Big Data Analytics

FOSCA GIANNOTTI AND LUCA PAPPALARDO

HTTP://DIDAWIKI.DI.UNIPI.IT/DOKU.PHP/BIGDATAANALYTICS/BDA/

DIPARTIMENTO DI INFORMATICA - Università di Pisa
anno accademico 2018/2019

Mobility Data Mining

PATTERNS&MODELS

Mobility Profiles

Derived patterns and models

- Combination & refinement of basic patterns and models

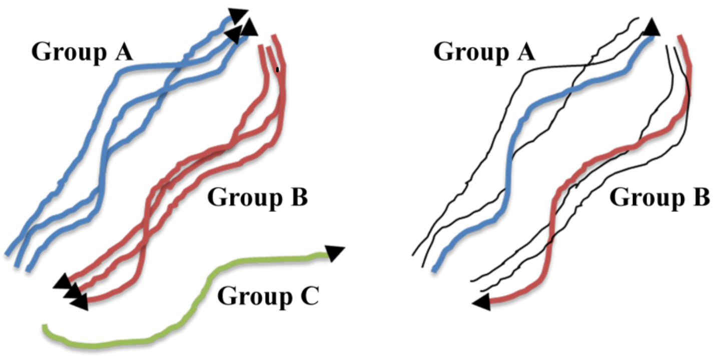

- Individual Mobility Profile: routine

User’s Mobility Profile

Given the user history as an ordered sequence of spatio-temporal points, we want to extract a set of *routines* in order to create the his\her *mobility profile.*

Where:

- A *Routine* is a typical local behavior of the user.
- A *Mobility profile* is the set of user’s routines

User's Mobility Profile

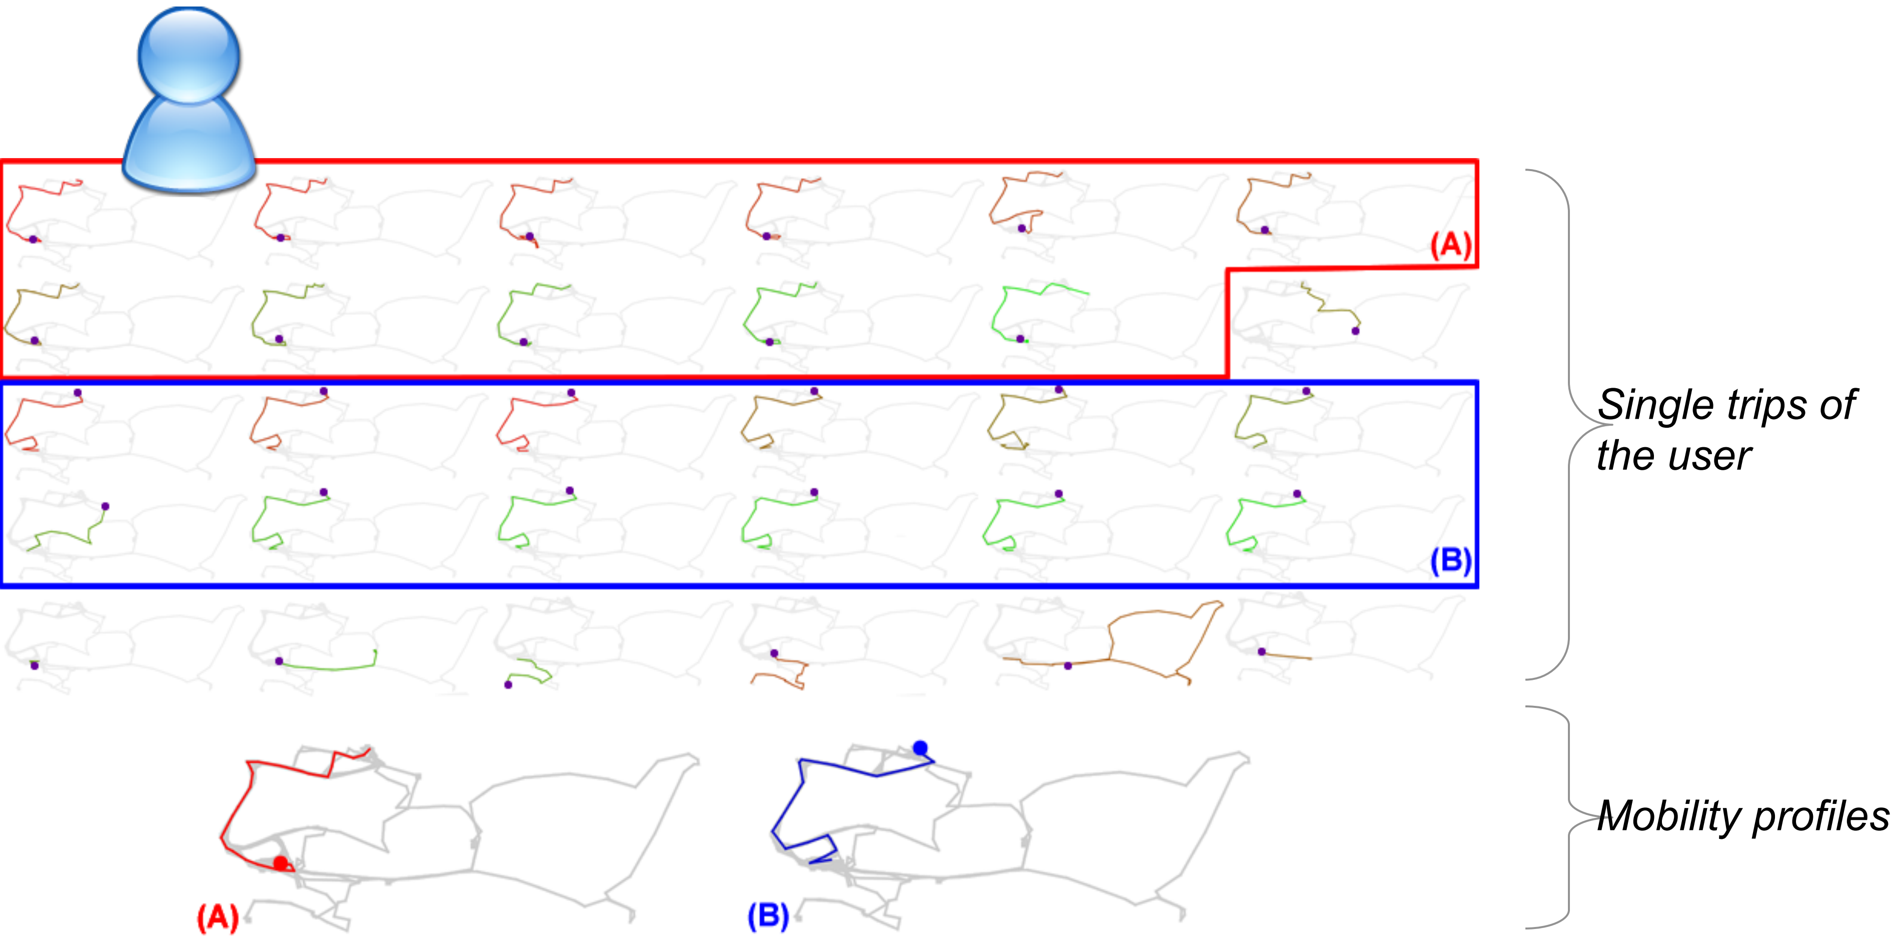

Discovering individual systematic movements

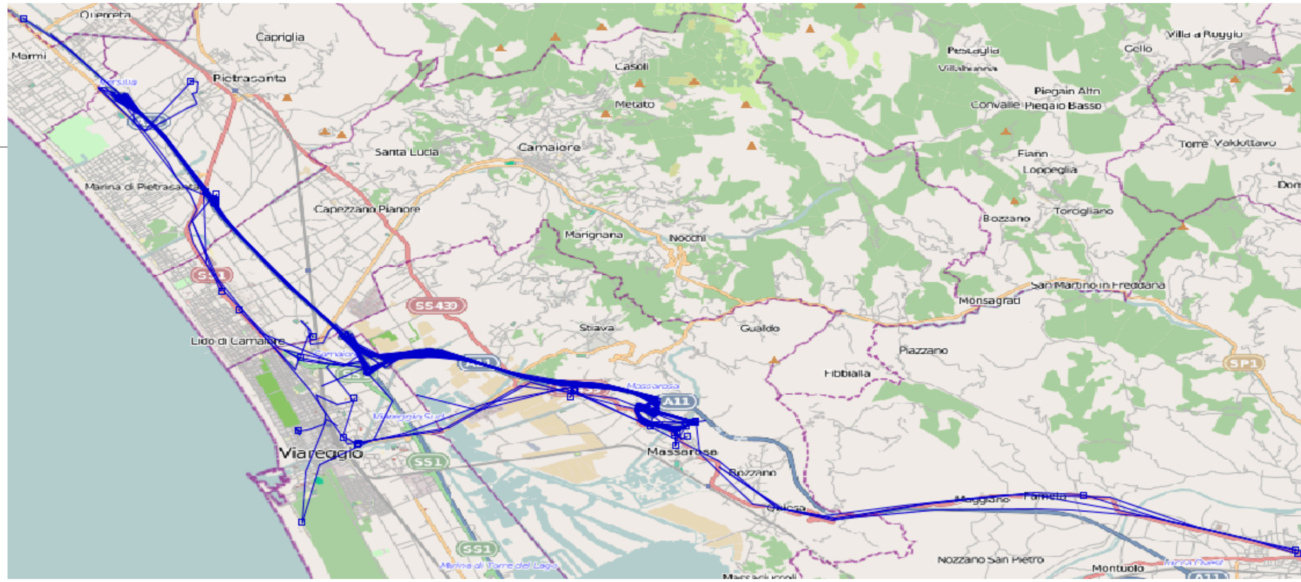

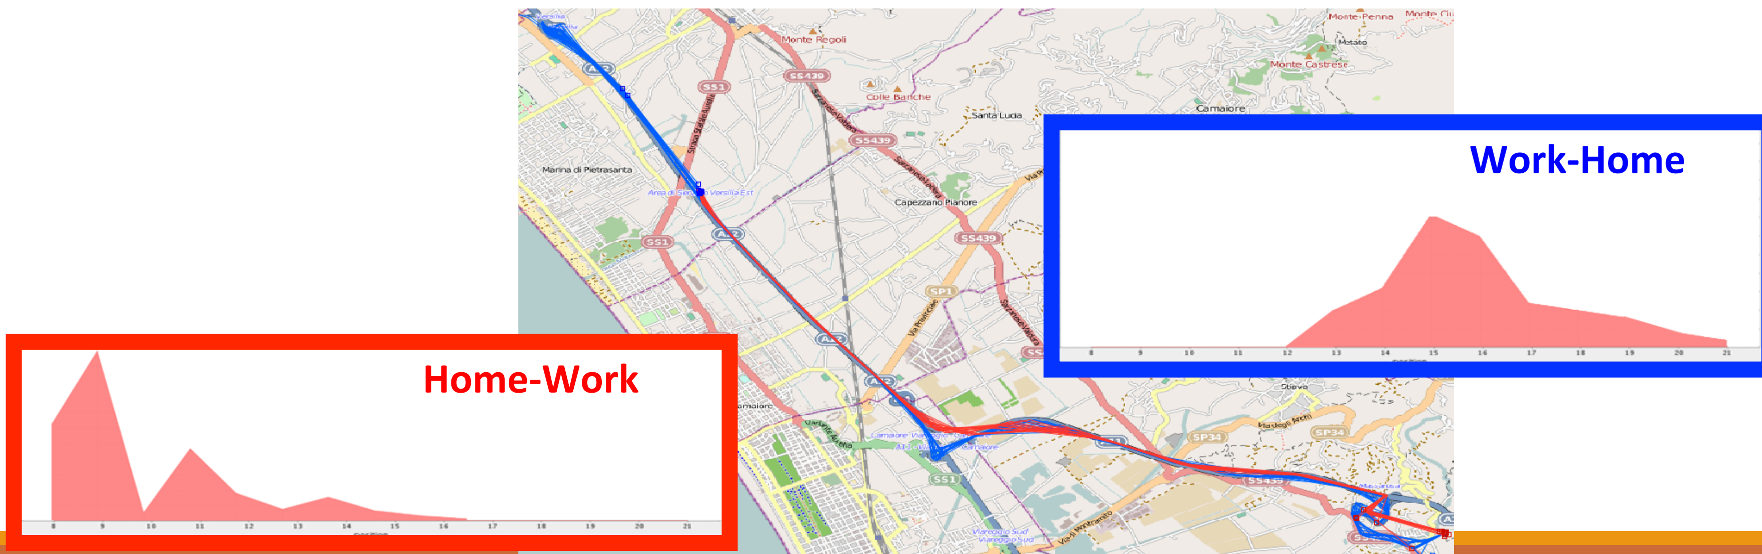

Derived patterns and models: mobility profiles

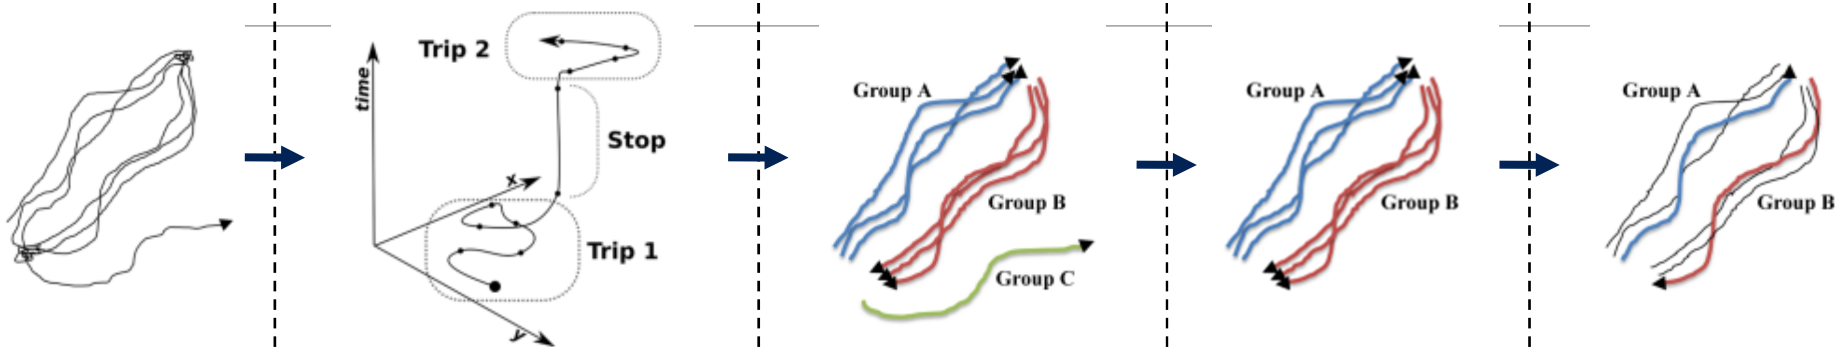

history Trips construction Grouping Pruning Profile extraction

of spatio-temporal points Cutting the user history when a stop is detected Performing a density based clustering equipped with a **spatio temporal** distance function The medoid of each group becomes user's **routines** and the all set become the u

Groups with a small
Number of trips are
Pruned

Spatial Tollerance
Temporal Tollerance
Spatio temporal distance

Stops Spatial Threshold
Stops Temporal Threshold

Support Threshold

What kind of distance?

- Start + End
  - Look for origin-destination pairs
- Route similarity
  - Look for recurrent paths followed
- Temporal dimension
  - Include time (of the day) to distinguish temporal regularity

What kind of representative?

- Classical average centroid cannot be applied
  - What is the centroid trajectory? Could make no sense
- Two practical solutions
  - Medoid: most central element of the cluster, e.g. minimized the sum of distances
  - Random: good enough if the clustering parameters are tight

Map Matching

Objective

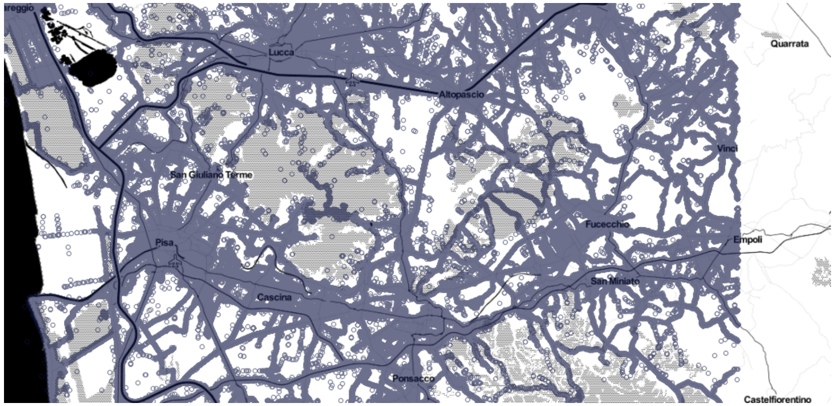

How to transform this…

Gps raw trajectories
Avg sampling rate 90 seconds
Affected by GPS positioning error

…into this?

Sequences of road segments crossed

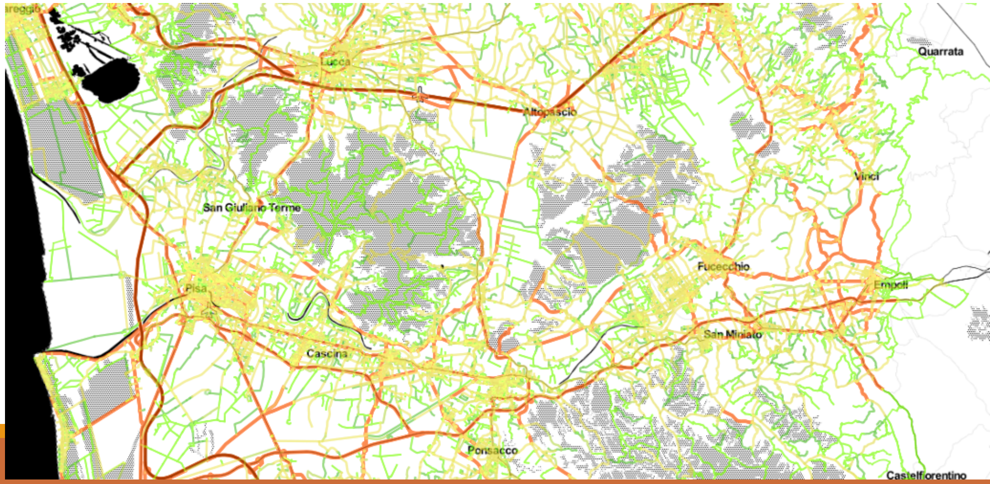

Objective

- Associate a sorted list of user positions to the road network on a digital map
- Two kinds of problems to solve
  - Map points to streets
  - Reconstruct path between points

Point mapping

- Determine which road segment a point belongs to
- Choose position within the segment

