TRB 2018 Notes

Matthew Battifarano
mbattifa@andrew.cmu.edu

January 2018

Contents

I had a fantastically fulfilling and enriching experience at the 2018 Transportation Research Board conference last week. This document is an attempt to collect my notes and thoughts on selected research from the sessions I was able to sit in on. There were over 5,000 presentations in nearly 800 lecterns, poster sessions, and workshops over the five-day conference. Please feel free to distribute and email me with questions, comments, or corrections at `mbattifa@andrew.cmu.edu`.

I want to thank David Abel[1] for the inspiration for this document. I used the structure of his notes from the NIPS 2017 conference[2] as a guide for my own.

Finally, I would like to thank my advisor, Sean Qian[3], the Traffic21 Institute[4], and my fellow students at the Mobility Data Analytics Center[5] at Carnegie Mellon.

1 Conference Summary and Highlights

As a student, my personal highlight was the one-on-one conversations I had with students and professors about the research they were presenting during the poster sessions. Each of these brief conversations gave me unique access to the research and, more importantly, to the thought and work that went into it.

My favorite sessions were,

[1] `http://cs.brown.edu/~dabel/`
[2] `http://cs.brown.edu/~dabel/blog/posts/misc/nips_2017.pdf`
[3] `https://www.cmu.edu/cee/people/faculty/qian.html`
[4] `https://traffic21.heinz.cmu.edu/`
[5] `http://mac.heinz.cmu.edu/research.html`

1. Network Equilibrium Modeling Subcommittee meeting. This meeting was a fantastic opportunity to hear about the future of network equilibrium research from both academia and industry. Lectern and poster sessions are devoted to research that has already been done, but committee meetings are all about ongoing and future work. Past work is published; future work is accessible *only* in forums like these. If you find yourself at TRB I highly recommend sitting in on a subcommittee meeting. See section 4.1.

2. Transportation Network Modeling poster session. Networks are fundamental to modeling transportation systems. I think some of the most interesting questions in transportation research are being addressed in this field. How is the existing network used? How will its travelers react if it is changed in a certain way? How will emerging technologies augment the network? See section 5.2.

3. Behavioral Route Choice poster session. Behavioral route choice is all about understanding and predicting individual and aggregate human transportation decisions. It is one of my research interests and an area which will continue to be transformed by emerging technologies. What does route choice even mean with autonomous vehicles on the road? See section 5.3.

4. Shared Mobility: Traditional and Emerging Modes lectern session. This was a fantastic series of talks imagining the future of transportation. A key take-away from the research presented here is that the future is wide-open with respect to how emerging transportation technology will affect us. See section 4.2

Throughout the sessions I attended there was an underlying sense of urgency to understand and model new modes of transportation: connected vehicles, automated vehicles, autonomous vehicles, ride-sharing services, ride-sourcing services, and bike-sharing networks. Around 80 sessions discussed one or more of these transformational technologies.[6] Researchers from across the subfields of transportation research seem to be asking the same fundamental question: "can existing models be extended to describe the behavior of emerging technologies in transportation?"

2 Sunday

2.1 Exhibition Hall Opening Reception

The exhibition hall housed over 200 exhibitors from industry and government.[7] Standouts included:

- DiDi Chuxing introduced the GAIA Initiative[8] to provide data and computing resources to academics in order to advance transportation research. DiDi processes over 25 million requests every day, making it the world's largest ride-sharing company.

- Einride displayed it's flagship autonomous electric truck, the T-Pod. Crucially, the T-Pod enables remote driving—a concept that has been successfully applied in unmanned aircraft for over two decades. As David Pickeral helpfully pointed out, this is a practical innovation that will ease the deployment of autonomous vehicles, and create new jobs along the way.[9]

3 Monday

3.1 Lectern 204: Traffic Equilibrium, Traffic Control, and Network Design

3.1.1 Yufeng Zhang: An Algorithm for Robust Shortest Path Problem with Travel Time Correlations [22]

The travel time on each link of a transportation network can be modeled as a random variable. Of course travel times on connected or nearby links are highly correlated, making statistical analysis difficult. This

[6] `https://youtu.be/PVD1QGZ_NN0?t=25s`
[7] `http://www.trb.org/AnnualMeeting/exhibitfloorplan.aspx`
[8] `https://outreach.didichuxing.com/research/opendata/en/`
[9] `https://www.linkedin.com/feed/update/urn:li:activity:6356528053098811392`

paper computes the mean-variance shortest path between two points in the road network. Optimization is a creative endeavor: much of the work involved is reformulating a complicated problem into something more easily solved. This paper is no exception.

