RUTGERS

CS 671: Data-Driven Cyber-Physical Systems for Smart Cities

Desheng Zhang

Computer Science
Rutgers University

Spring 2018

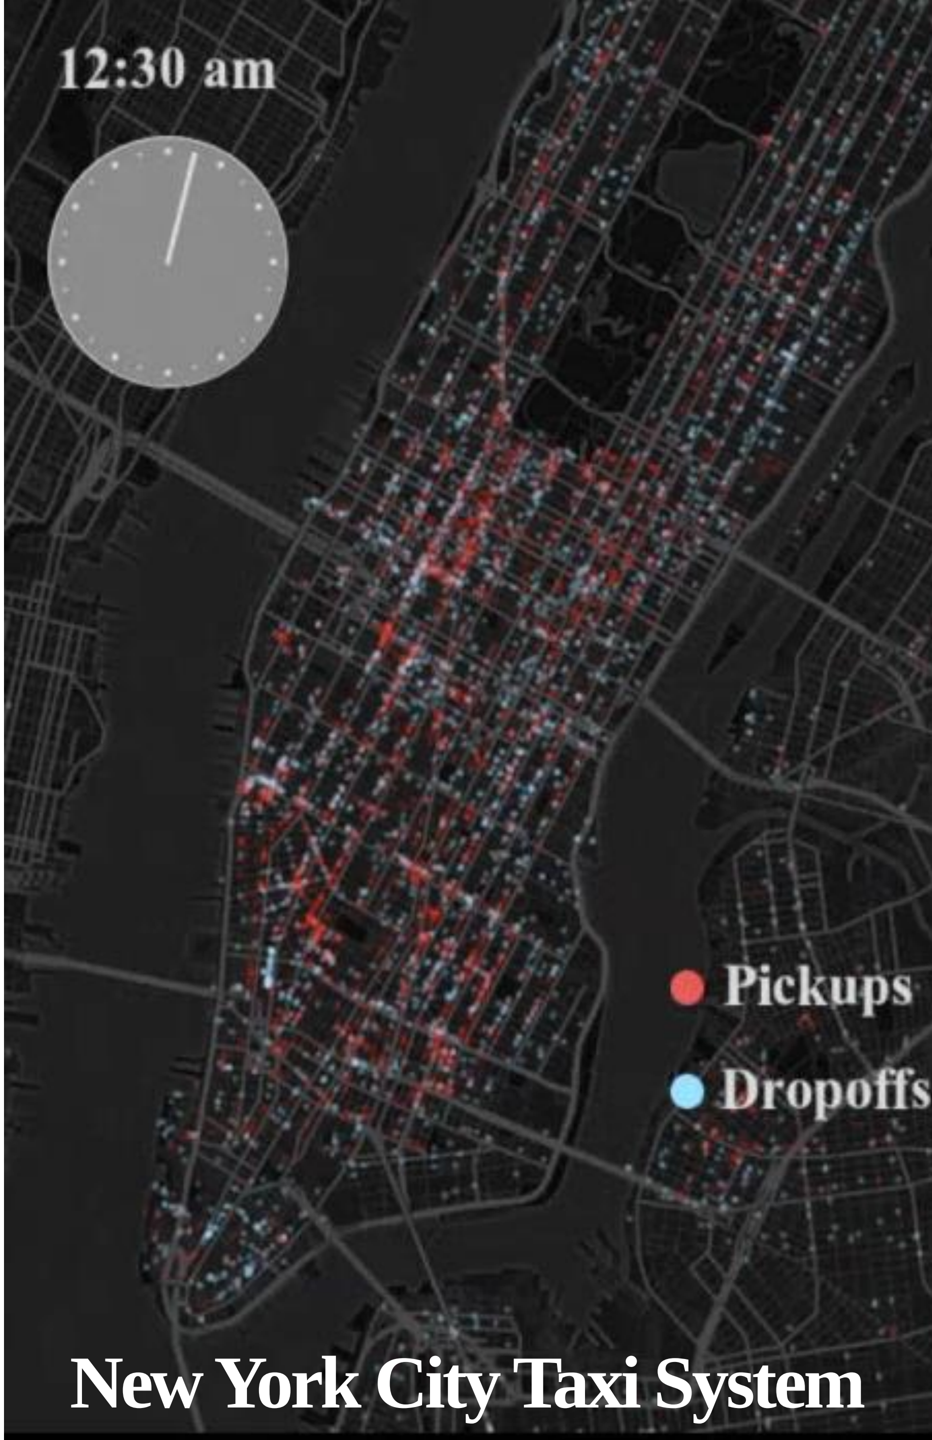

Who Am I

- Desheng Zhang
  - Assistant Professor
  - Department of Computer Science
  - Rutgers University
  - https://www.cs.rutgers.edu/~dz220/
  - Office: CoRE 307
  - Phone: 848-445-8307
  - Email: d.z@rutgers.edu (including CS671 in your subject!)

张德升

My background

China

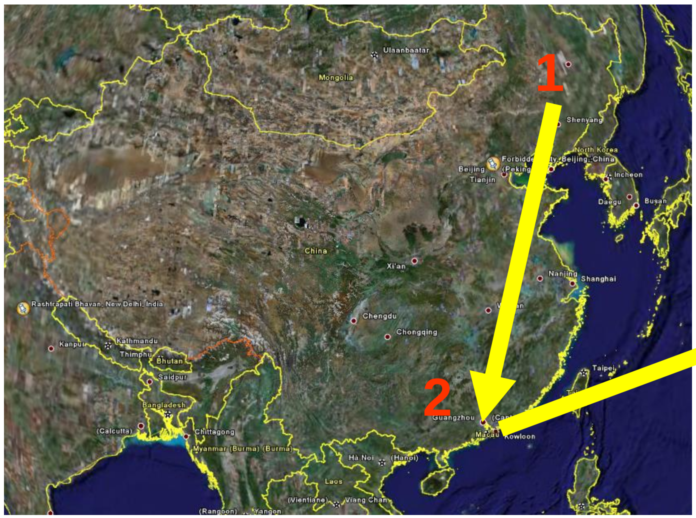

USA

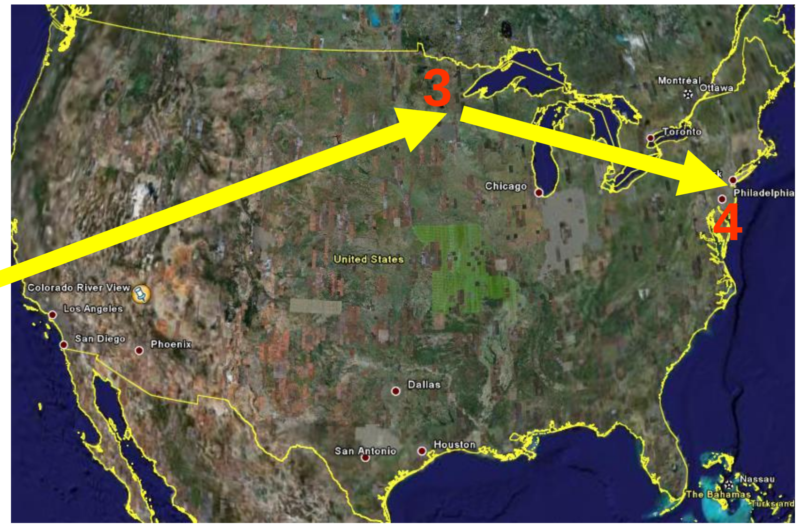

1. Bachelor & Master in CS at Heilongjiang University
2. Visiting Student at Shenzhen Institute of Advance Technology
3. Ph.D in Computer Science, University of Minnesota
4. Assistant Professor, Rutgers University

More about me

So, if I am not answering your email during a weekend, I am most likely…

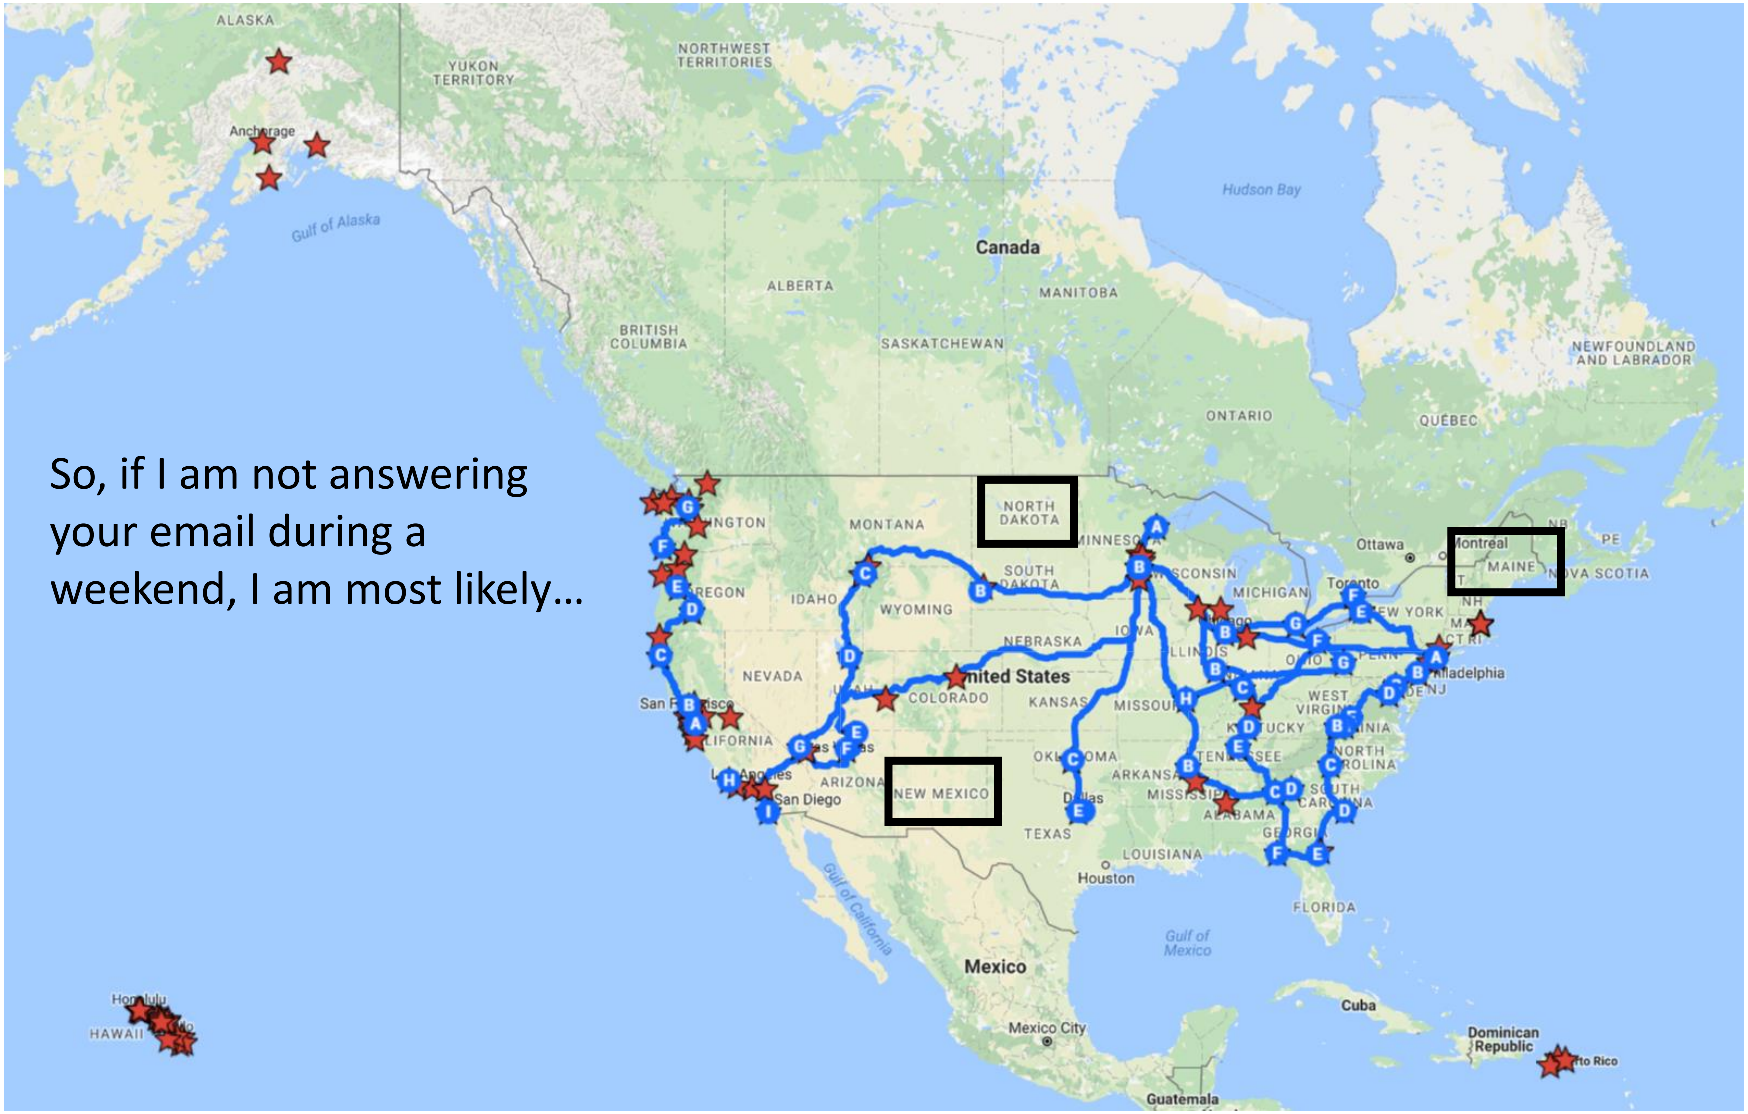

About you

- Talk to you neighbors
  - Introduce yourself to him/her for two minutes
  - (maybe you could find your project partner)

- Introduce yourself to the class for 30 seconds
  - Undergrad Major
  - Degree Seeking
  - Research interests

Outline

- **Logistics**
- Course Structure
- Overview of Smart Cities Research
- Class Schedule

Class Website

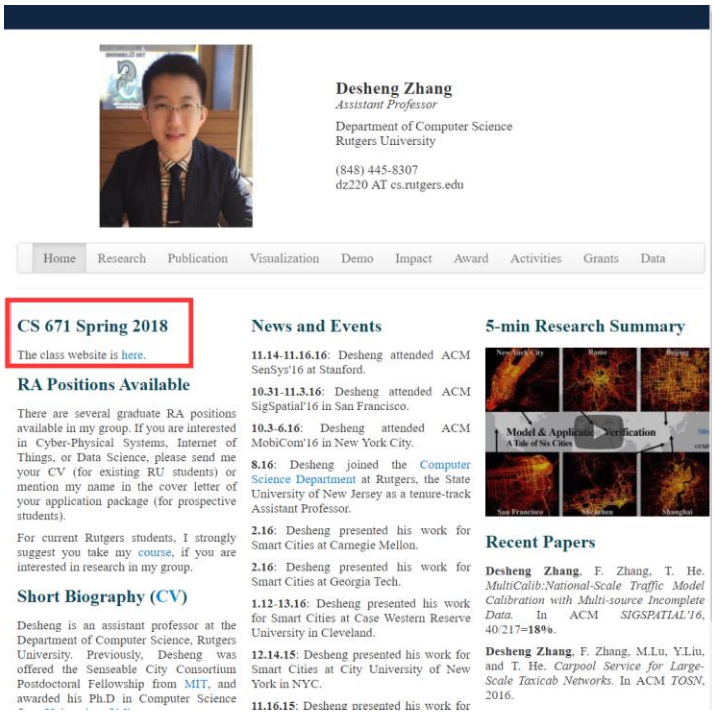

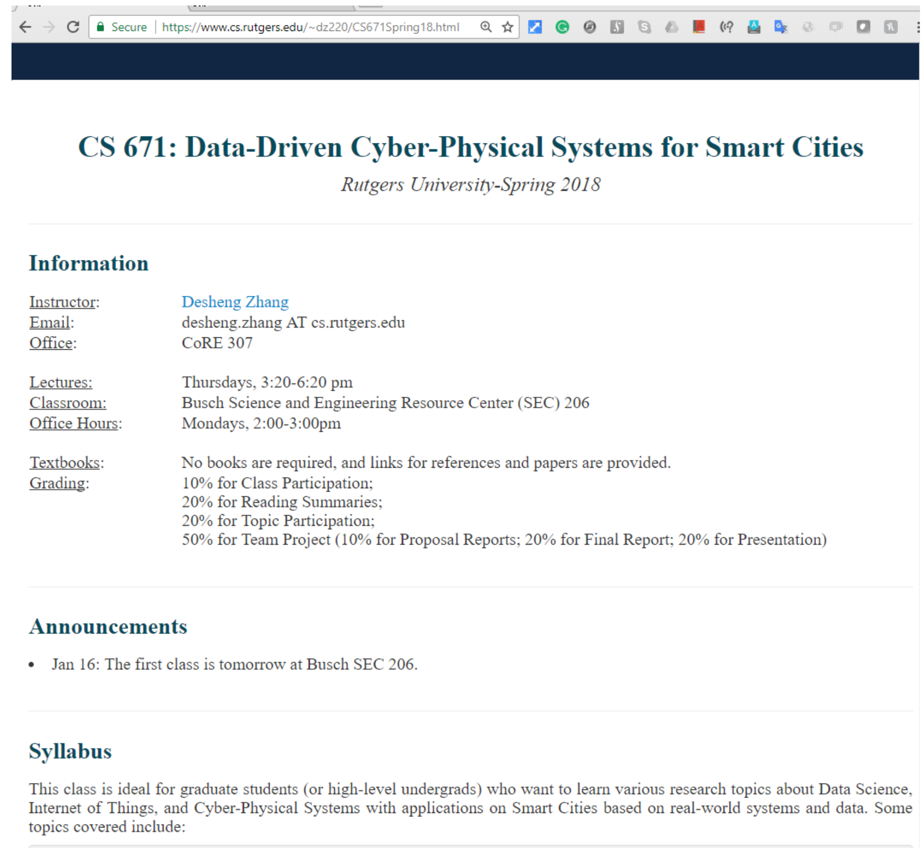

- On Sakai
- https://www.cs.rutgers.edu/~dz220/CS671Spring18.html

About this class

- Lectures:
  - Thursdays: 3:20-6:20 pm
- Location:
  - Science and Engineering Resource Center (SEC) 206
- Office Hours
  - Mon 2:00-3:00pm by appointments at CoRE 307

About this class

- Advanced Course on Smart Cities:
  - Reading
  - Presentations
  - Team Project

- Prerequisites:
  - Preliminary math knowledge
    - Calculus, Linear Algebra, & Probability
  - Skills for high-level programing languages are required.
    - C++, Java, R, Python or SAS

- No Textbooks are required:
  - links are provided

About this class

- Good for students who want to
  - Do Data-Driven **Research** on Data Science, Smart Cities, IoT
  - Have hands-on experiences about real data-driven projects
  - Improve their writing/presentation skills
- Not so good for students who want to
  - Practice Coding Skills
  - Do Hardware-related Projects

About this class

- Master Students:
  - This class **CAN** be used to satisfy **B requirements**
  - But each MSCS student should stop by the MSCS office (Hill 355) to make sure that his/her study plan is approved as per the graduation requirements.

- PhD Students:
  - This 67x class **CANNOT** be used to satisfy **B requirements**
  - Only Core Class CS 5xx can be used
  - Details on the CS website.

Outline

- Logistics
- **Course Structure**
- Overview of Smart Cities Research
- Class Schedule

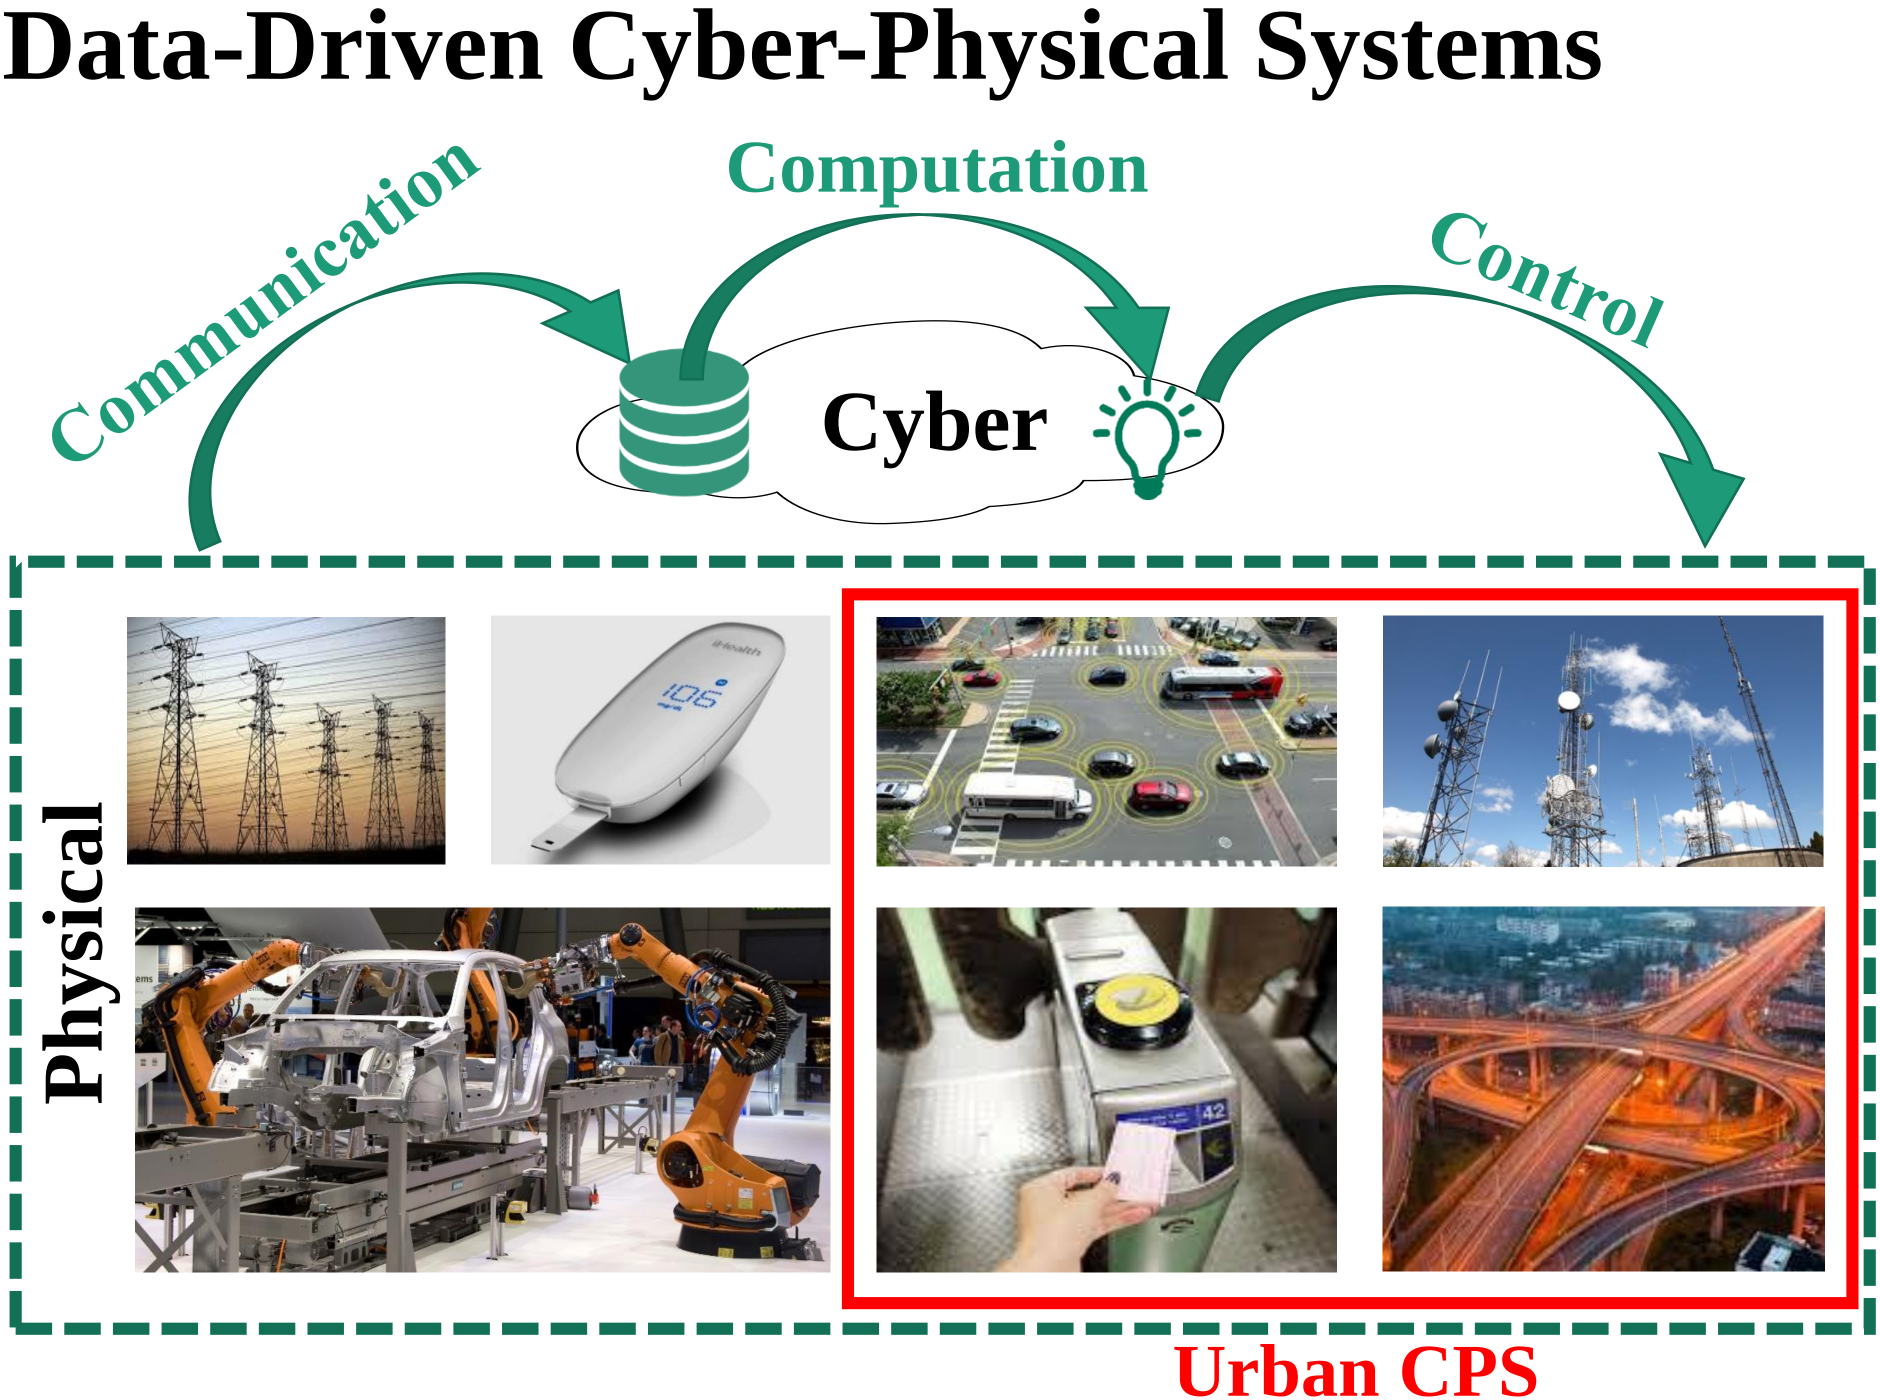
Data-Driven Cyber-Physical Systems
Communication
Computation
Control
Cyber
Physical
Urban CPS

Urban Sustainability

The White House unveils new US$160 million Smart Cities Initiative

17th September 2015 Tom Teodorczuk

Urban Systems

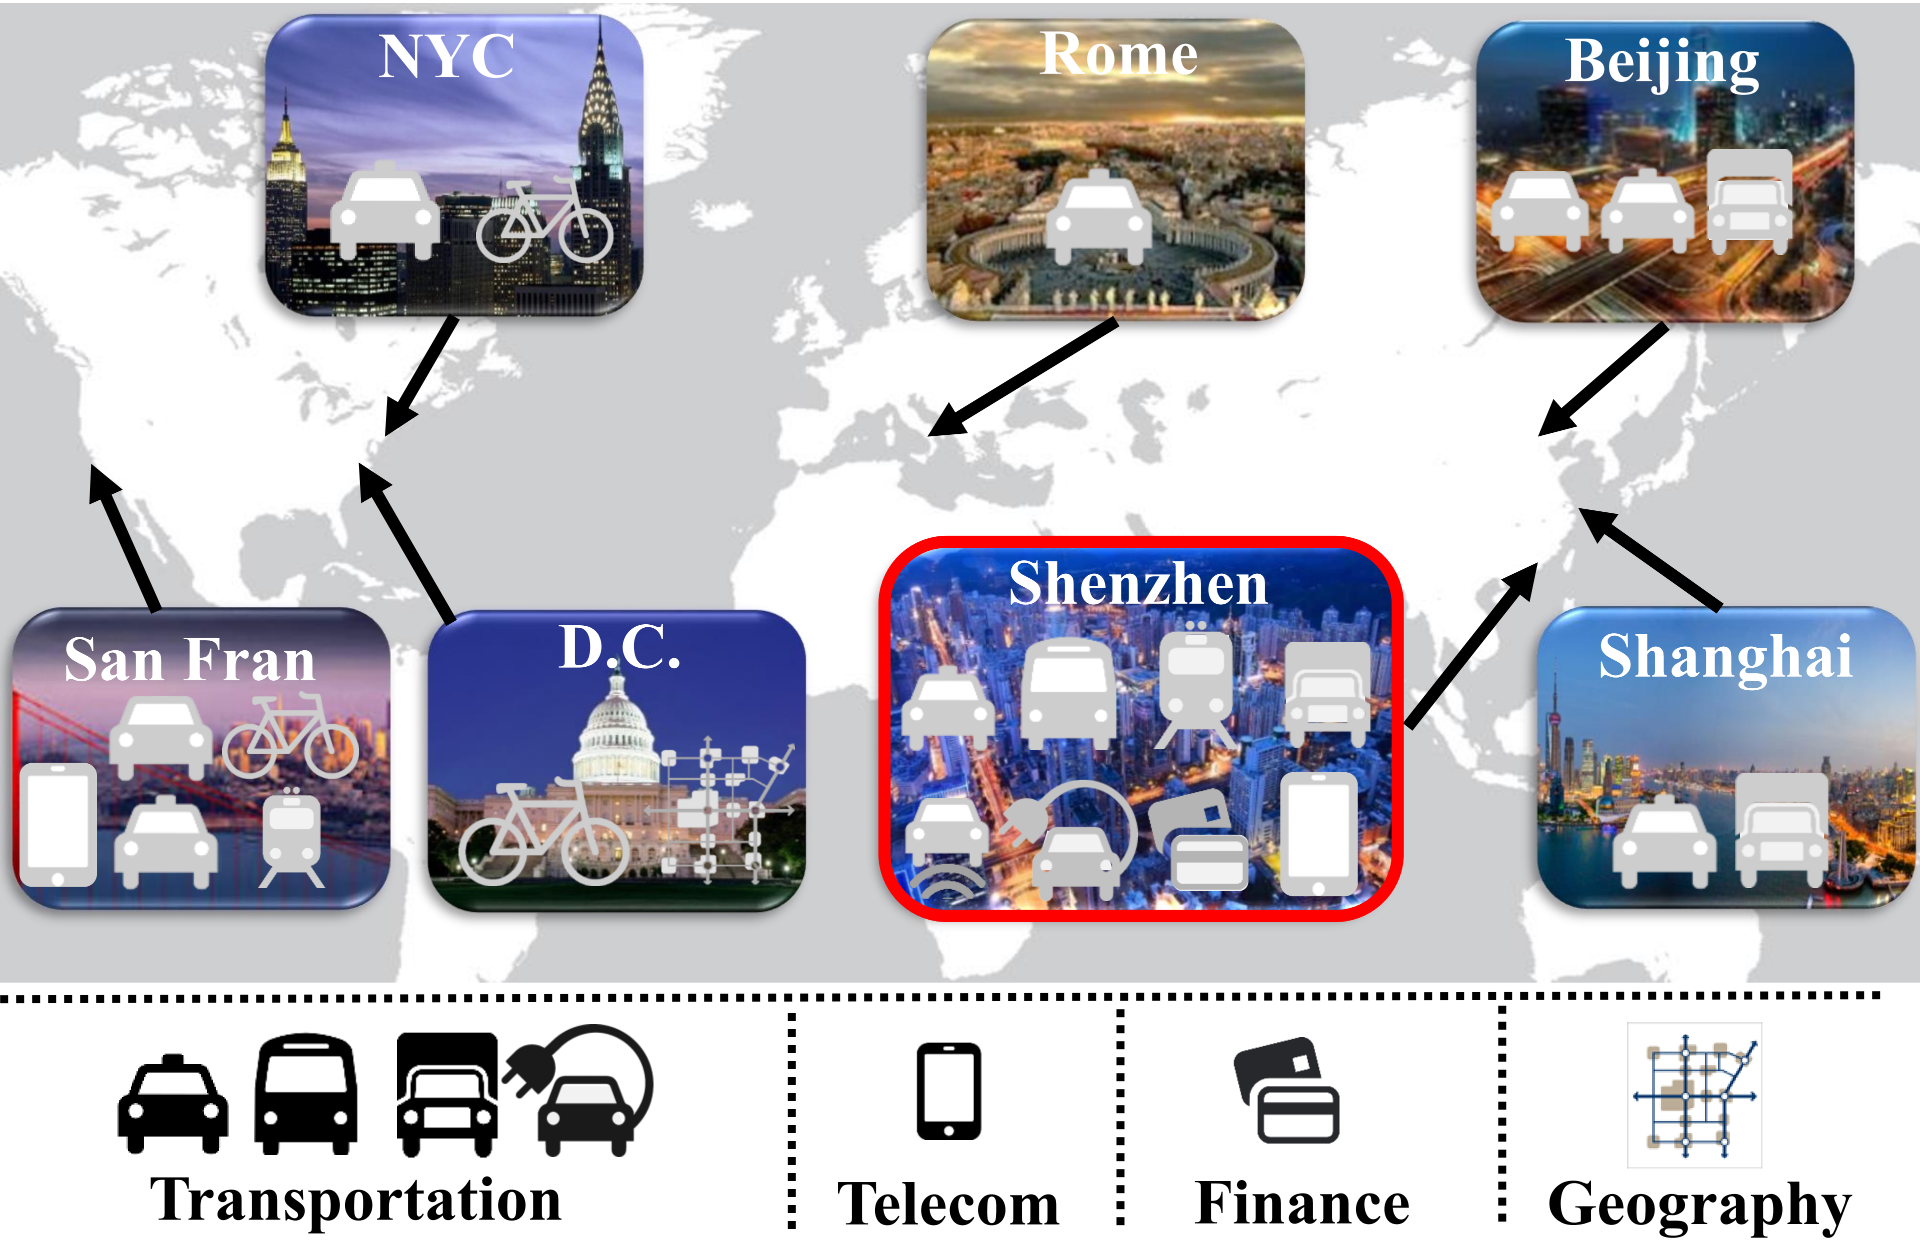

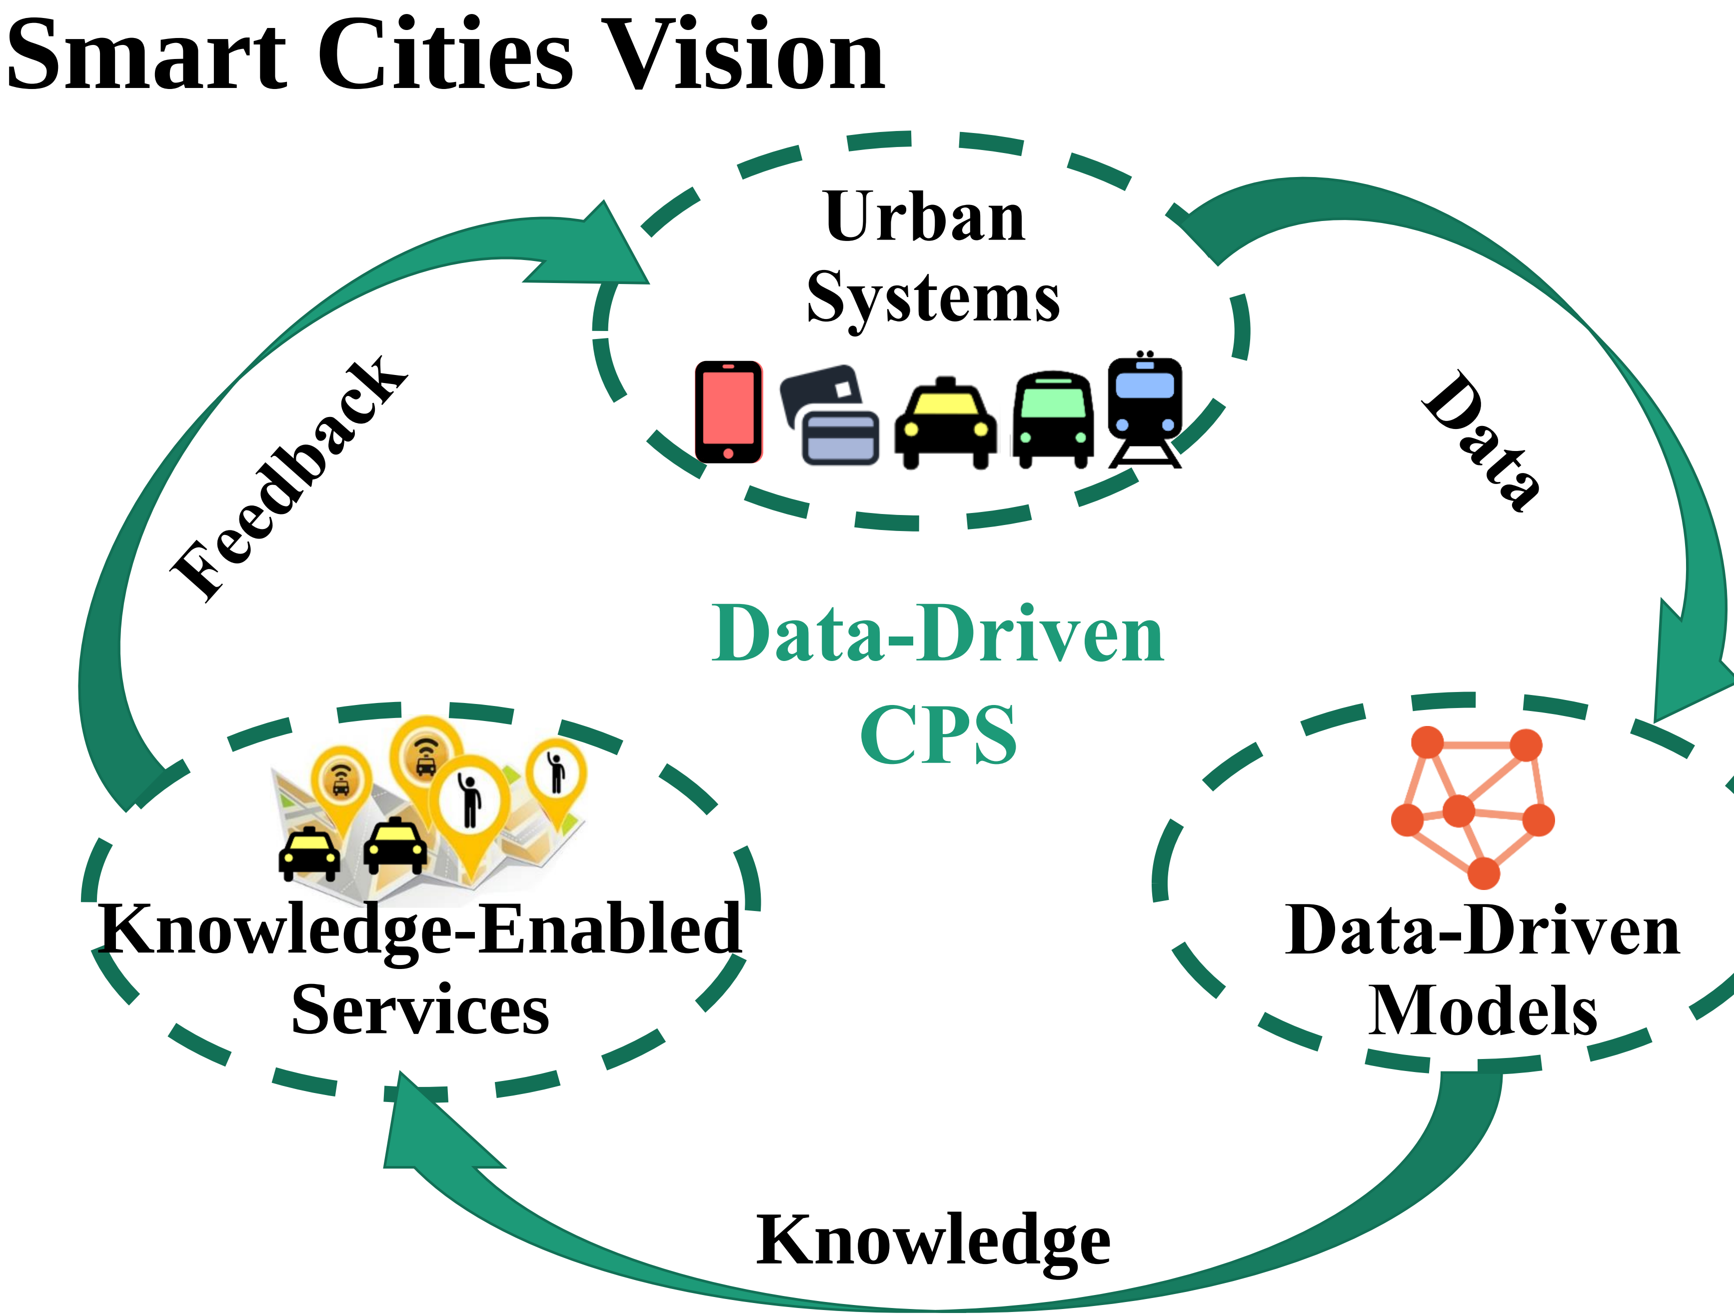
Smart Cities Vision
Urban Systems
Data
Feedback
Data-Driven CPS
Knowledge-Enabled Services
Data-Driven Models
Knowledge

Topics for 14 wks

- **Introduction (1 wk)**
- **Basic Topics (7 wks)**
  - Urban Sensing (1 wk)
  - Data Manage.&Processing (1 wk)
  - Data-Drive Modeling (3 wk)
  - Data Visualization (1 wk)
  - Data Predictive Control (1 wk)
- **Project Midterm Presentation**
  - 9th wk
- **Special Topics (4 wks)**
  - Novel Services (1 wk)
  - Dependency Analyses (1 wk)
  - Human-in-the-loop (1 wk)
  - Privacy and Security (1 wk)
- **Project Final Presentation**
  - 14th wk

Tentative Schedule

Week	Date	Topics and Reading Assignments
1	Jan 19	General Class Introduction Reading: • Urban Computing: Concepts, Methodologies, and Applications • Cyber-Physical Systems: Executive Summary
2	Jan 26	
3	Feb 2	
4	Feb 9	
5	Feb 16	
6	Feb 23	
7	Mar 2	
8	Mar 9	
Spring Break		
9	Mar 2[illegible]	Project Proposal Presentation No Reading Assignment
10	Mar 30	
11	Apr 6	
12	Apr 13	
13	Apr 20	
14	Apr 27	Final Project Presentation No Reading Assignment
	May 4	Final project papers are due on May 4th 11:59PM EST.

