Data-Driven Optimization of Automated Speed Enforcement Logistics

Mandana Hedayati Mobarakeh

A Thesis Submitted to the Faculty of Graduate Studies

In Partial Fulfillment of the Requirements

for the Degree of Master of Applied Science

Graduate Program in Civil Engineering

York University

Toronto, Ontario

November 17, 2024

Abstract

Canada's collision fatalities are about 2000 lives a year, decreasing in the last decade, reaching 1745 in 2020 due to initiatives like Vision Zero. Among municipalities' priorities is to enforce speed limits to reduce speeding-induced traffic fatalities, constituting 27 % of all traffic fatalities in Canada. An emerging strategy toward this goal is the deployment of , which detects violators through speed cameras positioned alongside designated roads. Empirical evidence from existing Automated Speed Enforcement (ASE) practices shows that the number of citations drops each month as driver become aware of camera locations and lower their driving speeds. Hence, ASE cameras are often relocated in cycles to expand their reach to more places and further deter speeding violations. The complexities of deployment lie in choosing camera locations and cycle duration, which have the highest deterrence impact on speeding during a planning period. This study proposes a data-driven model to classify camera site locations based on the effectiveness of ASE enforcement. Then, a Markov decision process optimization model is presented to find the optimal camera locations at each cycle and the length of the cycles for minimizing speed violations across the entire transportation network, considering limitations such as the number of available cameras.

Contents

List of Tables

List of Figures

1 Introduction

1.1 Research Motivation

According to the World Health Organization, approximately 1.25 million road traffic fatalities occur yearly, imposing a significant burden on individuals, properties, and the society as a whole [1]. Canadian road traffic collisions lead to 2,000 fatalities annually, a trend that has been declining over the last decade, reaching 1,745 in 2020, partly due to initiatives like Vision Zero. Vision Zero is a comprehensive road safety concept and strategy originating in Sweden in the 1990s and adopted by many cities worldwide. The underlying principle of Vision Zero is that human life and health take precedence over all other objectives of the road traffic system. Hence, it aims to eliminate all traffic fatalities and severe injuries by considering the limits of human behaviour [2]. Vision Zero also prioritizes institutional conditions and approaches that indirectly impact various stakeholders actions to enhance road transport system safety. The long-term objective of Vision Zero is to create a road transport system in which no one is killed or seriously injured due to traffic collisions [3, 4]. Traditional road safety activities have largely been based on behavioural science research, which concludes that a significant portion of road traffic collisions can be attributed to human error. The primary challenge in traditional safety studies is to prevent both conscious and subconscious faulty human actions [5]. Vision Zero, however, accepts that human beings make mistakes and focuses on preventing collisions from leading to death and serious injury. While collisions themselves can be accepted, their severe consequences cannot. Vision Zero emphasized that the primary cause of fatalities and serious injuries is the excessive energy transferred to individuals during a collision, surpassing the limits of what the human body can withstand — a concept based on research by the renowned American road safety expert William Haddon in the 1960s [6].

As a primary cause of traffic collisions, increasing speed increases both their severity and frequency. As a result, cities seek to gauge and enforce speed limits to reduce traffic fatalities related to violating the speed limit, which constitute 27% of all traffic fatalities in Canada [7]. An emerging approach to mitigating the impact of collisions related to violating the speed limit is Automated Speed Enforcement (ASE), which uses a camera and a speed measurement device to detect and capture images of vehicles exceeding the posted speed limit. ASE aims to work with traditional police enforcement to alter driver behaviour, reduce violating the speed limit, and enhance safety. ASE is primarily implemented using fixed cameras that are permanently installed at specific locations along roadways. These cameras continuously monitor vehicle speeds and can operate at designated times or around the clock without human oversight. After capturing an image of a vehicle exceeding the speed limit, the ASE system stores the image and associated data, such as the recorded speed, timestamp, and location, and sends it to a processing center. There, a provincial offenses officer reviews the image, and a ticket — including a digitized copy of the image and an enlarged view of the license plate — is mailed to the vehicle's registered owner.

While ASE relies on fixed cameras, mobile cameras can also support speed enforcement. These mobile cameras are mounted in police vehicles or set up on tripods by the roadside and are frequently relocated, often hourly, to different locations to maintain unpredictability in enforcement coverage. By combining the continuous monitoring capabilities of fixed cameras with the flexibility of mobile units, ASE and traditional methods provide a comprehensive approach to speed management, focusing on reducing collisions and improving road safety [8, 9]. It is noteworthy that ASE refers to implementing fixed cameras that do not need the supervision of a police officer. The ASE program of Ontario, Canada, was implemented following the 2017 Safer School Zones

Act (Bill 65), which amended the Highway Traffic Act to introduce ASE in school zones and community safety zones across the province [7]. In most cities throughout Ontario, , ASE cameras, fixed camera types, are relocated periodically every couple of months. This relocation strategy addresses two primary issues. First, prolonged camera installation at one location may lead to driver familiarity. Over time, drivers learn the locations of the cameras, reducing their speed only temporarily before the cameras and then violating the speed limit again after passing them, thus circumventing the intended enforcement. Second, the high number of community safety zones (CSZs) and road segments within each CSZ, combined with the costs associated with installing and operating each camera, make it impractical for municipalities to cover all segments continuously. To mitigate these issues, cities in Ontario regularly relocate most, if not all, of their ASE cameras. This approach allows them to cover more CSZs with a limited number of cameras and prevents drivers from becoming familiar with specific camera locations. The locations of the cameras are publicly available on each city's website, and cities are also required to post "ASE Coming Soon" signs one month before installing a camera at any new location. Table 1 provides detailed information on the Ontario municipality's ASE programs, including the start date of their ASE programs, the number of available cameras, and whether they rotate their cameras.

However, this relocation strategy presents significant challenges for cities regarding ASE scheduling. Specifically, it is challenging for municipalities to prepare a list of locations and determine the duration of enforcement at each site, ensuring that safety values will increase. This process requires careful selection of new locations and planning cycle times between relocations to maximize the safety benefits. To our knowledge, no specific model exists to optimize safety by optimally determining camera locations and cycle periods.

City	Start Date	Type of Enforcement	Cameras
Toronto	2020/6	Rotational	50
Brampton	2020/7	Rotational	50
Pickering	2021/8	Rotational	13
Durham region	2020/9	Fix - Rotational	4-4
Town of Ajax	2022/1	Rotational	3
Waterloo	2021/9	Fix	16
York Region	2020/11	Rotational	1
Hamilton	2021/10	Rotational	2
London	2021/9	Rotational	2
Mississauga	2021/6	Rotational	17
Ottawa	2020/7	Fixed	32
Peel Region	2020/9	Rotational	1

Table 1: Ontario ASE program summary such as ASE start date, type of enforcement(fixed at one locations, or rotate cameras every couple of months), and the number of cameras for each city

1.2 Research Objectivies

Building on the principles of Vision Zero and the importance of speed management strategies highlighted in the research motivation, this study aims to address the challenges associated with Automated Speed Enforcement (ASE) scheduling and optimization. While ASE has proven to be an effective tool for reducing the violations of speed limit and improving road safety, its implementation in Ontario, particularly the relocation strategy, presents significant challenges for municipalities. These challenges include selecting the most effective locations, determining appropriate enforcement durations, and ensuring that ASE deployment maximizes safety benefits across diverse Community Safety Zones (CSZs).

Currently,municipalities have prioritized safety zones and school zones based on collision history, vehicle speed data, 24-hour traffic volume, the percentage of students within walking distance, and requests from police and the public to minimize the risk of injuries and fatalities. However, this approach often fails to account for the varying effectiveness of ASE at different sites and does not provide a framework to predict outcomes at untested locations. Addressing these gaps requires a systematic and data-driven approach to optimize ASE deployment, considering both site-specific characteristics and broader scheduling strategies.

The following research objectives are designed to develop a comprehensive framework for improving ASE scheduling and placement:

1. **Analyze Relationships Between Physical and Traffic Characteristics and ASE Effectiveness**

   The first objective focuses on analyzing how the physical and traffic characteristics of road segments influence ASE effectiveness. Factors such as road design, traffic volume, speed limits, and proximity to schools will be examined to identify patterns that affect ASE performance. A predictive machine learning model will be developed to forecast ASE effectiveness at new or untested sites. This model will enable municipalities to estimate the potential impact of ASE based on site-specific characteristics, even without prior enforcement data, supporting more informed and targeted decision-making.

2. **Develop a Model for Optimized ASE Scheduling**

   The second objective aims to create an optimization model for ASE scheduling and placement. Using a Markov Decision Process (MDP), this model will identify the most effective locations for camera placement and determine the duration of enforcement at each site. The model will also account for constraints such as limited camera availability and the need to maintain unpredictability in enforcement patterns. This evidence-based approach will help municipalities strategically allocate resources, ensuring that ASE deployment maximizes safety outcomes while remaining efficient and adaptable.

3. **Develop a Practical, Data-Driven Tool for ASE Scheduling**

   The final objective is to integrate the insights from the previous objectives into a practical, user-friendly tool. This tool will allow municipalities to prepare optimized ASE schedules quickly and efficiently, tailored to their specific road and traffic conditions. Designed for scalability and adaptability, the tool will provide a comprehensive solution for municipalities aiming to improve their ASE programs. To demonstrate its practicality, the tool will be applied to the City of Vaughan as a case study, showcasing its effectiveness in real-world scenarios.

By addressing these objectives, this research seeks to provide a robust framework for ASE scheduling and placement, ensuring that enforcement strategies are both effective and efficient in reducing collisions related to violating the speed limit and enhancing road safety.

1.3 Research Structure

In this thesis, a thorough literature review has been conducted to examine how the effectiveness of ASE has been evaluated in various studies, focusing on the metrics used for these evaluations. The review also identifies the best practices in ASE implementation in North America and outlines the current ASE program regulations and practices in Ontario. In the data analysis chapter, the ASE ticketing data from the City of Toronto is analyzed to evaluate the effectiveness of the ASE program at different site locations. This analysis examines the correlation between site-specific physical traffic properties and ASE effectiveness. Subsequently, a machine learning prediction model is developed to forecast changes in the pattern of speed violations during enforcement based on the characteristics of the road sites. The optimization chapter introduces a Markov decision process model designed to optimize ASE scheduling. This model determines the optimal location and duration of enforcement for each ASE camera, with the aim of minimizing the number of speed violations throughout the network during the planning period by strategically distributing ASE cameras in community safety zones. Finally, the case study chapter illustrates the development of a data-driven tool in Mapbox. This tool integrates all data analysis and optimization processes into a single platform, providing a comprehensive solution for ASE implementation scenarios in the City of Vaughan. The summary of the six main steps throughout the research is as follows.

- **Step 1: Define a Safety Metric for Evaluating ASE Effectiveness**

  The first step in this study is to define a safety metric that accurately measures the effectiveness of Automated Speed Enforcement (ASE) programs. This metric will be established through a comprehensive literature review of existing road safety metrics and their applicability to ASE effectiveness. The chosen safety metric will provide a quantifiable measure of the change in safety performance during periods of ASE enforcement. By developing a reliable metric, this study aims to create a standard for evaluating the impact of ASE on overall road safety and compliance, forming the basis for further analysis and optimization.

- **Step 2: Develop a Framework to Assess ASE Effectiveness at Each Location**

  Building upon the defined safety metric, the next step involves creating a framework to assess the effectiveness of ASE at each site where it has been implemented. This framework will evaluate changes in the safety metric during ASE implementation, enabling a systematic assessment of ASE effectiveness across different locations. It will provide insights into which sites benefit the most from enforcement and which might require alternative strategies to enhance their safety impact.

- **Step 3: Identify Relationships Between Physical and Traffic Characteristics and ASE Effectiveness**

  With the framework for evaluating ASE effectiveness in place, this step will analyze the relationship between the physical and traffic characteristics of ASE sites and their measured effectiveness. This analysis will investigate factors such as road design, traffic volume, speed limits, and proximity to schools to understand how these variables influence ASE performance.

- **Step 4: Predict ASE Effectiveness for Future Implementations**

  The study will then develop a predictive machine learning model to forecast the effectiveness of ASE at new or untested sites. This model will allow cities to estimate the likely impact of ASE implementation at any location based on its unique physical and traffic characteristics. By providing a data-driven prediction of ASE effectiveness, this step will support better decision-making and resource allocation for future ASE deployments, ensuring that enforcement efforts are focused on areas where they are most likely to yield significant safety benefits.

