



Advanced Computing, Data Science, and Artificial Intelligence Research Opportunities for Energy-Focused Transportation Science

David Biagioni, John Farrell, Venu Garikapati, Peter Graf,
Nalinrat Guba, Yi Hou, Wesley Jones, Joe Severino,
Devon Sigler, Austin Todd, Juliette Ugirumurera,
Qichao Wang, and Stan Young

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy
Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-2C00-79589
July 2021



Advanced Computing, Data Science, and Artificial Intelligence Research Opportunities for Energy-Focused Transportation Science

David Biagioni, John Farrell, Venu Garikapati, Peter Graf,
Nalinrat Guba, Yi Hou, Wesley Jones, Joe Severino,
Devon Sigler, Austin Todd, Juliette Ugirumurera,
Qichao Wang, and Stan Young

Suggested Citation

Biagioni, David, John Farrell, Venu Garikapati, Peter Graf, Nalinrat Guba, Yi Hou, Wesley Jones, Joe Severino, et al. 2021. *Advanced Computing, Data Science, and Artificial Intelligence Research Opportunities for Energy-Focused Transportation Science*. Golden, CO: National Renewable Energy Laboratory. NREL/ TP-2C00-79589.
<https://www.nrel.gov/docs/fy21osti/79589.pdf>.

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/ TP-2C00-79589
July 2021

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Internal funding was provided by NREL. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

List of Acronyms

CAV	connected/autonomous vehicle
CNN	convolutional neural network
DRL	deep reinforcement learning
EEMS	Energy Efficient Mobility Systems
GPU	graphics processing unit
R-CNN	region-based convolutional neural network
RL	reinforcement learning
SAS	symbiotic autonomous systems
SSD	single-shot detection
YOLO	you-only-look-once

Table of Contents

List of Acronyms	iv
Table of Contents	v
List of Figures	v
List of Tables	v
1 Introduction	1
2 Research Needs	2
2.1 Using Large, Detailed Telemetry Data Sets for Activity-Based Demand Modeling	2
2.2 Streaming Real-Time, High-Fidelity Data for Optimal Route Planning	3
2.3 Computer Vision for Mobility Data Collection	4
2.4 Data Science Techniques for Modeling, Fusion, and Prediction	4
2.5 Artificial Intelligence Techniques for Automated, Real-Time Decision-Making	5
Internet of Things Data in Mobility	5
Mobility Sensor Health and Data Anomaly Detection	5
2.6 Simulation and Modeling to Provide Unprecedented Insight into Potential Futures	6
High-Fidelity Modeling for Dynamic Movement and Behavior	6
Tools for Cities	6
Optimal Resource Allocation and Infrastructure Investments	6
2.7 Simulation Techniques to Take Full Advantage of Artificial Intelligence	7
Simulation-Trained Deep Reinforcement Learning	7
Real-Time Control and Digital Twins	7
Digital Twins, Airports, and Multimodal Transportation	8
Regional Mobility Optimization	9
2.8 Scalable/Distributed Optimization and Control	10
Integration with the Electric Grid and Other Infrastructure Systems	10
Optimally Managing CAV Fleets	11
Optimal Traffic Control	12
2.9 Symbiotic Autonomous Systems	12
3 Summary of Research Needs	14
References	16

List of Figures

Figure 1. Display of live-feed TomTom relative speed profiles of Chattanooga, Tennessee, region with traffic congestion (left) and dispersion of congestion (right)	3
Figure 2. Display of energy density (gasoline gallon equivalent) per meter between morning off-peak (left) and morning peak (right) of Chattanooga, Tennessee, region	8
Figure 3. Impact of traffic management policy at DFW	9

List of Tables

Table 1. Summary of Research Needs to Accelerate Progress Toward a Vision of an Energy-Efficient and Electricity-Symbiotic Transportation System.	17
--	----

1 Introduction

To improve energy efficiency and electricity utilization of transportation, early-stage research and development must be pursued at the vehicle platform, traveler, and system levels. Significant changes in the mobility and electrification landscape are underway as a result of the advent of vehicle and infrastructure connectivity, autonomous driving, and rapid passenger- and freight-vehicle electrification. Meanwhile, new business models, such as first-and last-mile delivery, electric charging, multimodal freight, and mobility as a service, could transform the mobility landscape. These changes necessitate data infrastructure investments (e.g., secure-streaming data platforms driven by ubiquitous sensing and video analytics) as well as investments in critical capabilities for large-scale automated analysis and organization using modern machine learning, statistics, and artificial intelligence. Other chief needs include agile, large-scale storage that can be quickly searched and queried for relevant data to support validation and model development, data-sharing agreements, and formatting standards for key data types.

Data science, artificial intelligence, and advanced computing will play an increasingly important role in enabling energy-focused transportation science researchers to understand and identify the most important drivers and levers to improve the energy productivity of future integrated mobility systems. In this report, we outline and summarize energy-focused research needs and opportunities which depend on each of these three emerging and evolving computational technologies with a focus on accelerating and improving research via methods that rely on high-fidelity representations of reality.

Data Science includes the acquisition and management of large streaming and batch datasets and the development and application of advanced analysis and visualization workflows. Large, detailed telemetry data sets and complex streaming real-time data from vehicles and infrastructure and proliferation of advanced sensors are driving opportunities in computer vision, visualization, data fusion, and the application of different forms of artificial intelligence.

Artificial Intelligence includes advances in deep learning for supervised and unsupervised tasks as well as advances in reinforcement and other learning for rapid decision making. In conjunction with advanced computing, deep learning leverages data to learn things like traffic volume, speed and energy use across regions and the nation. Additional applications include understanding the performance of infrastructure like traffic signals. Deep reinforcement learning (DRL) combines the power of deep learning with reinforcement learning (RL) to create a computational framework that can solve previously unsolvable sequential-decision problems like optimal fleet routing and traffic signal timing.

Advanced Computing includes the utilization of accelerated computing (e.g. GPUs) and hybrid use of high-performance computing, cloud and edge computing. Emerging workflows include deep neural networks that are trained on high performance computers with data from the cloud and deployed on edge devices. Digital twins and simulations of the real world allow the training of artificial intelligence and the exploration of optimal actions during extreme events and other scenarios where real-world consequences need to be avoided or deployments don't yet exist.