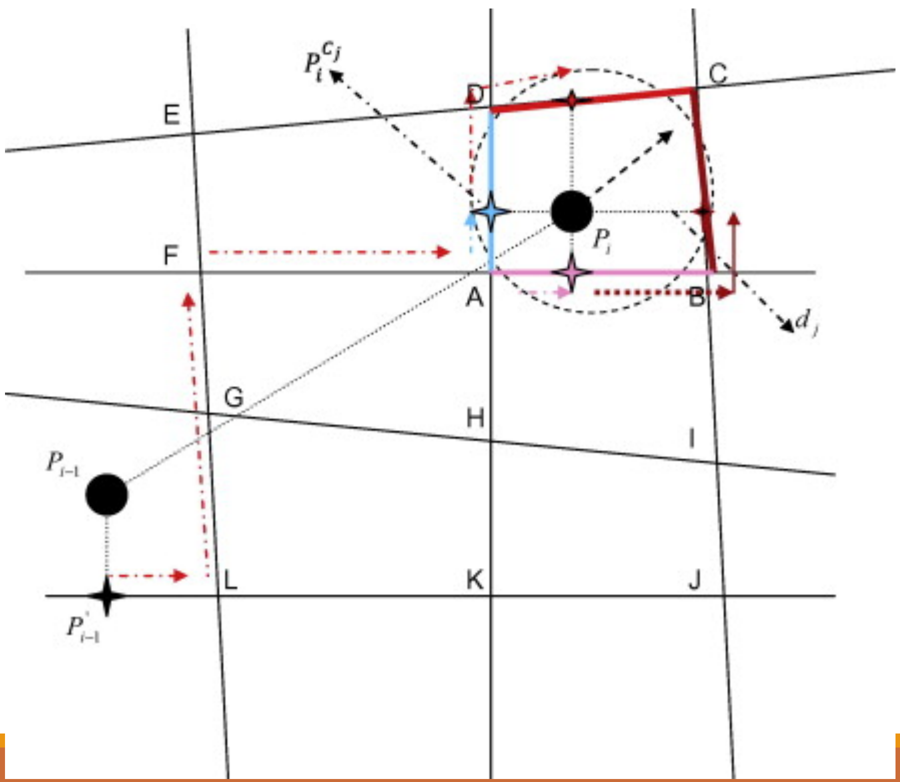

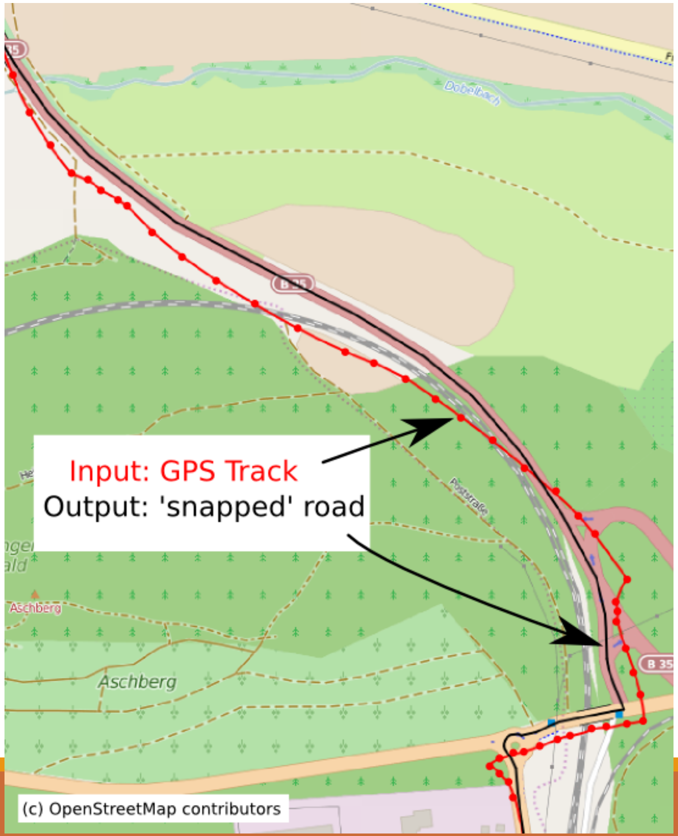

Objective

- Path reconstruction
  - Needed when gap between points is large

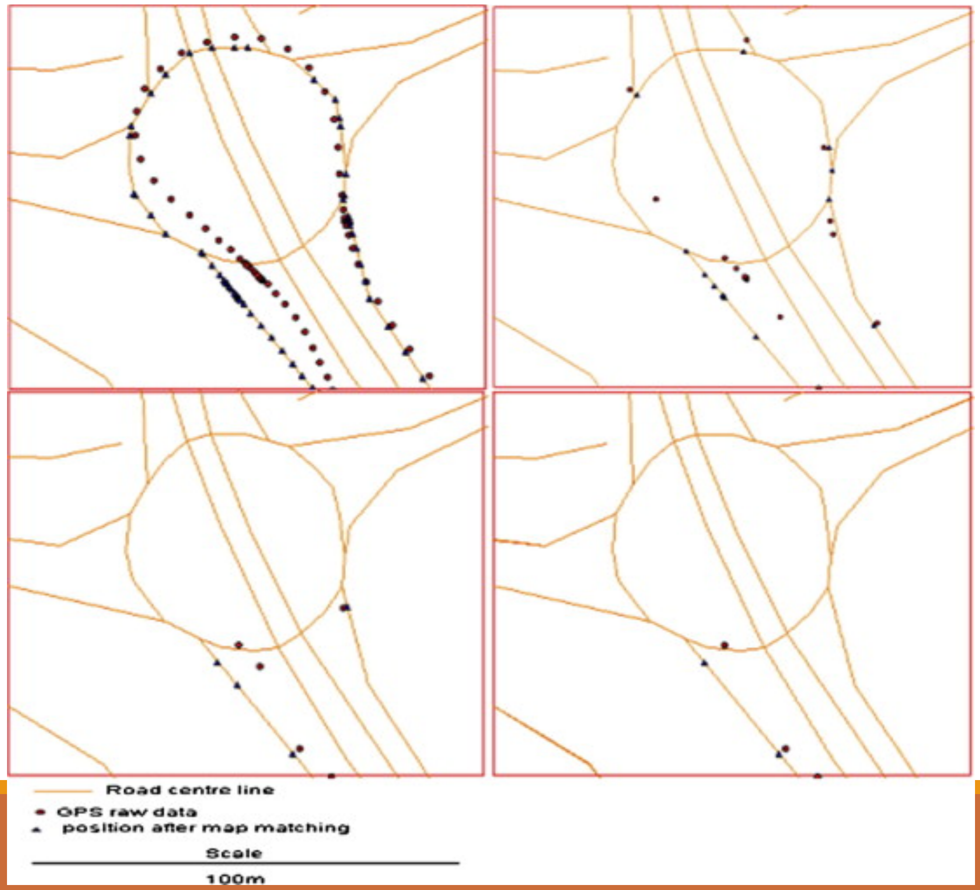

State of the art

- Map matching algorithms rely on shortest path between GPS points

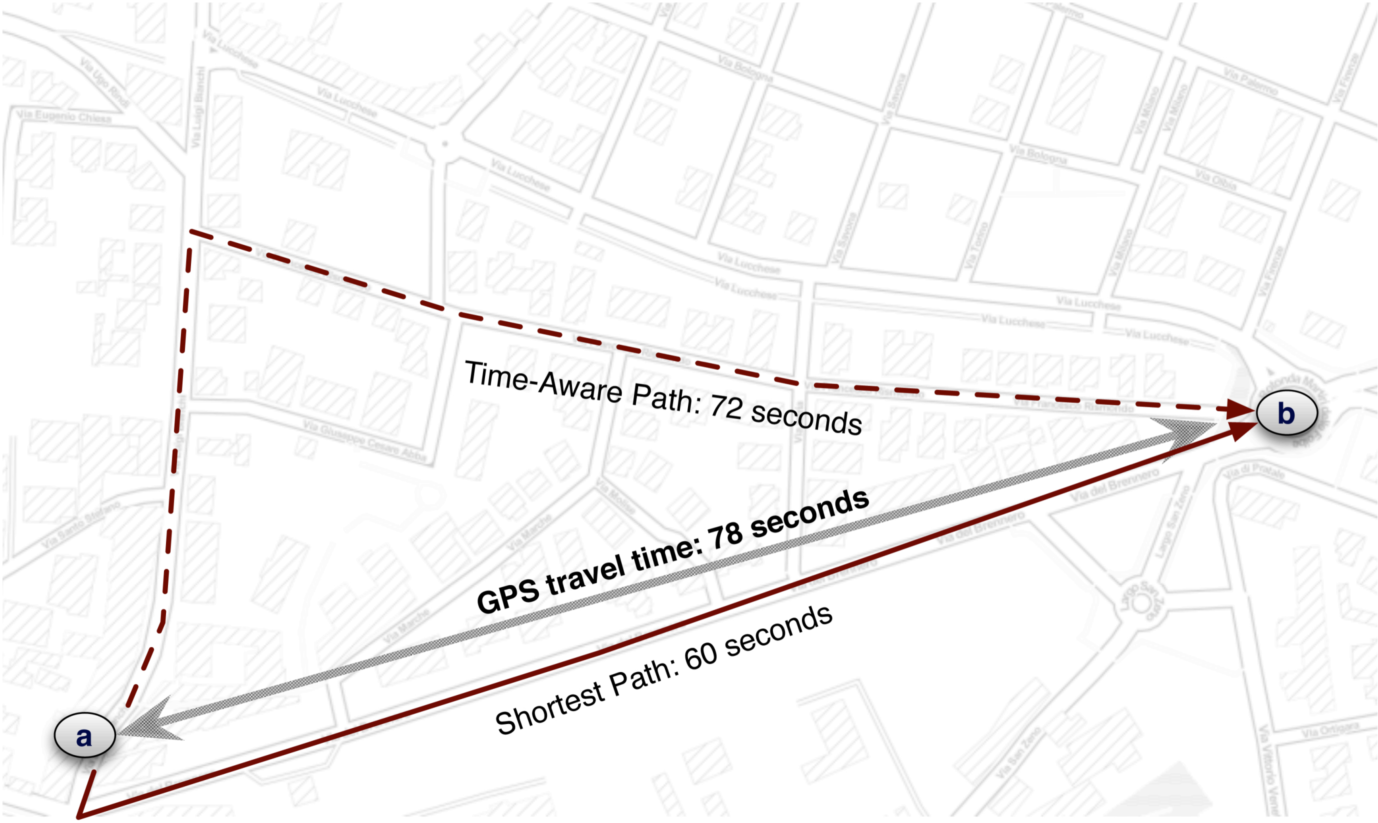

Result

- Trajectory → sequence of road segments

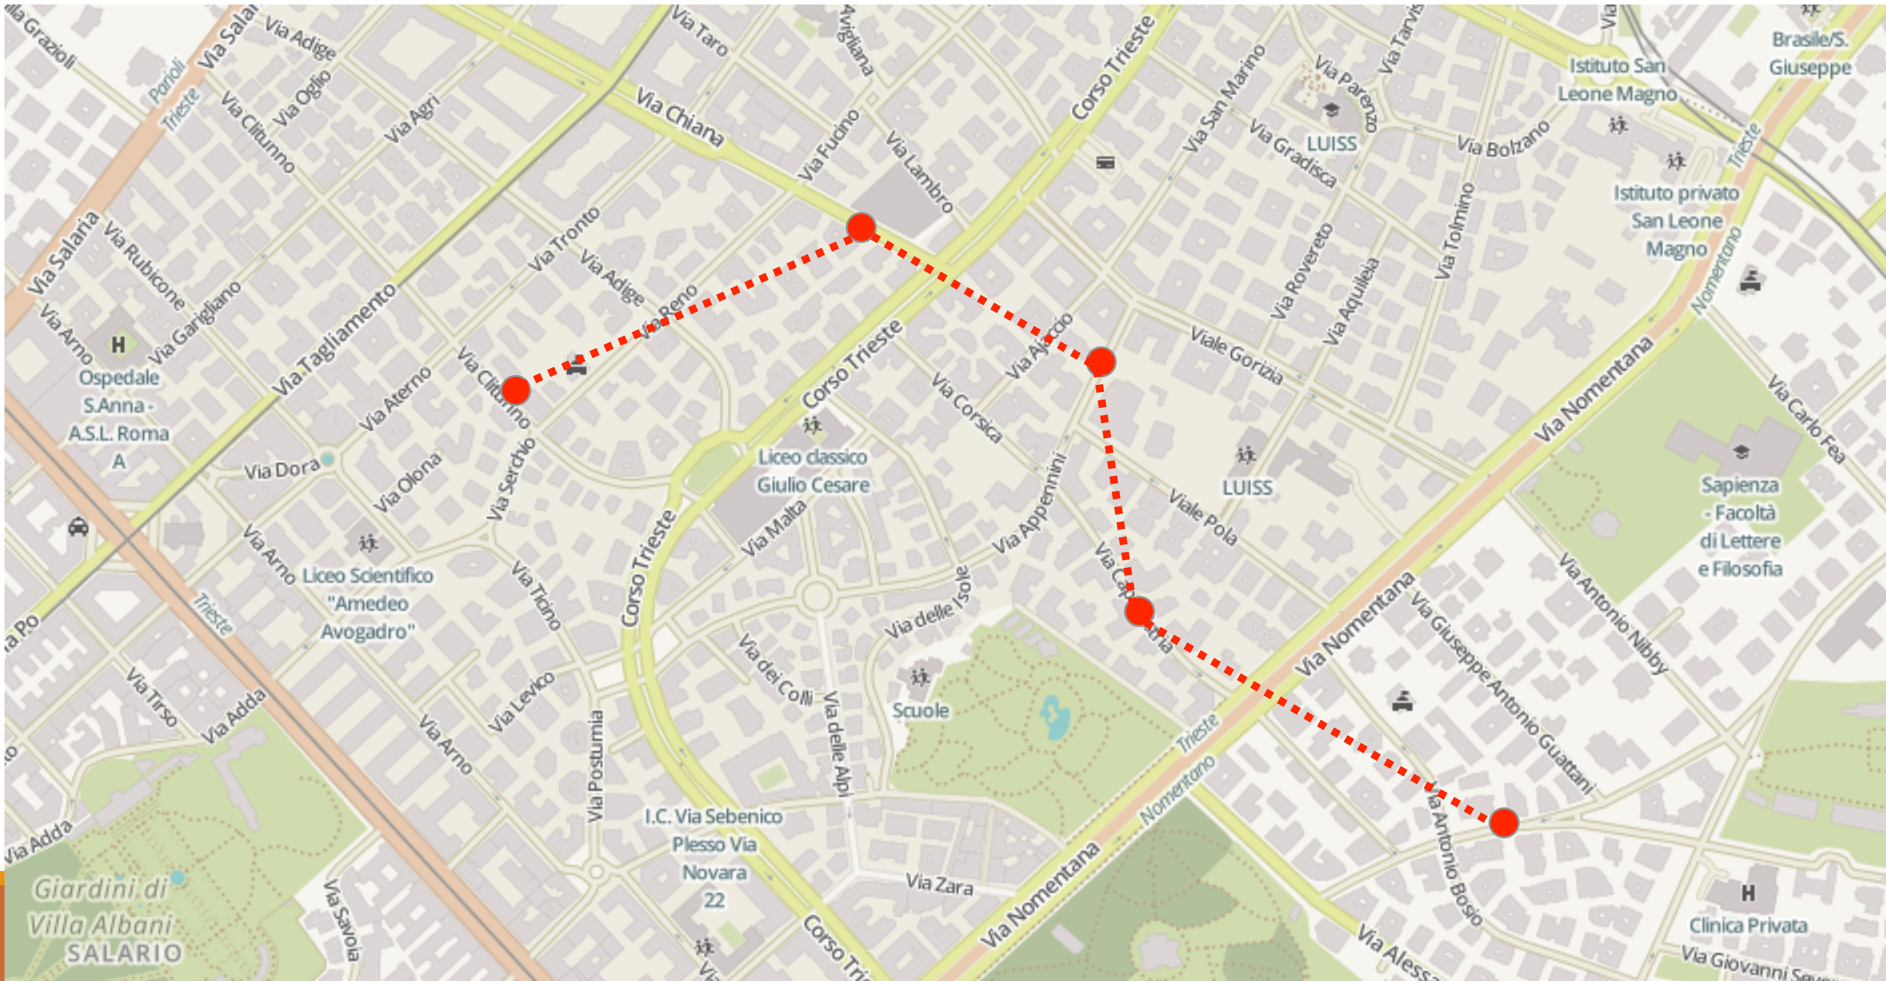

Result

- Trajectory → sequence of road segments

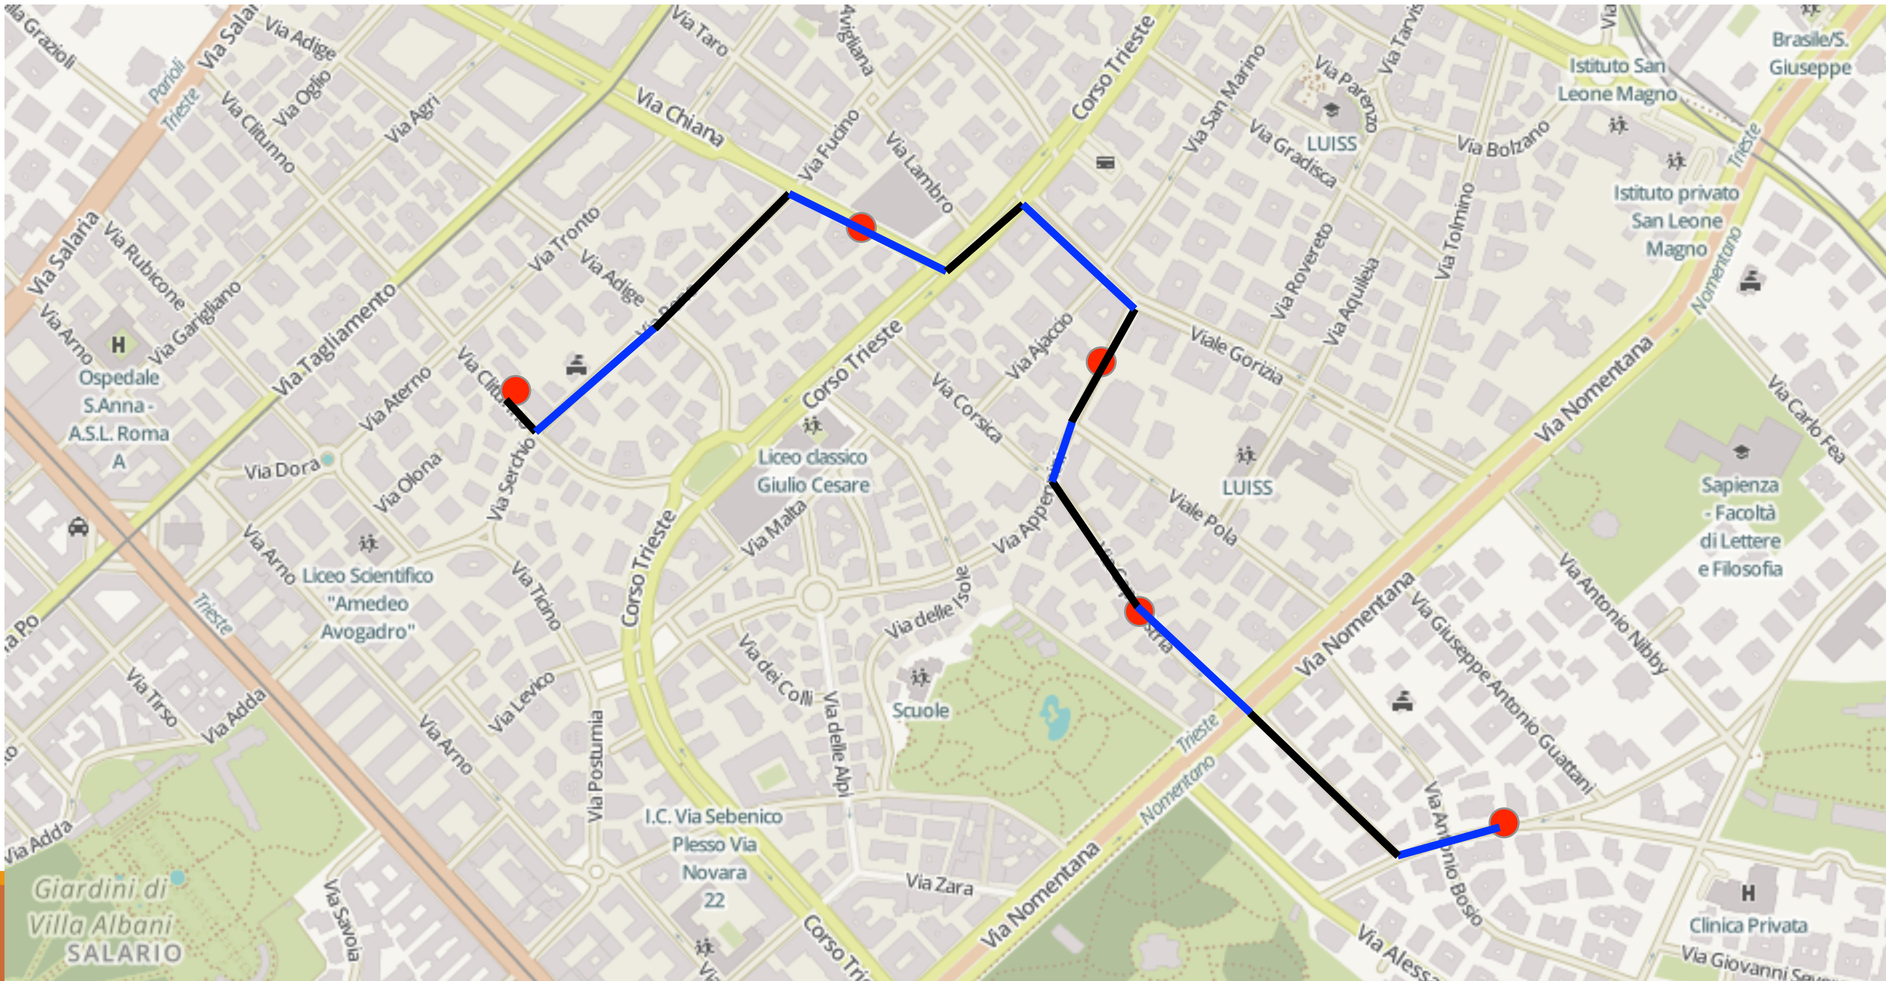

Shortest path vs real GPS time

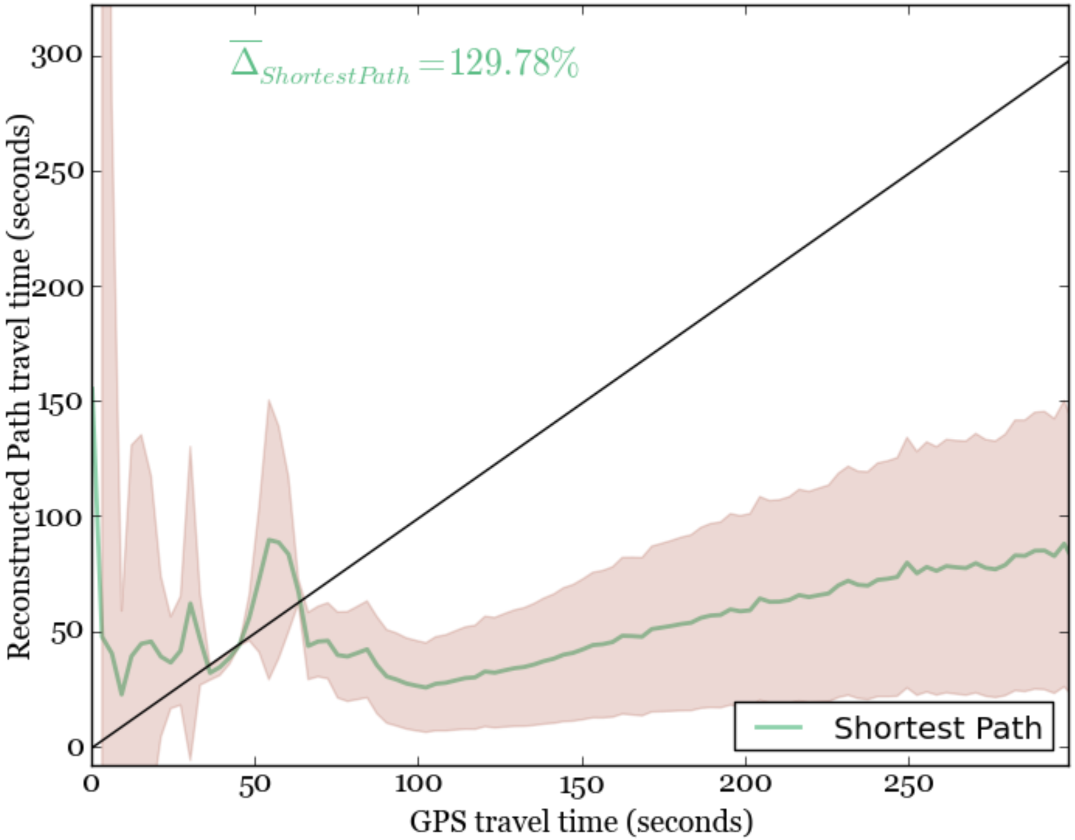

Matching GPS data with shortest path leads to significant differences w.r.t. real GPS travel time between two points

Alternative, Time-Aware approach

- **Given a road network with travel times for each edge, find the path that best fits given total travel time**
- **Satisfy some basic constraints, e.g. no useless turnarounds**

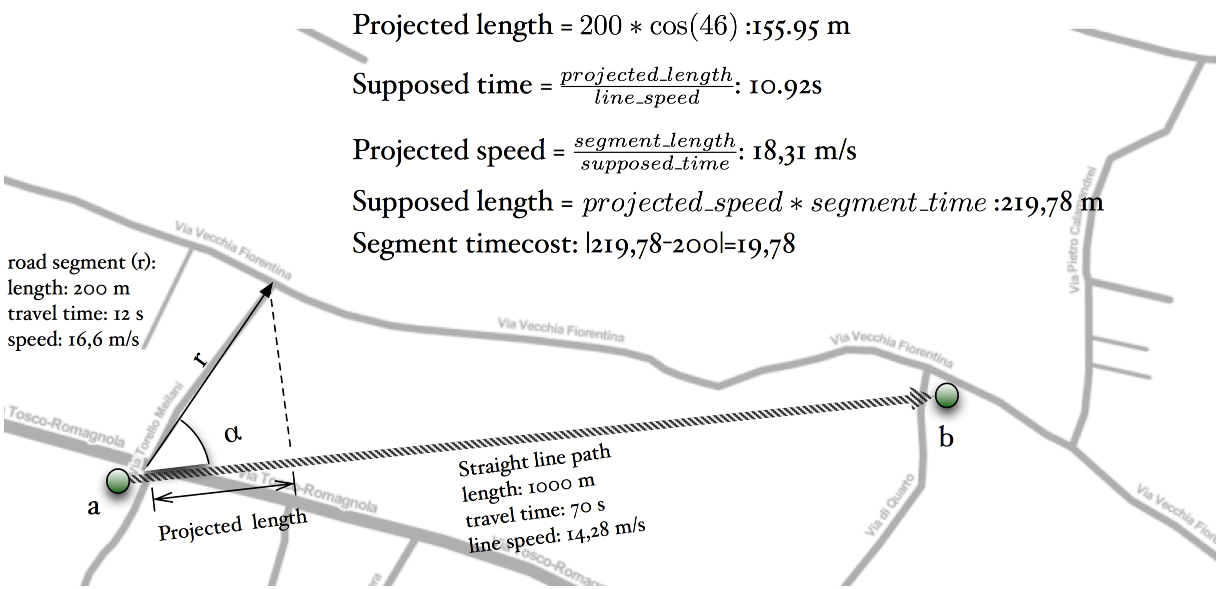

Workflow

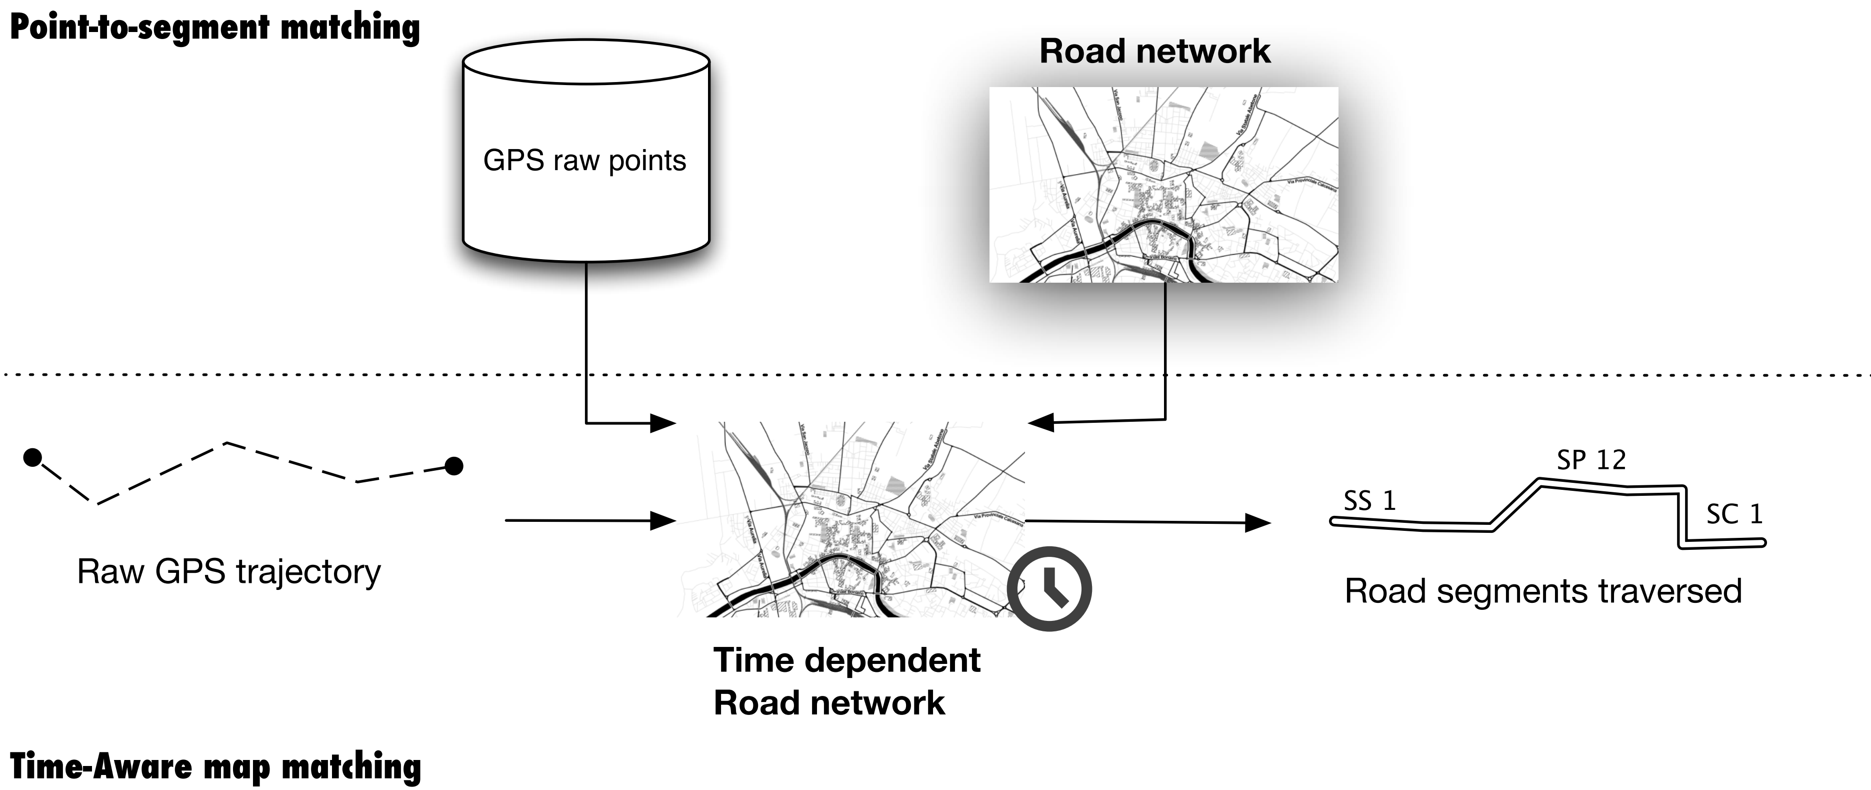

Point-to-segment matching
GPS raw points
Road network
Raw GPS trajectory
Time dependent
Road network
SS 1
SP 12
SC 1
Road segments traversed
Time-Aware map matching

Finding the travel time

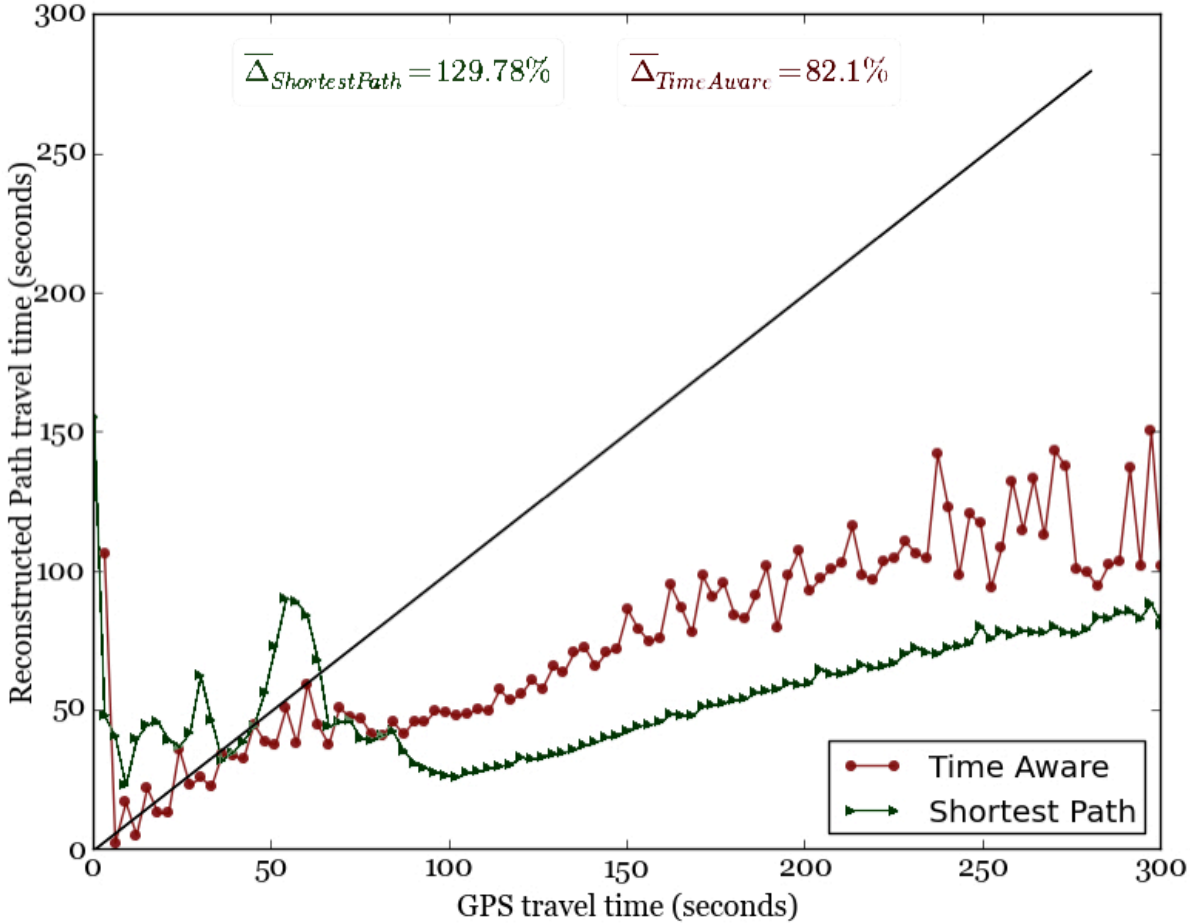

Effectiveness

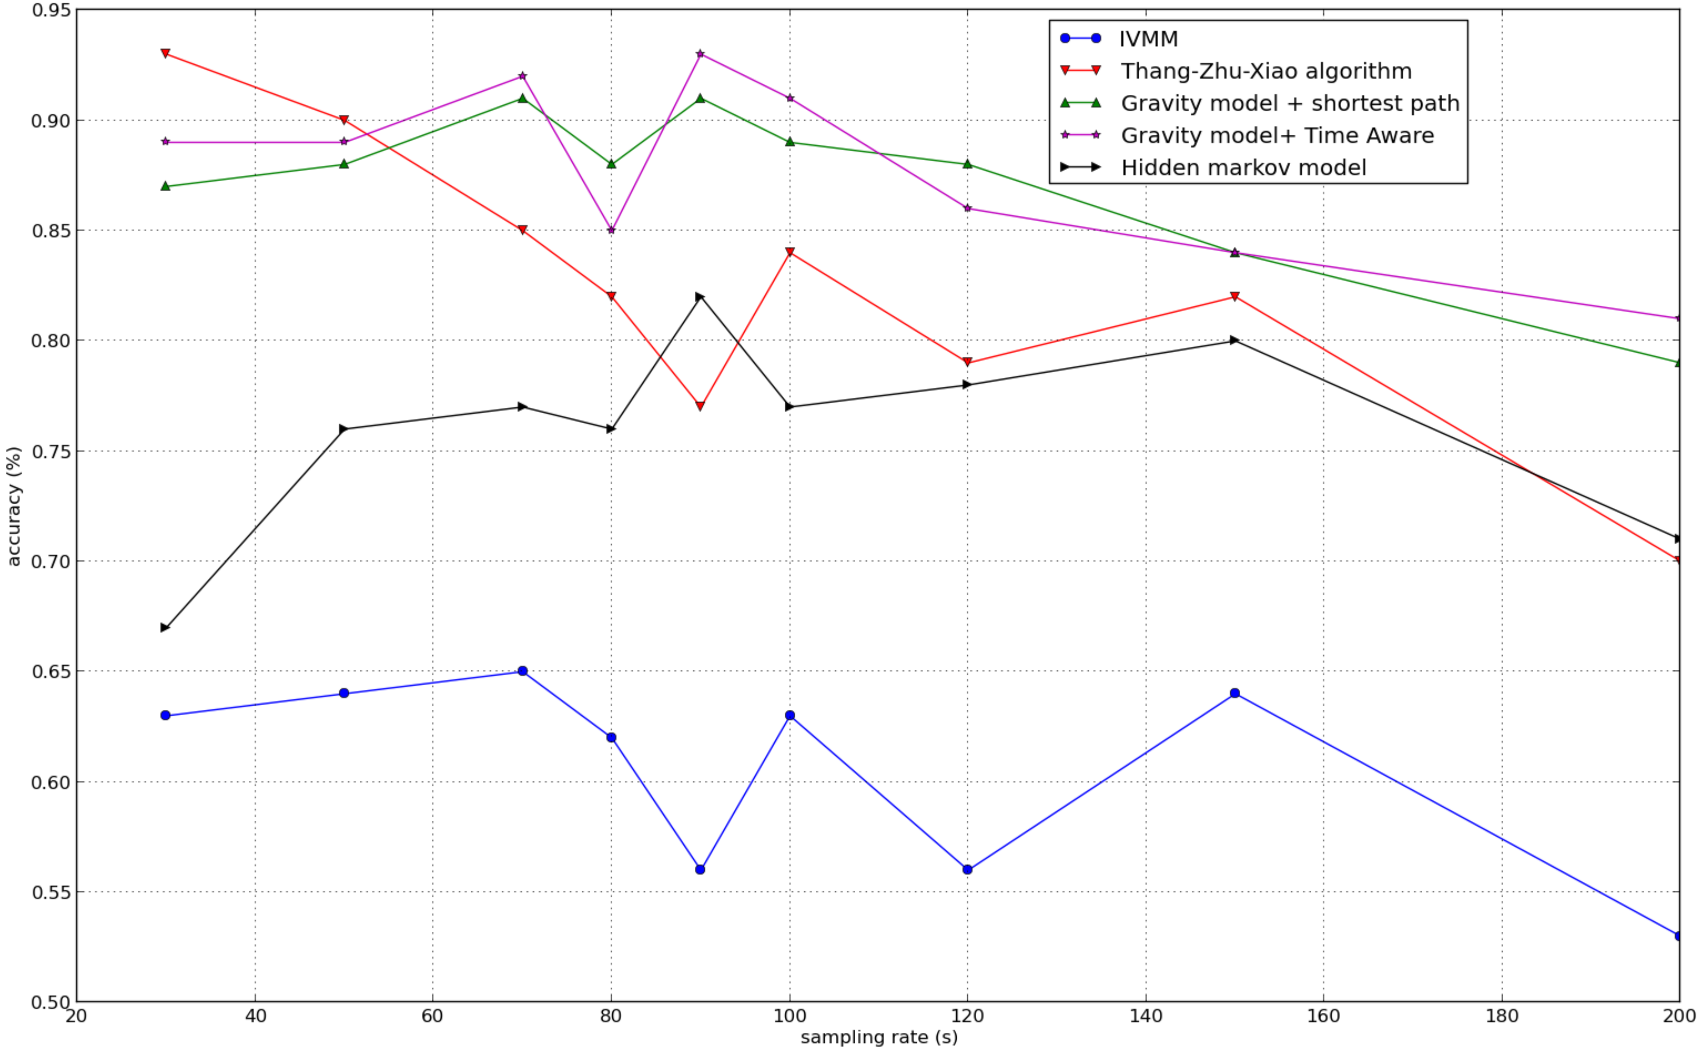

Activity Recognition

Objective / 1

- Infer the mobility mode
- Basic ML approach
  - Extract features of single points / segments of trajectory
  - Learn to classify by exploiting known examples

