Data-Driven Traffic Reconstruction and Kernel Methods for Identifying Stop-and-Go Congestion

Edgar Ramirez Sanchez*, **Shreyaa Raghavan***, **Cathy Wu**
Massachusetts Institute of Technology
Cambridge, MA 02139
{edgarrs, shreyaar, cathywu}@mit.edu

Abstract

Identifying stop-and-go events (SAGs) in traffic flow presents an important avenue for advancing data-driven research for climate change mitigation and sustainability, owing to their substantial impact on carbon emissions, travel time, fuel consumption, and roadway safety. In fact, SAGs are estimated to account for 33-50% of highway driving externalities. However, insufficient attention has been paid to precisely quantifying where, when, and how much these SAGs take place—necessary for downstream decision making, such as intervention design and policy analysis. A key challenge is that the data available to researchers and governments are typically sparse and aggregated to a granularity that obscures SAGs. To overcome such data limitations, this study thus explores the use of traffic reconstruction techniques for SAG identification. In particular, we introduce a kernel-based method for identifying spatio-temporal features in traffic and leverage bootstrapping to quantify the uncertainty of the reconstruction process. Experimental results on California highway data demonstrate the promise of the method for capturing SAGs. This work contributes to a foundation for data-driven decision making to advance sustainability of traffic systems.

1 Introduction

Transportation is the most significant contributor to greenhouse gas (GHG) emissions, accounting for 28 percent of the total in the U.S.. Traffic congestion on its own is estimated to cost the average driver around 2000 USD in large cities and, to the United States alone, 305 billion USD in 2017 (1). In particular, "stop-and-go"(SAG) traffic, a common phenomenon in congested traffic, happens when cars traveling on a highway periodically form waves that propagate down the highway, leading cars to constantly accelerate and brake unnecessarily. These SAG waves or events lead to up to 67 percent higher fuel consumption, longer travel times, higher GHG emissions levels, and could be a safety hazard (2). Studies have found that SAGs can be responsible for 33-50% of highway driving externalities. (3; 4).

However, a promising aspect of this phenomenon is its potential avoidance through proactive identification and subsequent mitigation measures, meaning that there is a significant portion of the congestion that is avoidable. Multiple studies explore how these SAG events form, propagate and dissipate (5; 6; 7; 8) , and others have shown potential measures to prevent or dissipate them using AVs (9; 2; 10), reinforcement learning (11), variable speed limits (12), etc. These could be empowered by ML approaches that are able to identify, predict, and extract features like congestion, safety-related events, SAGs, and others.

*These authors contributed equally to this work.

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2023.

Because of its large impact on the climate, the intersection of transportation, ML, and sustainability research is a promising line of research. However, these models require rich datasets to capture the nuanced dynamics of vehicles, yet there is very limited traffic data available. In the US, the largest highway dataset is the Freeway Performance Measurement System (PeMS), but even for sections of the highway with a dense distribution of sensors, the average distance between sensors is on the order of thousands of feet away, which is not enough for feature extraction and behavior-based analysis.

In response, this paper proposes the use of traffic reconstruction as a way to fill in the gaps and generate rich trajectory information to enable complex ML techniques. We apply this framework to a county in California and devise a kernel-based SAG identification method with bootstrapping to provide robust insights on traffic behavior.

2 Related Work

The study of stop-and-go events (SAGs) in traffic has a longstanding history, dating back to the 1960s (**?**). While much of the literature has traditionally focused on fundamental analyses, there has been a discernible shift toward identification methods (**?**), in part driven by challenges in obtaining sufficient data for large-scale identification. On one hand, theoretical studies encompass dynamic wave modeling, control mechanisms, instabilities, driving models, and other aspects that shed light on triggers, formation, propagation, and dissipation of these waves (**?** 6**? ? ? ? ? ? ? ?** ; 5; 7; 8).