The covaraince matrix is decomposed into a diagonal matrix of its eignvalues and an eigenvector matrix (e.g. $x^T \Sigma x = x^T V \Lambda V x$). Creating a new variable $y = V^T x$ adds and equivalent constraint, and the Lagrangian can be decomposed into two separable sub-problems minimizing over x and y respectively. This technique simplifies the solution procedure—the sub-problem over y is even convex! The algorithm is essentially a method of multipliers, alternatingly updating the Lagrange multiplier via a subgradient update and x and y via unspecified, but presumably, standard techniques.

3.1.2 Hanna Grzybowska: Applying Internet Protocol Random Early Detection Strategies to Real-Time Traffic Management in Transportation Networks [5]

Just as a road network can become congested with cars, the internet can become congested with packets. Random Early Detection (RED) is a method developed to detect and avoid congestion in computer networks. By monitoring traffic volume at each internet gateway, RED can drop packets in anticipation of a gateway reaching capacity. In this way packet loss is distributed more fairly than simply waiting until a gateway reaches capacity. An analogous process is proposed for transportation networks. Unlike packets, vehicles can not simply be dropped from the road network. Instead, vehicles are randomly notified of impending congestion on a link on their path so that they may avoid it by taking an alternate route.

Why not notify all drivers? This study highlights several previous studies which explore how information provision strategies impact the overall efficiency of the road network. In particular, correlations between links in the network, reliability of travel time information affect whether or not information provision is helpful to the network performance. This paper is part of ongoing efforts to develop optimal strategies for information provision.

3.1.3 Cesar Yahia: Network Partitioning Algorithms for Solving the Traffic Assignment Problem Using a Decomposition Approach [15]

The traffic assignment problem aims to determine the volume of traffic across each link in a road network given a set of origin destination pairs. Large networks pose computational challenges for algorithmic solutions because the number of paths a vehicle may take increases quickly with network size—$O(N!)$ in the worst case. By partitioning the network into smaller sub-networks and a master graph, computation can be distributed. Transportation networks are particularly well-suited for decomposition because they are constrained by geography to be sparsely connected. This paper evaluates two clustering-based methods of partitioning the network. The first minimizes the number of boundary nodes—nodes that bridge two clusters. Since overhead communication costs are a major concern in parallel computation, this is an intuitive choice. The second is a flow-weighted spectral clustering algorithm that minimizes the interflow between clusters. Minimizing interflow between clusters produces a more efficient partition for parallel computation because it reduces the overhead between traffic assignment subproblems.

3.1.4 Pinchao Zhang: A Subgradient Approach to Solve for Path-Based System Optimal Dynamic Traffic Assignment [20]

The path-based system optimal traffic assignment problem (SO-DTA) aims to find the set of paths through a road network that minimize the total travel time given a set of origin destination pairs. Typically, the problem can be solved using variational inequality (VI) algorithms. However, in some cases the travel time is not differentiable with respect to the path. This paper computes the subgradient or total travel time allowing the solution to the VI problem to be approximated numerically. This procedure converges to more accurate solutions than existing heuristic methods.

3.2 Lectern 406: Planning a Future with Autonomous and Connected Vehicles

The goal is to understand how top-line transportation statistics will change with the introduction of autonomous vehicle technology. A common theme in current AV research is how much we don't know and how much we don't know we don't know. The general public regards AVs with a degree of healthy skepticism: we've all seen apps crash on our phone, so why wouldn't they also crash in real life? Unfortunately, we don't have a good handle on the macro effects of autonomous vehicles because we don't have any understanding of how individuals will want to use them.