Learning Transportation Modes Based on GPS Trajectories

- Goal & Results: Inferring transportation modes from raw GPS data
  - Differentiate driving, riding a bike, taking a bus and walking
  - Achieve a 0.75 inference accuracy (independent of other sensor data)

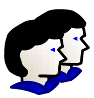

Users

GPS log

Infer model

Learning Transportation Modes Based on GPS Trajectories

- Motivation
  - For users:
    - Reflect on past events and understand their own life pattern
    - Obtain more reference knowledge from others' experiences
  - For service provider:
    - Classify trajectories of different transportation modes
    - Enable smart-route design and recommendation
- Difficulty
  - Velocity-based method cannot handle this problem well (<0.5 accuracy)
  - People usually transfer their transportation modes in a trip
  - The observation of a mode is vulnerable to traffic condition and weather

Yu Zheng, et al. Understanding Mobility Based on GPS Data. **UbiComp 2008**

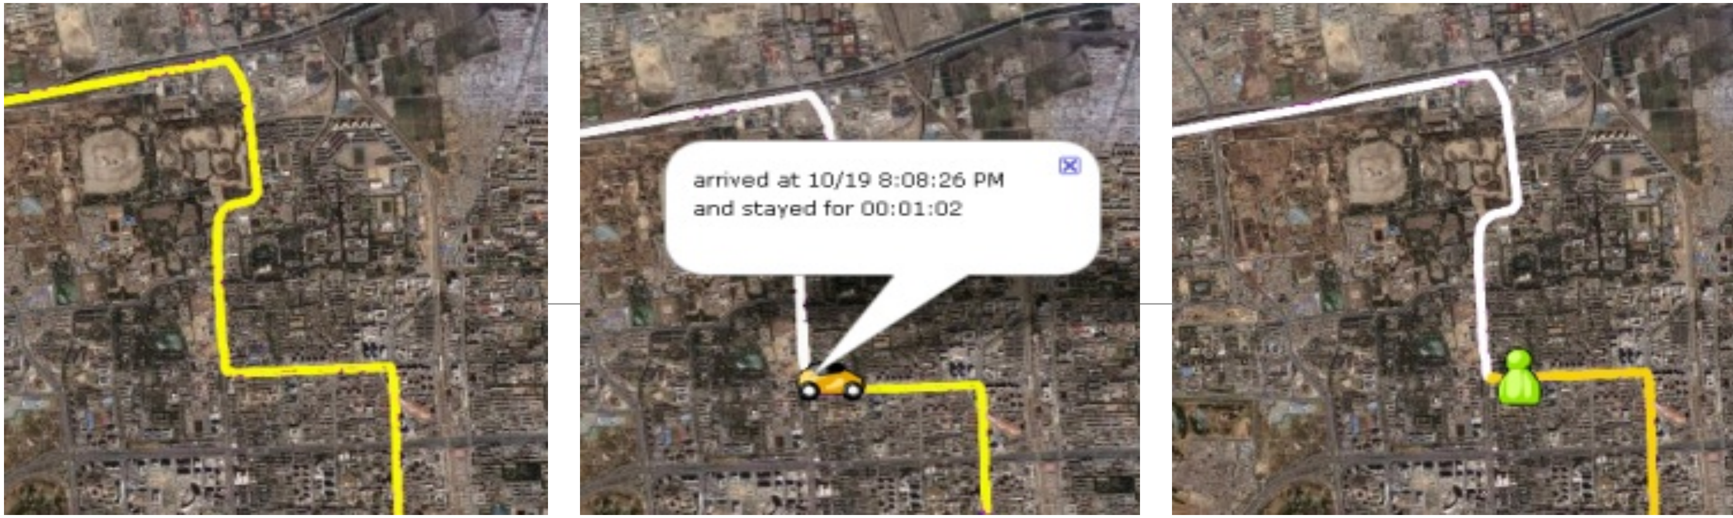
arrived at 10/19 8:08:26 PM
and stayed for 00:01:02

Learning Transportation Modes Based on GPS Trajectories

- Contributions and insights
  - A change point-based segmentation method
    - Walk is a transition between different transportation modes
    - Handle congestions to some extent
  - A set of sophisticated features
    - Robust to traffic condition
    - Feed into a supervise learning-based inference model
  - A graph-based post-processing
    - Considering typical user behavior
    - Employing location constrains of the real world
- WWW 2008 (first version)

Yu Zheng, et al. Understanding Mobility Based on GPS Data. **UbiComp 2008**

Architecture

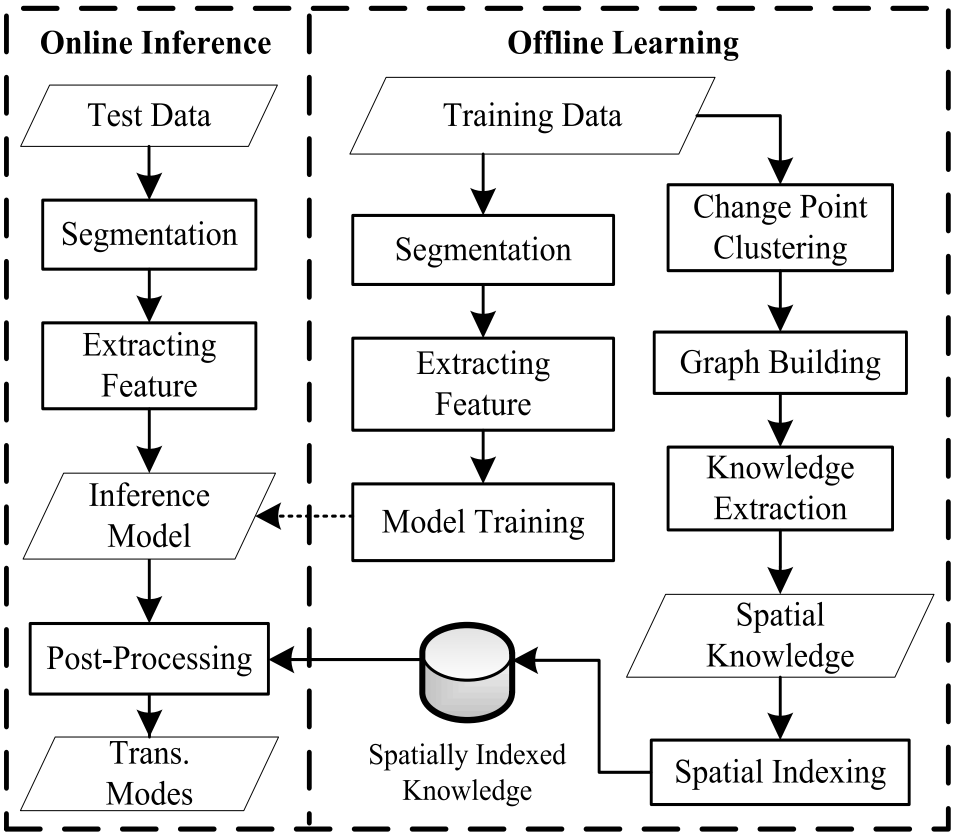

Walk-Based Segmentation

- Commonsense knowledge from the real world
  - Typically, people need to walk before transferring transportation modes
  - Typically, people need to stop and then go when transferring modes

Table I. Transition matrix among transportation modes

Transportation modes	Walk	Driving	Bus	Bike
Walk	/	41.1%	49.0%	9.0%
Driving	99.7%	/	0%	0.3%
Bus	98.7%	0.6%	/	0.6%
Bike	99.8%	0%	0.2%	/

Yu Zheng, et al. Understanding Mobility Based on GPS Data. **UbiComp 2008**

Walk-Based Segmentation

- Change point-based Segmentation Algorithm
  - ***Step 1***: distinguish all possible *Walk Points, non-Walk Points*.
  - ***Step 2***: merge short segment composed by consecutive *Walk Points* or *non-Walk points*
  - ***Step 3***: merge consecutive *Uncertain Segment* to *non-Walk Segment.*
  - ***Step 4***: end point of each *Walk Segment* are potential change points

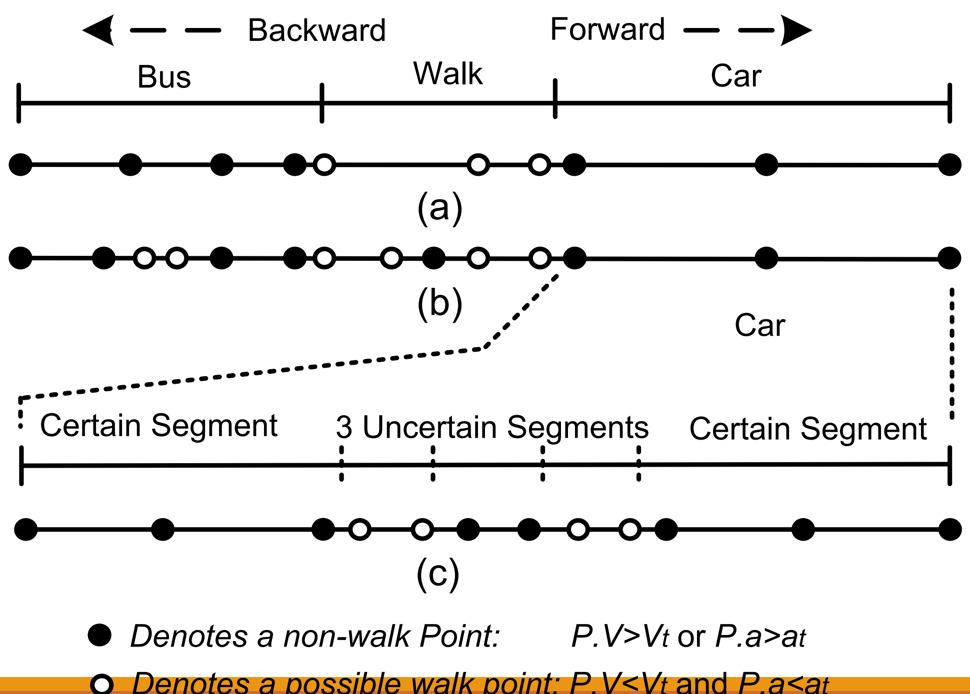

Feature Extraction (1)

- Features

Category	Features	Significance
Basic Features	Dist	Distance of a segment
	MaxV*i*	The *i*th maximal velocity of a segment
	MaxA*i*	The *i*th maximal acceleration of a segment
	AV	Average velocity of a segment
	EV	Expectation of velocity of GPS points in a segment
	DV	Variance of velocity of GPS points in a segment
Advanced Features	HCR	Heading Change Rate
	SR	Stop Rate
	VCR	Velocity Change Rate

Feature Extraction (2)

- Our features are more discriminative than velocity
  - Heading Change Rate (HCR)
  - Stop Rate (SR)
  - Velocity change rate (VCR)
  - >65 accuracy

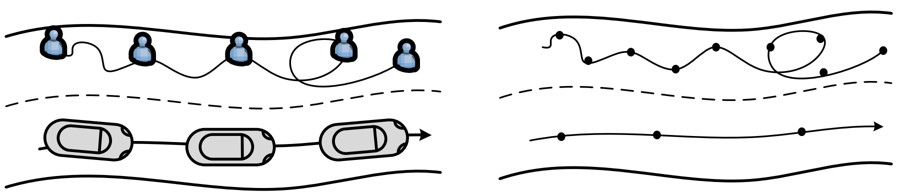

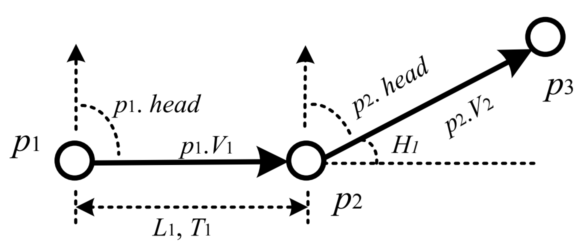

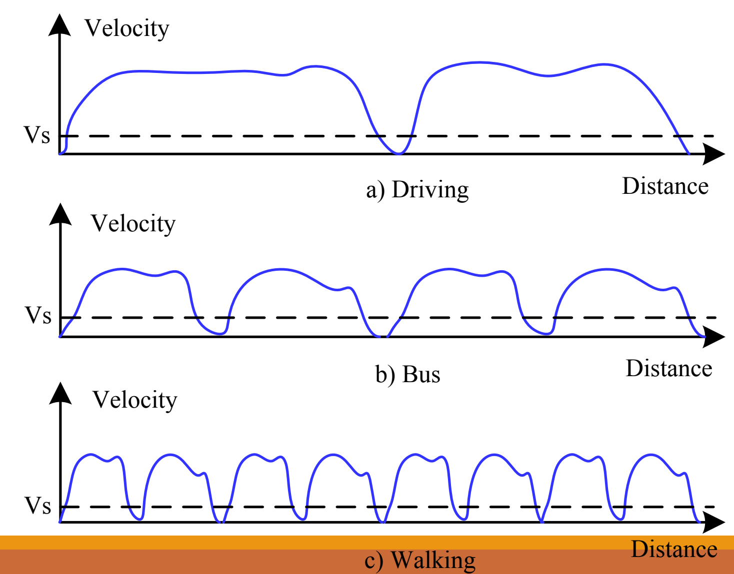

Objective

- Infer the purpose and/or activity performed of trips and locations
- Two approaches
  - Consider what kind of activities can be performed in that area
  - Consider how the user behaves (when and how he reaches the are, etc.)

Recognition through Points-of-Interest

Given a dataset of GPS tracks of private vehicles, annotate trajectories with the most probable activities performed by the user.

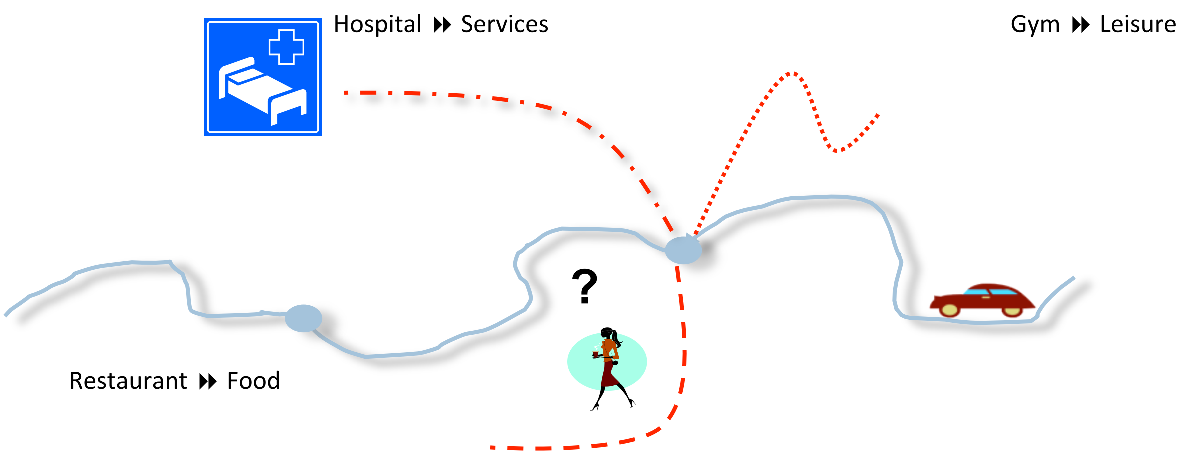

Associates the list of possible POIs (with corresponding probabilities) visited by a user moving by car when he stops.

A mapping between POIs categories and Transportation Engineering activities is necessary.

The enrichment process

- **POI collection**: Collected in an automatic way, e.g. from Google Places.
- **Association POI – Activity**: Each POI is associated to a ``activity''. For example Restaurant → Eating/Food, Library → Education, etc.
- **Basic elements/characteristics:**
  - C(POI) = {category, opening hour, location}
  - C(Trajectory) = {stop duration, stop location, time of the day}
  - C(User) = {max walking distance}
- **Computation of the probability to visit a POI/ to make an activity:** For each POI, the probability of ``being visited'' is a function of the POI, the trajectory and the user features.
- **Annotated trajectory**: The list of possible activities is then associated to a Stop based on the corresponding probability of visiting POIs

Input & Output

Lat; Lon
TimeStamp: Sun 10:55 – 12:05

- ✗ Bank: Mon – Fri [8:00 – 15:30]
- ✗ Dentist: Mon – Sat [9:00 – 13:00] [15:30 – 18:00]
- ✓ Church: Mon – Sat [18:00 – 19:00]
  Sun [11:00 – 12:00]
- ✗ Primary School: Mon – Sat [8:00 – 13:00]

Input & Output

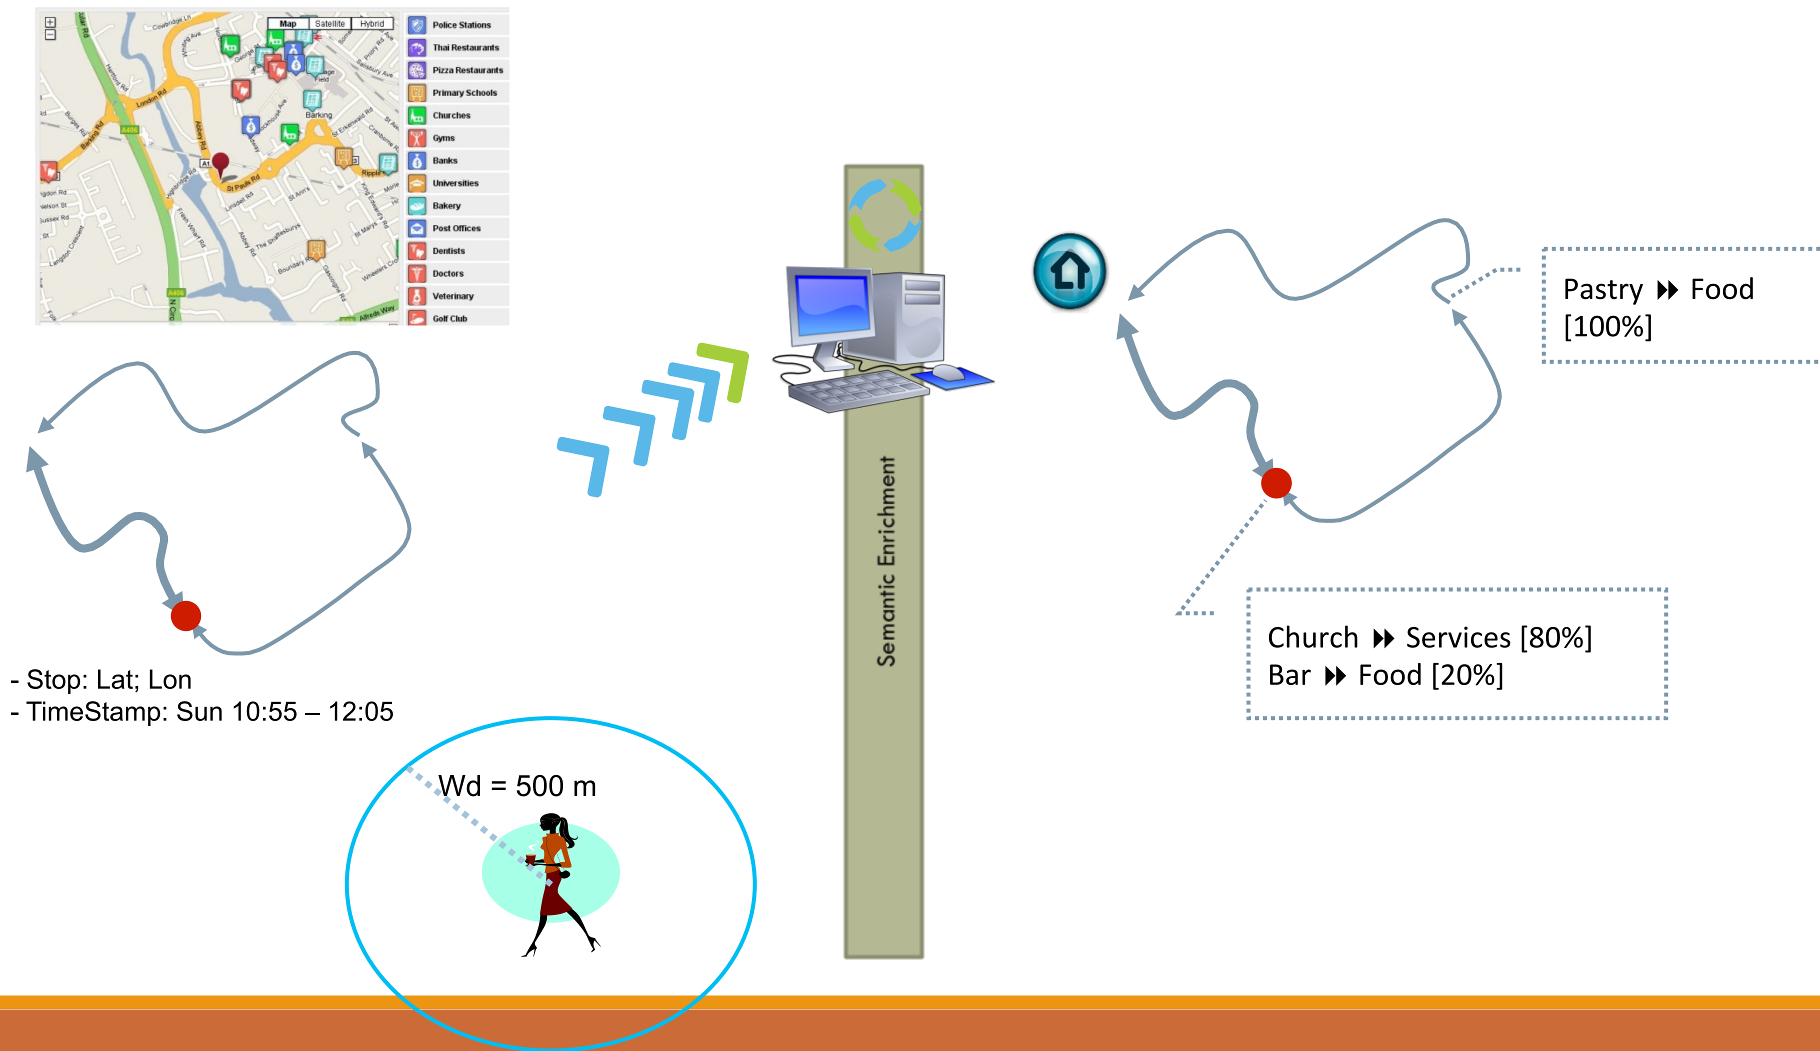

Extraction of personal places from Twitter trajectories in Dublin area

The points of each trajectory taken separately were grouped into spatial clusters of maximal radius 150m. For groups with at least 5 points, convex hulls have been built and spatial buffers of small width (5m) around them.
1,461,582 points belong to the clusters (89% of 1,637,346); 24,935 personal places have been extracted.