2 Research Needs

The strategic application of advanced computing, data sciences, and artificial intelligence can help accelerate progress toward a vision of an affordable, efficient, safe, and accessible transportation future coupled to electricity in which mobility is decoupled from emission generation and energy consumption. In this section, we describe nine categories of research that are driving a major increase in the use of advanced computing, data science and artificial intelligence that could potentially advance energy-focused transportation science. The research needs and opportunities are summarized in Table 1 at the end of the report (page 17) and include the following:

1. Utilizing large, detailed telemetry data sets for activity-based demand modeling
2. Streaming real-time, high-fidelity data for optimal route planning
3. Computer vision for mobility data collection
4. Data science techniques for modeling, fusion, and prediction
5. Artificial intelligence techniques for automated, real-time decision making
6. Simulation and modeling to provide unprecedented insight into potential futures
7. Simulation techniques to take full advantage of artificial intelligence
8. Scalable/distributed optimization and control
9. Symbiotic autonomous systems

2.1 Using Large, Detailed Telemetry Data Sets for Activity-Based Demand Modeling

Understanding emerging mobility trends in passenger travel and freight mobility requires identifying and exploiting the interconnectivity of automation, electrification, and vehicle connectivity. Advancements made using agent-based models for activity-based demand modeling help answer questions about how a road-user charge might impact travel decisions of individuals from varying socioeconomic backgrounds or whether flexible work schedules would impact the activity-travel patterns of individuals. However, even the most advanced agent-based models rely on traditional travel survey data (e.g., travel diary data), which is expensive and burdensome to collect. Hence, traditional travel surveys are done infrequently and often solicit responses from a small sample of the population.

Accessing location history and telemetry data from personal and corporate devices can address these shortcomings and provide a high-resolution, low-cost source of travel information. Despite some challenges concerning fusion of traditional, nontraditional, and multidomain data sets, high-resolution, high-quality data have the potential to enhance the capabilities of state-of-the-art activity-based models. Recent studies have used location-history data coupled with advanced machine-learning models to develop several components of travel demand, such as activity start time and duration, destination location for the activity, preferred travel mode choice, and activity sequencing (Vij and Shankari, 2015; Sadeghvaziri et al., 2016; Ruktanonchai et al., 2018; Zong et al., 2019). Advanced mathematical/computing challenges for demand modeling include the development of unified frameworks that can be applied to a wider range of shared and

multimodal scenarios and addressing a modeling boundary, including the spatial and temporal extent. In addition, improved computational infrastructure and expertise could potentially enable the agile processing of very large data sets, including the application of different computational and mathematical frameworks associated with data science activities.

2.2 Streaming Real-Time, High-Fidelity Data for Optimal Route Planning

The proliferation of vehicles with onboard sensors and data storage/transmission capabilities is contributing to massive surge in data pertaining to people, and vehicle movements. Companies are using streaming data to reduce fuel costs, improve customer service, and increase operational efficiency. Data collected through vehicle-to-vehicle and vehicle-to-infrastructure communication can save time and energy via coordinated signal timing and real-time traffic optimization.

Optimal route planning is essential for providing reliable service with strict timing constraints as well as for minimizing operating costs, energy consumption, and battery storage requirements. When road conditions and passenger load change in real time, route planning has to be done in a real-time fashion to adapt the route to changing conditions. Optimally managing fleets of autonomous and electric vehicles is particularly challenging and requires a cross-disciplinary effort that spans the broad areas of transportation system modeling, power systems modeling, optimization, and data analytics.

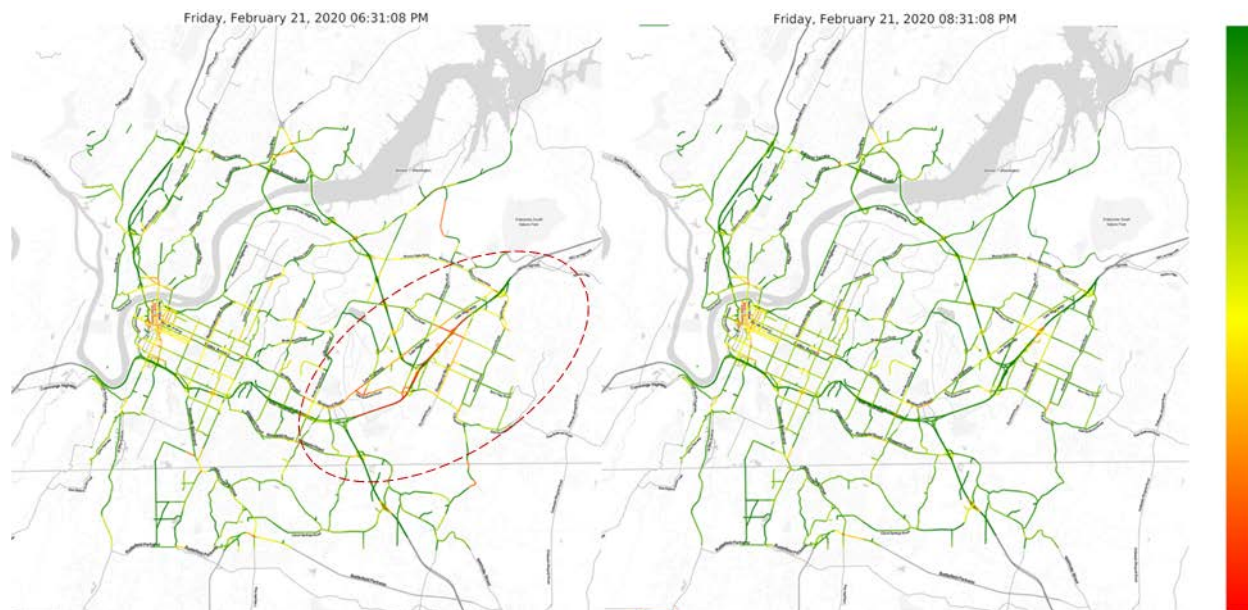


Figure 1. Display of live-feed TomTom relative speed profiles of Chattanooga, Tennessee, region with traffic congestion (left) and dispersion of congestion (right)

Critical research needs include the development of computationally efficient and distributed algorithms that can dispatch routes, charging commands, and service commands in real time, which is challenging given the large-scale nature of the models as well as the nonconvexity of the associated optimization problems. In addition, innovative online learning methods are required to process high volumes of data in real time, deal with missing measurements, and reliably predict both traffic conditions and power-system states. Figure 1 shows an example

traffic speed from TomTom (based on floating car or probe vehicle data) that is updated in real-time. While this data has intrinsic value in itself, process such real time data streams to predict impending traffic volume can greatly amplify the utility of such data streams.

