MOBICASE2011

Mobile Sensing for Behavior-aware Mobile Computing

Jiang Zhu
Carnegie Mellon University
Behavior-aware Mobile Computing Group (BEWARE)
jiang.zhu@sv.cmu.edu

Research Projects

- Natural language applications for mobile platforms
  - Mobile speech-to-speech translation (**Jibbigo**)
  - Mobile sign translation
  - Strategic reading
  - **Context-aware language processing** (Saluja et al. MT Summit 2011)
  - Supported by DARPA TransTAC, Transformative App and BOLT program
- Statistical natural language algorithms for non-language mobile applications
  - User behavior modeling (MobiCASE 2011)
  - Activity recognition and human behavior grammar induction (LREC 2010, ICDM 2011)
  - Anomaly detection (ICCCN 2011)
  - Geo-trace modeling (Pervasive 2011)
  - Supported by ARO/CyLab, Cisco, Google, Motorola, Yahoo!, Nokia Research, DAPRA TCTO

Mobile Sensing

- Smart phones are very popular now

Smartphone Penetration Worldwide, by Region and Country, 2009-2014 (% of total mobile handsets)

	2009	2010	2011	2012	2013	2014
North America						
US	32%	33%	37%	44%	51%	55%
Canada	30%	31%	34%	40%	47%	50%
Total	32%	33%	37%	44%	51%	54%
Western Europe						
Italy	36%	40%	47%	54%	63%	67%
Germany	17%	19%	22%	25%	29%	33%
France	16%	18%	21%	27%	29%	33%
UK	17%	18%	20%	23%	29%	32%
Rest of Western Europe	31%	36%	41%	49%	58%	64%
Total	25%	28%	32%	37%	44%	49%

* Study by Cisco

- Surprisingly, popular with low-income women too
  - Sunny Xun Liu and Ying Zhang , "A Little World in My Hand —The Use of Smartphones Among Low Income Minority Women," in the Proceedings of Annual Conference of the Association for Education in Journalism and Mass Communication (AEJMC 2011)

Mobile Sensing

- Smart phones come with embedded sensors
  - Accelerometers, gyroscope, magnetometer
  - GPS receiver
  - WiFi, Bluetooth, NFC
  - Microphone, camera,
  - Temperature, light sensor
  - "Clock" and "Calendar"
- Smart phones are connected with the Internet
  - Upload sensor data to the cloud
  - Viewing information computing on the server side

- Users carry the phone almost at all time.
  - Except for a certain group of users who prefer to keep their phones inside the purse
  - My phone "knows" where I am, what I am doing and my future activities.

Mobile Sensing Applications

- Activity logger
  - Mobile lifelogger
  - Activity summarization and time management
- Personalized behavior aware applications
  - Personalized TV
  - Forecasting user's behavior (projecting which ads to show on mobile)
- Personal behavior monitoring
  - Senior care, wellbeing and health monitor
  - Security and anomaly detection
- Social behavior understanding
  - Enhanced social graph
  - Better understanding of users' behavior (as compared to survey and questionnaires)
  - DARPA "Healing Hero" Project

Challenges

- Heterogeneous sensors
  - Different sensors have different meaning, format, sampling rate.
- Low-level activity recognition vs. high-level activity understanding
  - Low-level: walking, running, sitting
  - High-level: walk inside the house vs. walk in the park
- Ad-hoc approach vs. more generalized framework
- Very few labeled data
  - Annotating sensor data is time consuming and tedious, human annotators can only annotate high-level activities
- Sensibility of mobile phones
  - The sensors available and positions they are carried may not have enough information to differentiate 'driving' vs. 'in a car as passenger'
- Privacy
- Power cinsumption

Behavior Text: Core Research Idea

- Human behavior/activities share some common properties with natural languages
  - Meanings are composed from meanings of building blocks
  - Expressed as a sequence (time-series)
  - Exists an underlying structure (grammar)

Natural Language	activity language	Example
Word	Atomic Movement	Turn upper body left
Phrase	Movement	Stand up
Sentence	Action	Climb up stairs
Paragraph	Activity	Enter building, climb up stairs and walk into office
Document	Event	Left home and ride bicycle to campus arrived at my office at 2nd floor

- My phone is "writing" a book of my lives: Book of Joy
  - Apply Statistical Natural Language Processing to process the mobile sensor data
  - Information retrieval, machine translation, text categorization, summarization, prediction

Convert Multi-Sensory Data to Behavioral Text

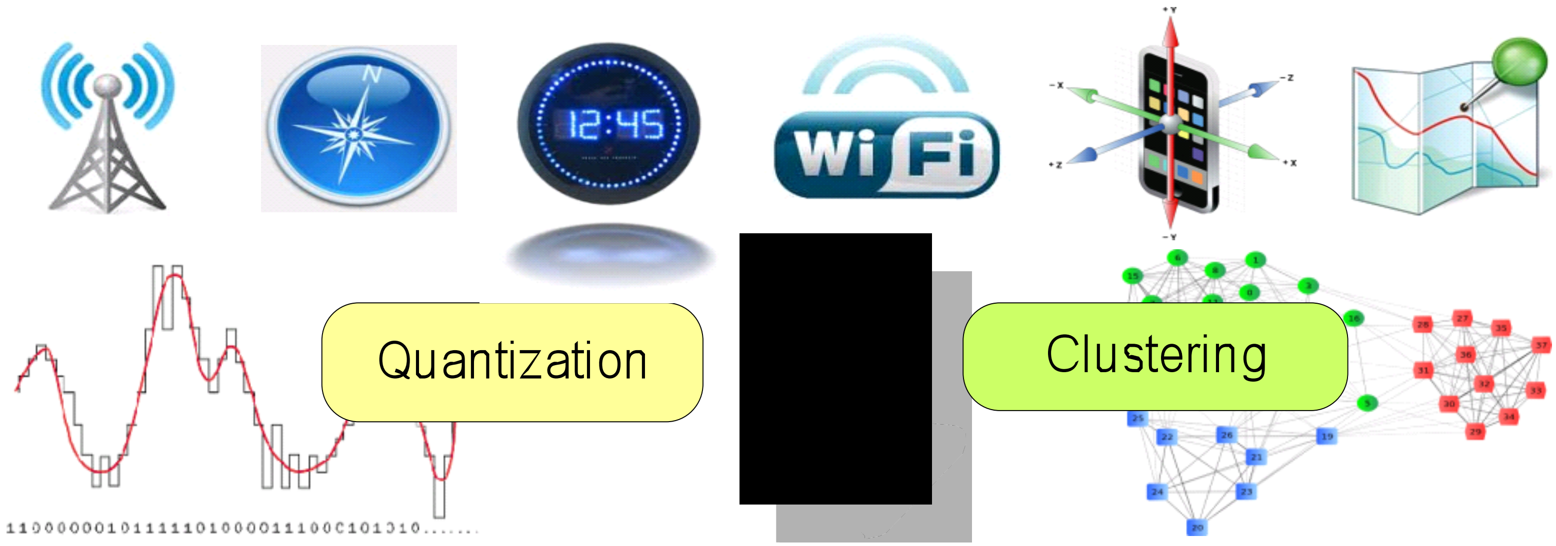

[31,271,37] [37,281,42] [37,276,47] [42,271,47] [42,266,53] [58,271,47] [53,271,47] [74,271,42] . . .