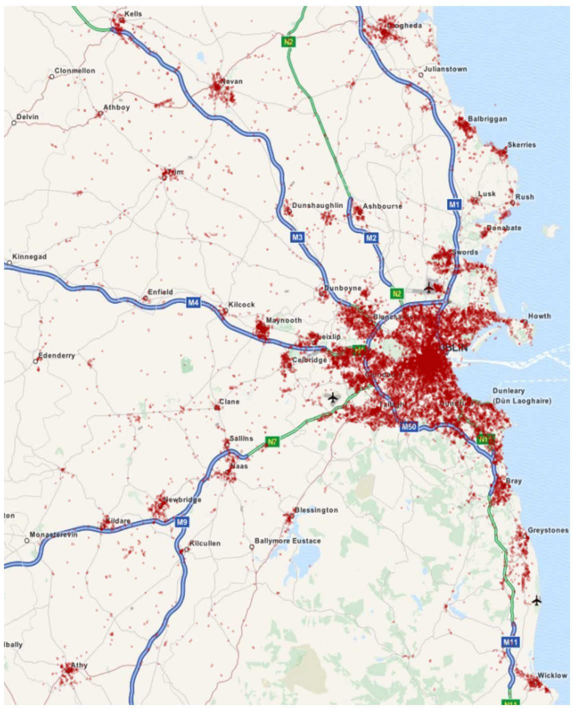

Examples of extracted places

N	N?	min	q1	med	q3	max	ave	stdd
5180	0	1.0	2.0	3.0	5.0	120.0	4.8	6.3

Statistical distribution of the number of places per person

Recognition of the message topics, generation of topical feature vectors, and summarization by the personal places

Topics have been assigned to 208,391 messages (14.3% of the 1,461,582 points belonging to the personal places)

Message	Features	topic=family: Occurrences of topic	topic=home: Occurrences of topic	topic=education: Occurrences of topic	topic=work: Occurrences of topic	t… O…
@joe_lennon I usually	education	0	0	1	0	
@joe_lennon together	education	0	0	1	0	
@jas_103 deadly; don	work	0	0	0	1	
Just got home and see	home	0	1	0	0	
So excited about my ne	sweets	0	0	0	0	
@Iamtcdizzy I haven't b	shopping	0	0	0	0	
Get in from my night ou	family;home;work	1	1	0	1	
Home again at 6pm! N	home	0	1	0		
Bussing it home for th			1	0		
Ah shite. It's been a p			0	0		
@ronanhutchinson be	education	0	0	1		

Get in from my night out; my dad gets home from work two minutes later. Great timing :)

…

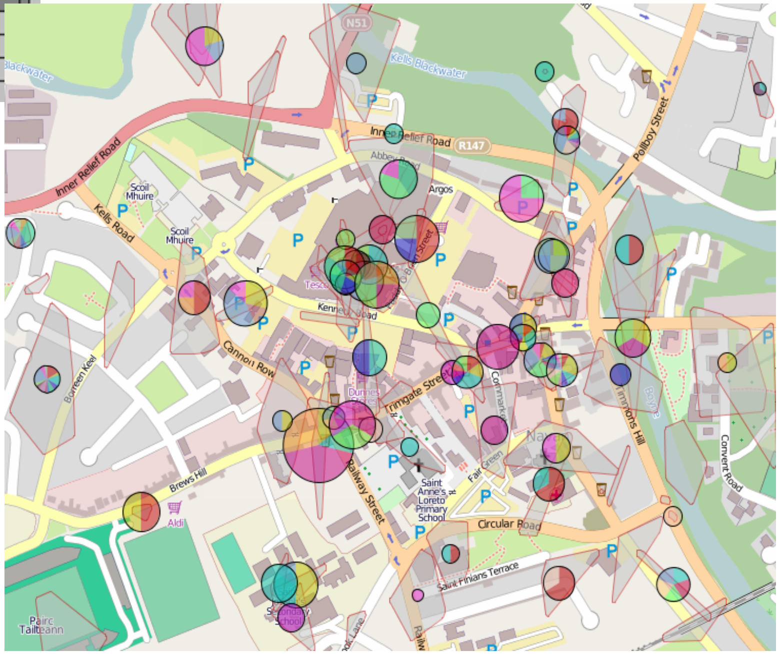

1) Some places did not get topic summaries (about 20% of the places)
2) In many places the topics are very much mixed
3) The topics are not necessarily representative of the place type (e.g., topics near a supermarket: family, education, work, cafe, shopping, services, health care, friends, game, private event, food, sweets, coffee)

Obtaining daily time series of place visits and comparison with exemplary temporal profiles

The daily time series of place visits have been obtained through aggregation of daily trajectories using only relevant places for each trajectory. The aggregation was done separately for the work days from Monday to Thursday, and for Saturday, Sunday, and Friday.

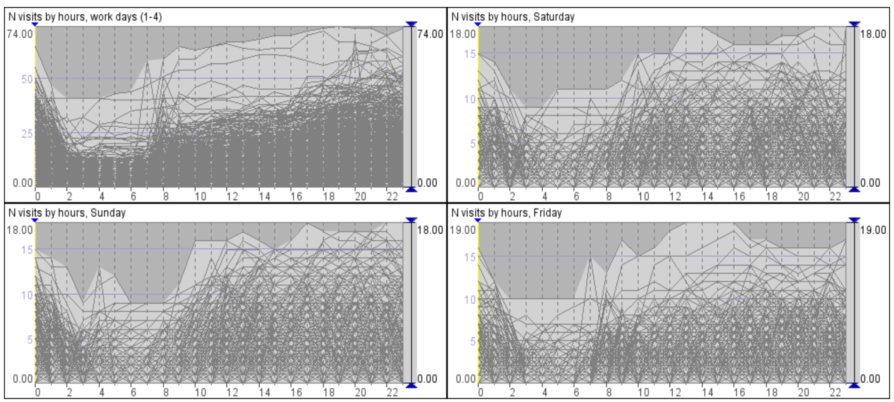

Exemplary temporal profiles of different activities
or visits to different types of places

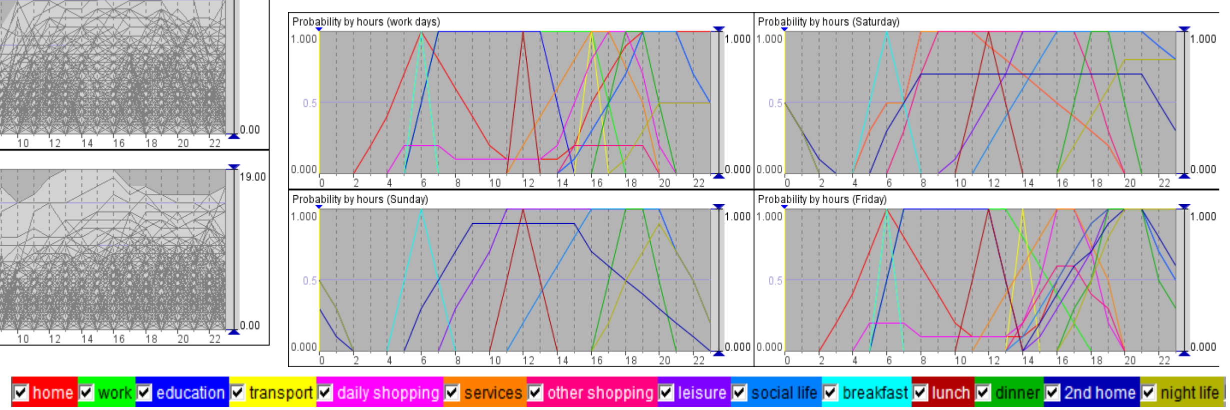

The time series of place visits are compared to the exemplary time profiles by means of the Dynamic Time Warping (DTW) distance function. Resulting scores: from 0 (no similarity) to 1 (very high similarity).
15,950 places (64% of all) have no similarity to any of the exemplary time patterns. 4,732 places (19%) have the maximal similarity score of 0.8 or higher; 4,179 of them (16.8% of all) were visited in 6 or more days.

Time series with high similarity to "work" (>=0.8)

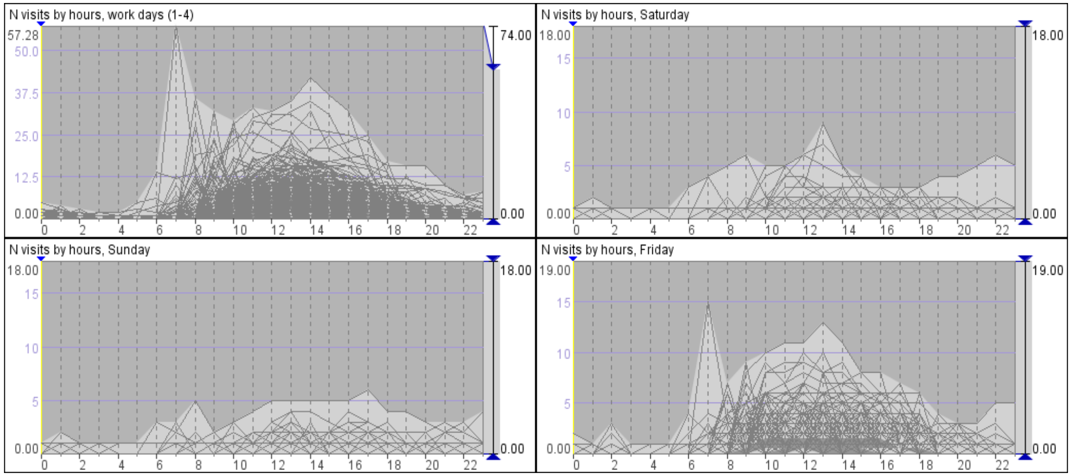

The time series similarity scores have been combined with the relative frequencies of the topics using a combination matrix

	home	work	education	transport	daily shopping	services	other shopping	leisure	social life	breakfast	lunch	dinner	2nd home	night life
family	1	0	0	0	0	0	0	0	1	0	0	0	1	-0.5
home	1	0	0	0	0	0	0	0	0	0	0	0	1	-0.5
education	0	0	1	1	0	0	0	-1	-1	-1	-1	-1	0	-0.5
work	0	1	0	1	0	0	0	-1	-1	-1	-1	-1	0	-0.5
local transport	0	0	0	1	0	0	0	0	0	-1	-1	-1	0	-0.5
regional transport	0	0	0	1	0	0	0	0	0	-1	-1	-1	0	-0.5
far transport	-1	-1	-1	1	-1	-1	-1	0	0	-1	-1	-1	0	-0.5
eatery	-1	-1	-1	-1	-1	-1	-1	0	0	0	1	0.5	-1	0
restaurant	-1	-1	-1	-1	-1	-1	-1	0	0.5	-1	0.6	1	-1	0.5
pub	-1	-1	-1	-1	-1	-1	-1	0.5	0.5	-1	-1	0.7	-1	1
cafe	-1	-1	-1	-1	-1	-1	-1	0.5	0.5	0.3	0.3	0.3	-1	0.2
shopping	-1	-1	-1	-0.5	1	0	1	0	0	0	0	0	-1	-0.5
daily shopping	-1	-1	-1	-0.5	1	0	0	0	0	0	0	0	-1	-0.5
other shopping	-1	-1	-1	-0.5	0	0	1	0	0	0	0	0	-1	-0.5
services	-1	-1	-1	-0.5	0	1	0	0	0	-1	-1	-1	-1	-0.5
health care	-0.5	-0.5	-0.5	-0.5	-0.5	0	-0.5	-0.5	-0.5	-1	-1	-1	-0.5	-0.5
fitness	0	0	0	0	0	0	0	1	0	-1	-1	-1	0	-1
wellness	0	0	0	0	0	0	0	1	0	-1	-1	-1	0	-1
nature	-0.5	0	0	0	0	0	0	1	0	0	0	0	-0.5	-1
culture	-0.5	-1	-1	-1	-1	-1	0	1	0	-1	-1	-1	-0.5	-1
sports	0	-1	-1	-1	-1	-1	-1	1	0	-1	-1	-1	0	-1
friends	0	0	0.5	0	0	0	0	0.5	1	0	0	0	0	0.5
game	0	-1	-1	-1	-1	-1	-1	1	0	-1	-1	-1	0	0
public event	-1	-1	-1	-1	-1	-1	-1	1	0	0	0	0	-1	0
private event	0	-1	-1	-1	-1	-1	-1	0	1	-1	-1	0	0	0.5
alcohol	0	-0.5	-0.5	-0.5	-0.5	-0.5	0	0.5	1	-0.5	0	0.5	0	1
food	0.5	-0.5	0	0	1	0	0	0	0.5	1	1	1	0.5	0.5
sweets	0	0	0	0	0	0	0	0.5	0.5	0.5	0	1	0	0.5
coffee	0	0	0	0	0	0	0	0	0.5	1	1	0	0	0
tea	0	0	0	0	0	0	0	0	0	0.5	0.5	0	0	0

1,520 places (6.1% of all). These places have also high similarity to "education", "transport", and "lunch".

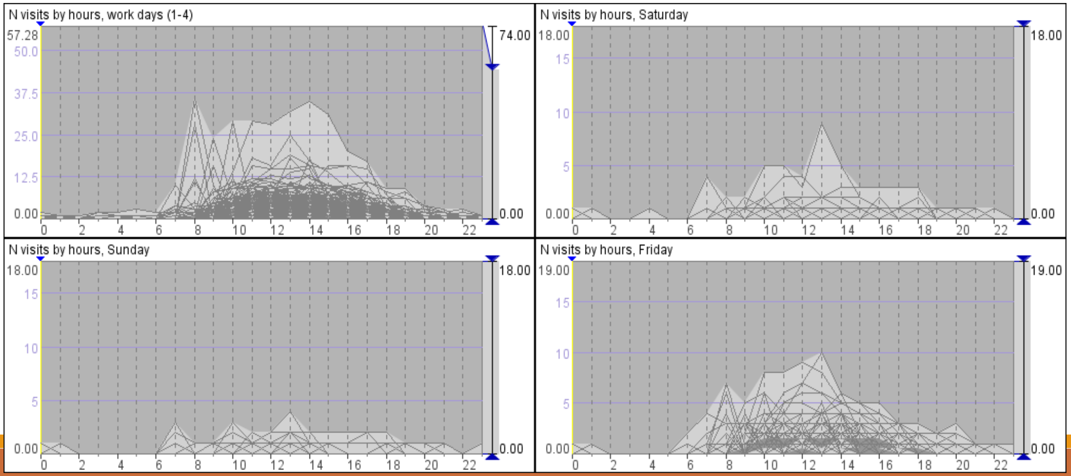

In 233 places out of the initial 1,520 (15%, 0.9% of all places) the similarity to the "work" profile has been reinforced based on the topic frequencies.

Classification of the places according to the highest combined score (minimum 0.8)

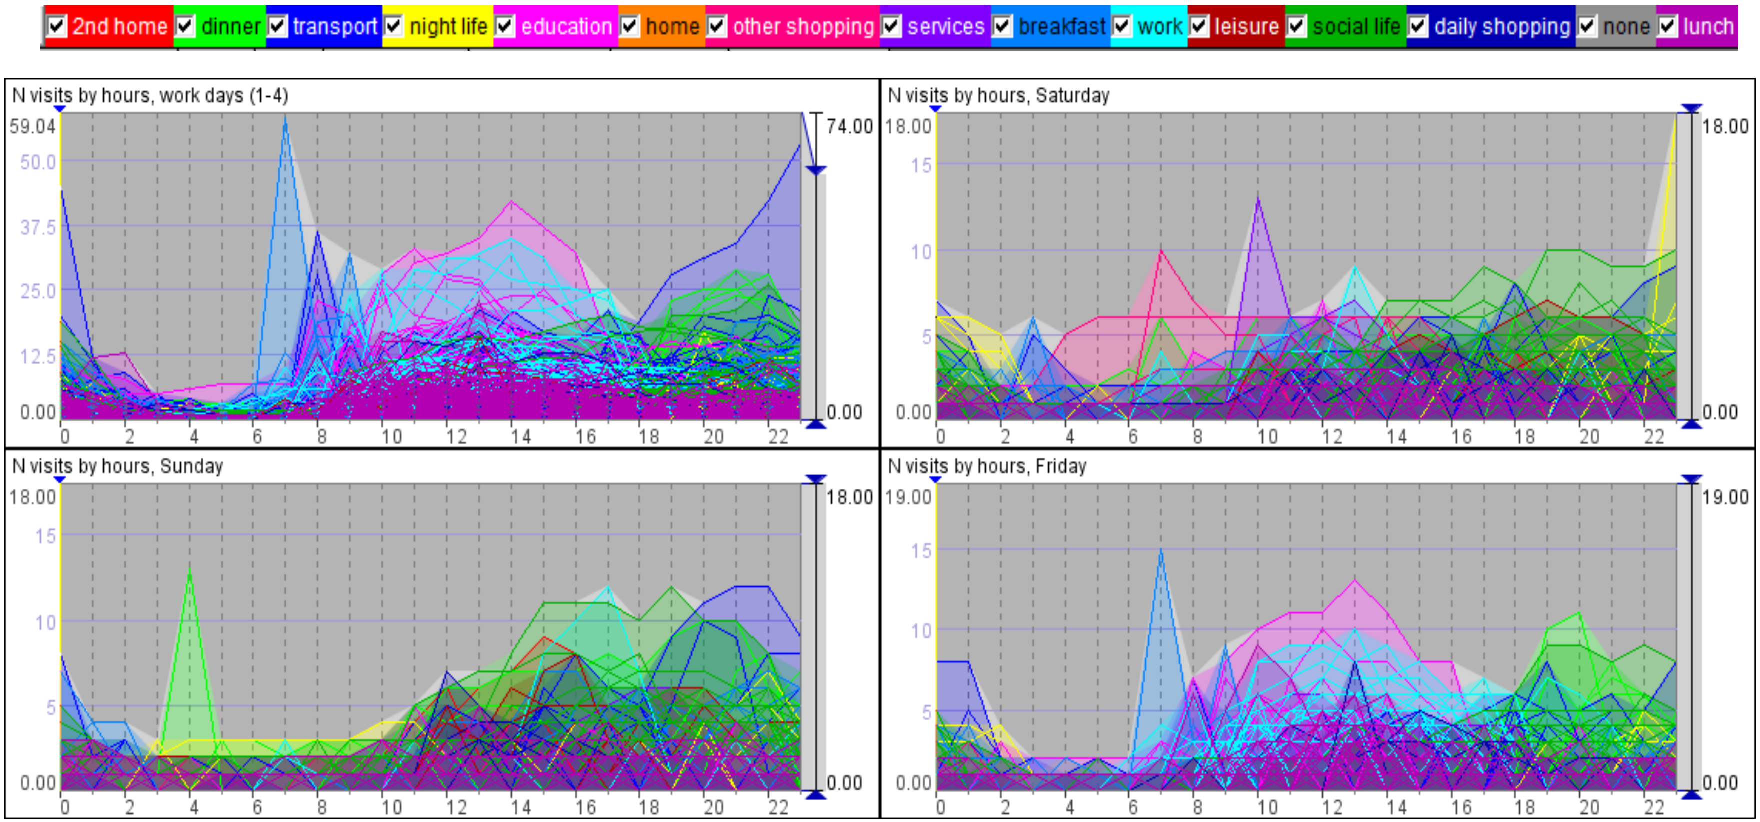

20,247 places (81.2%) are not classified; 4,688 (18.8%) are classified, of them 4,048 (16.2%) were visited in at least 6 d

Activity Recognition

Individual Mobility Networks

How to synthesize Individual Mobility?

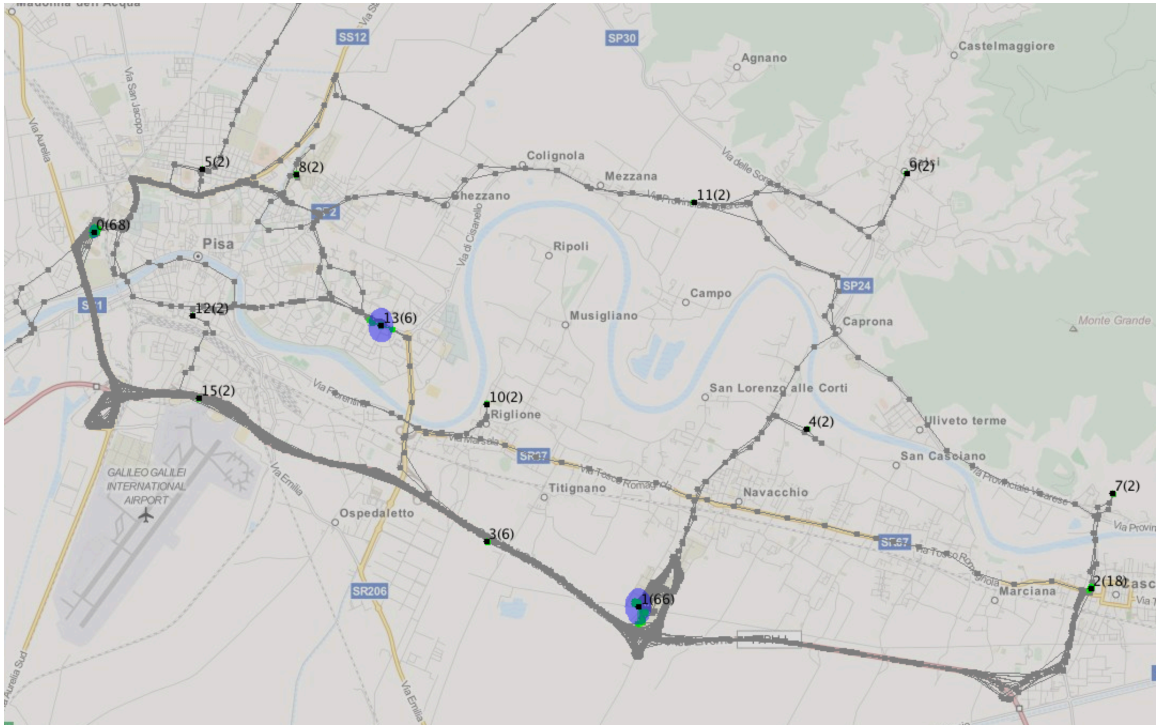

Mobility Data Mining methods automatically extract relevant episodes: locations and movements.

How to synthesize Individual Mobility?

- Basic approach: compute movement features of each trip
  - Length
  - Average speed or Duration
  - Bee-line length
  - Time of the day
  - ...

How to synthesize Individual Mobility?

- More advanced approach: consider overall mobility of the user
- First step: rank individual preferred locations

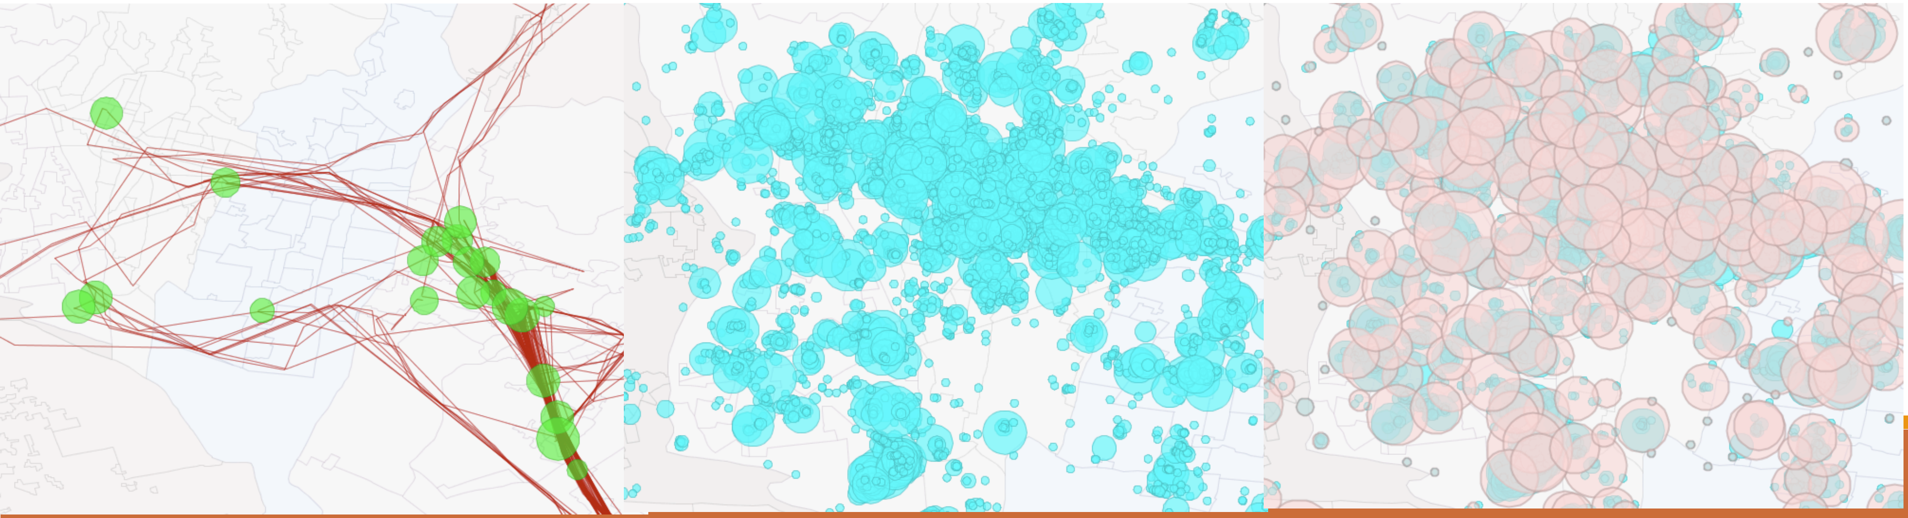

How to synthesize Individual Mobility?

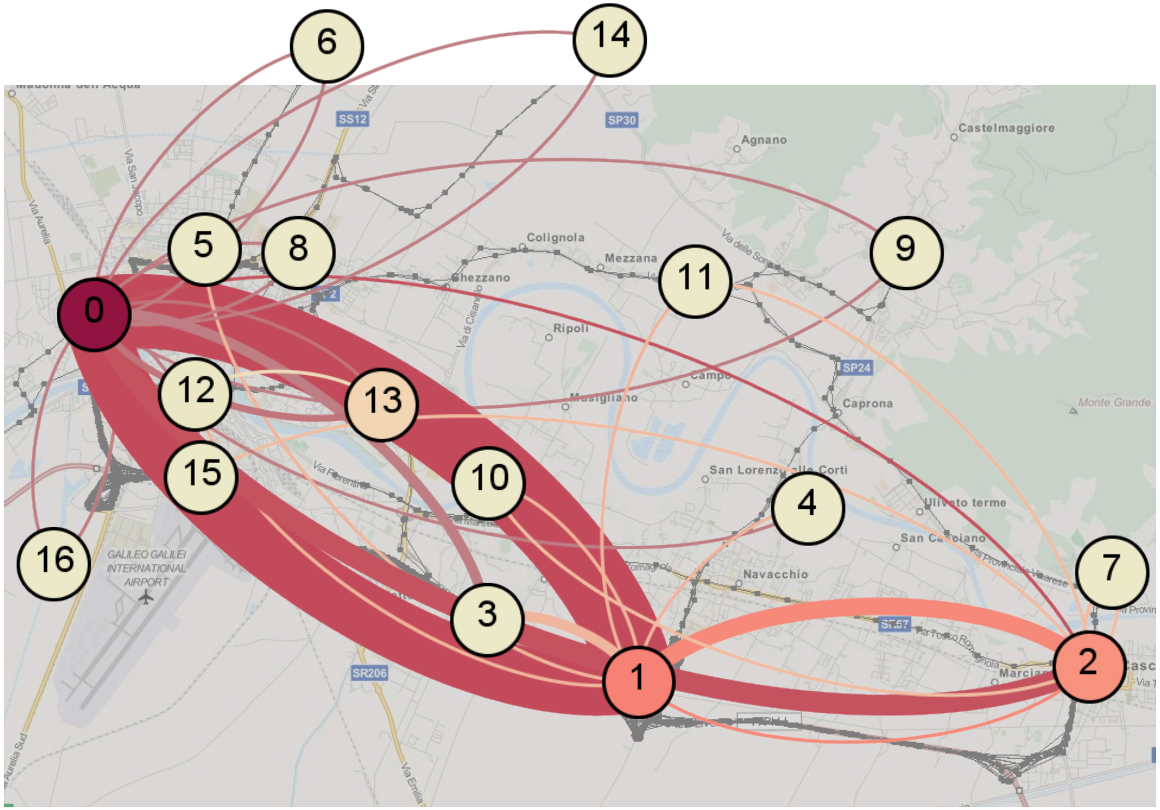

Graph abstraction based on locations (nodes) and movements (edges)

How to synthesize Individual Mobility?