On the other hand, empirical studies focus on identification through various techniques, such as time series and signal processing methods, to identify SAG occurrences in space and time from the observed traffic or vehicle-specific data (**? ? ? ? ? ? ? ?**). These approaches typically rely on rich trajectory data that only exists for very few stretches of highway, and only some of them have sufficient resolution to incorporate spatiotemporal features. In contrast, most studies use data from fixed and sparse sensors and identify SAGs using stationary wave processing techniques. In particular, wavelet transformation has been a popular approach for SAG identification (**? ?**), matching a Mexican hat profile to the signal to identify where this behavior occurs.

Finally, while research links stop-and-go waves to negative effects such as increased fuel emissions, safety risks, and driver distress (**? ? ?** 7), SAG quantification has not been conducted at scale. A unique example in quantifying the number of SAGs is a study estimating that 20% of traffic jams in the Netherlands could be stop-and-go waves (**?**). The significance of this estimation is underscored by its inclusion as a priority in the national plan of Intelligent Transportation Technologies (ITS) (**?**). Yet, this is an anomaly, as most countries, including the United States, do not treat stop-and-go mitigation as a priority due to a lack of quantitative evidence.

Unlike previous literature, our work makes use of both spatiotemporal features in traffic data and broadly available data, which is sparse in nature.

3 Traffic Feature Identification Framework

We propose a technique to augment raw traffic data with physics-based traffic simulations to build a richer, more continuous data representation, with the goal of enabling ML pipelines for stop-and-go identification. Our raw data source is the Caltrans Performance Measurement System (PeMS) dataset,

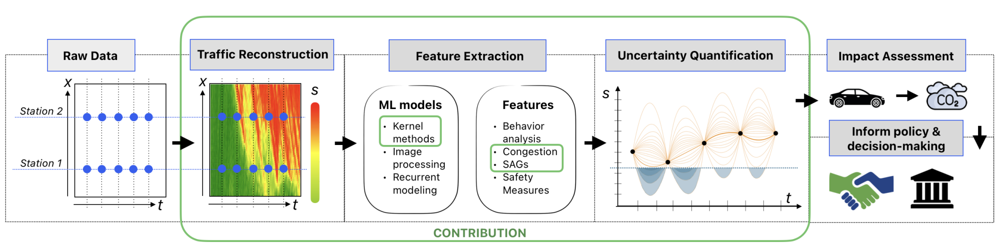

Figure 1: Pipeline for traffic feature identification. Our contribution is boxed in green.

which consists of roughly 39,000 induction loop detectors that capture traffic flow, occupancy, and speed aggregated every 30 seconds on most freeways in major metropolitan areas of California. The key concern with PeMS is that the dataset is too sparse to draw conclusions about traffic events or have a complete understanding of vehicle trajectories. As a result, we seek to fill in these data gaps via traffic reconstruction while using PeMS as our ground truth.

3.1 Traffic Reconstruction

We leverage a traffic reconstruction model built as part of a safety analysis framework, which simulates traffic scenarios at a fine-grain level to build individual vehicle trajectories. It does so by combining micro-level dynamics of driving behavior on California highways, captured by highway video data, with macro-level data obtained by induction loop detectors.

At the large scale, we represent the entire road network as a graph $G = (E, V)$. The goal is to recover the optimal flow through every edge e at time t such that the routes of the vehicles align closely to the ground-truth traffic data while still being realistic. We model this flow estimation problem as a maximum flow problem. At the smaller scale, we create a driver model that relies on the widely used "Intelligent Driver Model" (IDM) (13). In our traffic reconstruction framework, the driver model is calibrated to digital video footage of California highways to better capture behaviors commonly seen in that region (14). Both the driver model and the estimated vehicle routes are then combined in Simulation of Urban MObility (SUMO), a traffic simulator, to extrapolate detailed trajectories of traffic. The outcome is a rich representation of an entire highway network, that has information about every vehicle at every point in space and time. A more comprehensive overview of this traffic reconstruction method is currently in preparation.