⇓

CZ DG GI FK C BI CS DC HQ BX FI FI BX FI O . . .

CMU MobiSens: Social Sensing Platform

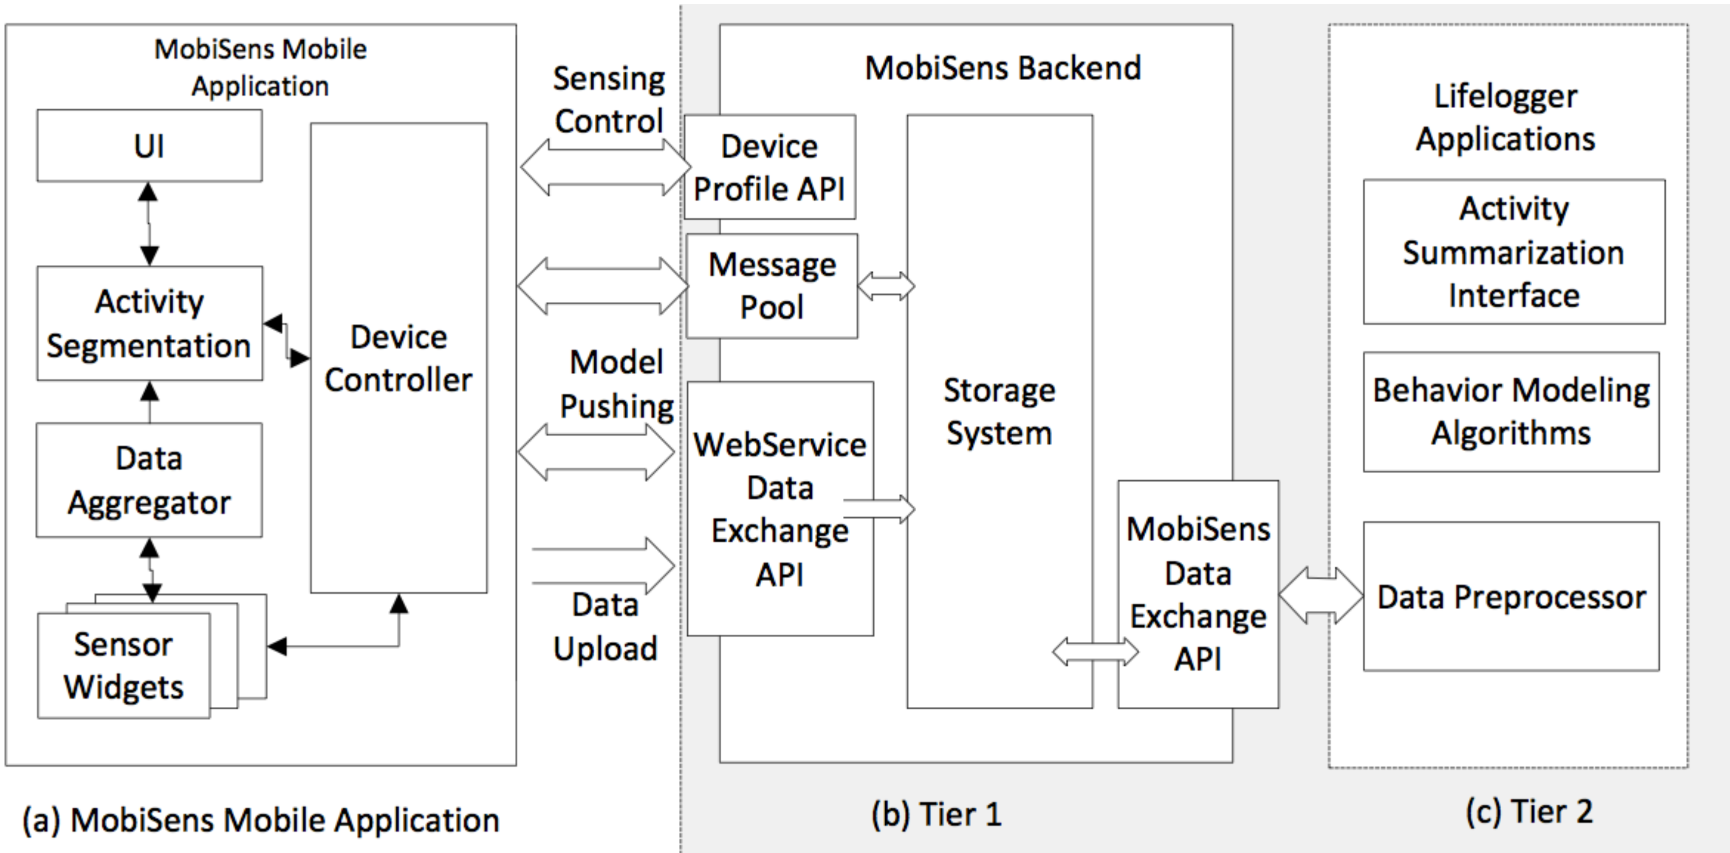

Tier 2 Example: Lifelogger (www.lifelogger.info)