- **Step 5: Develop a Model for Optimized ASE Scheduling**

  This step involves developing a comprehensive model to optimize the scheduling and placement of ASE cameras across a city's road network. Based on a Markov Decision Process (MDP), this model will determine the most effective locations for camera placement and the duration of enforcement at each site, considering various constraints, such as the availability of cameras and the need to maintain unpredictability in enforcement. This approach will help cities make evidence-based decisions on where and when to deploy ASE for maximum impact.

- **Step 6: Develop a Practical, Data-Driven Tool for ASE Scheduling**

  The final step of the study is to create a user-friendly, data-driven tool that cities can use to prepare ASE schedules quickly and efficiently. This tool will integrate all the findings from the previous steps, allowing city planners and traffic safety officials to generate optimized ASE schedules in seconds based on their city's specific road and traffic data. Designed to be adaptable to any urban environment, the tool will provide a scalable solution for cities seeking to improve their speed enforcement strategies and enhance road safety outcomes. As a case study, the tool will be applied to the City of Vaughan to demonstrate its practical application and effectiveness.

2 Literature Review

2.1 Collisions related to violating the speed limit

Empirical evidence from traffic collisions caused by violating the speed limit shows that higher vehicle speeds are strongly correlated with an increased probability and severity of accidents. While the relationship between traffic speed, accident frequency, and severity remains a topic of debate, higher vehicle speeds undeniably result in longer stopping distances, reducing the likelihood of stopping safely, particularly affecting less-protected road users such as pedestrians and cyclists. For instance, the UK Highway Code specifies typical stopping distances of 23 meters at 30 mph and 36 meters at 40 mph. Furthermore, the energy dissipated in an accident involving a car travelling at 40 mph is 78% higher than at 30 mph [10].

Numerous studies have sought to establish a logical relationship between speed and safety [11],[12], [13], [14]. In the 1960s, many studies concluded that traffic safety is highly related to speed variation, proposing a U-shaped relationship between accident rate and speed variation, concluding that the more the speed deviates from the traffic mean speed, the higher the likelihood of a collision. Lave et al. [11] reinforced this finding by estimating regression models that tested the relationship between the fatality rate, average speed, and the difference between the 85th percentile and average speed using cross-sectional data from 1981 and 1982. Lave concluded that there was no statistical evidence proving that average speed affects the fatality rate, advocating that traffic engineers and planners should focus on speed coordination rather than speed limitation.

Snyder [15] re-estimated Lave's model using a fixed-effect linear model with two measures for speed variability: the difference between the 85th percentile and median speed (for faster drivers) and the difference between the 15th percentile and median speed (for slower drivers). They concluded that both average speed and speed dispersion are important factors in highway fatalities, but the speed dispersion for faster drivers is particularly related to the fatality rate. Garber and Gadiraju [16] and Garber and Ehrhart [17] also found that accident rates increase with rising speed variance rather than with average speed. Levy and Asch [18] similarly concluded that a lack of speed coordination implies a more significant risk, with average speed also contributing to increased fatality risk depending on speed variance. Examining these studies reveals that focusing on speed coordination, managing speed variance, and implementing automated speed enforcement can significantly enhance road safety, reducing the frequency and severity of traffic collisions.

2.2 Speed Camera Effectiveness

The effectiveness of speed camera enforcement has been extensively evaluated using a variety of methodologies, metrics, and contexts to provide a comprehensive understanding of their impact on road safety.

Related studies leverage before-and-after analyses is involve comparing traffic conditions before and after the implementation of ASE to assess its impact. Studies, such as those conducted in Montgomery County, Maryland [19], and Scottsdale Loop 101 in Arizona [20], demonstrated significant reductions in violating the speed limit and collisions. These studies reported reductions of up to 70% in violating the speed limit and notable decreases in average speeds.

Empirical Bayes (E.B.) observational studies, like the evaluation on the Italian Motorway A1 [21], highlighted substantial total collision reductions of 31.2% , emphasizing ASE's effectiveness in mitigating high-severity collisions. The E.B. method combines historical collision data with observed data to provide a more accurate estimate of the treatment effect by accounting for regression-to-the-mean bias. This method is particularly useful in traffic safety studies because it helps correct the natural fluctuation in accident data over time, thus providing a more reliable measure of the impact of ASE.

Interrupted time series analyses, exemplified by the study in Seattle's school zones [22], showed a 53% reduction in speed violation rates and significant decreases in maximum and mean speeds, indicating the sustained impact of ASE over time. This method analyzes data points collected at multiple intervals before and after the intervention to identify changes attributable to the ASE implementation. It allows for the assessment of long-term trends and the effectiveness of interventions over time, accounting for any underlying trends or seasonal variations in traffic data.

As seen in Kunshan City, China [23], spatial modelling analyses revealed decreased regional collision risks and significant interaction effects with roadway and land use factors, suggesting that ASE systems can enhance safety across diverse urban environments. Spatial modelling uses geographic information systems (GIS) to assess the spatial distribution of collisions and the influence of various factors on traffic safety. This approach identifies high-risk areas and understands the spatial patterns of traffic collisions, providing valuable insights for targeted interventions.

Systematic reviews, including the one on Speed Cameras for the Prevention of Road Traffic Injuries and Deaths [24], reported reductions in collisions (5% to 69%), injuries (12% to 65%), and deaths (17% to 71%), consolidating global findings and underscoring the substantial safety benefits of ASE systems. Systematic reviews synthesize results from multiple studies to assess the effectiveness of the intervention. They are valuable for summarizing evidence and identifying general trends across different contexts and study designs.

Propensity score matching (PSM) methods used in the UK studies [25] accounted for confounding factors and provided robust evidence of significant accident reductions at camera sites. PSM involves matching treated and control groups based on propensity scores to reduce selection bias and better estimate the treatment effect. This method helps to create a more comparable control group by balancing the distribution of observed covariates, thus improving the validity of causal inferences. These studies cover various road types, from urban freeways and residential streets to school zones and motorways, demonstrating the broad applicability of ASE systems. The number of cameras deployed varied widely, from large-scale implementations with over 2,600 cameras in France [26] to smaller deployments in Seattle with just four cameras [22], highlighting the scalability of ASE systems.

Across these diverse methodologies and contexts, the findings consistently indicate that ASE systems effectively reduce violating the speed limit, collisions, injuries, and fatalities. Some studies report reductions of up to 80% in fatalities within specific zones [27]. This comprehensive evaluation underscores the critical role of ASE systems in modern traffic management strategies, offering substantial safety benefits and significantly reducing traffic-related collisions.

Table 2: Speed camera studies literature review

Author	**Methodology**	**Effectiveness Metric**	**Effectiveness**	**Road Type**	**Number of Cameras**
Thomas et al. [28]	Systematic review	Injurious collisions	20% to 25% reduction in injury collisions	Not specified	Not Applicable
Retting et al. [19]	Before-and-after analysis with comparison groups	Speed reductions	70% reduction in violating the speed limit over 10 mph with both signs and cameras, 39% with signs only	Residential streets and school zones	Not specified
Shin et al. [20]	Before-and-after analysis	Speed reduction and collision reduction	9 mph reduction in average speed, significant collision reduction	Urban freeway	Not specified
Carnis [26]	Empirical analysis	Fatalities and injuries	40% reduction in fatalities	Highway	Over 2,600

Author	Methodology	Effectiveness Metric	Effectiveness	Road Type	Number of Cameras
Gains et al. [29]	Analysis of areas within the program over four years	Speed reduction, casualty and death reduction	70% reduction at fixed sites, 22% reduction in personal injury collisions	Not Specified	38
Rodier et al. [30]	Literature review	Safety and financial effects	Varied, specific quantitative effects not provided	Residential streets	Not App-plicable
Montella et al. [21]	Empirical Bayes observational before-and-after study	Total collision reduction, specifically severe collisions and collisions at curves	31.2% total collision reduction	Motorway	Not specified
Quistberg et al. [22]	Interrupted time series analysis using multilevel mixed linear regression	Speed violation rates, hourly maximum violation speed, mean hourly speeds	53% reduction in speed violation rates, 2.1 mph reduction in maximum speed, 1.1 mph reduction in mean speed	School zones	4

Author	**Methodology**	**Effectiveness Metric**	**Effectiveness**	**Road Type**	**Number of Cameras**
Montella et al. [31]	Before-and-after analysis, empirical Bayes observational before-and-after study	Mean speed, 85th percentile speed, speed variability, proportion of drivers exceeding speed limits	26% reduction in speed variability, 32% reduction in total collisions, significant spillover effect	Urban motorway	1
Jeong-Gyu [27]	Before-and-after comparison	Speed means and variances, fatality rates	40% to 80% reduction in fatalities within 2 km of camera influence area	Highways	Over 1,000
Retting and Farmer [32]	Before-and-after analysis with comparison sites	Mean traffic speeds, proportion of drivers exceeding speed limits	14% reduction in mean speeds, 82% reduction in proportion of drivers exceeding speed limits by more than 10 mph	Residential streets, major arteries, highways, and school and work zones.	7
Wang et al. [23]	Spatial modeling analysis	Regional injury/PDO collision risks	26.1 Decreased regional collision risk	Various urban road types	Not specified

Author	**Methodology**	**Effectiveness Metric**	**Effectiveness**	**Road Type**	**Number of Cameras**
Wilson et al. [24]	Systematic review	Speed reduction, collision reduction, injury and death reduction	Reductions in collisions (5% to 69%), injuries (12% to 65%), deaths (17% to 71%)	Various road types	Not Applicable
Tay [33]	Analysis of speed enforcement data	Injury collision reduction	5.73 less collisions per 1000 issued tickets	Not Specified	Not specified
Li et al. [34]	Difference in difference (DID) based on propensity score matching (PSM)	Number of road accidents	26.78% personal injuries reduction in the early period and 12.18% in final study period	Various road types	771
Moon and Hummer [35]	ARIMA intervention analysis, before-and-after analysis with comparison sites	Long-term collision patterns, carryover safety effects	8% to 18% during the senforcement, and 17% to 21% after uninstalling	Urban corridors	Not specified

Author	Methodology	Effectiveness Metric	Effectiveness	Road Type	Number of Cameras
Pilkington and Kinra [36]	Systematic review	Road traffic collisions, injuries, deaths	Reductions in collisions (5% to 69%), injuries (12% to 65%), deaths (17% to 71%)	Various road types	Not Applicable
Mountain et al. [37]	Evaluation of various speed management schemes	Traffic speeds, accidents	22% reduction in personal injury accidents from speed cameras	Roads with 30 mph speed limit	Not specified
Li et al. [25]	Propensity score matching, empirical Bayes method	Number of accidents	Reduce in personal injury collisions (PICs) by 27.5% within 200 meters, 26.4% within 500 meters, and 18.5% within 1 km.	Various road types	771
Li and Graham [38]	Propensity score matching, local polynomial regression	Personal injury collisions	Reduction in collisions (10% to 40%)	Various road types	771

Author	**Methodology**	**Effectiveness Metric**	**Effectiveness**	**Road Type**	**Number of Cameras**
Novoa et al. [39]	Time series analysis	Number of collisions, number of people injured	30% reduction in collisions, 26% reduction in injuries on beltway	Urban arterial roads and beltway	Not specified
De Pauw et al. [40]	Before-and-after study with control for trend	Number of injury collisions, severe collisions, and fatalities	8% reduction in injury collisions, 29% reduction in severe collisions	Highways	65
Wegman and Goldenbeld [41]	Systmatic review	Casualty reduction	5% to 35% reduction in collisions	Various road types	Not specified
Factor et al. [42]	Before and after comparison with an experimental and a control group	Number of collisions	22% fewer collisions	Not Specified	144
Mountain et al. [37]	Before and after analysis	Collisiosns, Traffic speed	22% reduction in personal injuries	Roads with 30 mph speed limit	Not specified

2.3 ASE Scheduling

As mentioned previously, although it is possible to predict the probable number of violations of community safety zones during enforcement, finding the optimal scheduling plan to maximize safety benefits remains challenging. Only a few studies have been conducted on the best scheduling approach for speed cameras, and none have focused on fixed cameras.[43] suggested a ranking system in their paper for Alberta, prioritizing mobile speed camera enforcement locations. This system uses the Priority Index values calculated with the Equivalent-Property-Damage Only (EPDO) average collision frequency method, as defined by AASHTO in 2010. They proposed a one-month schedule targeting 110 possible sites, with 1080 visits distributed among them. The sites were categorized into three levels based on their Priority Index: Level 1 sites receive 1-9 trips, Level 2 sites receive 12-20 visits, and Level 3 sites receive 22-36 visits. The exact number of visits to each site within its level is randomly determined within the specified range. However, our study differs as it involves longer monthly enforcement periods, whereas Amy's study focused on short, hourly visits. In our approach, cameras are installed at a location for longer, potentially a week or a couple of months, depending on city regulations regarding camera relocation. This difference is crucial as our cameras are not designed for frequent relocation.