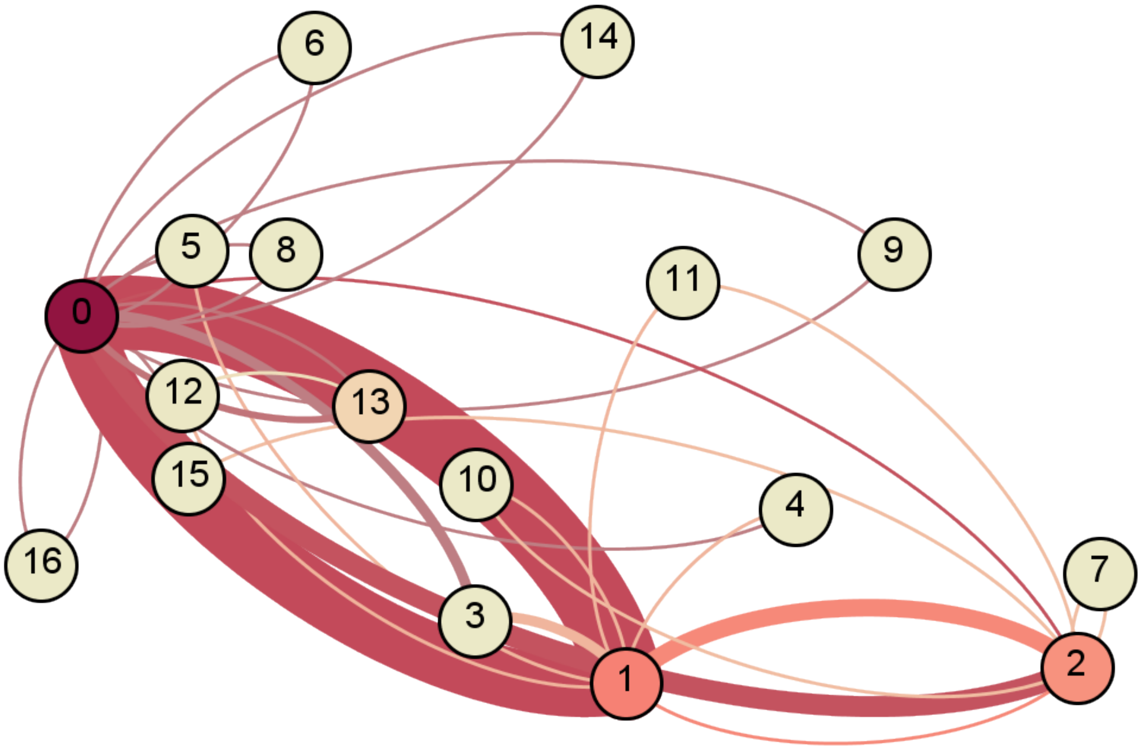

High level representation

Aggregation of sensitive data

Abstraction from real geography

From raw movement…

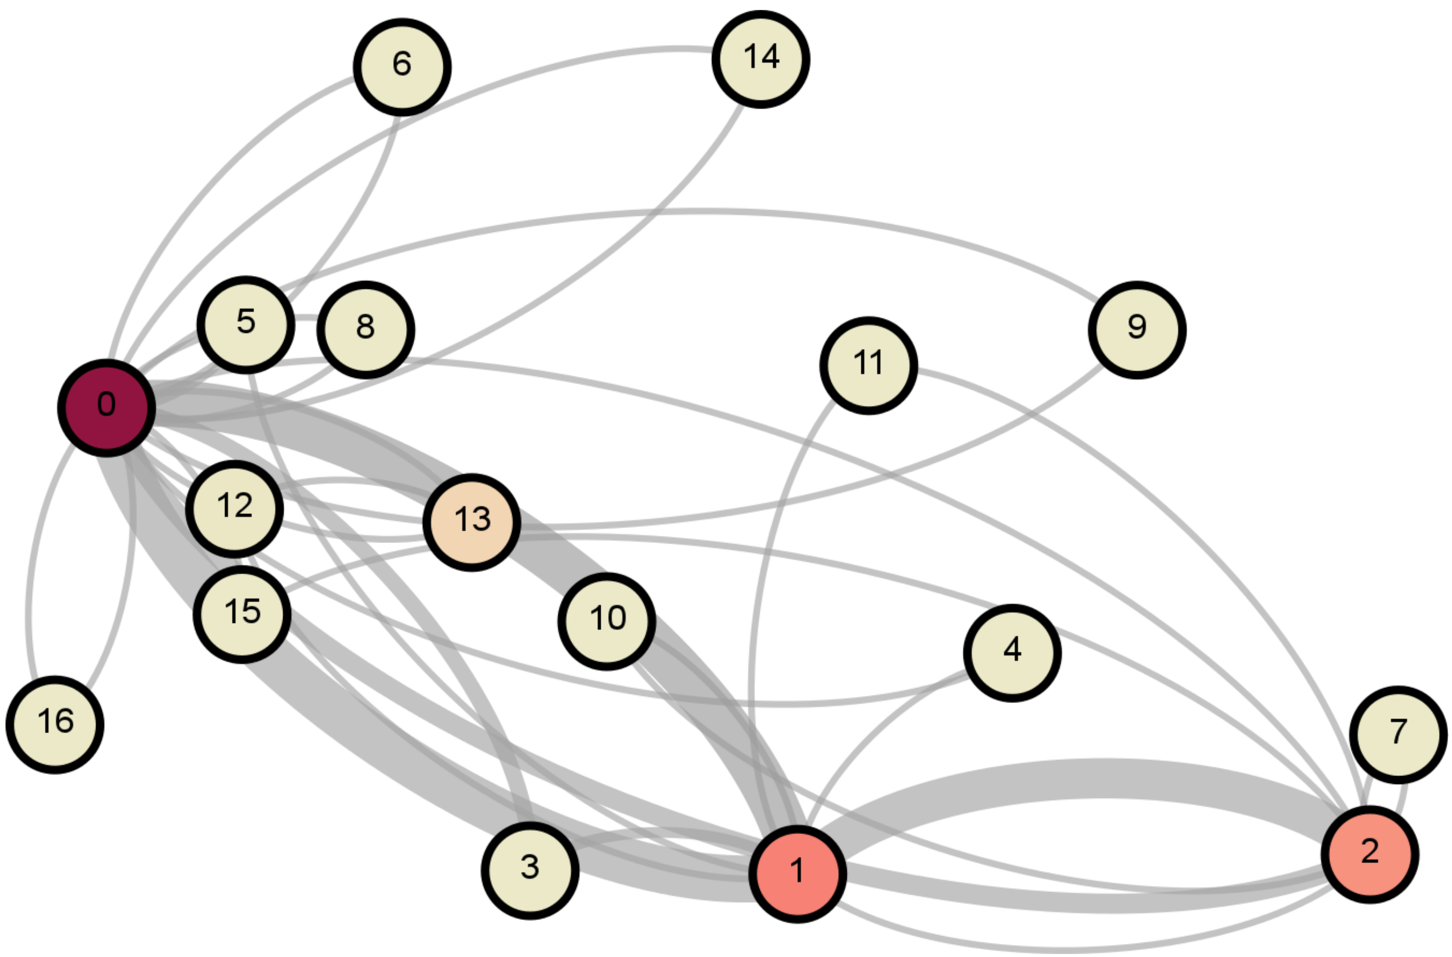

… to annotated data

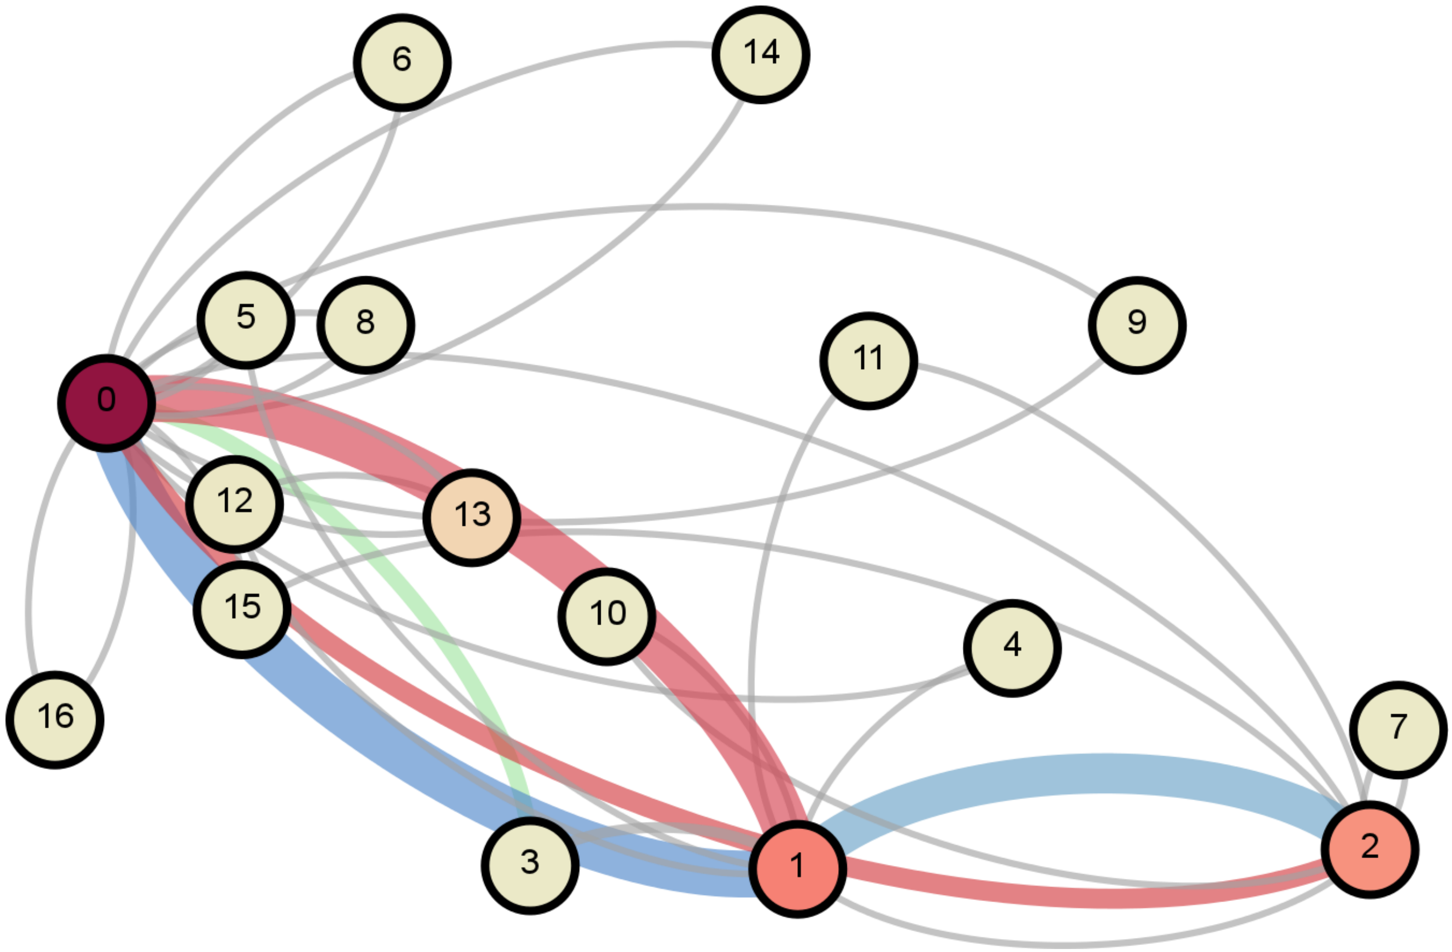

The ABC classifier

1) Build from data an
   Individual Mobility Network (IMN)

2) Extract structural features from the IMN

3) Use a cascading classification with
   label propagation (ABC classifier)

Extracting the IMN

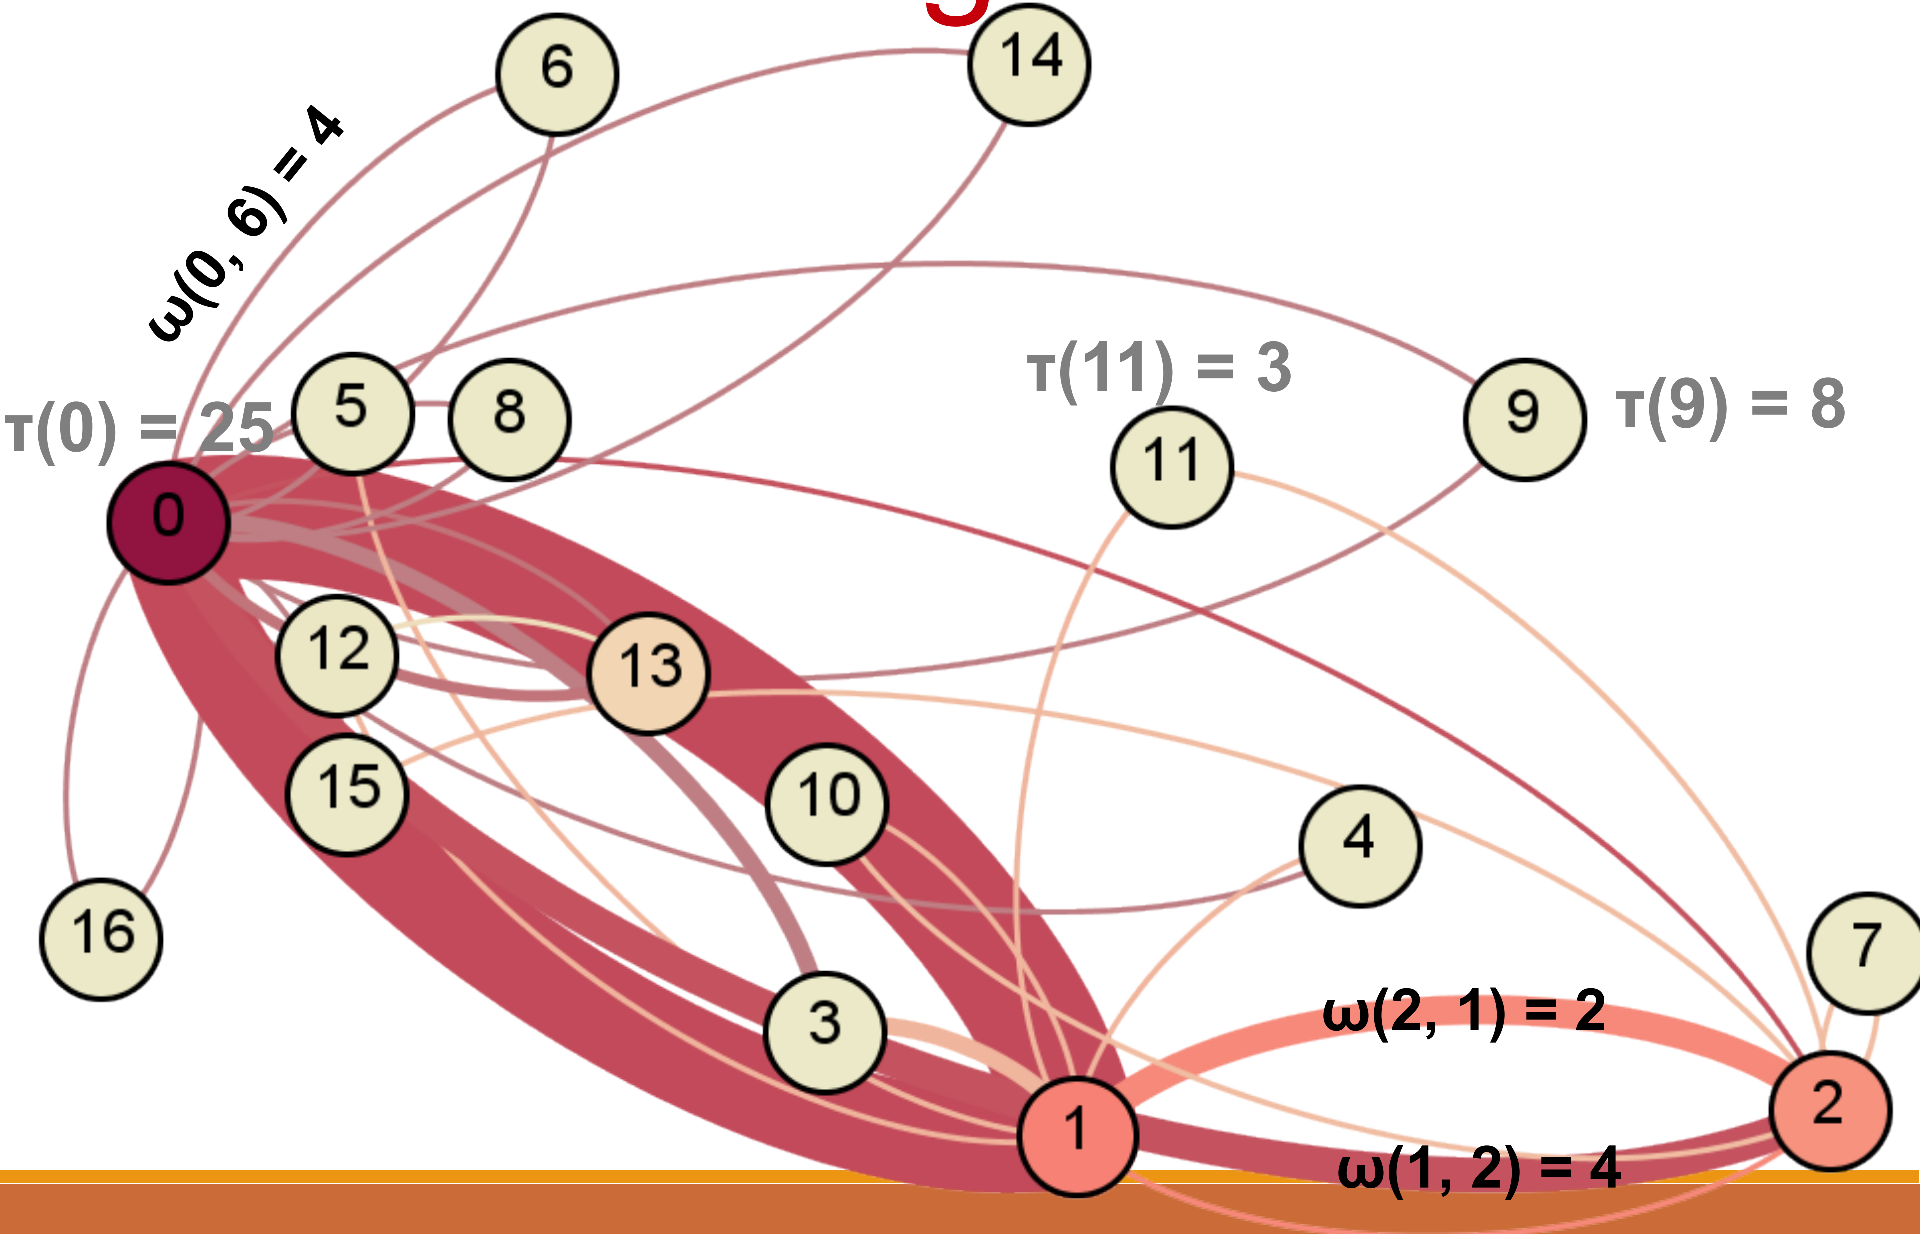

Extracting the IMN

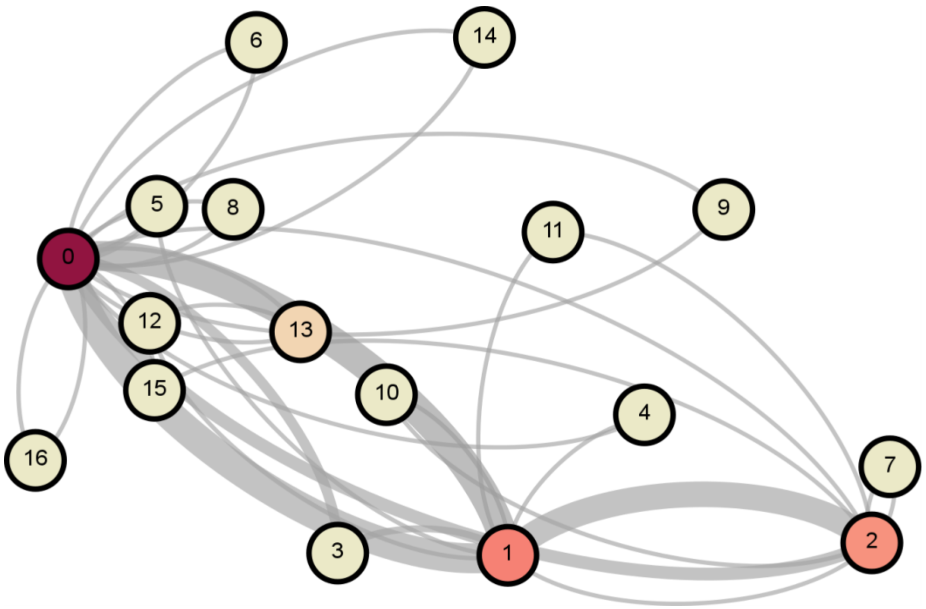

Trip Features
Length
Duration
Time Interval
Average Speed

Network Features	
centrality	clustering coefficient average path length
predictability	entropy
hubbiness	degree betweenness
volume	edge weight flow per location

Extracting the IMN

from	to	weight	ccFrom	ccTo	duration
1	2	2	0.22	0.12	10 min
1	2	2	0.22	0.12	5 min
2	1	3	0.12	0.22	4 min
2	1	3	0.12	0.22	6 min
2	1	3	0.12	0.22	4 min

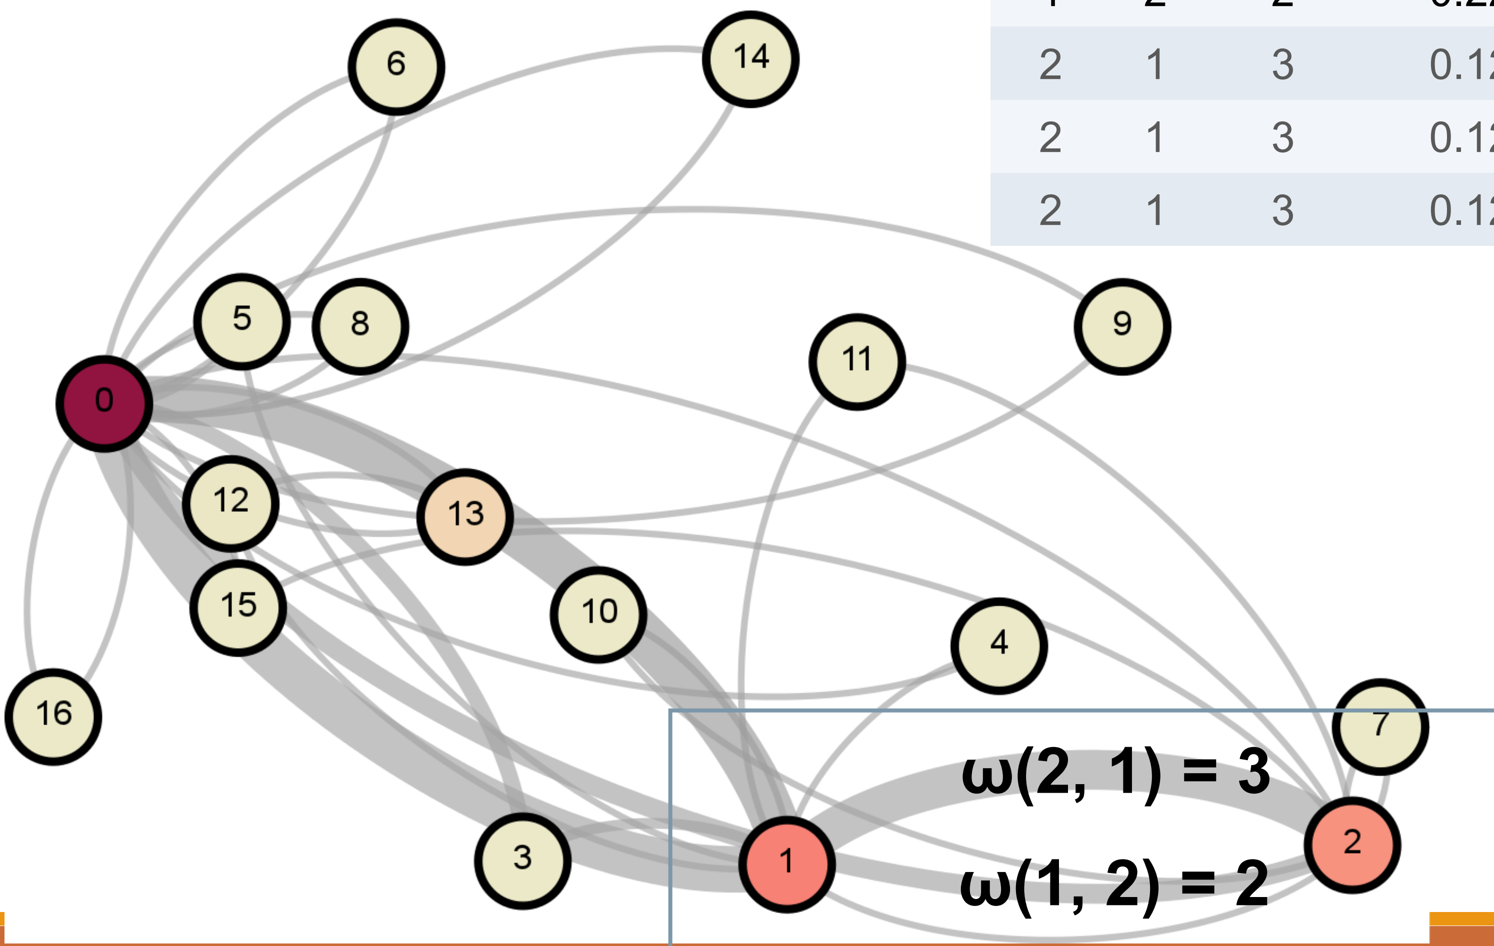

ABC Classifier

- Principles:
  - The activities of a user should be predicted as a whole, not separately
  - Some activities are easy to classify
  - Other activities might benefit from contextual information obtained from previous predictions
- E.g.: a place frequently visited after work might be more likely to be leisure / shopping

ABC Classifier

- Reduce the multi-class problem into several binary problems
- The binary classifiers are learnt in cascade
- The classification results of each step are used as source for later classifications