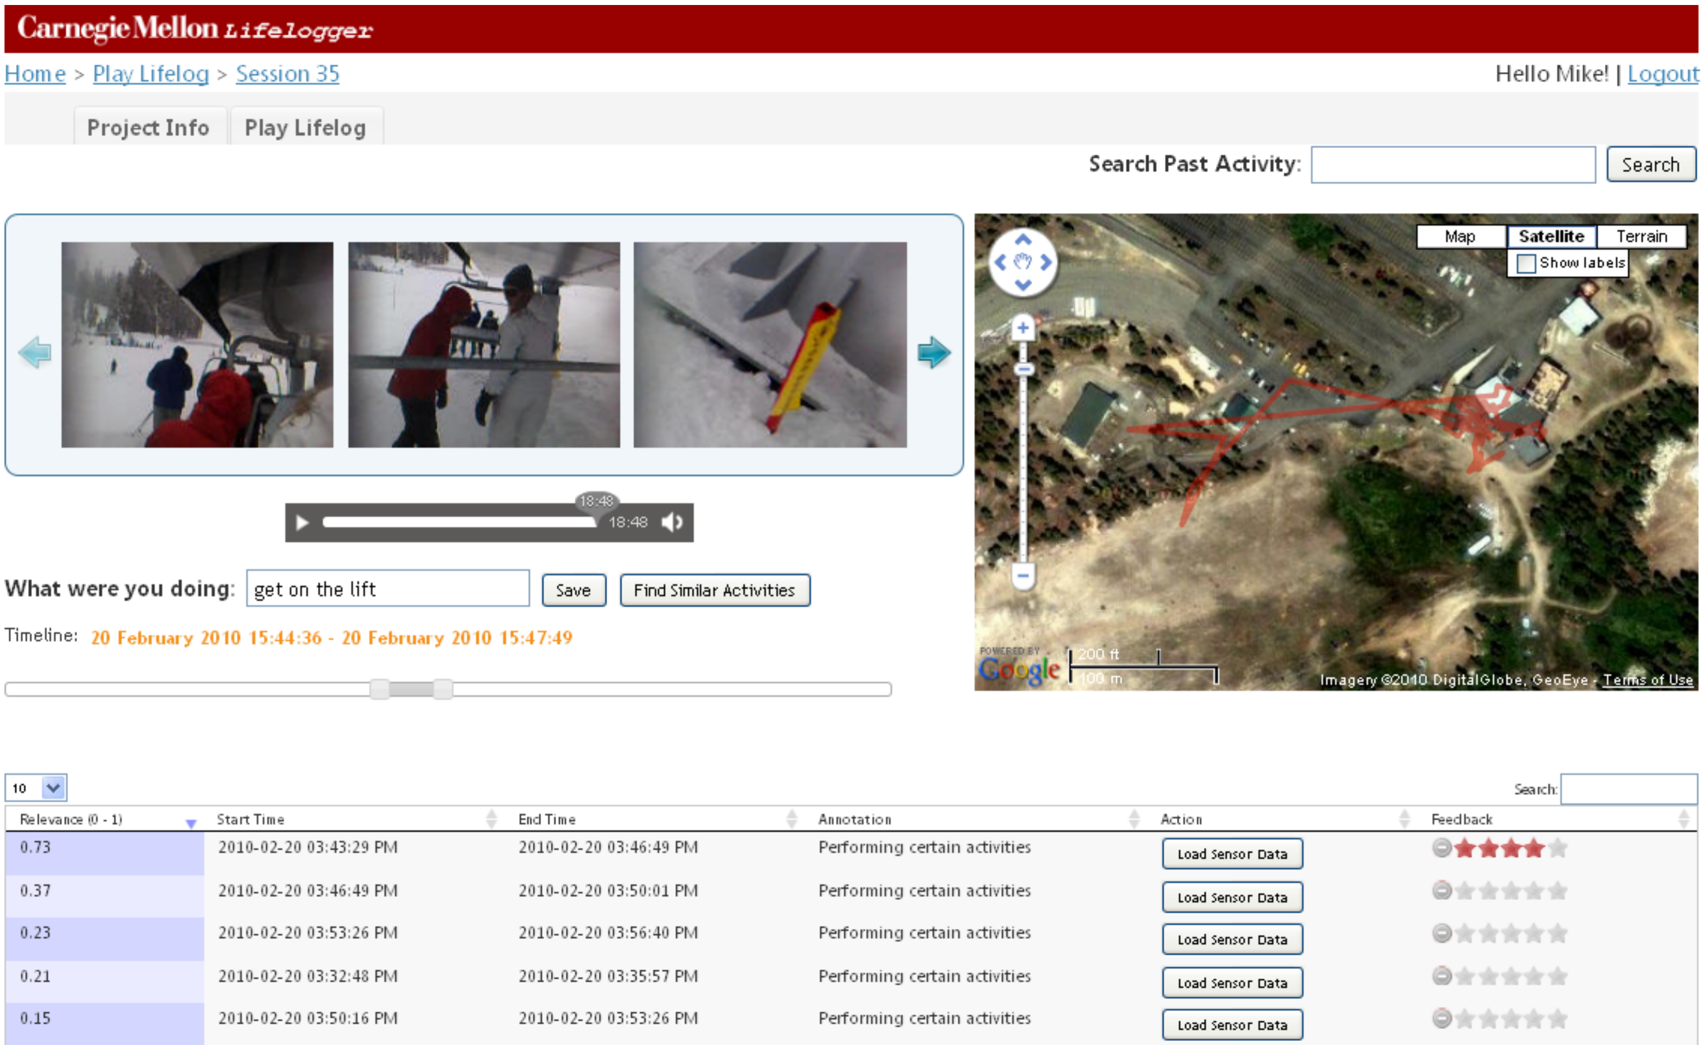

Social Sensing Platforms

- Simple Context by NRC
  - Collects all sensor readings
- Funf (MIT Media Lab)
  - Data collection only
  - Multiple data modalities.
  - All data stored locally
- CenceMe (Miluzzo et al.)
  - Two-tier build-in classifier (Can recognized a limited set of activities)
  - Multiple data modalities
  - Presence sharing

CMU MobiSens App

- Mobile App for Android (available on Android's Market)
  - Records all mobile sensor information
  - **On-device automatic activity segmentation and recognition**
  - View your friend's location and what they are doing now

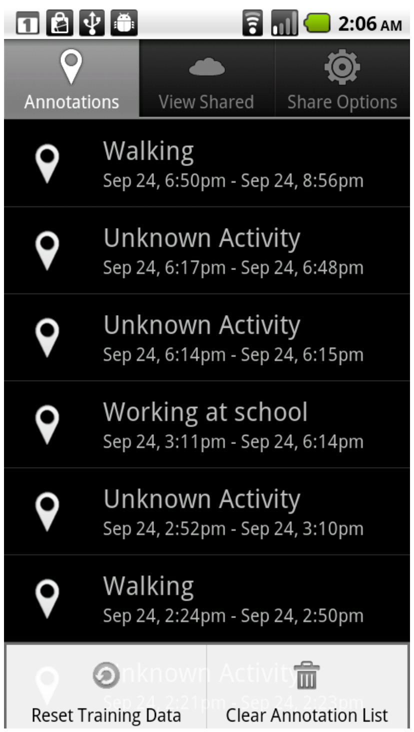

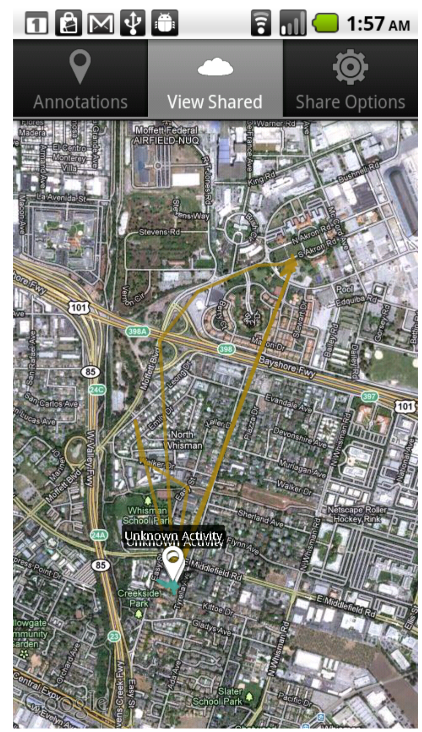

Logging all Mobile Sensors

Sensors	Information Sensed
3-axis accelerometer	motion
Magnetometer	azimuth value of heading direction
GPS coordinates	outdoor locations and trace
Rotation matrix[1]	orientation of the phone
Ambient light level	lighting level
Temperature sensor	environment temperature
WiFi RSS	indoor location and WiFi provider
Battery level	power consumption of the device
Process list	which application/service is running
Network stats	network traffic to/from applications

- Non-trivial task
  - Time stamp
  - Working around OS
  - Robustness
  - Transparency
  - Power consumption

Ask the User to Annotate Their Activity?