3.2 Uncertainty Quantification

With a continuous, time-space representation of traffic, we can now extract salient features. However, each reconstruction of traffic may have deviations, so relying on a single reconstruction for predictions could inflate or erase certain behaviors. This prompts us to question: given a single simulation, how certain are we in an observed event or behavior? Instead, we want a robust measure of confidence in our identification of features. To accomplish this, we use a technique similar to bootstrapping and leverage our traffic reconstruction simulator as a distribution sampler, which allows us to estimate the distribution of speeds for a point in time and space (15). We replicate the traffic reconstruction k times and build a distribution of average vehicle speed at position x and time t: $p_s(s|t, x)$. In each replication, we randomly sample (with replacement) stochastic elements of the reconstruction, such as speed deviations, offset in departure times, and vehicle driving dynamics.

Then, we define an event or feature of interest, A, such that $A(t, x) = 1$ if A occurs at (t, x), and $A(t, x) = 0$ otherwise. A is a function of speeds, $S_{t,x}$, such that $A(t, x) = f(S_{t,x})$. $S_{t,x} = \{s_{t+i,x+j} : \forall i \in [-m, m], j \in [-n, n]\}$ is a neighborhood of speeds surrounding $s_{t,x}$, where m and n are parameters that control how wide this neighborhood is. To provide a measure of certainty, we want to know the probability of A occurring at (t, x), which is $p_A(A|t, x)$. Since $A_{t,x} = f(S_{t,x})$, $p_A(A|t, x) = p_s(f(S_{t,x}, ..., \nu)|t, x)$. Hyperparameters for f are denoted as ν. This function f and collection $S_{t,x}$ differ based on the setting and behaviors of interest.

4 Application to Stop-and-Go Event Identification

Using the individual trajectories from the traffic reconstruction, we build a rich time-space representation of traffic that encodes meaningful features. This representation, known as a time-space diagram, is commonly used in the transportation community to visualize vehicle behaviors as they move through time (x-axis) and space (y-axis). The color of the diagram encodes the aggregated speed of the vehicles across all lanes. In our study, we reconstruct 100 variations of traffic for the entire Los Angeles County from 11:15 to 11:30 a.m.. We focus our analysis below on a 0.5 km segment of the 110-N freeway.

We characterize a stop-and-go event as high $\rightarrow$ low $\rightarrow$ high speeds. In a time-space representation, they appear as waves with a negative slope, as seen in Figure 2(b-e). We capture both the dip in speed over time and the backward movement over space by designing a 2-D kernel based on the Sobel operator for diagonal edge detection (16). To account for different lengths of SAGs, we parameterize

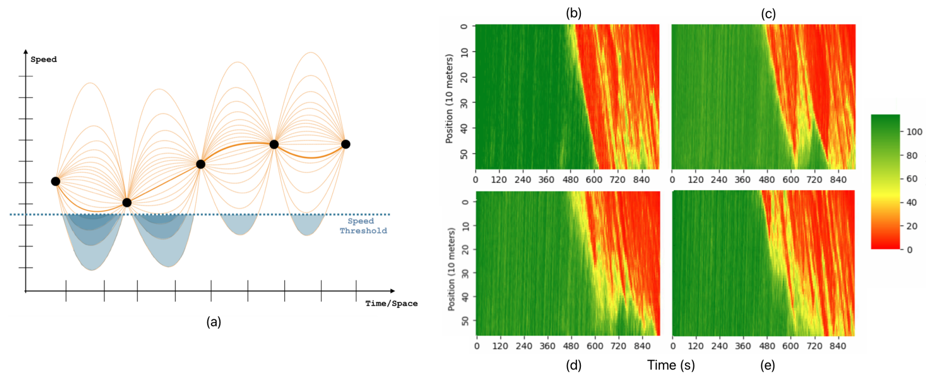

Figure 2: (a) Schematic of possible reconstructed trajectories that fill the gaps between ground-truth readings. These trajectories vary in each replication. (b-e) 4 varying time-space diagrams out of 100.