Intro (1 wk)

Basic Topics (7 wks)

Project Proposal Presentation (1 wk)

Special Topics (4 wks)

Project Final Presentation (1 wk)

Grades

- No Exams

1. 10% for Class Participation
2. 20% for Reading Summaries
3. 20% for Topic Presentation

Individual Based 50 %

1. 10% for Project Proposal Report
2. 20% for Project Final Report
3. 20% for Proposal or Final Presentation

Team Based 50%

Individual 1: Class Participation (10%)

- Coming to all the lectures
- Actively asking questions
- Participating discussions

Individual 2: Reading Summary (20%)

- Submitting 10 Reading Summaries
- Choosing 10 papers from assigned papers
    - Covering 10 out of 11 topics
    - 1 paper per topic
    - i.e., skipping one week
- A suggested format for summaries

Individual 2: Reading Summary: (20%)

- What is the key scientific question?
  - Why it is **hard or important** to solve this question?
- What are existing approaches?
  - Why they are not sufficient?
- What the authors’ idea?
  - Why this is new?
  - Better than existing ones?
  - How they evaluate it?

- What are the strengths of this paper?
- What are the weaknesses of this paper?
- How you can address these flaws?

Individual 3: Topic Presentation (20%)

- Every student will be assigned with a topic
  - A paper (a set of papers) from Assigned Reading List
  - Your own project is related to this topic

- Time 40 mins: 35 mins Talk & 5 mins Q&A

- Max 3 Presentations for Every Basic and Advanced Lecture

- Will send a poll to get five preferences of all students

- Coming to office hours before topic presentations

Team-based Project (50%)

- 10 Teams
  - 2-3 students per team
  - Assign based on topic interests (poll)
  - Form your own group if you really prefer

- A new team project related to smart cities
  - Common interests for all team members
  - Suitable scope for one semester
  - Be able to find data about it
  - Be able to evaluate it

Team-based Project (50%)

- Proposal Presentation (9^{th} wk)
  - 18 mins per team
  - Presented by first 1or 2 team members

10%

- Proposal Report (9^{th} wk)
  - 4 page-long double-column by ALL members

20%

- Final Presentation (14^{th} wk)
  - 18 mins per team
  - Presented by the rest of team members

20%

- Final Report (15^{th} wk)
  - 8 page-long double-column by ALL teams

Before Any Presentation

- Any Presentation
  - Topic, Proposal, or Final
- Go to office hours:
  - Mon: 2-3pm: Desheng: CoRE 307
- Discussions
  - Background
  - Scopes
  - Technical contents

Bonus!

- Data-Driven Visualization
- Data-Driven Evaluation
- Demo: Animation or Video
- Present a set of Papers in Topic Presentation
- Present Your Own Work in Topic Presentation

Project Contest

- Vote by all students based on
  - Proposal Presentation
  - Final Presentation
- Three Teams will win prizes
- Announce after the final presentation

Suggestions about the class

- Suggestions are welcome
- Reading List
- Topics Interested
- Team Formation

Summary

- Come to all topic presentations (10%)
- Submit 20 summaries (20%)
- A Topic Presentation (20%)

Individual Based 50 %

- A team-based project proposal report (10%)
- A team-based project final report (20%)
- A team-based project presentation (20%)

Team Based 50%

Useful References

- My website
- https://www.cs.rutgers.edu/~dz220/

- Mining of Massive Datasets - Stanford InfoLab
- http://infolab.stanford.edu/~ullman/mmds/book.pdf

- Urban Computing at Microsoft Research
- https://www.microsoft.com/en-us/research/project/urban-computing/

Questions?

Outline

- Logistics
- Course Structure
- **Overview of Smart Cities Research**
- Class Schedule

Cross-Domain Cyber-Physical Systems for Smart Cities

Desheng Zhang

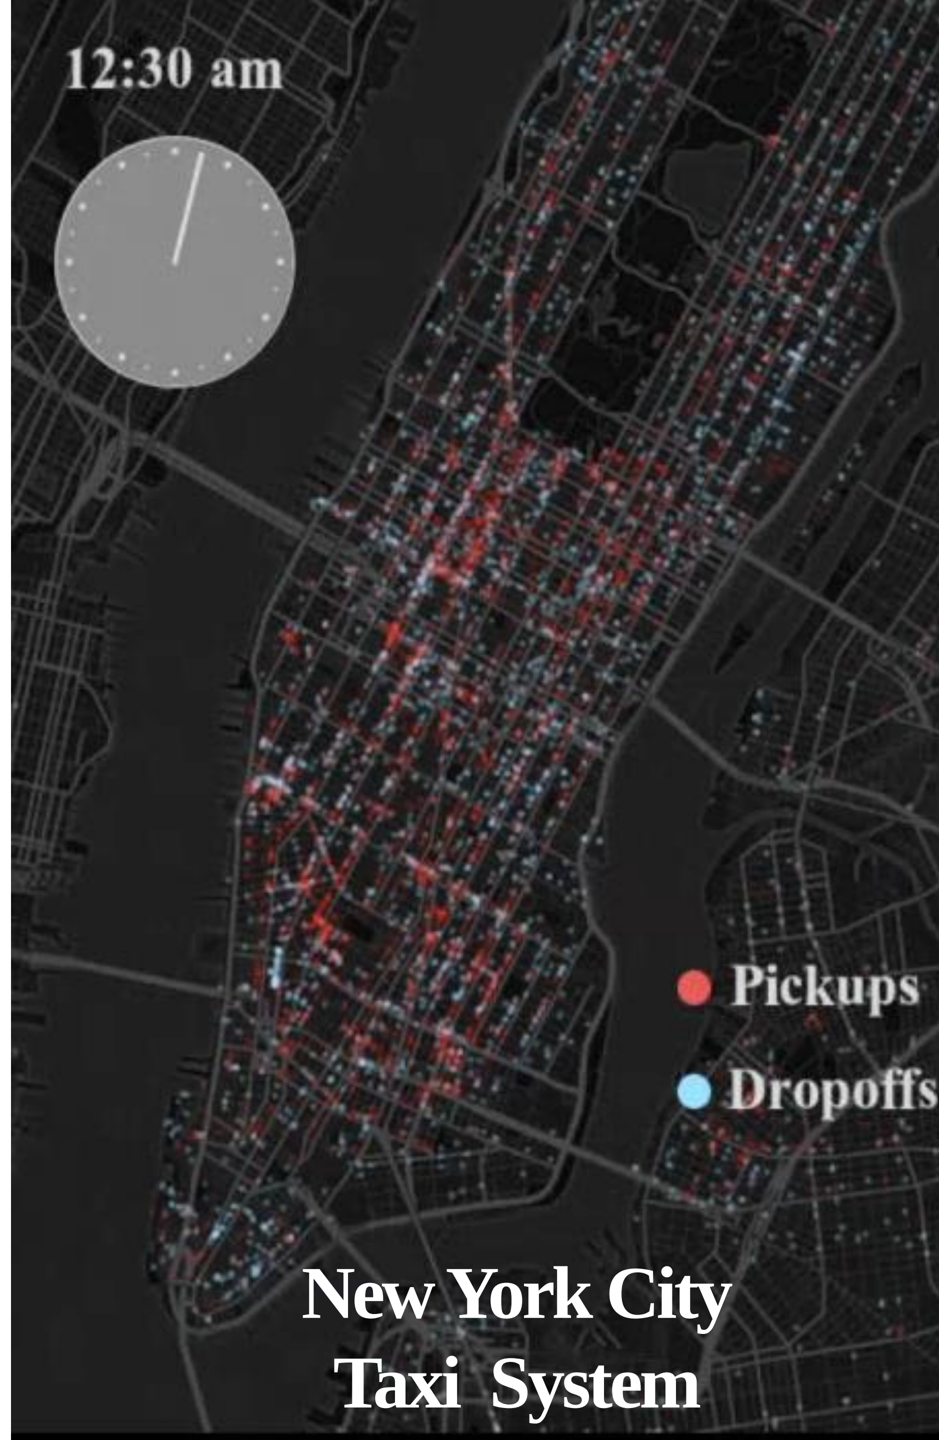

Cyber-Physical Systems

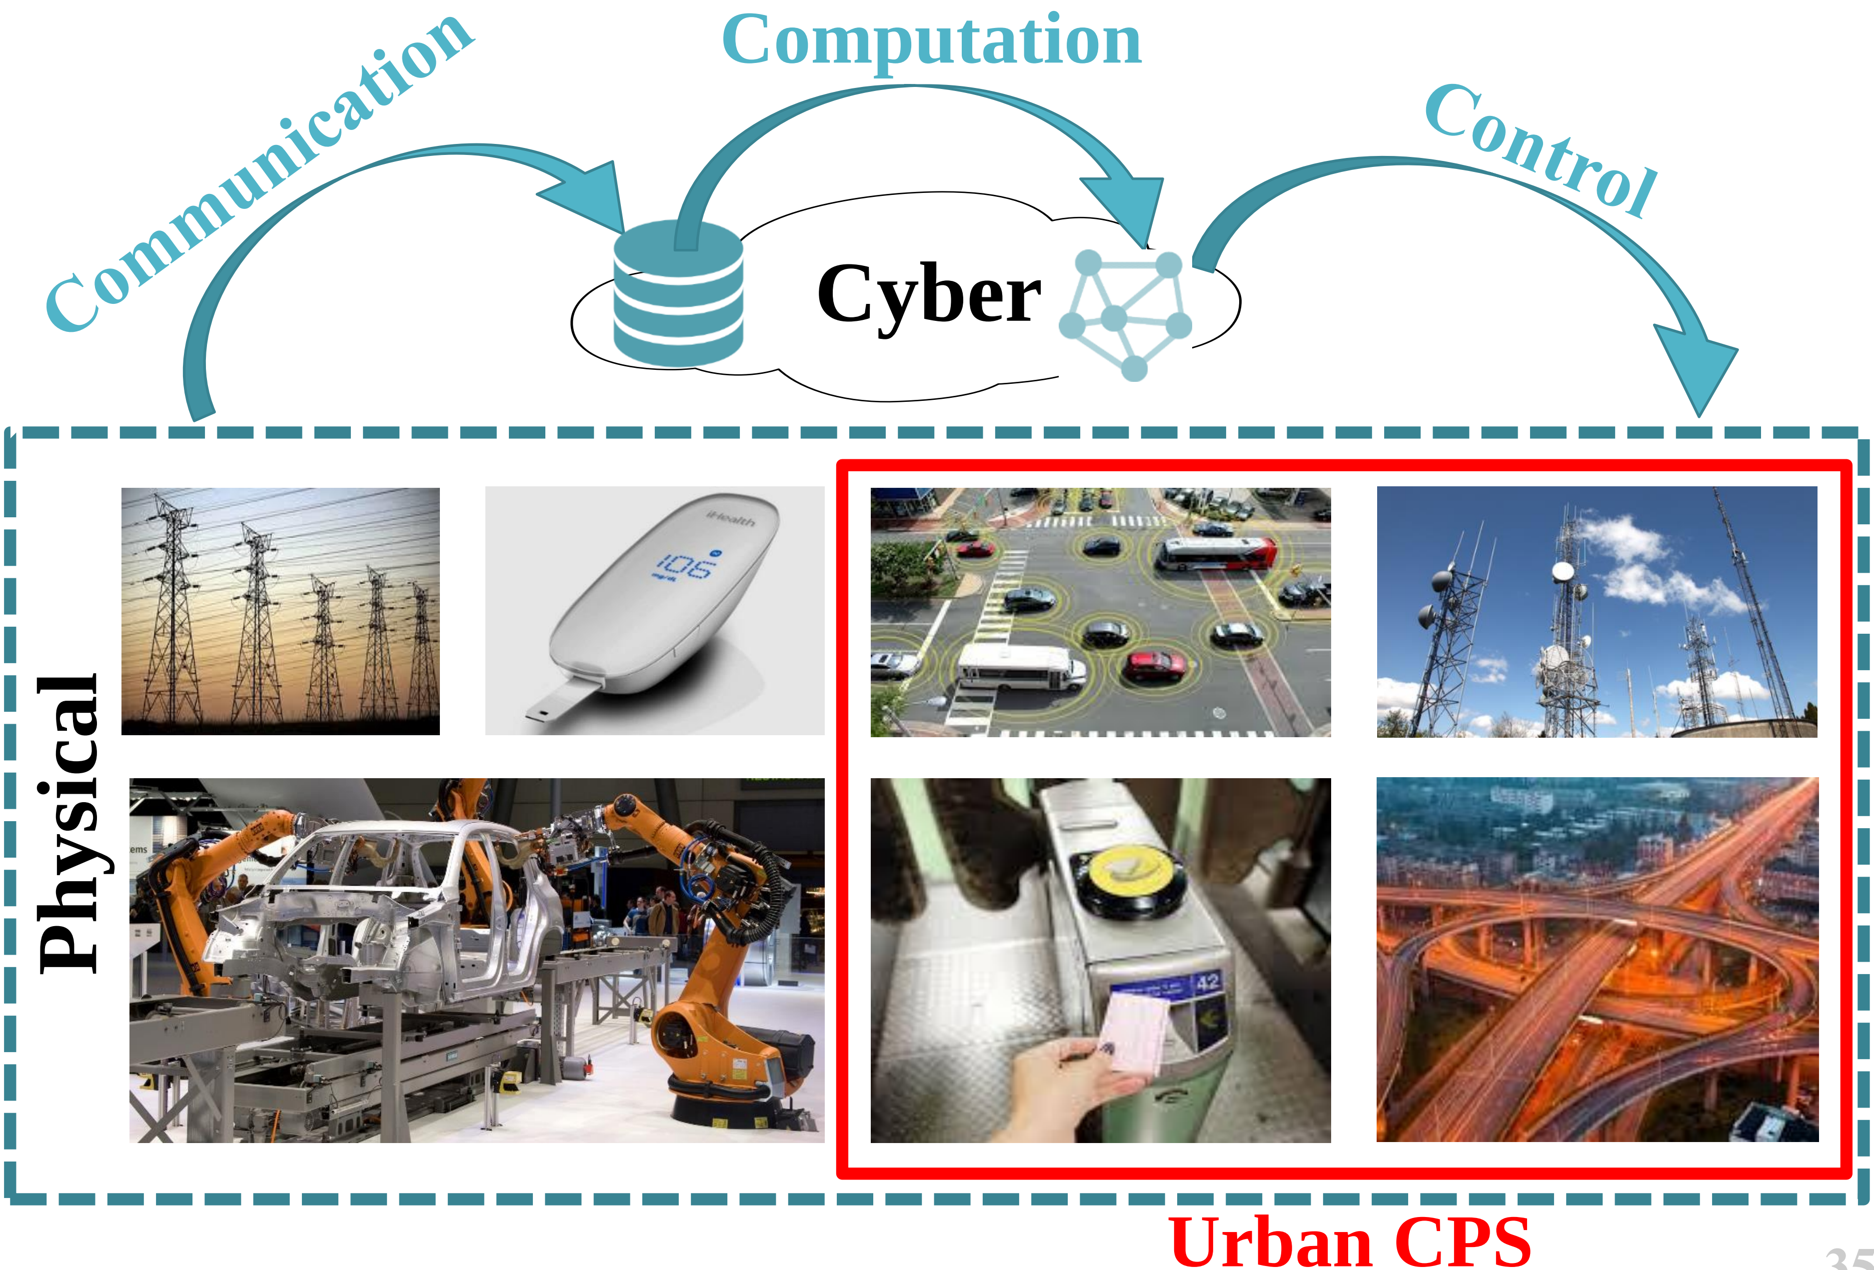

Why Urban?

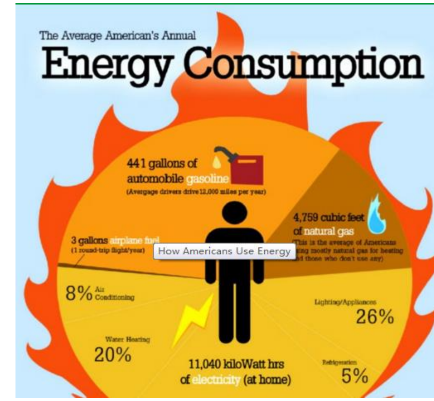

The White House unveils new US$160 million Smart Cities Initiative

17th September 2015 Tom Teodorczuk

- **$35M** for Smart Cities Grants by **NSF**
  - **$10M** for **CPS** in 2016

- **$70M** for Transportation and Energy, by
  - DoT, DoE, NIST…

Urban Systems

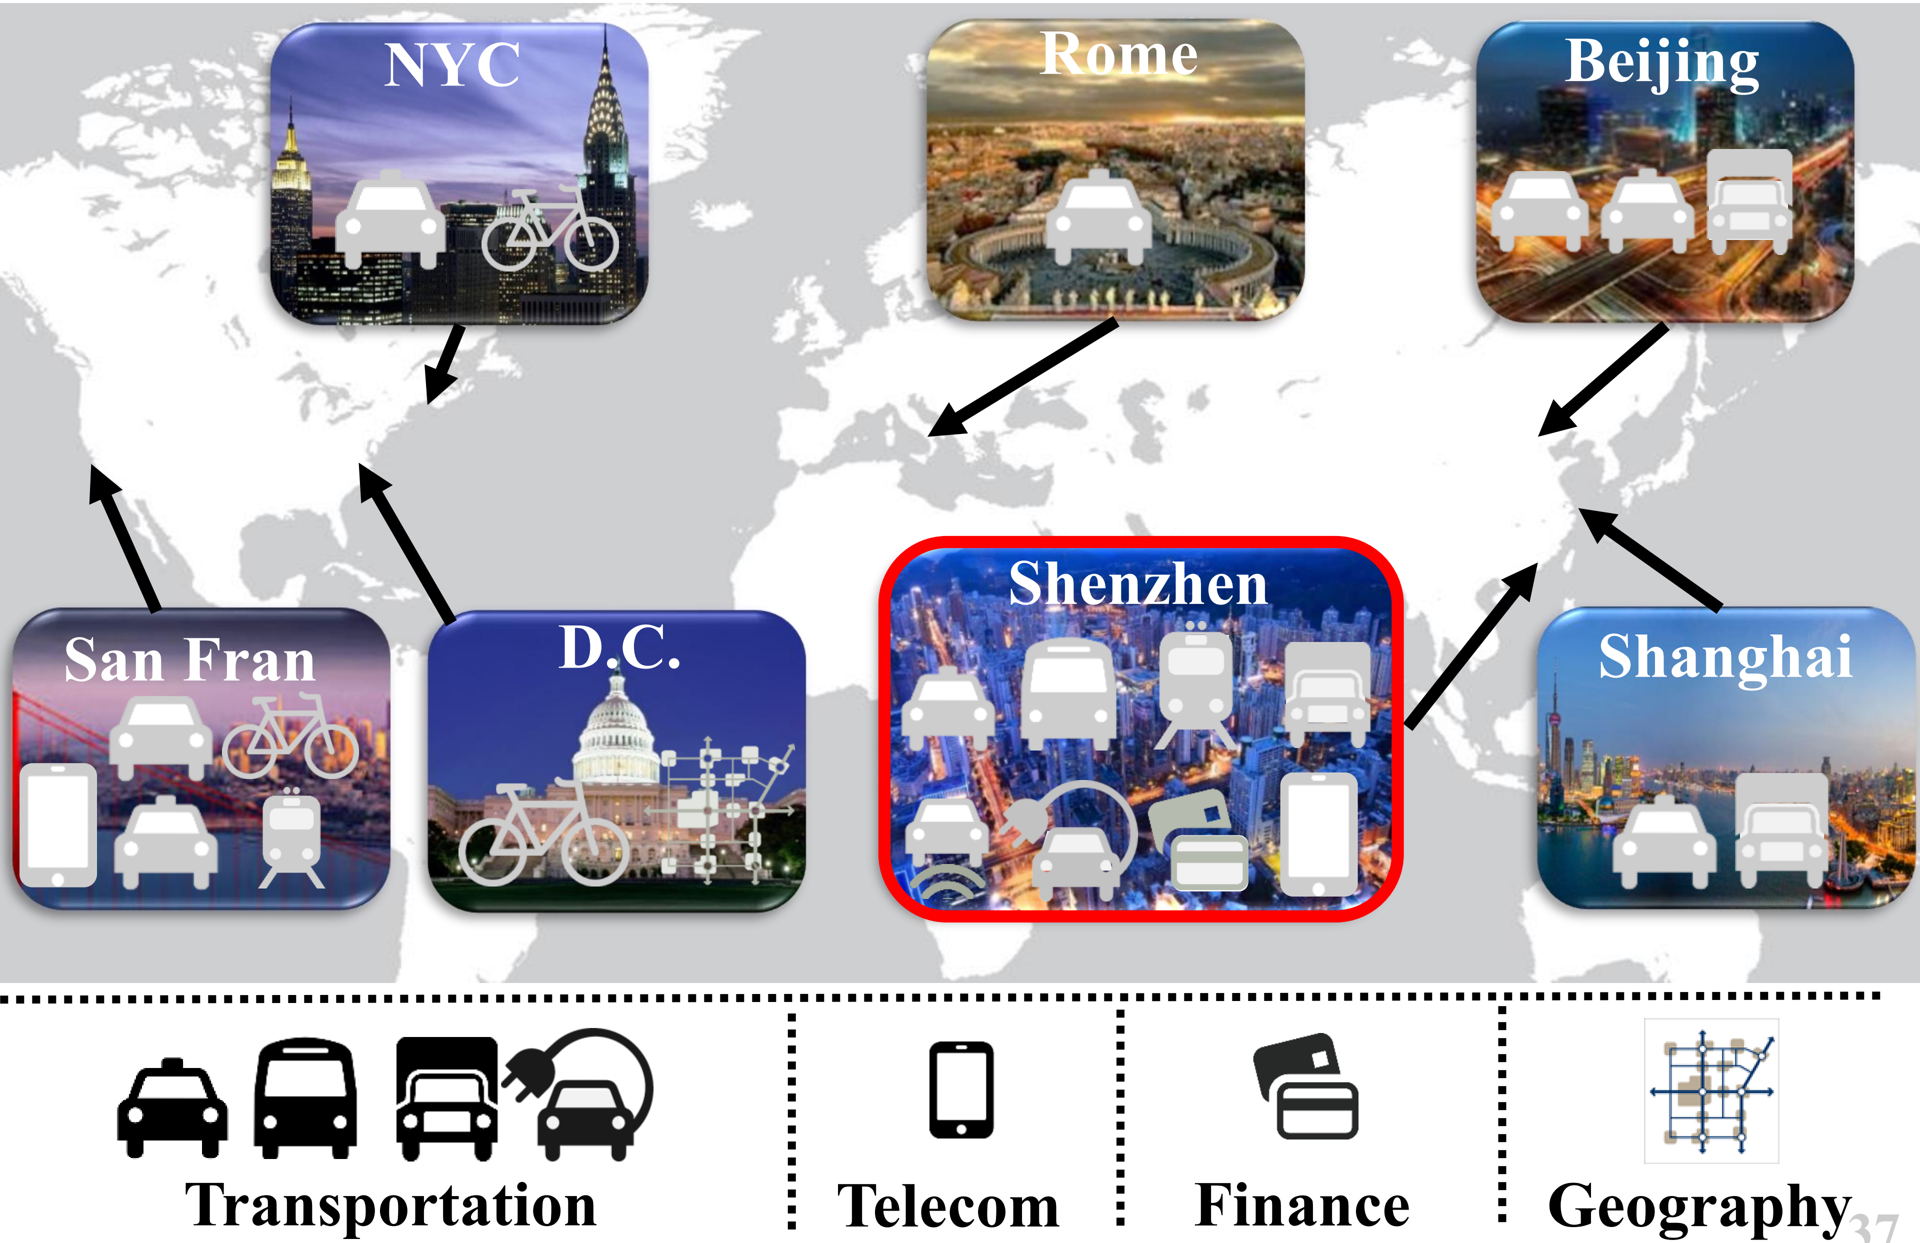

Smart Cities Vision

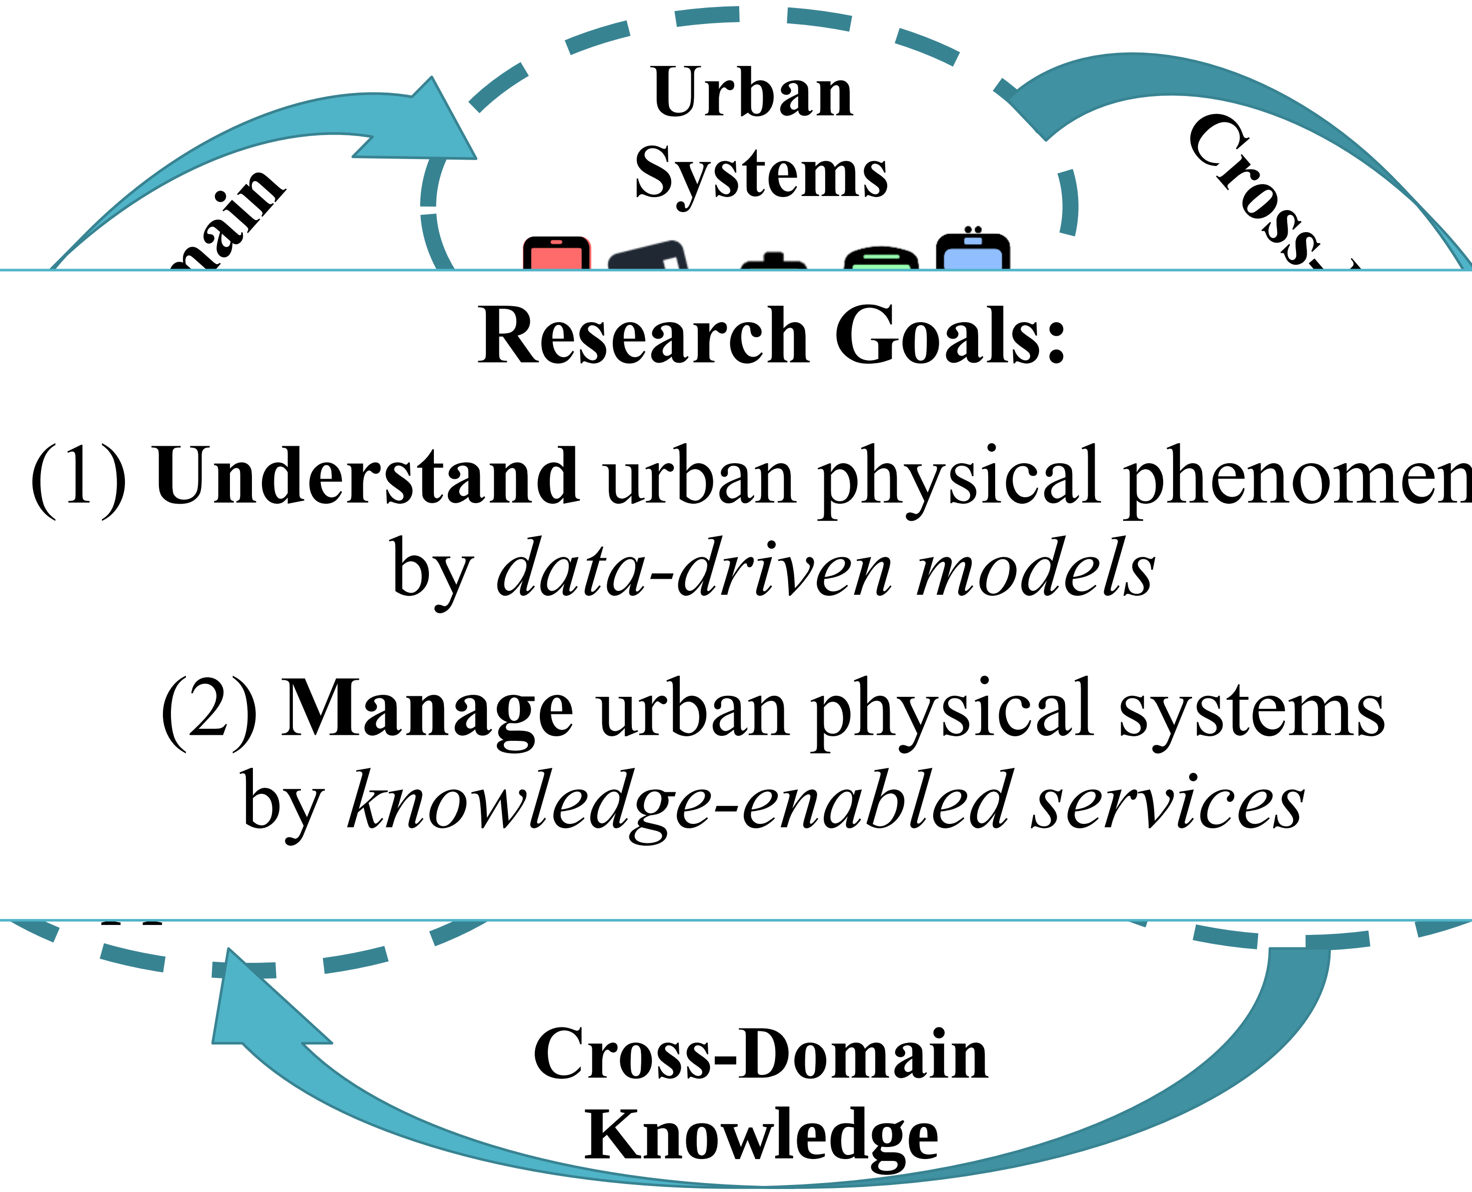

Research Goals:

(1) **Understand** urban physical phenomena
by *data-driven models*

(2) **Manage** urban physical systems
by *knowledge-enabled services*

An Overview

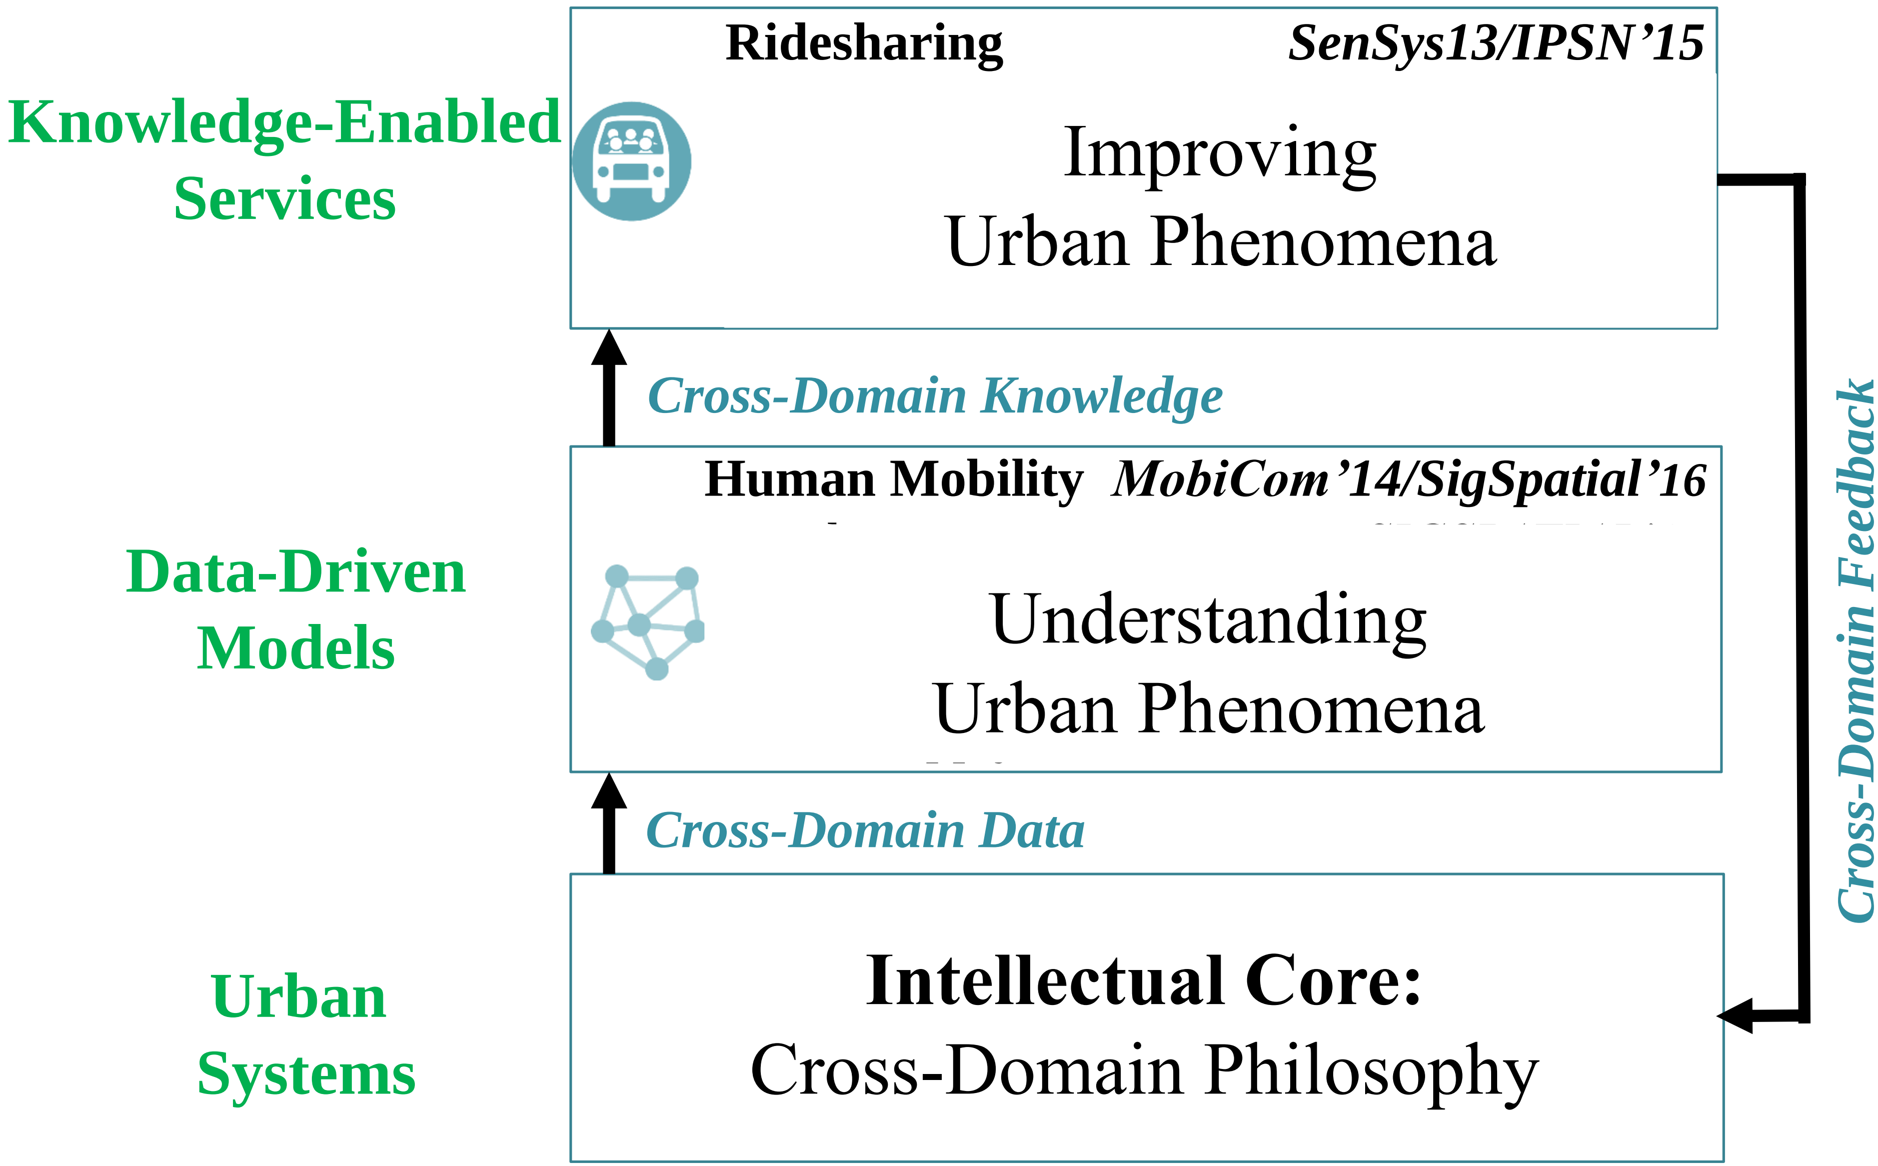

Knowledge-Enabled Services
Ridesharing
SenSys13/IPSN'15
Improving Urban Phenomena
Cross-Domain Knowledge
Data-Driven Models
Human Mobility
MobiCom'14/SigSpatial'16
Understanding Urban Phenomena
Cross-Domain Data
Urban Systems
Intellectual Core:
Cross-Domain Philosophy
Cross-Domain Feedback

Goal

Modeling Human Mobility

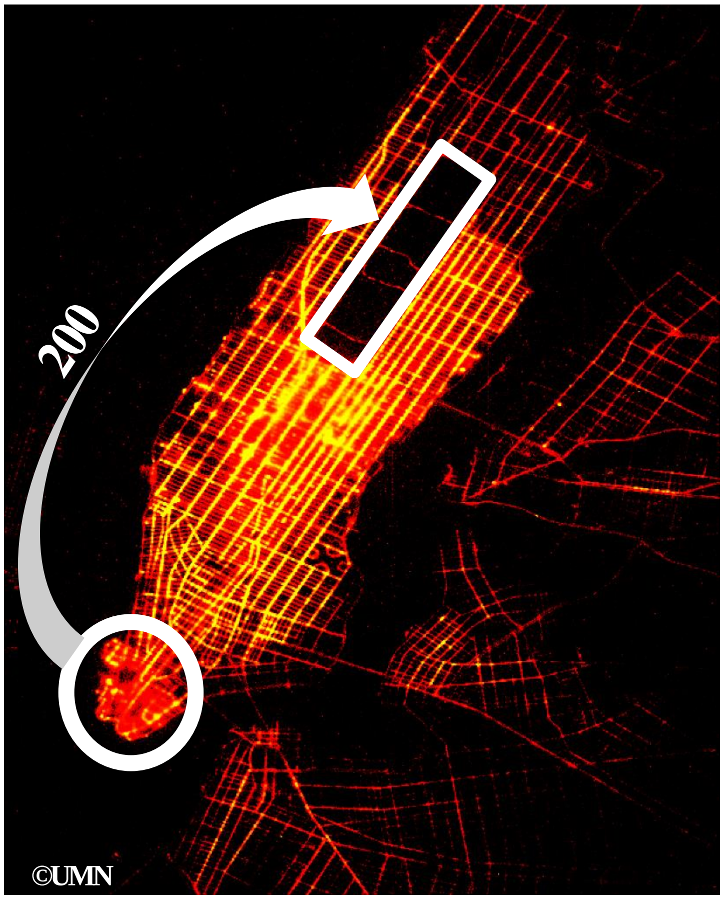

Urban Scale & Real Time

Various Applications

State of the Art

MobiCom'07	UbiComp'11	KDD'11	UbiComp'13
Zhang et al.	*Lathia et al.*	*Cho et al.*	*Ganti et al.*

Small Data Driven

or or or

Small Sample **Offline** **Single Domain Driven**

Biased Sampling & Overfitting

Opportunity: Cross-Domain Urban System Data

Shenzhem

Contribution: Multi-view Bounding

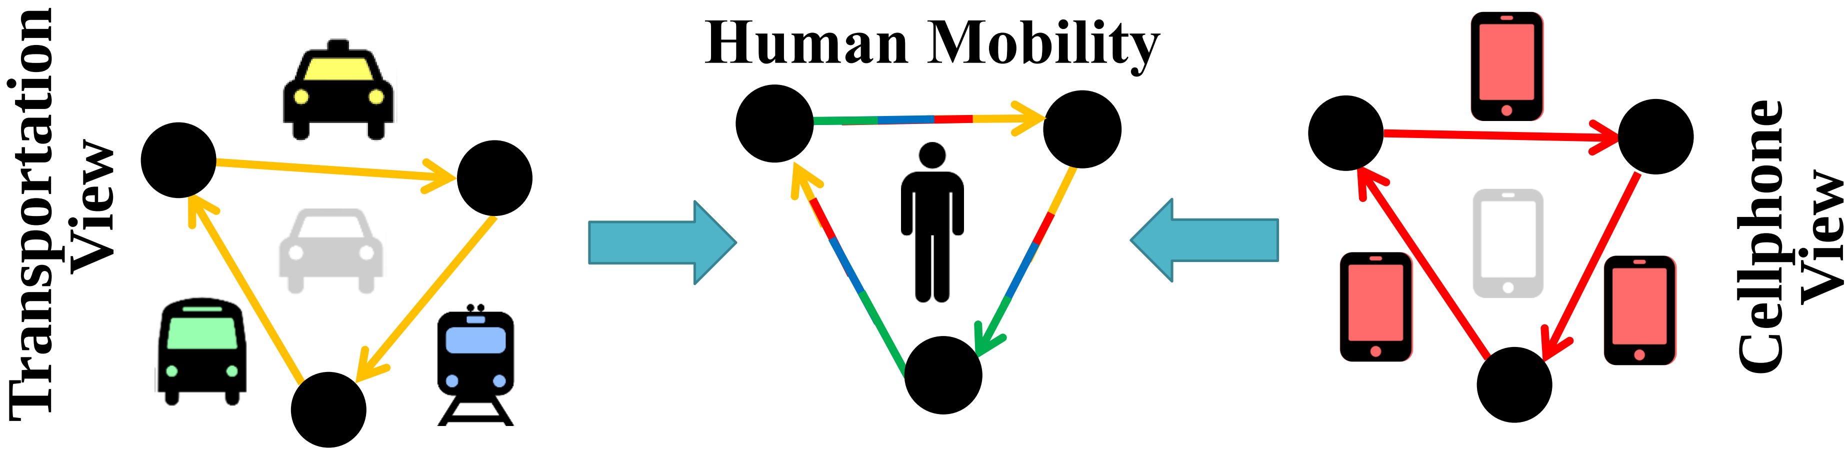

Data-Driven Model-Integration Technique

[MobiCom'14]

- Considering **Domains** as Views
- Quantifying **Biases** of Individual Views
- Bounding **Range** of Unknown Phenomena by Interdependent Views

Urban Scale

Online

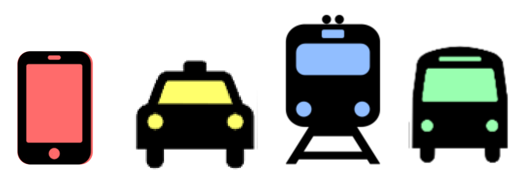

Cross Domain

Volume

Velocity

Variety

Technique: Context-Aware Multi-View Bounding

Spatiotemporal Context: A → B during t

?

