Mining Urban Data Workshop 3

KDD2018

August 20, 2018 – London, United Kingdom

kddlab.di.uoa.gr/mud3

www.vavel-project.eu

Opening	13:00 - 13:05
13:00 - 13:05	**Introduction by the Organizers**

Gennady Andrienko, Dimitrios Gunopulos, Vana Kalogeraki, Ioannis Katakis, Thomas Liebig, Jakub Marecek, Katharina Morik

Session A	13:05 - 14:30
13:05 - 13:45	(INVITED TALK) Andrew Tomkins (Google AI)
	Using Crowdsourced Data for Parking and Itinerary Generation

13:45 – 14:30	**Paper Presentations**

[1] Sudatta Mohanty and Alexey Pozdnukhov, **Graph CNN + LSTM Framework For Dynamic Macroscopic Traffic Congestion Prediction**

[2] Obaid Malik, Sarvapali Ramchurn and Alex Rogers, **A Multi-output Gaussian Process Model for Leak Detection and Quantification in Water Distribution Systems**

[3] Anthony Sicilia, Alexandros Labrinidis and Konstantinos Pelechrinis, **A Holistic Evaluation of Transit Supply and Demand using Network Analysis: The TDI Framework**

[4] Phillip Taylor, Nathan Griffiths and Alex Mouzakitis, **Selection of Compressible Signals from Telemetry Data**

Coffee Break	14:30 - 15:00
Session B	15:00 - 17:00
15:00 – 15:15	Dimitrios Gunopulos **The VaVeL Project: Challenges and Lessons for Smart Cities**

15:15 - 15:40	(INVITED TALK) Jakub Marecek (IBM Research AI)
	Recommender Systems and their Effects

15:40 - 17:00	**Paper Presentations**

[5] Akira Kinoshita, Atsuhiro Takasu and Jun Adachi, **Weather-Sensitive Road Segment Detection in A Snowy City**

[6] Nikos Zacheilas, Vana Kalogeraki and Thanasis Priovolos, **A Comparative Study of General-Purpose Pub/Sub Systems for Message Dissemination**

[7] Seyedsaeed Hajiseyedjavadi, Yu-Ru Lin and Konstantinos Pelechrinis, **Discovering functionality of urban regions by learning low-dimensional representations of a spatial multiplex network**

[8] Piotr Wawrzyniak and Jarosław Legierski, **On the Use of Apache SAMOA to Predict Delay and Position of Public Transport Vehicles**

[9] Archit Parnami, Prajval Bavi, Dimitris Papanikolaou, Srinivas Akella, Minwoo Lee and Siddharth Krishnan, **Deep Learning Based Urban Analytics Platform: Applications to Traffic Flow Modeling and Prediction**

[10] Peter Casey, Kevin Wilson and David Yokum, **A Cautionary Tail: A Framework and Case Study for Testing Predictive Model Validity**

[11] Jonathan Epperlein, Jaroslaw Legierski, Marcin Luckner, Jakub Marecek and Rahul Nair**, The Use of Presence Data in Modelling Demand for Transportation**

[12] Henry Dinhofer, Prateek Sappadla and Lakshminarayanan Subramanian, **Maximizing Coverage of EMS services in New York City using a Double Standard Model**

Speakers instructions: Each presentation will take up to 11 minutes (9 minutes presentation, 2 minutes questions)

Invited Speaker 1	
Andrew Tomkins (Google AI)	**Title**: *Using Crowdsourced Data for Parking and Itinerary Generation*

Bio: Andrew joined Google Research in 2009, where he serves as an engineering director, currently working on machine learning and understanding of geo data. Prior to these projects, he worked on measurement, modelling, and analysis of content, communities, and users on the World Wide Web. Before joining Google, he spent four years at Yahoo! serving as chief scientist of search, and eight years at IBM's Almaden Research Center, where he served as chief scientist on the WebFountain project. Andrew has authored over 100 technical papers and 90 issued patents. He received Bachelors degrees in Mathematics and Computer Science from MIT, and a PhD in CS from Carnegie Mellon University.

Abstract: Large-scale crowdsourced data is widely used in many domains, both online and offline. For example, systems for road navigation often use crowdsourced information about driver speed in order to estimate traffic conditions. In this talk, we'll discuss two other uses of crowdsourcing in a similar vein: estimating properties of the parking at a particular location, and generating high-quality travel itineraries within a city. The talk will describe the data and algorithms used for both of these problems, and will give some examples of actual products that use the technology.

Invited Speaker 2	
Jakub Marecek (IBM Research AI)	**Title:** *Recommender Systems and their Effects*

Bio: Together with some fabulous colleagues, Jakub Marecek develops solvers for optimisation and control problems at IBM Research -- Ireland. Jakub joined IBM Research from the School of Mathematics at the University of Edinburgh in August 2012. Prior to his post-doc in Edinburgh, Jakub had presented an approach to general-purpose integer programming in his dissertation at the University of Nottingham and worked in two start-up companies.

Abstract: Recommender systems are widely used in settings, where actions impact the recommendations, which in turn have impact on the actions. For an example of such a closed-loop setting, consider navigation systems, which use information about travel times to recommend a route. If the particular navigation system is used widely enough, the recommendation may impact the future traffic state, possibly rendering the recommendation suboptimal from both the point of view of the driver and the society as a whole. Similar effects can be illustrated on the recommendations of restaurants. If a small bistro without a table reservation system becomes top ranked, many customers may arrive at its door and get turned down, leading to poor reviews. Further, there can be issues related to priming, for example when the reviews suggest the place is not touristy. Several problems arise, including the recovery of unbiased user models in the presence of recommenders and developing recommenders that allow for some guarantees on the closed-loop behaviour of the system.