3.2.1 Fatemeh Nazari: Shared Mobility Versus Private Car Ownership: A Multivariate Analysis of Public Interest in Autonomous Vehicles [11]

Imagined futures of autonomous vehicles contain a mix of individually owned AVs, shared-use AVs (e.g., Car2Go, Uber), and shared ride AVs (e.g. Lyft Line). A stated preferences survey was conducted to explore the interest in one or more future AV services. This study mainly focused on consumer preferences regarding owned or shared autonomous vehicles. There are several interesting finding here, but overall it seems that interest in owning an AV was highly dependent on interest in owning a vehicle. Those who were familiar with or regularly used ride-sourcing or ride-sharing services were less likely to indicate interest in owning an AV while those who own a vehicle and rely on it for regular travel were more likely to indicate interest in AV ownership. As AVs become less hypothetical I have to wonder if attitudes will change; particularly once we know how much they will cost to own.

3.2.2 Wenwen Zhang: The Impact of Private Autonomous Vehicles on Vehicle Ownership and Unoccupied VMT Generation [21]

With conventional vehicles, fewer cars means fewer vehicle miles traveled (VMT). This is not the case with AVs because they drive themselves. For example, a two-car households might only need one AV to commute to and from work, but that AV would need to make two additional unoccupied trips. This flies in the face of what has quickly become conventional logic that AVs will universally improve traffic congestion and underscores the importance of optimal control policies for connected and automated vehicles.

3.2.3 T. Donna Chen: Impact of Ridesharing on Operational Efficiency of Shared Autonomous Electric Vehicles [2]

In contrast to privately owned AVs, shared ride AVs offer the potential to reduce both the number of cars *and* the vehicles miles traveled by aggregating riders into shared AVs. In a simulation, each AV replaces 13 conventional vehicles on average and about half of all trips were shared.

3.3 Poster 243: Statistical Methods in Transportation

3.3.1 Emmanuel Kidando: Exploring the Influence of Rainfall on a Stochastic Evolution of Traffic Conditions [9]

A Gaussian mixture model was used to cluster traffic into one of two regimes: "free-flow" and "congested." A Time-varying Markov Chain was used to fit the transition probabilities between the two regimes in the presence of rainfall.

The results are unsurprising—in the presence of rain, congestion is more likely to occur—but the simplicity and effectiveness of the method is inspiring.

3.4 Poster 255: Advances in Urban Freight Transportation Research and Practices

3.4.1 Chien-Lun Lan: A Feasibility Study for Last-Mile Synergies Between Passenger and Freight Transport for an Urban Area [12]

This project reminded me of an initiative we explored back when I worked at Bridj.[10] Passengers and freight often need to move around a city at complementary times; perhaps under-utilized public buses could be used to help move intra-city urban freight. A Mixed Integer Linear Program is proposed to determine the

[10] https://www.bostonglobe.com/business/2016/08/28/bridj-bring-boxes-bus/6vjAmqSiCRzpdTufTJ3SHO/story.html

optimal assignment of parcel to bus stop and route. Encouragingly, the authors find a significant reduction in both VMT and emissions under their scheme.

3.5 Poster 385: Artificial Intelligence and Machine Learning Tools for Estimation, Detection, and Prediction Applications in Transportation

3.5.1 Fabio Ramos: Gaussian Processes for Imputation of Missing Traffic Volume Data [13]

The Gaussian process is just one of those techniques that you can't help but admire for its simultaneous expressiveness and conceptual intuitiveness. For the uninitiated[11], a Gaussian process describes a time-varying[12] quantity using normal distributions at each time point. Moreover, it describes any pair (and any finite set) of time-points using a multivariate normal distribution (by way of a kernel function[13]), allowing it to capture correlations across time. Because each time-point is described by a normal distribution, the predicted value at every time-point comes with a ready-made confidence interval!

Here, the authors use a Gaussian process to fill in gaps of (artificially) missing data from a traffic-counting system. Expert knowledge of traffic patterns was encoded in the kernel function allowing the covariance matrix to reflect the periodicity inherent in traffic at different time scales and different times of day. This is a really interesting technique because it offers experts domain-oriented control of a vary powerful machine learning method.