2.3 Computer Vision for Mobility Data Collection

Cameras can be widely deployed for data collections as they are fairly low in cost. By leveraging state-of-the-art deep-learning, and computer-vision technologies, massive amounts of mobility data can be extracted from videos and images. For example, computer vision can help cities and transportation agencies collect infrastructure inventory data (e.g., parking space), monitor curbside activities (e.g., transportation network company pickup/drop-off, truck loading/unloading), perform traffic counts at intersections (e.g., vehicle, pedestrian, and bike counts), and determine vehicle composition (e.g., electric vehicle share, powertrain composition) and vehicle occupancy (e.g., in-vehicle passenger detection). These data support not only daily operation and long-term planning activities but also real-time decision-making. Most of computer-vision algorithms are computationally intensive. The U.S. Department of Energy’s high-performance computers and graphics processing unit, or GPU, clusters enable computer-vision application at large scale and with less time latency.

Several state-of-art computer-vision models, including single-shot detection (SSD), region-based convolutional neural networks (R-CNN), faster R-CNN, and you-only-look-once (YOLO), have been successfully used in areas such as autonomous driving, image recognition, and satellite image analysis. Research is needed to customize existing computer-vision techniques for mobility applications. For each new application, computer-vision models need to be retrained, which often requires an extensive set of labeled images. Hyperparameter tuning, which is the most critical step in model training, demands significant computational resources.

Computer vision holds the promise for achieving the ultimate vision of “cars that don’t crash.” Achieving this vision requires redundancy and perfect accuracy for safety-critical applications (e.g., an instrumented intersection that can detect a potential crash scenario and turn all the lights to red). Advances are also needed in the fusion of multimodal data (e.g., cameras, lidar, and radar), 3D dynamic predictions (e.g., velocity), increased computational speed, and acceleration of methods development using data generated through simulation.

2.4 Data Science Techniques for Modeling, Fusion, and Prediction

Developing and calibrating new agent-based models for cities or other geographical regions is currently challenging and time-consuming. Fusing traditional data sources with emerging data streams would both automate and significantly accelerate the development of high-fidelity models. Also, computational treatment of agents would need to be enhanced to account for behaviors associated with socioeconomic, and demographic attributes such as income, disability, education, and gender. In addition, our focus must evolve to include a more explicit accounting of passengers and freight/packages as well as security and resilience to identify ways to adaptively deal with technological convergence, surprises, and disruptions and to identify emergent properties. Future travel models need to be agile in capturing operations of existing modes and incorporating new modes as needed. The emergence of automation operational design domains (Czarnecki, 2018) will result in travel pattern shifts. And the earliest shifts will likely

involve the interstate highway system as vehicle travel increasingly challenges short-hop airlines, shifting more passenger miles to intercity interstate routes.

2.5 Artificial Intelligence Techniques for Automated, Real-Time Decision-Making

Internet of Things Data in Mobility

Within mobility, the Internet of Things is being realized, with vehicles serving as sensors. Crowd-sourced traffic data have already transformed traffic reporting. Over the last decade, “eyes in the sky” traffic reporting has been displaced by probe traffic data intelligence in which a small fraction of vehicles with either telematics or connected smartphone apps provide enough data to infer traffic flow at most locations. These low-penetration sampled data are being replaced by original equipment manufacturer-supported direct Internet of Things data feeds from vehicle fleets. Such data are already available from approximately 10% of vehicles, reporting position, speed, and other parameters at 3-second intervals. And the data are expected to get denser and faster, and to approach real-time reporting of all vehicles on all roadways. These expected trends represent a real-time large-data problem that will need to be addressed to monitor roadway conditions, optimize signal timing (adaptively), and reap the benefits of reduced delay and increased efficiency and safety.

The same spatial-sensing technologies that will eventually enable self-driving vehicles (e.g., lidar, video processing, and radar technologies) are driving a transformation in spatial intelligence. When applied to roadway infrastructure, sensors can provide dependable and redundant systems to deconflict intersections (preventing collisions and enabling safe, automated transit concepts), monitor pickup and drop-off zones to enable enforcement of curbside policy, and monitor areas (pedestrian crossings) for safe operations. For this to happen, data from sensors need to be filtered, fused, and operated upon and then processed, optimized, and reduced to edge-computing devices so that large-bandwidth sensor data do not have to be transmitted. High-performance computing is needed for the algorithms, AI, and machine learning to enable these processes.

Mobility Sensor Health and Data Anomaly Detection

Traffic sensors—including radar, infrared, Bluetooth, and cameras on the roadway systems—benefit traffic operation, planning, and decision-making. However, these traffic sensors need to be well maintained to serve accurately collection data. Lack of proper maintenance and severe weather conditions lead to unreliable traffic data that can potentially mislead decision makers. The large amount of data received daily by servers makes it impossible to manually identify false data. The current state of practice on sensor monitoring and data quality management for the departments of transportation is to send technicians to the field to check on sensors periodically and apply some heuristics to identify anomalies once data is collected. This strategy is costly, inefficient, and not accurate. To obtain good-quality traffic data from sensors, machine learning-based anomaly detection and health prognostic algorithms can serve as a cost-effective alternative for traffic sensor prognostic and data anomaly detection.

Current machine learning based anomaly detection algorithms including both supervised and unsupervised learning have been successfully applied to machinery (e.g. oil drilling equipment,

construction trucks, and airplane engines) health management. Few researchers explored machine learning to detect traffic sensor and data anomalies in transportation field. Significant research efforts are needed to customize existing machine learning algorithms to traffic data as anomalies in traffic can mean two things: 1) extreme high traffic or low traffic due to special events, incidents, or road closure, 2) extreme high traffic or low traffic due to sensor malfunction. The goal is to develop algorithms that differentiate the latter anomalies from the former.

In the next generation of transportation where connected and automated vehicles (CAVs) will be introduced to the road, sensor health and data quality will be of paramount importance since sensor malfunction in such scenario can cause severe consequences. Petabytes of data will be generated by CAVs in every single hour. Research is also needed to develop computational efficient anomaly detection algorithms with minimum time latency by leveraging high performance computers and GPU clusters.