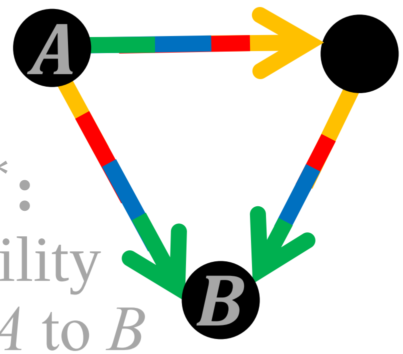

x^*: Mobility From A to B

K: Number of Views

x^k: Normalized Mobility Observed

w^k: Completeness Degree of a View k

$W = (w^1, \dots, w^K)$

?

Bounding **Range** of Unknown Phenomena by Multi-Views with Quantified Biases

View Completeness

Loss Function

Regularization Function

$$\min_{x^*, W} \mathbf{F}(x^*, W) = \underbrace{\sum_{k=1}^{K} [\; w^k \cdot \mathbf{D}\left(x^*, x^k\right) \;]}_{\text{Overall Weighted Deviation}}, \quad \text{s.t.} \quad \delta(W) = 1$$

High Penalty if a **more-complete** view deviates from the truth;

Low Penalty if a **less-complete** view deviates from the truth;

Implementation: Region Level

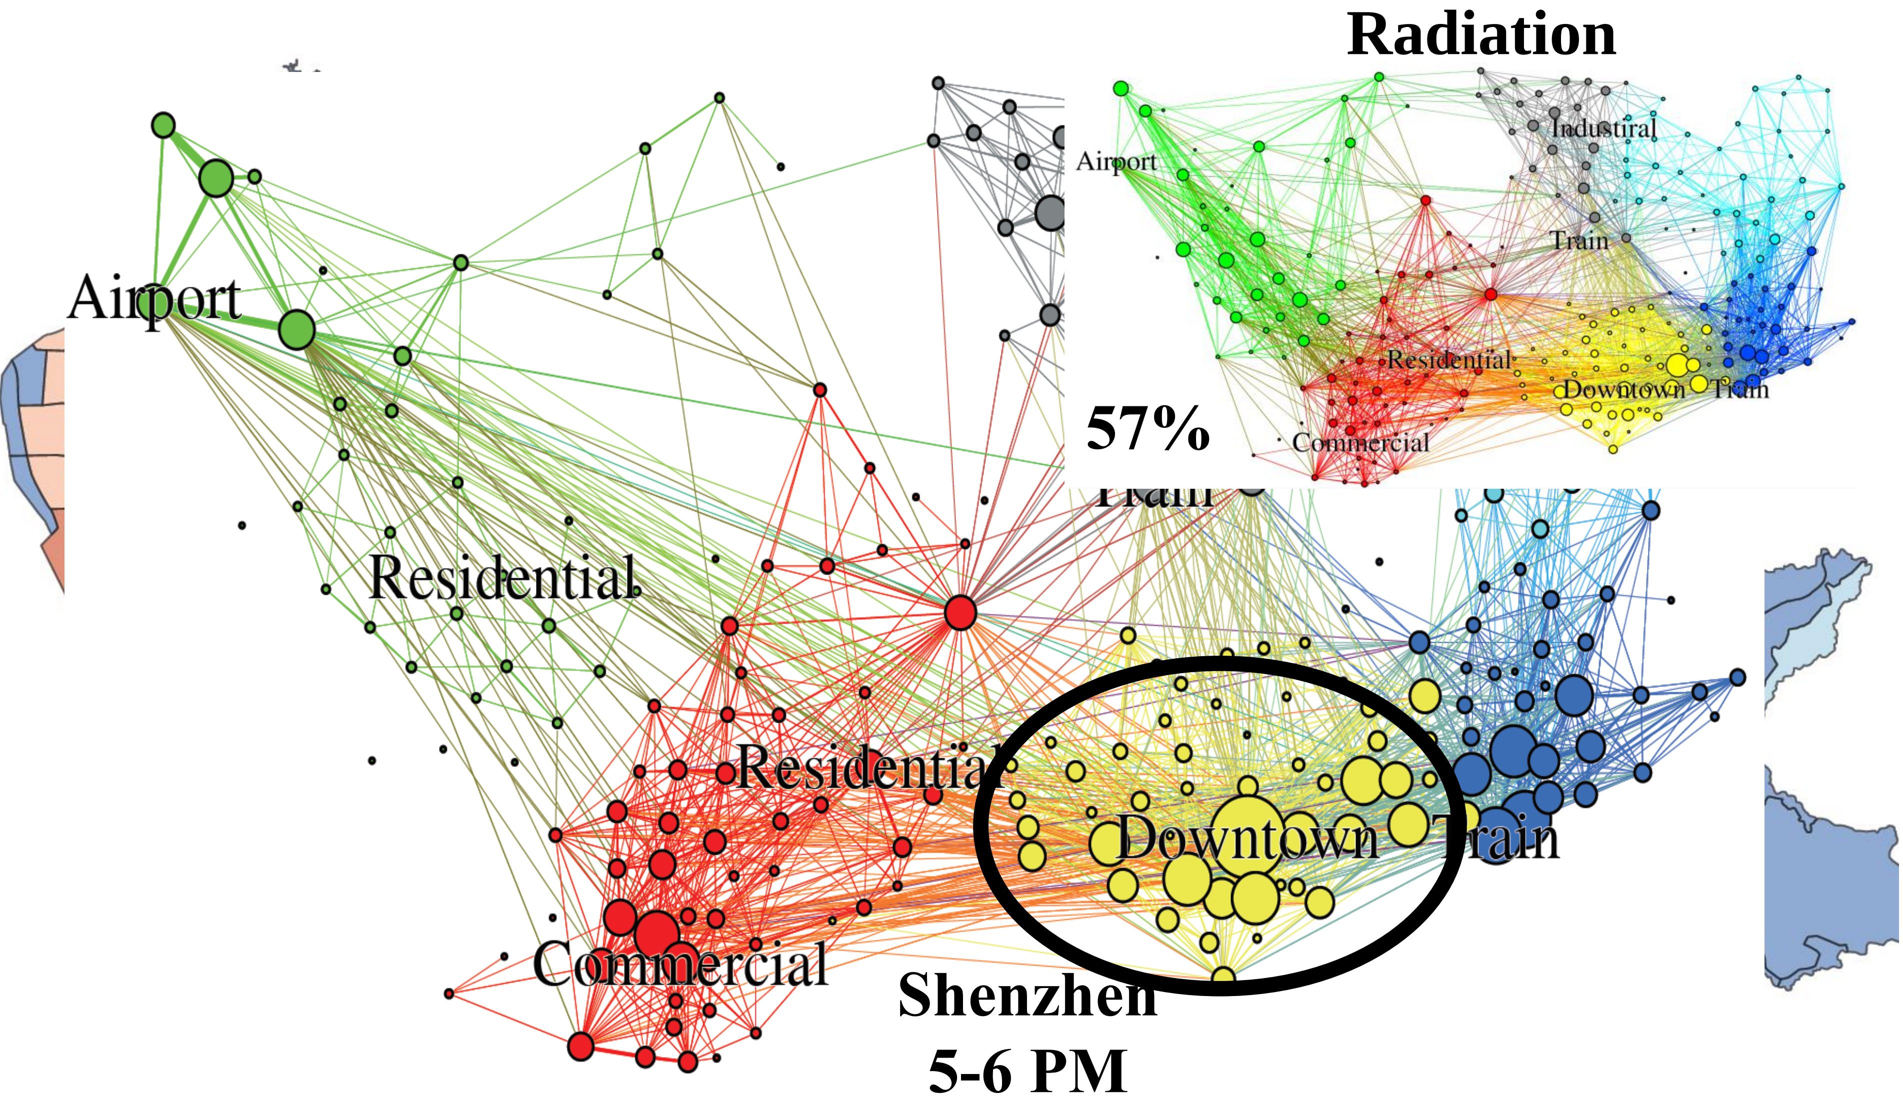

Vertex: Urban Region; **Edge**: Mobility Volume; **Color**: Urban Districts

Implementation: Street Level

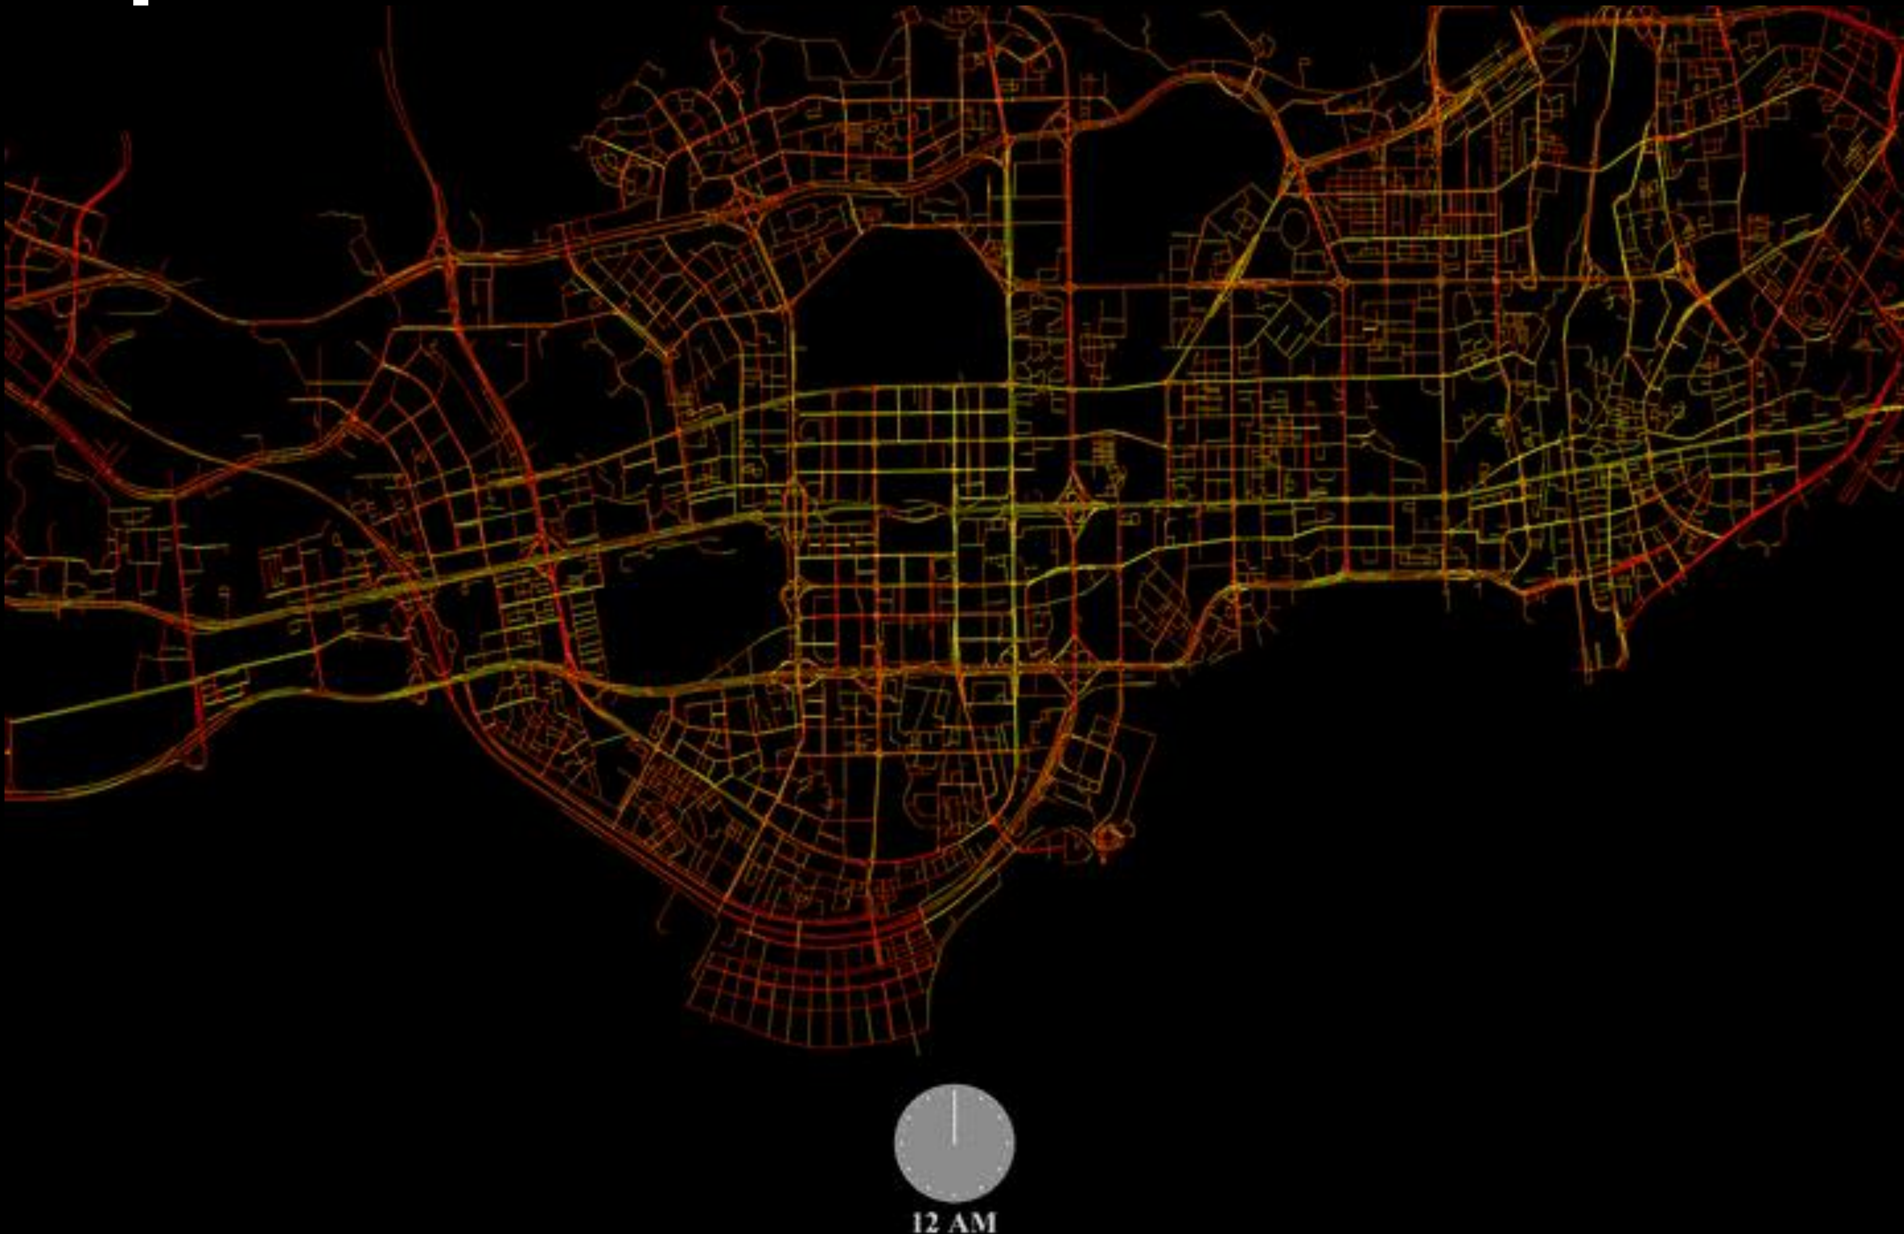

Broader Impact: Open Data

20 GB Cross-Domain Data

Taxi; Bus; Truck; Subway; Cellphone; Smartcards;

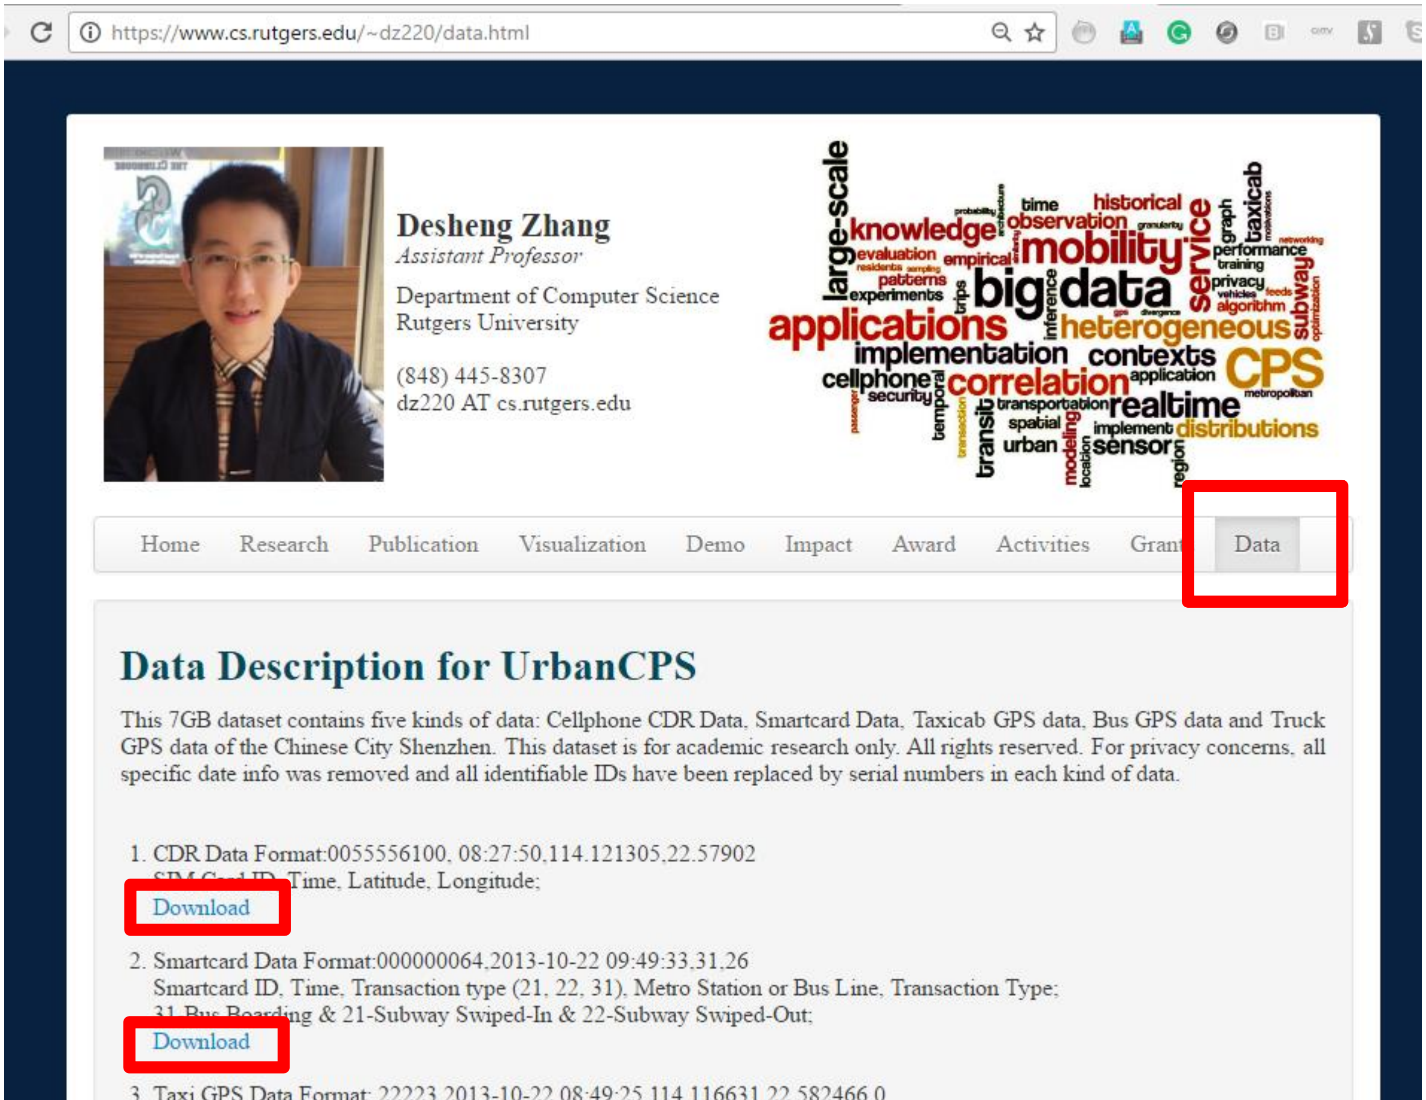

Big Picture: Understanding Real-time Urban Phenomena

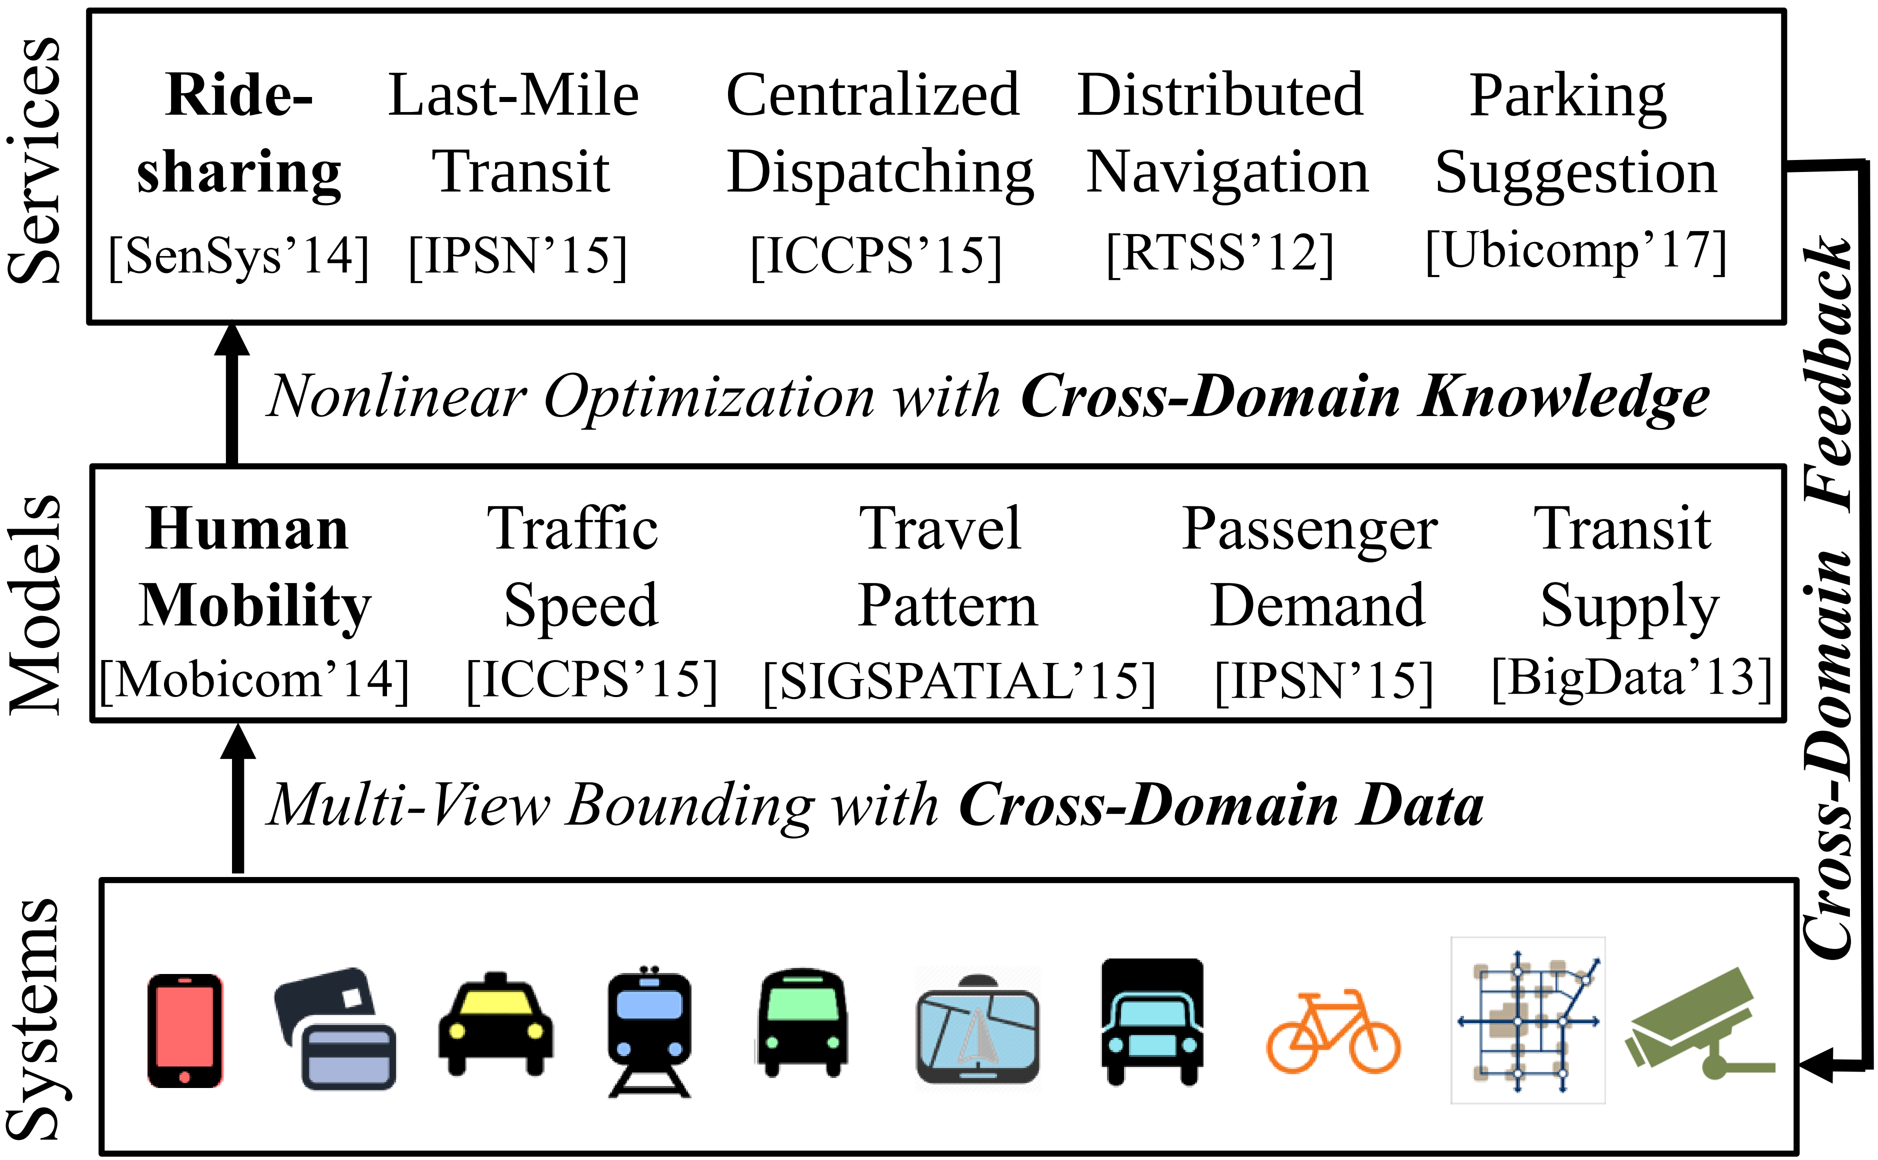

Interdependent Urban Phenomena

$$\rho = \frac{\text{Mobility Demand}}{\text{Transit Supply}}$$

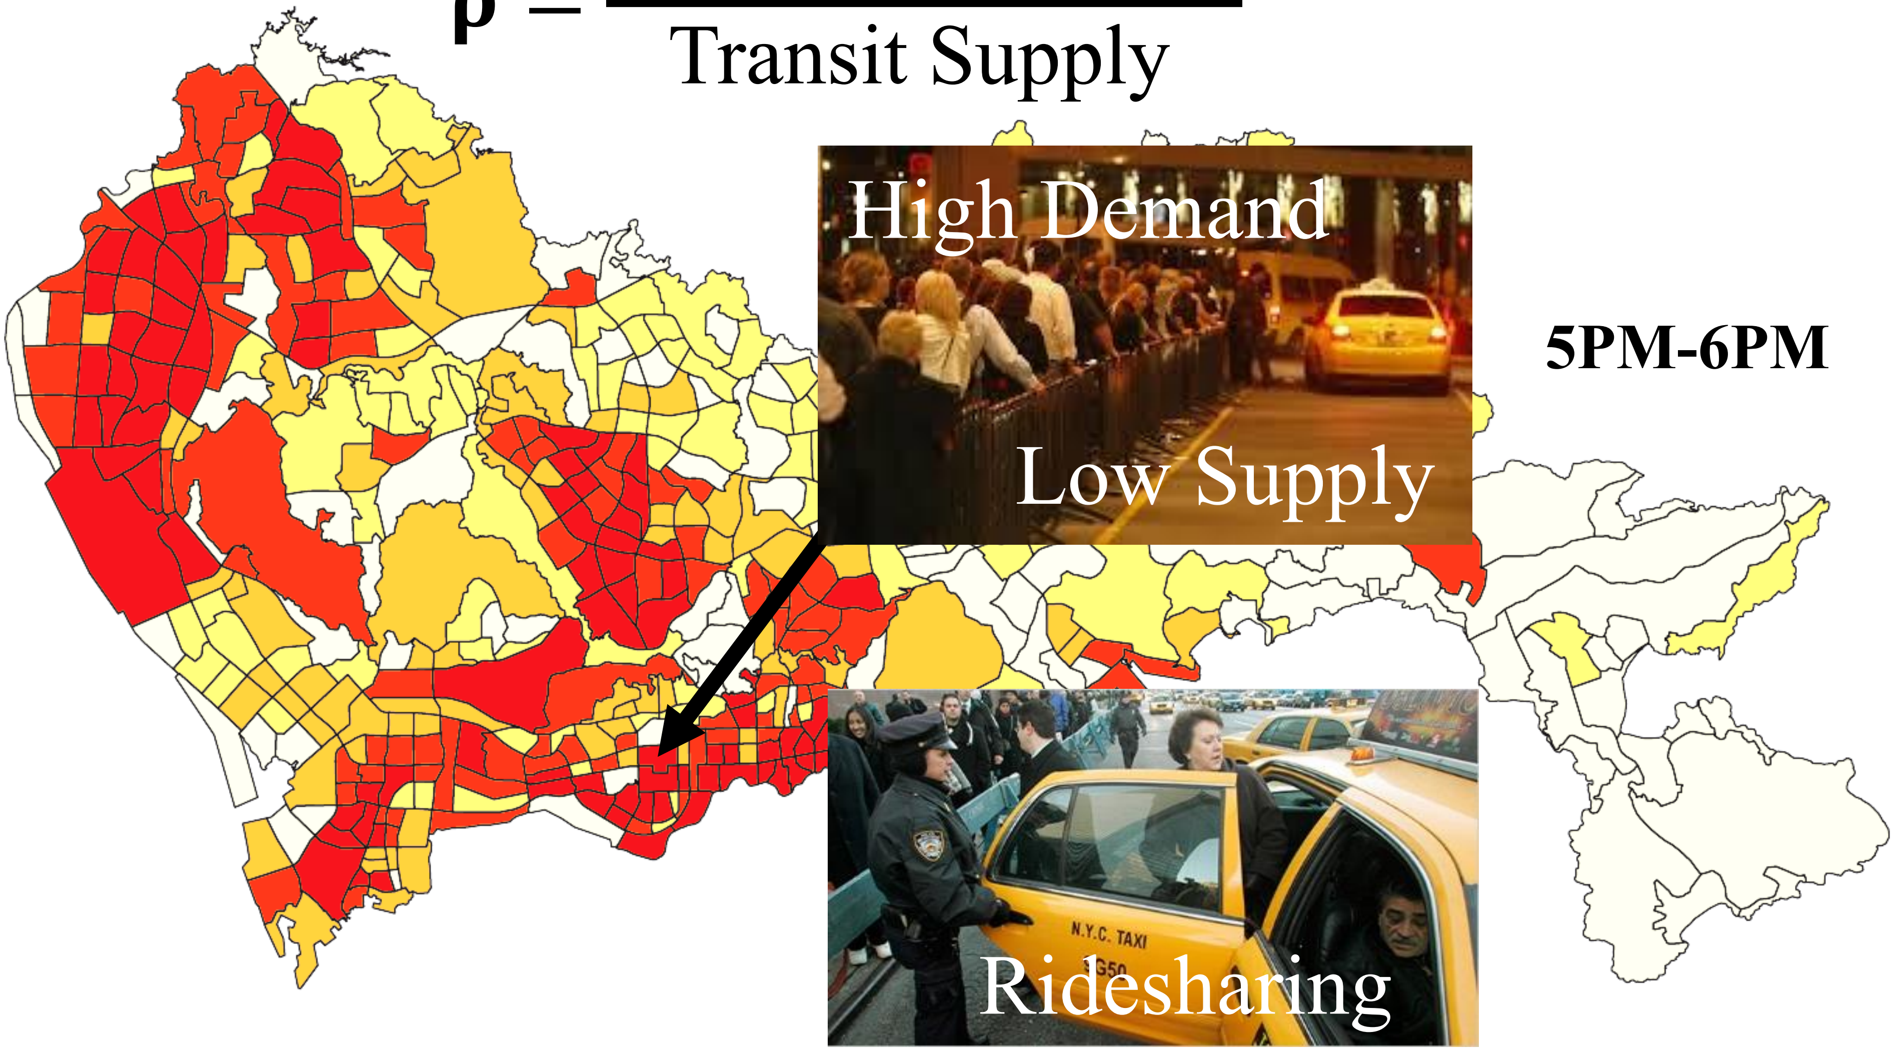

5PM-6PM

State of the Art & Limitations

1.3 Passenger Per Trip in NYC

70% Wasted

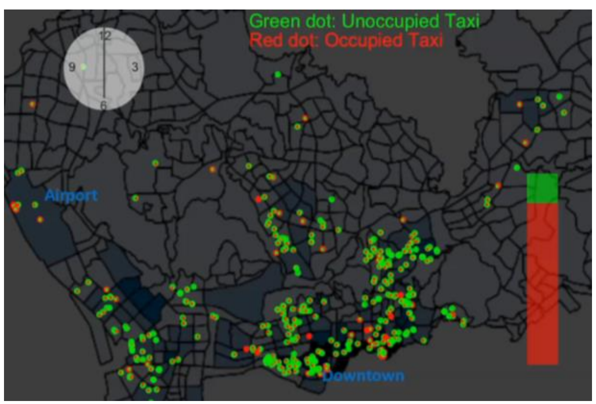

Similar Patterns

Sharing Capacity

T-Share (TKDE'14)

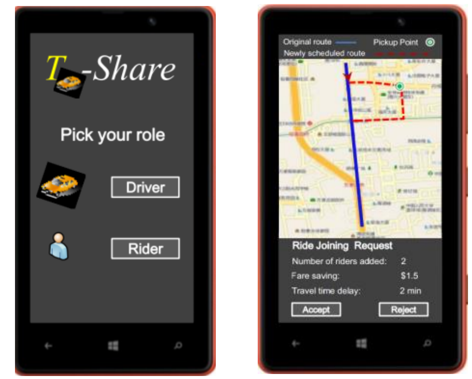

UberPool (2014.12)

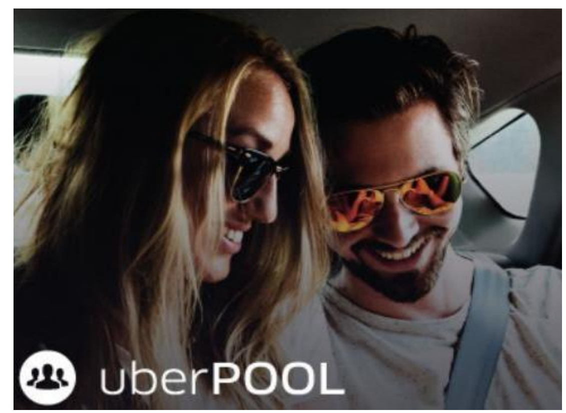

- Single-Domain Knowledge
- No Overall Urban Mobility Info

Opportunity: Cross-Domain Knowledge

Traffic Speed

Transportation & Geography
Finance & Telecom

Fare Model

Transportation & Finance

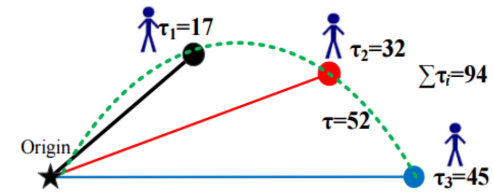

Minimizing Ridesharing Cost

Transportation: Mileage
Environment: Energy
Finance: Fare

coRide

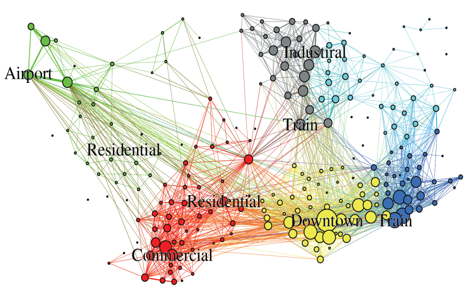

Mobility Demand

Telecom & Finance & Transportation

Transit Supply

Transportation & Geography

Contribution: coRide

- Cross-Domain Knowledge-Enabled Resource sharing Framework
  - Bounded Online Approximation
  - Real-world Implementation
  - Multi-Scale Multi-Site Evaluation
  - Potential Commercialization

Technique: 2 Approximation Algorithm

(1) **Passenger Assignment** with Minimum Spanning Tree (MST)

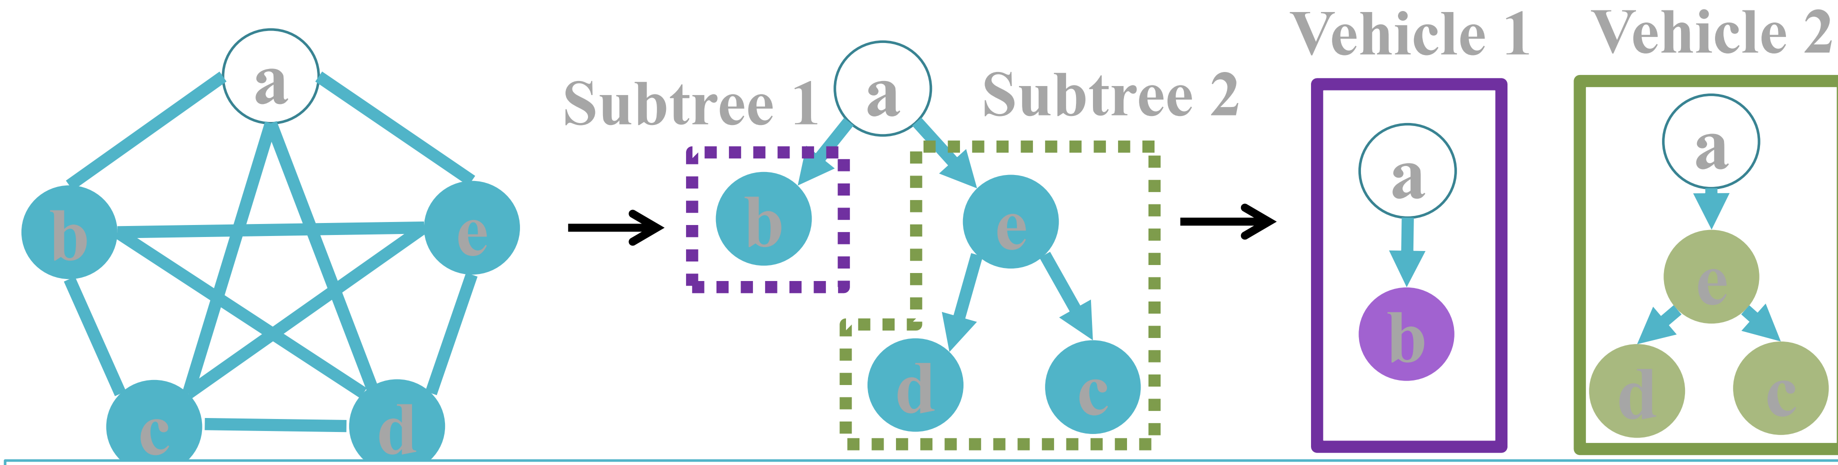

Mileage obtained by our *Online Approximation* is at most two times of *the Optimal Mileage*