3.5.2 Jiyu Zhang: Forecasting Subway Demand in Large-Scale Networks: A Deep Learning Approach [18]

This poster initially drew my eye because it applied two powerful techniques from neural network research to transit data in a unique way (as far as I'm aware). The authors transform ridership data into a heat map which is used as input to a Convolutional Neural Network[14] (CNN). The CNN extracts relevant spatial features from the ridership data which is in turn used as input to a Long Short-Term Memory Recurrent Neural Network[15] (LSTM) to capture temporal trends in ridership. The LSTM can then be used to predict future ridership for the next 10 minutes based on the previous 40.

The CNN serves to extract the spatial relationship of ridership between stations and the LSTM serves to describe the relationship of those features over time. The combination of the two was very interesting, but the transformation of the ridership into an image is counter-intuitive to me. Most pixels will not contain ridership and therefore are completely irrelevant yet are still modeled by the CNN.

Geometric Deep Learning[16], a relatively recent advancement in ML, aims to construct generalized CNN that can operate directly on the natural representation of domain, instead of 1- and 2-D lattices exclusively. Since transportation networks are naturally modeled as, well, networks, geometric deep learning seems particularly well-suited to exploit the structure of the network to extract higher-order features.

4 Tuesday

4.1 Meeting: Network Equilibrium Modeling Subcommittee

As I mentioned above, this was far and away my favorite hour and a half of the week. Network equilibrium modeling aims to understand the game people play in their heads and with each other when they travel: what is the best way for me to get from A to B and how many other people are going to come to the same answer as I will? A mind-bender not unrelated to the famous Yogi Berra-ism, "Nobody goes there anymore, it's too crowded."

[11] Katherine Bailey gives a fantastic and accessible introduction to Gaussian processes on her blog http://katbailey.github.io/post/gaussian-processes-for-dummies/

[12] it doesn't *have* to be time-varying specifically, just indexed by some set.

[13] http://www.cse.wustl.edu/~garnett/cse515t/spring_2017/files/lecture_notes/9.pdf

[14] https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/

[15] https://colah.github.io/posts/2015-08-Understanding-LSTMs/

[16] https://arxiv.org/abs/1611.08097

Equilibrium modeling is essential to models of transportation networks because it describes how people interact with it *and* react to other people interacting with it. Yet in spite of this, network equilibrium isn't on the tip of most people's tongues when it comes to transportation.

Jeff Ban[17] began the session by inviting Xuesong Zhou[18] to discuss the recent activities in the field. Professor Zhou raised three broad topics: computing network equilibrium in large-scale traffic networks efficiently, dynamic traffic assignment that can model emerging mobility-as-a-service technologies alongside one another, and stronger connections with industry.

Stephen Boyles[19] then highlighted the ongoing efforts across multiple research groups to develop robust routing algorithms for stochastic networks (see [21] and [23]).

Discussion turned to the modeling challenges posed by autonomous vehicles and other new and emerging transportation technologies. In particular, how does the notion of equilibrium change when humans are no longer driving, or even making routing choices? In terms of finding an equilibrium among a mix of conventional, autonomous, connected, and ride-sourced vehicles what quantity is being "equalized"?

4.2 Lectern 532: Shared Mobility: Traditional and Emerging Modes

4.2.1 Michael Hyland: Dynamic Autonomous Fleet Operations: Optimization-Based Strategies to Assign AVs to Immediate Traveler Demand Requests [6]

One of the most exciting aspects of the advent of autonomous vehicles is the merging of transportation modeling and real-time control. Much of transportation modeling relies of models of human behavior and we struggle to make accurate predictions because human behavior is not easily modeled. When autonomous vehicles enter the market, the decision-making model is known and can be modulated to achieve a macroscopic objective.

This paper proposes a linear program to match AVs to waiting passengers in real time—in much the same way Uber aims to match drivers and riders. There are two specific advantages of an AV fleet modeled in this paper. AVs will not opt-out of picking up a passenger like human drivers may so all assignments are immediately confirmed. Moreover, a users assigned AV may change without changing the user experience granting the system a degree of flexibility to change its mind if a better assignment emerges with future requests.