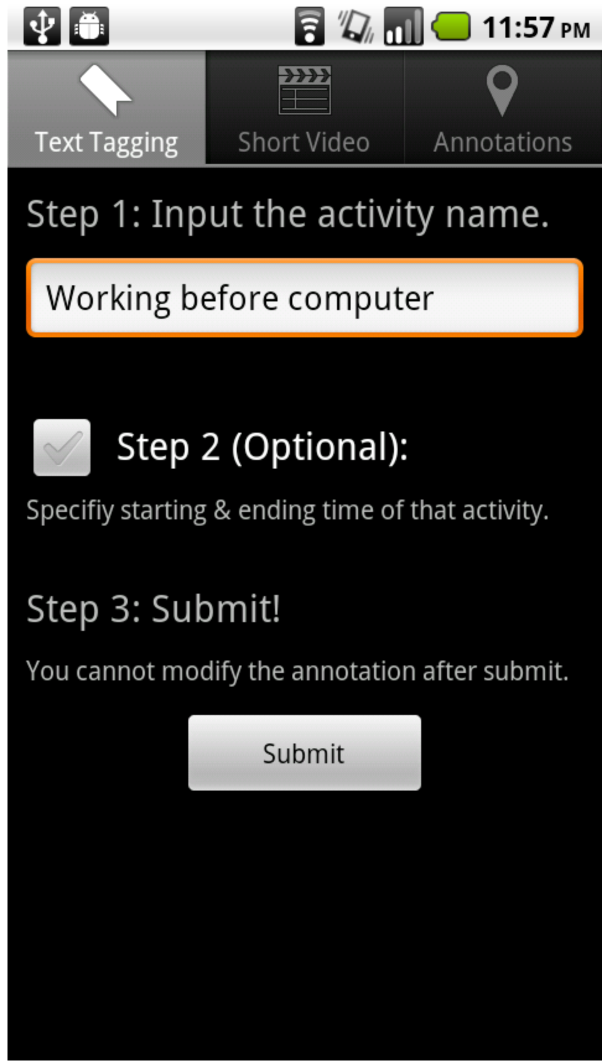

MobiSens Release One Labeling UI

Users won’t cooperate

- A result of self-report activity test:
  - Only 21 labels were submitted in 2 weeks.
  - 71% of these labels were done while the event was happening.
  - No one label events that happened more than 20 minutes ago.
  - No one labels events that were going to happen.
  - The users were less likely to label the same activity twice.