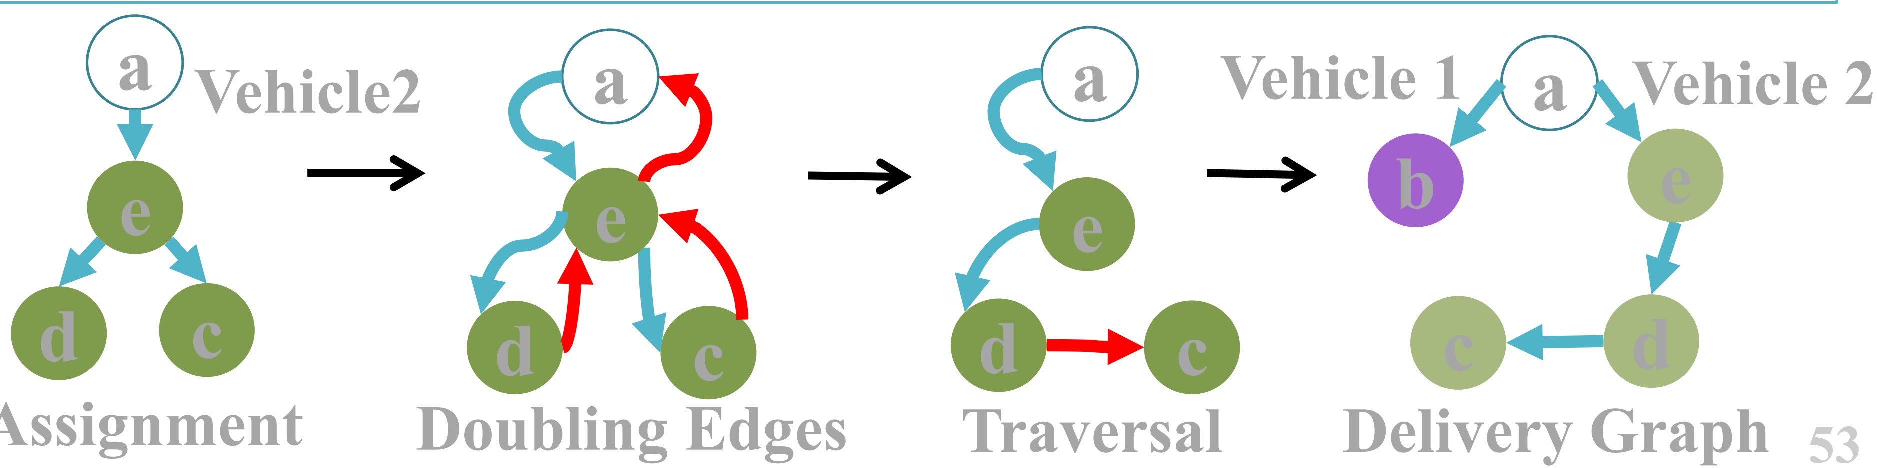

Ridesharing Evaluation (1/4): Region

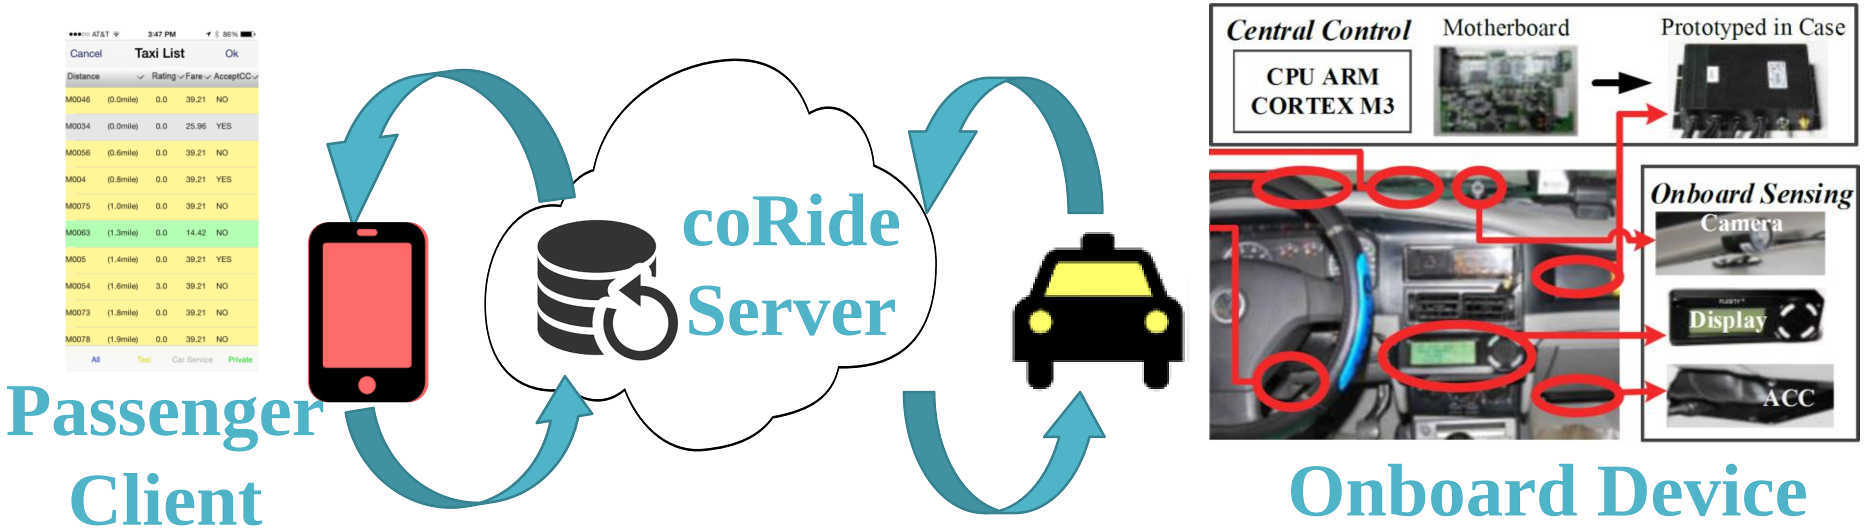

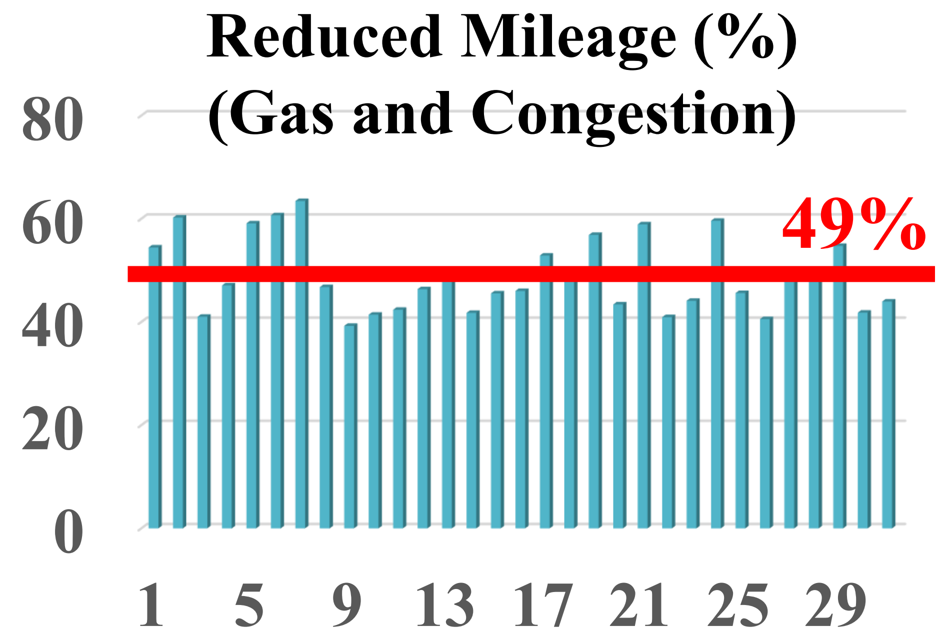

Ridesharing Evaluation (2/4): Urban

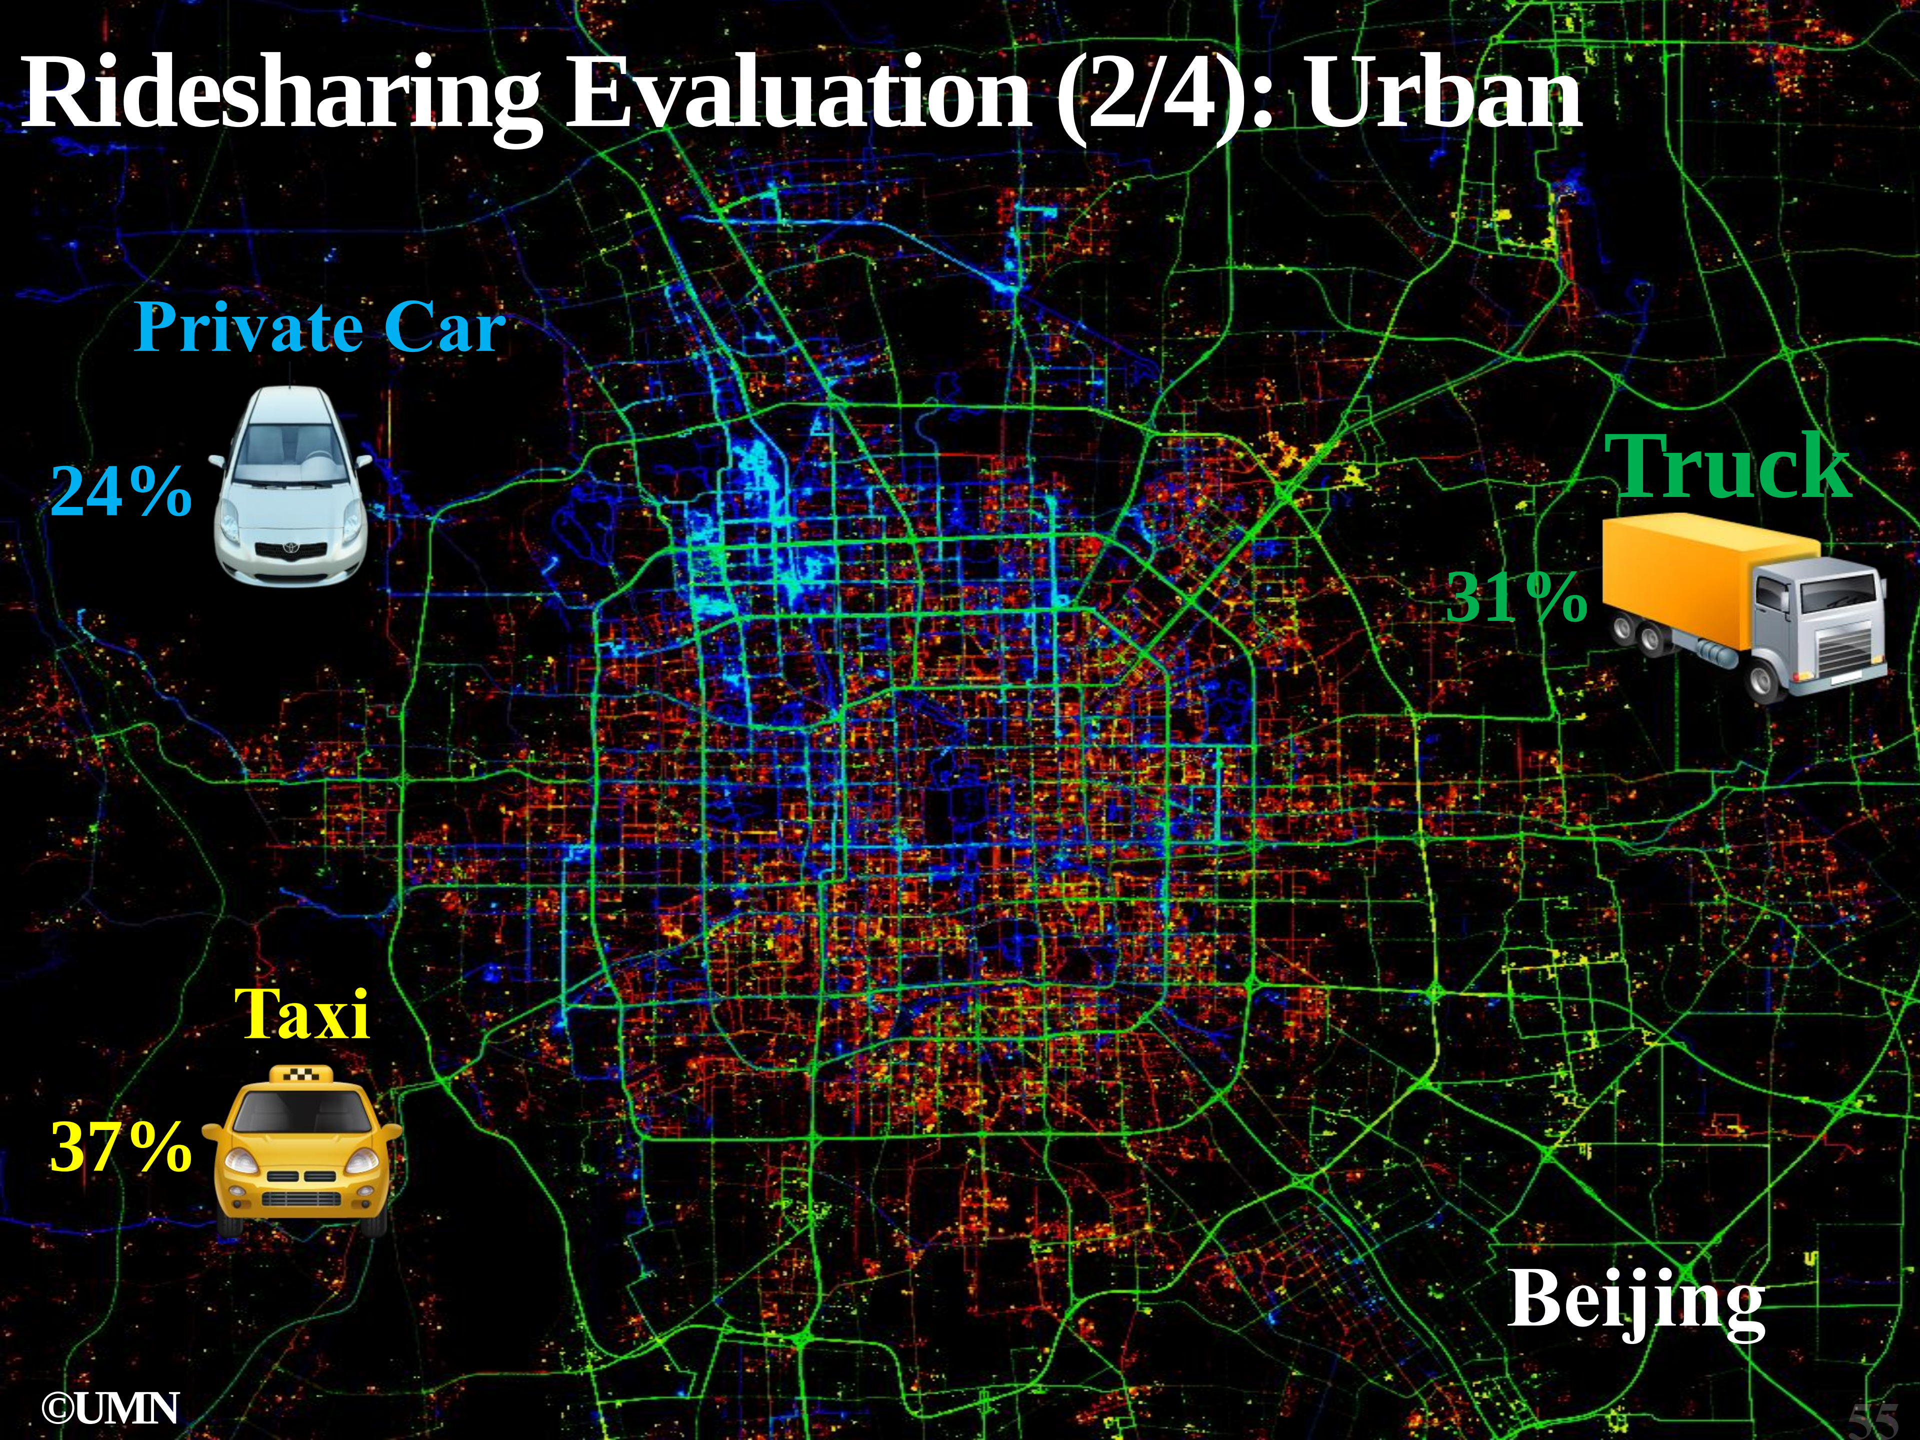

Ridesharing Evaluation (3/4): National

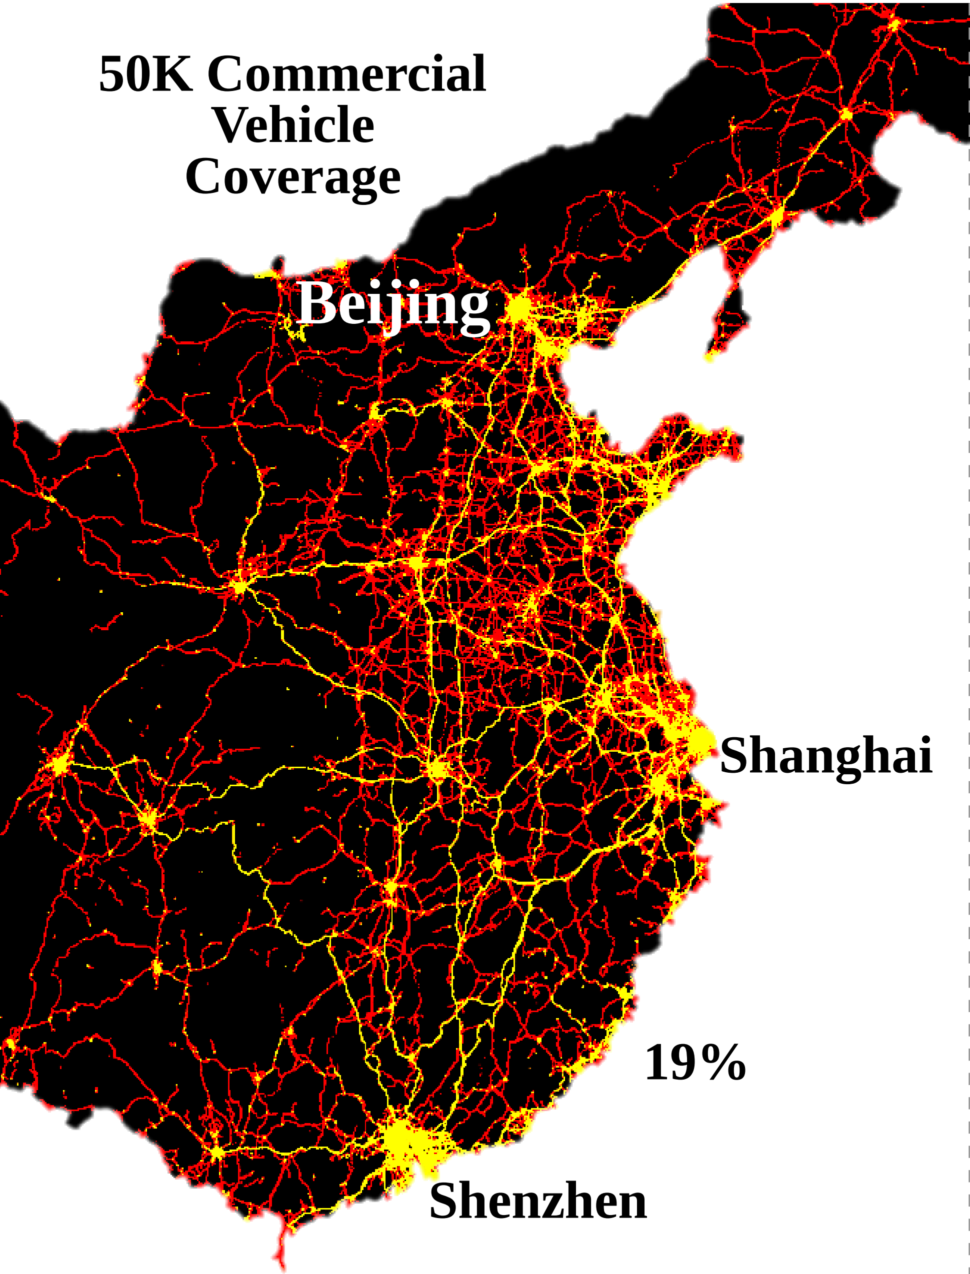

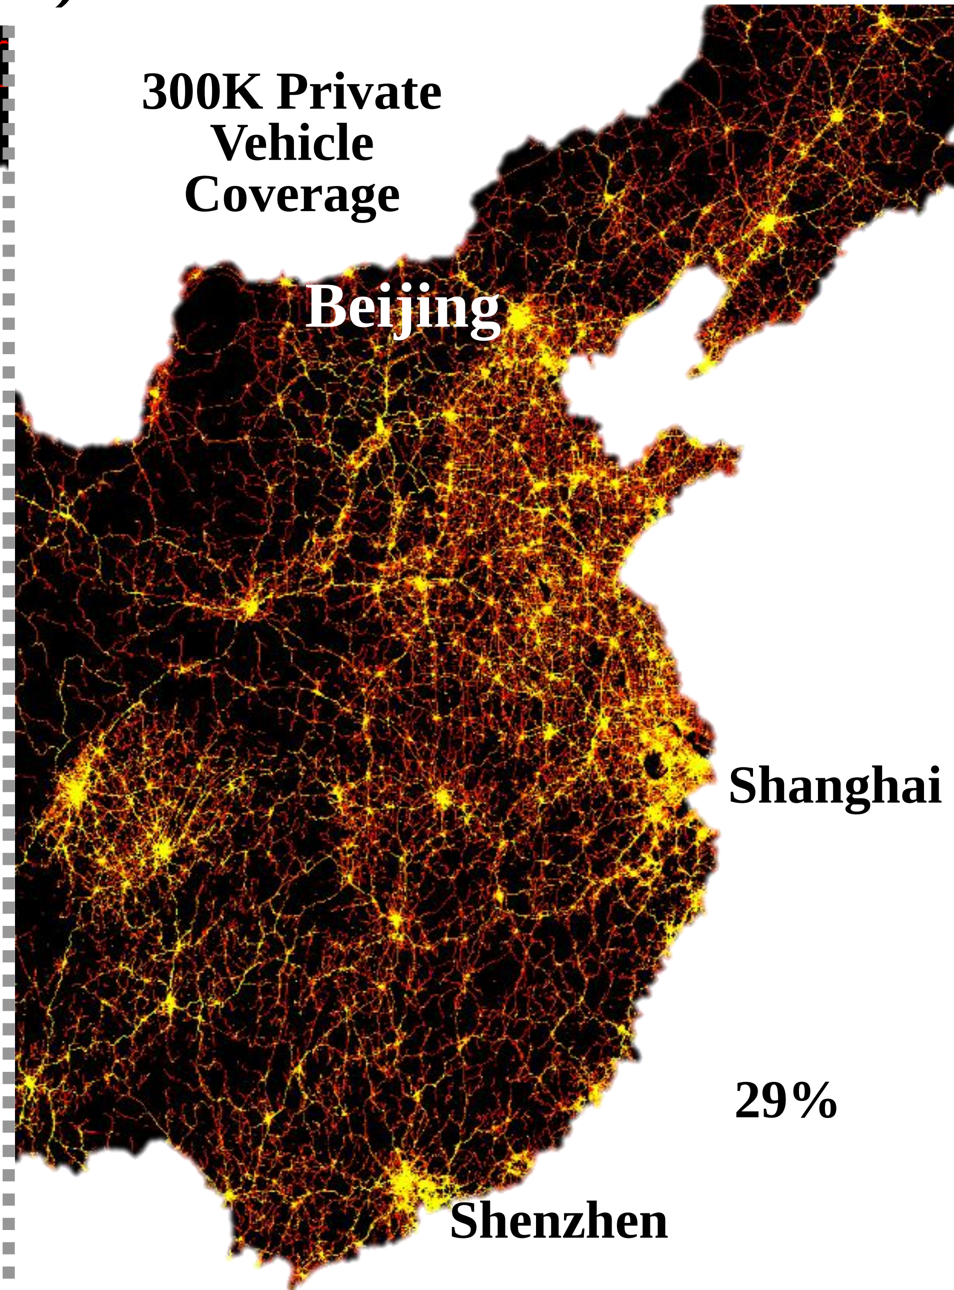

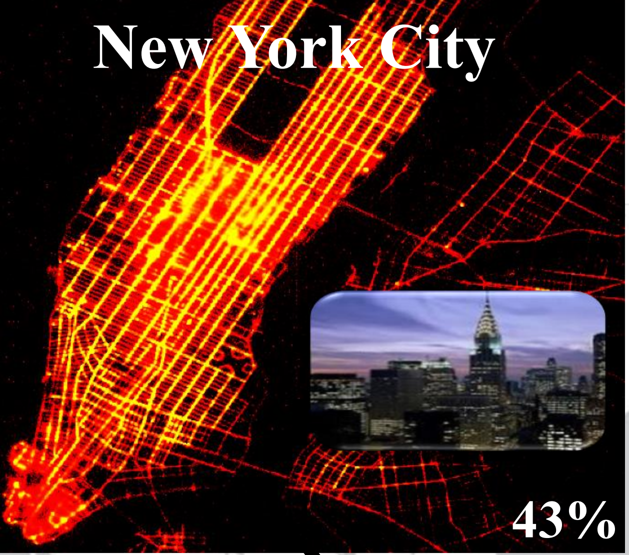

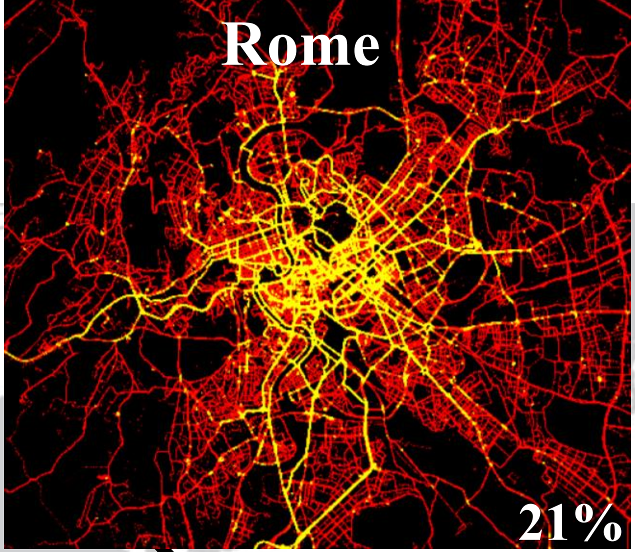

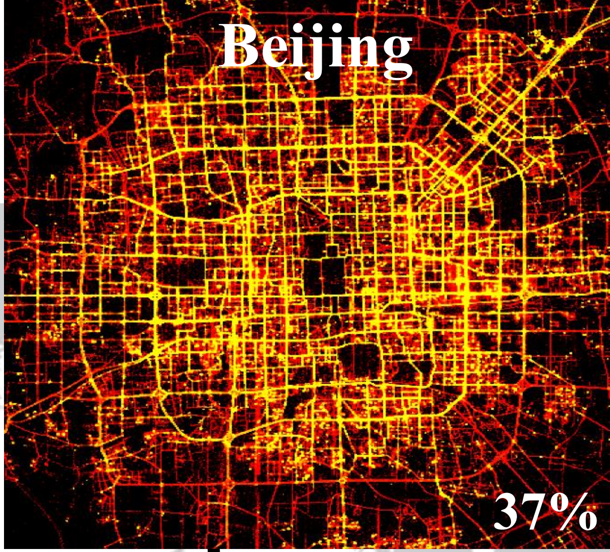

Ridesharing Evaluation (4/4): Worldwide

-- A Tale of Six Cities

[Video]