2.6 Simulation and Modeling to Provide Unprecedented Insight into Potential Futures

High-Fidelity Modeling for Dynamic Movement and Behavior

High-fidelity modeling of integrated transportation networks has strengthened our understanding of dynamic movement and behavior patterns, allowing us to project future energy use under alternative scenarios as new travel modes enter the mobility scene. Dynamic models with such capabilities are just now being developed with fine-grained geospatial (i.e., lat/longs of each trip) and temporal (i.e., trip start/end times modeled to the hour/minute) resolution. Such disaggregate resolution will help accurately quantify the congestion and energy related outcomes for various mobility futures. Combining these new modeling capabilities with basic urban science could enable researchers to explore the mobility/energy impacts of new modes and identify subtle interactions between mobility systems, ultimately facilitating the design of optimized mobility systems that are more sustainable and economically productive.

Tools for Cities

The current U.S. Department of Energy's Energy Efficiency and Mobility Systems (EEMS) program efforts focus on developing tools and computational capabilities primarily for researchers, but providing accessible tools to U.S. cities could deliver greater societal impact. Such tools could enable rapid impact assessments of emerging transportation options through visualization of scenarios, such as by identifying the impacts of new mobility options (e.g., connectivity, automation, and electrification) and business models (e.g., ride sharing) to improve efficiency, maximize mobility energy productivity, reduce congestion, decrease consumer costs, lower environmental impacts, and provide more inclusive and equitable transportation solutions.

Optimal Resource Allocation and Infrastructure Investments

The outcomes of real-world mobility experiments (for example, Innisfil town in Canada replaced public transit with Uber [Cecco 2019]) could be analyzed and modeled, and lessons learned from successes and failures could be applied to future analyses. Informing optimal resource allocation and infrastructure investment requires decision-making tools that model the mobility of passengers and freight in real time, for representative rural and metropolitan areas in the United States, over current and potential future road, and infrastructure networks.

2.7 Simulation Techniques to Take Full Advantage of Artificial Intelligence

Simulation-Trained Deep Reinforcement Learning

Deep reinforcement learning (DRL) combines the power of deep learning with reinforcement learning (RL) to create a computational framework that can solve previously unsolvable sequential-decision problems. Google’s AlphaGo Zero is one example of DRL. In this paradigm, the policies trained to solve the problem are learned without human-generated data sets or expert players. Instead, a simulation of the game or reality is constructed and the DRL algorithm is trained by repeatedly experiencing the simulated environment and learning, for a given state of the system, which actions lead to the highest expected rewards.

Research needs include extension to real-world, stochastic environments. Recent success stories for DRL have been in the domains of game playing (e.g., chess, Go, and StarCraft) and controlled robotics applications (e.g., autonomous vehicles). A basic RL research need for energy applications is to extend and apply these algorithms to real-world environments that are stochastic and have multiple decision-making agents, complex dynamics, and limited observability.

Additionally, many problems that arise in control and planning optimization are combinatorial in nature and thus computationally difficult to solve, which limits the extent to which they can be solved in real time or at large scales. As evidenced by a growing body of recent research (Silver et al. 2016, 2017), there is a fundamental need to explore the use of RL, in conjunction with other ML methods, to solve such problems.

The application of RL to real-world problems involving vehicle perception and control is at a very early stage, and it is unclear which parts of the connected/autonomous vehicle (CAV) application space are amenable to RL, from simple path following to communicating and cooperatively planning with other vehicles and infrastructure. Research is needed to identify relevant observation and action spaces, and reward functions. Efforts are also needed to select the distribution of scenarios for training vehicles (e.g., path following, car following, passing, merging, and all the ways to more complex multi-vehicle scenarios, eventually encompassing traffic-level simulations and more). The correct distribution (including sequence) of scenarios is critical to avoid “catastrophic forgetting.”

Real-Time Control and Digital Twins

Implementing real-time control at large scale is very challenging. The U.S. Department of Energy’s EEMS program is funding the development of “digital twins”—well-characterized, data-informed operational models that capture real-time traffic conditions in a given region via in-road sensors—for Los Angeles, California, and Chattanooga, Tennessee. These dynamic traffic simulations can be used to investigate the causes of traffic congestion and identify effective management strategies that city officials can apply. One of the promises of digital twins is the ability to learn optimal policies in simulation for operating a system based on real-time data. For example, Figure 2 depicts the energy density per meter on the Chattanooga road network during different time periods of the day. It can be observed that the energy density during off-peaks is much lower than that of the peak periods. Through digital twins it is not only

possible to generate such information using historic data but is also possible to generate energy density of alternate routing strategies (in real-time) in the event of unplanned network disruptions. AI and HPC capabilities would serve as the pillars on which digital twins can be mounted for efficient, real-time traffic operations.

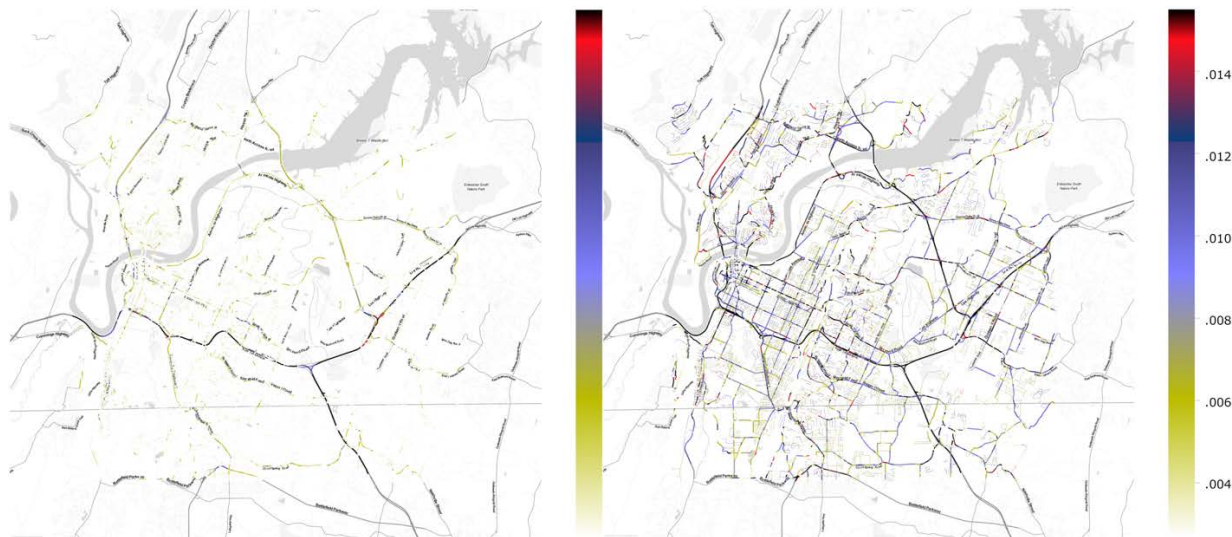


Figure 2. Display of energy density (gasoline gallon equivalent) per meter between morning off-peak (left) and morning peak (right) of Chattanooga, Tennessee, region