[44] used a nonlinear integer programming approach with an objective function to minimize enforcement costs while maximizing the effectiveness of enforcement on average speed. This formulation considered additional costs and benefits, such as the costs of patrolling vehicles and officers' salaries, life-saving and injury-reduction benefits, and collision costs.

[45] focused on determining where and when drivers speed by analyzing probe vehicle data from Prague expressway and collector road segments. After validating the data, they performed a de-

scriptive analysis of violating the speed limit, focusing on homogeneous road segments in individual hour intervals. Their statistical models showed that several factors contribute to violating the speed limit:

- AADT: Only marginal association with violating the speed limit.
- Speed limit: Higher violating the speed limit on 50 km/h segments compared to 70 or 80 km/h segments.
- Number of lanes: Higher violating the speed limit on 1+1 segments compared to 2+2 lanes.
- Median barrier: violating the speed limit increases in order: no barrier, solid barrier, cable barrier.
- Roadside activities: Less roadside activity correlates with higher violating the speed limit.
- Horizontal alignment: Higher violating the speed limit on curves compared to tangent segments.

[46] proposed a model combining particle swarm optimization and support vector machine (PSO–SVM) for traffic safety forecasting. Feature attributes included roadway track in use, mobile car retention quantity, population size, passenger turnover volume, and freight turnover volume, which were used as inputs for the PSO–SVM model. The decision-making attributes included the number of collisions, death toll, and number of injuries, which were used as outputs. Traffic safety data from China, spanning 1970 to 2006, was used to evaluate the forecasting ability of the proposed method. The experimental results showed that traffic safety forecasting using the PSO–SVM model outperformed the BP neural network model in finding the best location for ASE implementation. This

study aims to explore the relationship between ASE site characteristics and the effectiveness of ASE systems in improving road safety. Specifically, it investigates how various factors, such as traffic volume, collision history, and road design, influence the success of ASE in reducing violating the speed limit and related collisions. Additionally, this study develops a comprehensive model that generates a detailed ASE schedule, which determines the optimal locations for camera placement and the enforcement periods at each site. This model maximizes safety benefits while adhering to regulatory requirements and budget constraints. By optimizing camera deployment and scheduling, the model seeks to enhance the overall efficacy of ASE programs in reducing traffic violations and improving road safety across multiple jurisdictions.

2.4 Best Practices of ASE

Speed enforcement cameras have been more widely used in Australia and Britain compared to other countries. In 1991, Victoria, Australia, deployed 54 mobile cameras to monitor speed violators for an average of 4,200 hours per month. Studies have affirmed the effectiveness of Victoria's speed cameras in reducing both the number and severity of accidents, not only at camera locations but throughout the road network. Other Australian states either followed Victoria's approach or increased the use of fixed and uncovered speed cameras [47]. Britain began using speed cameras at a few sites in 1991, expanding to approximately 4,500 cameras by 2000, most of which were fixed cameras that were regularly relocated. British regulations mandate officials to warn drivers about the presence of speed enforcement cameras, requiring warning signs ahead of both fixed and mobile cameras. Although speed cameras have significantly contributed to road safety in Europe and Australia, the United States and Canada have been more hesitant to adopt this approach [48]. New York, Chicago, and Toronto are examples of North American cities implementing ASE 3. In

2021, New York City's ASE assessment showed a 73% reduction in violating the speed limit at ASE locations. Chicago's 2017 ASE assessment reported an 11% decrease in severe collisions at speed camera locations compared to a 2% increase citywide. In Canada, Saskatchewan implemented ASE as a pilot program in school zones from 2014 to 2017, resulting in a 63% reduction in speed-related accidents, 51% fewer injuries, and up to a 17% drop in average vehicle speeds. A 2016 ASE pilot in Quebec reported a 13.3 km/h reduction in average speed and a 15% to 42% reduction in the number of accidents at ASE sites. The Photo Enforcement Safety Program evaluation in Winnipeg showed a 1% to 15% reduction in average speed and a 14% to 65% reduction in the number of violating the speed limit vehicles [7]. ASE in North America:

Many North American cities, including New York, Chicago, and Toronto, use various approaches to address violating the speed limit issues, among which ASE is a common method. In New York City, ASE implementation was part of a larger plan to encourage safe driving, including speed limit changes, installation of speed humps and cushions, street improvements, the collaboration between NYPD and DOT, and driver education using advertisements. Similarly, Chicago's approach to speed management involves road design & engineering Changes, ASE and traffic management. New York City's ASE program was initially piloted in 20 school speed zones in 2013 and later expanded to 140 in 2014 as part of the Vision Zero program. In 2019, the law was amended to increase the program's coverage and hours, and 750-speed cameras were installed in June 2020, covering all school zones and further expanding the program's impact. The December 2021 ASE assessment showed a 73 percent reduction in violating the speed limit at fixed camera locations. In Chicago, the city's first speed cameras were installed in 2013; by 2019, there were about 161 cameras around the city. ASE assessment in 2017 showed an 11% reduction in severe collisions at speed camera locations, compared to a 2% increase citywide.

City	Start Date	Enforcement Type	Cameras	Site Criteria
New York	2013	Fixed	750	School Zones
Chicago	2013	Fixed	161	Children’s safety Zones
Washington D.C.	2001		85	
Philadelphia	2020	Fixed	10	-

Table 3: Compare ASE Practices in New York , Chicago and Toronto

3 ASE Ticketing Data Analysis

The categorization was performed by analyzing trends in speeding citations issued at each location. Locations were classified into three categories based on the observed changes in the number of speeding tickets over time: 'Moderately effective' , where there was a noticeable but limited reduction in speeding violations; 'Highly effective' , where there was a substantial and sustained decrease in speeding violations; and 'Not effective', where there was no reduction or even increase in speeding violations during the enforcement period. These categorized effectiveness levels were then used to develop a predictive model that estimates the likely effectiveness category of each location using road and traffic metrics. Consequently, with the road metrics of any network, it is possible to forecast the trend of ticket issuance during enforcement periods in any previously unanalyzed network. The dataset for this study encompasses the ASE citation data from the City of Toronto from June 2020 to the end of May 2023. This data is publicly accessible through the City of Toronto's open data portal. It includes information on the enforcement start and end dates, the number of tickets issued monthly during enforcement periods for each location, and the specific site locations where enforcement cameras were installed.

We also introduced an additional column to the dataset labelled "cycle" to indicate the installation cycle of the enforcement cameras at each location. Since the locations were organized based on the start date of enforcement, we manually determined the cycle number, primarily relying on this start date for classification. However, some locations underwent enforcement during multiple distinct cycles, reflecting periods of active enforcement, followed by gaps without enforcement, and then a return to active enforcement. For instance, site code A069 was enforced across three different cycles: the second, the fifth, and towards the end of the sixth cycle. The cycle numbers

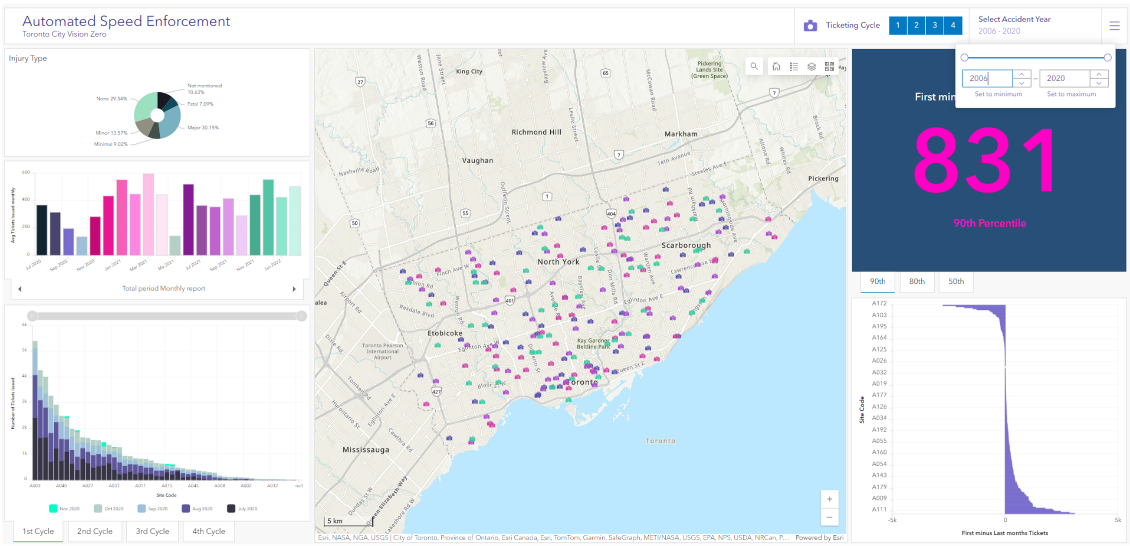

Figure 1: Toronto ASE dashboard displays the locations of ASE cameras at each cycle

are approximated because the relocation of cameras does not occur simultaneously across all sites.

3.1 Data Preparation

3.1.1 ArcGIS Dashboard

For visual analysis and geographic representation, ArcGIS Online was used to geocode the locations identified in the ASE dataset. This process involved using the ”location” column from the dataset to plot each ASE camera site accurately on a map. This initial step allowed for visualizing the spatial distribution of enforcement cameras across the City of Toronto and laid the groundwork for further analysis. After geocoding the ASE locations, an analytical dashboard using ArcGIS was developed, as seen in Figure1. This dashboard was designed to display the precise locations of ASE cameras and incorporate data on fatal and severe injuries from speeding-related collisions in Toronto since 2006. This integration of ASE and collision data provided a comprehensive overview of traffic safety in the city. The ASE dashboard, developed in the summer of 2022, showcases data from July 2020 to February 2022. It features several interactive elements that enhance user engagement and data exploration: [1]

- Filtering Capability: Users can filter the displayed camera locations based on the installation cycle. This allows for a targeted examination of enforcement effectiveness over time.

- Heat Map of Collisions: The background of the dashboard includes two heat maps that visualize the concentration of Killed and Serious Injuries (KSI) and speeding-related collisions through the city of Toronto, offering insights into high-risk areas as illustrated in Figure2 and Figure3.

- Yearly Collision Filter: A dual-thumb slider and text input boxes enable users to filter collision data by specific year ranges, facilitating temporal analysis of traffic safety trends.

- Monthly Ticketing Trends: On the right side of the dashboard, graphical representations display the average number of tickets issued city-wide for each month within the covered period. These visualizations are available for each enforcement cycle, revealing patterns such as a decreasing trend in the first cycle, fluctuating trends in subsequent cycles, and noticeable increases in certain periods as seen in figures 4a and 4b.

The dashboard also provides a detailed breakdown of violations and enforcement effectiveness. For example, discrepancies in the average number of tickets issued across different cycles highlighted the variability in enforcement impact. Several factors contribute to this variability:

- Partial Month Enforcement: The calculation of average monthly tickets did not account for cameras that were not operational for the entire month, leading to potential overestimations of enforcement effectiveness.