Example

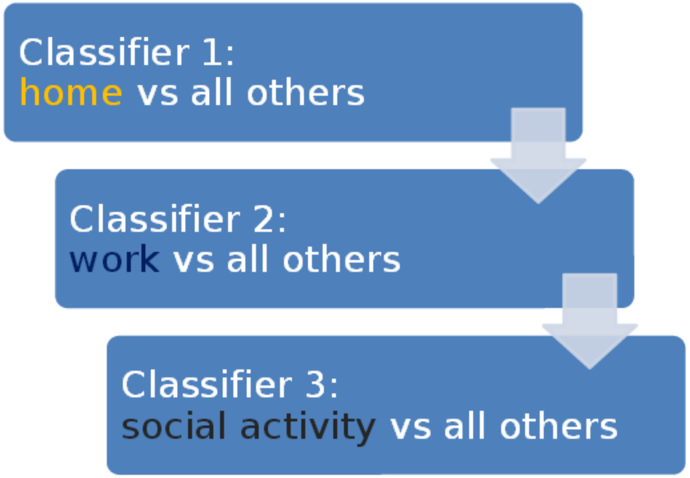

ABC Classifier

- Inspired by Nested Cascade Classification

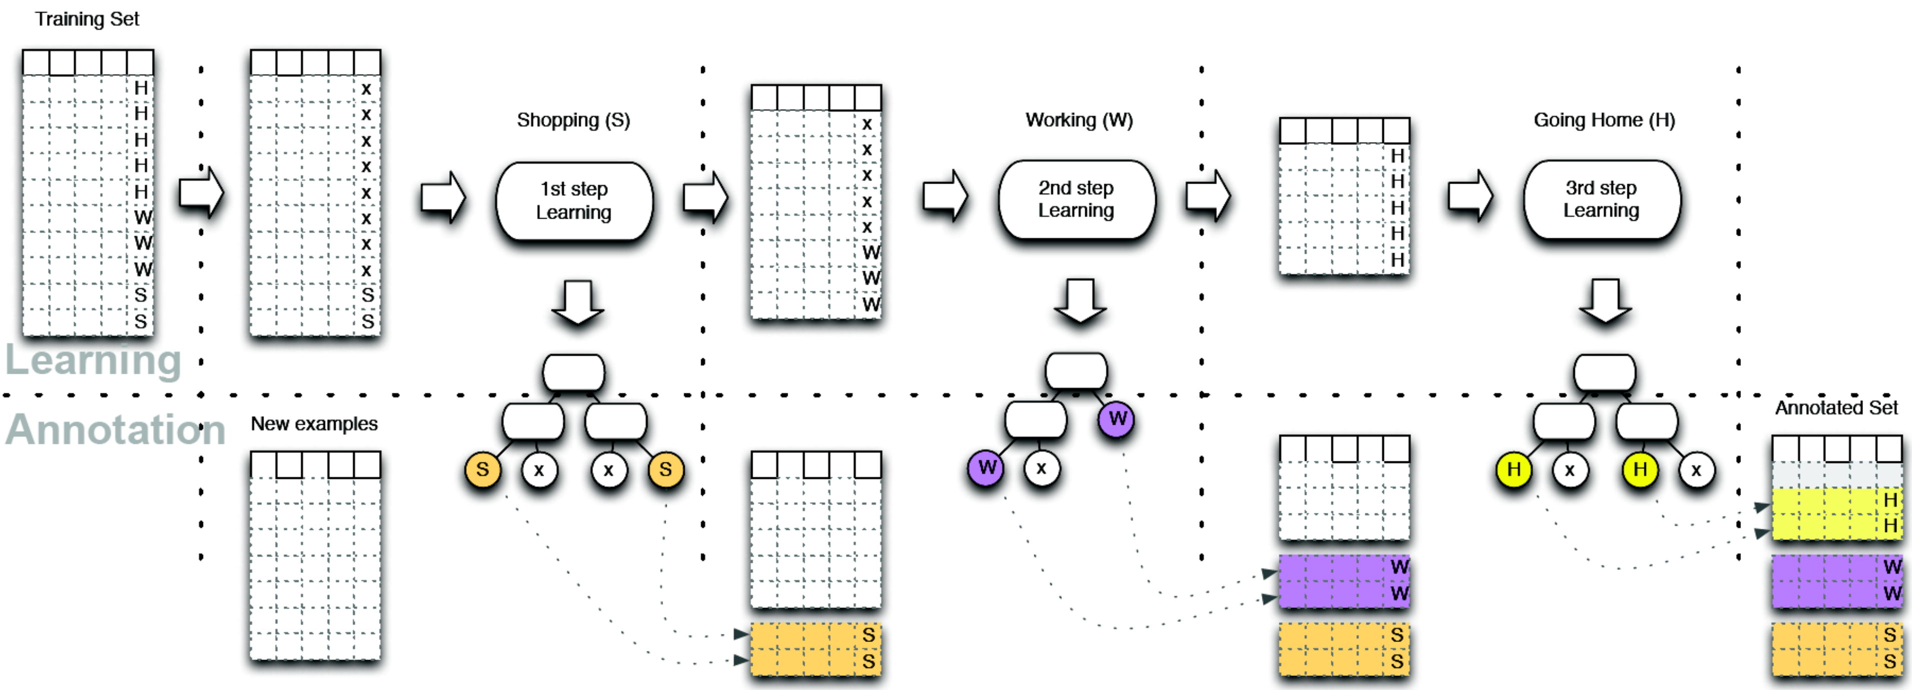

ABC Classifier

- Inspired by Nested Cascade Classification

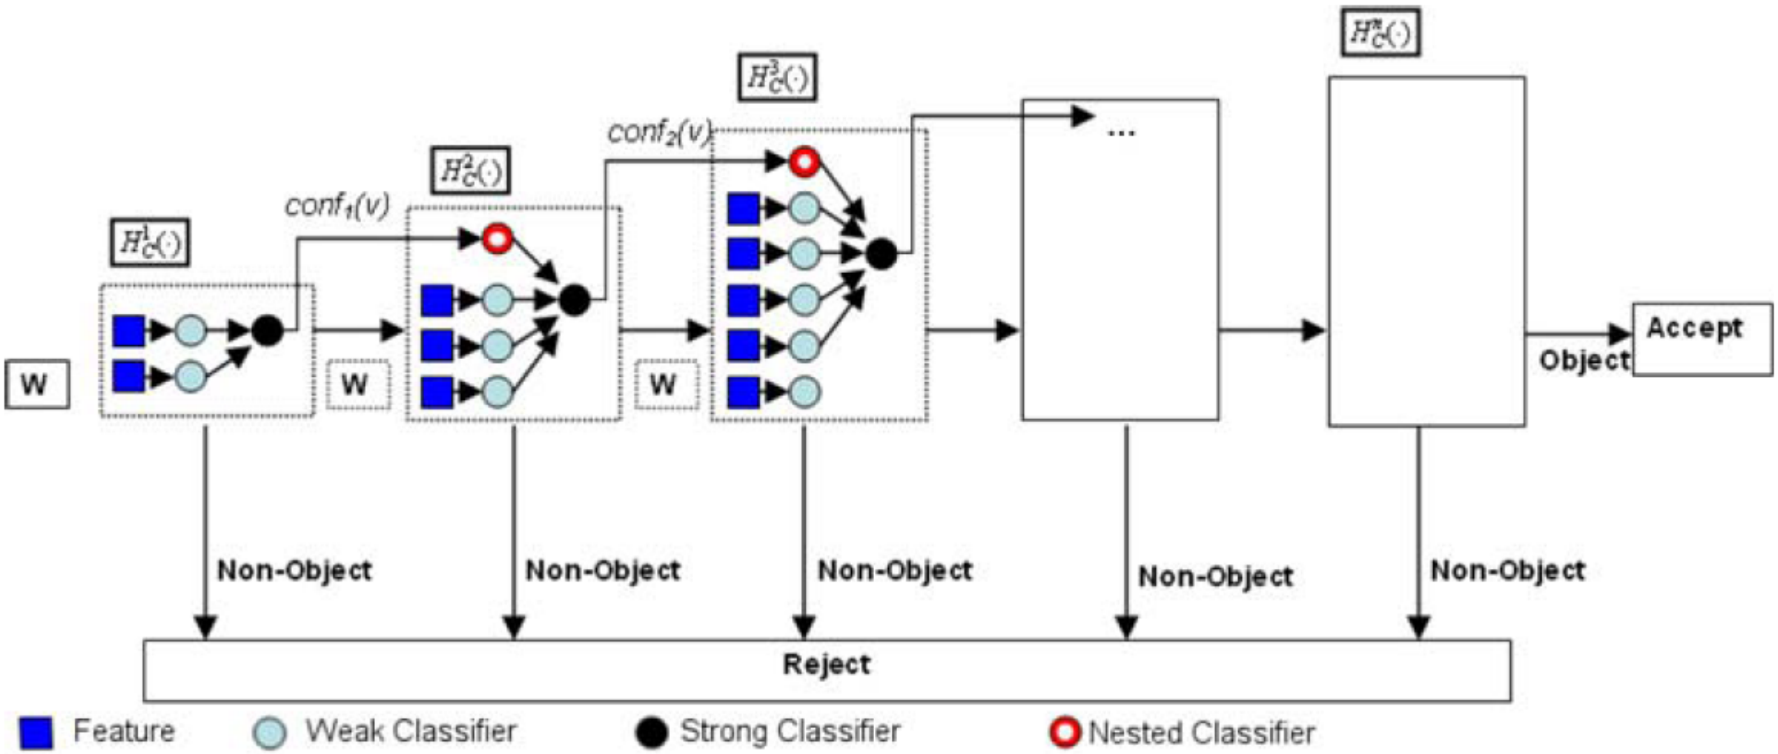

ABC Classifier

- After recognizing an activity (e.g. work), we use this information to enrich the features of the yet-unclassified trips
- E.g. add a feature describing whether the remaining trips are adjacent to the previous activity
  - Are there direct trips from work to the new place?

