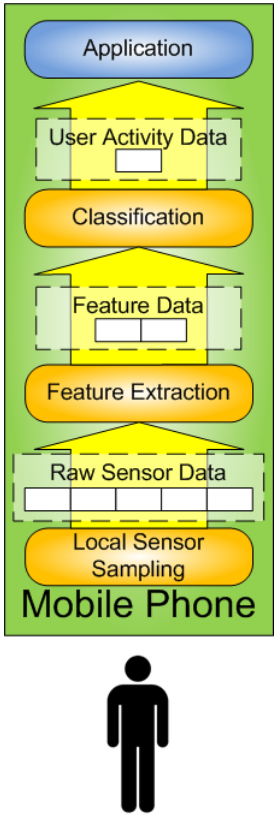

- HAR using mobile sensing devices is an established field.
- Sensor sampling yields discrete measurements of continuous signals
- Windowing allows signal features to be extracted
- Machine learning matches patterns in features to activity labels
- So how do we apply this to groups of individuals?

Group Activity Recognition (GAR)

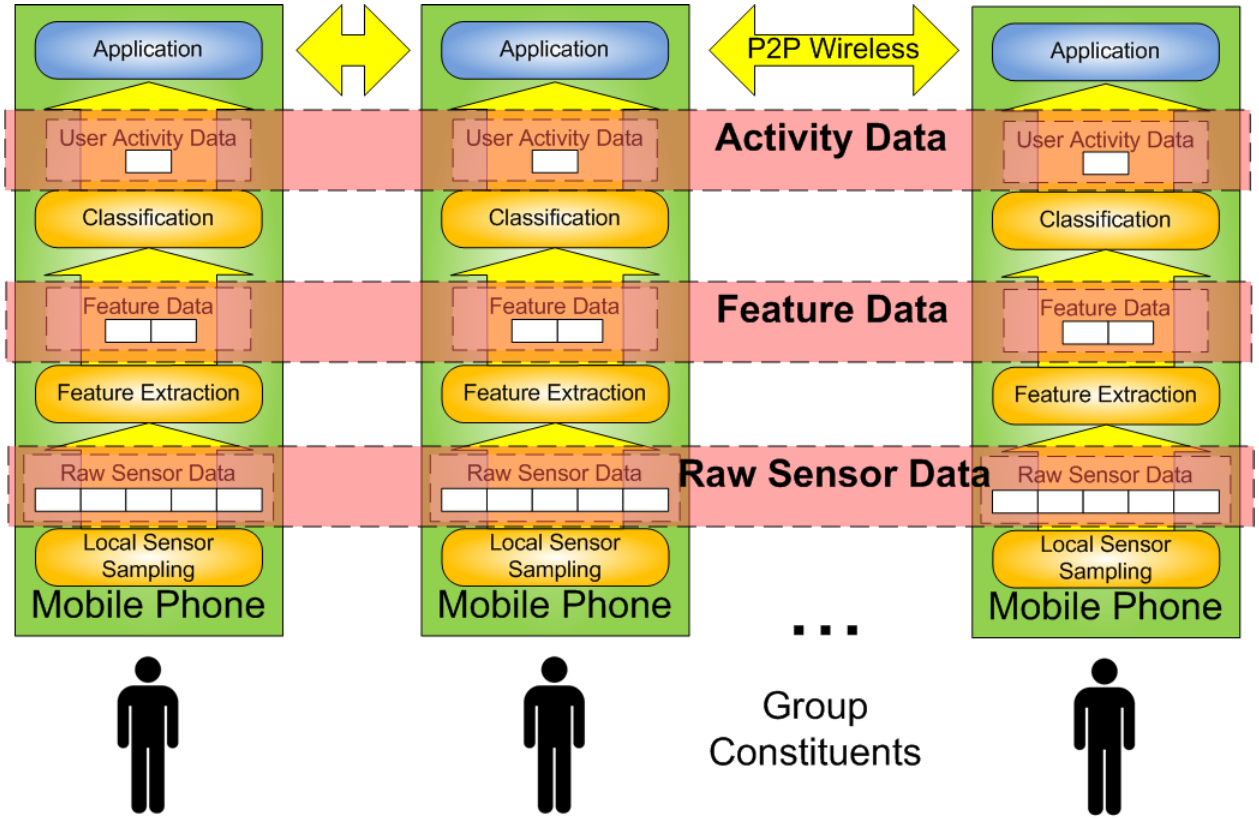

- Single-user data must be fused
- Low abstraction
  - high costs
  - high accuracy
- High abstraction
  - Lower costs
  - but accuracy?
- Where is the tradeoff?