[1] The dashboard can be accessed at: ASE Dashboard

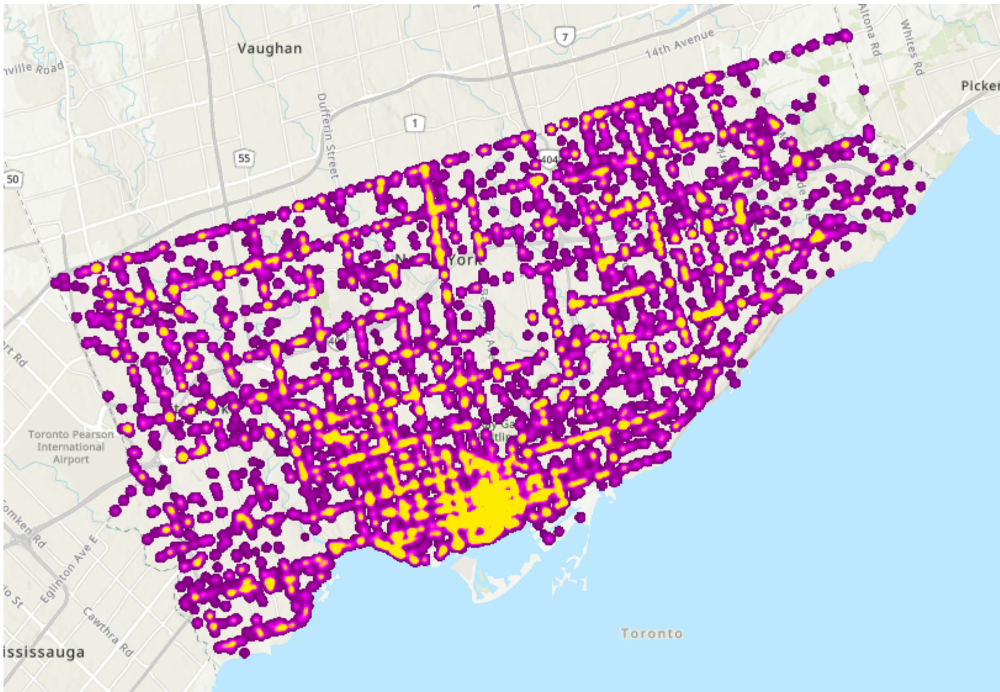

Figure 2: Toronto KSI collisions heatmap, based on 2016 to 2019 data

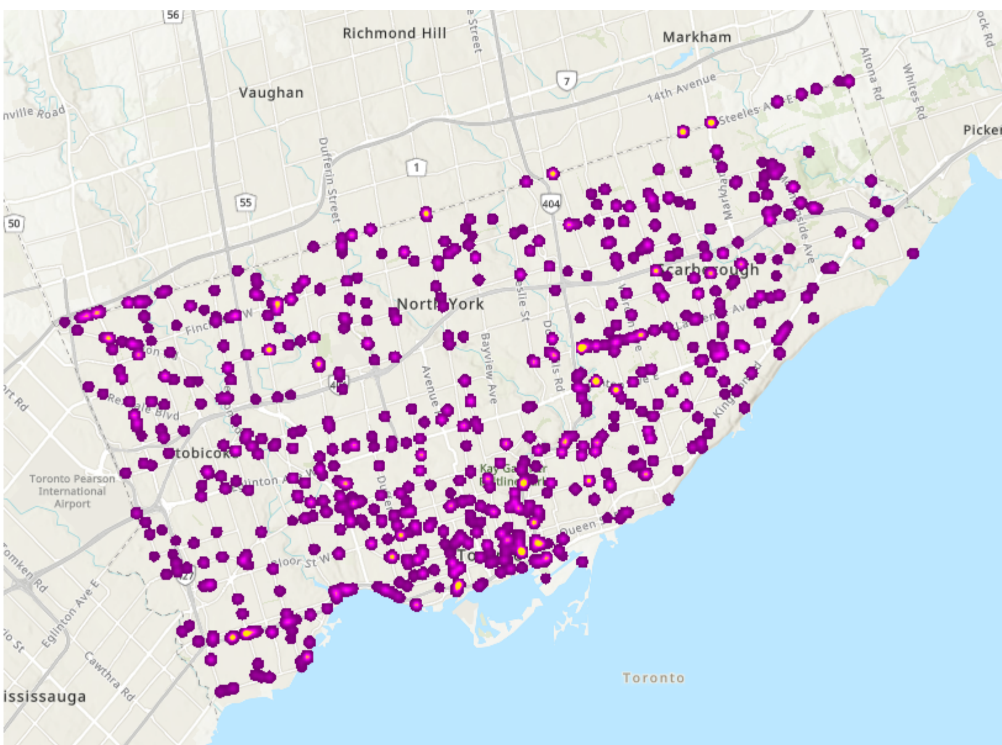

Figure 3: Toronto speeding collisions heatmap, based on 2016 to 2019 data

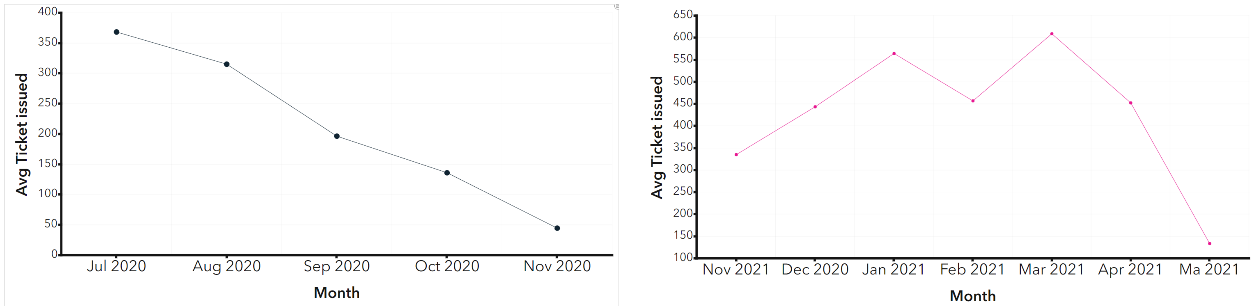

(a) Average number of tickets issued per month during first cycle

(b) Average number of tickets issued per month during second cycle

Figure 4: Visualization of average number of tickets per month during 1st and 2nd cycle

- Cycle Definition and Enforcement Start Dates: Due to significant variations in enforcement start dates and cycle durations across different locations, some sites may have been active for only part of an assigned cycle, affecting the analysis of ticketing trends based on the cycles.

The dashboard visually represents the total number of violations captured for each site code, sorted by the total number of violations. Violations per month are colour-coded to illustrate the proportion of violations occurring each month. A pie chart at the top right corner of the map details the proportion of each type of injury resulting from speeding collisions, providing a stark visual reminder of the human impact of traffic violations.

The ASE dashboard's interactive features and comprehensive data presentation significantly enhance our understanding of enforcement effectiveness and traffic safety trends. This visualization tool not only aids in the immediate analysis but also informs our approach to data cleaning and further analytical steps, ensuring a robust foundation for subsequent research phases.

3.1.2 Integrating Traffic Data

The enhancement of each ASE site code with comprehensive road and traffic characteristics enriches the dataset used in this study, allowing for a more nuanced analysis of factors influenc-

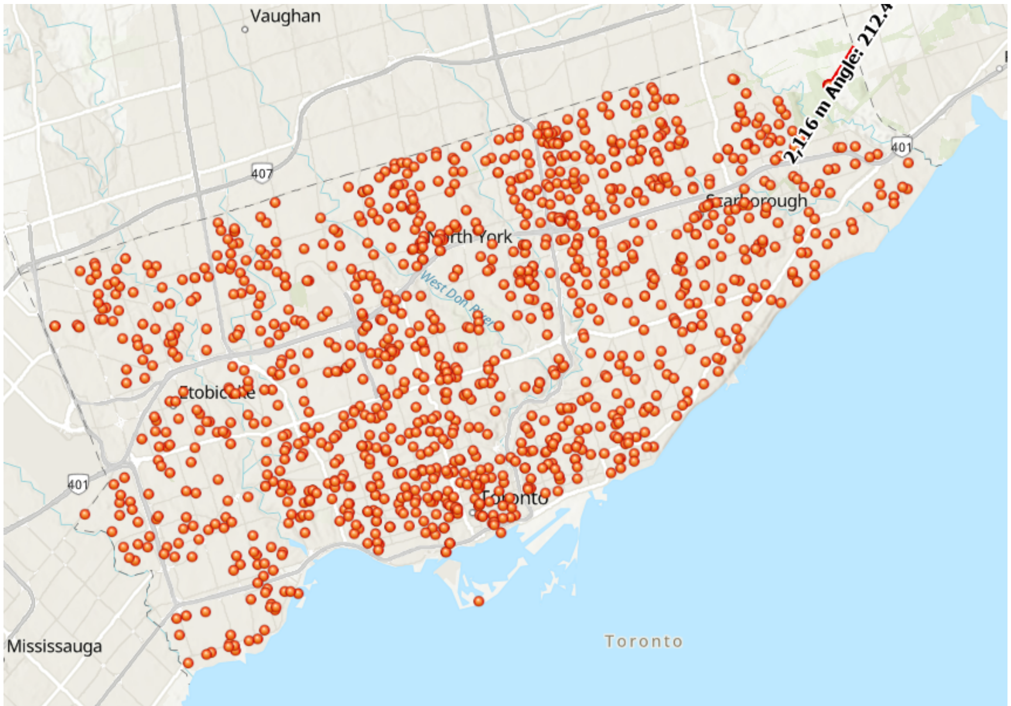

Figure 5: Locations of schools throughout city of Toronto

ing ticketing patterns and road safety. This section details the process of integrating various road and traffic data metrics with each ASE site location, leveraging the spatial analysis capabilities of ArcGIS Pro for precise data alignment and analysis. Incorporating Toronto's road and traffic data—including speed limits, road types, the number of lanes, traffic volumes, proximity to schools, and records of traffic collisions—involved aggregating information from diverse sources. This multi-faceted dataset was attached to each site location through a spatial join operation in ArcGIS Pro, ensuring that each ASE site was analyzed within its surrounding road environment. One of the key metrics evaluated was the presence of educational institutions within a 1-kilometer radius of each ASE site. Using the within-the-radius feature of the spatial join tool allowed for the precise calculation of schools near each enforcement camera, reflecting the potential for higher pedestrian activity and the critical need for speed regulation. Figure 5 displays the geocoded locations of schools throughout the city of Toronto.

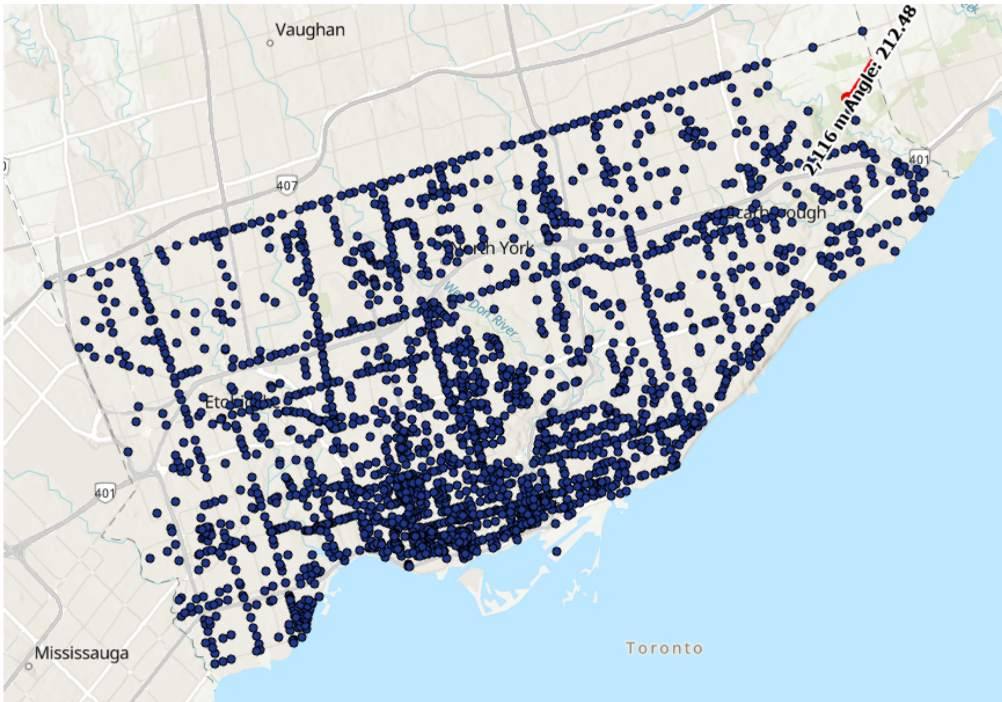

Figure 6: Locations of traffic volume data available throughout city of Toronto

Traffic volume counts were another essential component of this analysis. Initially gathered from the City of Toronto, the dataset encompassed traffic flows at 2360 intersections, offering a detailed breakdown of transportation modes—including bicycles, buses, cars, trucks, pedestrians, and other modes—approaching from each direction. These counts, measured in 15-minute intervals, provided an in-depth view of traffic dynamics at various times of the day. By aggregating these interval counts to sum traffic flows for entire intersections and then obtaining the hourly rate, we established a comprehensive metric for traffic volume. This methodology allowed for assessing traffic density and flow around ASE sites, with average hourly volumes obtained for areas within a 1-kilometer radius of each location. Such an approach affords insights into the relationship between traffic volume and the likelihood of speeding violations. Figure 6 displays the geo-coded locations of the traffic counts data. Furthermore, data on severe and fatal injury collisions between 2016 and 2019 were compiled, focusing on incidents within 1 to 1.5 kilometres of each ASE site. This historical collision data provides a critical backdrop for understanding the safety challenges of these areas and the potential impact of ASE. Additional road characteristics, such as the number of lanes and road type for each segment, were sourced from OpenStreetMap. While comprehensive, this dataset had gaps in information for some road segments, necessitating careful consideration in the analysis. Speed limit data, vital for understanding enforcement contexts, was retrieved from Can Map route logistics. Notably, anomalies in speed limit records for two road segments—listed initially at 10 km/h—were corrected to 30 and 50 km/h after cross-referencing with Google Street View, ensuring the accuracy of enforcement zone data. The addition of road and traffic characteristics to each ASE site code through spatial analysis techniques enriches the dataset, paving the way for a detailed investigation into the factors influencing speeding enforcement effectiveness and road safety. This integrated approach not only enhances the depth of the study but also supports the development of targeted, data-driven strategies for improving traffic safety in Toronto.