Digital Twins, Airports, and Multimodal Transportation

Complementary efforts are underway at two of the world’s largest multimodal transportation hubs, Dallas-Fort Worth and O’Hare airports, where digital twins of airport passenger and freight traffic are under development to support long-term planning efforts. These models simulate the impacts of future capacity expansion scenarios and identify options that maximize the value of passenger and freight mobility per unit of energy and cost. Figure 3 shows the impact of a simple traffic management policy on traffic congestion at DFW using high-fidelity microscopic simulation. The simulation is based on an Open Street Map of DFW shown in Figure 3.a. The simulation represents a high-volume day in June 2018 at DFW, where an additional 72,000 additional vehicles travel to/from the airport. Figure 3.b shows the traffic congestion observed at DFW without any traffic management intervention. The implemented simple intervention policy assumes that vehicles going to DFW are informed of anticipated congestion at their destination terminal using a ML-based demand prediction model (Lunacek et al. 2021). These vehicles are then advised to go to another terminal or parking lot to drop off their passenger to avoid delays. As depicted in Figure 3.c, this simple traffic management policy reduces the overall congestion levels observed at DFW. The simulation results demonstrated a 34% reduction in fuel consumption and 38% reduction in delay due to this simple intervention policy. The underlying analysis frameworks can be readily applied at other regionally integrated transportation hubs such as train stations, rapid transit stations, seaports, and truck terminals.

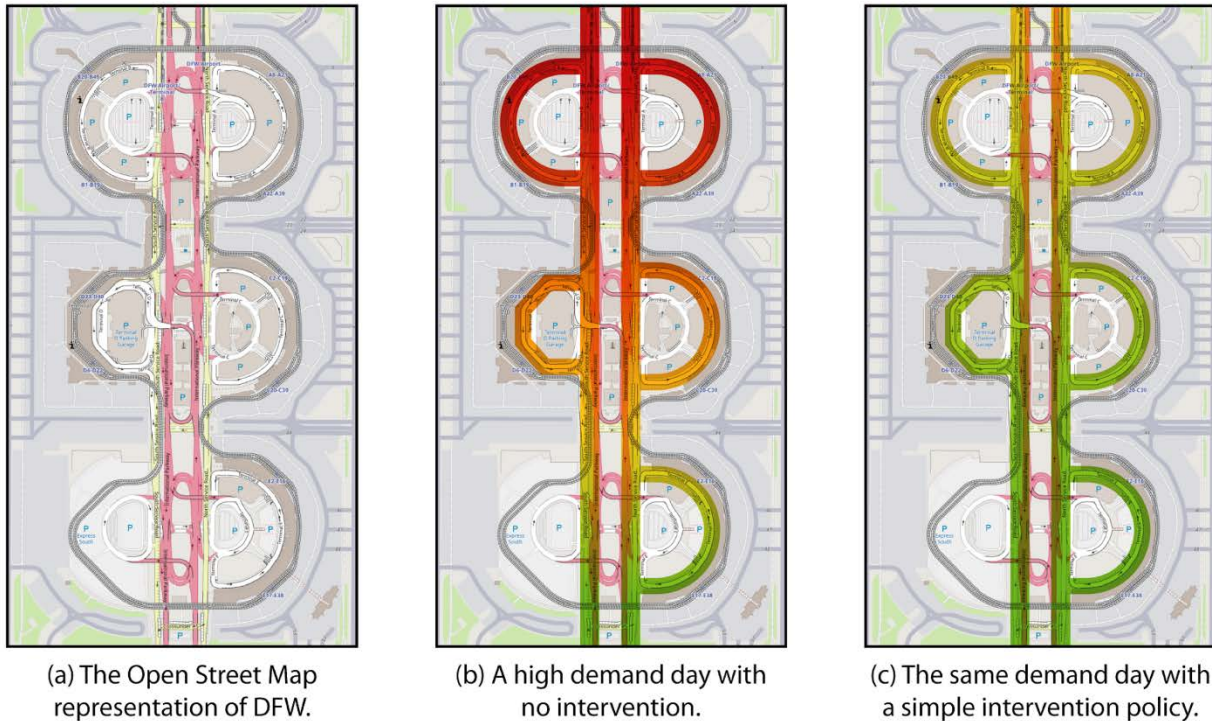


Figure 3. Impact of traffic management policy at DFW

The microsimulation model is based on the Open Street Map of DFW shown in 2a. The high demand period we simulated creates a network that experiences congestion during peak traffic periods as shown by red coloring in 2b. Applying a simple policy that diverts some of the traffic to other terminals eliminates the heavy congestion as seen in 2c.

Regional Mobility Optimization

Modeling an entire mobility system is analogous to vehicle modeling. At the top integrated level, each part or subsystem within the vehicle is its own agent, perhaps a mesoscopic model of how a subsystem works. Regional mobility models work this way. The time and spatial aggregations are chosen both for computational tractability as well as appropriate scientific abstraction for realistic results. Although vehicle interactions, such as at a traffic signal, can be modeled second to second, such a fine temporal resolution contributes less to the understanding of regional flows than a coarser-resolution model that might be based simply on a volume-to-capacity ratio.

The second-to-second microsimulations might inform behavioral models of vehicles or drivers, or they might help optimize a subsystem of a vehicle algorithm; however, second-to-second simulations add no meaningful granularity to larger regional traffic flow models compared to simpler mesoscopic representations of an intersection. However, a large-scale mesoscopic regional model integrated into a high-resolution model (high spatial and/or temporal resolution) for a district, intersection, autonomous vehicle mode, or other subcomponent provides the opportunity for optimization of the subsystem within the larger context of the overall system. This multi-resolution aspect (i.e., embedded microsimulation within a regional-scale model) is needed with respect to optimization of modes, infrastructure, and electrification (i.e., charging), and to inform design.