The real-time optimal control of vehicle decisions is a huge area for research because there is a critical trade off between complexity of the optimization problem and runtime. Although operations research has been very successfully applied in transportation, for example to schedule and route airline fleets, it has typically increased complexity at the expense of runtime. Flights are scheduled months in advance; no one would turn down software that takes an addition hour but finds a better schedule.

4.2.2 Meng Li: A Restricted Path-Based Ridesharing User Equilibrium[10]

Being new to the concept of equilibrium in transportation networks, I was captivated by the way this research combined mode and route-choice into a single user equilibrium problem. Typically, traffic equilibrium problems aim to find the assignment of paths to drivers that satisfies the equilibrium condition. This paper takes it a step farther by attempting to find the assignment of one of three modes—solo driver, rideshare driver, or, rideshare passenger—*and* path to a person. This is particularly relevant as mode choices increase; if I want to travel, do I take public transit, car2go, Uber, or bike share? There is plenty here to explore.

Additionally, this paper takes another look at Braess's paradox[20]: that adding an additional road segment to certain networks may make congestion worse. They show, interestingly, that the paradox does not occur if the additional road segment applies a toll to single-occupancy vehicles. With connected vehicles, or even app-based ridesharing, it may be feasible to implement such a system in real-world networks.

[17] `https://www.engr.washington.edu/facresearch/newfaculty/2016/XuegangBan`
[18] `https://isearch.asu.edu/profile/2182101`
[19] `http://www.caee.utexas.edu/faculty/directory/boyles`
[20] `https://supernet.isenberg.umass.edu/facts/braess.html`

4.3 Lectern 687: Advancing Theory and Application of Large-Scale Urban Traffic Network Models

4.3.1 Mehmet Yildirimoglu: Hierarchical Management of Heterogeneous Large-Scale Urban Networks via Path Assignment and Regional Route Guidance [17]

As connected and autonomous vehicles gain market share, traffic managers will have the ability to make much finer adjustments to traffic flow than they do now: perhaps the route of an individual vehicle. With power comes complexity and the ability to make coordinated fine adjustments in real-time is essential to making productive use of these new technologies. This research proposes a hierarchical model of urban traffic networks that allows efficient network-wide optimization of traffic flows. Similar to [15], this approach partitions the network into regions. At the region level, an optimal region-to-region flows are computed. Within each region vehicles are assigned to paths in order to achieve the optimal region-to-region flows. In this way the optimization avoids the complexity of modeling each vehicle and the within-region assignments can be done completely in parallel.

4.4 Poster 520: Advances in Intelligent Connected and Automated Transportation

4.4.1 Weinan Gao: Data-Driven Cooperative Adaptive Cruise Control of Buses on the Exclusive Bus Lane of the Lincoln Tunnel Corridor [3]

Adaptive cruise control is a very active area of research as the technology is attainable in the near term. This research addresses two orthogonal impediments to its implementation. First, traditional optimal control rely on a model of the system known as the plant. If the plant model is wrong, the control is suboptimal. Here, the authors estimate the system parameters directly from the observed vehicle data thereby avoiding the process of specifying am a priori model of vehicle dynamics. Further, control systems are often centralized. Adaptive cruise control is a local problem; centralized control both scales poorly with the fleet size and offers little benefit. The proposed method avoids the communication overhead by utilizing vehicle-to-vehicle communication to solve the control problem locally.

5 Wednesday

5.1 Lectern 869: Connected Multimodal Transportation System Modeling and Simulation

5.1.1 Mohsen Ramezani: Lane Distribution Optimization of Autonomous Vehicles for Highway Congestion Control [16]

Increased roadway capacity is one of the more popular perceived benefits of autonomous and connected vehicles. It's easy to imagine: autonomous vehicles can organize themselves into tightly packed platoons that can move as a unit. One can also imagine optimizing the micro-interactions that happen constantly between vehicles on the road. This research explores an optimal highway lane-changing strategy in the presence of vehicles merging from an on-ramp.

