Learning a Pareto Frontier for Safety and Performance in Autonomous Mobility

MMI Project Summary (PI: Prof. Cathy Wu)

1 Introduction

In this MMI project, we tackle the challenges associated with large-scale data-driven evaluation of traffic safety in simulation environment. This capability has the potential to help us understand how emerging technologies will affect safety-performance trade-off without mass-scale deployment. The ultimate goal is to find the strategy where safety and performance can both be optimized for the era of autonomous vehicles, but we first must need to grasp a clear understanding of what the trade-off between two metrics is for human driving vehicles.

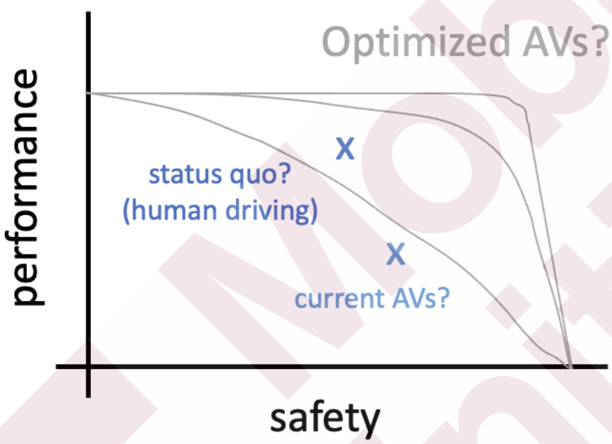

Figure 1: Understanding the status quo of trade-off between safety and performance

Traditional safety evaluation approaches rely on historical crash data, a reactive strategy that suffers from issues of scalability and fails to generalize to new technologies or scenarios. Recent works have started to explore the potential of Surrogate Safety Measures (SSMs) within simulated traffic environments, a method that leverages the statistics of near-crashes to infer safety levels. However, these works often focus on small-scale scenarios and fail to account for factors including temporal effects, weather, and driver behavior.

To enhance the validity of SSMs for extended analysis, we developed a novel framework that integrates micro-level driver behaviors with macro-level traffic states, incorporating external factors such as weather and geographical variations. By utilizing traffic data from California's Performance Measurement System (PeMS) and driving data from Next Generation Simulation (NGSIM), we performed a large-scale traffic simulation accounting for 21,918,675,426 vehicle miles traveled on highways within six counties in August 2018. Then, we conducted a comprehensive set of six variations of regression tests that correlate the SSMs extracted from the simulation with recorded crash statistics.

Our results indicate a significant correlation between SSM counts and the number of actual crashes, validating the predictive potential of these measures. However, we also observed that changing the thresholds for counting SSMs did not show a clear trend in the resulting correlation coefficients, suggesting limitations in the currently available public data for establishing a robust link between simulated SSMs and real-world crashes.

This insight underscores the need for improved data collection and simulation techniques to better represent and accurately understand the safety levels for current human-driving performance. We emphasize that as mobility systems evolve, especially with the advent of new technologies, there is a pressing need for updated safety evaluation methods.

2 Methodology Overview

The overall framework consists of data collection, traffic reconstruction, deriving SSMs and crashes, and analzying with statistical models. An overview of this framework is presented in Figure 2.