The ABC classifier

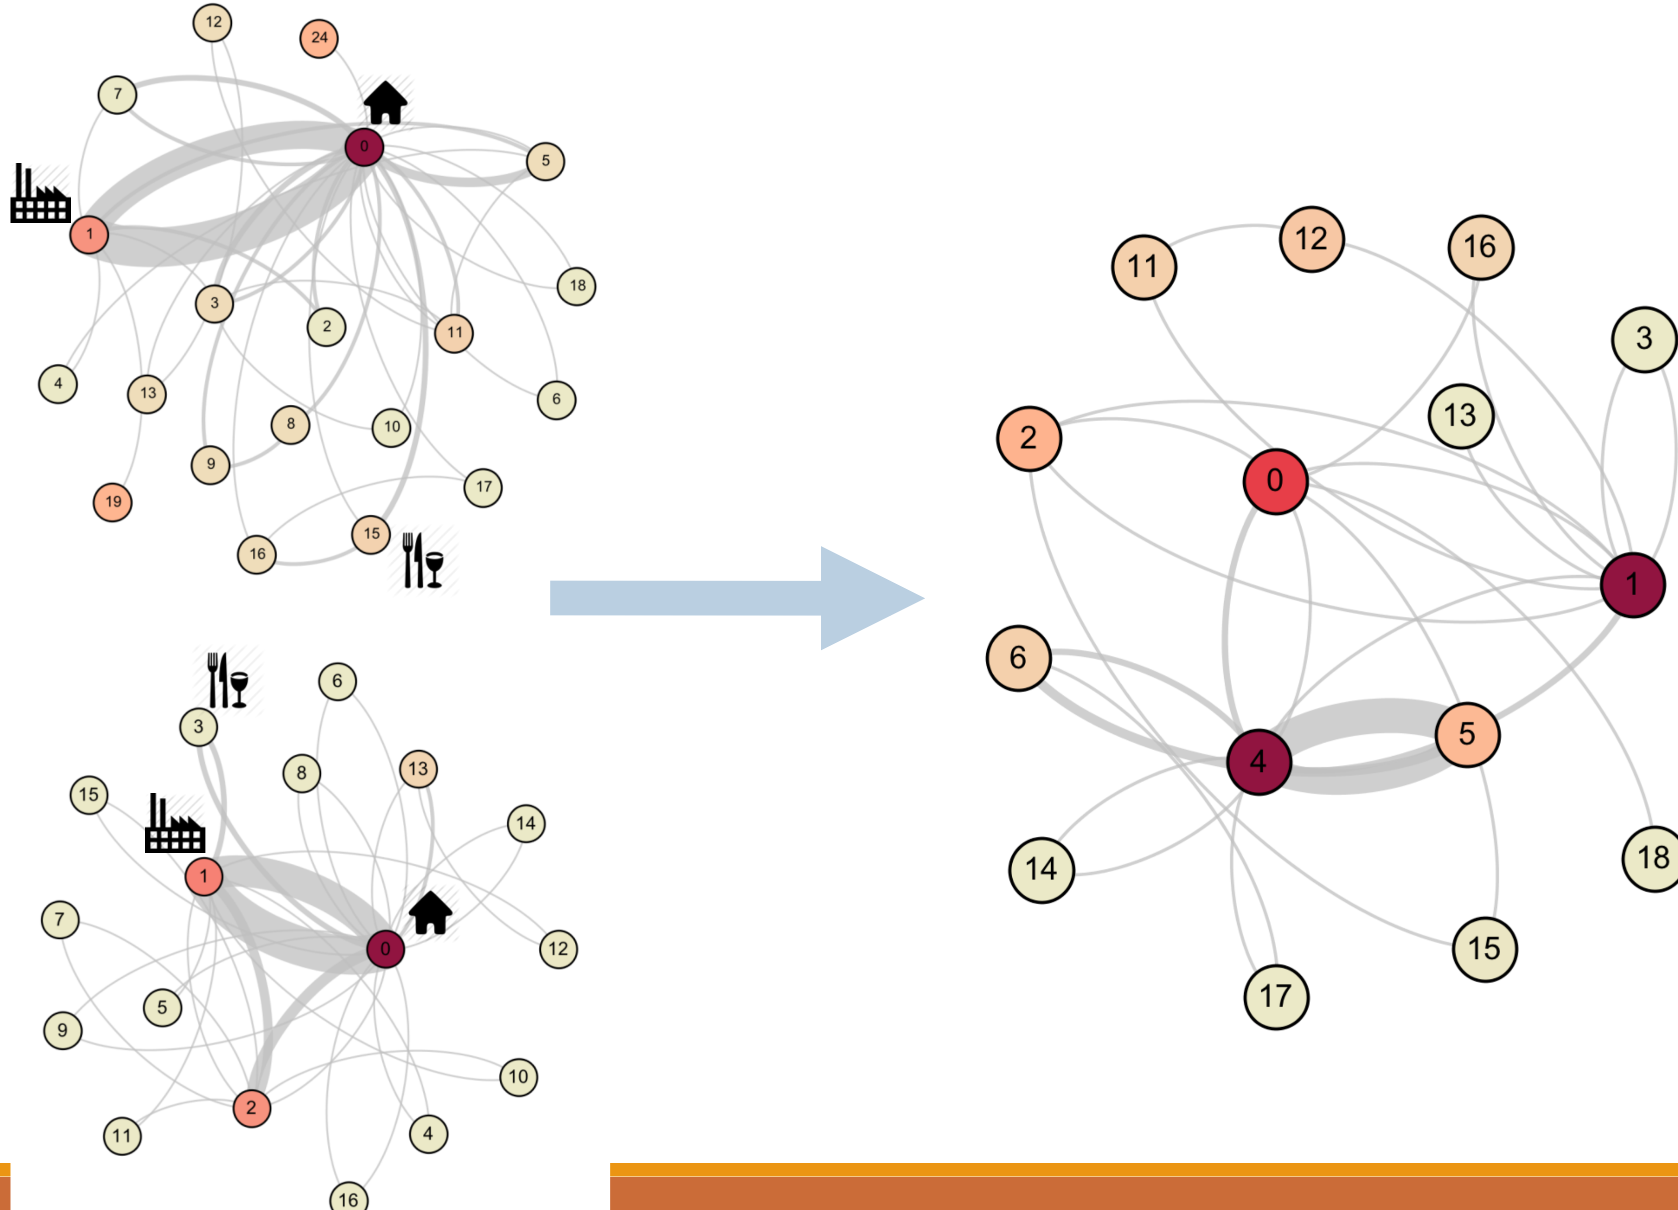

The ABC classifier

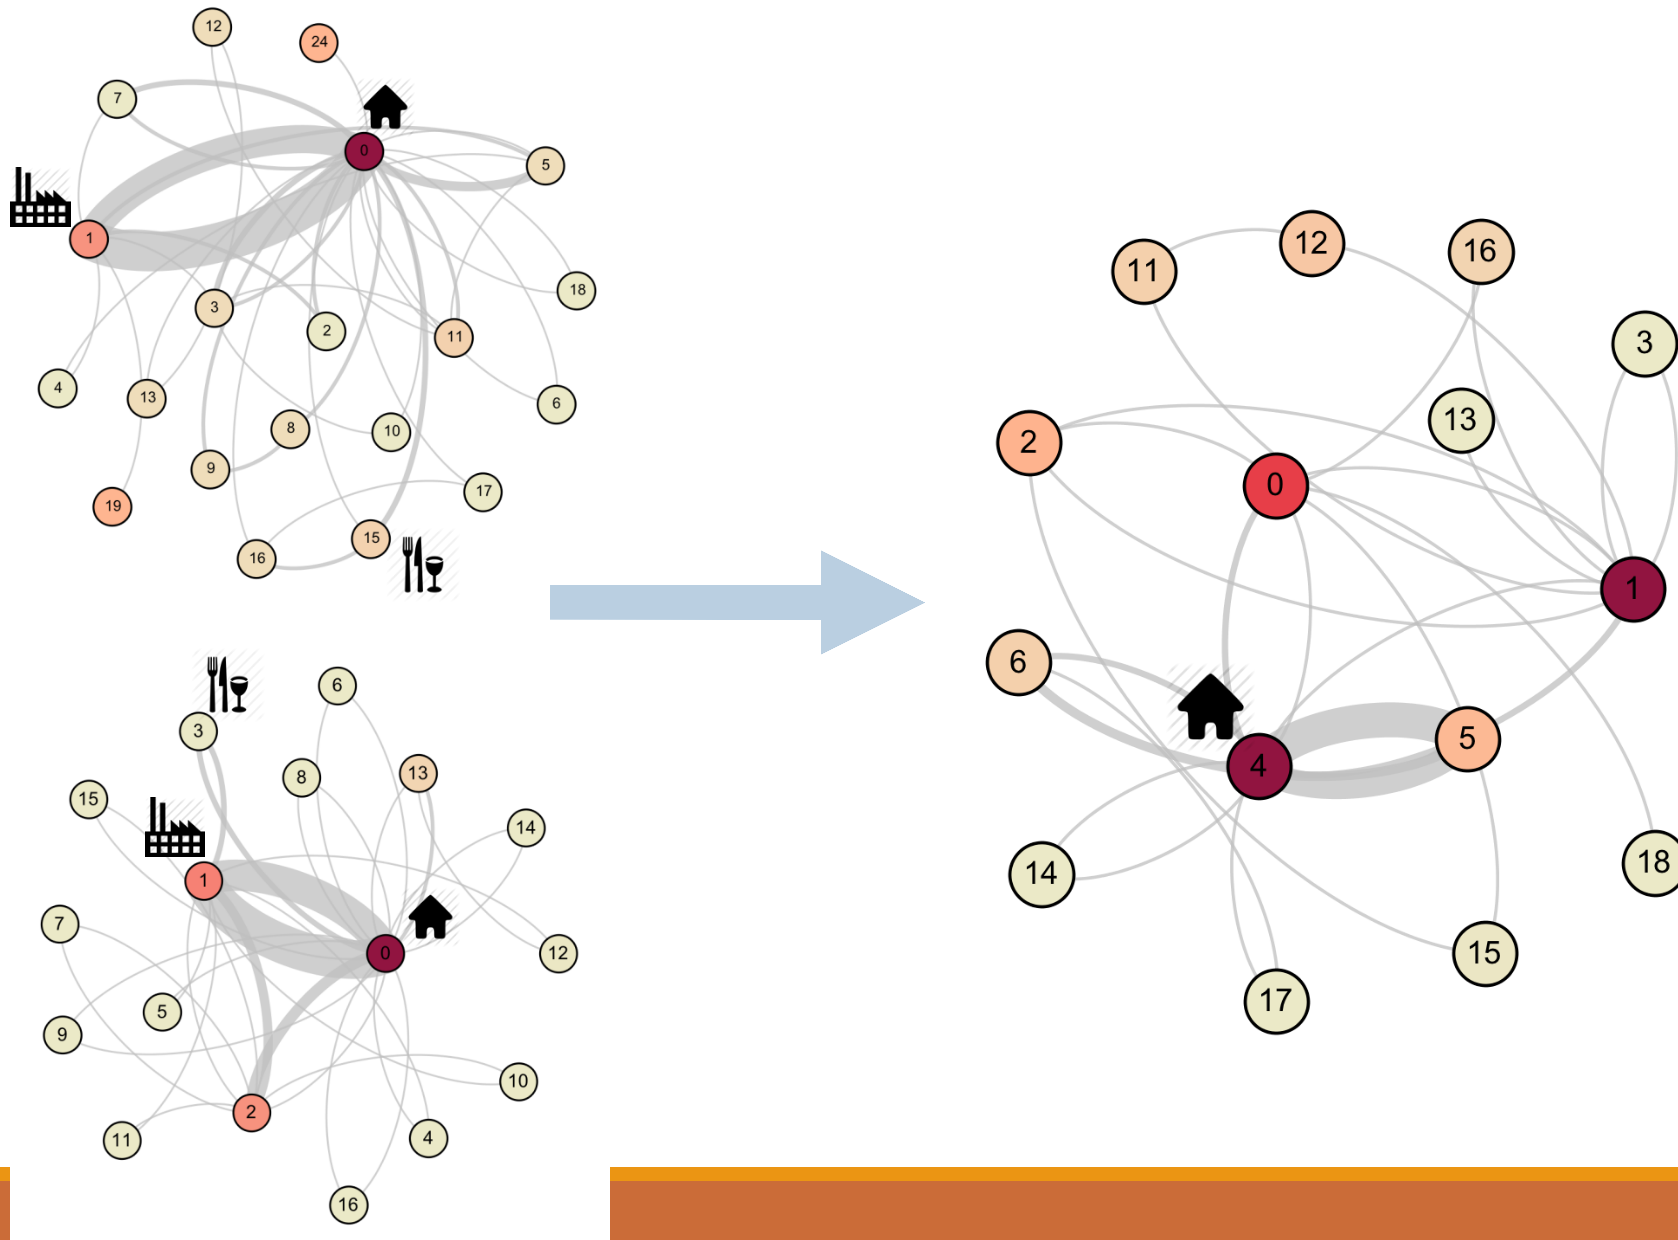

The ABC classifier

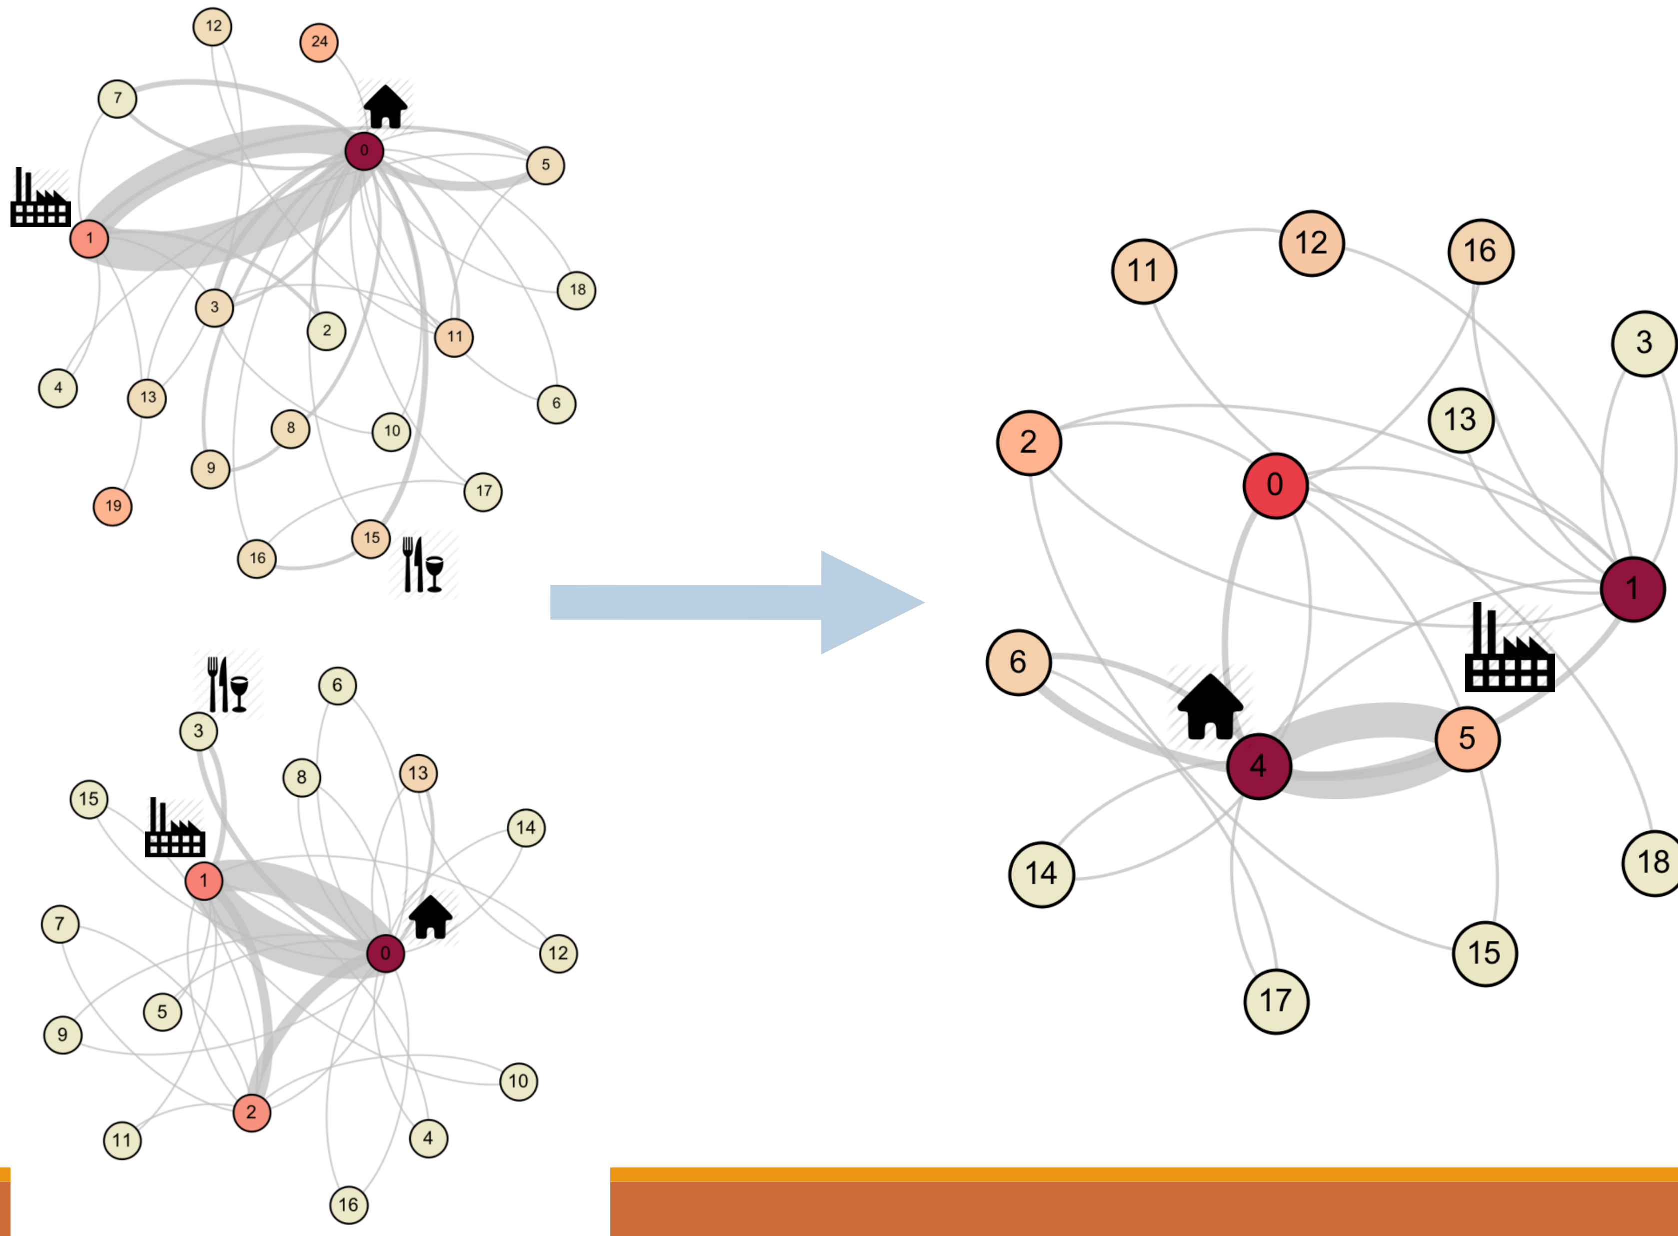

The ABC classifier

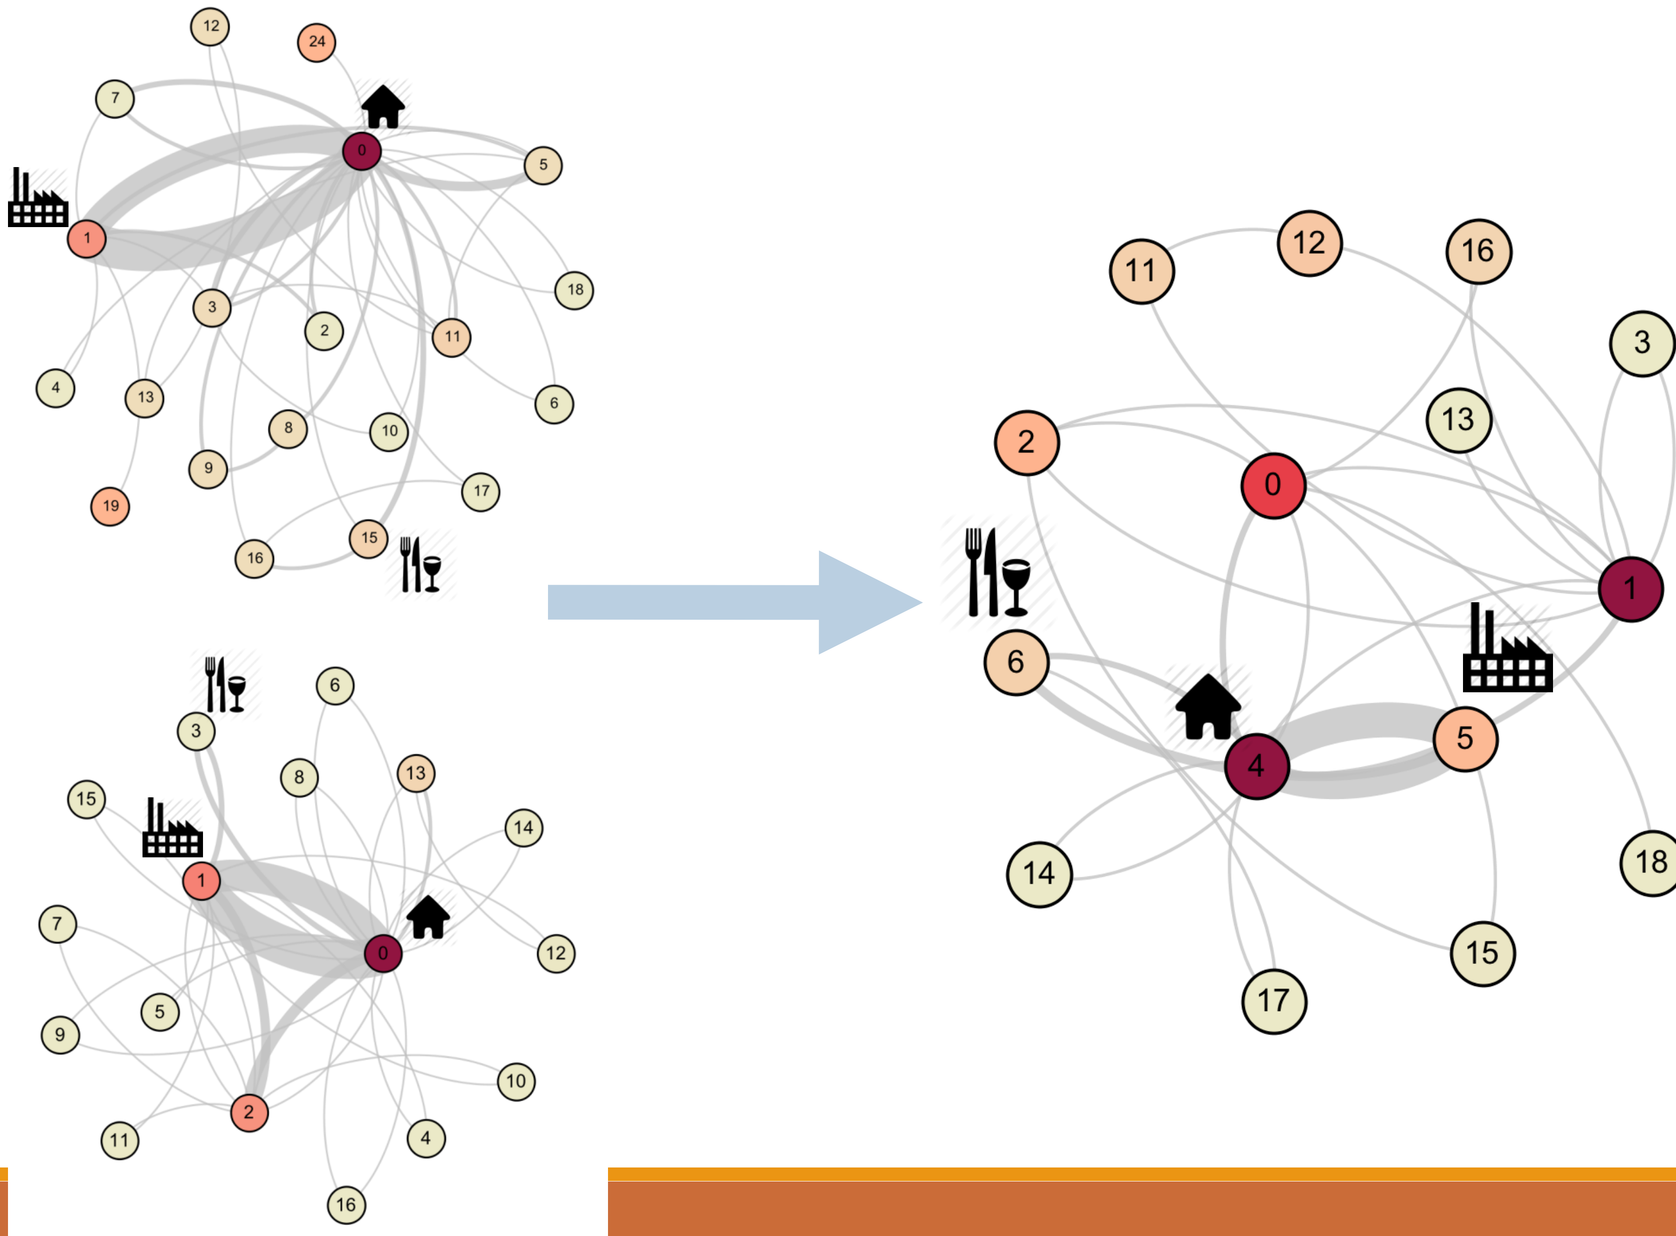

The ABC classifier

Reduce the multi-class problem into several binary problems

The binary classifiers are learnt in cascade

The classification results of each step are used as source for later classifications

Example

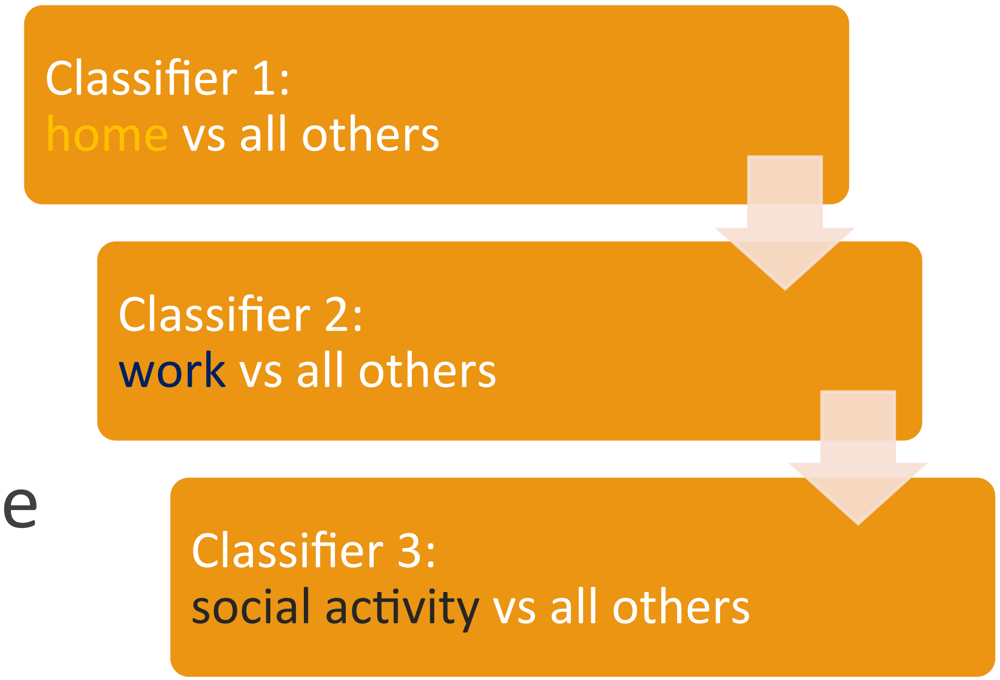

ABC Classfier

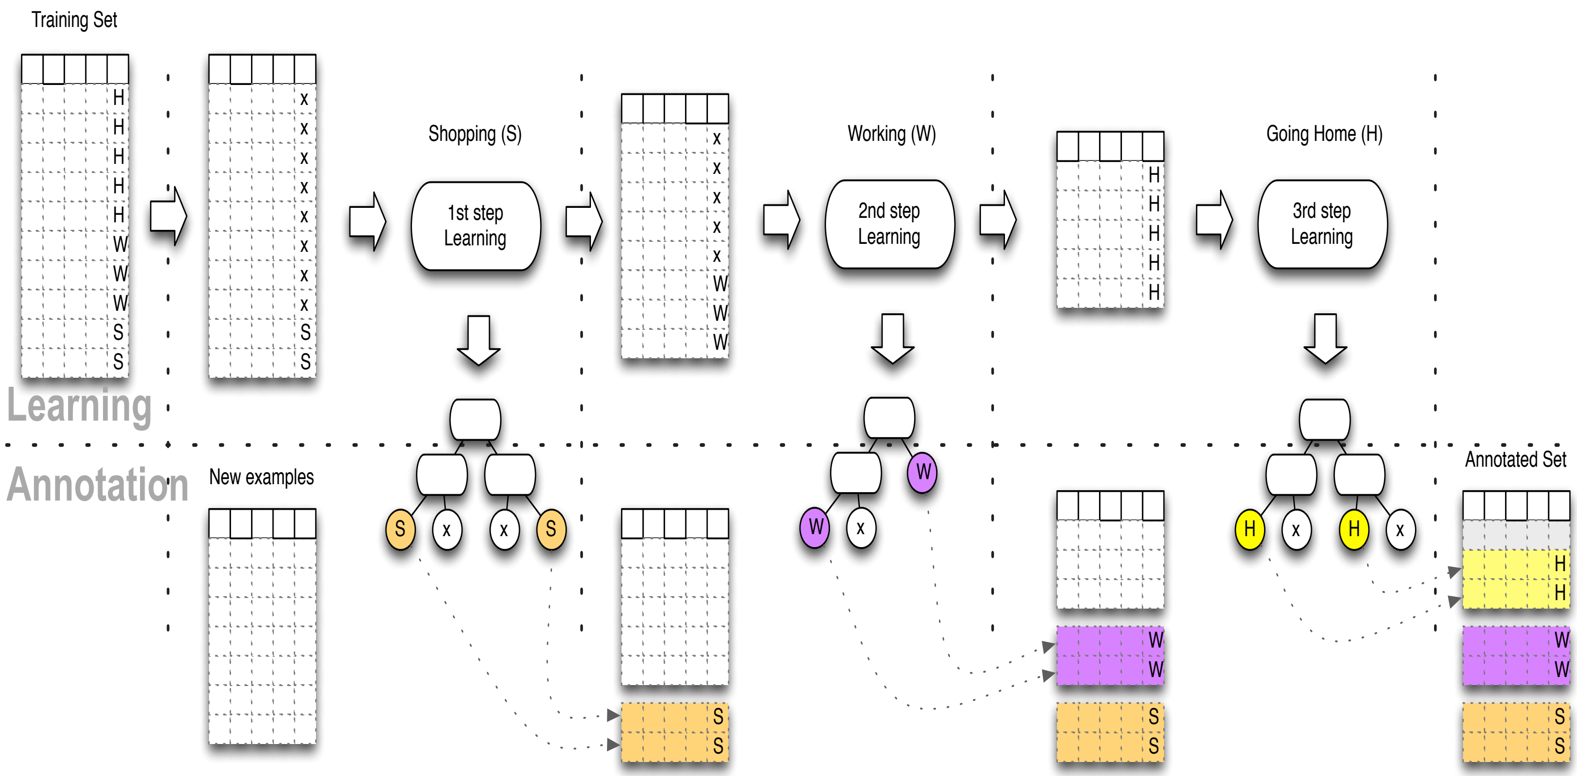

Experiments

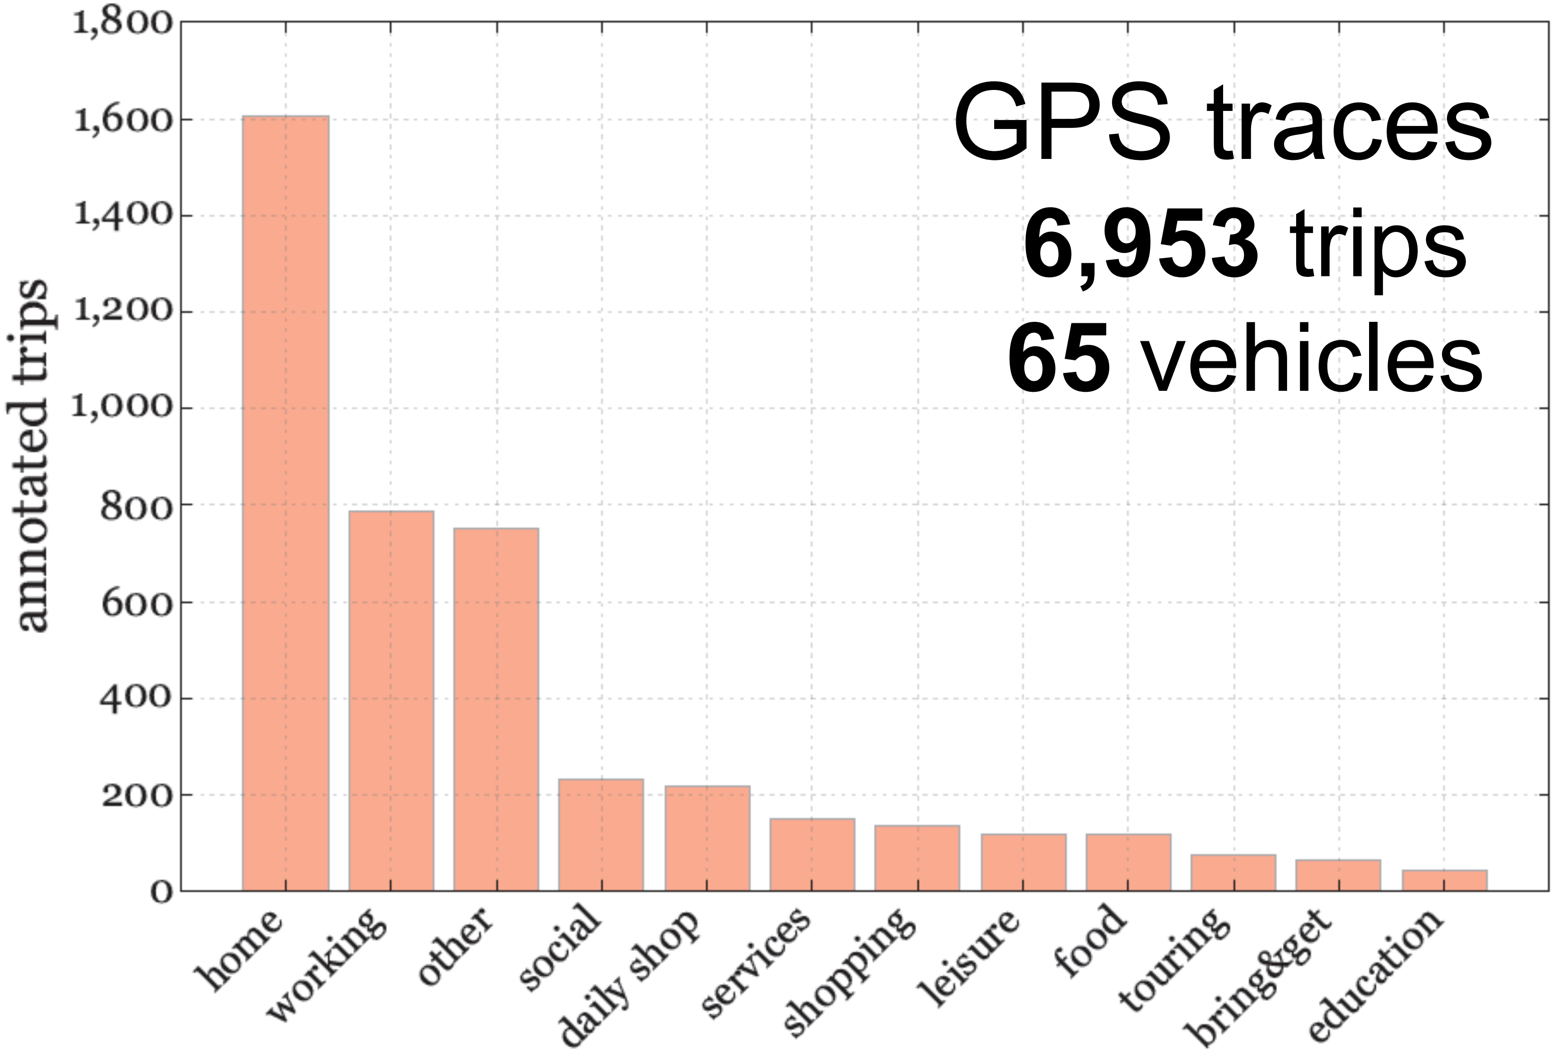

Experiments

- Is the order of activities in the learning relevant?

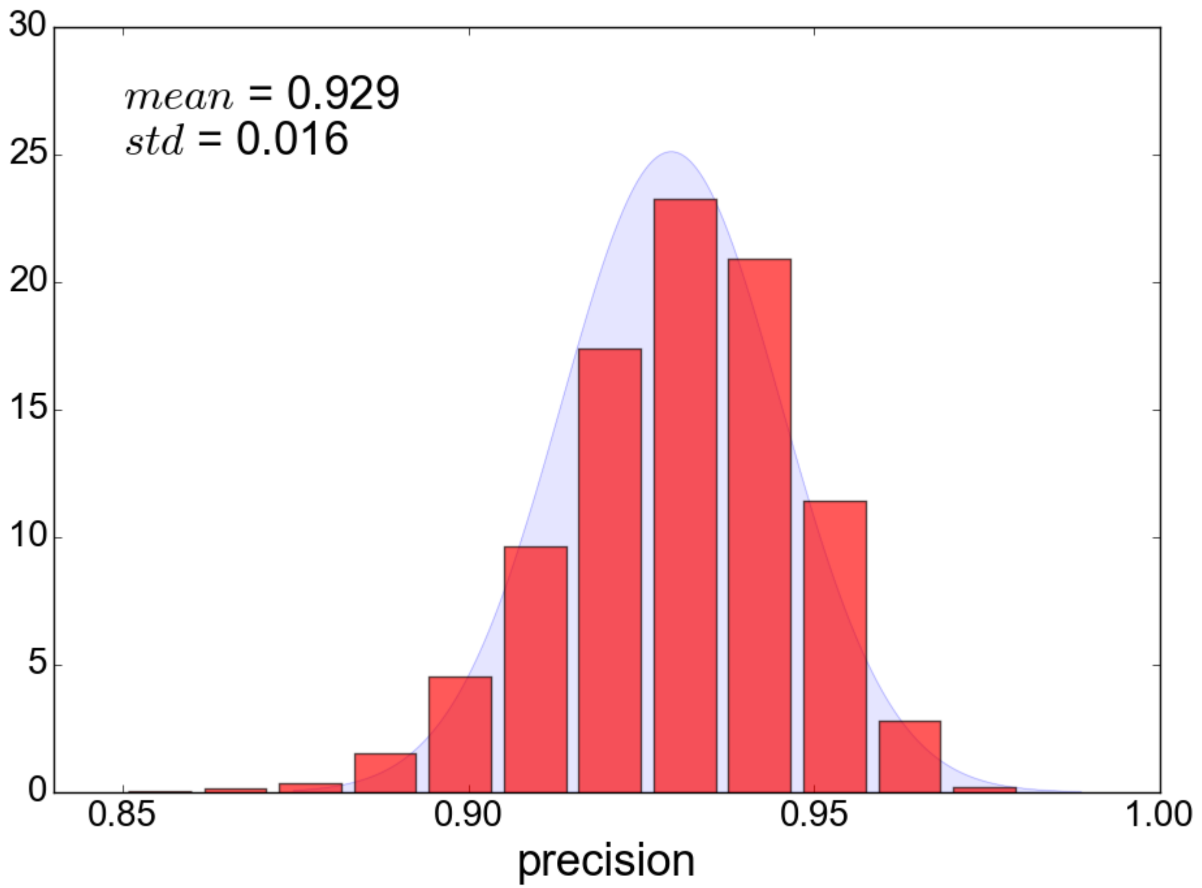

Semantic Mobility Analytics

Temporal Analysis

- Pisa traffic

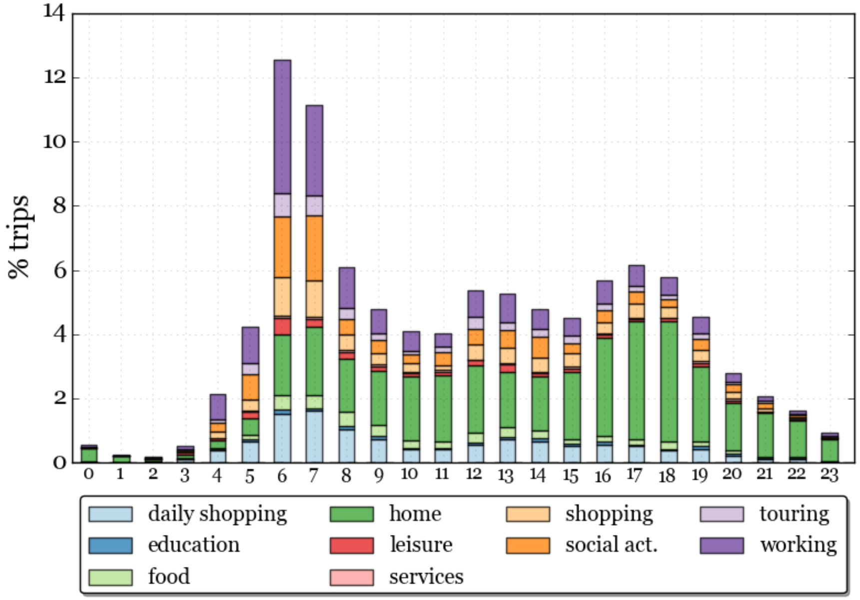

In

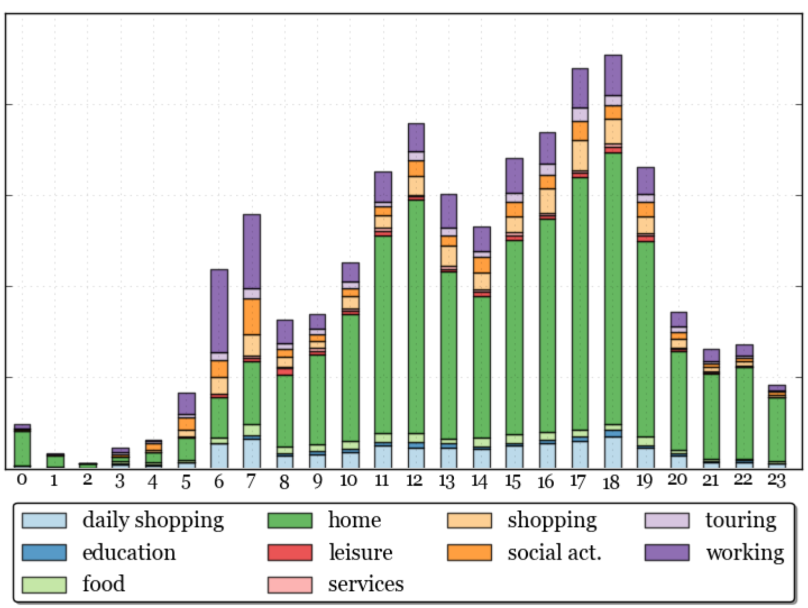

Out

Semantic Mobility Analytics

Temporal Analysis

- Calci traffic

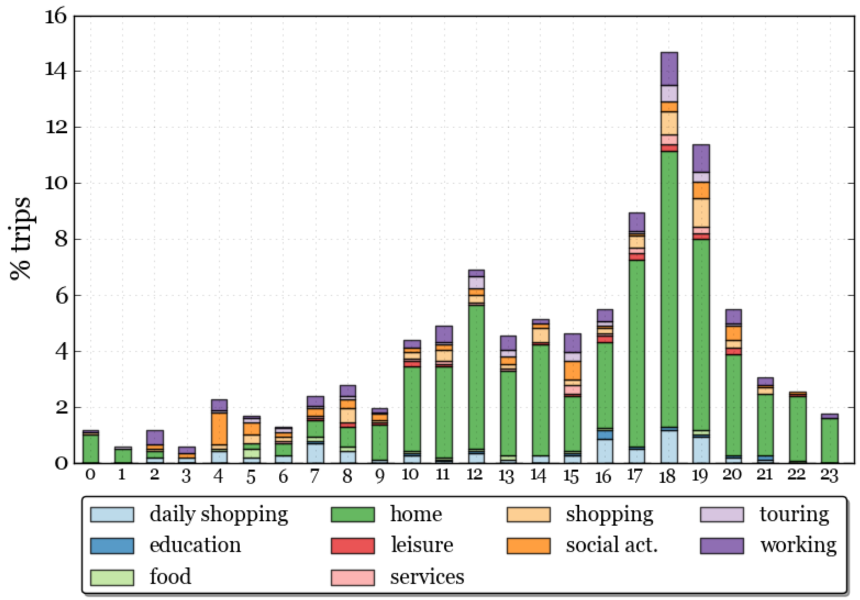

In

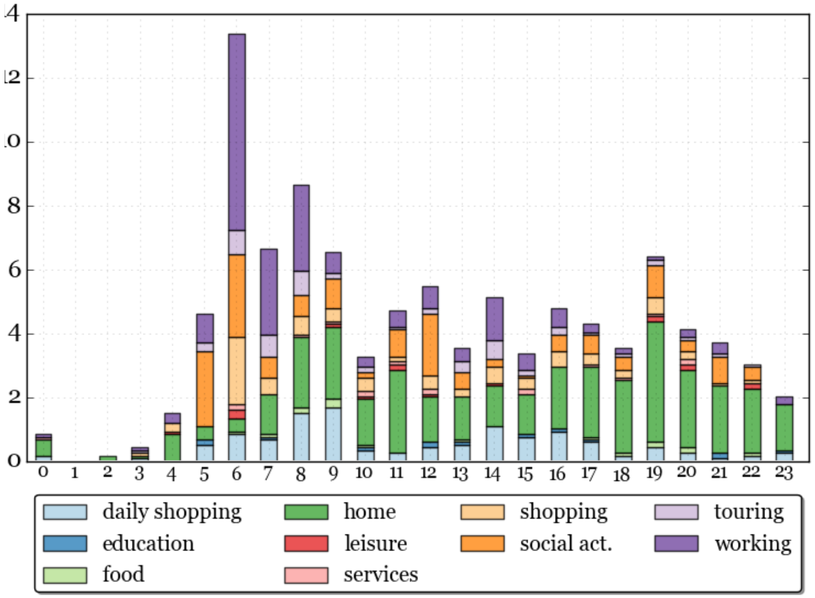

Out

Semantic Mobility Analytics

Temporal Analysis

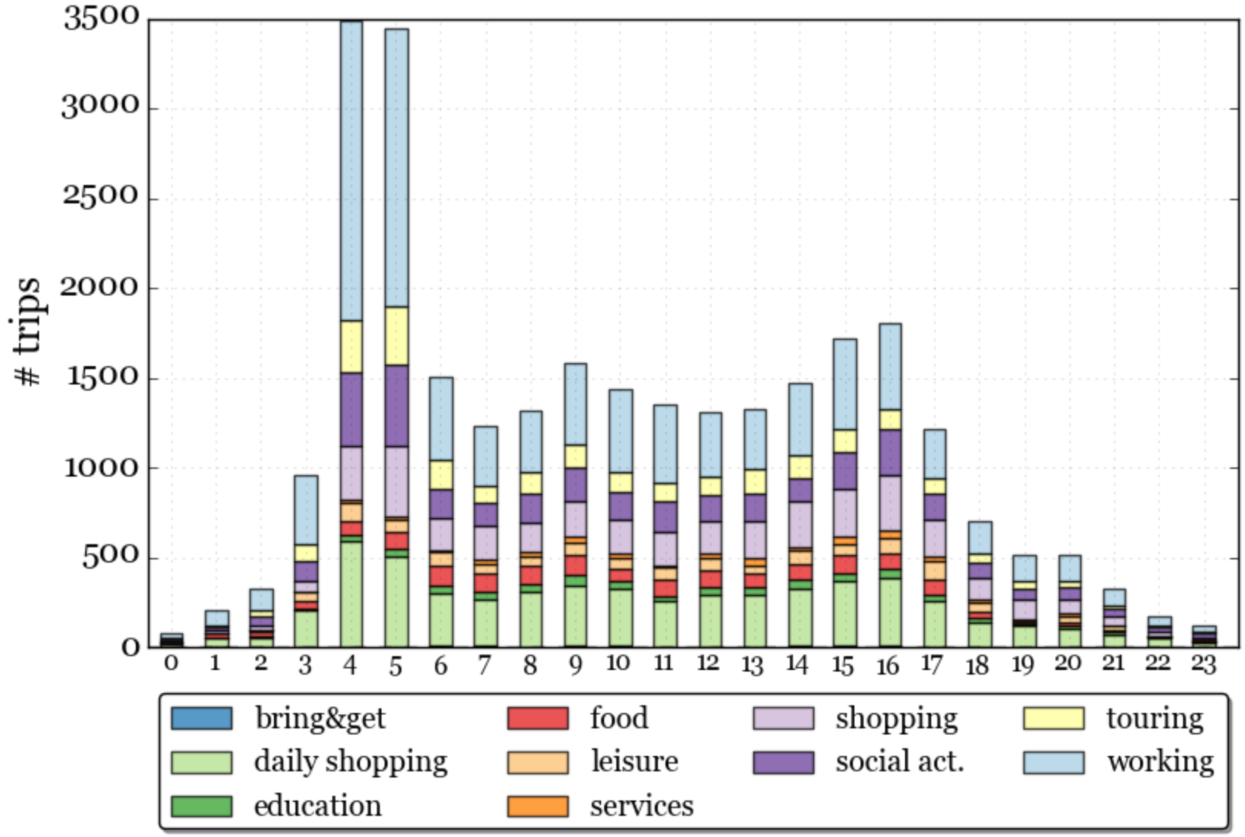

The Purpose of Motion

Given a small training mobility set annotated with activities (home, work, leisure,bring&get,eating-out..)