The authors propose a bi-level method: a proactive control method to optimize the distribution of cars with respect to lanes in real-time and a reactive control method in the vicinity of the on-ramp to coordinate specific lane changes to prevent interference with a merging vehicle.

There are two important (and related) themes that are again found in this work. First, optimal control is computed and actuated in real-time. Second, optimal control is piecewise: optimal control means a different thing immediately before an on-ramp than it does on uninterrupted stretches of highway. Concretely, algorithms that attempt to optimize global state are likely to be too computationally complex to satisfy the run-time constraints required for real-time control. However, the prospect of a patchwork of local optimal control schemes is somewhat nightmarish—to the software engineer responsible for implementing it and to the researcher responsible for analyzing its performance.

5.1.2 Kyungwon Kang: Modeling Driver Merging Behavior: A Repeated Game-Theoretic Approach [8]

At a high level, I think a completely autonomous future is intuitive because its easy to imagine how two AVs might communicate in an intelligent way. What is almost entirely opaque is how AVs will interact with human drivers. Broadly speaking this research highlights how important models of human behavior are in a semi-autonomous world. The authors extend their previous research [7] in which they introduce the game theoretic approach to simulate a repeated game as a way to model that drivers involved in a merge are continuously evaluating and making decisions. Encouragingly, the repeated game more accurately predicted merge behavior than a single game.

5.1.3 Amir Ghiasi: Markov Chain-Based Mixed Connected Automated Traffic Capacity Analysis [4]

Mixed traffic analysis has traditionally considered the traffic dynamics of traffic with a mix of vehicle classes (e.g. trucks and passenger vehicles). With the rise of AVs, similar analysis must be performed with a mix of human and autonomous vehicles. This research models highway capacity as a function market penetration and platooning intensity of connected automated vehicles expressed very elegantly in the transition probabilities of a Markov chain. The results of the model suggest that, counter-intuitively, higher levels of CAVs and inclination to platoon can *increase* congestion in certain circumstances.

5.2 Poster 864: Transportation Network Modeling

5.2.1 Fatemeh Fakhrmoosavi: Decomposition of Stochastic Time-Varying Networks for the Path-Finding Problem Considering Travel Time Correlations and Heterogeneity of Users [1]

Similar to [22], this paper proposes a shortest path algorithm for road networks where the travel time on each link is stochastic and possibly correlated with nearby links. Further, the authors devise a mechanism to extract a sub-network for the queried O-D pair. Deterministic optimistic and pessimistic travel times are computed using the minimum and maximum link travel times respectively. Nodes are then removed from the network if the optimistic travel time from the origin to destination through that node is greater than the pessimistic travel time from that node to the destination. This technique is reminiscent of the bounding principle of the branch and bound[21] algorithm in optimization theory. Stochastic shortest path is then done on the sub-network.

The decomposition to a smaller network dramatically improves running time by reducing the dimensionality of the problem. Moreover, the authors showed that the extracted network can be made smaller with minimal degradation in accuracy by removing nodes more aggressively.

5.3 Poster 863: Behavioral Route Choice

5.3.1 Kenan Zhang: Mitigating the Impact of Selfish Routing: An Optimal Ratio Control Scheme (ORCS) Inspired by Autonomous Driving [19]

Selfish routing is a non-cooperative travel behavior that contributes to congestion yet is locally optimal from the perspective of each individual agent. Conventional passenger vehicles lack mechanisms for cooperative route choice behavior. However, the rise of connected vehicles, ride-sourcing/ride-sharing platforms, and autonomous vehicles offers new opportunities for coordinated control of route-choice which can be used to reduce congestion. This paper proposes a method to determine the optimal percentage of controlled vehicles between each O-D pair in a road network. By minimizing the weighted sum of system performance (total travel time) and control intensity (total amount of controlled flow).

Their results indicate that there can be a significant reduction in congestion with a relatively small number of cooperative vehicles. This is particularly relevant because it offers a path for cities to institute congestion policies on a subset of the fleet that is easiest to control like ride-sourcing services or autonomous

[21] `https://www.cs.cmu.edu/~aarti/Class/10725_Fall17/Lecture_Slides/Discrete_Optimization.pdf` (pages 19-29)

vehicle fleets. This is particularly relevant given research like [21] and it wouldn't be difficult to imagine cities implementing policies to require unoccupied AVs to route themselves cooperatively.