the width of the kernel. An example of the kernel K with width $= 4$ is as follows:

$$K(4) = \begin{bmatrix} 0 & -1 & -1 & -1 & -1 & 0 & 2 & 2 \\ 2 & 0 & -1 & -1 & -1 & -1 & 0 & 2 \\ 2 & 2 & 0 & -1 & -1 & -1 & -1 & 0 \end{bmatrix}$$

We convolve this kernel $K(w)$ across neighborhoods of speeds, $S_{t,x}$, in the time-space diagram and normalize these values from -1 to 1. We define this convolution as the function $f_{SAG}(S_{t,x}, w|t, x) = norm_{(-1,1)}(K(w) * S_{t,x}(w, 1))$, which we also refer to as the "Kernel Activation Value". The Kernel Activation Value provides a measure of how strongly the behavior at (t, x) matches the kernel.

Then, we define $C_{t,x}$ as the event that a SAG occurs at (t, x), which is 1 when $f_{SAG}(S_{t,x}, w|t, x) \geq \epsilon$ and 0 otherwise. This ϵ controls how steep the stop-and-go behavior should be in order to be considered a SAG event. For our study, we use a threshold of $\epsilon = 0.30$. We build an indicator variable $C^j_{t,x}$ defined by the occurrence of a SAG in trial i at (t, x) Then, we estimate the probability of a SAG occurring at (t, x) by using the $C^j_{t,x}$ as samples of the probability distribution:

$$p(C_{t,x} = 1|t, x) = p_s(f_{SAG}(S_{t,x}, w|t, x) \geq \epsilon|t, x) = \frac{1}{100} \sum_{i=1}^{100} C^i_{t,x}(w, \epsilon)$$

Using the continuous time-space representation and our kernel-based method, we are able to accurately detect and visualize stop-and-go traffic in time and space, as shown by the blue diagonal lines in the first row of Figure 3. The probabilities in the second row of Figure 3 create hotspots that we know are more prone to SAGs. As a result, we can quantify how certain we are that a SAG detected in a single trial is not simply variance or noise. This probabilistic analysis allows us to extrapolate the observed traffic patterns in data to the real world and more confidently use the results of this framework for downstream applications.

5 Conclusion

This paper presents a novel traffic reconstruction framework, offering a rich spatio-temporal representation that effectively and robustly detects traffic behaviors, particularly stop-and-go events. Our representation enables the creation of comprehensive traffic datasets, while the success of our kernel method in capturing SAG patterns demonstrates the potential for more complex ML methods to leverage this data. Through SAG identification at scale, we can quantify how much of the congestion is avoidable, and even pinpoint when and where it occurs. This informs ML pipelines, smart infrastructure, and decision-makers, among others, and allows them to intervene on high-emission events, such as SAGs. Future research directions include the development of sequential ML tools for time-series analysis, extending to other features such as safety measures, causal analysis via individual trajectories, behavior analysis, and prediction through supervised learning. Additionally, scaling the SAG identification process for the full California highway system and for longer periods of time is a promising step towards at-scale SAG intervention.

