Mining Data From Mobile Devices

Algorithms: Sensing

Spiros Papadimitriou, Tina Eliassi-Rad

Overview

- **Indoor localization**
- Low-level activity detection
  - High-level: in next part (context)

Indoor localization
Overview

- GPS not available indoors
- Cell tower triangulation is far too coarse
- Use other signals instead to infer location or proximity

RSSI (Received Signal Strength Indicator) | Localization

Accelerometer | Proximity

Audio (ambient noise) | Mapping

⋮ | ⋮ ?

WiFi-based localization
[Ladd et al., IROS 2002]

- Signal propagation and attenuation models work reasonably outdoors
- Indoors, signal strength is determined by building geometry, which may change (doors open/close, movements, etc) – very hard to model
- *Supervised* learning approach:
  - Bayesian localization framework
  - Sensor fusion using HMM (people walk slow enough)
- Accuracy of up to 1m

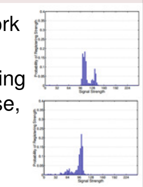
Actual RSSI distributions

Andrew M. Ladd, Kostas E. Bekris, Guillaume Marceau, Algis Rudys, Dan S. Wallach, Lydia E. Kavraki: *Using wireless Ethernet for localization.* IROS 2002: 402-408.

WiFi SLAM (GP-LVM)
[Ferris et al., IJCAI 2007]

- SLAM: Simultaneous Localization and Mapping (robotics)
- Gaussian Process (GP):

$$Y_i = f(X_i) + \varepsilon$$

RSSI position

- Latent Variable Model (LVM): position is unobserved
- Assumptions:
  - Similar locations have similar RSSI – squared exponential kernel
  - Instead of odometry data, assume "office building" constraints on:
    - Distance between successive positions
    - Change in orientation between successive positions
    - Alignment of parallel line segments

Brian Ferris, Dieter Fox, Neil D. Lawrence: *WiFi-SLAM Using Gaussian Process Latent Variable Models.* IJCAI 2007: 2480-2485

WiFi SLAM (GP-LVM)
[Ferris et al., IJCAI 2007]

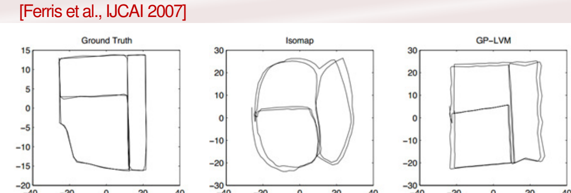

- Initialization: Isomap
- Localization error ~ 4 ± 0.6m

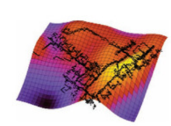
GP model for RSSI of one AP

Brian Ferris, Dieter Fox, Neil D. Lawrence: *WiFi-SLAM Using Gaussian Process Latent Variable Models.* IJCAI 2007: 2480-2485

WiFi GraphSLAM

[Huang et al., ICRA 2011]

- Assumption that similar RSSI corresponds to similar locations is not true in signal-sparse environments
- Additionally, can use phone sensors for basic odometry

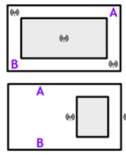

Joseph Huang, David Millman, Morgan Quigley, David Stavens, Sebastian Thrun, Alok Aggarwal: *Efficient, generalized indoor WiFi GraphSLAM.* ICRA 2011: 1038-1043

WiFi GraphSLAM

[Huang et al., ICRA 2011]

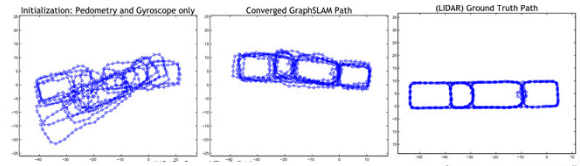

- No trajectory priors!
- Ground truth: LIDAR (~10cm accuracy)
- Mean error ~ 2.2m
  - ~7.1m with pedometry and gyro only

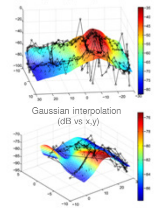

Joseph Huang, David Millman, Morgan Quigley, David Stavens, Sebastian Thrun, Alok Aggarwal: *Efficient, generalized indoor WiFi GraphSLAM.* ICRA 2011: 1038-1043

RSS-space tasks

- May not need mapping to (Euclidean) spatial coordinates
  - Proximity detection
  - Geo-fencing
- Substantially simplifies processing
- Still need some way to filter noise, esp. if a large number of APs and/or measurements are not available

[e.g.] Christos Laoudias, George Constantinou, Marios Constantinides, Silouanos Nicolaou, Demetrios Zeinalipour-Yazti, Christos G. Panayiotou: *The Airplace Indoor Positioning Platform for Android Smartphones.* MDM 2012.

