ESRC Strategic Network:
Data and Cities as Complex
Adaptive Systems
(DACAS)

CASE STUDY REPORT

CASE STUDY 06.A
THE CO-EVOLUTION OF TRADITIONAL (TAXI) AND ICT ENABLED (UBER) TRANSPORT SYSTEMS (NEW YORK DATASET) + AUTONOMOUS MOBILITY

CUAUHTEMOC ANDA

JOINT DACAS
ICTP-SAIFR
WORKSHOP

20–24 JUNE 2016

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL

DACAS

ICTP SAIFR

ESRC Strategic Network: Data and Cities as Complex Adaptive Systems (DACAS)

Workshop on Modelling Urban Systems

The co-evolution of autonomous mobility-on-demand (AMoD) systems and human-driven vehicles (Case Study Report)

Date: 20-24 June 2016
Place: Sao Paulo, Brazil

Leader
Cuauhtemoc Anda
Singapore-ETH Future Cities Lab

Support:
Ulysses Sengupta, Manchester Metropolitan University
Ana Bazzan, UFRGS
Nir Oren, University of Aberdeen

Introduction

Mobility-on-demand (MoD) services under the flag of Uber's disruptive business model have made a breakthrough in the way we travel in big cities. New door-to-door mobility options have been added to the available options from which a commuter can choose. These new transportation systems leverage on information and communication technologies, enabling service requests from a smartphone and online payments.

However, urban mobility is just entering a *tipping point*[1] with four mayor innovations in the car manufacturing industry coming through: in-vehicle connectivity, electric vehicles, car sharing, and autonomous driving. Big firms such as Google and Tesla have started testing their autonomous vehicles (AVs) with plans from Google to make them available to the public in 2020[2]. In addition, smaller competitors like *nuTonomy* a spinup from MIT has recently partnered up with Singapore's Land Transport Authority (LTA) to deliver a fleet of autonomous taxis by 2018[3].

It is then in the research community outlook, not only to finalise the development of the self-driving technology, but also address the implications in a broader scale, the city scale. What sort of interactions from self-driving and manned vehicles can we expect from a higher perspective? Understanding the general panorama of both systems competing in the same urban space to service the cities' travelling needs will make us realise the real contributions of AV's

[1] http://www.mckinsey.com/business-functions/sustainability-and-resource-productivity/our-insights/urban-mobility-at-a-tipping-point

[2] http://www.ibtimes.com/google-inc-says-self-driving-car-will-be-ready-2020-1784150
[3] http://www.digitaltrends.com/cars/nutonomy-driverless-taxi-singapore/

in a multi-modal transport setting, and if we can consider AVs as part of the sustainable transport ecosystem.

In this perspective we analysed autonomous mobility-on-demand (AMoD) systems/services constituted by a fleet of AVs. AMoD has emerged as a promising solution for urban transportation. Compared to prevailing systems, AMoD promises sustainable, affordable personal mobility through the use of self-driving shared vehicles. For instance, AMoD systems are a future prospect to solve the last-mile problem, balance the overall efficiency of the transportation network, and due to its autonomy, overcome the fundamental load balancing problem of conventional mobility-on-demand services.

Hence, the current case study analyses the interactions and evolution of three players: the commuters, the traditional manned-vehicle services (e.g. taxi), and the inclusion of an AMoD service. All parties are represented as agents with their own utility/fitness functions they seek to optimise. The simulation will output the emergent behaviour derived from the different dynamics and interactions of the three players and their inherent evolution towards their optimisation in time, cost, and space. This will ultimately help us understand from the city level perspective the interactions between self-driving and man-driven vehicles and to assess under what conditions AMoD systems are a sustainable option.