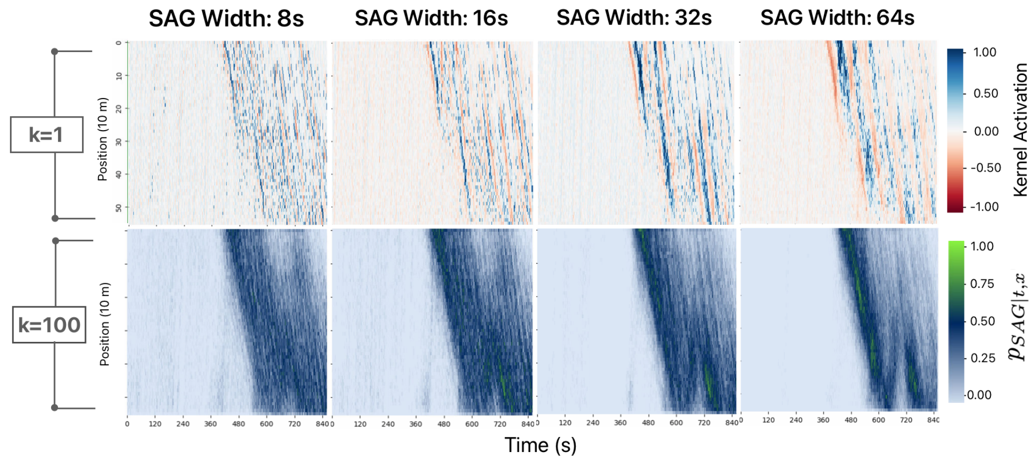

Figure 3: f_{SAG} for kernels of different widths applied to a single trial (same trial as Figure 2(b)) are shown in the first row. Values classified as SAG ($> \epsilon$) are boosted for visual reference. The second row shows the probability (combined across all 100 simulations) of a SAG occurring.

References

[1] E. P. Agency, *Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2021*. Jul 2023.

[2] R. E. Stern, S. Cui, M. L. Delle Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, H. Pohlmann, F. Wu, B. Piccoli, *et al.*, "Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments," *Transportation Research Part C: Emerging Technologies*, vol. 89, pp. 205–221, 2018.

[3] K. Goldmann and G. Sieg, "Quantifying the phantom jam externality: the case of an autobahn section in germany," *European transport research review*, vol. 13, no. 1, pp. 1–15, 2021.

[4] K. Goldmann and G. Sieg, "Economic implications of phantom traffic jams: evidence from traffic experiments," *Transportation Letters*, vol. 12, no. 6, pp. 386–390, 2020.

[5] C. Balzotti and E. Iacomini, "Stop-and-go waves: A microscopic and a macroscopic description," in *Mathematical Descriptions of Traffic Flow: Micro, Macro and Kinetic Models*, pp. 63–78, Springer, 2021.

[6] J. A. Laval and L. Leclercq, "A mechanism to describe the formation and propagation of stop-and-go waves in congested freeway traffic," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1928, pp. 4519–4541, 2010.

[7] S. Oh and H. Yeo, "Impact of stop-and-go waves and lane changes on discharge rate in recovery flow," *Transportation Research Part B: Methodological*, vol. 77, pp. 88–102, 2015.

[8] A. Portz and A. Seyfried, "Analyzing stop-and-go waves by experiment and modeling," in *Pedestrian and Evacuation Dynamics*, pp. 577–586, Springer, 2011.

[9] S. Almatrudi, K. Parvate, D. Rothchild, and U. Vijay, "Using automated vehicle (av) technology to smooth traffic flow and reduce greenhouse gas emissions," 2022.

[10] R. E. Stern, S. Cui, M. L. Delle Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, H. Pohlmann, F. Wu, B. Piccoli, *et al.*, "Dissipation of stop-and-go waves via control of autonomous vehicles: Experimental results: Data," 2017.

[11] A. R. Kreidieh, C. Wu, and A. M. Bayen, "Dissipating stop-and-go waves in closed and open networks via deep reinforcement learning," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1475–1480, IEEE, 2018.

[12] M. Wang, W. Daamen, S. P. Hoogendoorn, and B. Van Arem, "Connected variable speed limits control and car-following control with vehicle-infrastructure communication to resolve stop-and-go waves," *Journal of Intelligent Transportation Systems*, vol. 20, no. 6, pp. 559–572, 2016.

[13] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical review E*, vol. 62, no. 2, p. 1805, 2000.

[14] C. Zhang and L. Sun, "Bayesian calibration of the intelligent driver model," *arXiv preprint arXiv:2210.03571*, 2022.

[15] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*. CRC press, 1994.

[16] N. Kanopoulos, N. Vasanthavada, and R. L. Baker, "Design of an image edge detection filter using the sobel operator," *IEEE Journal of solid-state circuits*, vol. 23, no. 2, pp. 358–367, 1988.