2.8 Scalable/Distributed Optimization and Control

Integration with the Electric Grid and Other Infrastructure Systems

The transportation system of the future will not exist in isolation. Instead, it will be closely integrated with behavior and decision-making associated with advanced machines, the grid, communications networks, and buildings. A better understanding of the relationships among mobility, economic productivity, and quality of life is needed and calls for tight integration and quantification of social sciences so that quality of life can be modeled and optimized. Problems confronting the development of “Smart Cities,” the “Smart Grid,” and “Smart Transportation” are intrinsically coupled and must be solved together. Co-simulation of these interrelated, complex systems is required for systems optimization and to fully leverage renewables on the grid. Large-scale co-simulation offers opportunities to develop large, reusable, multi-systems models leveraging high-performance computing capabilities. Yet developing accurate city-scale models exceeds (or at minimum greatly challenges) our current computational capabilities. Additionally, algorithms and modeling frameworks are needed that can leverage the results of large-scale co-simulations of an ecosystem of interacting systems to enable optimal codesign and improvement of that ecosystem moving forward.

Distributed and online optimization, adaptive control, and machine learning are becoming increasingly inadequate in high-dimensional regimes (multiple physical and temporal scales). There is a crucial need to advance the foundational science for real-time monitoring, optimization, and control of large-scale mobility systems governed by complex physical, behavioral, and computational interdependencies at multiple spatiotemporal scales. Cloud and edge computing support distributed optimization and control. Efforts in this area could be anchored on a rigorous mathematical underpinning that offers a common and powerful analytical framework for modeling, computation, communication, optimization, and control. And research is into the foundational mathematics of nonlinear controls, optimization theory, big-data analytics, and complex system theory could help advance several strategically important domains relevant to energy-focused transportation science; these domains share the opportunity to use real-time data to observe, control, and optimize high-dimensional energy systems.

The basic research directions and fundamental scientific underpinnings needed to address such challenges include:

- **Data Analytics:** develop methods to use real-time vehicle and infrastructure data (addressing access and privacy) to enable automated and distributed decision-making from machine-learning techniques.
- **Optimization Theory:** develop computationally affordable and stable algorithms that can be implemented in real time and distributed fashions.
- **Nonlinear Control Theory:** develop scalable, decentralized, and distributed controls that account for inherently asynchronous operations resulting from communications delays, losses, and distributed (asynchronous) control actions.
- **Complex Systems:** develop modeling and simulation methods that enable analysis of the interdependencies of many high-dimensional, nonlinear, interacting systems (e.g., energy, transport, buildings, and communications) at various temporal and spatial scales.

Realizing synergistic integration necessitates breakthroughs in the optimization and control of highly distributed transportation and vehicle systems. Specific near-term opportunities exist to drive the development of this fundamental work and validate the results in relevant real-world applications.

Optimally Managing CAV Fleets

Optimally managing CAV fleets requires a cross-disciplinary approach spanning transportation systems modeling, power systems modeling, building operations modeling, optimization, and data analytics.

Critical research needs include the development of models for CAVs that account for the coupling of electric transportation systems and power networks, which comprehensively capture quality-of-service requirements for travelers, CAV operators, and utility companies. Given the large-scale nature of these models as well as the complexity associated with optimization problems, critical research needs also include the development of computationally efficient and distributed algorithms that can dispatch routes, charging commands, and service commands in real time.

Research is needed to advance core optimization theory and machine learning-based data-analytics tools and to develop adaptive optimization and routing algorithms for CAVs that account for network conditions, passenger loads, states-of-charge, and other aspects in real time. Moreover, to predict both traffic conditions and power-system states, efforts are needed to advance online learning theory by developing algorithms to process high volumes of data coming from different sources (e.g., vehicle sensor data, power system metering units, and traffic system sensors).

The performance and evaluation of pertinent optimization and learning tools requires significant computational power. This is especially true when dealing with realistic scenarios with large vehicle fleets and stochastic models for road conditions, passenger loads, and power network conditions. Solving machine-learning and optimization problems in real time will require the computational power of the national laboratories' supercomputers. Distributed methods can use the distributed computing power of supercomputers to efficiently implement these algorithms and enable real-time (and even faster) simulation and evaluation.

The primary goal of fleet management is to ensure the optimal, resilient, and reliable distributed operation of passenger and freight vehicles in a computationally affordable manner. A variety of techniques for modeling complex systems could be evaluated to determine the best way to capture critical information at multiple spatial and temporal scales. Such modeling could strike a balance between accurately representing system physics and couplings and addressing the need for computationally affordable optimization tasks. Ultimately, a synergistic framework could be developed in which control, optimization, and data analytics were unified under the same distributed mathematical formalism. Achieving this will require a significant integration effort, where the operating principles of vehicles, buildings, and autonomous energy systems merge under a cohesive reference architecture.

Optimal Traffic Control

Current traffic-control methods are derived from certain views or models of traffic systems. Traditional traffic-systems models have certain assumptions (e.g., homogeneous traffic composition and traffic arrivals following uniform distribution or Poisson processes). With better observability enabled by better traffic state sensing infrastructure such as edge computing and higher penetration rate of probe vehicles, traffic systems can be modeled with higher fidelity for control purposes. Also, the CAV technologies can enable (or augment) stationary control through infrastructure (e.g., signalized intersections, ramp metering, and variable speed limits) with nonstationary control through connected vehicles. Therefore, complex optimal control based on more-realistic traffic systems models with more forms of control could be rederived with better performance in the real world.

Traffic systems can be viewed as combinations of several subsystems (e.g., an arterial includes several intersections, and a regional traffic network has several arterials and highways). Control of one subsystem will affect the dynamics of another subsystem. Traditionally, traffic controls have been studied for small areas because the data collection was expensive and time-consuming.¹ Also, solving a large system without breaking it down to smaller subsystems and losing the interactions among the subsystems requires too many computational resources. Conventional optimization methods may be unable to solve the large systems optimal control problems.

Stochasticity is a nature of transportation systems. Uncertainty from people’s activities, weather, and human’s reactions to the environment all contribute to the uncertainty in the number of vehicles on the roads and the movements of these vehicles. Real-time sensing enables us to provide online feedback for optimal traffic control to address the uncertainty in the traffic systems.