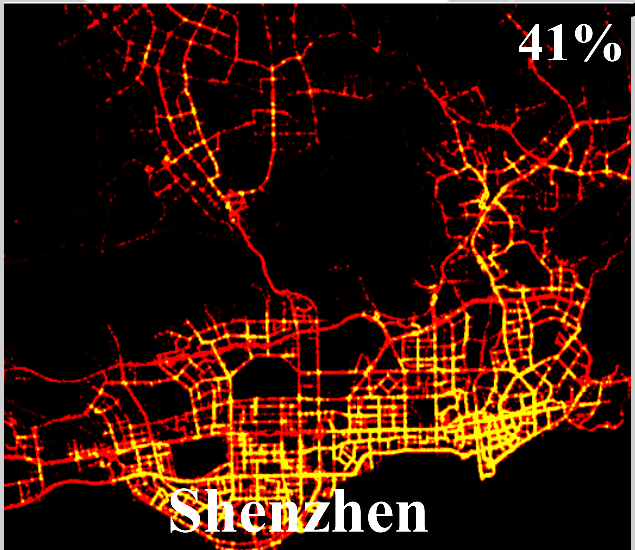

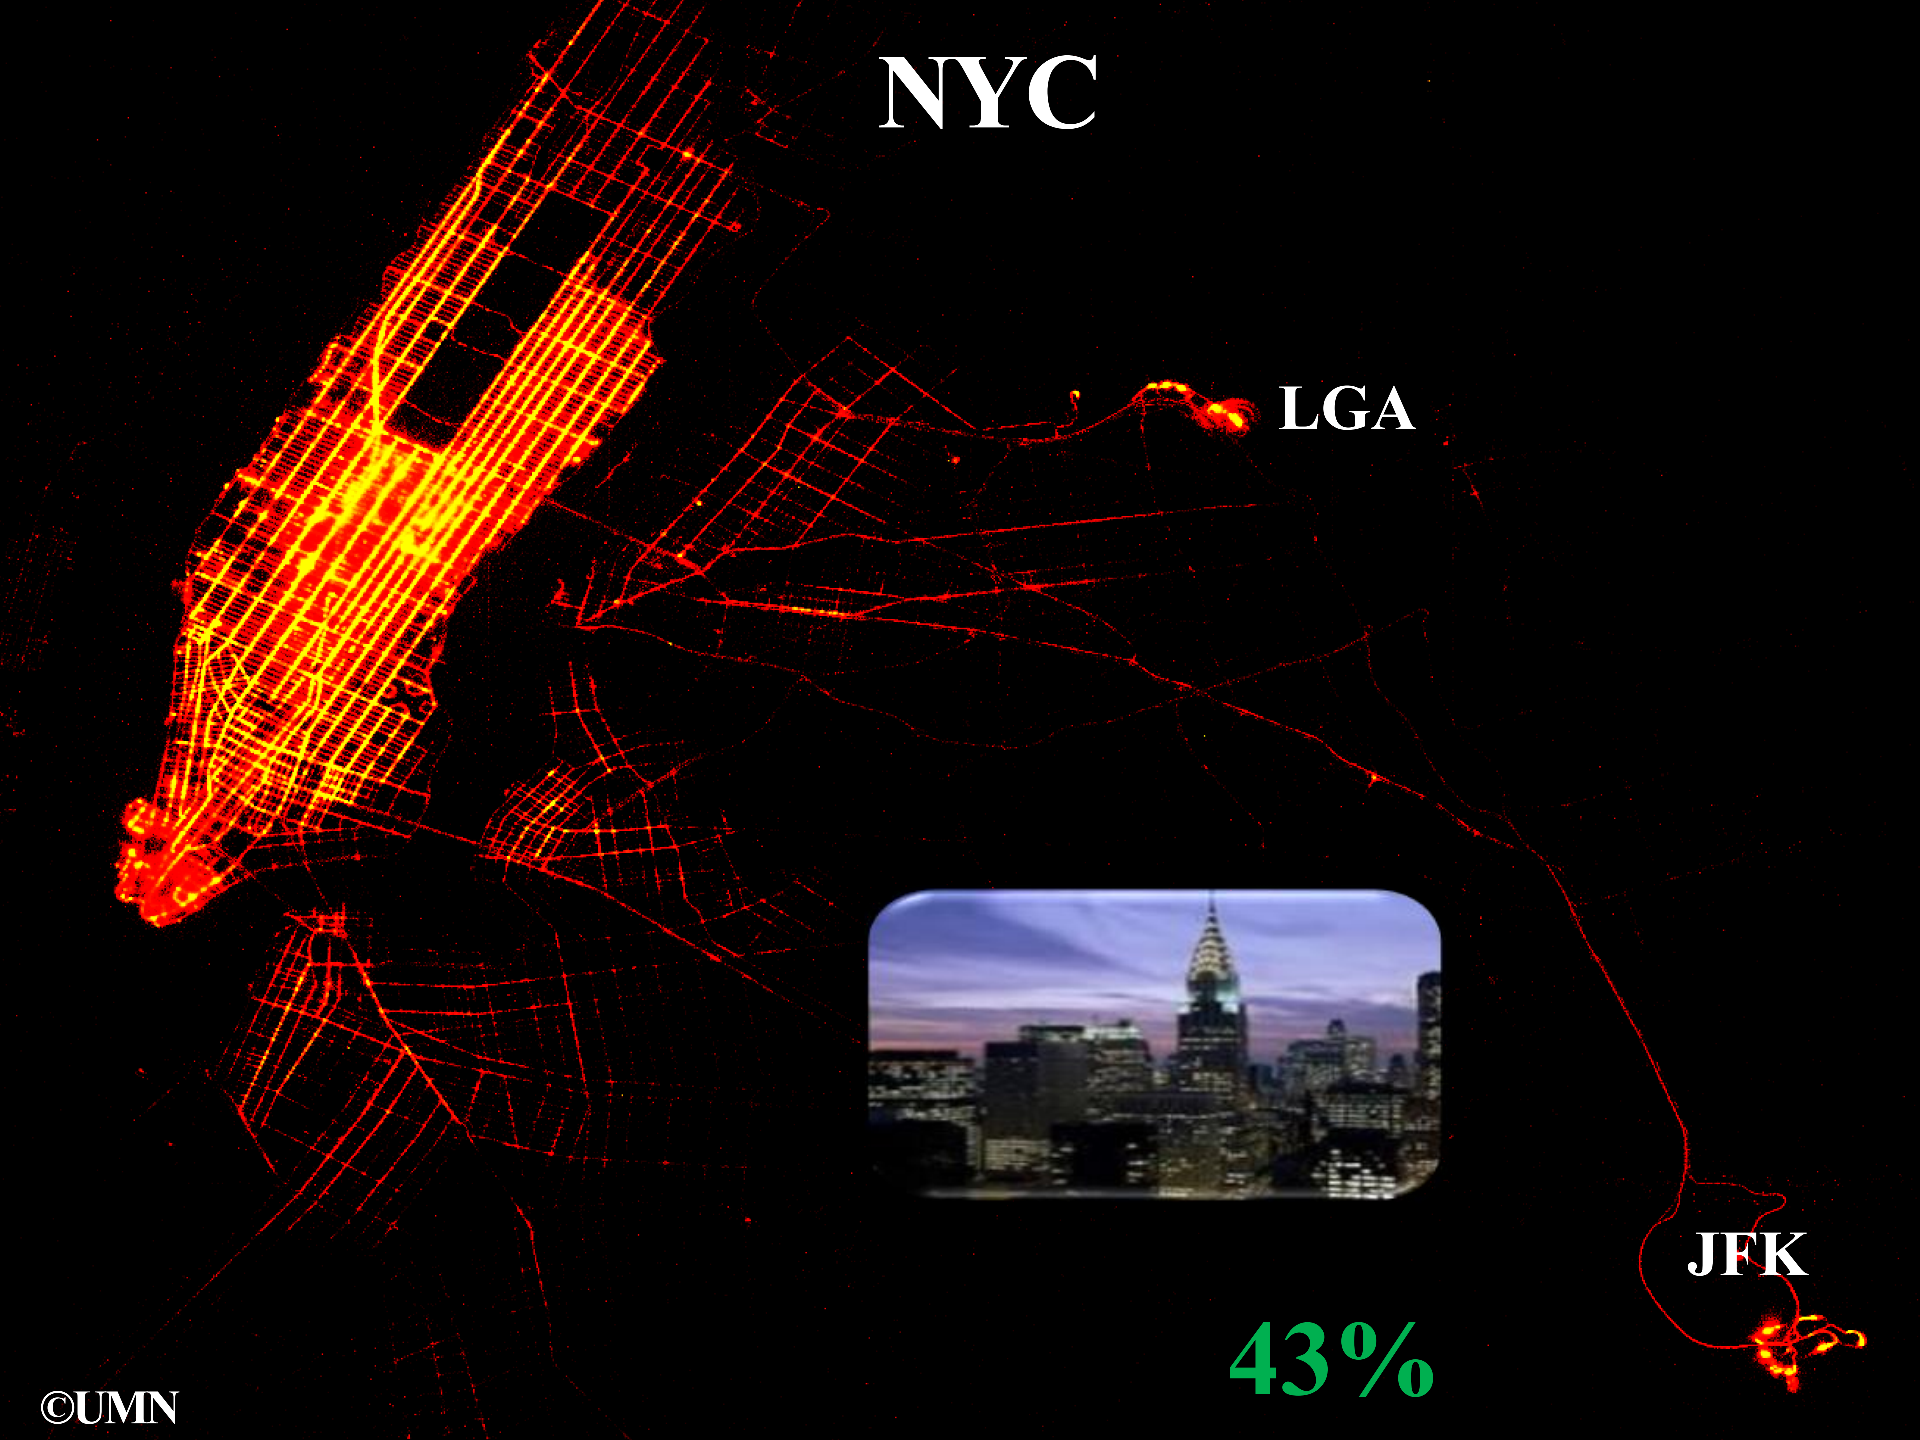
NYC
LGA
JFK
43%
©UMN

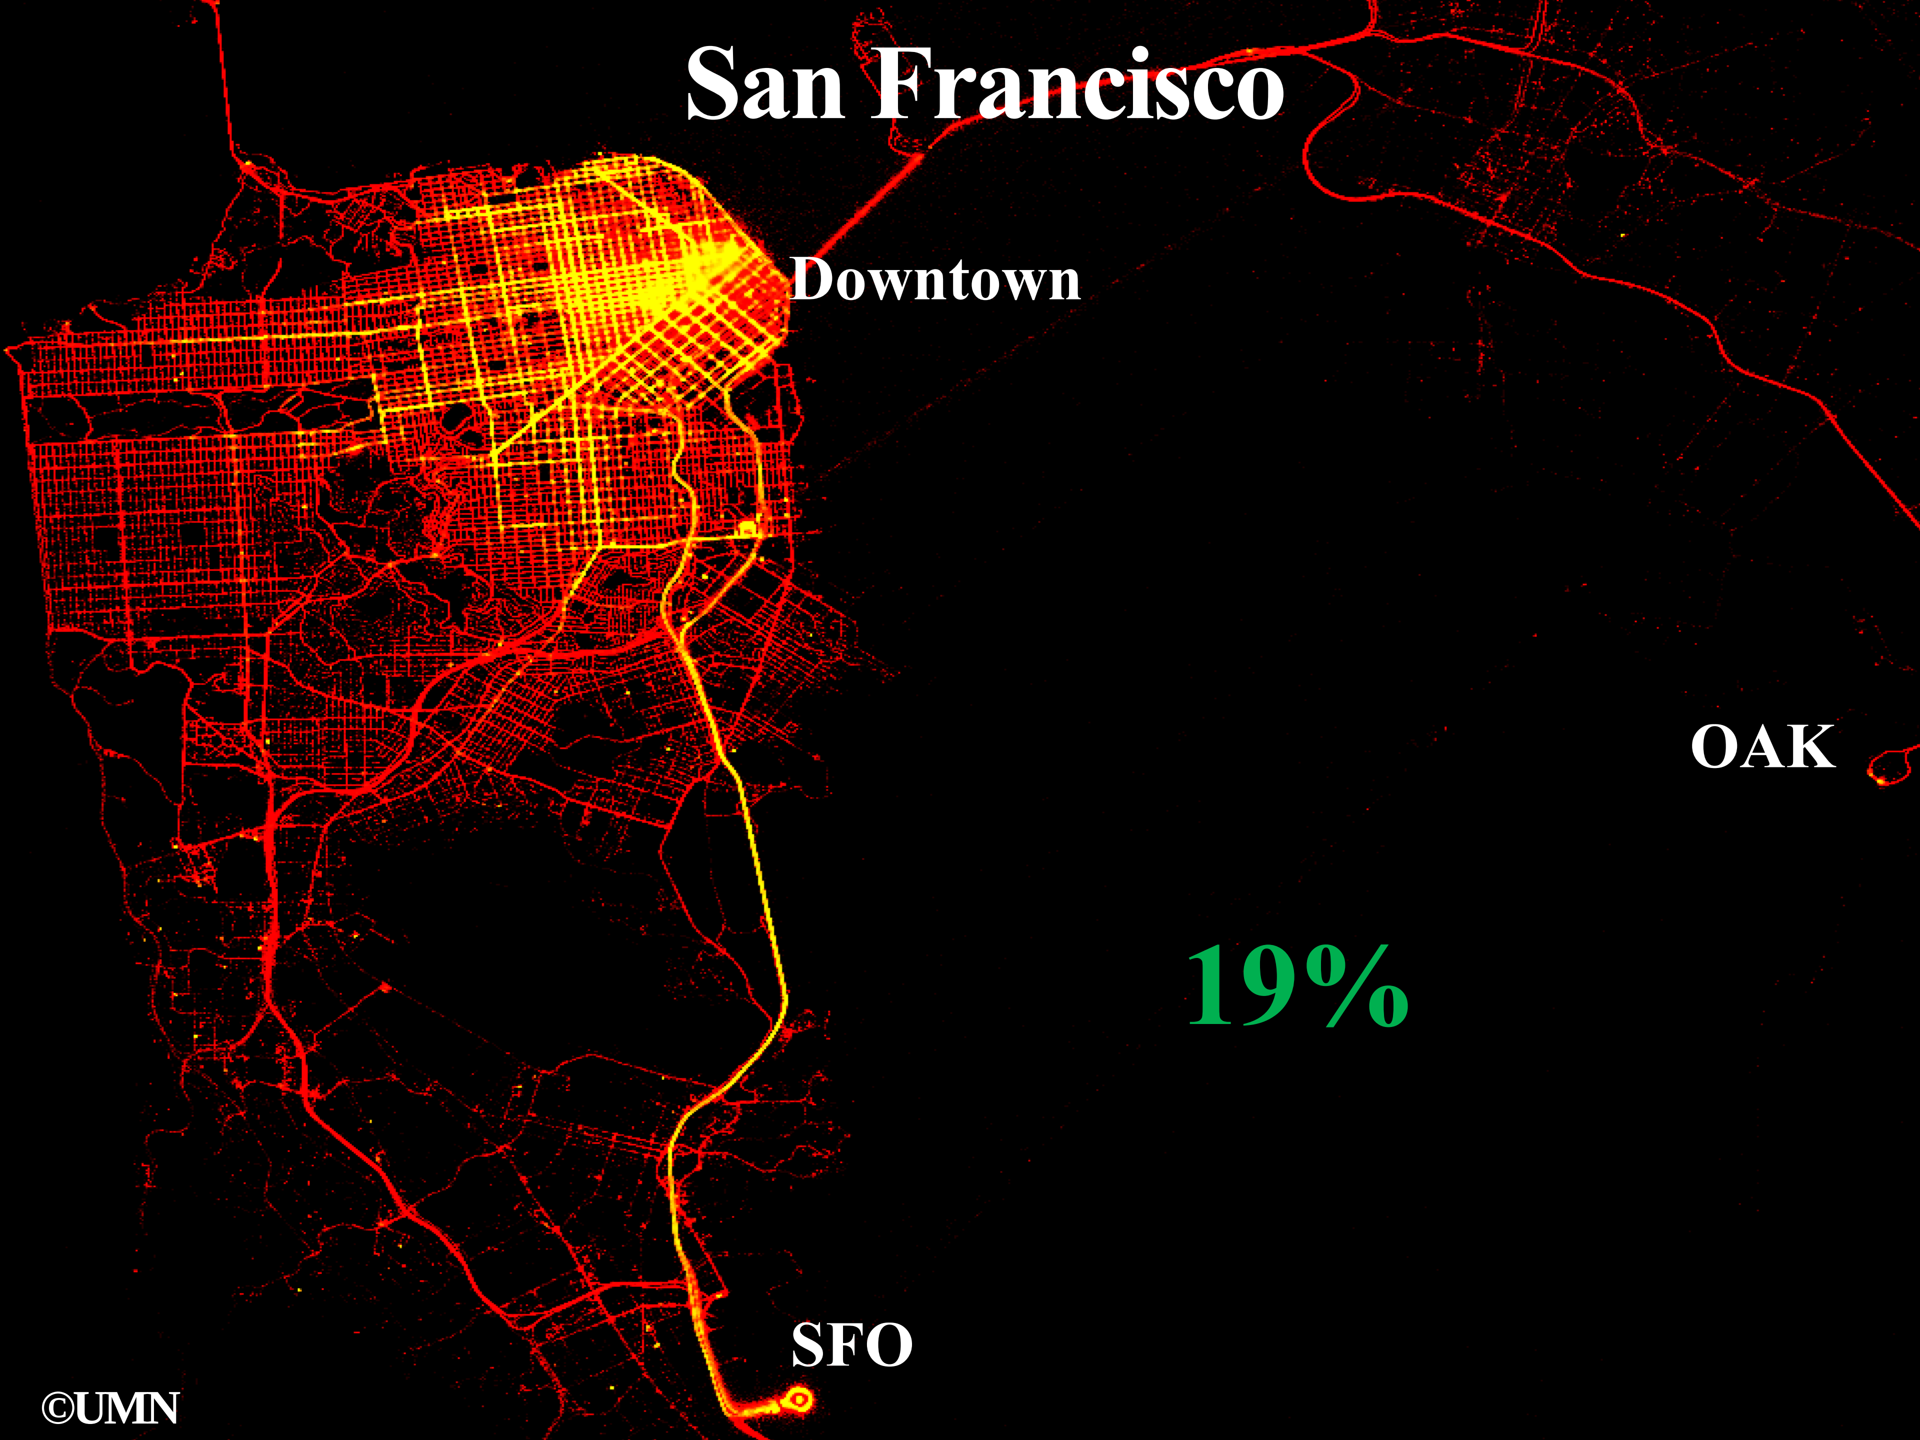
San Francisco
Downtown
OAK
19%
SFO
©UMN

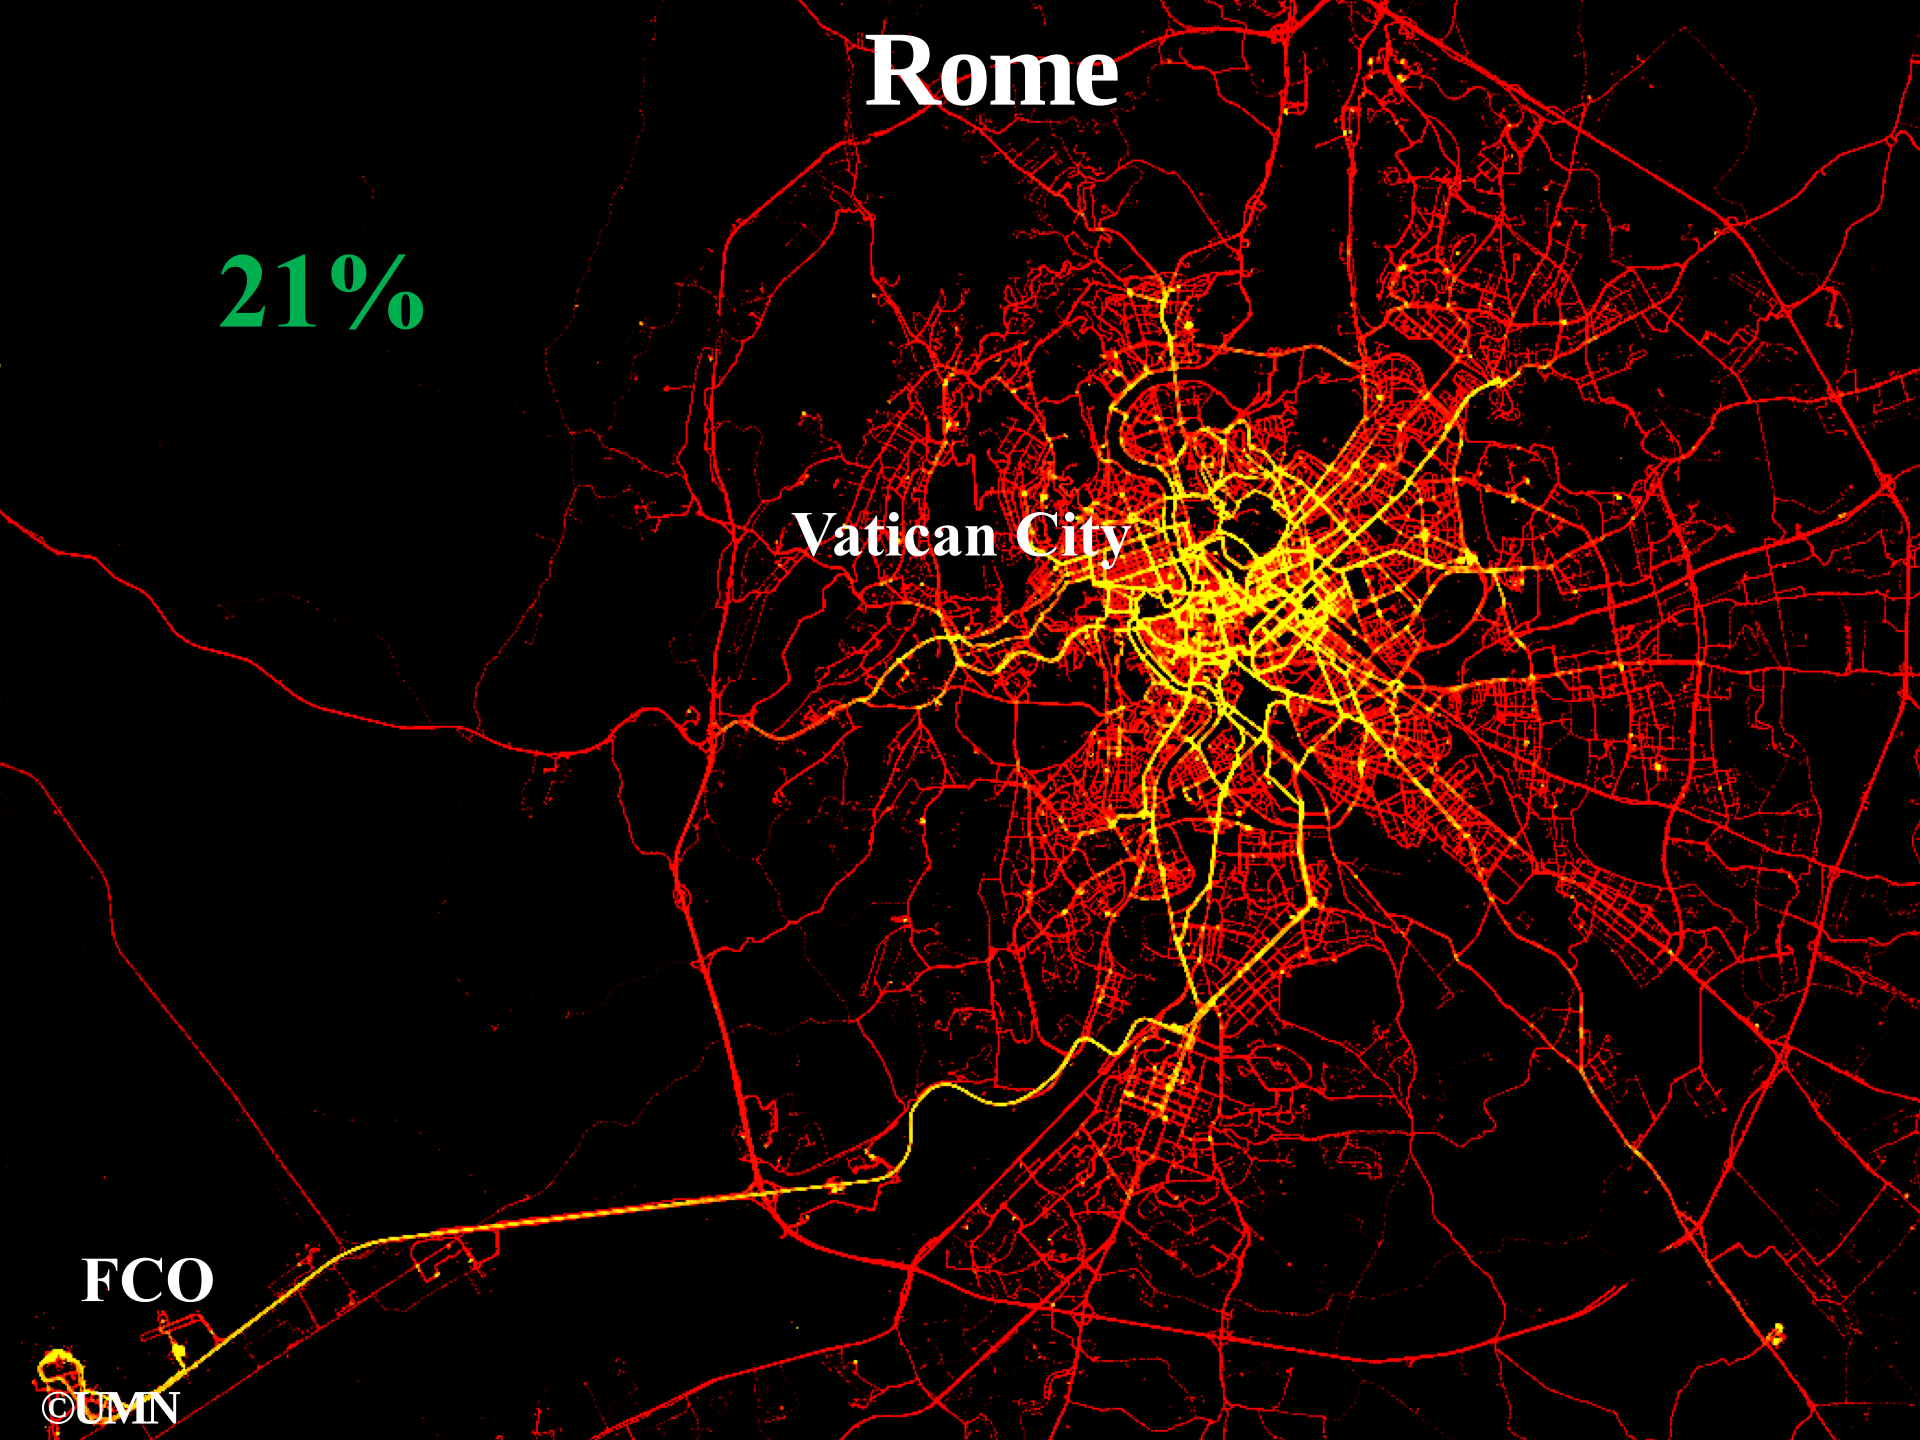
Rome
21%
Vatican City
FCO
©UMN

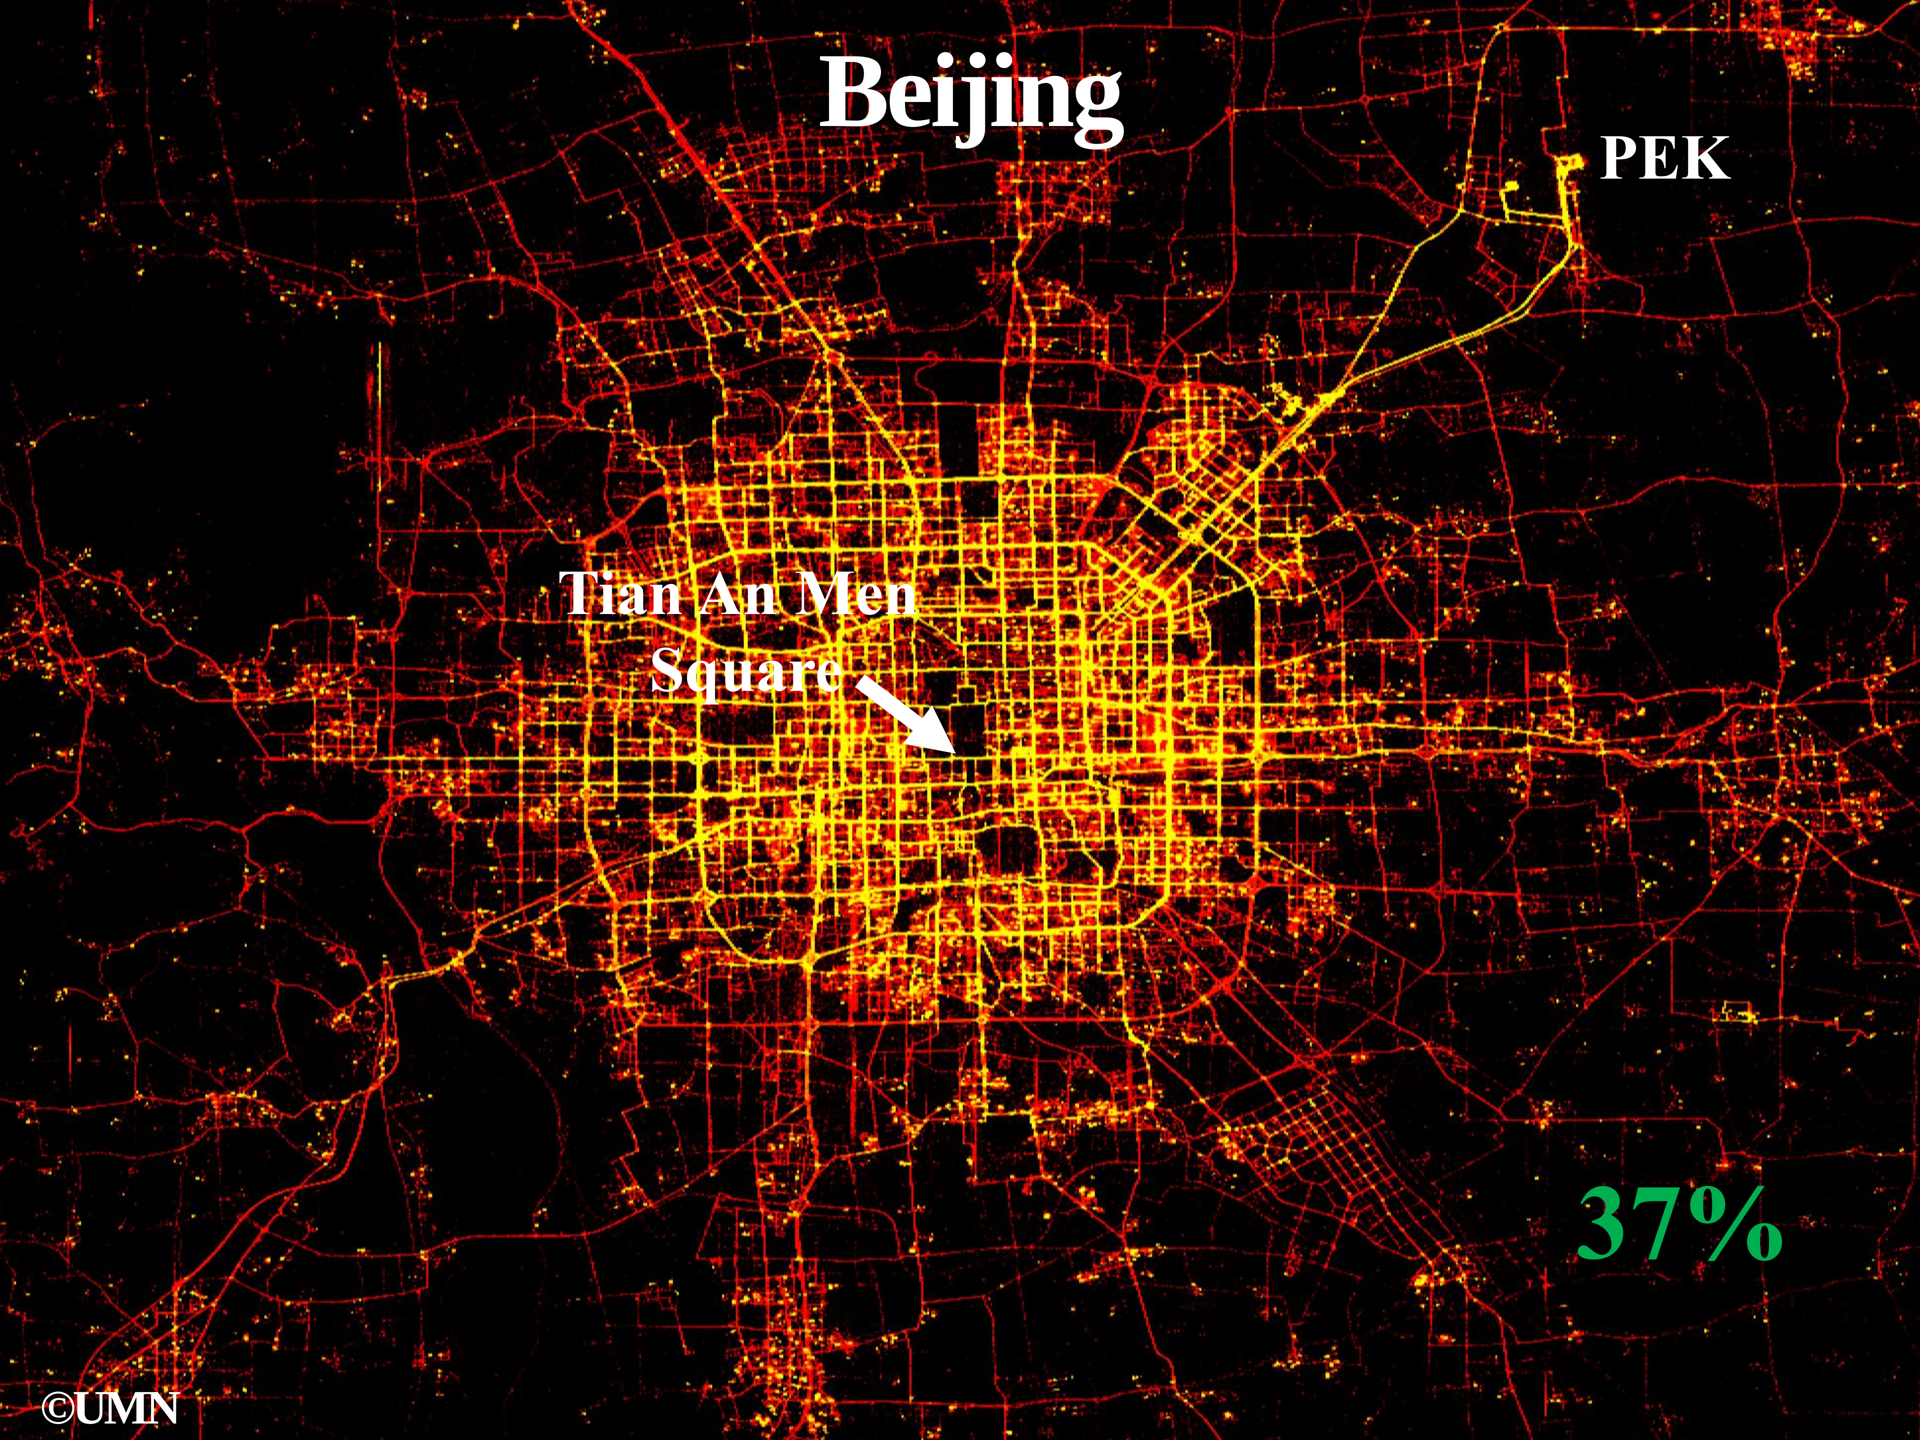

Beijing
PEK
Tian An Men
Square
37%
©UMN

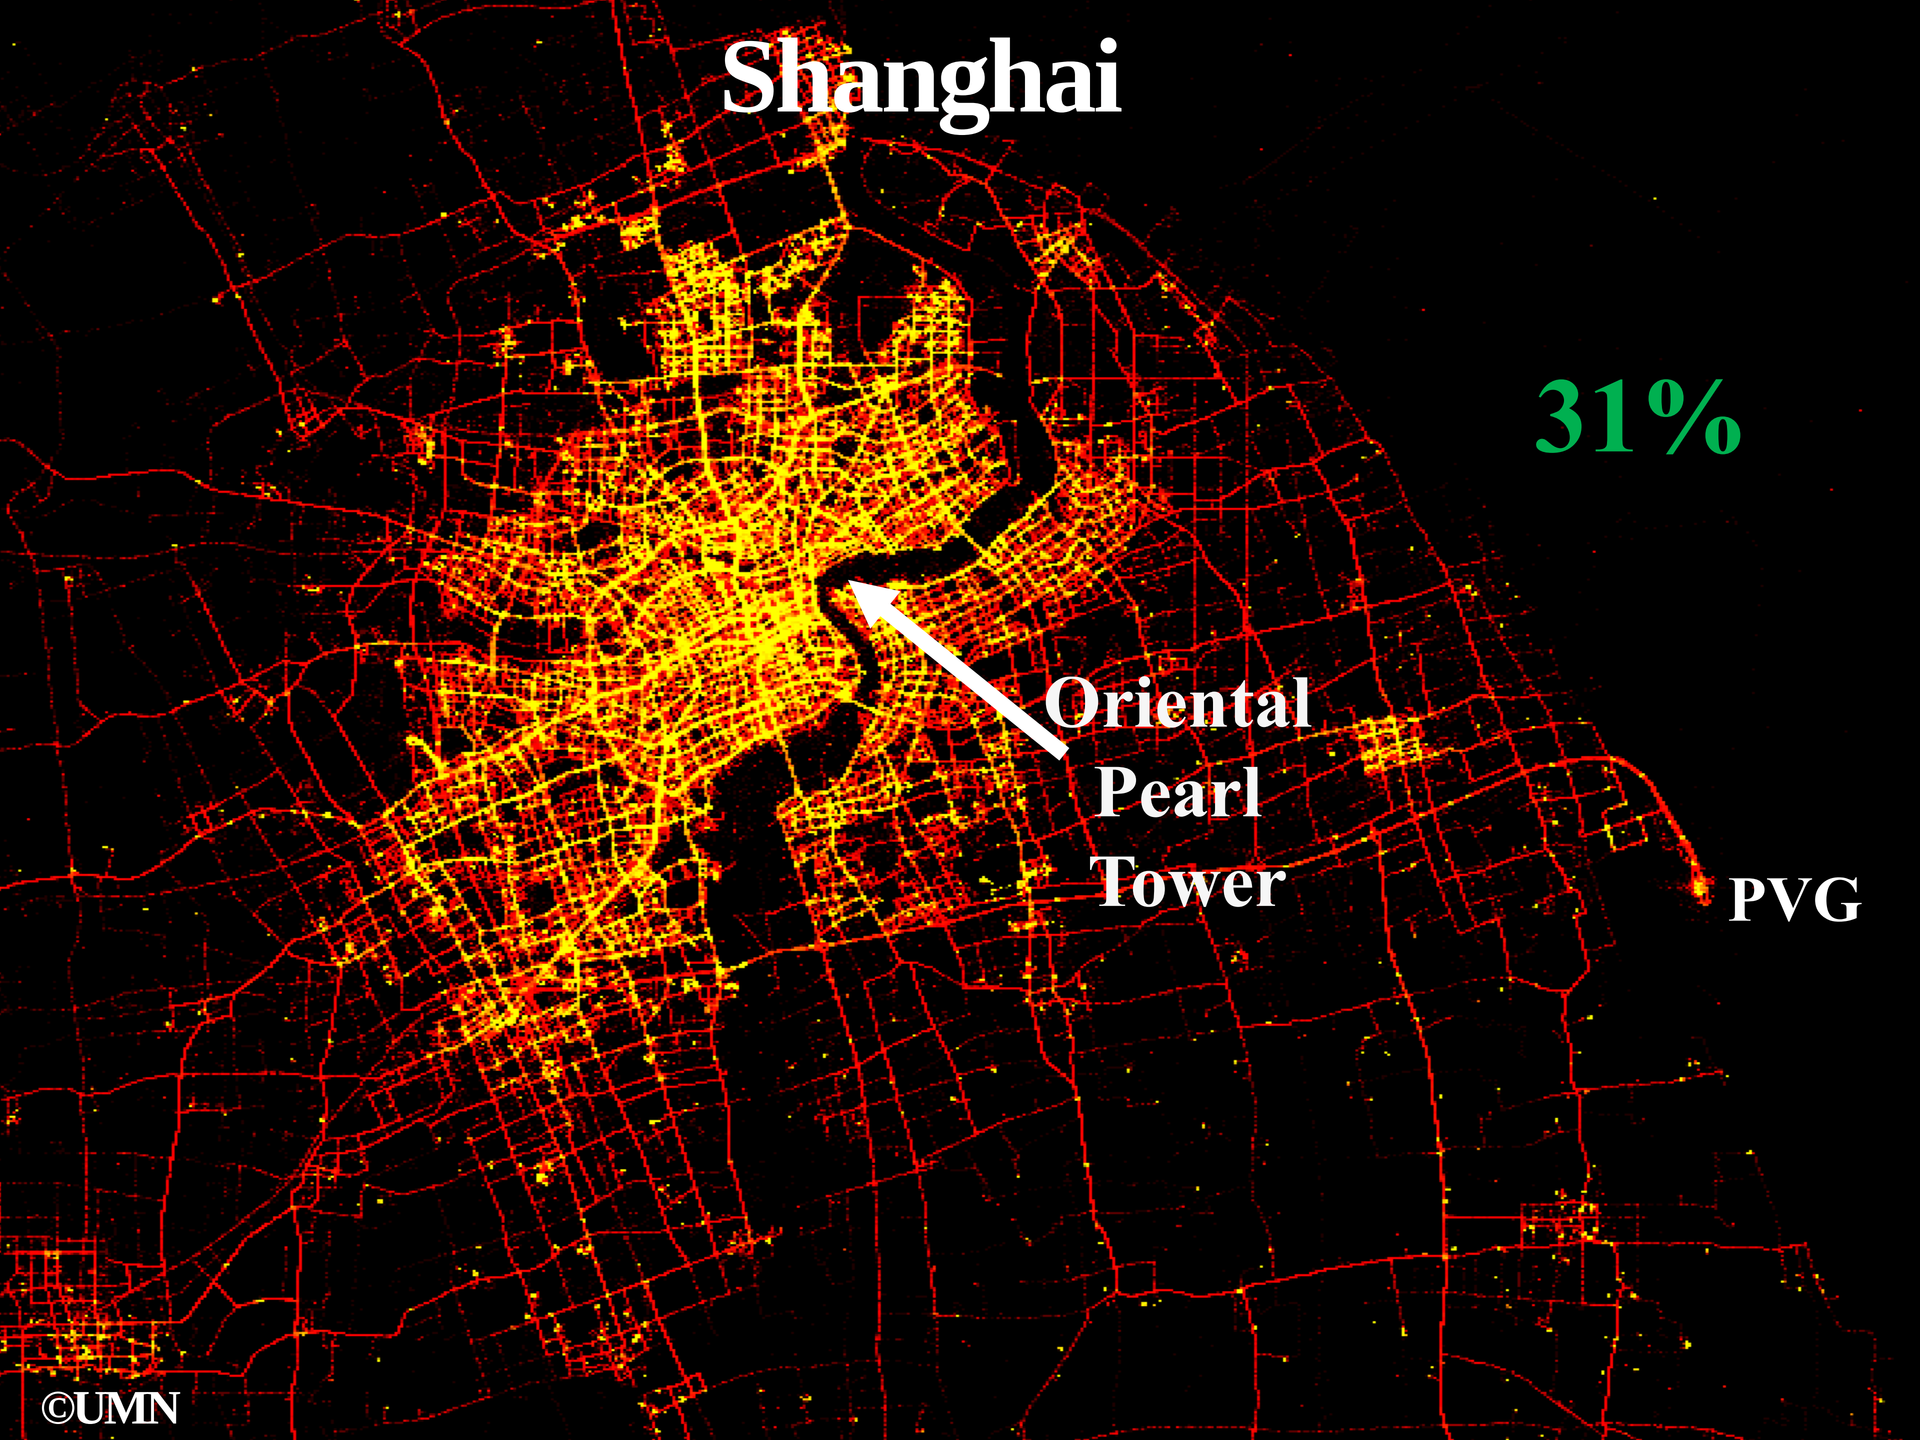
Shanghai
31%
Oriental
Pearl
Tower
PVG
©UMN


Shenzhen
Urban
Suburban
Electric
SZX
Downtown
41%
©UMN

Real World Impact 1: Dallas Taxi System

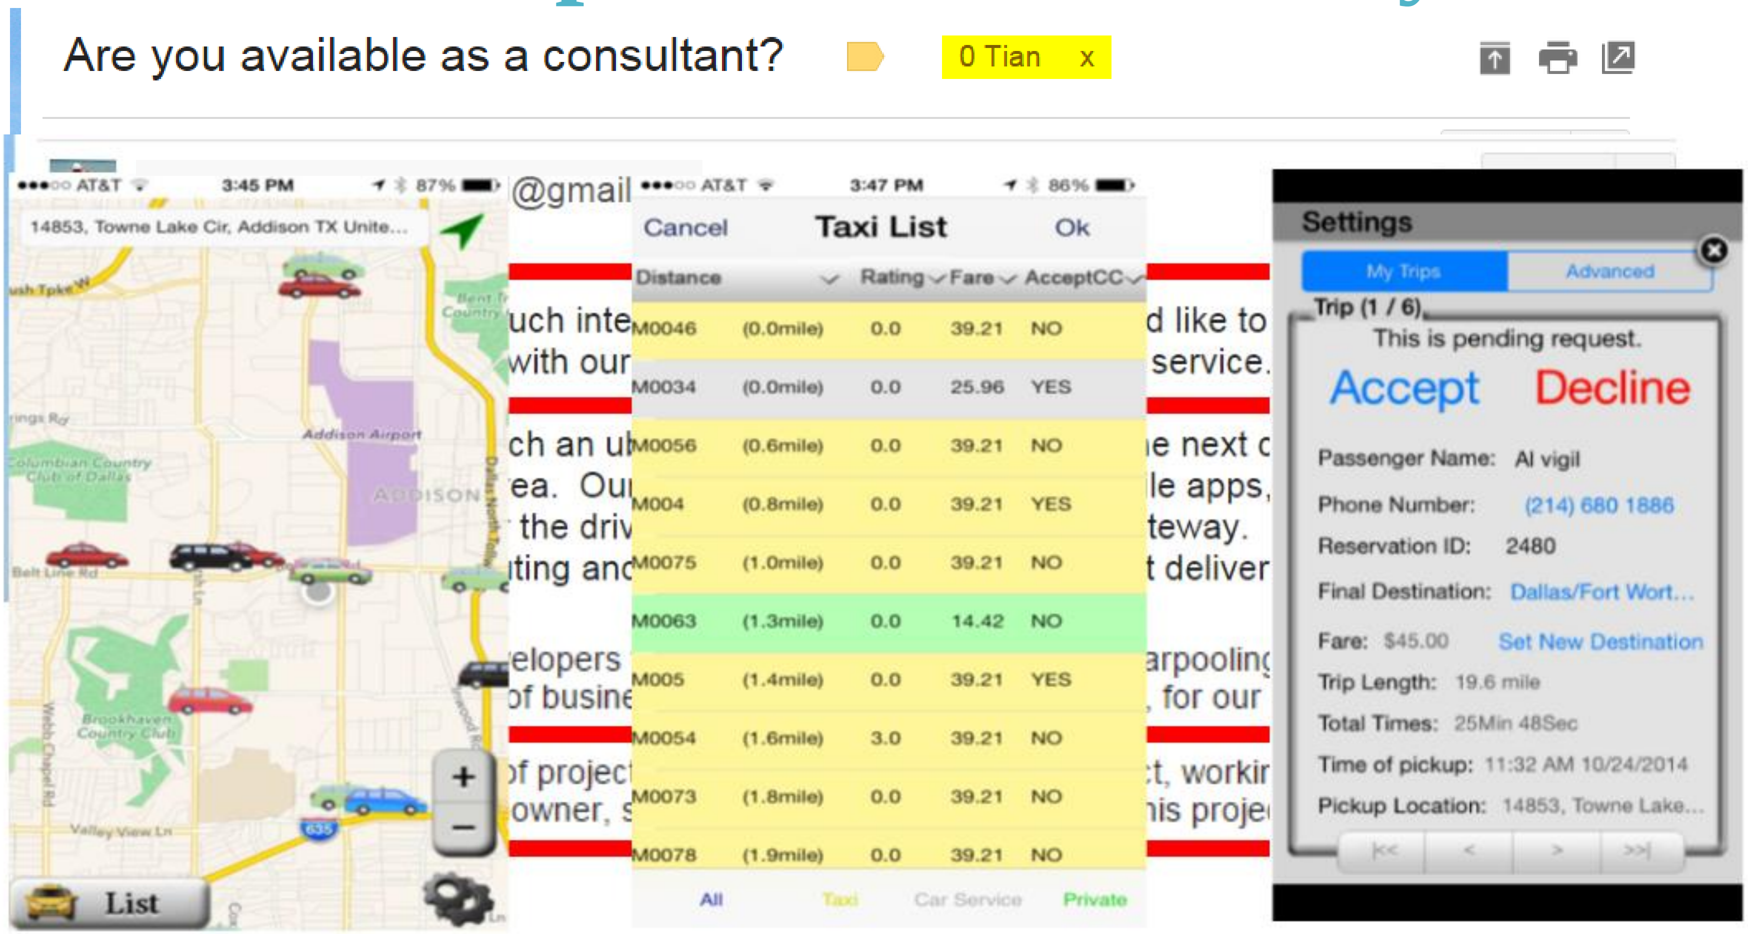

ARSENLIMO

- Working with Mr. Vigil in Pchomes. Inc
  - Practical Ridesharing in Dallas

Real World Impact 2: London Food Takeaway

A Ridesharing Problem for Food Takeout with Unique Challenges

- Traffic, Weather, Capacity, Timing, Food Types, Uncertain Origins

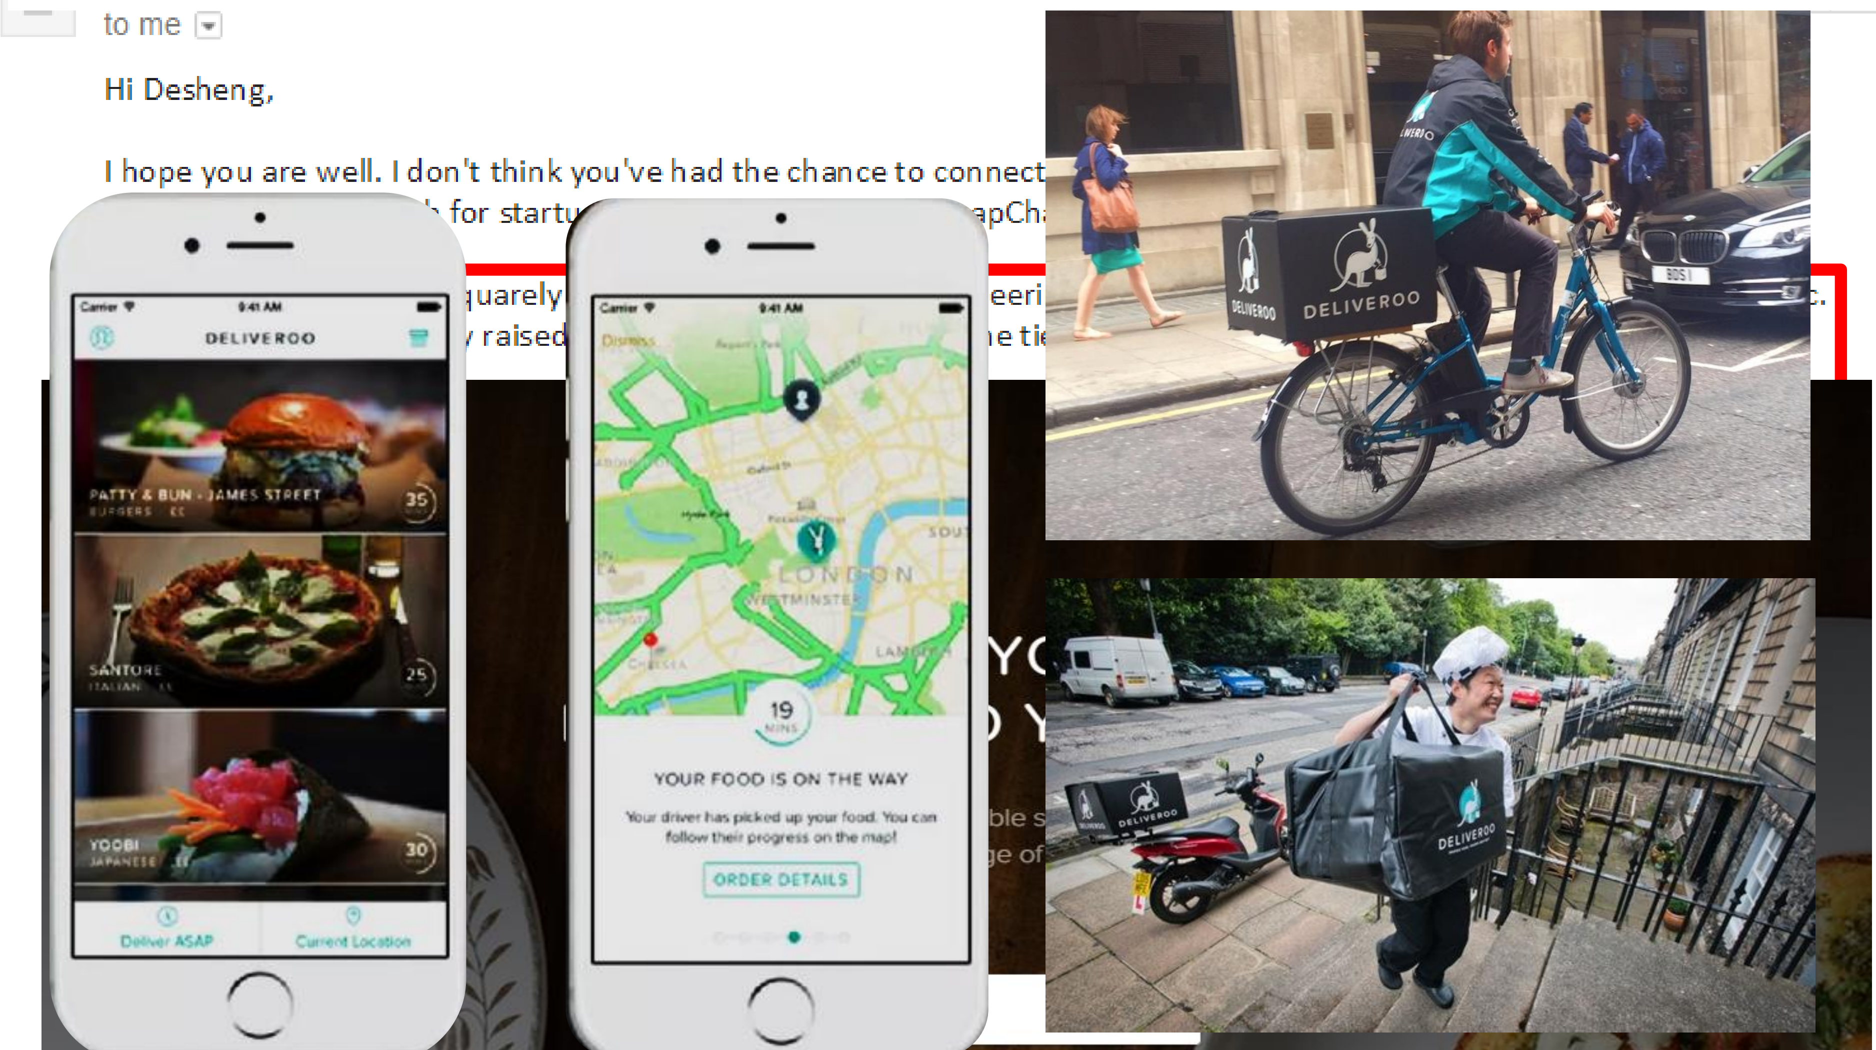

Big Picture: Improve Urban Efficiency with Sharing Economy

Contributions

- Resource Sharing Optimization Framework
- Cross-Domain Knowledge from Various Data-Driven Models
- Groundwork for Broader Logistics & Commercialization