Activity Changes Overtime

AE EF2 GT AE R TGH IK LF MN NB AE EF TR AE GG YTF PO QK XZ TG YY JN JN I

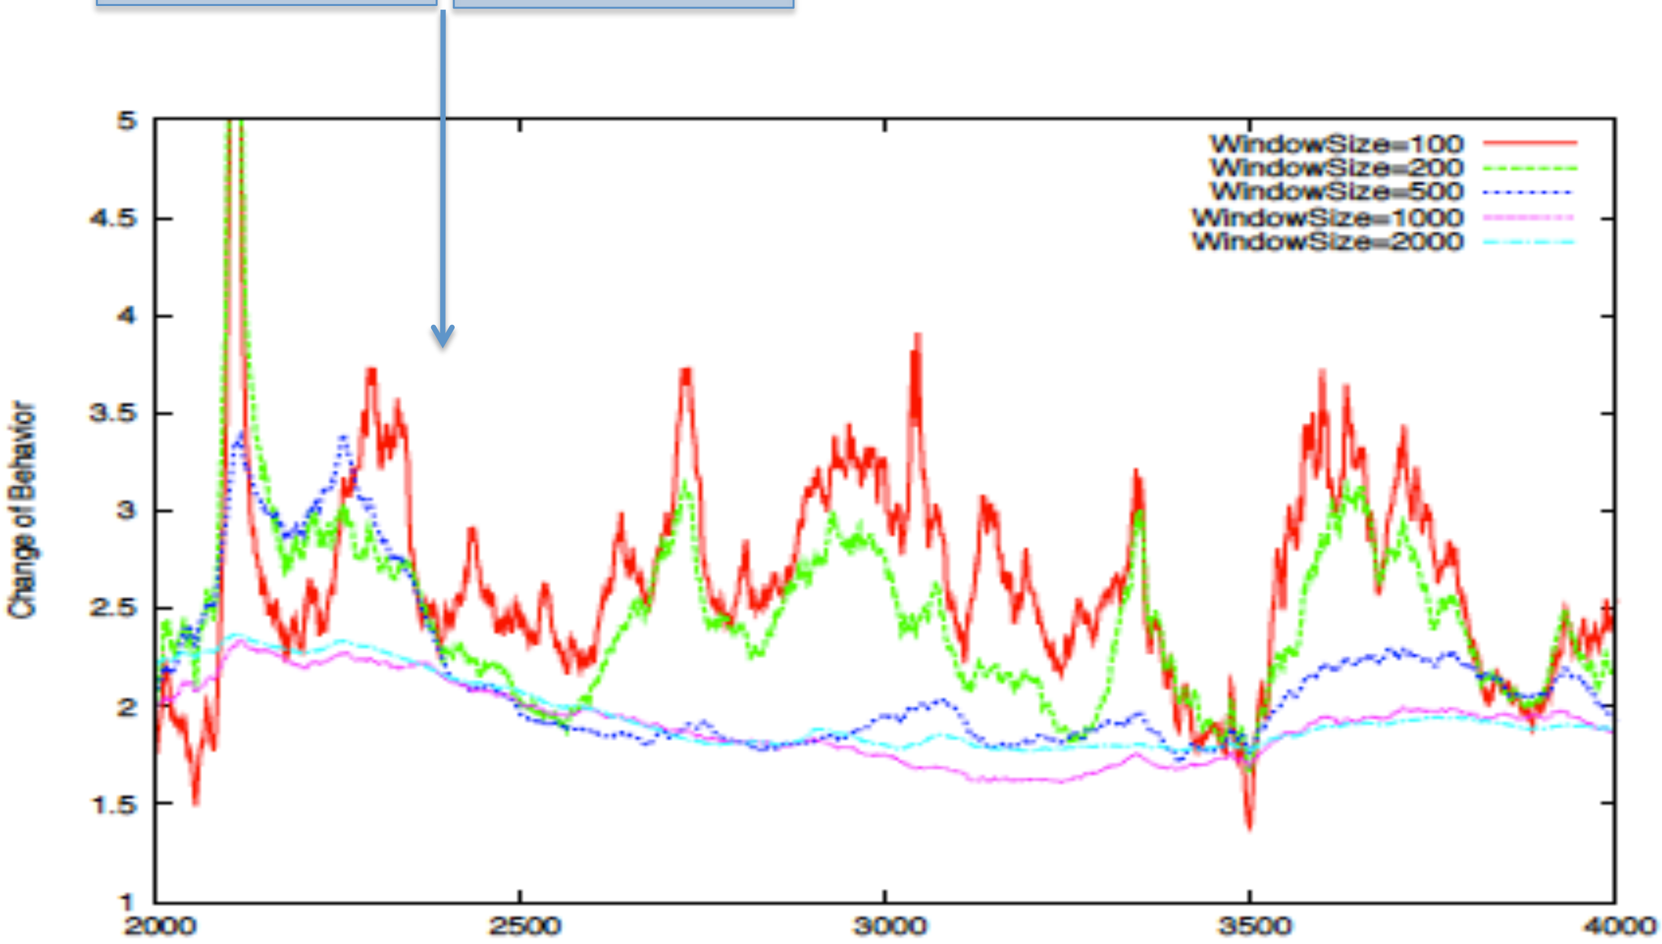

Segmentation on Labeled Activity

$$F = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

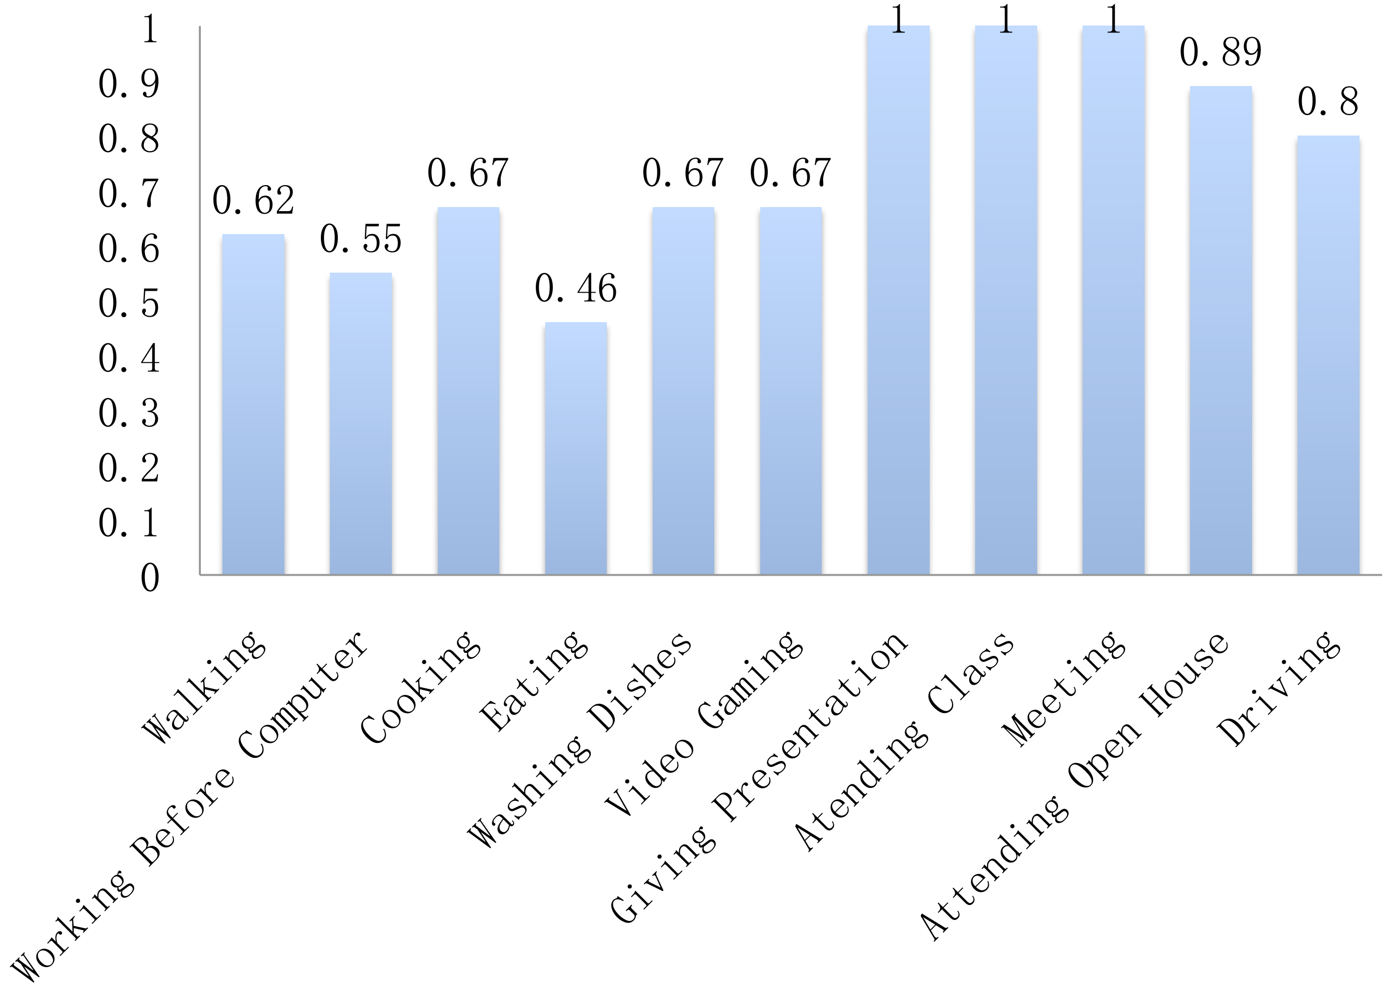

Light-weight Activity Recognition on Mobile Devices

- Recognize activities that users have already annotated before
- Light-weight algorithm for mobile devices: KNN
- Convert ngram frequencies into a vector

  A_1 = ABAA = {$freq_A$,$freq_B$,$freq_{AB}$,$freq_{BA}$,$freq_{AA}$}

  A_2 = BBAA = {$freq_A$,$freq_B$,$freq_{BB}$,$freq_{BA}$,$freq_{AA}$}

- Cosine similarity between annotated example A1 and new sensor input A2

$$Sim(A_1, A_2) = \frac{A_1 \bullet A_2}{\| A_1 \| \| A_2 \|}$$

Automatic Activity Segmentation/Recognition Significantly Increases Users' Annotation

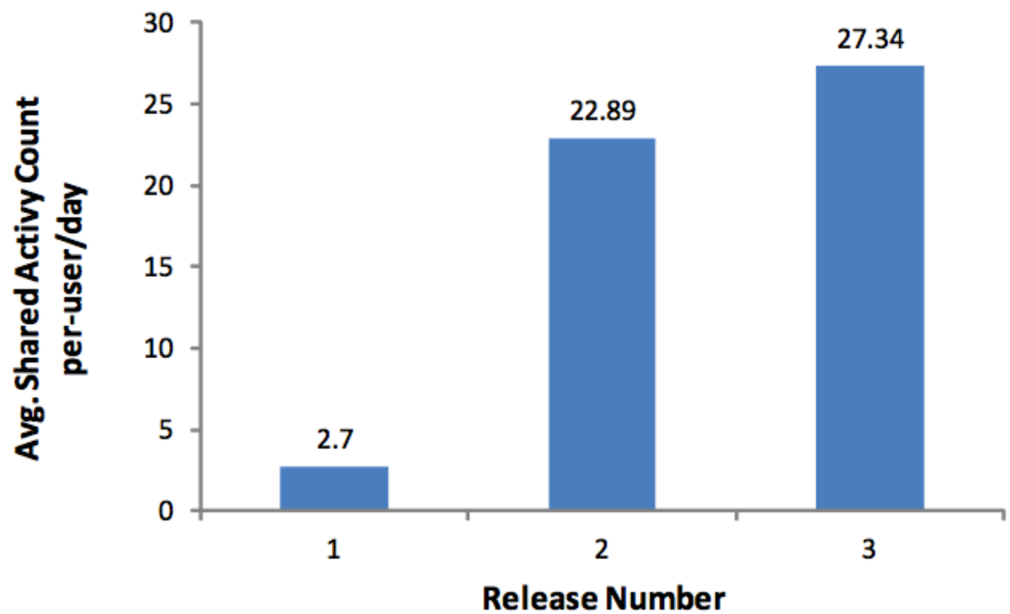

- Release 1: user specify time range and what they were doing.
- Release 2: unsupervised activity segmentation, semi-supervised activity recognition with user correction
- Release 3: users can share their activities among friends

- 3 months since its release: downloaded to 234 devices, 9102 hours of data from 65 users were collected. 4,000 activities annotated.