3.1.3 Initial Data Cleaning

In normalizing and cleaning our dataset, a critical issue regarding the enforcement start and end dates across various locations was addressed. Notably, enforcement did not consistently commence on the first day or conclude on the last day of the month, leading to disproportionately lower ticket counts in the initial and final months of enforcement for many locations. Uniform ticket distribution within each month was assumed to rectify this discrepancy and achieve a more accurate representation of ticketing patterns. For each location and month, the daily average of violations was calculated by dividing the total number of tickets by the actual enforcement days. This approach allowed us to standardize the data, making it more reflective of the proper enforcement activity. However, challenges arose with missing start or end date records for some site codes. In cases where the ticket count for the first recorded month was significantly lower than the subsequent month, we inferred that enforcement likely commenced towards the month's end and opted to exclude this data from the analysis. When the data appeared to have a logical number of tickets for the first month, we proceeded under the assumption that enforcement began on the first recorded day of ticketing for that location. Similarly, we determined the end of enforcement based on the last day tickets were issued, utilizing the actual calendar days to ensure precision. For example, site A001's data indicated enforcement from July 6, 2020, to November 11, 2020. The noticeably lower ticket count in the final month suggested reduced enforcement days, validating our approach of daily violation recalculations for a more logical analysis. Figures 7a and 7b illustrate the adjustment process, showing how the tickets for site A001 in November are excluded from the analysis after averaging the number of tickets per day due to the significantly lower count compared to other months.

A	B	C	D	E	F	G	H	I
Site Code	**Location***	**Enforcement Start Date**	**Enforcement End Date**	**Jul-20**	**Aug-20**	**Sep-20**	**Oct-20**	**Nov-20**
A001	Royalcrest Rd. Near Cabernet Circle	6-Jul-20	11-Nov-20	1,117	479	487	354	39
A002	Harefield Dr. Near Barford Rd.	6-Jul-20	25-Oct-20	43	21	10	16	-
A003	Renforth Dr. Near Lafferty St.	6-Jul-20	25-Oct-20	2,428	1,665	1,037	268	-

(a) screen shot of the original ASE ticketing excel file

Site Code	Country	City	lon	lat	Location*	Dur	Enforceme	Enforcement End	20-Jul	20-Aug	20-Sep	20-Oct	20-Nov	20-De
A001	Candana	Toronto	-79.6043	43.75271	Royalcrest	128	7/6/2020	11/11/2020	36.03226	15.45161	16.23333	11.41935		
A002	Candana	Toronto	-79.5692	43.72828	Harefield D	111	7/6/2020	10/25/2020	1.387097	0.677419	0.333333	0.516129		
A003	Candana	Toronto	-79.5901	43.66718	Renforth D	111	7/6/2020	10/25/2020	78.32258	53.70968	34.56667	8.645161		
A004	Candana	Toronto	-79.5333	43.69211	Trehorne D	111	7/6/2020	10/25/2020	27	13.3871	7.966667	5.419355		

(b) Calculation of the average number of violations per day for each month and removing unreliable cells

A	B	N	O	P	Q	R	S	T	U	V	W
Site Code	centreline	length	SPEED_LIM	Shape__Le	Col_1.5	Col_1	volume	Violations	Vio_rate	Y	T
A001	2:	253.376	50	356.1697	41	14	75972.29	23429.57	0.245831	36.03226	1
A001	2:	253.376	50	356.1697	41	14	75972.29	23429.57	0.245831	15.45161	2
A001	2:	253.376	50	356.1697	41	14	75972.29	23429.57	0.245831	16.23333	3
A001	2:	253.376	50	356.1697	41	14	75972.29	23429.57	0.245831	11.41935	4

(c) site code A001 is repeated four times to represent the four months of enforcement activity for fitting regression model

Figure 7: Toronto ASE ticketing data cleaning and preparation process

After recalculating daily violation rates and removing unreliable data points (as demonstrated in the transition from the initial dataset to the adjusted one displayed in figures7a and 7b), we compiled summary statistics to gain insights into ticketing trends across Toronto. Our findings revealed:

- Mean Daily Ticket Count: Fluctuated over time within an approximate range of 10 to 20 tickets per day, with notable peaks and troughs, underscoring the variability in enforcement effectiveness.
- Maximum Daily Ticket Count: Exhibited significant month-to-month variation, peaking in January 2021 with over 200 tickets per day. This highlighted periods of heightened enforcement or increased violation rates.
- Minimum Daily Ticket Count: Remained consistently low, indicating a persistent presence of locations with minimal ticket issuance each day.
- Median Daily Ticket Count: Showed moderate fluctuations, suggesting a balanced distribution of higher and lower enforcement sites.

This refined data normalization and cleaning approach, grounded in clear assumptions and careful recalculations, provides a reliable foundation for analyzing enforcement patterns and effectiveness. The resulting summary statistics give a more detailed view of the dynamics of speed violations during ASE enforcement in Toronto.

3.2 Evaluate ASE effectivity

The subsequent phase of our analysis, which aimed at assessing the effectiveness of ASE across various site codes, involved developing a framework to determine whether ASE has been effective at one site. To this end, we used linear regression analysis for each site code, which allowed us to explore the relationship between time (as the enforcement month number) and the enforcement impact (measured by average violations per day). For this regression analysis, we prepared a dataset where each entry corresponds to a specific duration of enforcement for a given site code. Expanding the original dataset, each site code's data entry was replicated according to the months with recorded violations. For instance, as illustrated in Figure 7c, the entry for site code A001 is repeated four times to represent the four months of enforcement activity documented in Figure 7b. In this new dataset, two critical columns were added to each row: the first, denoted as 'X,' represents the duration of enforcement, serving as a temporal marker of the enforcement period. The second, labelled 'Y,' quantifies the average number of violations per day for the corresponding month. This approach facilitates a detailed month-by-month analysis of enforcement effectiveness, with 'X' incrementally increasing with each subsequent month of enforcement and 'Y' reflecting the observed enforcement outcomes.

3.2.1 Fit Linear Regression

After organizing the data into this structured format, linear regression models were fitted for each site code individually. This process involved classifying the expanded dataset entries by site code to ensure that the regression analysis accurately reflects the enforcement dynamics unique to each location. For sites experiencing multiple enforcement cycles, such as A069, the 'X' value resets to 1 with each new cycle, based on the premise that sufficient time elapsed between cycles to consider them as independent enforcement periods.

The linear regression analysis yielded several key metrics for each site code: the slope of the regression line, indicating the trend in enforcement effectiveness over time; the y-intercept, offering a baseline measure of violations; the coefficient of determination (R^2), providing insights into the model's explanatory power; and the Mean Squared Error (MSE), assessing the model's prediction accuracy all available in appendix B.

3.2.2 Clustering ASE Effectivity

In this analysis, we aimed to evaluate the effectiveness of Automated Speed Enforcement (ASE) at various locations based on ticket issuance data over time. Fitting linear regression models for each site allowed us to quantify trends in ticket counts during enforcement. The regression slope provided a clear metric for determining ASE effectiveness: a negative slope indicated a reduction in speed violations (effective ASE), while a positive slope suggested an increase in violations (ineffective ASE). This regression-based approach provided a standardized way to compare enforcement outcomes across sites.

However, not all locations with negative slopes exhibited the same level of effectiveness. Some sites showed steep declines in violations, while others demonstrated more gradual reductions as demonstrated in Figure 8 . To better understand these variations and tailor predictive models to different types of locations, we needed to group sites with similar characteristics. Clustering was essential for identifying patterns in ASE effectiveness and ensuring that sites with comparable enforcement outcomes were analyzed together. This step also facilitated feature engineering, enabling us to investigate the specific road attributes contributing to ASE effectiveness within each group.

We considered several clustering methods to categorize sites with negative regression slopes: percentile-based clustering, hierarchical clustering, density-based clustering (DBSCAN), and K-means. Percentile-based clustering used predefined thresholds to divide sites into categories based on slopes and intercepts. Although straightforward, this method was rigid and did not capture the natural structure of the data, leading to over-simplified classifications. Hierarchical clustering provided a visual dendrogram to explore relationships between sites, but was computationally intensive and less practical for larger datasets. DBSCAN, which identifies clusters based on dense regions of data points, was robust against outliers and worked well for irregular cluster shapes. However, It required fine-tuning parameters like ε (neighborhood radius) and *MinPts* (minimum points), which proved challenging for our dataset, where clusters were defined more by linear trends than density. Ultimately, we selected K-means clustering as the most appropriate method. It efficiently grouped sites based on regression slopes and intercepts, minimizing variance within clusters while allowing natural patterns in the data to emerge. The elbow method determined the optimal number of clusters, ensuring that the classification balanced detail with simplicity. Unlike percentile-based methods, K-means did not impose arbitrary thresholds and were computationally more efficient than hierarchical clustering. Furthermore, K-means aligned better with linear regression-based trends in ASE effectiveness than DBSCAN, providing clear and interpretable results. By combining regression analysis with K-means clustering, we created a robust framework to analyze the effectiveness of ASE and identify key factors that influence its outcomes.

The initial classification examined regression slopes and intercepts, assigning each site a code corresponding to a specific cluster. The data set was classified into two groups based on the slope of the fitted regression lines in Figure 8:

- Locations with a positive slope were labeled ineffective, indicating increased violations after deployment of the ASE.
- Conversely, locations with a negative slope, which indicated a decrease in speed violations, were deemed effective.