Problem Definition

Parting from the travel demand data generated in the New York City Taxi and Uber datasets, we would like to understand the indirect interactions of AMoD systems and conventional transportation services (taxi) on servicing the NYC's travel demand. Find the optimal parameters and operation policies for an AMoD fleet and analyse the evolution of users' choice in a mixed-traffic setting. From the fleet management perspective, find out the optimal number of dispatch units needed to cover the demand of NYC, and find the optimal coordinated dispatching, route choice, and rebalancing policies for the AMoD fleet. Furthermore, we would like to explore under what market penetration rates, are AMoD systems improving the overall transportation performance, and if the inclusion of AVs can be considered part of sustainable solution.

Hence, the aims of the proposed case scenario are two-fold:

1) Explore suitable fleet management strategies for AMOD systems. Find optimal control strategies for request assignment, route choice and the rebalancing problem. And also on the economic dimension: fleet sizing and financial analyses.
2) On a higher level (city level) assess the interactions with manned vehicles and how the competition between both services affects the user's choice and the transportation system overall performance.

Concept Transferability

1. Multi-agent based simulations

Agent-based models (ABM) from the Complex Systems domain are computational models to simulate the actions and interactions of autonomous agents with a view of assessing their effects on the system as a whole. In the concept of multi-agent simulations, the goal is to search for explanatory insight into the collective behaviour of agents obeying simple rules. They are a type of micro scale model that simulates the simultaneous operations and interactions of multiple agents in an attempt to re-create and predict the appearance of complex phenomena. In this sense, a key notion is that simple behavioural rules generate complex behaviour.

For large-scale transportation systems, activity based models can be implemented through multi –agent base simulations. Agents represent people willing or in need to travel to perform activities. Thus, each agent has a set of daily activities. They are able to make their own decisions such as route choice, mode choice and time of departure. In the case study, user agents, will be represented by the demand generated from the NYC Taxi and Uber datasets. The study will explore the evolution of the agents' decision towards which mode of transport maximises their own utility (mainly driven by waiting time and cost) between AMoD services and conventional personal mobility services (e.g. Taxis, Uber).

For the supply side, we need another type of agents to represent the AMoD fleet and the fleet of Taxis. On a mesoscopic level the main difference in the salient behaviour between the two is not the model but the control strategies (e.g. closed-loop, coordinated, intelligent) for the assignment, routing and rebalancing tasks. Hence, the dynamics and control of an AMoD system can be divided in two tasks: firstly, when an AMoD unit is active, the route can be selected to balance the overall system performance by routing vehicles through less-congested roads; and secondly, when the unit is not in use (i.e. idle state), the AMoD system can redistribute AVs to better meet demand through an automated balancing strategy. These control strategies should be built along the coordination capabilities of the AV fleet to respond to a diversity of stochastic events in an uncertain and complex setting. Interactions between the AMoD, the conventional taxi service and the NYC travel demand will result in emergent phenomena such as traffic jams.

For this task, the case study can be implemented in one of the two most widely used available open frameworks to develop multi-agent transportation simulations on a large scale: MATSim[4] (java-based) or SimMobility[5] (c++-based).

2. Reinforcement Learning

Once the multi-agent scenario has been set, the second challenge to be solved is the fleet management strategy implementation for the AMoD system. By leveraging on the self-driving technology and coordinated algorithms, the control strategies for the AMoD system make the significant difference in the mesoscopic simulation in comparison to the fleet of conventional taxis. Hence, to find optimal policies for the AMoD fleet management tasks, routing and

[4] http://www.matsim.org/
[5] https://its.mit.edu/research/simmobility

rebalancing, a promise approach is the Reinforcement Learning (RL) paradigm in the context of optimal control, and specifically, multi-agent reinforcement learning.