Experiment Hardware: Wireless Sensing

Karlsruhe Institute of Technology

- Open-source, open-hardware sensor node project: www.jennisense.teco.edu
- ContikiOS ported to the Jennic wireless microcontroller from NXP
- Sensing
  - ADXL335 3D acceleration sensor
  - Sampled at 33 Hz
  - (Current version: 3D Acc./Gyro/Compass, light, temp, pressure, infrared distance, time-of-flight)
- Feature extraction
  - Window size of 0.5s w/ 50% overlap
  - Mean and variance only
- Single-user activity recognition
  - Supervised
    - kNN (k=10, no weighting)
    - DT (C4.5)
    - nB (no kernel estimation, single Gaussian)
  - Unsupervised
    - K-means clustering, hard, top 1
    - Uses subtractive clustering for cluster identification

P2P Architecture: Smart-Mugs and Neo

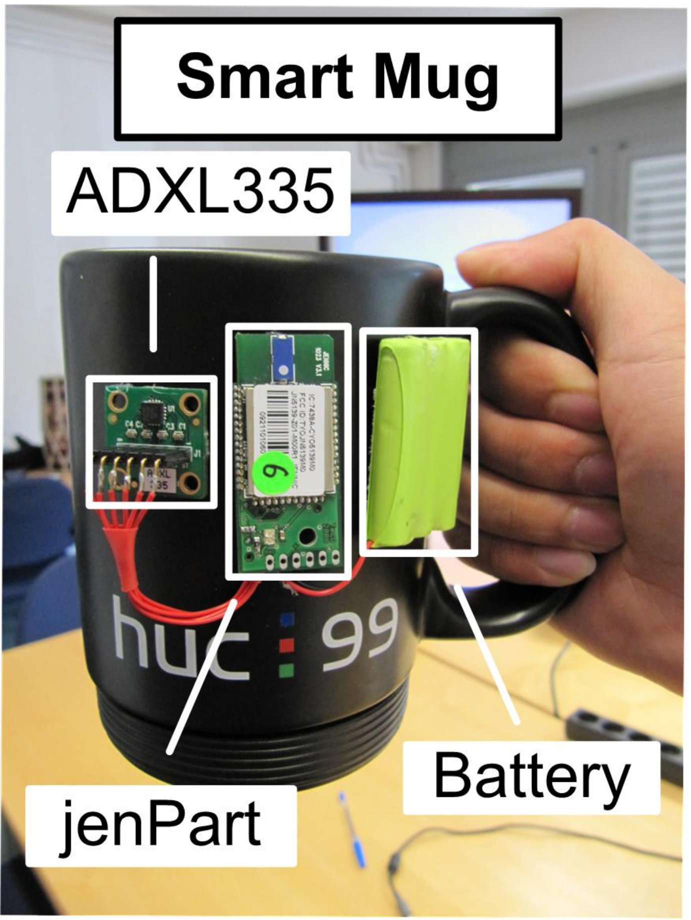

Coffee Cup/ jenPart WSN

802.15.4

Mobile Phone

System:
- jenPart sensor node
- ConTiki OS

Tasks:
- Sensor sampling
- Local feature extraction
- Local activity recognition

System:
- Neo Freerunner
- Debian Linux

Tasks:
- Global multi-user group activity recognition
- Global and local training
- Visualization
- Administration