We utilized the K-means algorithm, an unsupervised learning technique, to categorize site codes

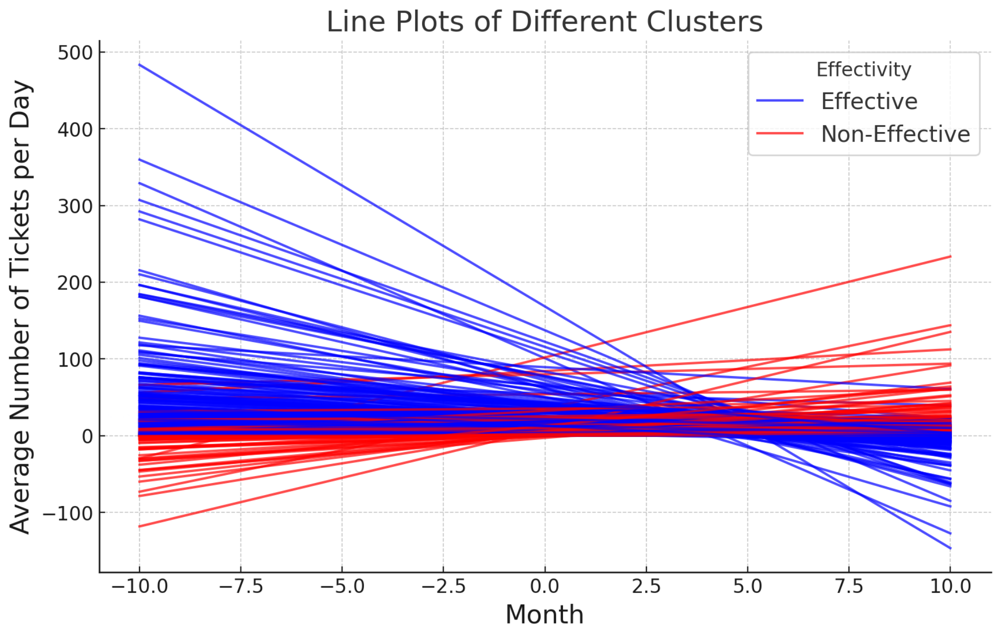

Figure 8: Fitted linear regression lines clustered into two groups

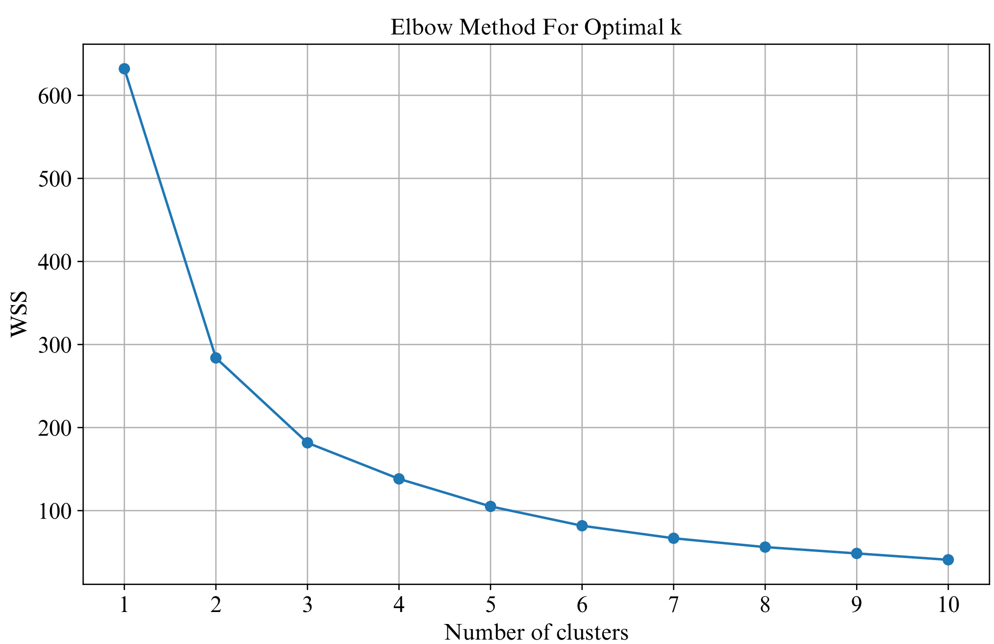

Figure 9: implement elbow method to determine the optimal number of clusters using the WSS

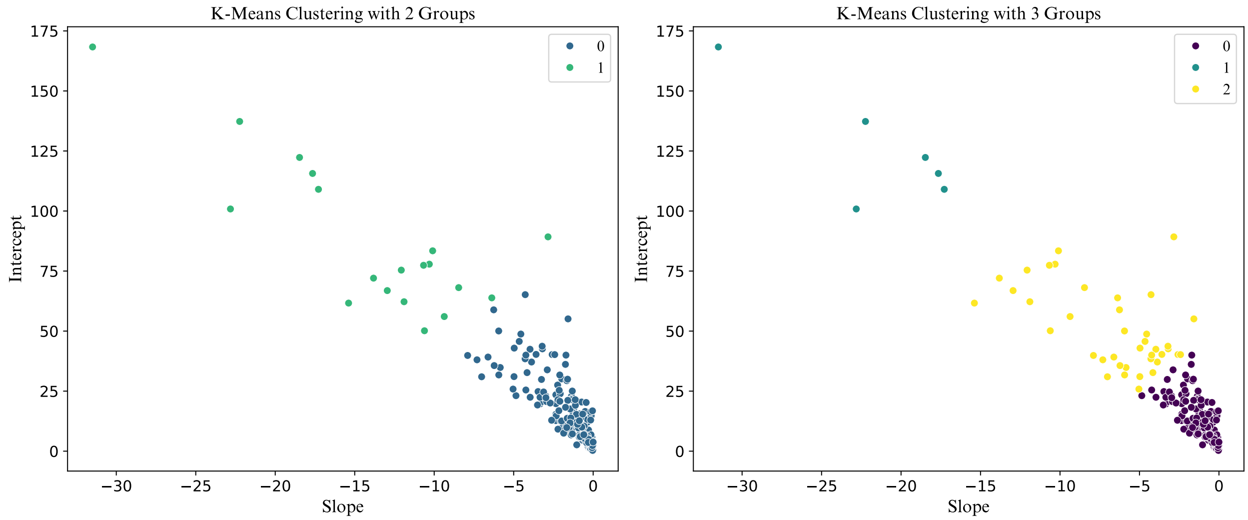

Figure 10: Kmeans to cluster negative slope locations for k=2 and k=3

based on the magnitude of the regression slopes and the initial number of violations. K-means minimized variances within clusters by iteratively assigning data points to the nearest centroids and updating centroids based on the mean location of the assigned points. The optimal number of clusters, k, was determined using the elbow method. The Elbow Method is a technique used to determine the optimal number of clusters in a dataset by analyzing the Within-Cluster Sum of Squares (WSS)[2]. It involves plotting the WSS against the number of clusters (k) to form an ”elbow-shaped” curve. The WSS decreases as the number of clusters increases because each cluster will be smaller and more compact. However, after a certain point, the decrease in WSS decreases significantly, creating an ”elbow” in the plot. The optimal number of clusters is identified at this elbow point, where the addition of more clusters no longer results in a substantial reduction in WSS, balancing the trade-off between minimizing variance within clusters and avoiding overfitting with too many clusters. This method plots the sum of squared distances from points to their respective cluster centroids against different k values to identify the elbow point, as plotted in Figure 9, where

[2] **Within-Cluster Sum of Squares (WSS)** quantifies how compact the clusters are; a lower WSS indicates that the points are closer to their cluster centroids, which is a desirable property in clustering.

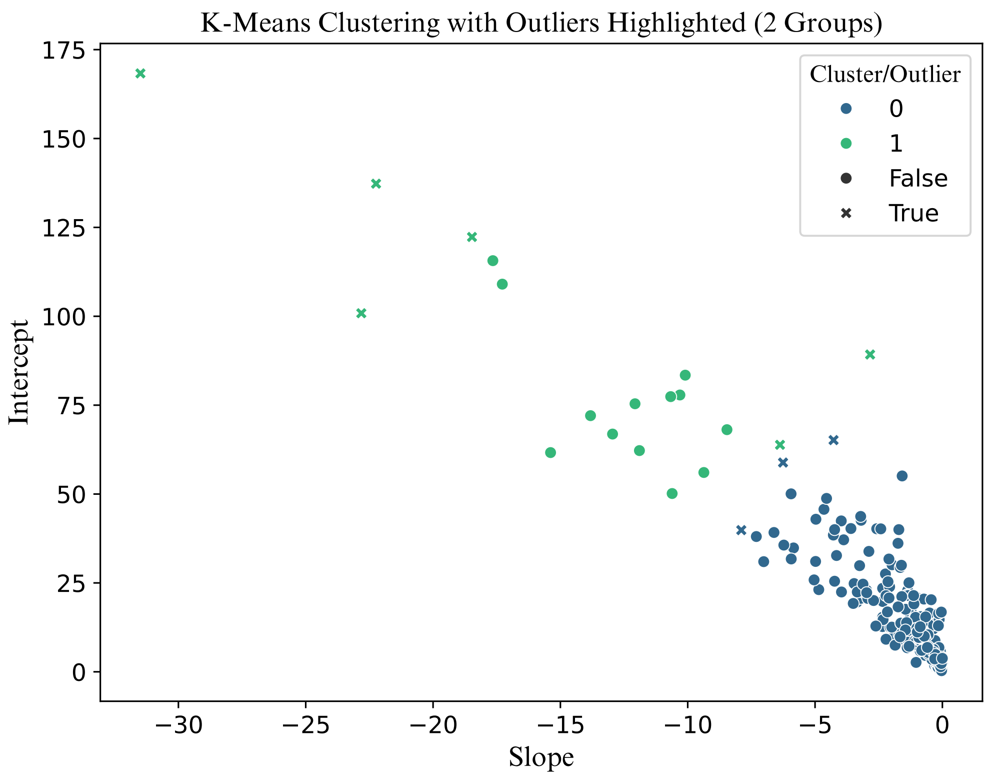

Figure 11: Implement Silhouette analysis to determine outliers in Kmeans clustering for K=2

increases in k yield diminishing returns. We explored k values of 2 and 3 and decided on k=2 due to the dataset's size, as k=3 resulted in sparser clusters in 10.

To validate the chosen clusters for k=2, a Silhouette analysis was conducted, which assesses cluster cohesion and separation shown in Figure11. Based on this clustering, the city of Toronto ASE site location will classified into three groups as displayed in 12. This analysis helped identify outliers and confirm the fittingness of site codes within their clusters. This validated clustering allowed us to classify ASE effectiveness into three levels: Not Effective, Moderately Effective, and Highly Effective. By systematically categorizing road segments and employing targeted feature engineering based on our clustering results, our methodology enhances the prediction accuracy of ASE's impact. This approach facilitates a detailed understanding of how ASE affects road safety and ensures that our predictive models are well-suited for real-world applications.

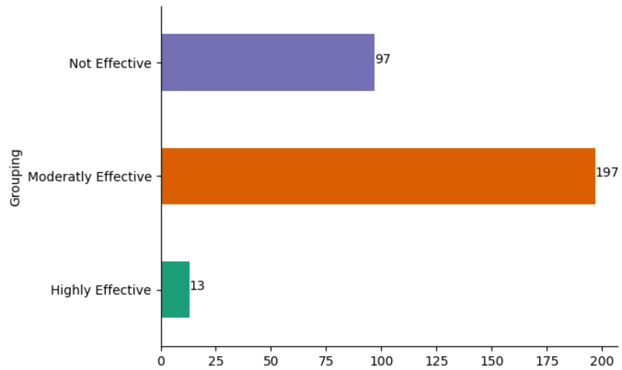

Figure 12: Density of each cluster among Toronto ASE locations

3.3 Data Preprocessing, Feature Engineering, and Modeling

In this section, steps include preprocessing, feature engineering, and applying various machine learning models to predict the ASE location's predictivity clusters as defined in the previous section. In Figure13, the overall architecture of the model, including the data preprocessing steps, feature engineering, and the application of various machine learning models, is shown. In the next step, each component will be described in full detail.