Solving complex optimal traffic control problems for larger traffic systems in real time will require advanced computing methods and resources. An ideal optimal control strategy will require heavier efforts on offline parameter seeking and lighter efforts on online optimal control. Distributed methods could be used to solve each subsystem’s optimal control problem with extra constraints to steer the results toward the optimal solution for the bigger system.

2.9 Symbiotic Autonomous Systems

Technologies and platforms are emerging today within a new field of science referred to as symbiotic autonomous systems (SAS), enabled by advances in sensors, ubiquitous communications, distributed computing, and deep-learning algorithms. These smart, context-aware SAS tools are expected to increasingly augment human decision-making and physical capabilities, to the point where “these tools will become a seamless extension of our body and mind.”² This transformation has already begun—for example, people routinely use smartphones to choose the fastest route home. As with any new technology, SAS brings potential benefits as

¹ See Section 7.3 (<https://ops.fhwa.dot.gov/publications/fhwahop08024/chapter7.htm#7.3>) of the “Traffic Signal Timing Manual,” U.S. Department of Transportation Office of Operations, <https://ops.fhwa.dot.gov/publications/fhwahop08024/index.htm>.

² Roberto Saracco, Raj Madhavan, S. Mason Dambrot, Derrick de Kerchove, and Tom Coughlin, *Symbiotic Autonomous Systems: An FDC Initiative*, IEEE White Paper, (IEEE, November 2017). <https://digitalreality.ieee.org/images/files/pdf/sas-white-paper-final-nov12-2017.pdf>.

well as potential risks, and it gives rise to many questions about how such systems will impact mobility and energy use. To strategically advance this emerging scientific field, new tools and methods must be developed to:

- Intimately integrate oft-siloed technical arenas such as artificial intelligence, robotics, power systems engineering, spatial sensing and intelligence, building science, and automated electric vehicles
- Address controls and optimization across disparate time and length scales while accommodating stochastic inputs into decision-making and control.
- Interweave deep expertise in the behavioral and social-ecological sciences to capture the human side of these coevolving systems.

3 Summary of Research Needs

This report outlines several research opportunities and needs in the areas of advanced math/computing, data science, and artificial intelligence that must be met to accelerate progress towards a vision of energy-efficient and electricity-symbiotic transportation system. These needs include:

1. Utilizing large, detailed telemetry data sets for activity-based demand modeling
 - Unified mathematical frameworks that can be applied to a wider range of “shared and multimodal” scenarios and addressing a modeling boundary, including spatial and temporal extent.
 - Computational infrastructure and expertise to enable the agile processing of very large data sets, including the application of different computational and mathematical frameworks.
2. Streaming real-time, high-fidelity data for optimal route planning
 - Computationally efficient and distributed algorithms that can dispatch routes, charging commands, and service commands in real time.
 - Innovative online learning methods are required to process high volumes of data in real time, cope with missing as well as anomalous measurements, and reliably predict both traffic conditions and power-system states.
3. Computer vision for mobility data collection
 - Customizing existing computer-vision techniques for mobility applications, especially for hyperparameter tuning, which is the most critical step in model training and demands huge computational resources.
 - Fusion of multimodal data, dynamic predictions, and accelerating methods development through the use of simulation-generated data.
4. Data science techniques for modeling, fusion, and prediction
 - Approaches to fuse traditional data sources with emerging data streams that automate and dramatically accelerate the development of high-fidelity models.
 - Enhanced computational treatment of agents to account for behaviors associated with demographic and socioeconomic attributes such as gender, income, disability, and education.
5. Artificial intelligence techniques for automated, real-time decision making
 - AI and machine-learning techniques to utilize real-time sensor data with edge-computing devices.
 - Machine learning-based fault detection and prognostic algorithms to serve as an alternative for sensor anomaly detection to obtain quality traffic data from sensors.
6. Simulation and modeling to provide unprecedented insight into potential futures
 - Transportation network models integrated with basic urban science able to model completely new modes and identify subtle interactions between mobility systems.

- Tools for cities to enable rapid impact assessments of the impacts of new mobility options and business models.
 - Tools that model the mobility of passengers and freight in real time, for representative rural and metropolitan areas in the United States, over current and potential future road and infrastructure networks.
7. Simulation techniques to take full advantage of artificial intelligence
- Integration of high-resolution subcomponents (district, intersection, etc.) into a large-scale mesoscopic regional model that allows optimization of the subsystem within the larger context of the overall system.
 - Deep reinforcement-learning algorithms extended to real-world stochastic environments that have multiple decision-making agents, complex dynamics, and limited observability.
 - Reinforcement learning combined with other machine-learning methods to solve control and planning related optimization problems that are combinatorial in nature, beyond the use of traditional software and algorithms, to enable solutions in real time and/or at large scales.
8. Scalable/distributed optimization and control
- Algorithms and modeling frameworks that can leverage the results of large-scale co-simulations of an ecosystem of interacting systems to enable optimal codesign.
 - Improved approaches for real-time monitoring, optimization, and control of large-scale mobility systems, governed by complex physical, behavioral, and computational interdependencies at multiple spatiotemporal scales.
 - Improved understanding of the foundational mathematics of nonlinear controls, optimization theory, big-data analytics, and complex system theory.
 - Scalable, decentralized, and distributed controls that account for inherently asynchronous operations resulting from communications delays, losses, and distributed (asynchronous) control actions.
 - CAV models that account for the coupling of electric transportation systems and power networks, and comprehensively capture quality-of-service requirements for travelers, CAV operators, and utility companies.
 - Computationally efficient and distributed algorithms that can dispatch routes, charging commands, and service commands in real time.
 - Adaptive optimization and routing algorithms for CAVs that account for network conditions, passenger loads, states-of-charge, and other factors in real time.
9. Symbiotic autonomous systems
- New tools and methods that tightly couple artificial intelligence, robotics, power-systems engineering, spatial sensing, and intelligence, building science, and automated electric vehicles.
 - Controls and optimization schemes across disparate time and length scales that accommodate stochastic inputs into decision making and control.