System operational modes

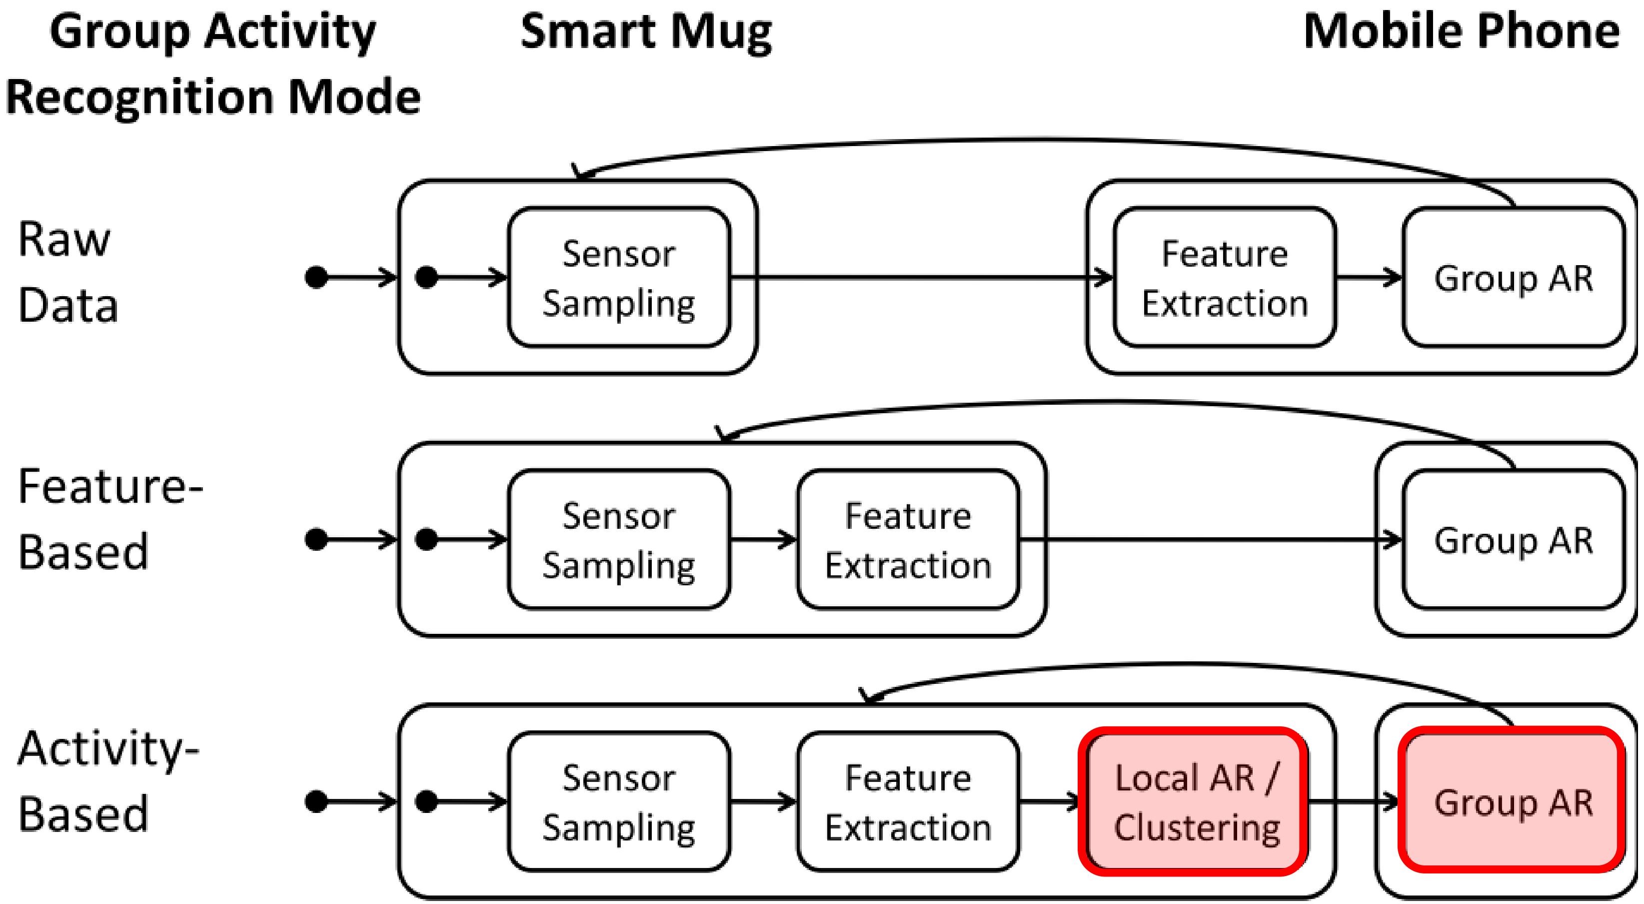

- Doubly-labeling problem

Experiment

- Evaluate GAR rates and power consumption using different data abstraction levels
  - Raw sensor data
  - Sensor signal features
  - Local activities
- Raw sensor data and feature based GAR accuracies identical (feature selection)
- Using local activities = doubly labeling
  - Separate local and global training phases
  - Local clustering (unsupervised)
- Group activities:
  - Meeting, Presentation, Coffee break
- Single-user activities:
  - Mug on table, holding in hand, gesticulating, drinking
- 3 subjects, 45 mins, 22,700 vectors