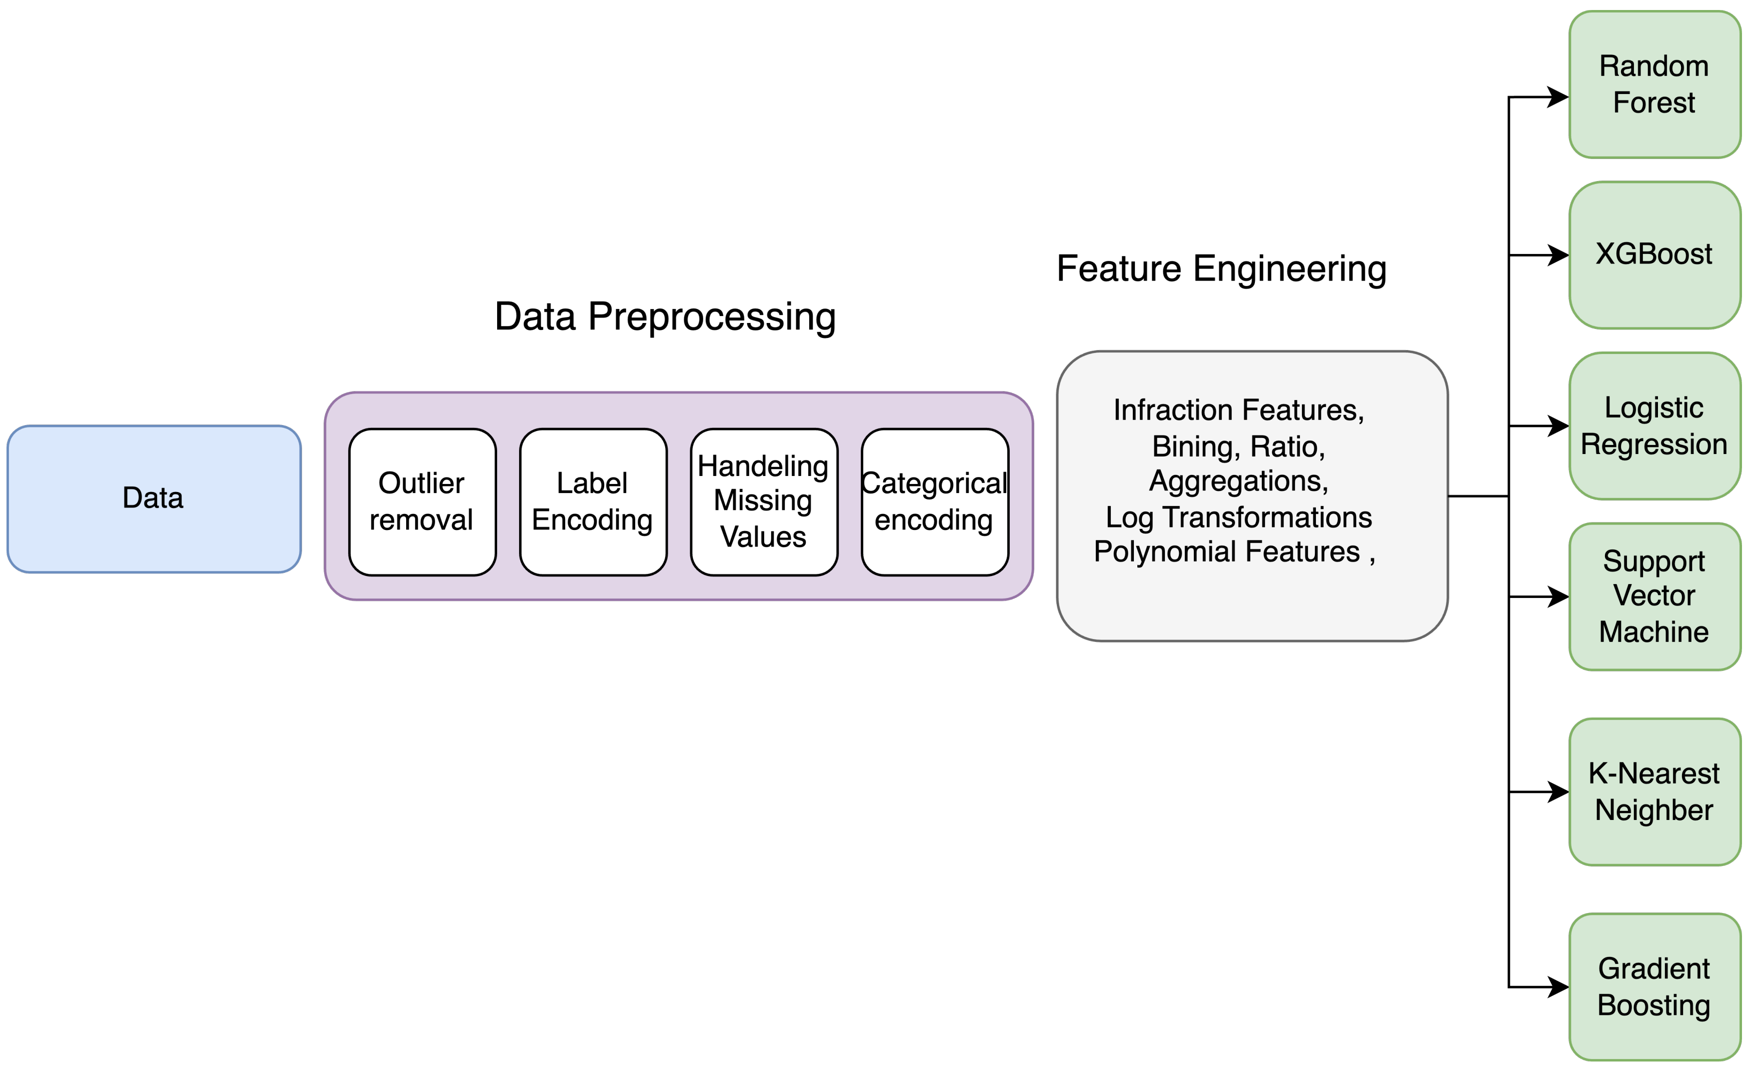

Figure 13: Overall Architecture of the Model

3.3.1 Data Cleaning and Preprocessing

In this study, before the data gets into the machine learning, a preprocessing step is applied:

1. **Outlier Removal**: First, to ensure the integrity of our dataset, rows that are considered outliers were removed.

2. **Label Encoding**: The 'effectiveness' column, which categorizes the effectiveness of ASE locations, was transformed into numeric labels using *LabelEncoder*. This transformation was necessary because machine learning algorithms require numeric input. The labels were assigned as follows:

    - Moderately effective was labelled as 0.
    - Highly effective was labelled as 1.
    - Not effective was labelled as 2.

3. **Handling Missing Values**: Missing values can lead to biases and inaccuracies in the model, so it is crucial to address them effectively. To handle missing values in our dataset, we adopted a strategy based on the average values for each *road_type*. For example, if a road segment had a missing value in the *number_of_lanes* column, we filled it with the average number of lanes for that specific road type. For instance, the average number of lanes for residential streets in our dataset is 2; if a residential street segment had a missing *number_of_lanes* value, we imputed it with 2. Similarly, if the *speed_limit* column had missing values and the average speed limit for a particular road type was 50 km/h, we used 50 km/h to fill in those gaps. By using the average values specific to each *road_type*, we ensured that

the imputed data was realistic and contextually appropriate, thereby maintaining the integrity of our dataset and improving the accuracy of our analysis.

4. **Categorical Encoding**: Categorical features such as *one_way* and *road_type* were converted to numeric values. Encoding categorical features is crucial because machine learning algorithms cannot directly process categorical data. By converting these categories into numeric codes, we enable the algorithms to interpret and utilize these features effectively.

To enhance the predictive power of the model, several new features were generated from the existing features. These generated features help capture complex patterns and relationships in the data that may not be immediately apparent from the raw variables. Here is a detailed explanation of each type of feature and its importance:

1. **Interaction Features**: Interaction features are created by combining two or more existing variables. They are useful because they can capture the combined effect of multiple factors, which might be more predictive than the individual factors alone.

   - speed_limit_x_lanes: This feature is the product of the speed limit and the number of lanes. It captures the interaction between the road's speed limit and its capacity (number of lanes). For example, a high-speed road with many lanes might have different traffic dynamics and enforcement effectiveness compared to a high-speed road with few lanes.
   - *traffic_volume_x_violations*: This feature is the product of traffic volume and the number of violations. It indicates how traffic density interacts with the frequency of violations. Roads with high traffic volume and high violations may require different enforcement strategies than less busy roads with fewer violations.

2. **Binning**: Binning converts continuous variables into categorical bins. This can simplify the model and capture non-linear relationships.

    - *traffic_volume_bin*: Traffic volume was divided into three categories. This helps in understanding how different levels of traffic volume impact the effectiveness of enforcement, which might not be linear.

    - *violations_bin*: Violations were also binned into three categories. This helps to categorize locations based on their violation rates, making it easier to compare and analyze patterns across different levels of violations.

3. **Ratios**: Ratio features provide relative measures, offering insights into proportional relationships between variables.

    - *violations_to_traffic_volume*: This ratio indicates the number of violations relative to the traffic volume, providing a normalized measure of enforcement effectiveness that accounts for the volume of traffic.

    - *speeding_collision_ratio*: This ratio combines the speeding collisions within 1 km and 1.5 km to traffic volume. It helps to understand the proportion of speeding collisions in relation to the traffic density around ASE sites.

4. **Aggregates**: Aggregated features summarize data by combining multiple observations, providing a higher-level view.

    - *total_speeding_collisions*: This is the sum of speeding collisions within 1 km and 1.5 km. It aggregates the total number of speeding incidents near each enforcement site, providing a measure of overall safety in those areas.

- *average_speeding_collisions*: This is the average of total speeding collisions. It offers a normalized measure that can be used to compare different sites regardless of their size or traffic volume.

5. **Log Transformation**: Log transformation is applied to variables to reduce skewness and normalize distributions, which can improve model performance.

   - *log_traffic_volume*: Applying a log transformation to traffic volume helps in handling wide ranges of values and reducing the impact of extremely high traffic volumes.
   - *log_violations*: Similarly, log-transforming the number of violations helps in normalizing the data, making the distribution more symmetrical and the model more robust.

6. **Polynomial Features**: Polynomial features are created by combining variables up to a specified degree, capturing non-linear relationships.

   - Polynomial features were generated using *PolynomialFeatures* from sklearn. This technique captures interactions between multiple variables and their non-linear effects, which can significantly improve the model's ability to predict complex patterns in the data.

Each of these engineered features plays a crucial role in enhancing the model's ability to learn from the data. By incorporating these features, we can capture more intricate patterns and relationships, leading to improved predictive performance and a deeper understanding of the factors influencing ASE effectiveness. The following are a series of visualizations that provide insights into features derived from the generated dataset.

- **Histogram of Speed Limit with Number of Lanes**: This histogram shows the distribution of the interaction feature *speed_limit_x_lanes*. This feature captures the combined effect of a road's speed limit and the number of lanes, which is crucial for understanding traffic dynamics and enforcement effectiveness. As shown in Figure 14, the majority of data points fall within a specific range, indicating common combinations of speed limits and lane numbers across different locations. This concentration suggests that certain road configurations are more prevalent, which could influence traffic patterns and enforcement outcomes.

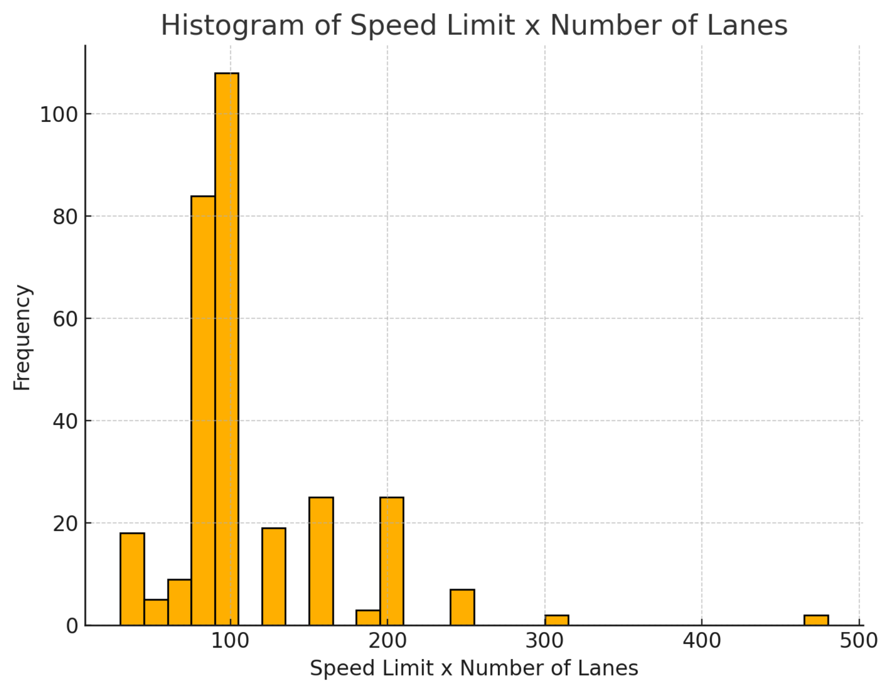

Figure 14: Histogram of Speed Limit $\times$ Number of Lanes

- **Histogram of Violations**: This histogram illustrates the distribution of traffic violations recorded across different locations. Understanding the frequency and distribution of violations is fundamental for analyzing the impact of ASE on road safety. As depicted in Figure 15, the data shows a wide range of violation counts, highlighting the variability in enforcement outcomes. The skewed distribution with a long tail indicates that while most locations have relatively few violations, a few sites experience a very high number of violations, which could be hotspots for enforcement.