Reinforcement Learning is very closely related to the theory of classical optimal control, as well as dynamic programming, stochastic programming, simulation-optimization, stochastic search, and optimal stopping (Powell, 2012). Both RL and optimal control address the problem of finding an optimal policy (often also called the controller or control policy) that optimises an objective function (i.e., the accumulated cost or reward), and both rely on the notion of a system being described by an underlying set of states, controls and a plant or model that describes transitions between states. However, optimal control assumes perfect knowledge of the system's description in the form of a model (i.e., a function T that describes what the next state of the robot will be given the current state and action). For such models, optimal control ensures strong guarantees which, nevertheless, often break down due to model and computational approximations. In contrast, reinforcement learning operates directly on measured data and rewards from interaction with the environment. Reinforcement learning research has placed great focus on addressing cases which are analytically intractable using approximations and data-driven techniques. The goal of reinforcement learning is to discover an optimal policy $\pi*$ that maps states (or observations) to actions so as to maximize the expected return J, which corresponds to the cumulative expected reward. (Kober, Bagnell, & Peters, 2013)

Why Reinforcement Learning?
For the case study, Reinforcement learning is an appropriate method since we want to find optimal policies for routing and the rebalancing problem under complex scenarios, for which we don't possess a model of the world. Thus, from the events experienced by the fleet of AVs, we would like the AMoD systems to learn what are the best actions under different stochastic conditions as to optimised travel times, energy consumption and maximised profit.

Data Availability

In order to derive the travel demand for New York City, two open datasets are available to set up the initial conditions for our multi-agent simulation framework.

- NYC Taxi Data

The official TLC trip record dataset contains data for over 1.1 billion taxi trips from January 2009 through June 2015, covering both yellow and green taxis. Each individual trip record contains precise location coordinates for where the trip started and ended, timestamps for when the trip started and ended, plus a few other variables including fare amount, payment method, and distance travelled.

Link: https://github.com/toddwschneider/nyc-taxi-data

- Uber Data

Data covering nearly 19 million Uber rides in NYC from April–September 2014 and January–June 2015. In particular Uber provides time and location for pickups only, not drop offs.

Link: https://github.com/fivethirtyeight/uber-tlc-foil-response

Literature Review

Modelling, simulation and control of AMoD systems

Autonomous on-demand (AMoD) systems, constitute a transformative and rapidly developing mode of transportation wherein robotic, self-driving vehicles transport passengers in a given environment. Specifically, Pavone (2014) addresses AMoD systems along three dimensions: (1) modeling, that is analytical models capturing salient dynamic and stochastic features of customer demand, (2) control, that is coordination algorithms for the vehicles aimed at throughput maximization, and (3) economic, that is fleet sizing and financial analyses.

Pavone (2014), Zhan and Pavone (2014), Zhang, Spieser, Frazzoli, and Pavone (2015) propose a dedicated model of an AMoD system, in which a spatial queue model is used to manage the stochastic travel requests generated in a defined map. And thus the problem to find an optimal policy for the AV's to serve the requests is turn into an operations research problem: joint task allocation and scheduling problem. They proposed a closed-loop control policy based on the Dynamic Traveling Repairman problem and on the Dynamic Traffic Assignment problem. Firstly, they transform the AMoD system into a closed Jackson network with respect to the vehicles, where the equilibrium distribution of a queueing network is possible to compute as the product has a product-form solution. Secondly, to rebalance the vehicles to ensure even vehicle availability, the strategy was to add virtual customer streams. The model and control strategy was tested for the cities of New York and Singapore.

A more recent approach for the AV fleet control is the one by Zhang, Rossi, and Pavone (2016). They proposed a model predictive control (MPC) from the control theory domain for a station-based management of an AMoD system. They firstly built a discrete-time model of the dynamics of the AV fleet, and secondly designed a model predictive control algorithm for the optimal coordination of the AMoD system. At each optimisation step, the vehicle scheduling and routing problem is solved as a mixed integer linear program where the decision variables are binary variables representing whether a vehicle will wait at a station, service a customer, or rebalance to another station. One of the major contributions was the inclusion of charging constraints associated with using electric vehicles.