Experiment

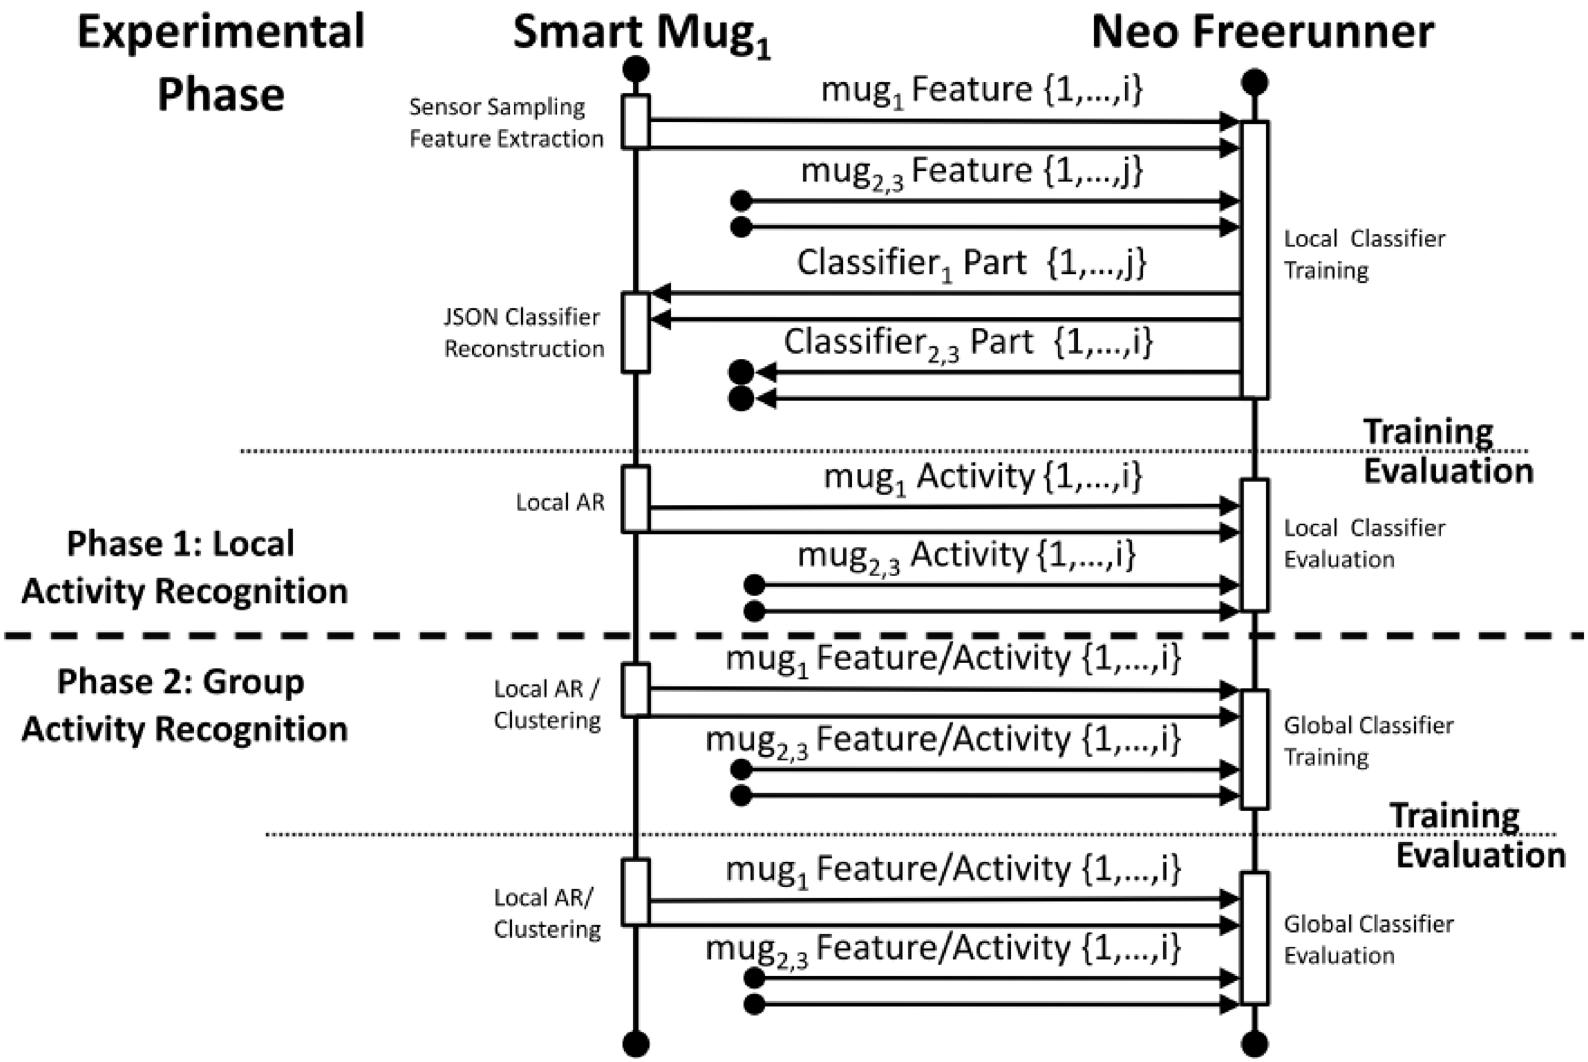

Experimental Phase
Smart Mug$_1$
Neo Freerunner
mug$_1$ Feature {1,...,i}
Sensor Sampling
Feature Extraction
mug$_{2,3}$ Feature {1,...,j}
Local Classifier Training
Classifier$_1$ Part {1,...,j}
JSON Classifier Reconstruction
Classifier$_{2,3}$ Part {1,...,i}
Training
Evaluation
mug$_1$ Activity {1,...,i}
Local AR
Local Classifier Evaluation
Phase 1: Local Activity Recognition
mug$_{2,3}$ Activity {1,...,i}
mug$_1$ Feature/Activity {1,...,i}
Phase 2: Group Activity Recognition
Local AR / Clustering
mug$_{2,3}$ Feature/Activity {1,...,i}
Global Classifier Training
Training
Evaluation
mug$_1$ Feature/Activity {1,...,i}
Local AR/ Clustering
Global Classifier Evaluation
mug$_{2,3}$ Feature/Activity {1,...,i}

Single-User HAR

a) Local Activities (Averaged Over Nodes)

Data Basis	DT		kNN		nB	
	Acc.	F-meas.	Acc.	F-meas.	Acc.	F-meas.
Features	0.958	0.958	0.954	0.955	0.941	0.943

- In total 9 classifiers, 3 per node
- Values averaged over nodes
- High results - indicates simple classification problem
- Little variance over nodes and classifiers

Global GAR Results

b) Global Activities

Data Basis	DT Acc.	DT F-meas.	kNN Acc.	kNN F-meas.	nB Acc.	nB F-meas.
Features	0.962	0.962	0.894	0.898	0.565	0.593
Clusters	0.762	0.764	0.597	0.605	0.491	0.494
Activities	0.507	0.524	0.424	0.484	0.491	0.505

- Feature-based recognition provides decent results – information is there!
  - But (very) naïve Bayes fails – multiple clusters