Future Urban CPS Vision

3-5 Years

- Cross-Domain Modeling & Apps for Mobility
- Smart Cities Initiative
- Real-world Deployment for Research Impact

5-10 Years

- Data-Driven **Physical** Phenomenon Modeling & Apps
- Domains: **Energy, Pollution, Safety, Privacy**
- Homes, Buildings, and Cities, i.e., a **Connected Smart World**

Summary: Cross-Domain CPS for Smart Cities

Cross-Domain Knowledge
Optimization Framework
Potential Co
ercialization

Applications

Cross-Domain Feedback
Positive Feedback to
Improv
tems

Intellectual Core:
Cross-Domain Cyber-Physical Systems
Generalization:
Autonomous Driving;
Energy, Pollution, Safety;
Healthcare, Manufacturing

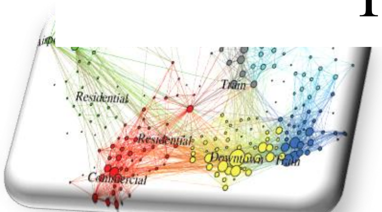

Models

Systems

Cross-Domain Data
Multi-View Bounding
Real-world Implementation

Bike Systems

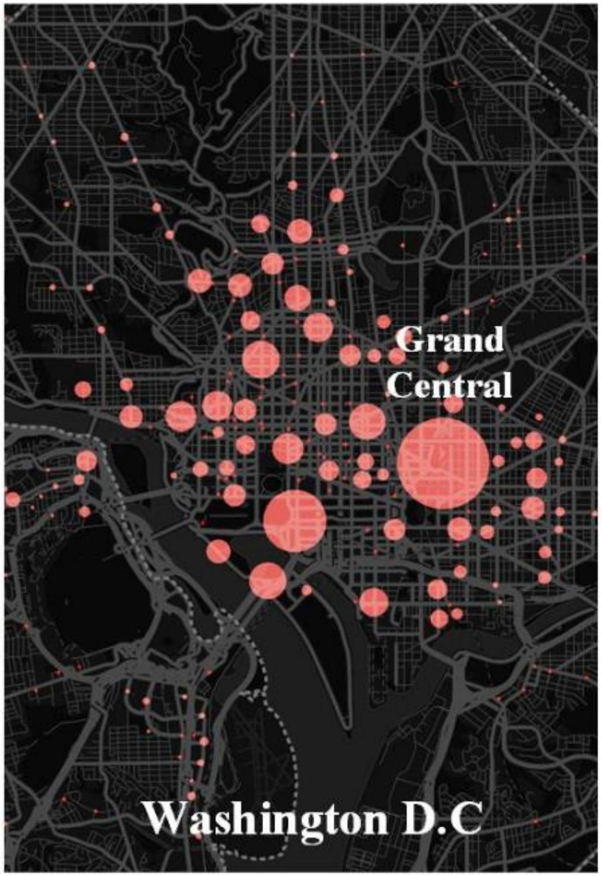

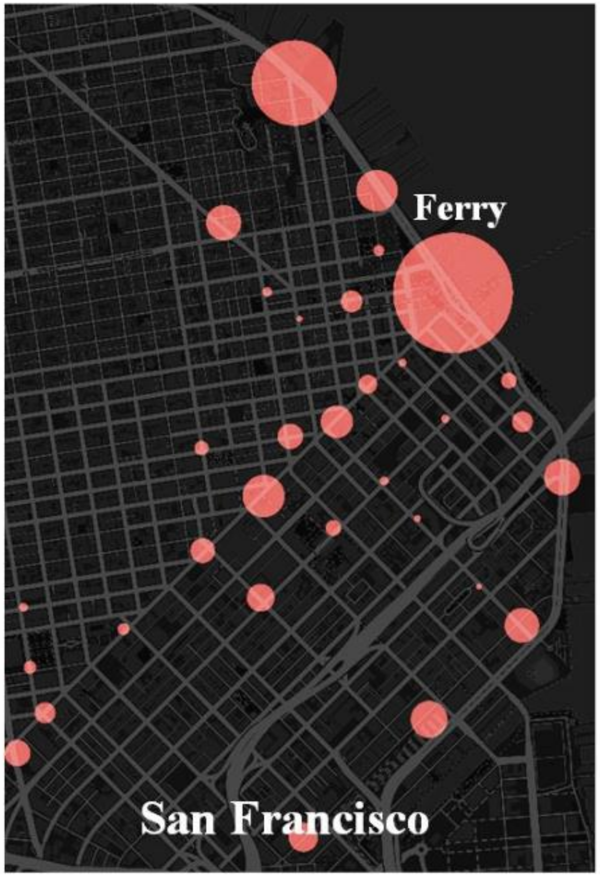

Taxi Systems

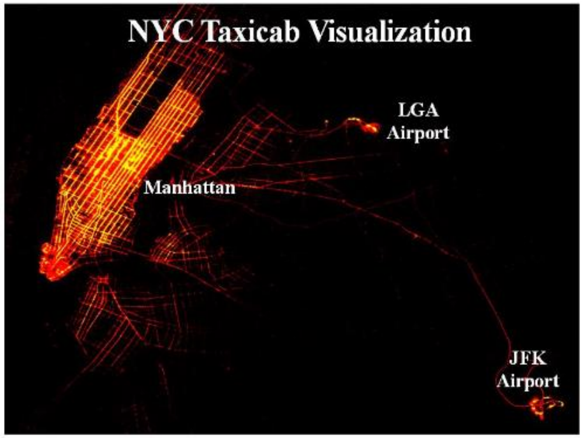

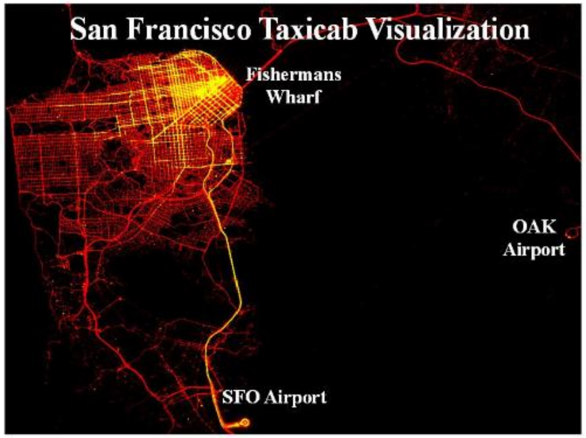

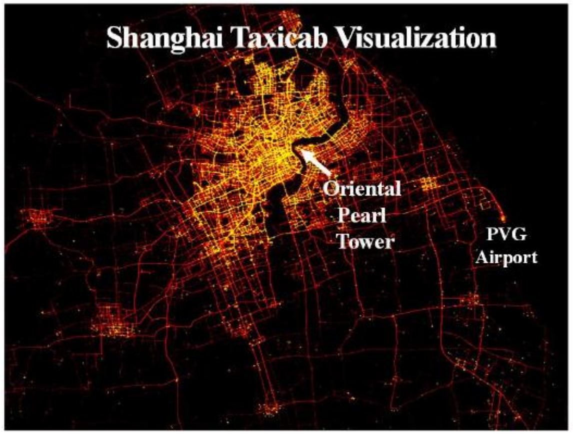

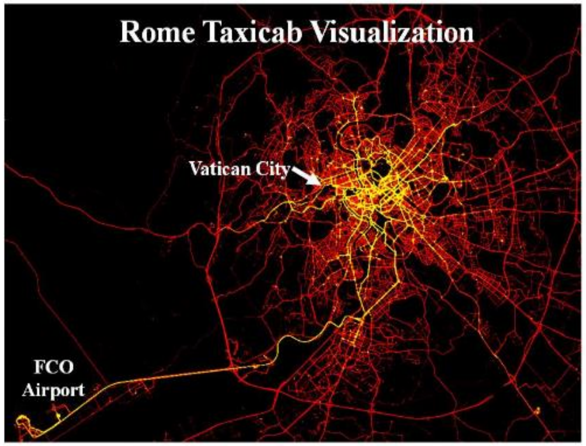

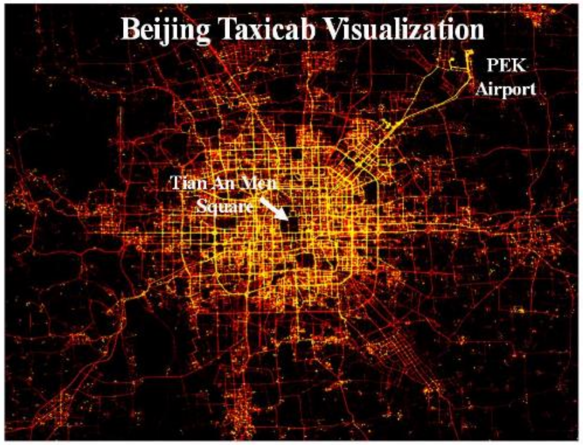

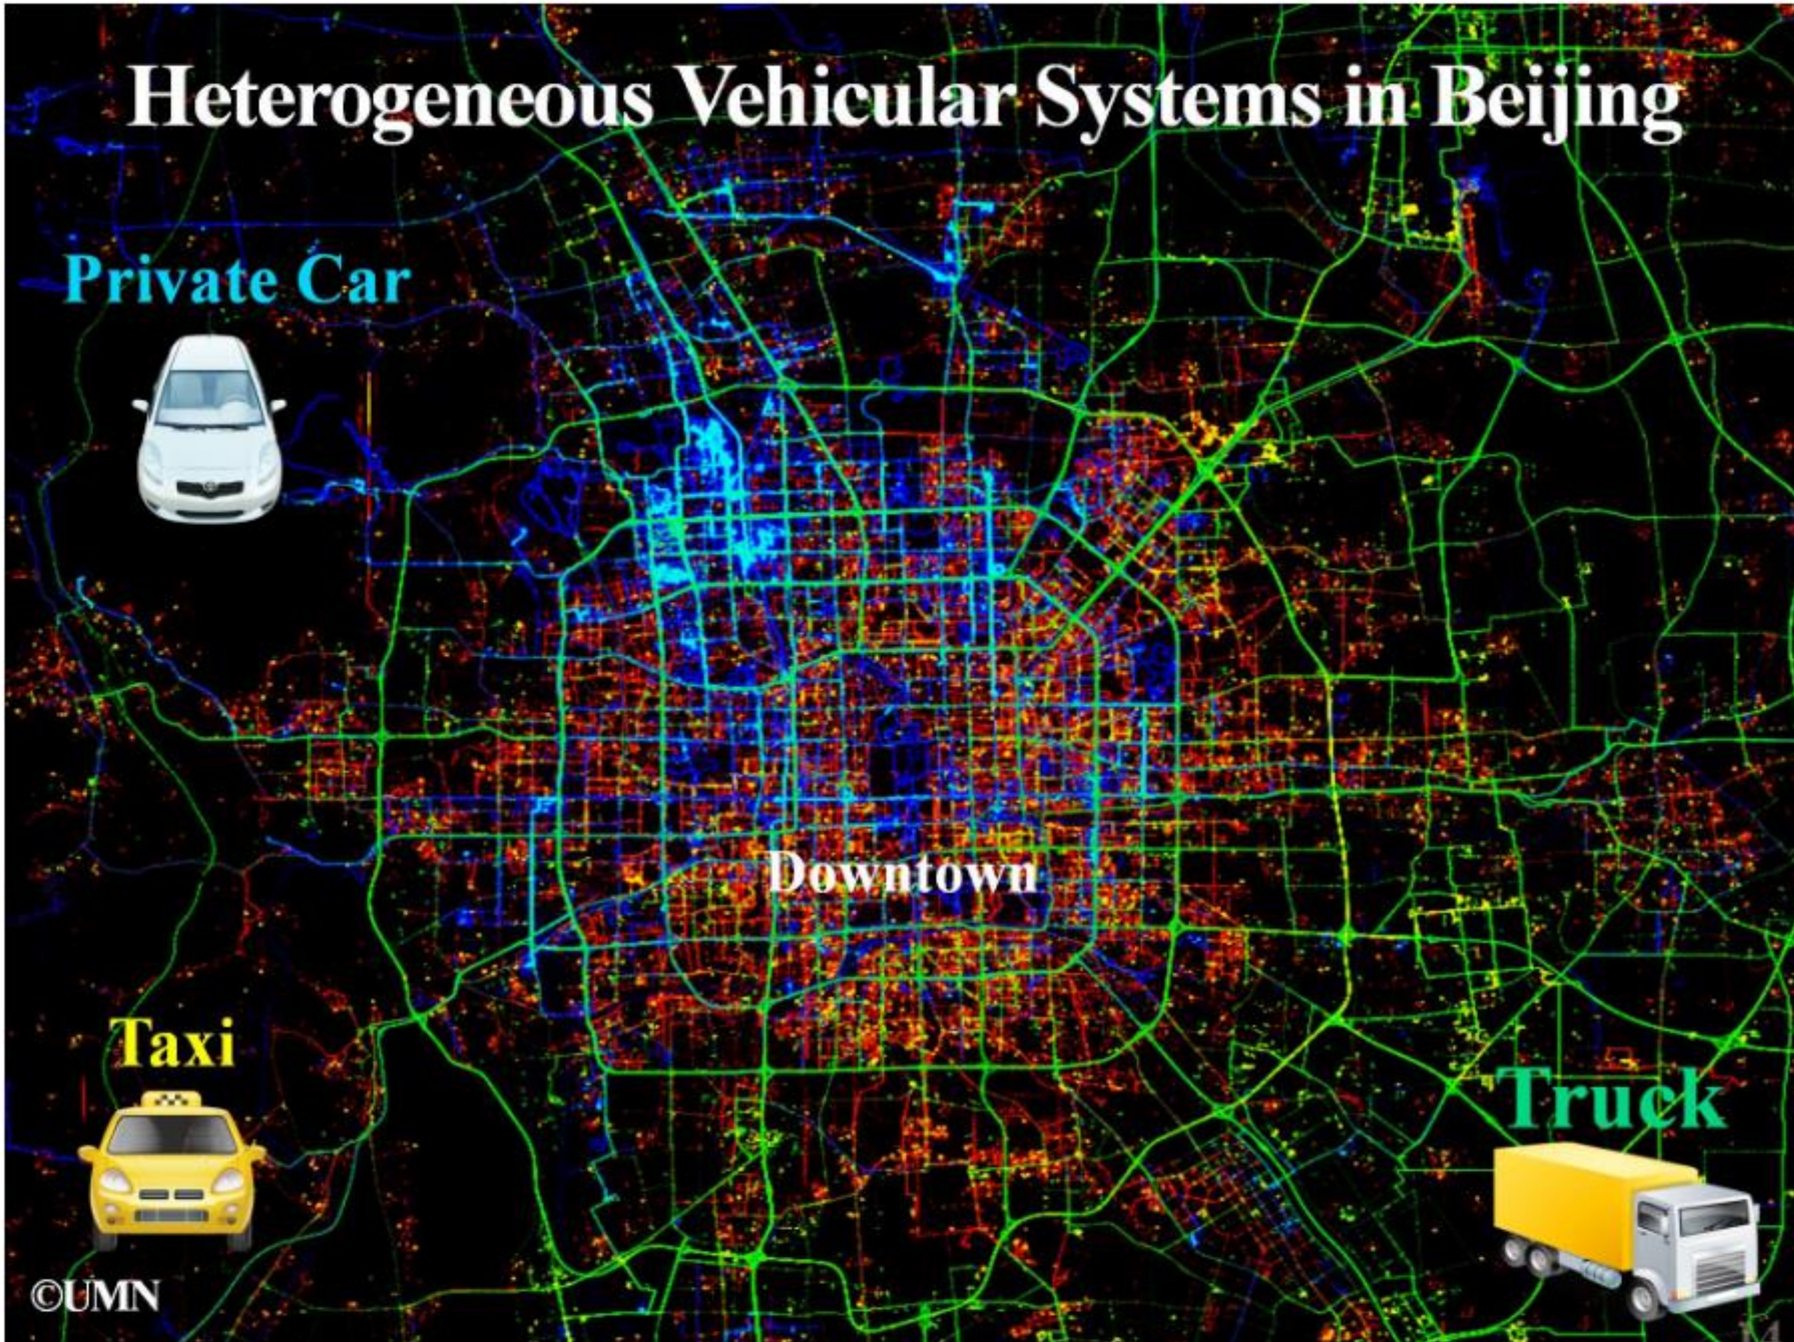
Heterogeneous Vehicular Systems in Beijing
Private Car
Downtown
Taxi
Truck
©UMN
14

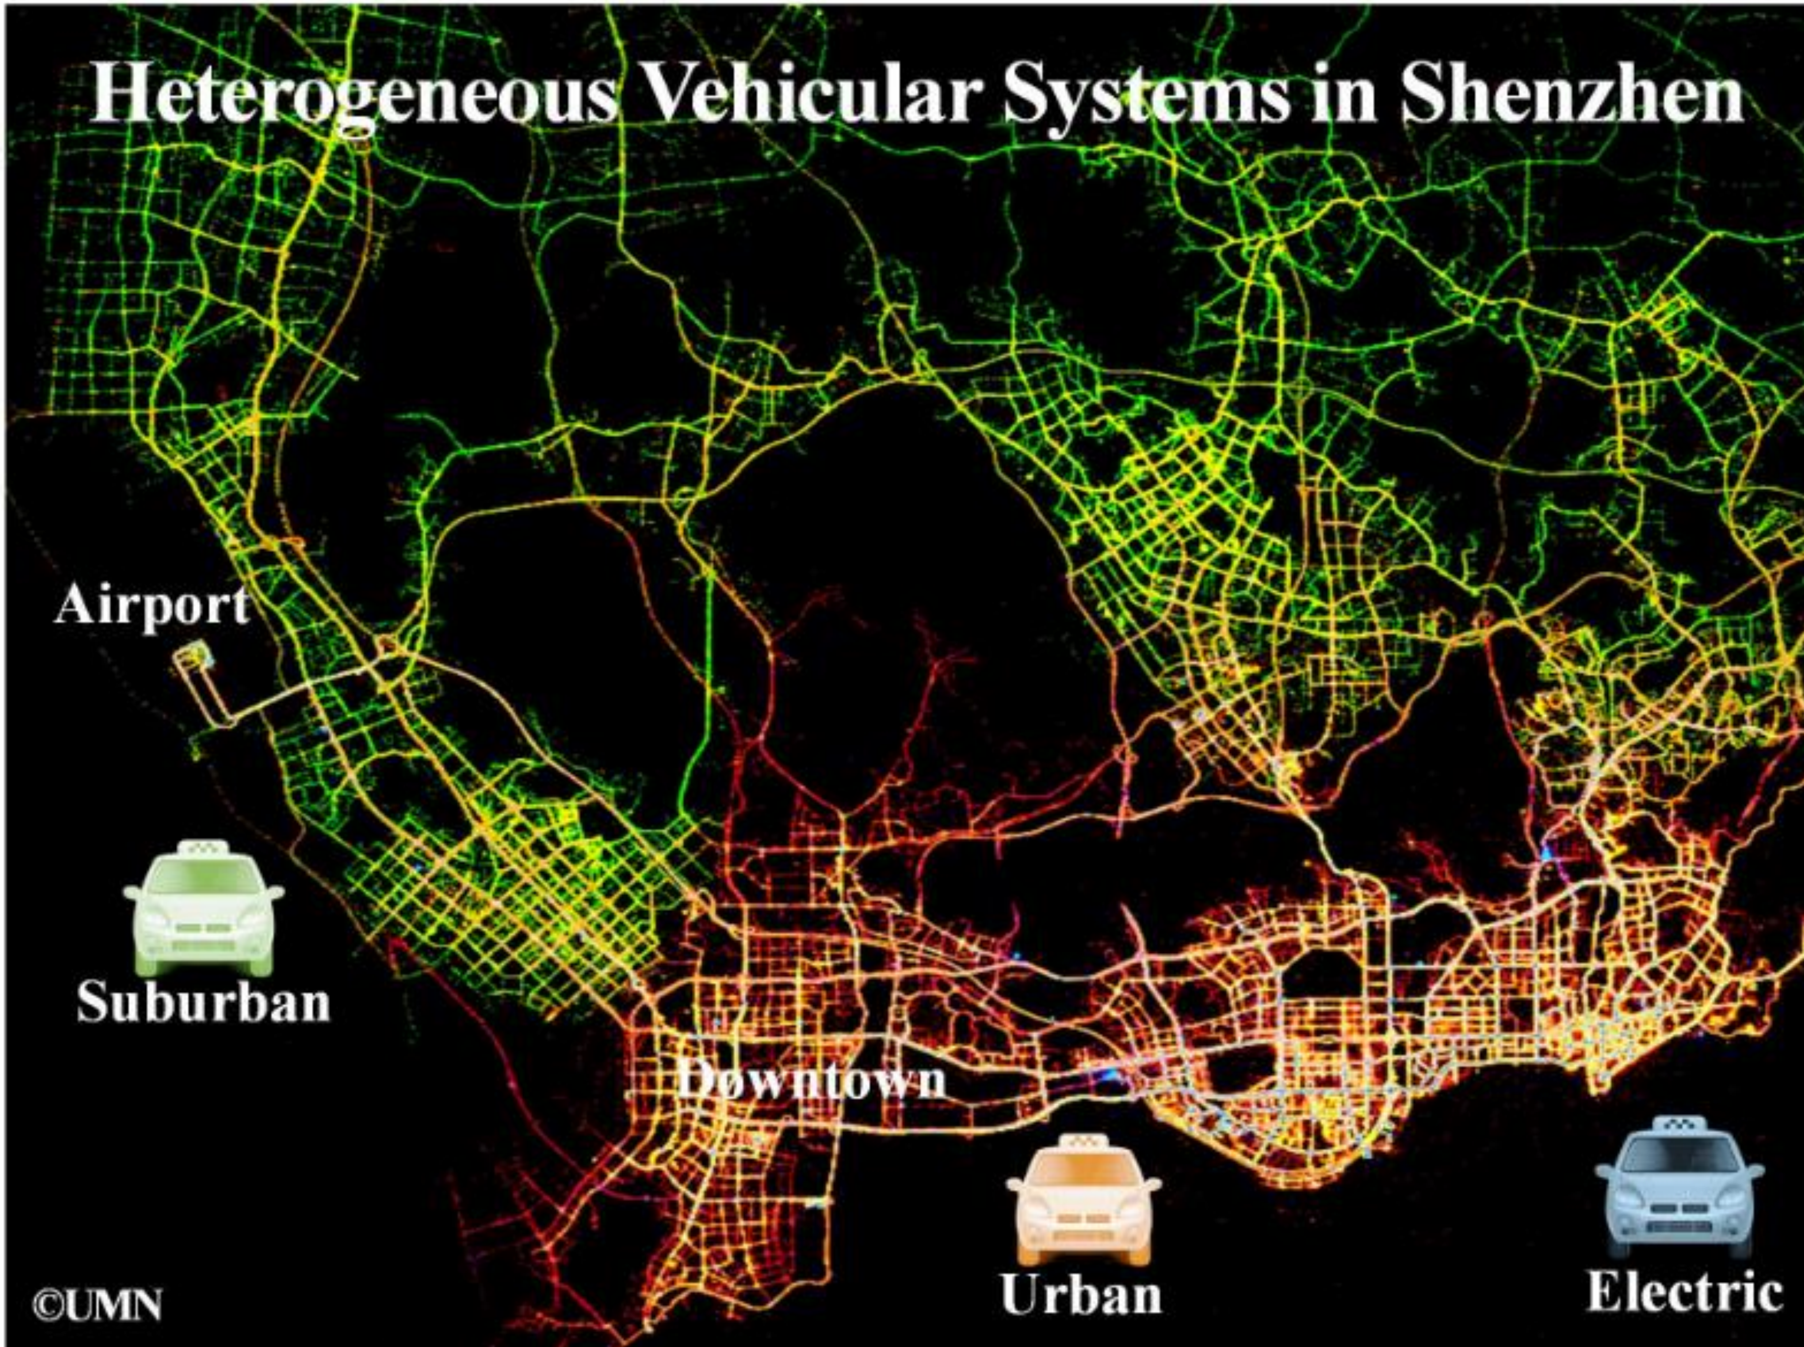
Heterogeneous Vehicular Systems in Shenzhen
Airport
Suburban
Downtown
Urban
Electric
©UMN

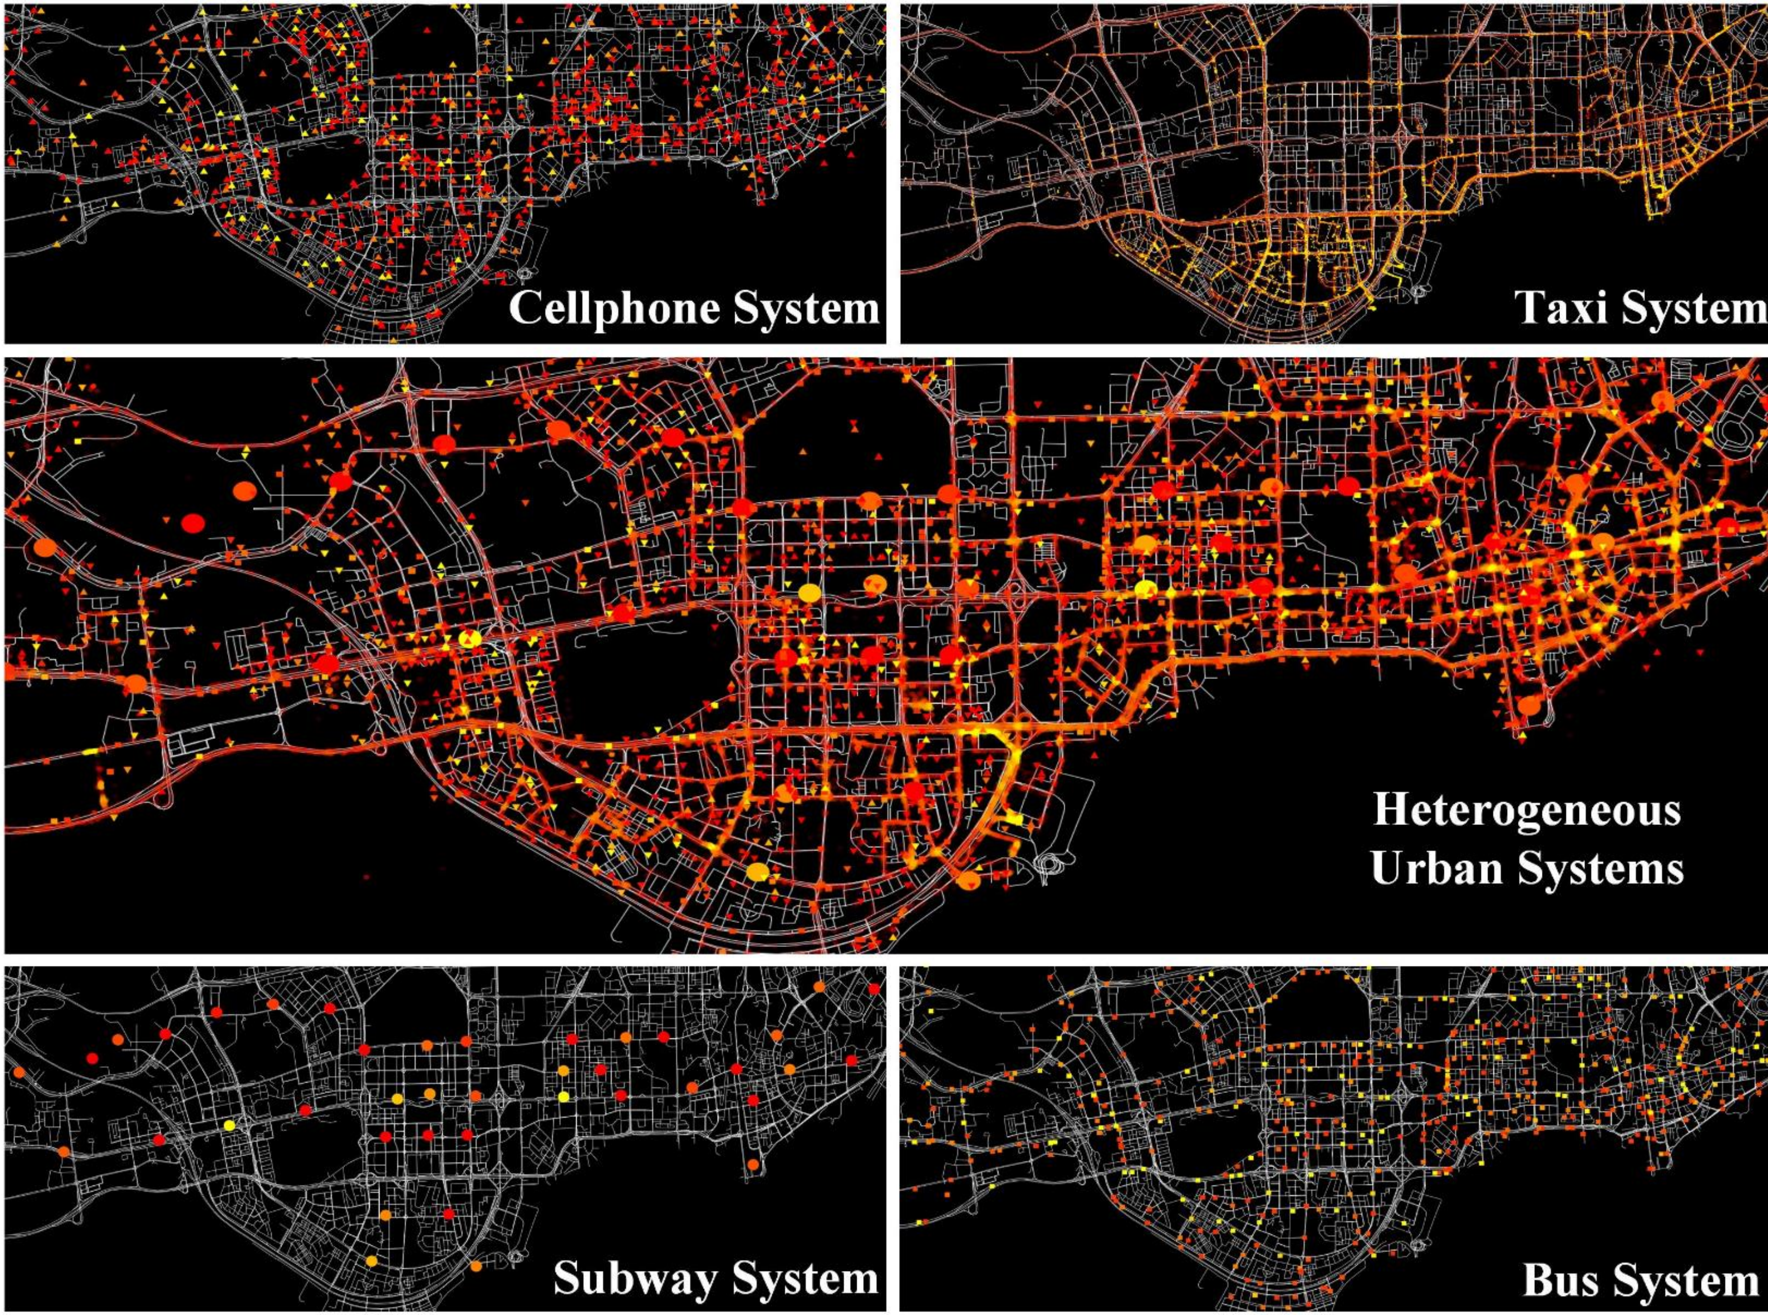
Cellphone System
Taxi System
Heterogeneous
Urban Systems
Subway System
Bus System