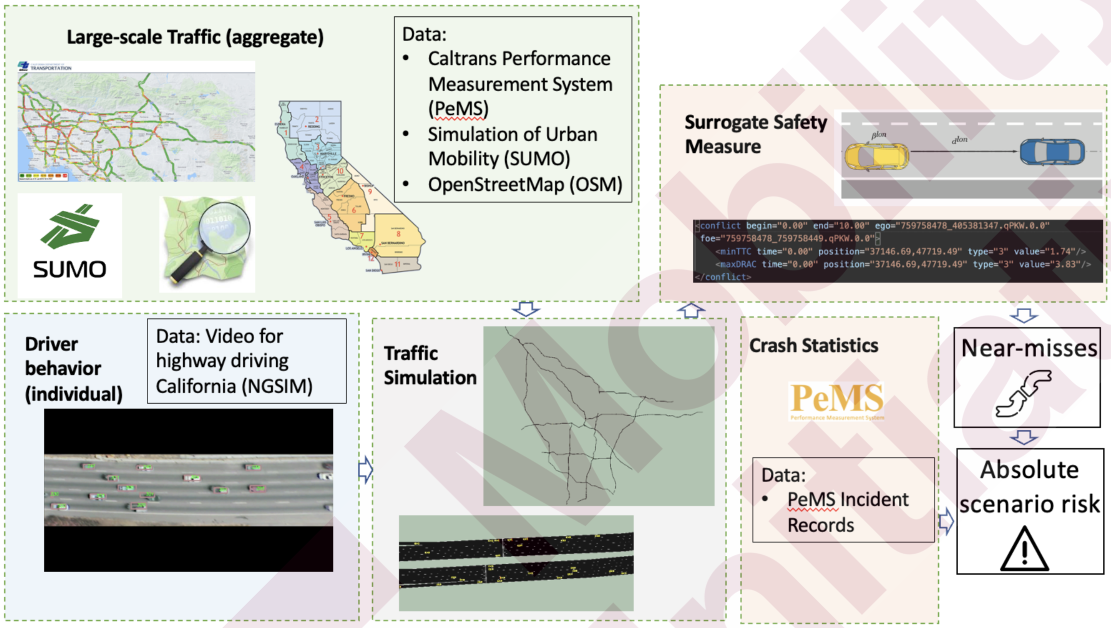

Figure 2: Pipeline for Validating SSM

2.1 Data Collection

To model driver behavior, we first gather data from the widely recognized NGSIM US-101 dataset. This dataset records vehicle trajectories along a section of U.S. Highway 101 (Hollywood Freeway) in Los Angeles, CA, specifically from 7:50 AM to 8:35 AM on June 15, 2005. It captures the movement of 3,233 vehicles at a 10 Hz frequency, with each vehicle observed for an average of 366.5 seconds during this 45-minute period. Although dated, this dataset remains relevant for understanding driving behavior on Californian highways and aligns with the traffic state data used for our modeling.

After extracting parameters for the Intelligent Driver Model (IDM) from the NGSIM dataset, we integrate a month's worth of data from the Caltrans Performance Measurement System (PeMS) collected in August 2018 across six counties: Los Angeles, San Joaquin, San Bernardino, Sacramento, Santa Clara, and Orange. PeMS uses about 40,000 detectors to monitor traffic in real-time every 30 seconds, which we map to the OpenStreetMap road network to simulate traffic accurately. This simulation encompasses a total of 21,918,675,426 vehicle miles traveled, providing a robust data set for statistical analysis.

Additionally, we use traffic incident data from the PeMS database, which includes 16,875 recorded incidents across the mentioned counties during the study period.

As illustrated in Figure 3, the analysis assesses the correlation between SSMs and crash counts. The shaded blocks

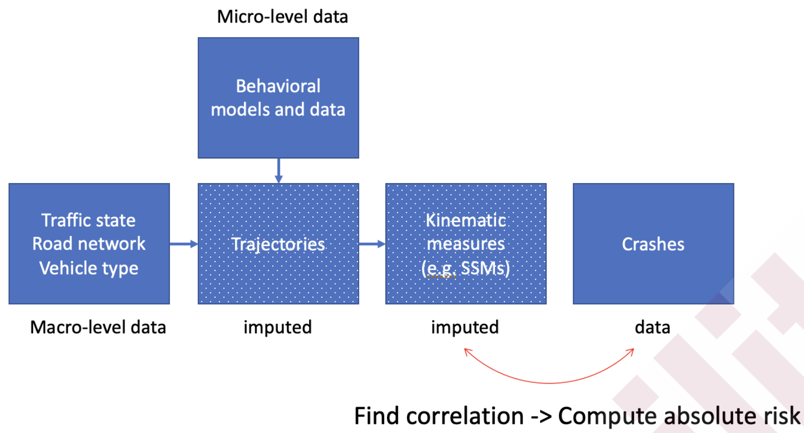

Figure 3: Data Flowchart

represent imputed data, and the other blocks denote raw data, as discussed above.