- Using classified activities produces low GAR rates
  - Data analysis: users could not reproduce own behavior – min/max, variance
- Clustering produces promising results!
  - Hard, top-1 clustering not optimal for kNN, nB
  - Soft clustering approaches should improve on this.

Power Consumption

Mode	Data Volume (B/s)	Neo Freerunner Avg(P) (W)	Smart Mug Avg(P) (mW)	Smart Mug E_{Tx} (mJ)
Raw Data	404.25	1.771	24.574	1.012
Features	107.25	1.723	24.233	0.909
Classes/Clusters	12.375	1.700	23.140	0.605

- Significant reductions in transmitted data volume
- Small reductions in total device power consumption
  - Due to scenario, low sample rate, small number of features and sensors, etc.
- Better indicator is how much energy is spent on communication
  - Still doesn't quit scale with volume
  - Due to packet overheard/scenario paramters

Summary

- HAR can be used to recognize group activities
- Abstracting to features yields 96% recognition, saves 10% transmission energy
- Abstracting to local activities saves 33% more energy, but creates labeling issues
    - Users cannot reproduce behavior under different conditions (50% acc. using activities)
    - Clustering promising (76% with room for improvement)
- Conditions for GAR are different than HAR
    - More distinct clusters due to multi-user (nB results)
- Future work
    - Explore other labeling approaches
    - Soft probabilistic clustering
    - Distribute GAR classification as well

That's All

- Thank You!
  - Questions?