5.3.2 Xuesong Zhou: Enabling Transportation System Intelligence: Hierarchical Estimation of Traffic Network State and Behavior Coefficients Using a Computational Graph Approach Modeling Framework [14]

This paper borrows from deep learning architectures to create a computational graph model to estimate unobservable network state variables (trip generation, O-D matrices, and path/link flows) from a synthesis of observable data sources: travel surveys, cellphones, and road sensors. The input layer generates trips, the second layer represents O-D pairs to which the trips are distributed, the third layers represents the paths taken between origins and destinations, and the fourth and final layer represents road network links. The connections between each layer represent the spatial distribution of trips, the route choice of each trip, and the traffic flow across each link based on the route respectively. As in the artificial neural networks that inspired it, prediction errors are backpropogated through the entire network. At each layer observed data from a different source is used to compute the estimation error: survey data at the trip generation layer, cellphone data at the route choice layer, and sensor data at the link flow layer.

I was struck by how nicely this model is able to integrate disparate data sources together into a single model of the road network. The data used in this model describe fundamentally different things, yet are fit together in a principled way to describe the system as a whole. As the volume and richness of collected data increases, so too will the need for models that can piece them together to form a cohesive representation of system state.

6 Closing Thoughts

The conference was an invaluable experience and I've returned to school energized about by ongoing projects and how I might better integrate them with the current movement within the field.

Several threads emerged from the research I was able to see at TRB. Whether or not autonomous and connected vehicles fulfill their promise to improve congestion is entirely dependent on how they are used—and how they interact with human drivers. Of particular importance are the optimal control algorithms that are in place. Optimal control must happen in real time so reducing computational complexity by decomposing large problems is essential. At the same time, the interactions between the patchwork of optimal control policies at work (will Ford AVs use the same algorithms as GM AVs?) will need to be modeled at the macroscopic level in order to identify where adverse interactions of control policies are occurring.

References

[1] Fatemeh Fakhrmoosavi, Ali Zockaie, Khaled Abdelghany, and Hossein Hashemi. Decomposition of Stochastic Time Varying Networks for the Path Finding Problem Considering Travel Time Correlations and Heterogeneity of Users. In *97th Annual Meeting Transportation Research Board*, pages 1–20, 2018.

[2] J. Farhan and T. Donna Chen. Impact of Ridesharing on Operational Efficiency of Shared Autonomous Electric Vehicle Fleet. In *97th Annual Meeting Transportation Research Board*, 2018.

[3] Weinan Gao, Zhong-Ping Jiang, Kaan Ozbay, and Jingqin Gao. Data-driven Cooperative Adaptive Cruise Control of Buses on the Exclusive Bus Lane of the Lincoln Tunnel Corridor. In *97th Annual Meeting Transportation Research Board*, 2018.

[4] Amir Ghiasi, Omar Hussain, Zhen (Sean) Qian, and Xiaopeng Li. Makov-chain Based Mixed Connected Automated Traffic Capacity Analysis Part I : Analytical Modeling. In *97th Annual Meeting Transportation Research Board*, 2018.

[5] Hanna Grzybowska, Steven Willmott, and S Travis Waller. Applying Internet Protocol Random Early Detection Strategies. In *97th Annual Meeting Transportation Research Board*, 2018.

[6] Michael Hyland and Hani Mahmassani. Dynamic Autonomous Vehicle Fleet Operations: Optimization-Based Strategies to Assign AVs to Immediate Traveler Demand Requests. In *97th Annual Meeting Transportation Research Board*, pages 1–11, 2018.

[7] Kyungwon Kang and Hesham A. Rakha. Game Theoretical Approach to Model Decision Making for Merging Maneuvers at Freeway On-Ramps. *Transportation Research Record: Journal of the Transportation Research Board*, (2623):19–28, 2017.