2.2 Traffic Reconstruction

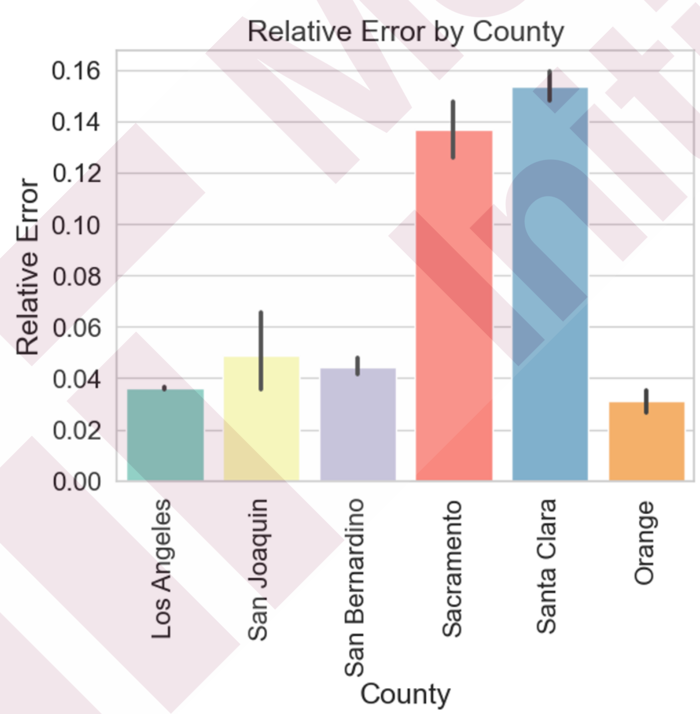

Figure 4: The relative error in reconstructed traffic flow for each county

To create detailed traffic simulations that reflect both individual vehicle driving behavior and real-world traffic conditions, we developed a method that captures both the micro-level behaviors of drivers and the macro-level traffic flows. We first utilized data from PeMS sensors (i.e., induction loop detectors), which monitor various real-time traffic metrics, though they do not specify traffic volume on specific routes. To address this limitation, we developed an imputation system aimed at estimating traffic flow throughout the entire network, ensuring that our simulations faithfully mirror actual traffic behaviors. In addition, we fine-tuned driver models to emulate real-world driving tendencies, employing a

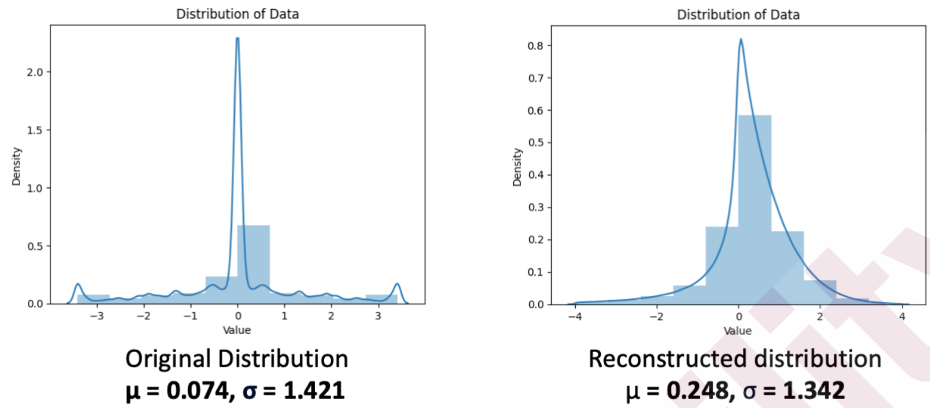

Figure 5: A comparison between distributions for vehicle accelations

Bayesian calibration method to accommodate the inherent variability in driver behavior. By integrating these comprehensive simulations with empirical data, we substantially bolstered the precision of traffic reconstruction. We present the relative error in terms of traffic flow in Figure 4 and a comparison between the distributions for vehicle acceleration in original NGSIM data and reconstructed data in Figure 5.