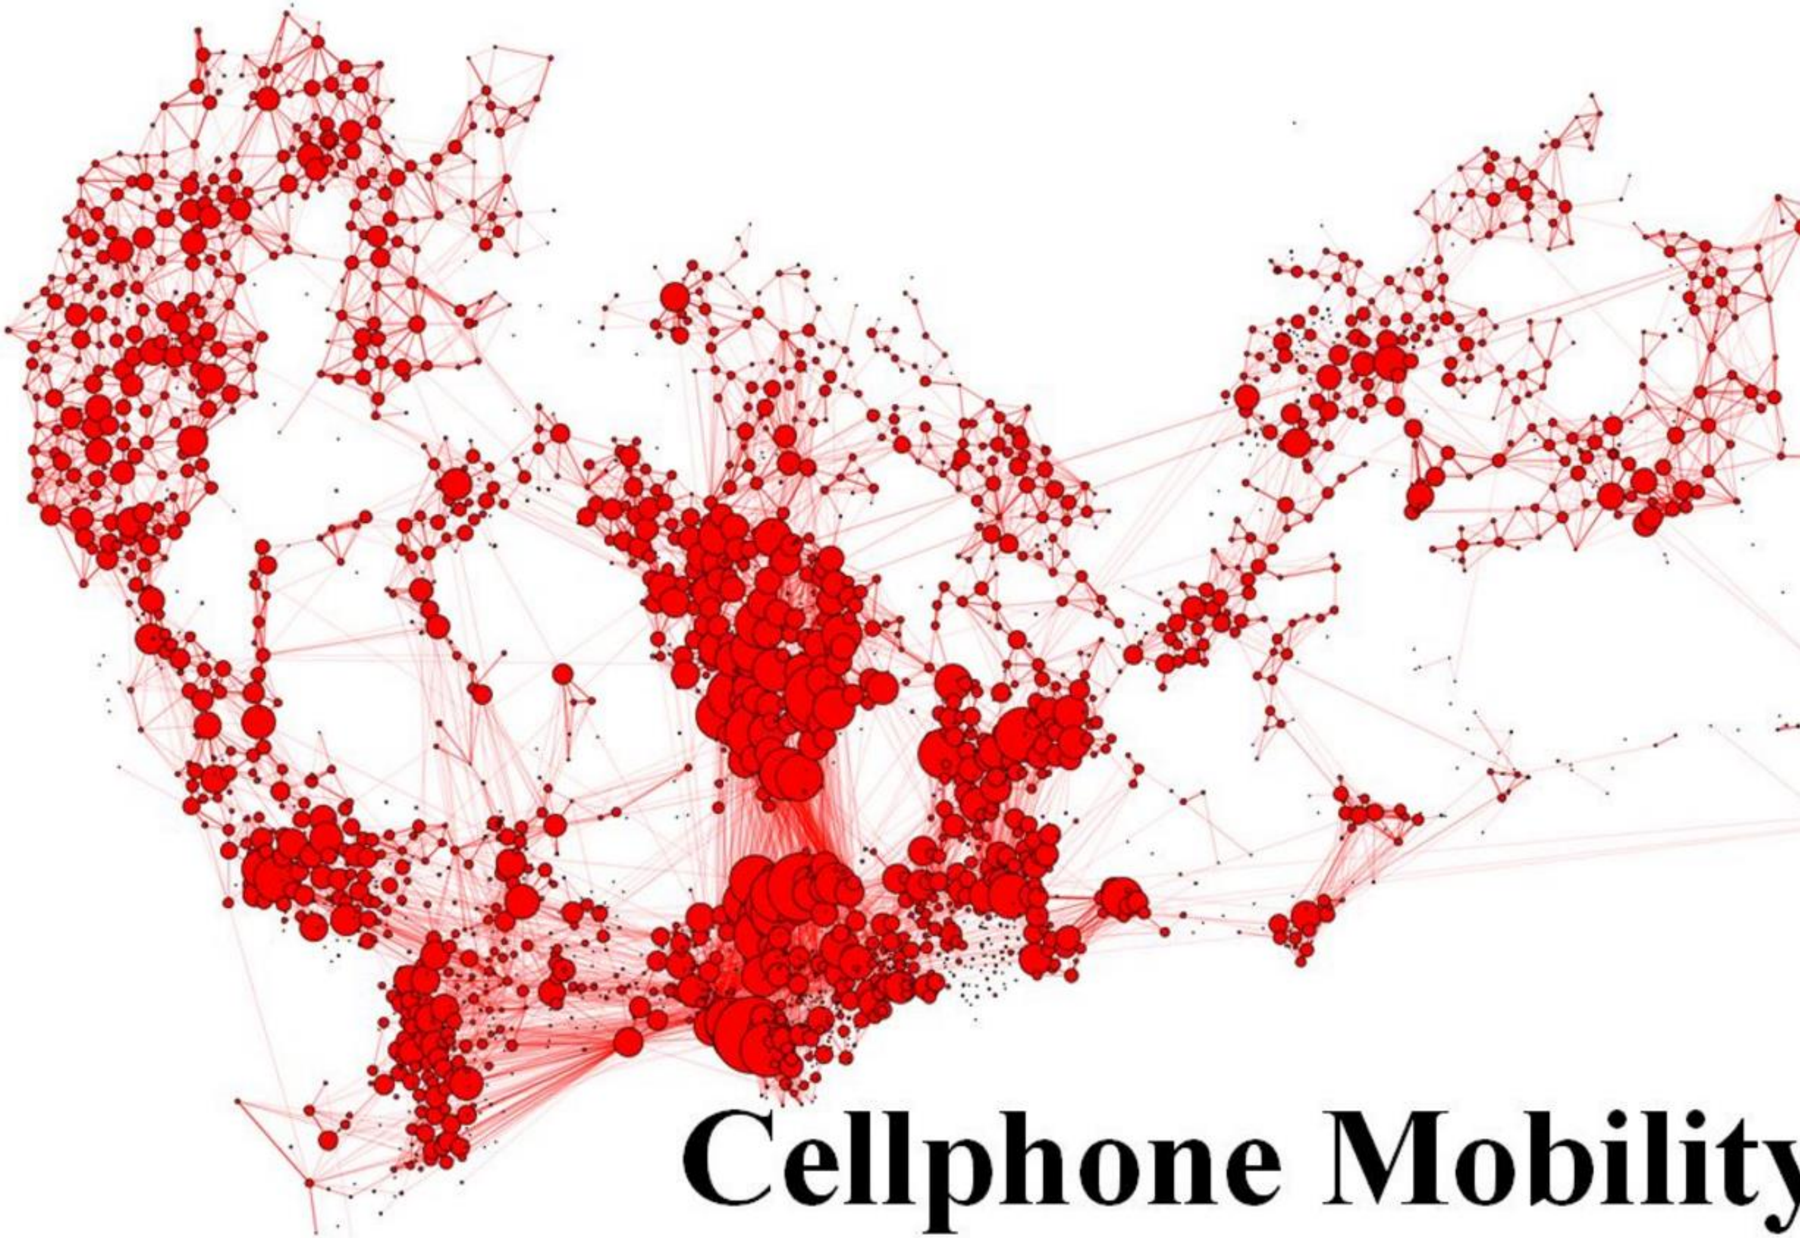

Cellphone Mobility

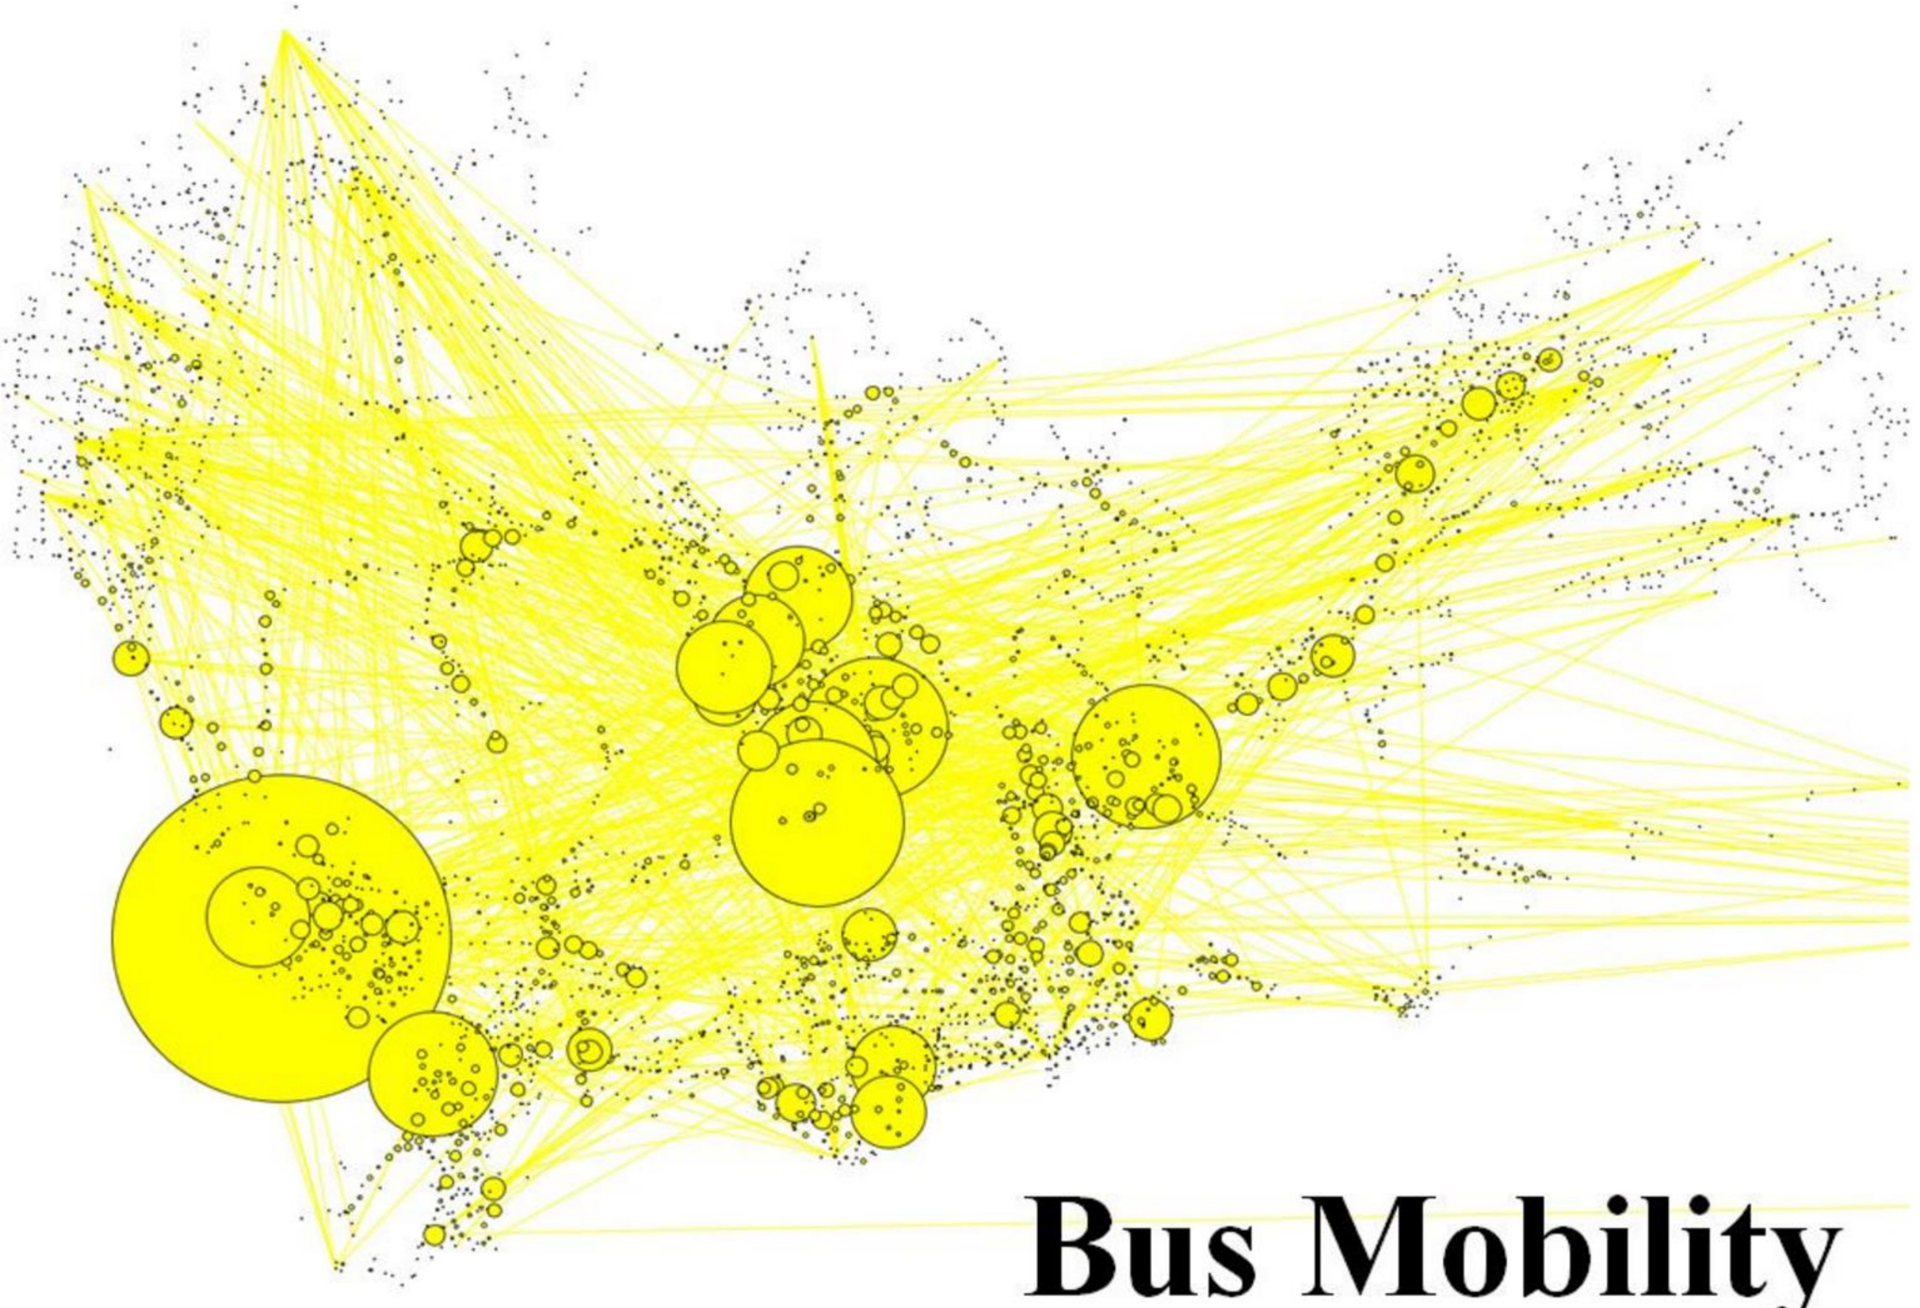

Bus Mobility

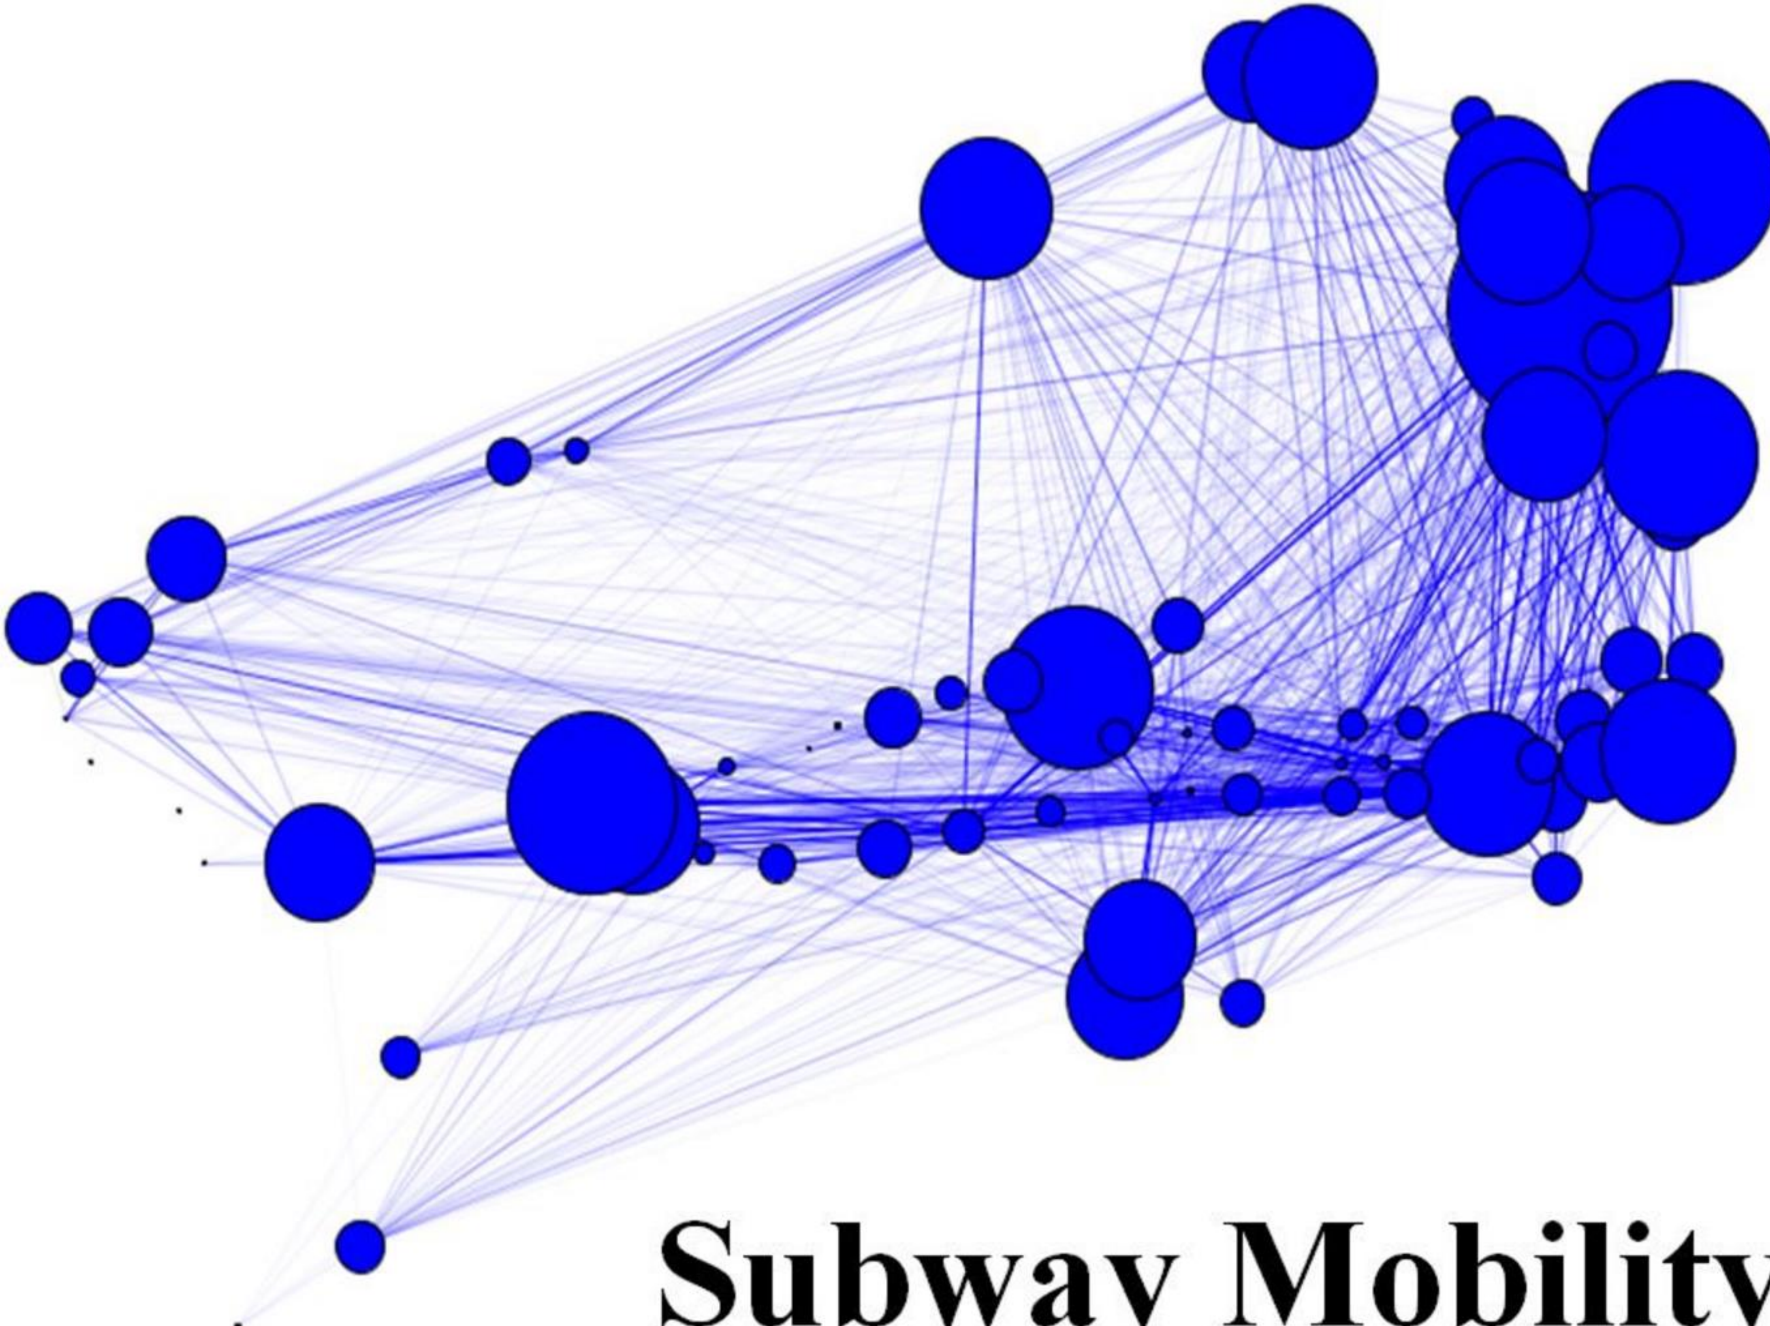

Subway Mobility

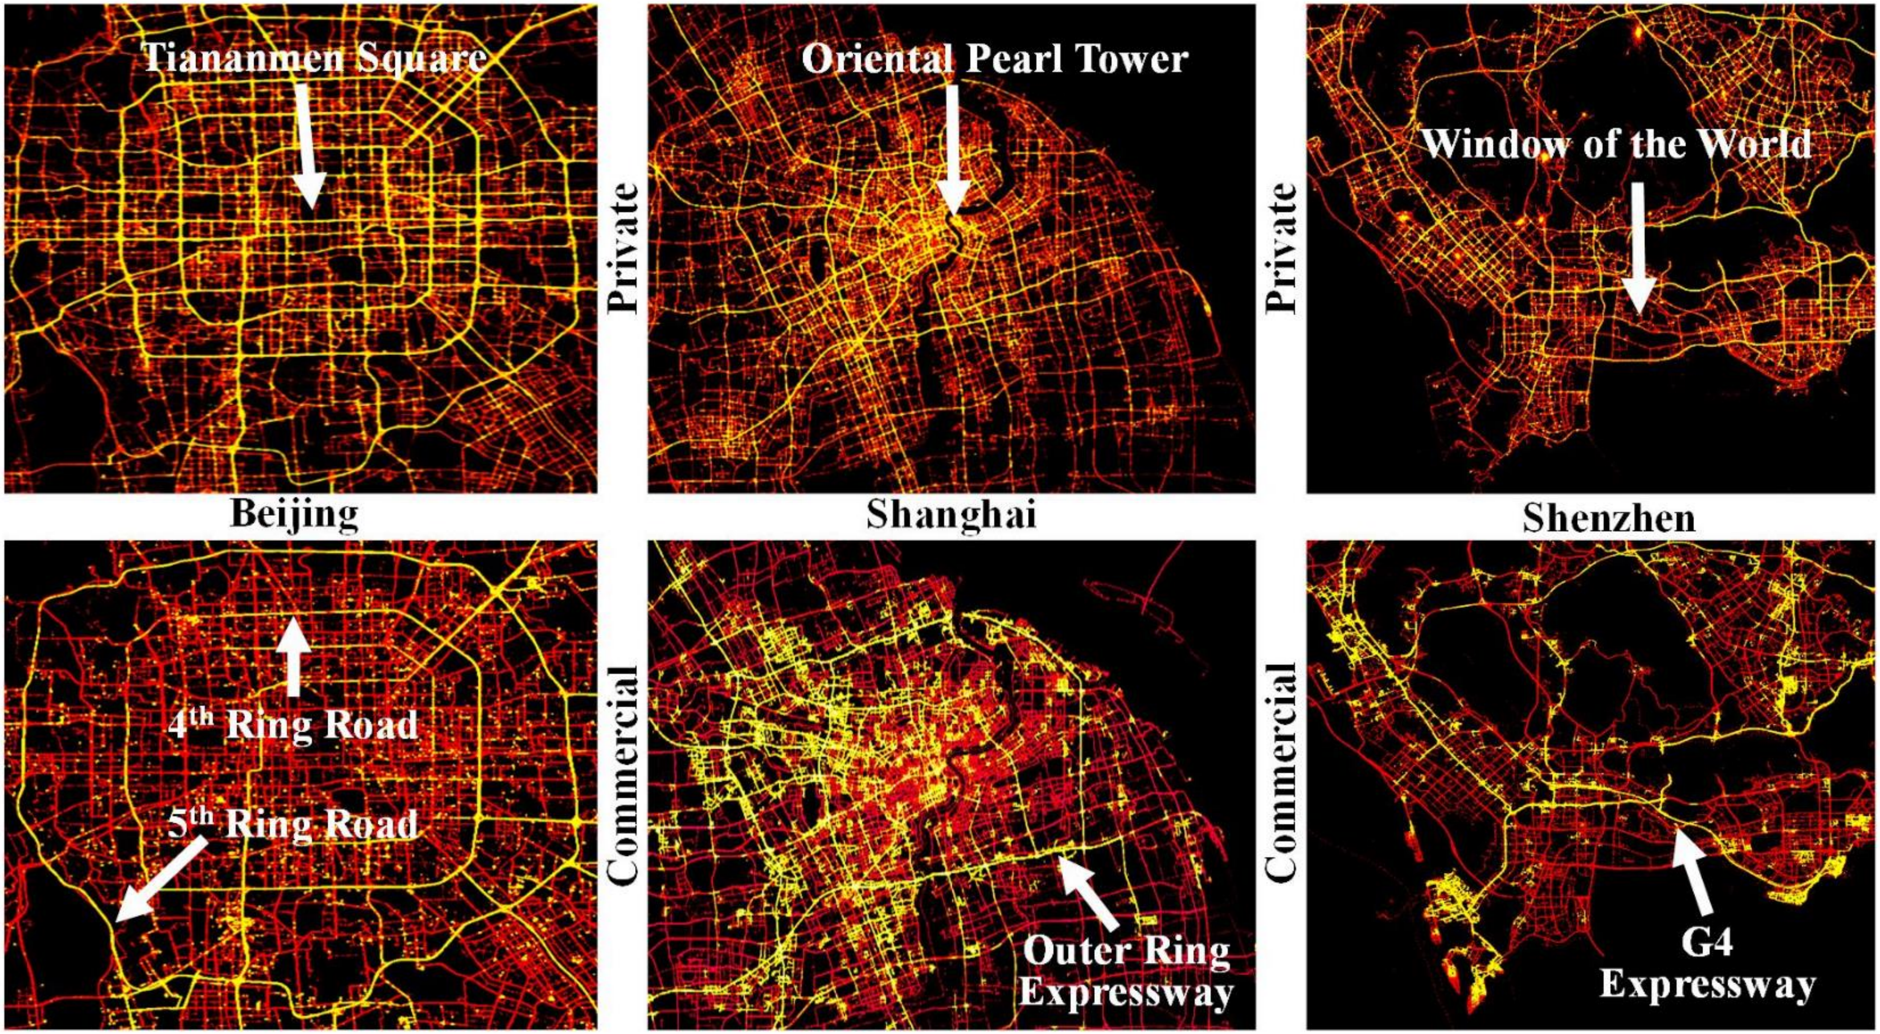

National Heterogeneous Systems

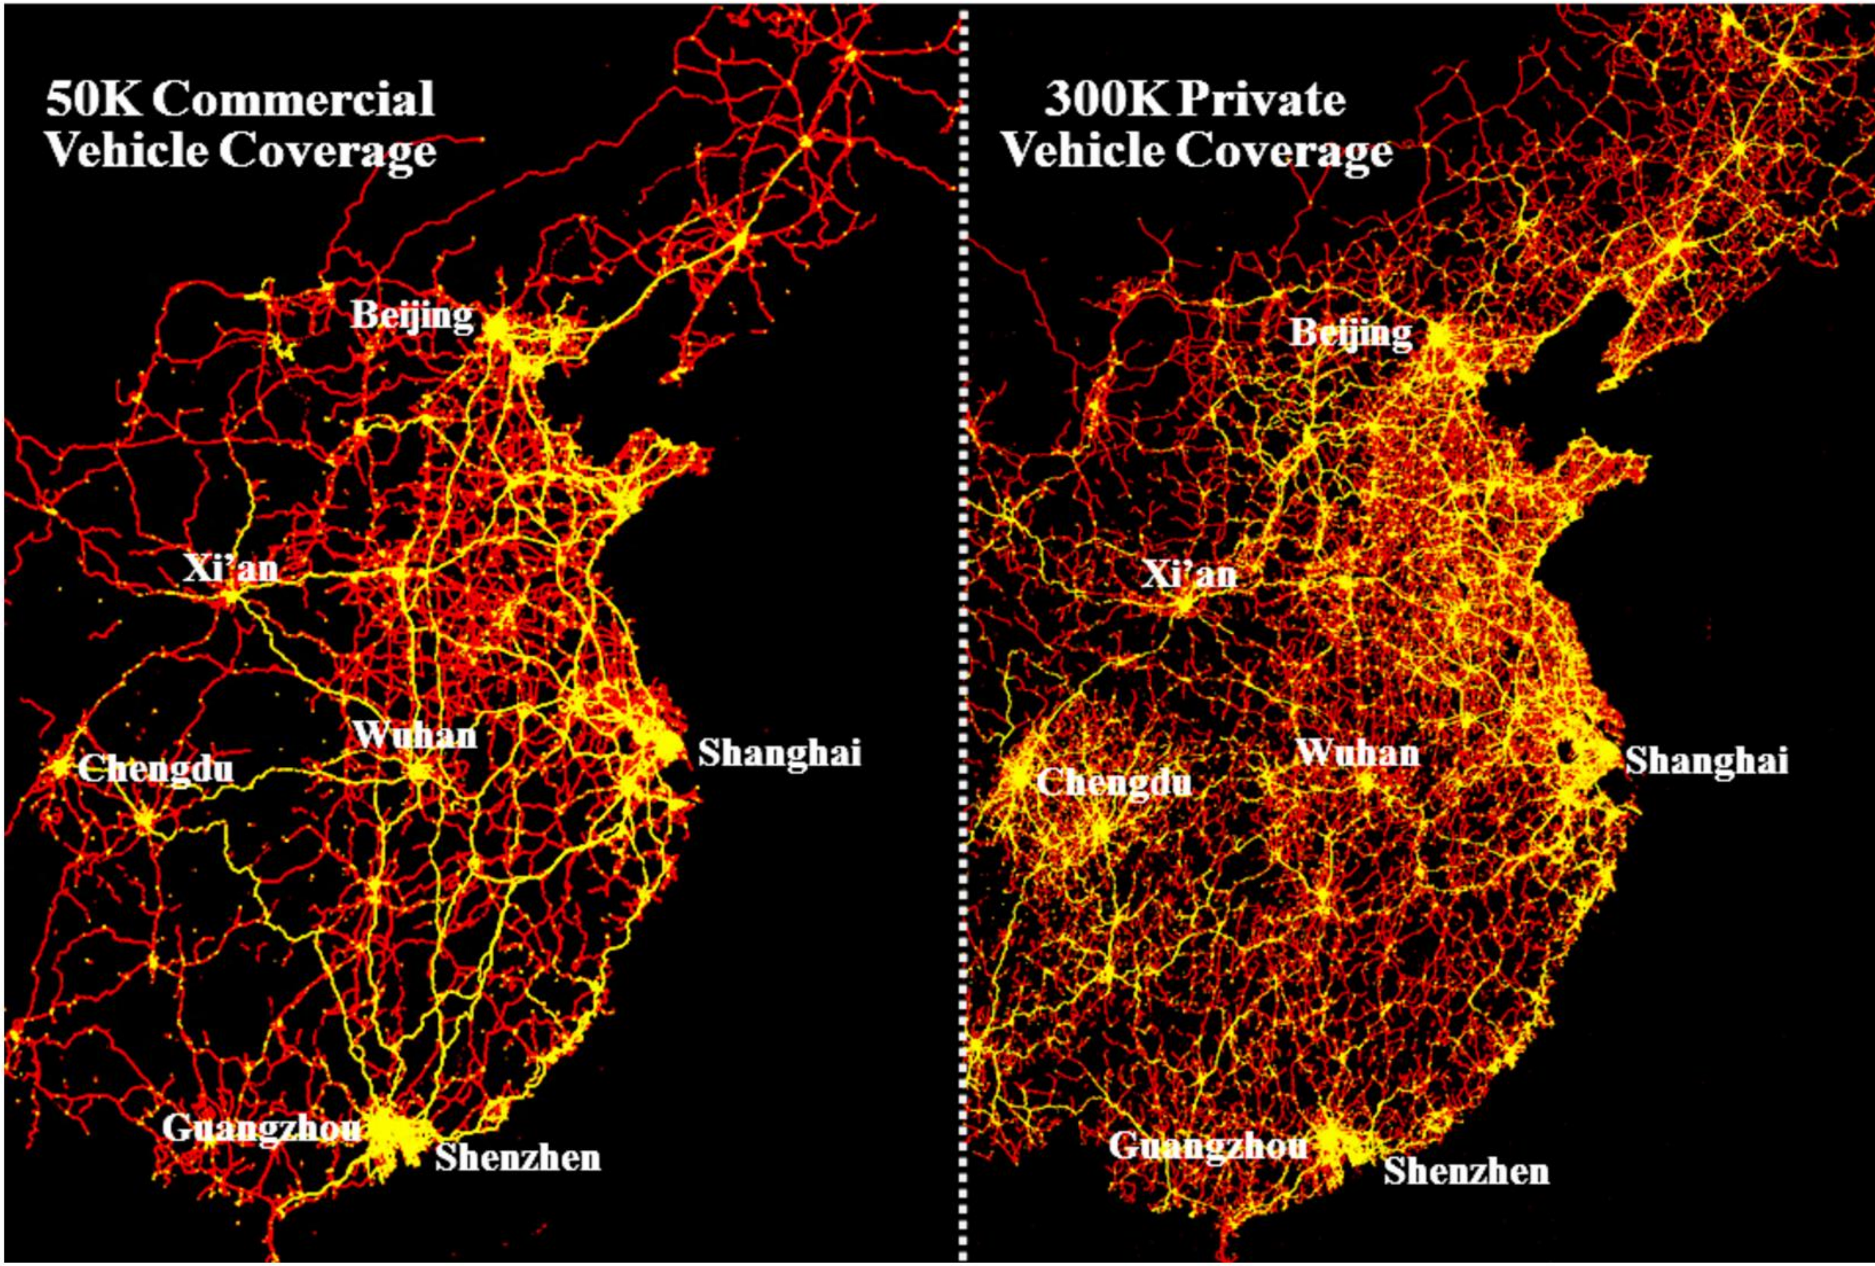
50K Commercial
Vehicle Coverage
300K Private
Vehicle Coverage
Beijing
Xi'an
Wuhan
Chengdu
Shanghai
Guangzhou
Shenzhen
Beijing
Xi'an
Wuhan
Chengdu
Shanghai
Guangzhou
Shenzhen

Thanks

Data and More Work @ https://www.cs.rutgers.edu/~dz220/

Outline

- Logistics

- Course Structure

- Overview of Smart Cities Research

- **Class Schedule**

1. Overview

Week	Date	Topics and Reading Assignments
1	Jan 19	General Class Introduction Suggested Reading: • Urban Computing: Concepts, Methodologies, and Applications • Cyber-Physical Systems: Executive Summary • Research Directions for the Internet of Things • Systems Computing Challenges in the Internet of Things

Urban Computing: Concepts, Methodologies, and Applications

YU ZHENG, Microsoft Research
LICIA CAPRA, University College London
OURI WOLFSON, University of Illinois at Chicago
HAI YANG, Hong Kong University of Science and Technology

(a) Motivation: Big cities, data and challenges

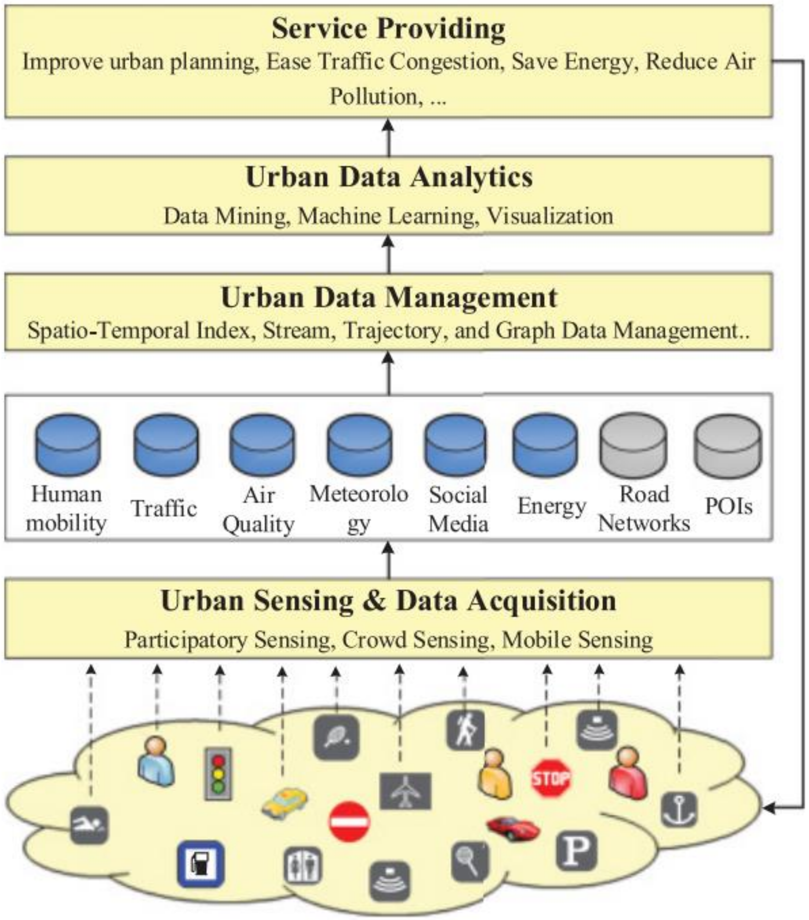

2. Urban Sensing

2	Jan 26	Urban Sensing Assigned Reading: • How Long to Wait?: Predicting Bus Arrival Time with Mobile Phone based Participatory Sensing • ParkNet: Drive-by Sensing of Road-Side Parking Statistics • VTrack: Accurate, Energy-aware Road Traffic Delay Estimation Using Mobile Phones • Discovering Regions of Different Functions in a City Using Human Mobility and POIs Suggested Reading: • People-Centric Urban Sensing

ParkNet: Drive-by Sensing of Road-Side Parking Statistics

(a)

(b)

(c)

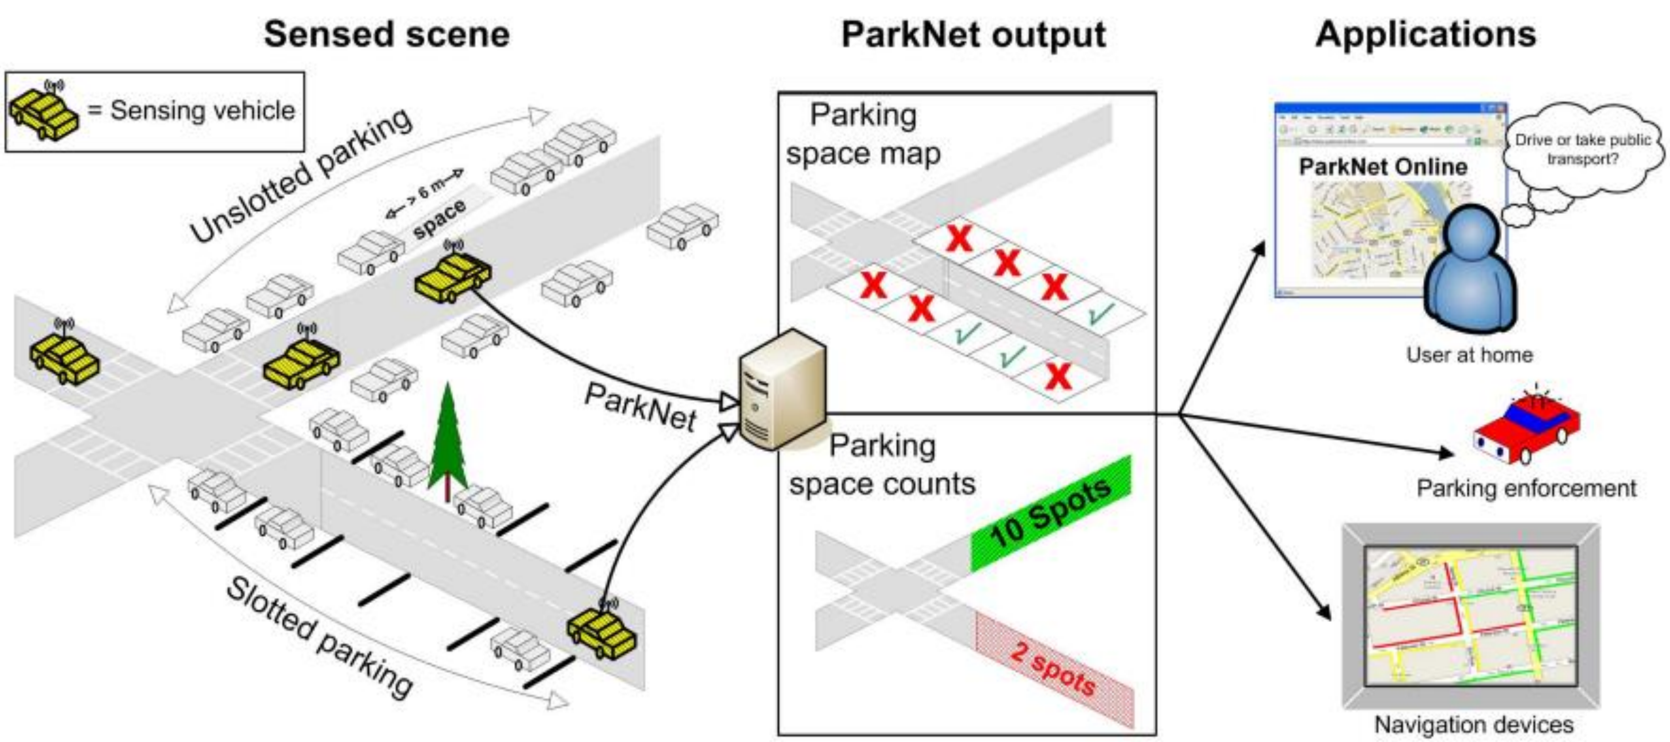

3. Data Management & Processing

3	Feb 2	Data Management & Processing Assigned Reading: • Managing Massive Trajectories on the Cloud • TrajStore: An Adaptive Storage System for Very Large Trajectory Data • Naiad: A Timely Dataflow System • Pyro: A Spatial-Temporal Big-Data Storage System Suggested Reading: • SpatialHadoop: A MapReduce Framework for Spatial Data

Managing Massive Trajectories on the Cloud

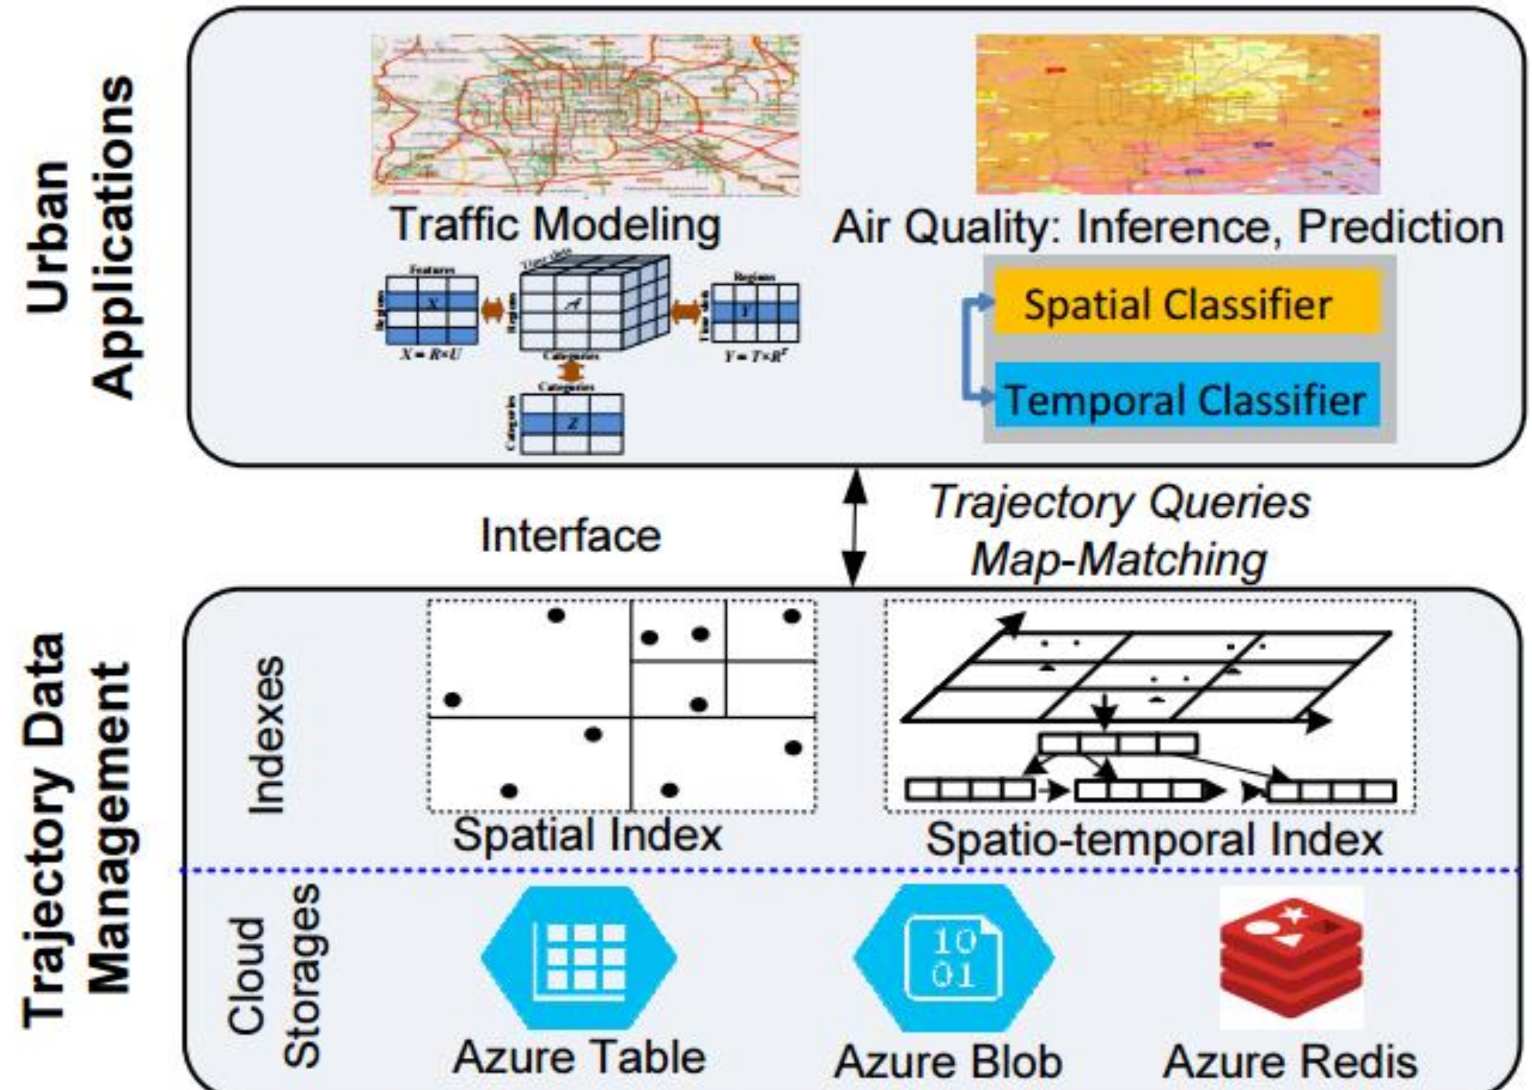

Figure 1: Motivation Scenarios.

