

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 7

Human mobility at the aggregate level (descriptive)

- Billions of anonymized Call Detail Records (CDRs) from a cellular network
- Characterized **daily travel**, **carbon emissions**, **number of workers** and **event goers**, and **traffic volumes** of hundreds of thousands of people

Richard A. Becker, Ramon Caceres, Karrie Hanson, Sibren Isaacman, Ji Meng Loh, Margaret Martonosi, James Rowland, Simon Urbanek, Alexander Varshavsky, Chris Volinsky: **Human mobility characterization from cellular network data**. Commun. ACM 56(1): 74-82 (2013).

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 8

Human mobility (descriptive)

- M. Kim, D. Kotz, S. Kim: **Extracting a mobility model from real user traces**. In InfoCom 2006:1-13
- K. Lee, S. Hong, S. Kim, I. Rhee, S. Chong: **Slaw: A new mobility model for human walks**. In InfoCom 2009, 855-863
- Z. Li, B. Ding, J. Han, R. Kays, P. Nye: **Mining periodic behaviors for moving objects**. In KDD 2010, 1099-1108
- M. Martino, F. Calabrese, G. Di Lorenzo, C. Andris, L. Liu, C. Ratti: **Ocean of Information: Fusing Aggregate & Individual Dynamics for Metropolitan Analysis**. IUI 2010: 357-360
- M. Kim, D. Kotz, D: **Identifying unusual days**. Journal of Computing Science and Engineering 5(1), 2011:71-84
- ...

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 9

Work in human mobility

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 10

Human mobility at the individual level (predictive; short-term; GPS)

- “Where are you going to be in the next hour?”**
 - Successful techniques: **hidden Markov models**, **random walk based formalisms**
 - Performance around **3-5 km off**; classification **accuracy low 90%**
 - Learning from GPS alone
 - D. Ashbrook, T. Stamer: **Using GPS to learn significant locations and predict movement across multiple users**. Personal Ubiquitous Comput. 7, 2003:275-286.
 - L. Liao, D. Fox, H. Kautz: **Location-based activity recognition using relational Markov networks**. In IJCAI 2005.
 - J. Krumm, E. Horvitz: **Predestination: Inferring destinations from partial trajectories**. In UbiComp 2006: 243-260.
 - B. Ziebart, A. Maas, A. Dey, J. Bagnell: **Navigate like a cabbie: Probabilistic reasoning from observed context-aware behavior**. In UbiComp 2008:322-331.
 - B. Shaw, J. Shea, S. Sinha, A. Hogue. **Learning to Rank for Spatiotemporal Search**. WSDM 2013: 717-726.
- ...

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 11

Human mobility at the individual level (predictive; long-term; GPS)

- “Where are you going to be 285 days from now at 2PM?”**
- FarOut
 - Identifies periodicity via Fourier analysis (mapping time to frequency)
 - Uses PCA for pattern extraction
 - Utilizes PCA-based classification
- Data: 32K days worth of GPS data across 703 subjects (½ people; ½ cars)
- High variance in area across subjects
 - From 30 to more than 10^8 km²
 - Surface area of earth = 5.2×10^8 km²
- Number of contiguous days = 7 to 1247 ($\mu = 45.9$; $\sigma = 117.8$)
- Captures both continuous (raw GPS) and discretized (triangular cells) data
- Each subject has a matrix D, where each row is a day.
 - Performance **continuous rep.**: 1 km off; baseline 2.5km off
 - Performance **discrete rep.**: **80% accuracy up to 80 weeks into the future**; baseline ~60%

Adam Sadilek and John Krumm: **Far Out: Predicting Long-Term Human Mobility**, AAAI 2012.

Figure 3: Our continuous vector representation of a day d consists of the median latitude and longitude for each hour of the day (00:00 through 23:59), binary encoding of the day of week, and a binary feature signifying whether a national holiday falls on d .

Figure 4: Our cell-based vector representation of a day d encodes the probability distribution over dominant cells conditioned on the time within d , and the same day-of-week and holiday information as the continuous representation (last 8 elements).