Overview

- Indoor localization
- **Low-level activity detection**

Activity detection

- Active area of research (esp. in networking / ubi-comp)
  - Could probably do an entire tutorial on just this ☺
- Becoming mainstream: Android APIs

Low-level (this part):
- Am I standing/falling, walking/driving, etc?
- How is my mood?
- …

High-level / “context” (later):
- Out with friends, looking for a restaurant
- Commuting to work, drive or take public transit?
- …

Fall detection

[Sposaro, Tyson, EMBS 2009]

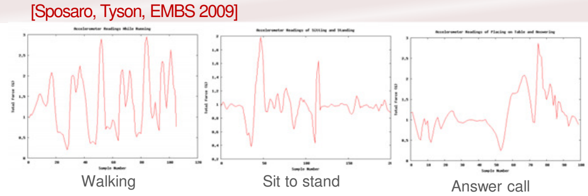

- Use accelerometer data to detect fall
- Relatively simple threshold approach:
  - Acceleration exceeds a threshold
  - Followed by a “still” period
  - In a 90º changed orientation
- Challenges: false positives
- A *lot* of work in this area

[e.g.] Frank Sposaro, Gary Tyson: *iFall: An Android Application for Fall Monitoring and Response.* IEEE EMBS 2009..

Public dataset (Reyes-Ortiz et al., donated 2012): http://archive.ics.uci.edu/ml/datasets/Human+Activity+Recognition+Using+Smartphones

Classification for activity detection

- How about detecting more than a single type of event? E.g.,
  - Standing vs. sitting
  - Walking vs. cycling vs. driving
- Challenges:
  - Phone data can be very noisy (loose phone, many factors)
  - Efficiency of on-phone classification
  - Population variances (one size does not fit all)
  - Sampling frequency (power draw)
  - …

[background] Anguita, D., Ghio, A., Pischiutta, S., Ridella, S.: A Hardware-Friendly Support Vector Machine for Embedded Automotive Applications, IJCNN 2007.

Population-based activity detection

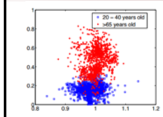

"Walking" features: top-2 PCs

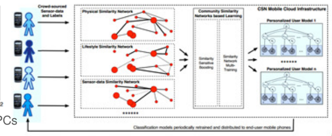

Architecture overview

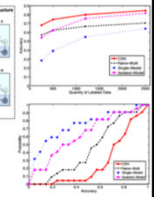

- Collecting data requires substantial effort
- What if we could also use data from other "similar" people?
- "Community Similarity Networks" – 3 similarity measures:
  - Physical similarity (age, height, weight, well-being measures)
  - Lifestyle similarity (mobility patterns, activity distributions)
  - Sensor data similarity ("set": duplicate elimination)
- Performs better than out-of-the box semi-supervised methods

Nicholas D. Lane, Ye Xu, Hong Lu, Shaohan Hu, Tanzeem Choudhury, Andrew T. Campbell, Feng Zhao, *Enabling Large-scale Human Activity Inference on Smartphones using Community Similarity Networks (CSN)*, UbiComp 2011

Mood detection

[LiKamWa et al. 2013]

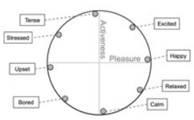

Mood model

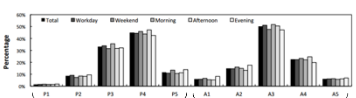

Displeased → Pleased   Inactive → Active

- Can cellphone usage patterns reveal user's mood?
  - Communication patterns
  - App usage patterns
- Accuracy 66%, improved up to 93% over time
  - Vs. self-reported, 32 users

Data Type	Histogram by:	Dimensions
Email contacts	# Messages	10
	# Characters	10
SMS contacts	# Messages	10
	# Characters	10
Phone call contacts	# Calls	10
	Call Duration	10
Website domains	# Visits	10
Location Clusters	# Visits	10
Apps	# App launches	10
	App Duration	10
Categories of Apps	# App launches	12
	App Duration	12
Previous 2 Pleasure and Activeness Averages	N/A	4

Multi-linear regression + (greedy) feature selection

Robert LiKamWa, Yunxin Liu, Nic Lane, and Lin Zhong, *MoodScope: Building a Mood Sensor from Smartphone Usage Patterns*, MobiSys 2013

Android activity recognition API

- Classifier [?] now available in Android APIs
- Apps can request one-shot estimates or event stream

```
public class ActivityRecognitionIntentService extends IntentService {
  ...
  @Override
  protected void onHandleIntent(Intent intent) {
    ...
    if (ActivityRecognitionResult.hasResult(intent)) {
      ActivityRecognitionResult result =
              ActivityRecognitionResult.extractResult(intent);
      DetectedActivity mostProbableActivity =
              result.getMostProbableActivity();
      int confidence = mostProbableActivity.getConfidence();
      int activityType = mostProbableActivity.getType();
      // IN_VEHICLE, ON_BICYCLE, ON_FOOT, STILL, UNKNOWN, TILTING, ...
...
```

http://developer.android.com/training/location/activity-recognition.html

Recap

- Indoor localization (focus: WiFi RSS-based)
- Low-level activity detection

Other:

- Localization using other modalities (e.g., ambient noise: Color app)
- Face detection and recognition (e.g., screen unlock)
- Power consumption logging and mining
- …

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

MINING DATA FROM MOBILE DEVICES

Algorithms: Sensing

Spiros Papadimitriou, Tina Eliassi-Rad

License