4. Data-Driven Modeling: Human Mobility

4	Feb 9	Data-Driven Modeling: Human Mobility Assigned Reading: • Human Mobility Modeling at Metropolitan Scales • Inferring Human Mobility Patterns from Taxicab Location Trace • Reconstructing Individual Mobility from Smart Card Transactions: A Space Alignment Approach • Mobility Modeling and Prediction in Bike-Sharing Systems Suggested Reading: • Trajectory Data Mining: An Overview • A survey on Human Mobility and its Applications • Human Mobility Characterization from Cellular Network Data • A Review of Urban Computing for Mobile Phone Traces

Human Mobility Modeling at Metropolitan Scales

Sibren Isaacman*, Richard Becker†, Ramón Cáceres†, Margaret Martonosi*,
James Rowland†, Alexander Varshavsky†, Walter Willinger†
*Princeton University, Princeton, NJ, USA †AT&T Labs, Florham Park, NJ, USA
isaacman@princeton.edu, {rab,ramon}@research.att.com, mrm@princeton.edu,
{jrr,varshavsky,walter}@research.att.com

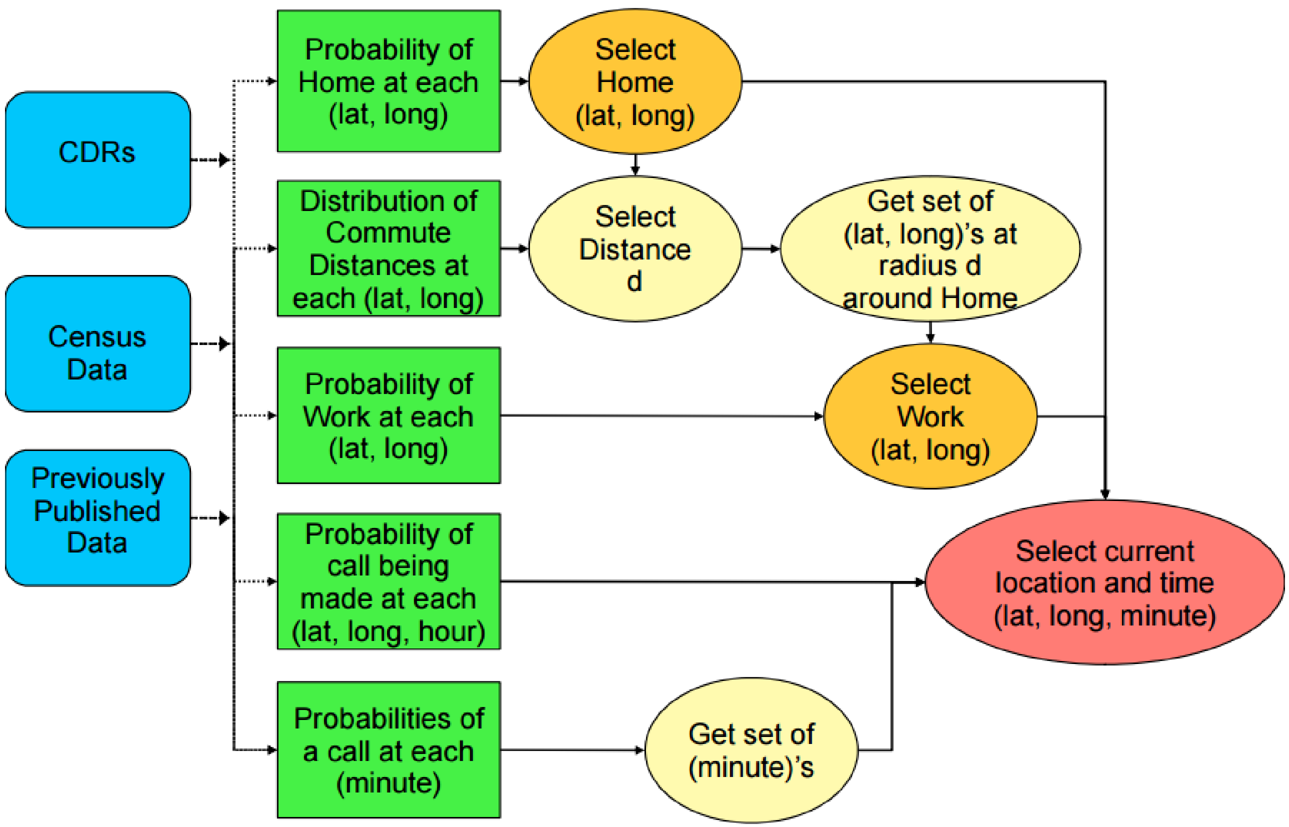

5. Data-Driven Modeling: Urban Phenomena

5	Feb 16	Data-Driven Modeling: Urban Phenomena Assigned Reading: • Real-Time Trip Information Service for a Large Taxi Fleet • City-Scale Traffic Estimation from a Roving Sensor Network • A Cost-Effective Recommender System for Taxi Drivers • Inferring Gas Consumption and Pollution Emissions of Vehicles throughout a City

City-Scale Traffic Estimation from a Roving Sensor Network

Javed Aslam
College of Computer and Information Science
Northeastern University
jaa@ccs.neu.edu

Sejoon Lim
CSAIL
Massachusetts Institute of Technology
sjlim@csail.mit.edu

Xinghao Pan
DSO National Laboratories
pxinghao@dso.org.sg

Daniela Rus
CSAIL
Massachusetts Institute of Technology
rus@csail.mit.edu

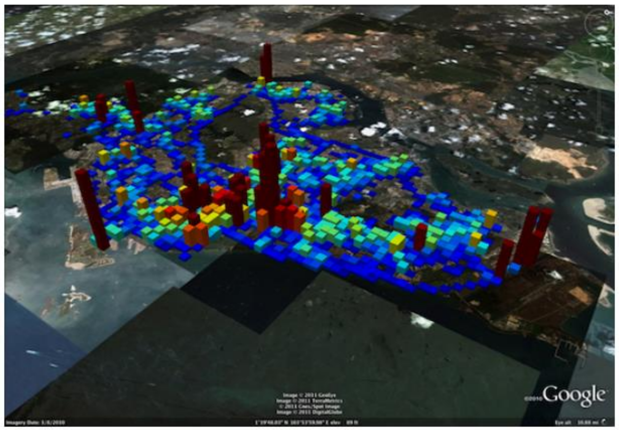

(a) 3∼5 am

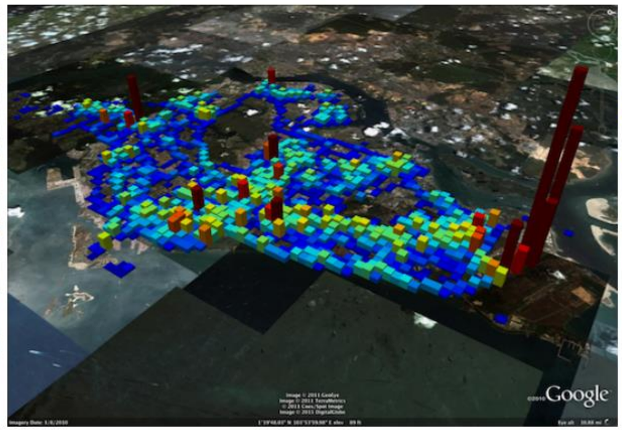

(b) 5∼7 am

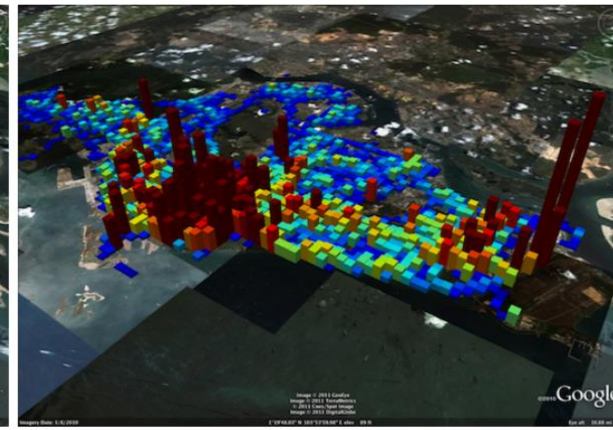

(c) 8∼10 am

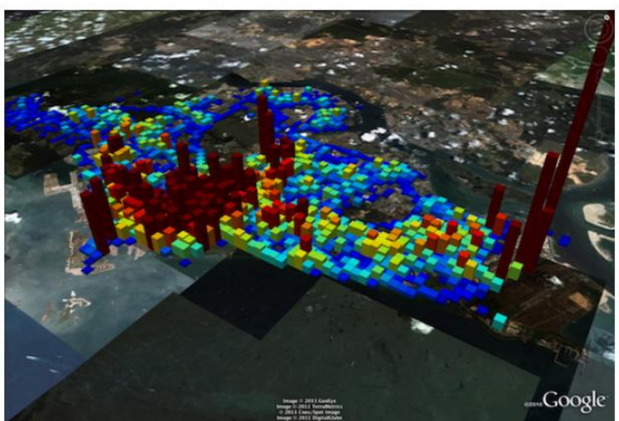

(d) 1∼3 pm

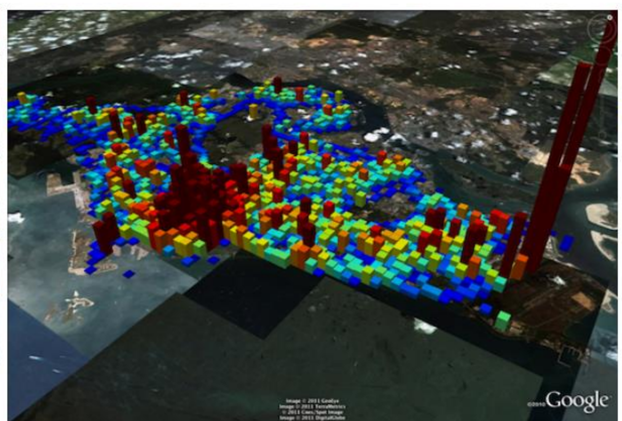

(e) 5∼7 pm

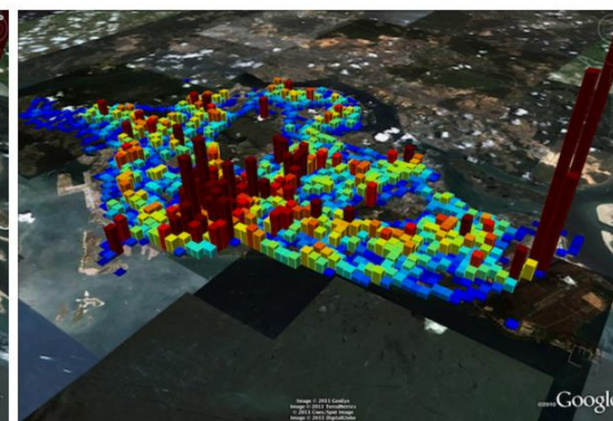

(f) 9∼11 pm

6. Data-Driven Modeling: Data Fusion

6	Feb 23	Data-Driven Modeling: Data Fusion Assigned Reading: • Diagnosing New York City's Noises with Ubiquitous Data • U-Air: When Urban Air Quality Inference Meets Big Data • Exploiting Geographic Dependencies for Real Estate Appraisal • Transfer Knowledge between Cities Suggested Reading: • Methodologies for Cross-Domain Data Fusion: An Overviewa

Diagnosing New York City's Noises with Ubiquitous Data

Yu Zheng[1], Tong Liu[1,2], Yilun Wang[1], Yanmin Zhu[2], Yanchi Liu[3], Eric Chang[1]
[1]Microsoft Research, Beijing, China
[2]Shanghai Jiao Tong University, Shanghai, China
[3]Information Systems Department, New Jersey Institute of Technology, Newark, NJ, United States
{yuzheng, v-tongli, v-yilwan, echang}@microsoft.com; yzhu@cs.sjtu.edu.cn; yanchilyc@gmail.com

7.Data Visualization

7	Mar 2	Data Visualization Assigned Reading: • Visual Exploration of Big Spatio-Temporal Urban Data: A Study of New York City Taxi Trips • TelCoVis: Visual Exploration of Co-occurrence in Urban Human Mobility Based on Telco Data • Origin-Destination Flow Data Smoothing and Mapping • Stacking-Based Visualization of Trajectory Attribute Data Suggested Reading: • Visual Analytics in Urban Computing: An Overview

Visual Exploration of Big Spatio-Temporal Urban Data: A Study of New York City Taxi Trips

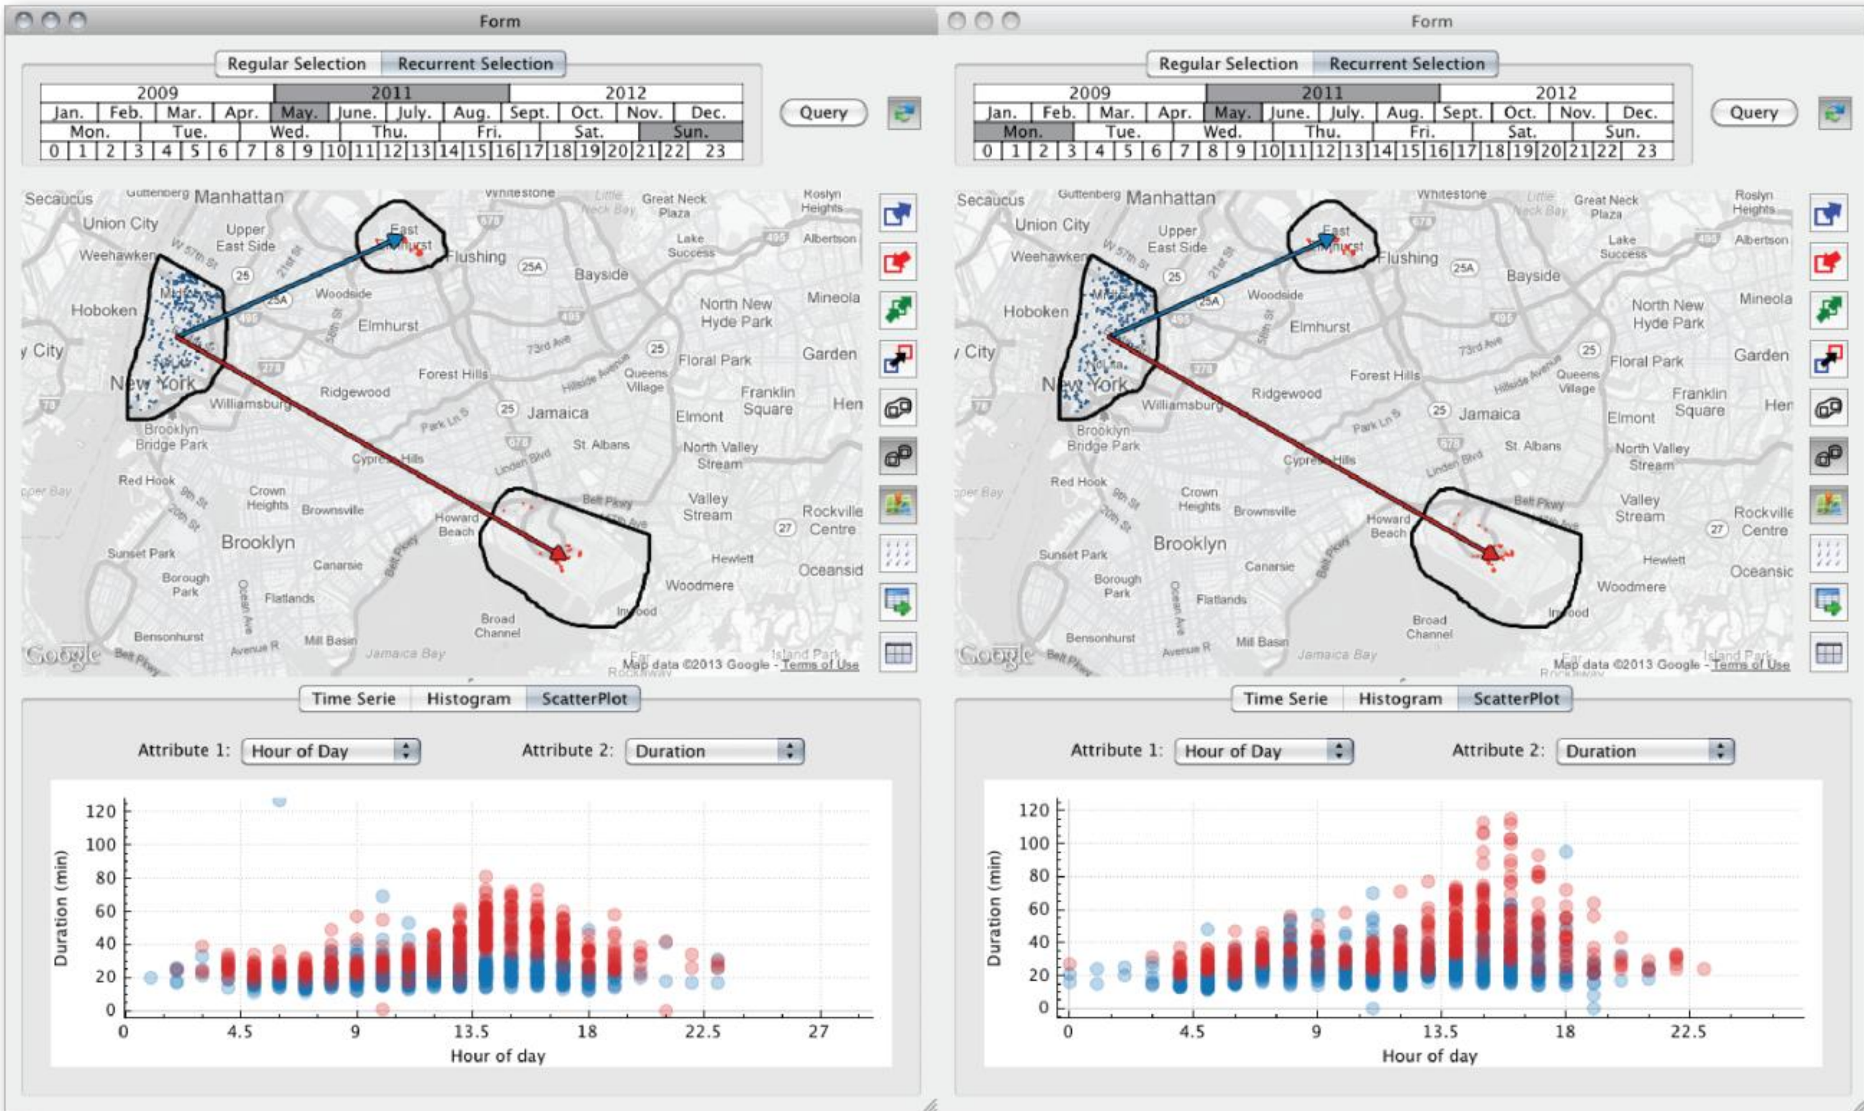

8.Data Predictive Control

8	Mar 9	Data Predictive Control Reading: • Exploiting Heterogeneous Human Mobility Patterns for Intelligent Bus Routing • Rebalancing Bike Sharing Systems: A Multi-source Data Smart Optimization • T-Finder: A Recommender System for Finding Passengers and Vacant Taxis • Taxi Dispatch with Real-Time Sensing Data in Metropolitan Areas

Rebalancing Bike Sharing Systems: A Multi-source Data Smart Optimization

Junming Liu[1], Leilei Sun[2], Weiwei Chen[3], Hui Xiong[1]*
[1]Management Science and Information Systems, Rutgers University, USA, {jl1433, hxiong}@rutgers.edu
[2]Institute of Systems Engineering, Dalian University of Technology, China, leisun@mail.dlut.edu.cn
[3]Supply Chain Management, Rutgers University, USA, wchen@business.rutgers.edu

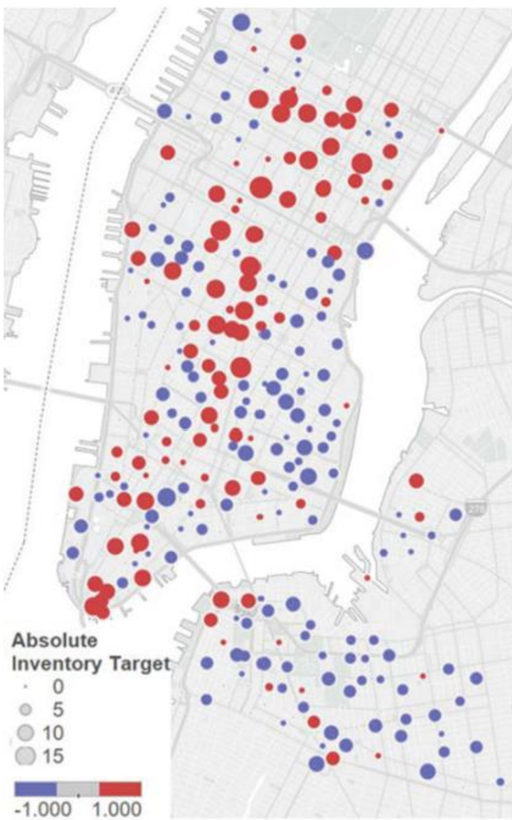

(a) Station Target Distribution

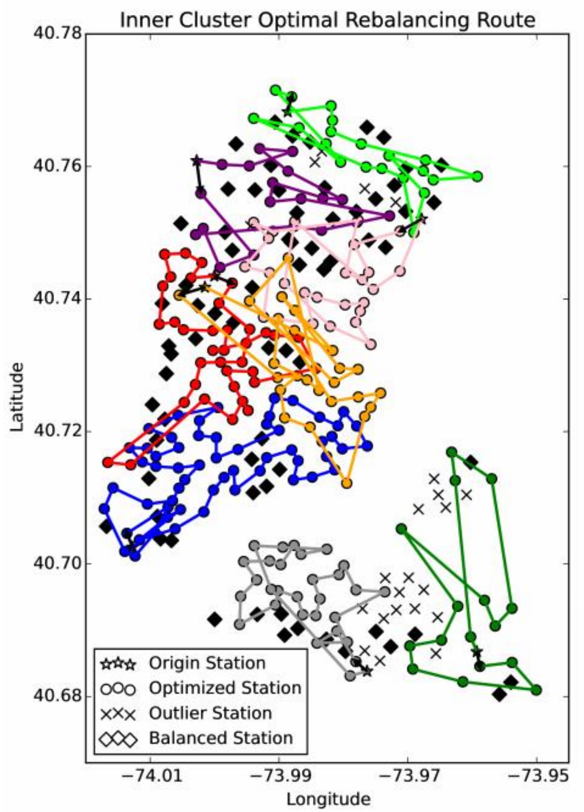

(b) Route with outliers ($VN = 8$)

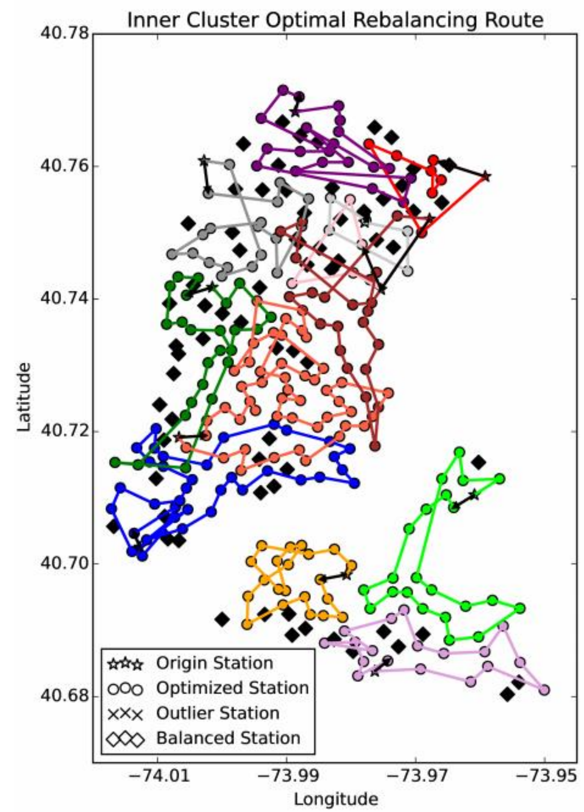

(c) Route without outliers ($VN = 12$)

9. Project Proposal Presentation

- After Spring Break
- No New Topics

10. Novel Services

10	Mar 30	Novel Services Assigned Reading: • Catch Me If You Can: Detecting Pickpocket Suspects from Large-Scale Transit Records • A Taxi Driving Fraud Detection System • CrowdAtlas: Self-Updating Maps for Cloud and Personal Use • Growing the Charging Station Network for Electric Vehicles with Trajectory Data Analytics Suggested Reading: • Modeling and Probabilistic Reasoning of Population Evacuation During Large-scale Disaster

Catch Me If You Can: Detecting Pickpocket Suspects from Large-Scale Transit Records

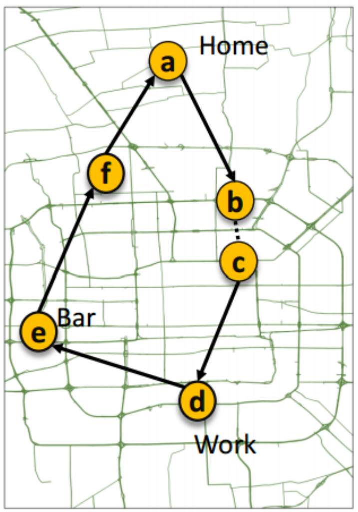

(a) Example Activities

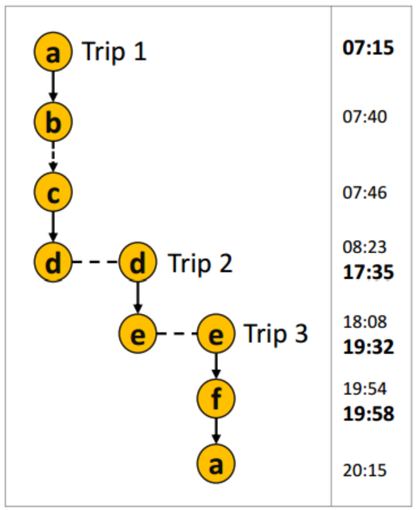

(b) Trips

Smart Card ID	Route Number	Boarding Station	Boarding Time	Exiting Station	Exiting Time
4322	Route 52	a	07:15	b	07:40
4322	Route 26	c	07:46	d	08:23
4322	Route 11	d	17:35	e	18:08
4322	Route 11	e	19:32	f	19:54
4322	Route 16	f	19:58	a	20:15

(c) Transit Records

(a) all passengers

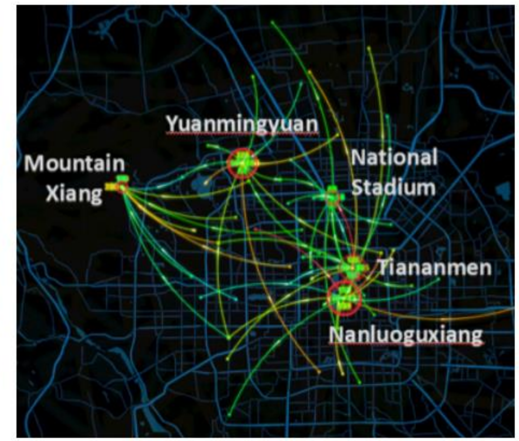

(b) visitors

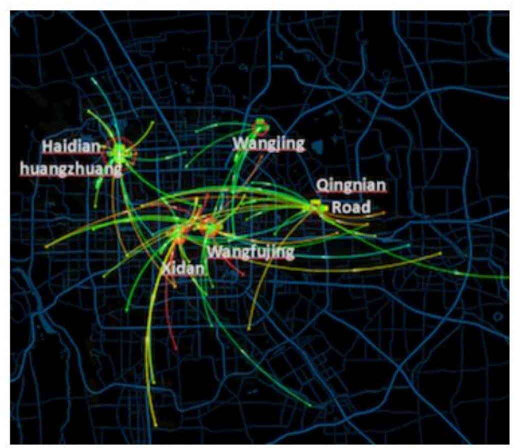

(c) shoppers

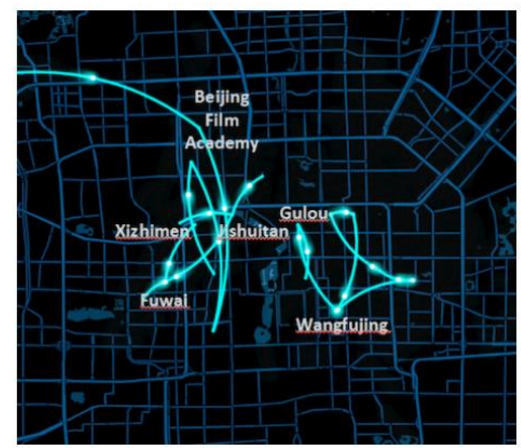

(d) thieves

11. Conflict and Dependency Analyses

11	Apr 6	Conflict and Dependency Analyses Assigned Reading: • Detection of Runtime Conflicts among Services in Smart Cities • DepSys: Dependency Aware Integration of Cyber-Physical Systems for Smart Homes • Scalable Social Sensing of Interdependent Phenomena • Catastrophic cascade of failures in interdependent networks Suggested Reading: • The fragility of interdependency • Identifying, understanding, and analyzing critical infrastructure interdependencies

Detection of Runtime Conflicts among Services in Smart Cities

M. N
J. Stankovic *
nd Information Technology
iversity,
Sweden
Email: {meiyi, pre
neberg@eit.lth.se

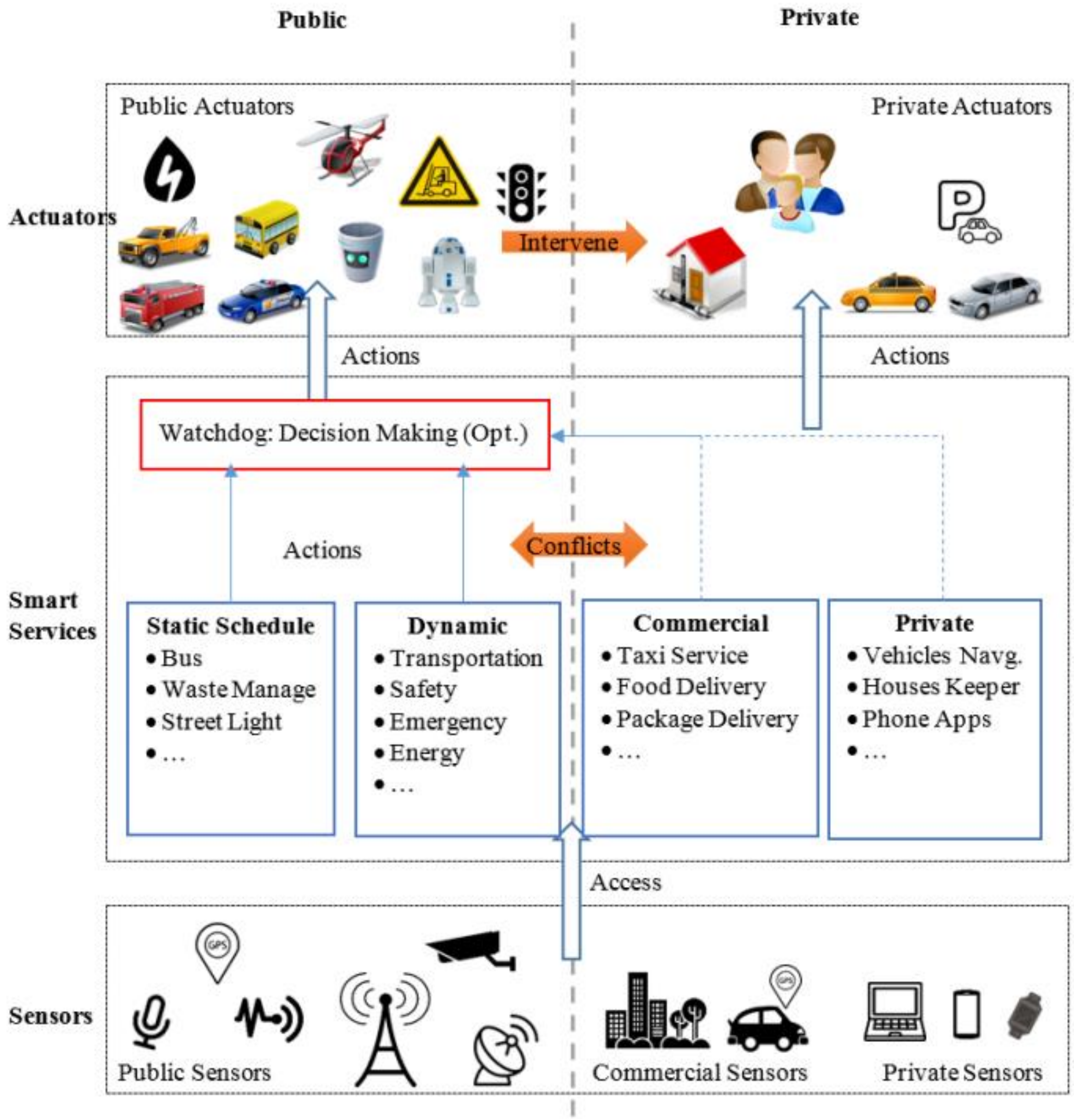

12. Human-in-the-loop

12	Apr 13	Human-in-the-loop Assigned Reading: • Taxi Driving Behavior Analysis in Latent Vehicle-to-Vehicle Networks: A Social Influence Perspective • Using Humans as Sensors: An Estimation-theoretic Perspective • Human Mobility, Social Ties, and Link Prediction • Friendship and Mobility: User Movement In Location-Based Social Networks

Taxi Driving Behavior Analysis in Latent Vehicle-to-Vehicle Networks: A Social Influence Perspective

Tong Xu[1], Hengshu Zhu[2], Xiangyu Zhao[1], Qi Liu[1]
Hao Zhong[3], Enhong Chen[1], Hui Xiong[3]

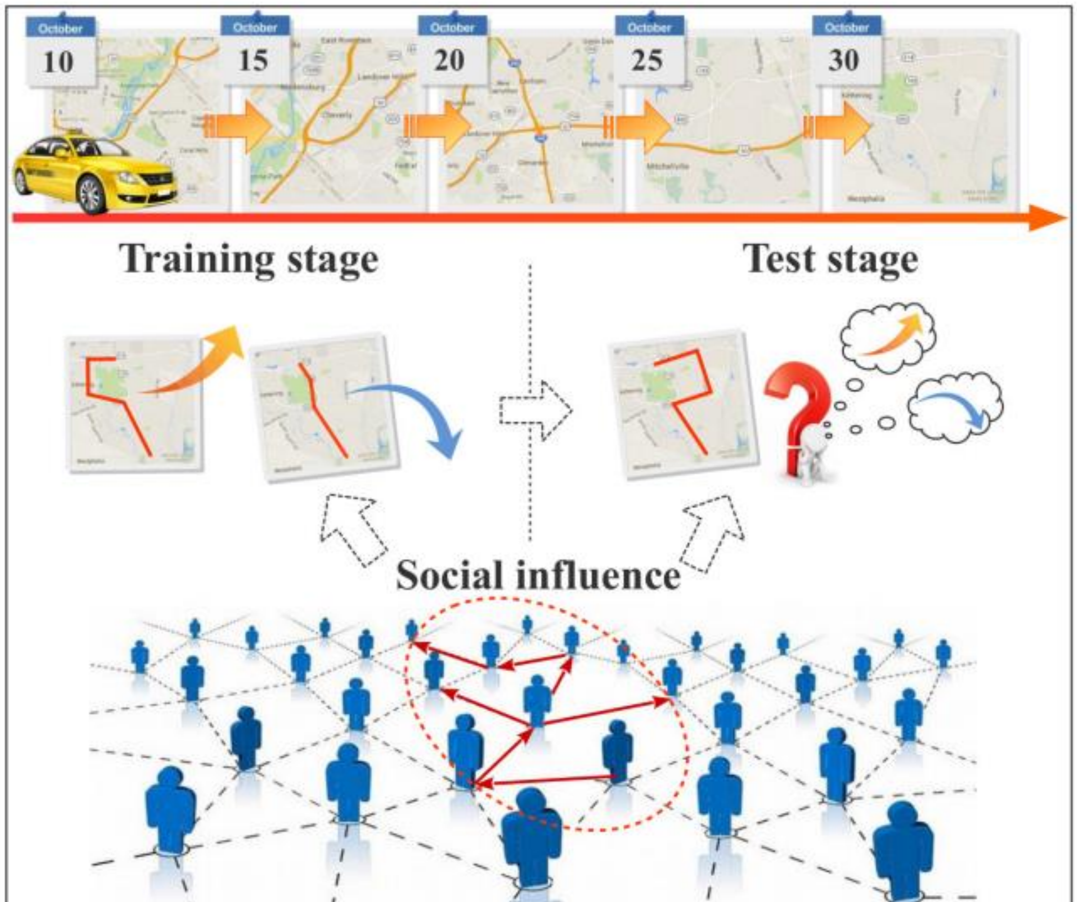

13. Privacy and Security

13	Apr 20	Privacy and Security Assigned Reading: • DP-WHERE: Differentially Private Modeling of Human Mobility • Elastic Pathing: Your Speed is Enough to Track You • VPriv: Protecting Privacy in Location-Based Vehicular Services • ZUbers against ZLyfts Apocalypse: An Analysis Framework for DoS Attacks on Mobility Systems Suggested Reading: • Differentially Private Transit Data Publication: A Case Study on the Montreal Transportation System • Anonymization of Location Data Does Not Work: A Large-Scale Measurement Study

Elastic Pathing: Your Speed is Enough to Track You

Xianyi Gao, Bernhard Firner, Shridatt Sugrim,
Victor Kaiser-Pendergrast, Yulong Yang, Janne Lindqvist
Rutgers University

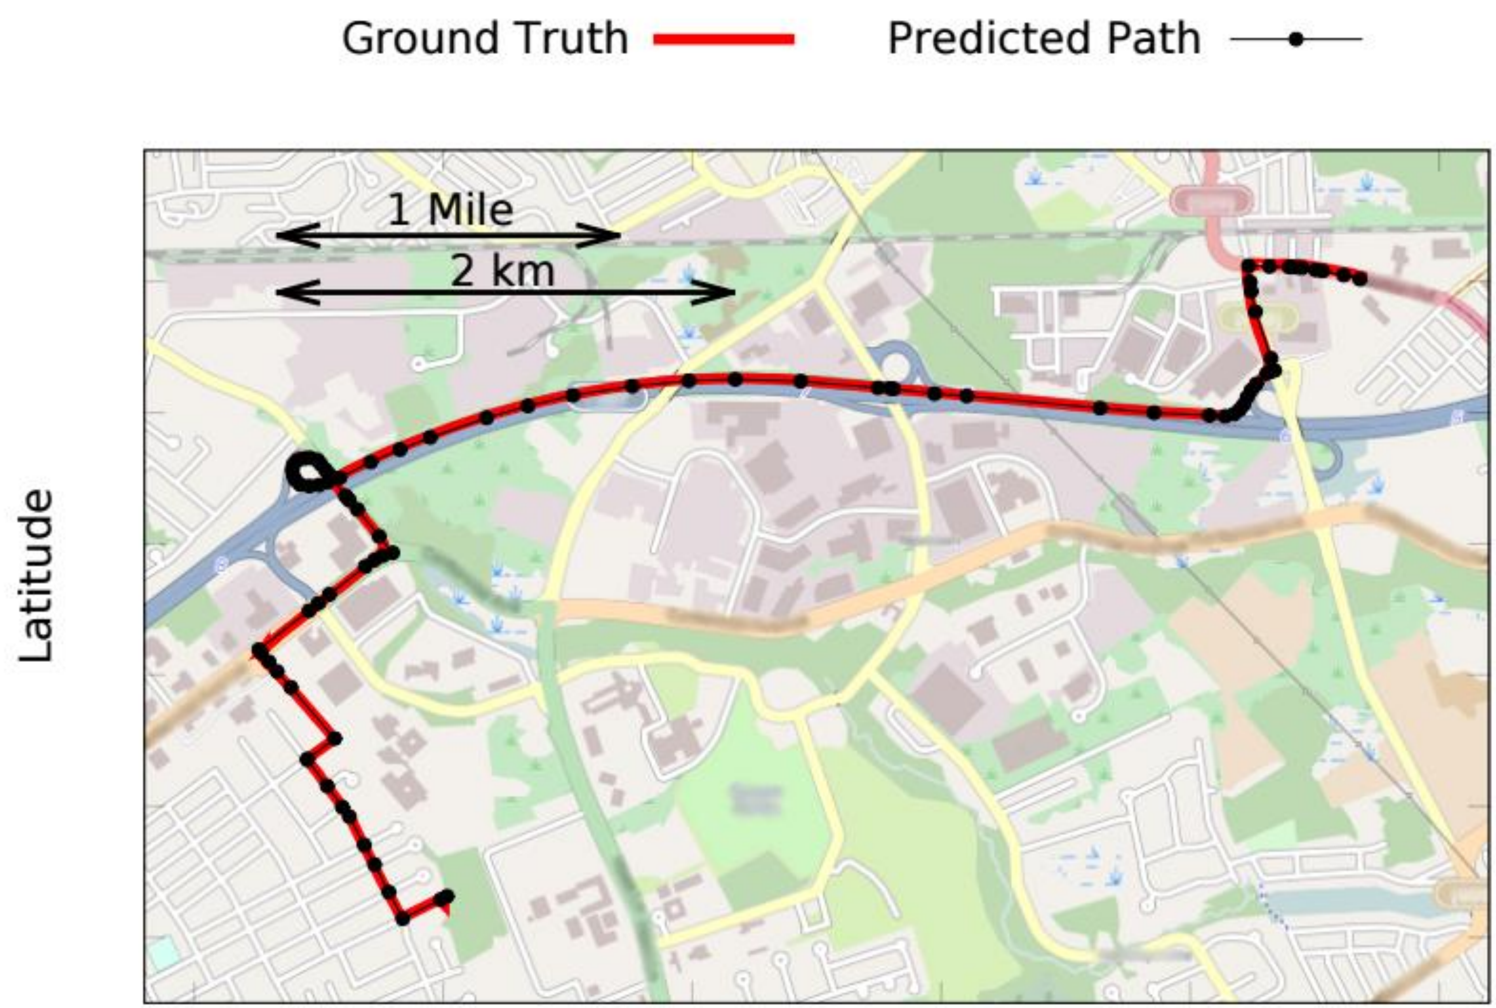

14. Final Presentation

- No New Topics

Questions?

Details about Grades

- **individual-based component 50% in total:**
  - **Class Participation 10%:** Please show up at all presentations and actively participate in the discussion after the presentation.
  - **Reading Summary 20%:** For 11 regular lectures (2^{nd}-8^{th} week and 10^{th}-13^{th} week), please read the introduction sections of all 4 papers in this lecture, and then select 2 papers from them to read all sections and write 2 summaries for each paper in the format given in the lecture slide. You can skip two summaries without hurting your grade, i.e., submit 20 summaries in total. Summaries are due in the beginning of class.
  - **Topic Presentation 20%:** Present a paper related to a topic in smart cities from the paper list. A topic presentation consists of 35 min talk plus 5 min Q&A.
- **Team-based component 50% in total:**
  - **Proposal Report 10%:** 4-page report describing the background, the problem, the existing solutions, and the proposed solution.
  - **Final Report 20%:** 8-page report: including 4 proposal report plus your implementation, evaluation, and conclusion.
  - **Proposal or Final Presentation 20%:** A student will present either in the proposal presentation or final presentation from 15 to 20 mins based on overall team size.

Details about Topic Selection

- Every student submits 5 of their preferred topics in the descending order among the following 10 topics.
- We will assign papers to you based on the topics you picked. Please send me by Jan 25.

Data Management and Processing,	**Data Predictive Control**
Data-Driven Modeling: Human Mobility	**Novel Services**
Data-Driven Modeling: Urban Phenomena	**Dependency Analyses**
Data-Driven Modeling: Data Fusion	**Human-in-the-loop**
Data Visualization	**Privacy and Security**

Team Formation:

- You are encouraged to find your own team members and let us know **by Jan 30**.
- If you cannot find anyone from the class, we will form teams **by Feb 2** based on your interests in topic selection sent to us previously.
- It is preferred that we have a mix-and-match for PhD students, master students, and undergrads in every team. In general, no more than one Ph.D is in the same team.