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 12

Human mobility at the individual level (predictive; long-term but < 24 hours)

- N. Eagle and A. Pentland: **Eigenbehaviors: Identifying structure in routine**. Behavioral Ecology & Sociobiology 63(7), 2009:1057-1066
- Predictions up to 12 hours into the future
- Class labels: {Home, Elsewhere, Work, No Signal, Off}.
- PCA based classification
- 79% accuracy

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 19

Spatial search on a LBSN (using all available data, GPS, check-ins, etc)

- Best models of social venue search: LambdaMart with a laundry list of features:
 - Spatial score, timeliness, popularity, here now, personal history, creator, mayor, friends here now, personal history w/ time of day

Model	P@1
Random	0.009
Spatial only	0.201
User history only	0.358
Popularity only	0.143
Linear regression: spatial + temporal	0.230
Linear regression: spatial + temporal + popularity	0.251
Linear regression: all features	0.434
Coordinate ascent: all features w/ nonlinear pairs	0.493
LambdaMART: all features	0.531

• B. Shaw, J. Shea, S. Sinha, A. Hogue. [Learning to Rank for Spatiotemporal Search](#). WSDM 2013: 717-726.
 • Q. Wu, C. J. Burges, K. M. Svore, and J. Gao. [Adapting boosting for information retrieval measures](#). Information Retrieval, 13(3):254-270, June 2010.

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 20

Location & Activity

- MobileMiner [Wang et al. SIGMOD 2009]

The diagram shows a 'Profile Mining Platform' that receives 'Moving record stream' and 'Calling record stream' from mobile devices. It processes these through 'Novel data mining techniques' including 'Sequence Mining', 'Clustering', and 'Graphic Mining'. The platform also interacts with 'Domain Knowledge'. The results are used in the 'Application Layer' for 'Mobile Customer Segmentation', 'Social Community Discovery', 'Churn Prediction', and 'Customer Re-entry Analysis'.

Tengjiao Wang, Bishan Yang, Jun Gao, Dongqing Yang, Shiwei Tang, Haoyu Wu, Kedong Liu, Jian Pei: MobileMiner: A real world case study of data mining in mobile communication. SIGMOD 2009: 1083-1086

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 21

Context includes...

- Location/Local**
 - What resources are nearby?
 - Where are you?
- Social**
 - Who are you with?
- Activity**
 - What are you doing?

A Venn diagram with three overlapping circles labeled 'Location', 'Social', and 'Activity'. The 'Location' circle is shaded yellow, and the intersection of all three circles is shaded red.

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 22

Location & Social & Activity

- Lots of work analyzing Twitter data in this space
- A recent best paper winner is from U. of Rochester
- Flap uses a dynamic Bayesian network per user to predict his/her locations given location of friends, time of day, type of day
- Experiments on over 4M tweets from users in LA and NYC
- It can correctly place a user within a **100m radius** (location prediction) with **up to 85% accuracy** (link prediction)
- Adam Sadilek, Henry A. Kautz, Jeffrey P. Bigham: [Finding your friends and following them to where you are](#). WSDM 2012:723-732 (best paper).
- Adam Sadilek, Henry Kautz, Jeffrey P. Bigham: [Modeling The Interplay of People's Location, Interactions, and Social Ties](#). IJCAI 2013.

Figure 2: Flap's visualization of a sample of geo-active friends in NYC. Red links between users represent friendships.

RUTGERS Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 23

Context includes...

- Location/Local**
 - What resources are nearby?
 - Where are you?
- Social**
 - Who are you with?
- Activity**
 - What are you doing?

A Venn diagram with three overlapping circles labeled 'Location', 'Social', and 'Activity'. The 'Location' circle is shaded yellow, and the intersection of all three circles is shaded red.

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY Mining Data from Mobile Devices / Papadimitriou, Eliassi-Rad 24

MINING DATA FROM MOBILE DEVICES

Algorithms: Location & Context

Spiros Papadimitriou, Tina Eliassi-Rad
[\[spiros.papadimitriou,tina.eliassi@rutgers.edu\]](mailto:spiros.papadimitriou,tina.eliassi@rutgers.edu)

License



These slides are made available under a Creative Commons Attribution-ShareAlike license (CC BY-SA 3.0):

<http://creativecommons.org/licenses/by-sa/3.0/>

You can share and remix this work, provided that you keep the attribution to the original authors intact, and that, if you alter, transform, or build upon this work, you may distribute the resulting work only under the same or similar license to this one.

© 2013, 2015 Spiros Papadimitriou, Tina Eliassi-Rad