However, while Pavone (2014), Zhan and Pavone (2014), Zhang, Spieser, Frazzoli, and Pavone (2015) work was tested in large-scale, the control approach cannot cope with realistic phenomenon such as congestion and interaction with human-driven vehicles. They conclude that future research should come up with efficient control algorithms for increasingly more realistic models. On the other hand, the model predictive control from Zhang, Rossi, and Pavone (2016) although it takes into account more realistic constraints (e.g. charging stations), it still needs to be scale up for city-wide systems, since the computational complexity of the mixed integer linear program scales exponentially with the number of stations and vehicles. Moreover, they conclude that the inclusion of the congestion aspect and optimal coordination algorithms in an intermodal system are a matter of future research.

With the proposal of a multi-agent simulation framework, the case study can address the limitations described with a city-scale multimodal platform, in which the interactions from human-driven vehicles and AVs, as well as congestions are capable to emerge. As for the AMoD's control strategy the inclusion of a Reinforcement Learning (RL) approach would account for the uncertainties in the real world, for both the routing and rebalancing policies.

RL paradigm leverages the robotics domain and capabilities of the AV fleet, where it has been proved successful for the solution of optimal problems under uncertain environments.

Agent-based simulation for AMoD

For large-scale transportation systems, the inclusion of a fleet of taxis and autonomous taxis as another class of agents has been already implemented in both MATSim and SimMobility. Hörl, Erath, and Axhausen (2016) have included the modelling infrastructure for the simulation of AVs in MATSim. As for SimMobility, Marczuk et al. (2015) included a module for AVs in the short-term simulator, which simulates the individual decisions and the transportation network at the sub-second level. In both cases they have initially tested the simple cases for the fleet management tasks: FIFO service and nearest assignation, routing according to the network free speeds, return to the original station or wait at drop-off rebalance strategies. However, the bases are settled for the proposal and inclusion of more efficient control algorithms in the large-scale, multimodal setting offered by both platforms.

Reinforcement Learning

Even though a formal control strategy for an AMoD system through reinforcement learning has not been proposed yet, several studies can be used as starting point for the case study development.

As an anchor point, the work by Bazzan (2008) on opportunities for multiagent systems and multiagent reinforcement learning in traffic control, highlights how can open problems in traffic enginnering can be approach through coordination algorithms and reinforcement learning for multiagent systems. Later on, Tavares and Bazzan (2012) proposed a reinforcement learning approach for route choice in traffic scenarios, which relies solely on drivers' experience to guide their decisions. Experimental results demonstrate that reasonable travel times can be achieves and vehicles can be distributed themselves over the road network avoiding congestion. Along the routing choice line, Cox, Jennings, and Krukowski (2013) adapted the Q-Routing algorithm, originally developed for packet routing in communication networks and inspired by the Q-learning algorithm, for an adaptive routing strategy of AVs shortest path planning in congested networks.

Another focus of reinforcement learning in vehicle trajectories is the optimisation of a taxi route to maximised the revenue (Wang and Lampert, 2014). Although the optimisation process is set for one vehicle, if extended to a fleet of taxis, the overall fleet's efficiency could be greatly improved (less particles emitted), drivers would minimise their vacancy time and working hours, and passenger could see better Taxi coverage.

In the case study, the main motivation towards using RL is to leverage the capabilities of the AV fleet when encounter with a networked, heterogeneous, stochastic decision problem with uncertain information. In a coordinated and self-organised manner, RL is an important piece of the puzzle to learn optimal policies under uncertain events and complex scenarios. In addition, this will allow us to explore the possibility to balance the global transportation system utility using the AMoD system.

Research Questions

- Would AMoD systems decrease congestion? In general, do AMoD systems represent an economically viable, sustainable, and societally-acceptable solution to the future of personal urban mobility?

- Can the inclusion of an AMoD system be used to reach the transport system overall optimal performance?

- For which percentage of AVs inclusion can we start having increments in the overall efficiency of the network without sacrificing individual utility too much?

- How does the user's choice evolve when an AMoD system is included based on lower pick-up times, but for some penetration rates of AVs longer travel times?