References

- Cecco, Leyland. 2019. “The Innisfil Experiment: The Town that Replaced Public Transit with Uber,” *The Guardian* (July 16, 2019). <https://www.theguardian.com/cities/2019/jul/16/the-innisfil-experiment-the-town-that-replaced-public-transit-with-uber>.
- Czarnecki, K., 2018. Operational Design Domain for Automated Driving Systems. <https://doi.org/10.13140/RG.2.2.18037.88803>
- Lunacek, Monte, Lindy Williams, Joseph Severino, Karen Ficenec, Juliette Ugirumurera, Matthew Eash, Yanbo Ge, and Caleb. Phillips. 2021. “A Data-Driven Operational Model for Traffic at the Dallas Fort Worth International Airport,” *Journal of Air Transport Management* 94 102061, <https://doi.org/10.1016/j.jairtraman.2021.102061>.
- Ruktanonchai, Nick Warren, Corrine Warren Ruktanonchai, Jessica Rhona Floyd, and Andrew J. Tatem. 2018. “Using Google Location History Data to Quantify Fine-Scale Human Mobility.” *International Journal of Health Geographics* 17(1): 1–13. <https://doi.org/10.1186/s12942-018-0150-z>.
- Sadeghvaziri, Eazaz, Mario B. Rojas IV, and Xia Jin. 2016. “Exploring the Potential of Mobile Phone Data in Travel Pattern Analysis.” *Transportation Research Record* 2594(1): 27–34. <https://doi.org/10.3141/2594-04>.
- Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, et al. 2016. “Mastering the Game of Go with Deep Neural Networks and Tree Search,” *Nature* 529: 484–489. <https://doi.org/10.1038/nature16961>.
- Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, et al. 2017. “Mastering the Game of Go Without Human Knowledge,” *Nature* 550: 354–359. <https://doi.org/10.1038/nature24270>.
- Vij, Akshay, and Kalyanaraman Shankari. 2015. “When is Big Data Big Enough? Implications of Using GPS-Based Surveys for Travel Demand Analysis.” *Transportation Research Part C: Emerging Technologies* 56: 446–462. <https://doi.org/10.1016/j.trc.2015.04.025>.
- Zong, Fang, Yongda Tian, Yanan He, Jinjun Tang, and Jianyu Lv. 2019. “Trip Destination Prediction Based on Multi-Day GPS Data.” *Physica A: Statistical Mechanics and its Applications* 515: 258–269. <https://doi.org/10.1016/j.physa.2018.09.090>.

Table 1. Summary of Research Needs to Accelerate Progress Toward a Vision of an Energy-Efficient and Electricity-Symbiotic Transportation System.

Research Need	Examples of Research
Use of large, detailed telemetry data sets for activity-based demand modeling	<ul style="list-style-type: none"> • Unified mathematical frameworks that can be applied to a wider range of “shared and multimodal” scenarios and addressing a modeling boundary, including spatial and temporal extent • Computational infrastructure and expertise to enable the agile processing of very large data sets, including the application of different computational and mathematical frameworks
Streaming of real-time, high-fidelity data for optimal route planning	<ul style="list-style-type: none"> • Computationally efficient and distributed algorithms that can dispatch routes, charging commands, and service commands in real time • Innovative online learning methods to process high volumes of data in real time, deal with missing measurements, and reliably predict both traffic conditions and power system states
Computer vision for mobility data collection	<ul style="list-style-type: none"> • Customized existing computer-vision techniques for mobility applications, especially for hyperparameter tuning, which is the most critical step in model training and demands huge computational resources. • Fusion of multimodal data, 3D dynamic predictions, and acceleration of methods development using simulation-generated data
Data science techniques for modeling, fusion, and prediction	<ul style="list-style-type: none"> • Approaches to fuse traditional data sources with emerging data streams that automate and dramatically accelerate the development of high-fidelity models • Enhanced computational treatment of agents to account for behaviors associated with demographics and socio-economic attributes such as gender, income, disability, and education
Artificial intelligence techniques for automated, real-time decision making	<ul style="list-style-type: none"> • AI and machine-learning techniques to utilize real-time sensor data with edge-computing devices • Machine learning-based fault detection and health prognostic algorithms to serve as an alternative for sensor anomaly detection to obtain quality traffic data from sensors
Simulation and modeling to provide unprecedented insight into potential futures	<ul style="list-style-type: none"> • Integration of transportation network models and basic urban science that can model completely new [transportation/mobility] modes and identify subtle interactions between mobility systems • Tools for cities to enable rapid impact assessments of the impacts of new mobility options and business models • Decision-making tools that model the mobility of passengers and freight in real time, for representative rural and metropolitan areas in the United States, over current and potential future road and infrastructure networks

Research Need	Examples of Research
Simulation techniques to take full advantage of artificial intelligence	<ul style="list-style-type: none"> • Integration of high-resolution subcomponents (e.g., districts and intersections) into a large-scale mesoscopic regional model that allows optimization of the subsystem within the larger context of the overall system. • Deep reinforcement-learning algorithms extended to real-world stochastic environments that have multiple decision-making agents, complex dynamics, and limited observability • Reinforcement learning combined with other machine-learning methods to solve control and planning optimization problems that are combinatorial in nature, and are beyond the use of traditional software and algorithms, to enable solutions in real time and/or at large scales
Scalable/distributed optimization and control	<ul style="list-style-type: none"> • Algorithms and modeling frameworks that can leverage the results of large-scale co-simulations of an ecosystem of interacting systems to enable optimal codesign • Improved approaches for real-time monitoring, optimization, and control of large-scale mobility systems, governed by complex physical, behavioral, and computational interdependencies at multiple spatiotemporal scales • Improved understanding of the foundational mathematics of nonlinear controls, optimization theory, big-data analytics, and complex system theory • Scalable, decentralized, and distributed controls that account for inherently asynchronous operations resulting from communications delays, losses, and distributed (asynchronous) control actions • CAV models that account for the coupling of electric transportation systems and power networks, and comprehensively capture quality-of-service requirements for travelers, CAV operators, and utility companies • Computationally efficient and distributed algorithms that can dispatch routes, charging commands, and service commands in real time • Adaptive optimization and routing algorithms for CAVs that account for network conditions, passenger loads, states-of-charge, and other factors in real time
Symbiotic autonomous systems	<ul style="list-style-type: none"> • New tools and methods that intimately integrate artificial intelligence, robotics, power-systems engineering, spatial sensing and intelligence, building science, and automated electric vehicles • Controls and optimization schemes across disparate time and length scales that accommodate stochastic inputs into decision making and control.