[8] Kyungwon Kang and Hesham A. Rakha. Modeling Driver Merging Behavior: A Repeated Game Theoretical Approach. In *97th Annual Meeting Transportation Research Board*, 2018.

[9] Emmanuel Kidando, Ren Moses, Angela E Kitali, Valerian Kwigizile, Sia Macmillan Lyimo, Deo Chimba, and Thobias Sando. Exploring the influence of rainfall on a stochastic evolution of traffic conditions. In *97th Annual Meeting Transportation Research Board*, pages 1–17, 2018.

[10] Meng Li, Xuan Di, Henry Liu, Ann Arbor, and Hai-Jun Huang. A Restricted Path-Based Ridesharing User Equilibrium. In *97th Annual Meeting Transportation Research Board*, 2018.

[11] Fatemeh Nazari, Mohamadhossein Noruzoliaee, and Abolfazl (Kouros) Mohammadian. Shared Mobility vs . Private Car Ownership: A Multivariate Analysis of Public Interest in Autonomous Vehicles. In *97th Annual Meeting Transportation Research Board*, 2018.

[12] Moschoula Pternea, Chien-Lun Lan, Ali Haghani, and Shih Miao Chin. A Feasibility Study for Last-Mile Synergies Between Passenger and Freight Transport for an Urban Area. In *97th Annual Meeting Transportation Research Board*, 2018.

[13] Fabio Ramos, Heudson Mirandola, Douglas Picciani, Glaydston Mattos Ribeiro, Ivani Ivanova, Saul Germano Rabello Quadros, Romulo Dante Orrico Filho, Leonardo R. Perim, and Carlos A. Abramides. Gaussian Processes for Imputation of Missing Traffic Volume Data. In *97th Annual Meeting Transportation Research Board*, 2018.

[14] Xin Wu, Jifu Guo, Ruixi Lin, Kai Xian, and Xuesong Zhou. Enabling Transportation System Intelligence: Hierarchical Estimation of Traffic Network State and Behavior Coeffients Using a Computational Graph Approach Modelling Framework. In *97th Annual Meeting Transportation Research Board*.

[15] Cesar N Yahia, Venktesh Pandey, and Stephen D Boyles. Network Partitioning Algorithms for Solving the Traffic Assignment Problem Using a Decomposition Approach. In *97th Annual Meeting Transportation Research Board*, 2018.

[16] Eric Ye and Mohsen Ramezani. Lane Distribution Optimisation of Autonomous Vehicles for Highway Congestion Control. In *97th Annual Meeting Transportation Research Board*, 2018.

[17] Mehmet Yildirimoglu, Isik Ilber Sirmatel, and Nikolas Geroliminis. Hierarchical Management of Heterogeneous Large-Scale Urban Networks Via Path Assignment and Regional Route Guidance. In *97th Annual Meeting Transportation Research Board*, 2018.

[18] Jiyu Zhang, Xiaolei Ma, Chuan Ding, Yupeng Wang, and Jianfeng Liu. Forecasting subway demand in large-scale networks: a deep learning approach. In *97th Annual Meeting Transportation Research Board*, 2018.

[19] Kenan Zhang and Yu Nie. Mitigating the impact of selfish routing: An optimal-ratio control scheme (ORCS) inspired by autonomous driving. In *97th Annual Meeting Transportation Research Board*, 2018.

[20] Pinchao Zhang and Zhen (Sean) Qian. A subgradient approach to solve for path-based system optimal dynamic traffic assignment. In *97th Annual Meeting Transportation Research Board*, 2018.

[21] Wenwen Zhang, Subhrajit Guhathakurta, and Elias B Khalil. The Impact of Private Autonomous Vehicles on Vehicle Ownership and Unoccupied VMT Generation. In *97th Annual Meeting Transportation Research Board*, 2018.

[22] Yufeng Zhang and Alireza Khani. An Algorithm for Robust Shortest Path Problem with Travel Time Correlations. In *97th Annual Meeting Transportation Research Board*, 2018.

[23] Shuaidong Zhao and Kuilin Zhang. A data-driven dynamic route choice model under uncertainty using connected vehicle trajectory data. *97th Annual Meeting Transportation Research Board*, 2018.