- In the competition game between AMoD vs manned MoD services, what is the user's expectancy on improved waiting and travel times? Based on these factors, how do the users' choice evolve in terms of the preferred service selection?

- What is an optimal fleet size for the AMoD system under competition for a fixed demand? (see Boesch, Ciari, & Axhausen, 2016 for fleet size and quality of service)

- What is an optimal route choice algorithm for a fleet of AVs?

- What is an optimal rebalancing strategy? What are the trade-offs between different strategies?

Research Output

A platform to evaluate coordinated control strategies for AMoD systems with complex stochastic interactions in a large-scale multimodal setting.

Future research possibilities

1. Trip sharing and autonomous vehicles. How to include efficient algorithms and simulations of sharing trips.

2. Exploring pricing mechanism boundaries for the optimality of both service providers (cost) and users (cost and waiting time).

References

Bazzan, A. L. (2009). Opportunities for multiagent systems and multiagent reinforcement learning in traffic control. *Autonomous Agents and Multi-Agent Systems*, *18*(3), 342-375.

Boesch, P. M., Ciari, F., & Axhausen, K. W. (2016). Required autonomous vehicle fleet sizes to serve different levels of demand. In *Transportation Research Board 95th Annual Meeting* (No. 16-2682).

Godfrey, G. A., & Powell, W. B. (2002). An adaptive dynamic programming algorithm for dynamic fleet management, I: Single period travel times.*Transportation Science*, *36*(1), 21-39.

Hörl, S., Erath, A., Axhausen, K. (2016). Implementation of an autonomous taxi service in a multi-modal traffic simulation using MATSim. *In press.*

Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 0278364913495721.

Marczuk, K. A., Hong, H. S. S., Azevedo, C. M. L., Adnan, M., Pendleton, S. D., & Frazzoli, E. (2015, July). Autonomous mobility on demand in SimMobility: Case study of the central business district in Singapore. In *2015 IEEE 7th International Conference on Cybernetics and Intelligent Systems (CIS) and IEEE Conference on Robotics, Automation and Mechatronics (RAM)* (pp. 167-172). IEEE.

Michalski, R. S., Carbonell, J. G., & Mitchell, T. M. (Eds.). (2013). *Machine learning: An artificial intelligence approach*. Springer Science & Business Media.

Pavone, M. (2014). Autonomous mobility-on-demand systems for future urban mobility. In *Autonomous Driving* (pp. 387-404). Springer Berlin Heidelberg.

Powell, W. B. (2012). AI, OR and control theory: A rosetta stone for stochastic optimization. Technical report, Princeton University.

Tavares, A. R., & Bazzan, A. L. C. (2012). Reinforcement learning for route choice in an abstract traffic scenario. In *VI Workshop-Escola de Sistemas de Agentes, seus Ambientes e aplicaçoes (WESAAC)* (pp. 141-153).

Wang, J., & Lampert, B. (2014) Improving Taxi Revenue with Reinforcement Learning.

Zhang, R., & Pavone, M. (2014). Control of robotic mobility-on-demand systems: a queueing-theoretical perspective. *The International Journal of Robotics Research*, *35*(1-3), 186-203.

Zhang, R., Spieser, K., Frazzoli, E., & Pavone, M. (2015, July). Models, algorithms, and evaluation for autonomous mobility-on-demand systems. In*2015 American Control Conference (ACC)* (pp. 2573-2587). IEEE.

Zhang, R., Rossi, F., & Pavone, M. (2015). Model Predictive Control of Autonomous Mobility-on-Demand Systems. *arXiv preprint arXiv:1509.03985*.

Zolfpour-Arokhlo, M., Selamat, A., Hashim, S. Z. M., & Afkhami, H. (2014). Modeling of route planning system based on Q value-based dynamic programming with multi-agent reinforcement learning algorithms. *Engineering Applications of Artificial Intelligence*, *29*, 163-177.