Power Consumption Optimization

- Sampling accelerometers and WIFI is cheap
- GPS sampling is big power drainer
- Sampling at a flat rate (3 minutes/request) can only maintain 26 hours.
  - Using Droid
  - No phone/data functions

Dynamic GPS Sampling

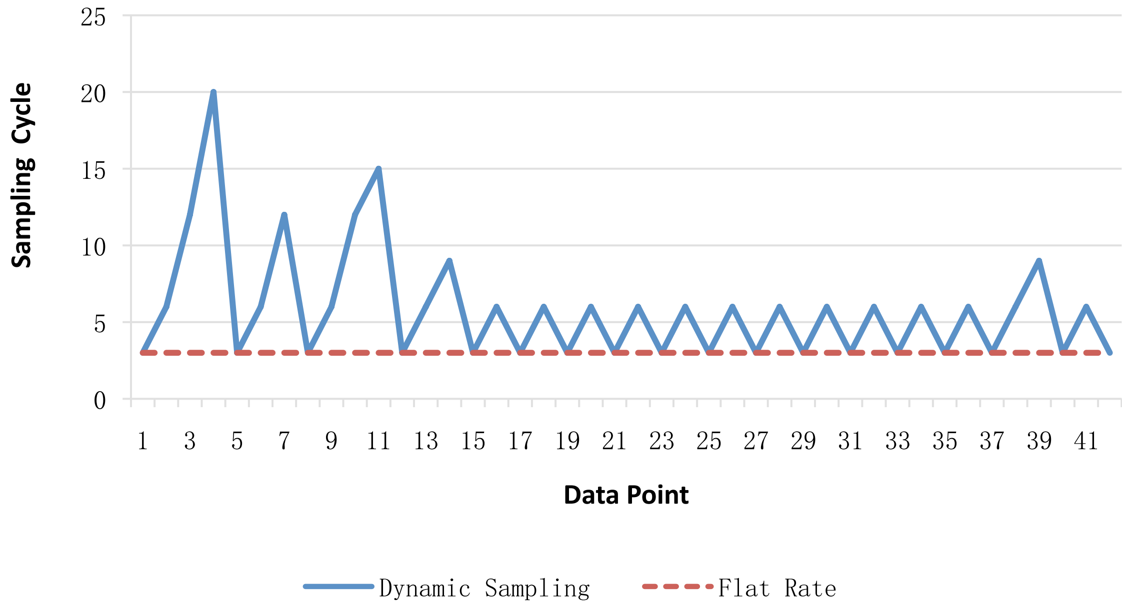

- Adjust the sampling rate based on GPS location changes
- Using WIFI to detect if users' location is changed
- Increased battery life by 25%-40%
- Empirically 12 hours of battery life with normal usage.

Hierarchical Activity Segmentation and Recognition (MobiSens Backend)

- Hierarchically detect the "change" of topic over the behavior text
- Annotate the activity based on similarities to labeled

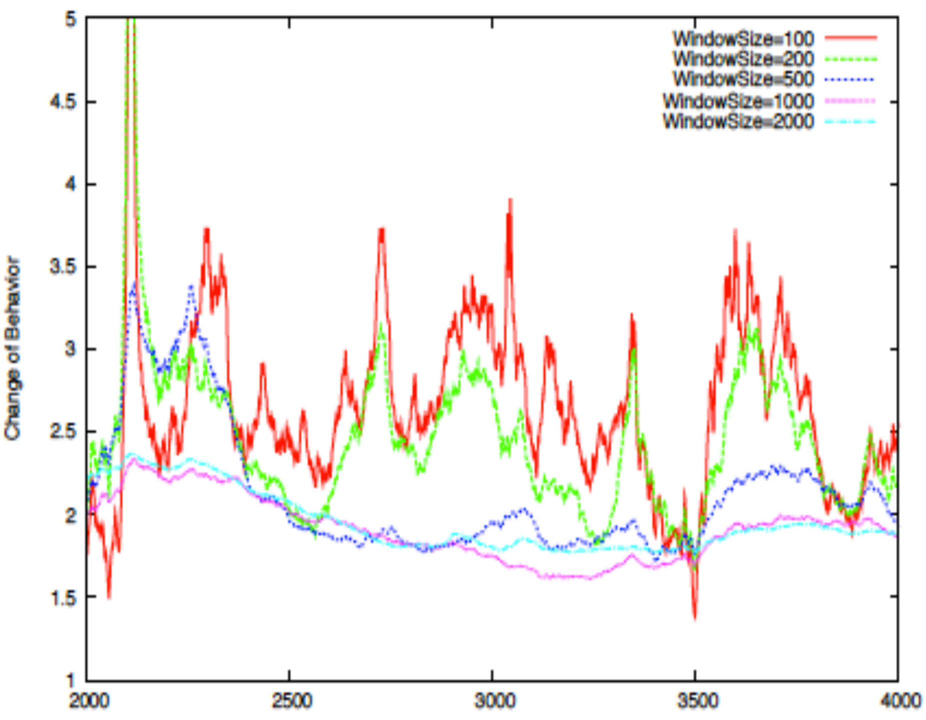

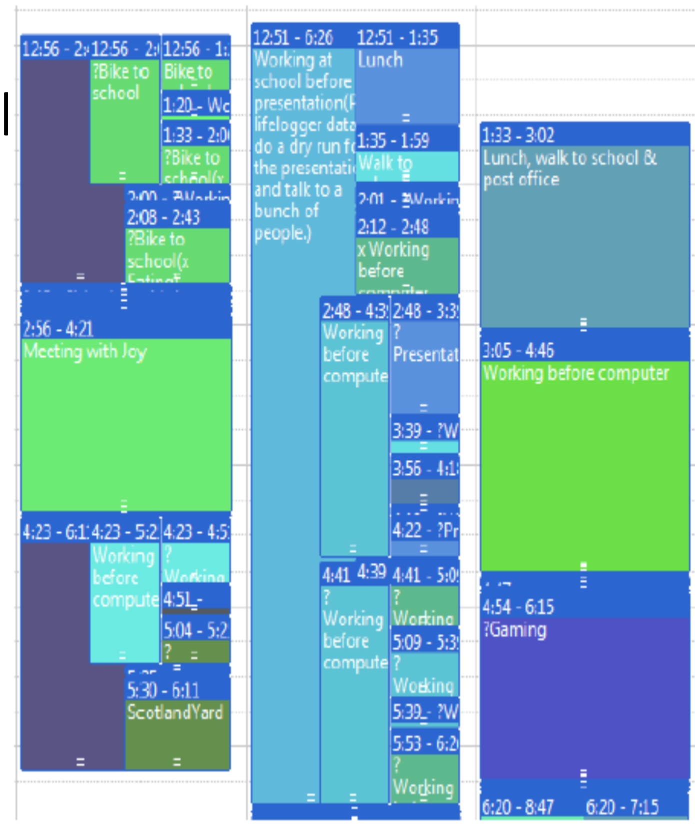

High-level Activity Recognition through Unsupervised Grammar Induction for Human Activity