2.3 Statistical Models

We extract SSMs with varying thresholds from the above simulation and correlate them with corresponding crash data. For analysis, we build a poisson regression model. Our dependent variable in this model is the accumulated number of crashes. On the other hand, SSM aggregates, bracketed within diverse thresholds, serve as the explanatory variables. In order to control for geographical and temporal variations, additional fixed-effect variables are incorporated for each geographic region and temporal interval (categorized as weekdays or weekends). The complete formulation of the regression model is as follows:

$$\log(Y^{it}) = \beta_0 + \beta_1 X^{it}_{a<SSM<b} + \beta_2 q^{it} + \beta_3 W^{it} + \zeta_i + \gamma_t + \epsilon_{it} \tag{1}$$

where

- i denotes the region.
- t denotes the time period.
- Y^{it} denotes the aggregate number of crashes within region i during time period t.
- $X^{it}_{a<SSM<b}$ denotes the aggregate surrogate safety measure recorded within region i during time period t, with pre-defined thresholds a and b.
- q^{it} denotes the traffic flow within region i during the time period t.
- W^{it} denotes a multi-dimensional vector representing the weather information within region i during time period t.
- ζ_i is the geographic fixed-effect term.

- γ_t is the temporal fixed-effect term.
- ϵ_{it} is the noise term, which is assumed to follow a Normal distribution with mean 0 and standard deviation 1, i.e., $\epsilon_{it} \sim N(0,1)$.

Regarding calculating the aggregate SSM in our experiment, we have selected Time-to-collision(TTC) and Deceleration-rate-to-avoid-crash(DRAC) as our primary SSM metrics. To illustrate, if we establish a threshold of 3 seconds for TTC, we can define $X^{it}_{0<SSM<3}$ as the total count of instances in which a vehicle experiences a TTC of less than 3 seconds. As for the DRAC metric, where a higher value indicates a more perilous situation, we define the SSM threshold as a, and use ∞ to represent b.

3 Experiments and Results

As shown in Table 1, our primary model indicates that for the subject SSMs, there is a significant positive correlation with the number of crashes, demonstrating that SSMs can serve as strong indicators of roadway safety. However, although we expected the correlation to decrease with less strict thresholds for SSMs, this was not observed in our numerical results.

Table 1: Results for our proposed model

Time to Collision			DRAC		
Threshold	Coef	P-value	Threshold	Coef	P-value
1	0.000054	1.59E-13	5	6.77E-05	3.45E-10
1.5	0.000033	4.62E-14	4.5	5.75E-05	2.69E-10
2	0.000025	4.25E-14	4	4.907E-05	2.04E-10
2.5	0.000021	5.40E-14	3.5	4.23E-05	2.04E-10
3	0.000017	4.87E-14	3	3.62E-05	2.07E-10

4 Significance and Impact

In this study, we delve into a longstanding oversight in research: the suitability of simulated environments for deriving Safety Surrogate Measures (SSM) and informing roadway safety analyses using publicly available data. With meticulous efforts to utilize diverse data sources and construct models with precision, we uncover a significant correlation between SSM counts and crash numbers. However, altering SSM thresholds does not reveal statistically significant trends. The absence of this analysis in prior studies suggests that existing simulated roadway safety assessments might inadequately equip us to evaluate the safety impacts of emerging mobility technologies accurately. Our investigation exposes several challenging limitations regarding open-access data, notably limited access to detailed driving data, potentially resulting in a disparity between actual and simulated driver behavior. Our reliance on public data from the Performance Measurement System (PEMS), the largest global open dataset for highway driving, reveals a lack of detailed crash data, hindering the identification of crashes directly linked to SSM metrics. Moreover, the absence of fine-grained driving trajectories poses integration challenges. An alternative approach using naturalistic driving data, such as the Second Strategic Highway Research Program (SHRP2), promises more accurate SSMs with precise vehicle trajectories. Although accessing such data is difficult, its limited applicability to macro-level traffic scenarios underscores the importance of developing robust methodologies to advance safety analysis.