- Build from data the Individual Mobility Network (IMN)
- Extract structural features from the IMN
- Use a cascading classification with label propagation (ABC classifier)

The purpose of motion: Learning activities from individual mobility networks
S Rinzivillo, L Gabrielli, M Nanni, L Pappalardo, D Pedreschi, F Giannotti, DSAA 2014

Trajectory Prediction

Individual and Collective Profile

- **Individual Profile**
- Input: Individual Data
- Output: Individual Patterns

- **Collective Profile**
- Input: Collectivity Data
- Output: Collective Patterns

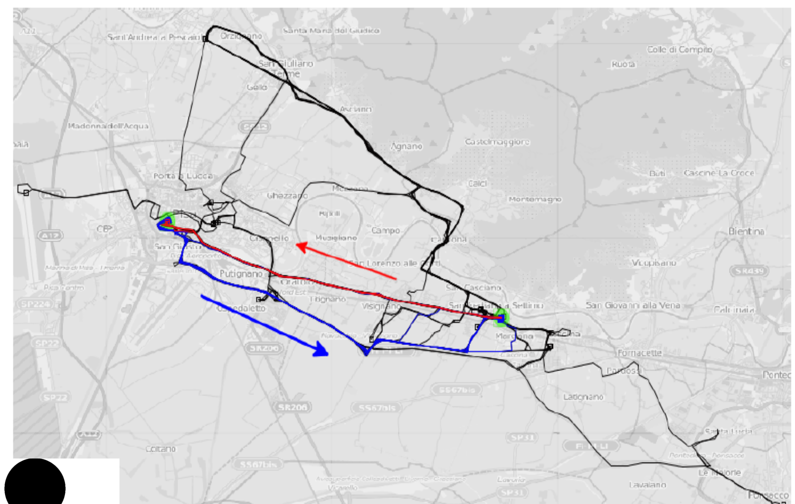

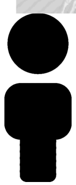

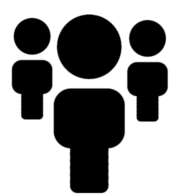

Collective prediction using t-patterns

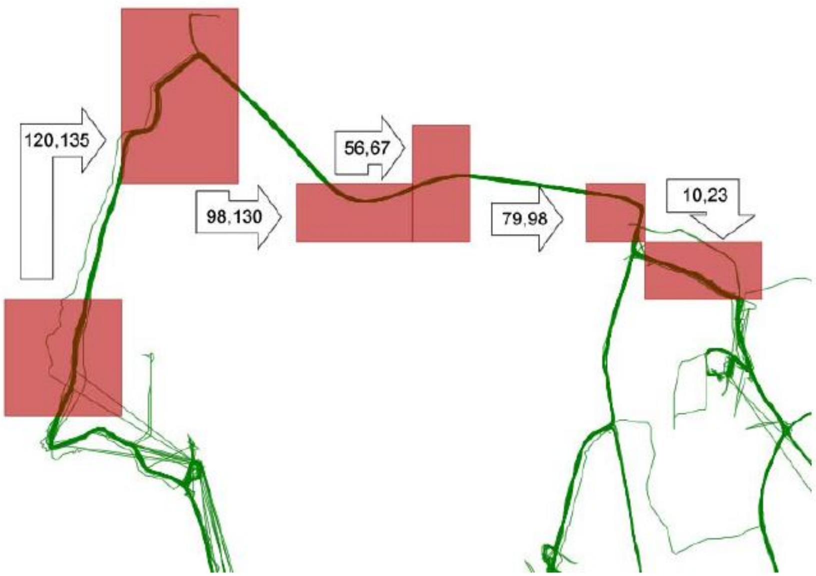

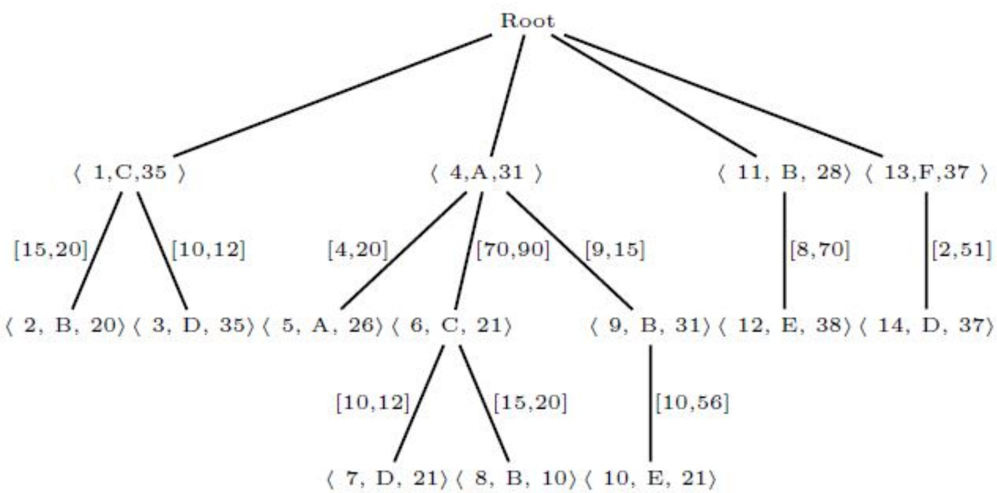

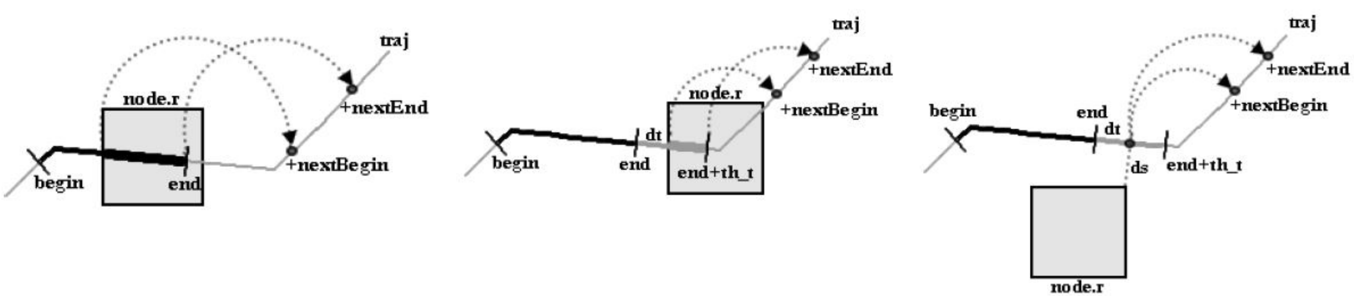

p_score = node.support

p_score = node.support/β*d_t

p_score = node.support/(β*d_t+α*d_s)

A. Monreale, F. Pinelli, R. Trasarti, and F. Giannotti. **Wherenext: a location predictor on trajectory pattern mining.** 2009.

Mobility Profiling

A concise model ables to describe the user's mobility in terms of representative movements, i.e. routines.

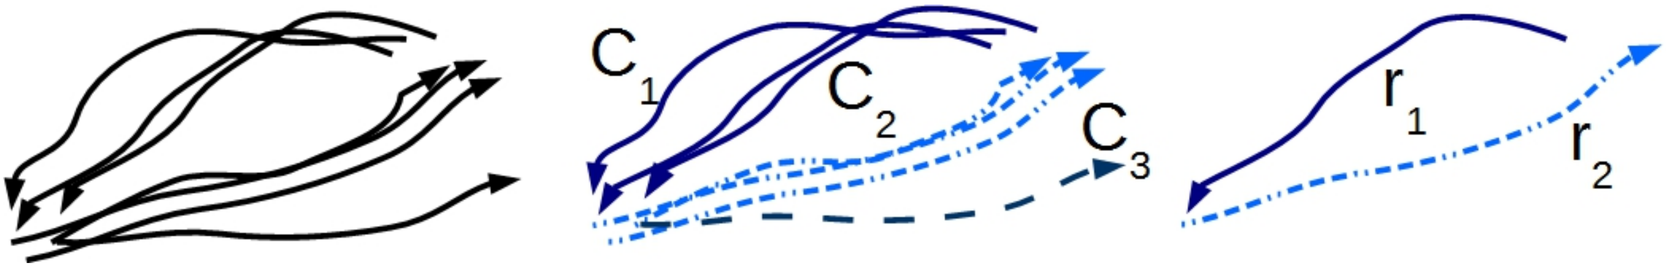

This model is called Mobility Profile.

Mining mobility user profiles for car pooling.

Trasarti, Pinelli, Nanni, Giannotti. KDD 2011

MyWay prediction in a nutshell

Use the mobility profile to predict the user's movements. If it is not able to produce a prediction, a collective predictor is used.
The collective predictor is built using the mobility profiles of the crowd.

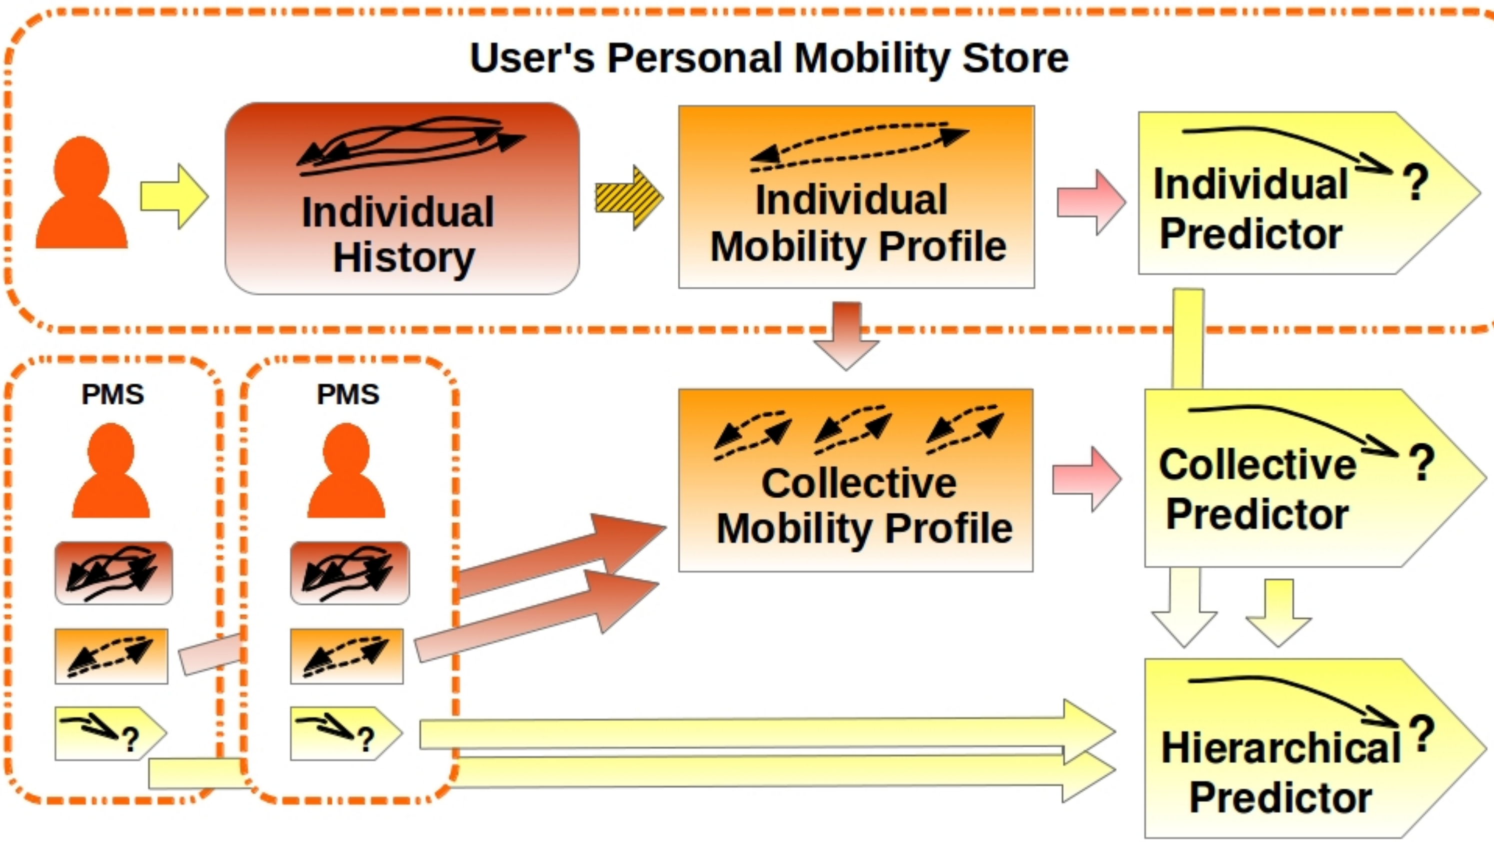

Experimental setting

Starting from a dataset of 1 month of movements, 5.000 users and 326.000 trajectories.

Divide the training set, i.e. 3 weeks and as test set the remaining last week.

The trajectories in the test set are cut to become the queries for the predictor.

- Cuts tested: first 33% or 66% of the trajectories.

Extracting the Mobility Profiles

Quality assessment of profiles

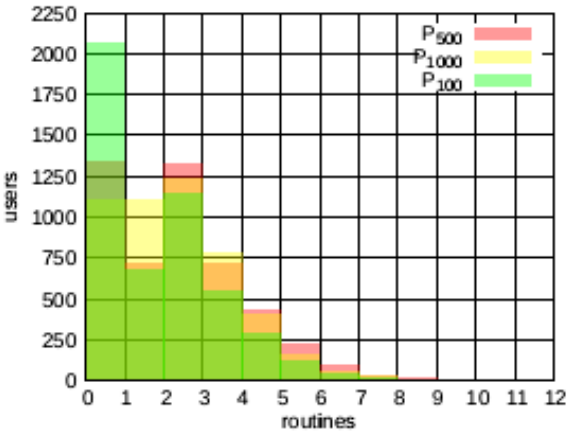

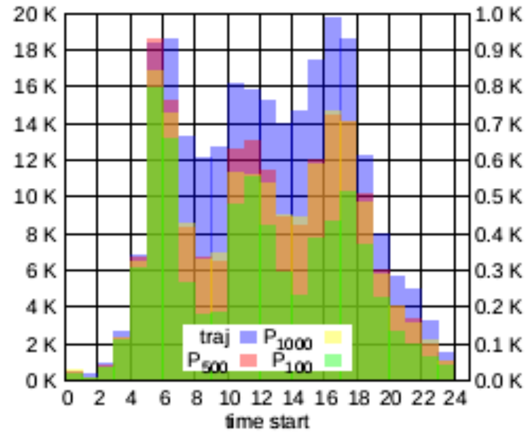

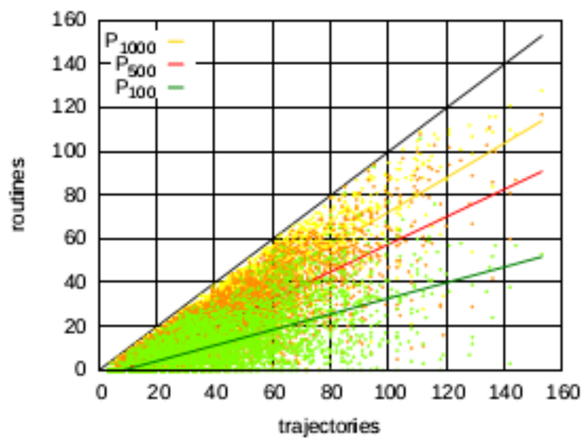

- Routines per user distribution (left)
- Trajectories and routines time start distribution (right)
- Dataset coverage (bottom)

Results

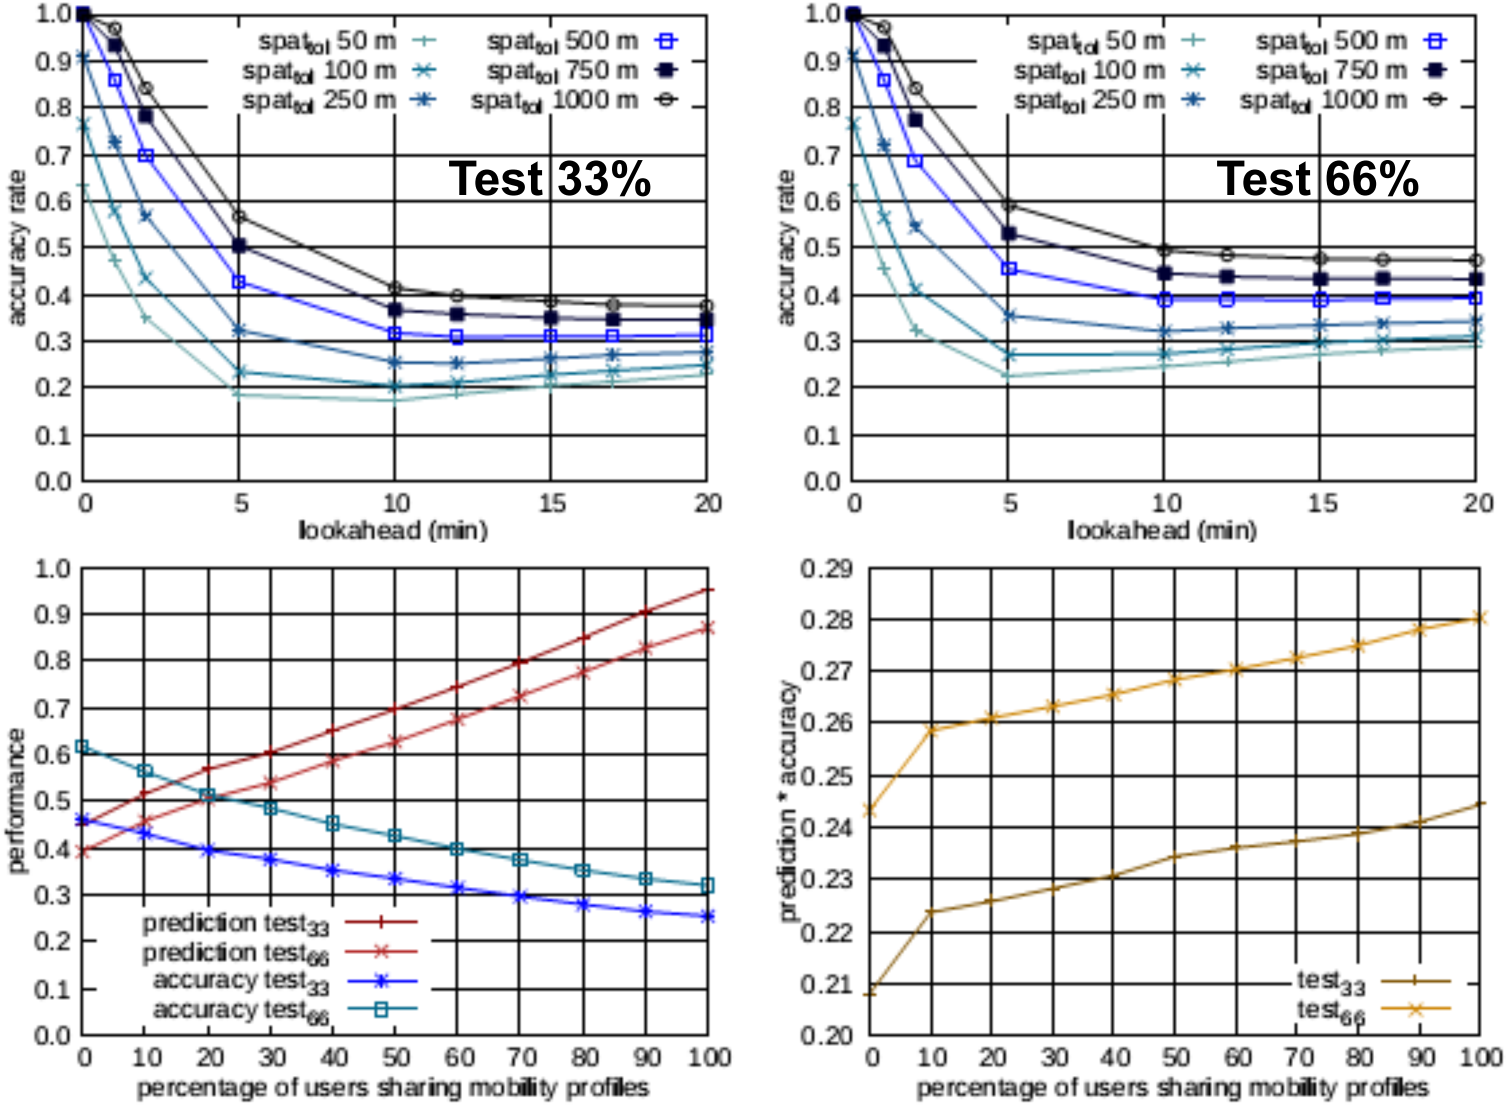

Test 33%
Test 66%
accuracy rate
lookahead (min)
spat_tol 50 m
spat_tol 100 m
spat_tol 250 m
spat_tol 500 m
spat_tol 750 m
spat_tol 1000 m
performance
percentage of users sharing mobility profiles
prediction test33
prediction test66
accuracy test33
accuracy test66
prediction * accuracy
test33
test66

Key publications

Returners and explorers dichotomy in human mobility. Luca Pappalardo, Filippo Simini, Salvatore Rinzivillo, Dino Pedreschi, Fosca Giannotti & Albert-László Barabási Nature Communications 6, Article number: 8166 (2015) doi:10.1038/ncomms9166 (2015),

Myway: Location prediction via mobility profiling . R Trasarti, R Guidotti, A Monreale, F Giannotti, Information Systems, 10, 2015

Small area model based estimators using Big Data Sources. Giusti, Marchetti, Pratesi, Salvati, Pedreschi, Giannotti, Rinzivillo, Pappalardo, Gabrielli.. Journal of Official Statistics, 31(2) 2015.

Unveiling mobility complexity through complex network analysis, R Guidotti, A Monreale, S Rinzivillo, D Pedreschi, F Giannotti, Social Network Analysis and Mining 6 (1), 59, 2016

Towards user-centric data management: individual mobility analytics for collective services. R Guidotti, R Trasarti, M Nanni, F Giannotti, Proceedings of the 4th ACM SIGSPATIAL , 2016

*An analytical framework to nowcast well-being using mobile phone dat*a. L Pappalardo, M Vanhoof, L Gabrielli, Z Smoreda, D Pedreschi, F. Giannotti, International Journal of Data Science and Analytics 2 (1-2), 75-92, 2017

Never drive alone: Boosting carpooling with network analysis. R Guidotti, M Nanni, S Rinzivillo, D Pedreschi, F Giannotti,. Information Systems 64, 237-257, 2017

Scalable and flexible clustering solutions for mobile phone-based population indicators, A Lulli, L Gabrielli, P Dazzi, M Dell'Amico, P Michiardi, M Nanni, L Ricci, International Journal of Data Science and Analytics 4 (4), 285-299

Short bibliography

- G. Andrienko, N. Andrienko, S. Rinzivillo, M. Nanni, D. Pedreschi, and F. Giannotti. Interactive visual clustering of large collections of trajectories. In Visual Analytics Science and Technology, 2009. VAST2009. IEEE Symposium on , pages 3{10, 2009.
- F. Giannotti, M. Nanni, D. Pedreschi, F. Pinelli, C. Renso, S. Rinzivillo, and R. Trasarti. Unveiling the complexity of human mobility by querying and mining massive trajectory data. VLDB J. , 20(5):695{719, 2011.
- Fosca Giannotti, Mirco Nanni, Fabio Pinelli, and Dino Pedreschi. Trajectory pattern mining. In KDD , 2007.
- M.Nanni, R.Trasarti, G.Rossetti, and D.Pedreschi. Ecient distributed computation of human mobility aggregates through user mobility proles. In UrbComp12 , 2012.

Mobility data: Modeling, Managing and understanding, Cambridge press.

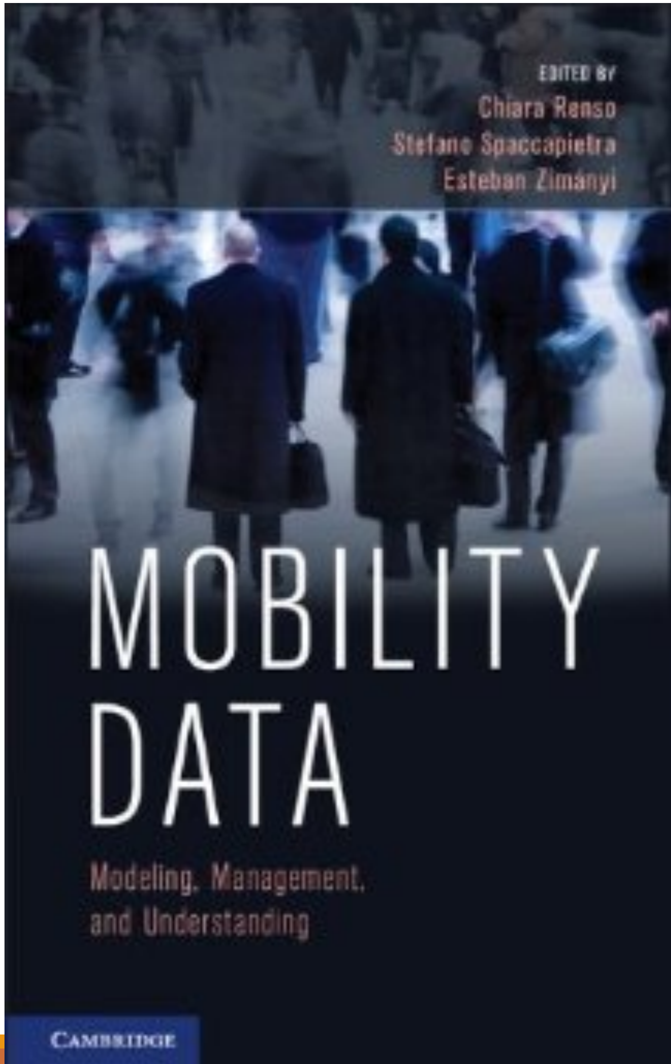

Giannotti · Pedreschi (Eds.)

Mobility, Data Mining and Privacy

The technologies of mobile communications and ubiquitous computing pervade our society, and wireless networks sense the movement of people and vehicles, generating large volumes of mobility data. This is a scenario of great opportunities and risks: on one side, mining this data can produce useful knowledge, supporting sustainable mobility and intelligent transportation systems; on the other side, individual privacy is at risk, as the mobility data contain sensitive personal information. A new multidisciplinary research area is emerging at this crossroads of mobility, data mining, and privacy.

This book assesses this research frontier from a computer science perspective, investigating the various scientific and technological issues, open problems, and roadmap. The editors manage a research project called GeoPKDD, Geographic Privacy-Aware Knowledge Discovery and Delivery, funded by the EU Commission and involving 40 researchers from 7 countries, and this book tightly integrates and relates their findings in 13 chapters covering all related subjects, including the concepts of movement data and knowledge discovery from movement data; privacy-aware geographic knowledge discovery; wireless network and next-generation mobile technologies; trajectory data models, systems and warehouses; privacy and security aspects of technologies and related regulations; querying, mining and reasoning on spatiotemporal data; and visual analytics methods for movement data.

This book will benefit researchers and practitioners in the related areas of computer science, geography, social science, statistics, law, telecommunications and transportation engineering.

ISBN 978-3-540-75176-2

› springer.com

Giannotti
Pedreschi (Eds.)

Mobility, Data Mining
and Privacy

Fosca Giannotti
Dino Pedreschi (Eds.)

Mobility, Data Mining and Privacy

Geographic Knowledge Discovery

Springer