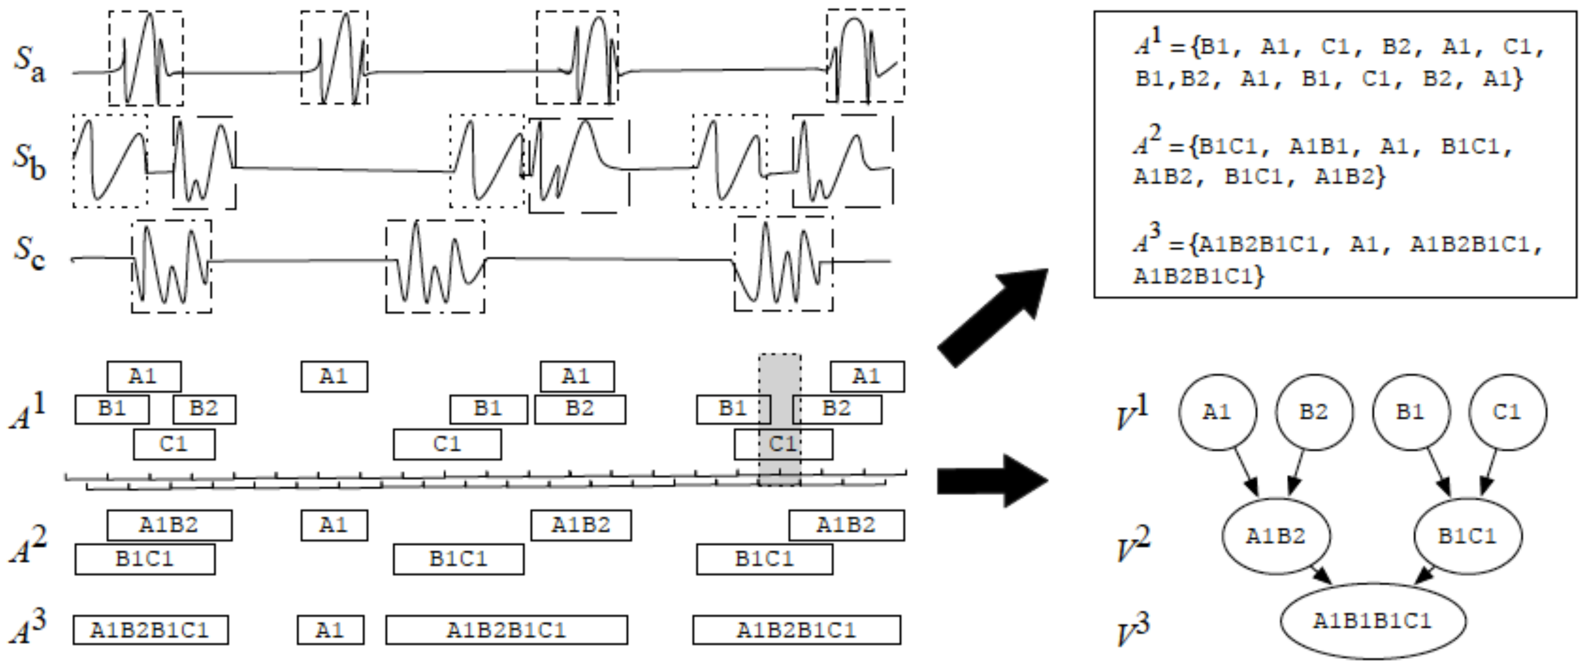

- Recursively identify *constituents* in human activity
- Two constituents are merged as a new category if they have similar *content* and *context* (similar to the Context Constituent Grammar induction framework by Dan Klein).
- Inducted grammar is used to recognize high-level meaning activities.

* Peng et al., "Helix: Unsupervised Grammar Induction for Structured Human Activity Recognition", ICDM 2011

Lifelogger: Demo

Application 1: Activity Recognition through Language Modeling

- Generative language model: P(English sentence) given a model
  - P("President Obma has signed the Bill of … "| Politics) >> P("President Obma has signed the Bill of … " | Sports)
  - LM reflects the n-gram distribution of the training data: domain, genre, topics.
- With labeled behavior text data, we can train a LM for each activity type: "walking"-LM, "running"-LM and classify the activity as $i^* = \arg\max_i P(t|a_i)$

	Predicted Activity		
	walking	running	cycling
walking	95%	1%	4%
running	4%	94%	2%
cycling	2%	0%	98%

Chen et al., 2010 LREC

Application 2: Anomaly Detection

- Is this play Shakespeare's work?

- Comparing the play to Shakespeare's known library of works,
- Track words and phrases patterns in the data
- P (unknown play| known Shakespeare's work)
  - Greater than threshold? Yeah, we found another Shakespeare's work
  - Or, Plagiarism

Application 2: Anomaly Detection

- Trained with users' "normal" behavior data and flag on events when probability of the observed data is too low
  - Elderly users accidental fall (Baki et al., 2009)
  - Abnormal geo-trace: mobile phone stolen, elderly lost (Buthpitya et al., 2011; Zhu and Zhang, 2011), both with about 86% accuracy.

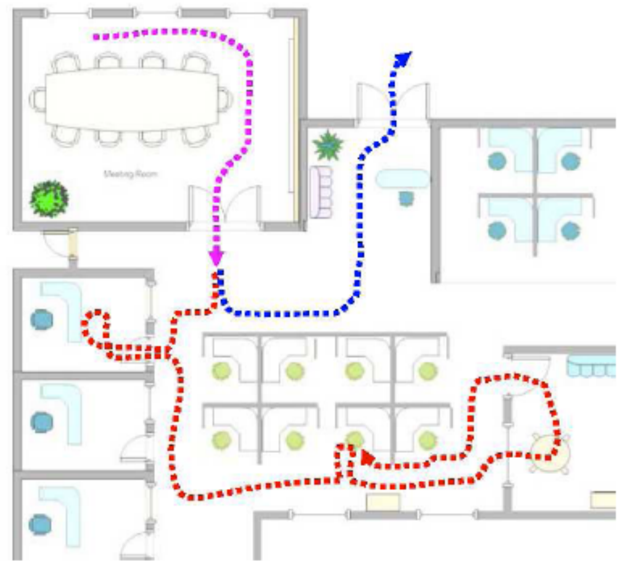

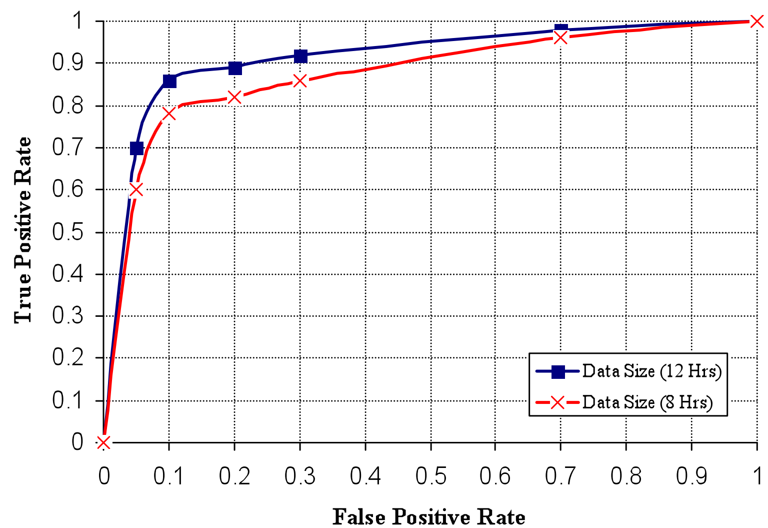

Application 3: SensCare

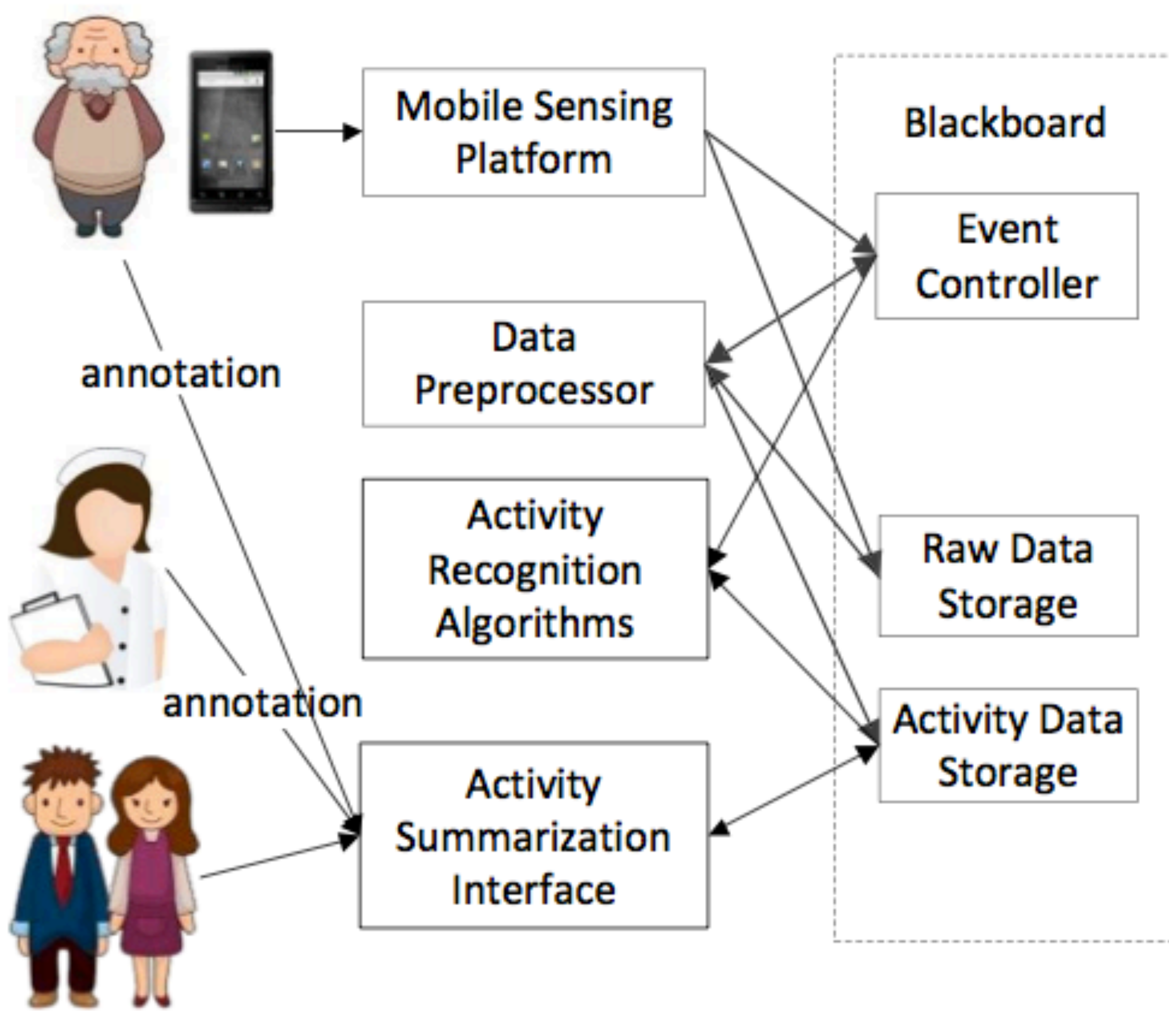

Conclusion

- Behavior text representation as an effective method to model human behavior
- Despite challenges, great opportunities and exciting applications
- CMU MobiSens: a social sensing platform.
- On-going work
  - Societal-scale sensing
  - Behavior-aware mobile applications

References

- Huan-Kai Peng, Pang Wu, Jiang Zhu, and Ying Zhang, "Helix: Unsupervised Grammar Induction for Structured Human Activity Recognition," to appear in the Proceedings of The IEEE International Conference on Data Mining series (ICDM), Vancouver, Canada, Dec 11-14, 2011.
- Avneesh Saluja, Ian Lane and Ying Zhang , "Context-aware Language Modeling for Conversational Speech Translation," to appear in the Proceedings of Machine Translation Summit XIII, Xiamen, China, Sep 19-23, 2011.
- Pang Wu, Huan-Kai Peng, Jiang Zhu and Ying Zhang , "SensCare: Semi-Automatic Activity Summarization System for Elderly Care," to appear in the Proceedings of The 3rd International Conference on Mobile Computing, Applications, and Services (MobiCASE 2011), Los Angeles, CA, October 24-27, 2011.
- Thanh-Le Nguyen and Ying Zhang , "Probabilistic Infrastructureless Positioning in the Pocket," to appear in the Proceedings of The 3rd International Conference on Mobile Computing, Applications, and Services (MobiCASE 2011), Los Angeles, CA, October 24-27, 2011.
- Jiang Zhu and Ying Zhang , " Towards Accountable Mobility Model: A Language Approach on User Behavior Modeling in Office WiFi Networks ," To appear in the Proceedings of The IEEE International Conference on Computer Communications and Networks (ICCCN 2011) , Maui, Hawaii, July 31 - August 4, 2011.
- Senaka Buthpitiya , Ying Zhang, Anind Dey and Martin Griss, "n-gram Geo-Trace Modeling," In the Proceedings of the Ninth International Conference on Pervasive Computing , San Francisco, CA, June 12-15, 2011.
- Peng-Wen Chen, Snehal Kumar Chennuru, Senaka Buthpitiya and Ying Zhang, "A Language-Based Approach to Indexing Heterogeneous Multimedia Lifelog," In the Proceedings of The Twelfth International Conference on Multimodal Interfaces (ICMI) , Beijing ,China , November 8-12, 2010.
- Snehal Kumar Chennuru, Peng-Wen Chen, Jiang Zhu and Ying Zhang, " Mobile Lifelogger - recording, indexing, and understanding a mobile user's life ," In the Proceedings of The Second International Conference on Mobile Computing, Applications, and Services , Santa Clara, CA, Oct 25-28, 2010.
- Peng-Wen Chen, Snehal Kumar Chennuru and Ying Zhang, "A Language Approach to Modeling Human Behavior," In the Proceedings of the seventh international conference on Language Resources and Evaluation (LREC 2010) , Malta, May 17-23, 2010.