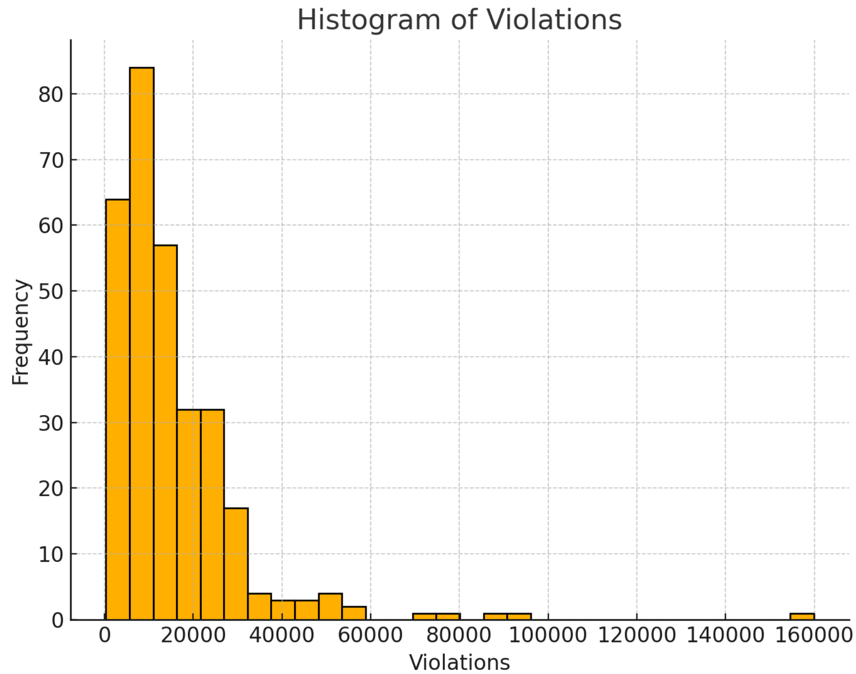

Figure 15: Histogram of Violations

- **Histogram of Log Traffic Volume**: The log-transformed traffic volume helps normalize the data and reduce skewness. This histogram displays the distribution of $log\_traffic\_volume$, providing a more symmetrical view of traffic data. As shown in Figure 16, this transformation is essential for improving the performance of predictive models by stabilizing variance and making patterns more discernible. The central peak indicates the most common levels of traffic volume, which helps in understanding typical traffic conditions.

- **Box Plot of Violations across Effectiveness**: This box plot compares the distribution of violations across different categories of enforcement effectiveness. It provides a clear visualization of how violation counts vary based on the effectiveness of enforcement measures. As illustrated in Figure 17, the box plot highlights median values, quartiles, and potential out-

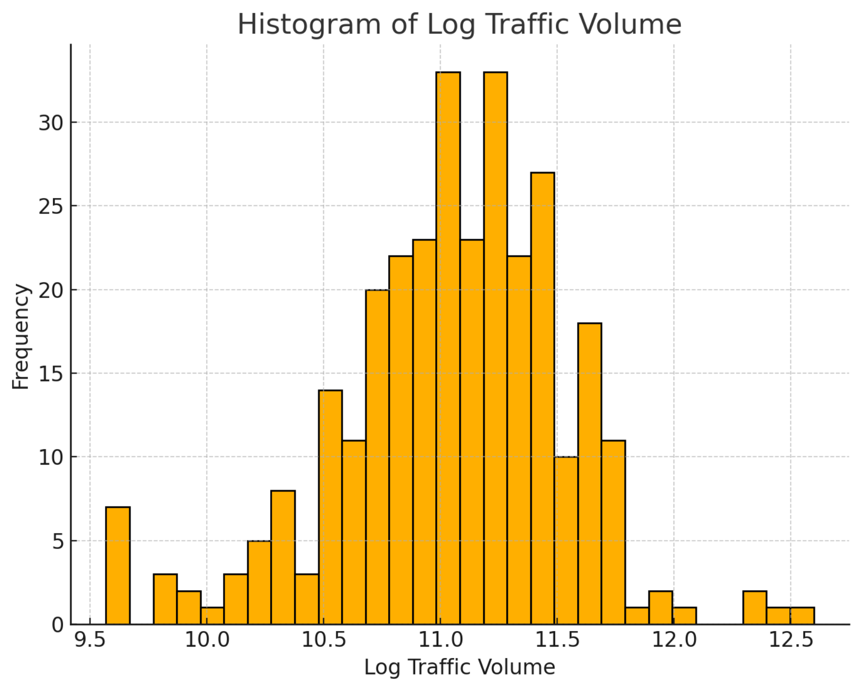

Figure 16: Histogram of Log Traffic Volume

liers, offering insights into the central tendency and dispersion of violations. The presence of outliers, especially in the 'Highly effective' category, suggests that some locations have exceptionally high or low violations compared to the typical range.

- **Bar Plot of Average Violations per Effectiveness Category**: This bar plot shows the average number of violations for each category of enforcement effectiveness. By comparing these averages, we can identify trends and assess the relative impact of different enforcement strategies. As shown in Figure 18, the plot reveals which categories are associated with higher or lower average violation counts, aiding in the evaluation of ASE effectiveness. It is evident that 'Not effective' sites tend to have higher average violations, indicating a need for improved enforcement measures in these areas.

These visualizations provide a comprehensive overview of the key features and their distributions,

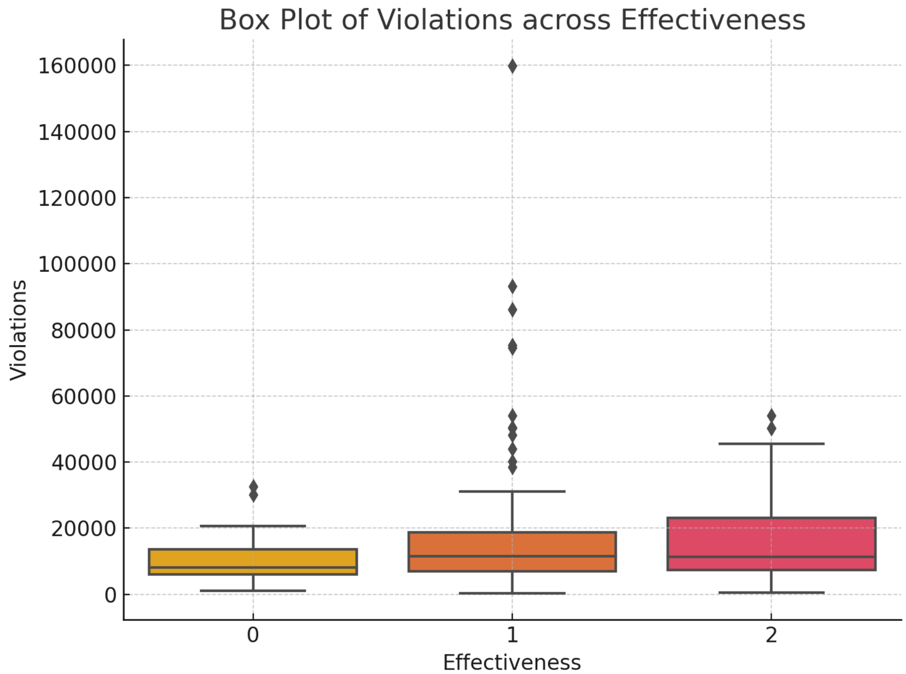

Figure 17: Box Plot of Violations across Effectiveness

helping to elucidate patterns and relationships within the dataset. Getting insight from this visualization, a correlation analysis was conducted to identify and remove features that either do not significantly contribute to the target variable or are redundant due to high correlation with other features. The steps involved are as follows:

1. **Correlation Matrix**: We calculated the correlation matrix to understand the relationships between all the features in the dataset. The correlation matrix is a table that shows the correlation coefficients between variables. Each cell in the matrix represents the correlation between two variables. For instance, a value close to 1 indicates a strong positive correlation, a value close to -1 indicates a strong negative correlation, and a value around 0 indicates no correlation. By examining the correlation matrix, we can identify which features are highly correlated with each other and with the target variable.

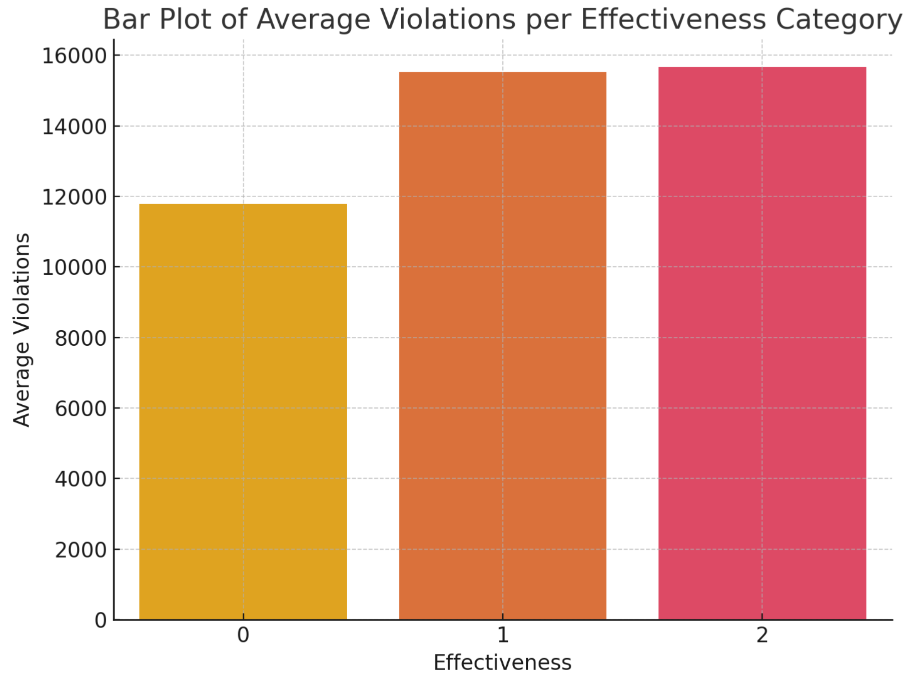

Figure 18: Bar Plot of Average Violations per Effectiveness Category

2. **Low Correlation Removal**: Features that had an absolute correlation of less than 0.05 with the target variable were removed. This is because such low correlation indicates that the feature does not have a significant relationship with the target variable and is unlikely to contribute useful information for predicting the target.

3. **High Correlation Removal**: Features with a correlation above 0.9 with other features were also removed to reduce multicollinearity. Multicollinearity occurs when two or more features are highly correlated with each other, leading to redundancy. This can cause issues in the model by making it difficult to determine the individual effect of each feature. For example, if *traffic_volume* and *log_traffic_volume* have a correlation coefficient of 0.95, they provide nearly identical information. In such cases, we keep one feature and remove the other. This step ensures that the model remains simple, interpretable, and less prone to overfitting.

By performing this correlation analysis and feature selection, we streamline the dataset, removing unnecessary features while retaining those that provide meaningful information for predicting the effectiveness of ASE locations.

3.3.2 Modeling and Evaluation

The goal of this study was to classify road segments into three categories of ASE effectiveness— Not-Effective, Moderately-Efective, and Highly-Effective— using their physical and traffic characteristics. This multi-class classification task required models that could handle diverse data distributions, capture non-linear relationships, and balance interpretability with predictive power. The selected classifiers—Random Forest, XGBoost, Logistic Regression, Support Vector Machine (SVM), and K-Nearest Neighbors (KNN)—were chosen based on their strengths in handling these requirements and their established performance in similar transportation and classification problems each described in detail in the appendix B.

1. **Random Forest Classifier**

   Why Chosen: Random Forest is an ensemble method that combines multiple decision trees to improve accuracy and reduce overfitting. It is robust to noisy data and can model non-linear relationships effectively. Its ability to quantify feature importance provides insights into how road and traffic characteristics influence ASE effectiveness.

   Strengths: Handles multi-class problems well, allows understanding of feature contributions, and performs reliably on medium to large datasets[49].

2. **XGBoost Classifier**

   Why Chosen: XGBoost is a gradient boosting algorithm known for its high efficiency and predictive performance. It is particularly effective at capturing complex patterns in the data, making it suitable for modeling non-linear interactions between road characteristics and ASE effectiveness.

Strengths: Provides excellent accuracy, has built-in mechanisms for handling class imbalance, and is highly scalable for large datasets [50].

3. **Logistic Regression**

   Why Chosen: Logistic Regression was selected as a baseline model due to its simplicity and the ease of understanding how features influence predictions. While it assumes linear relationships between features and the target, it provides a clear benchmark for comparison with more complex models.

   Strengths: Easy to interpret, computationally efficient, and provides probabilistic outputs, which are useful for understanding the certainty of predictions [51].

4. **Support Vector Machine (SVM)**

   Why Chosen: SVMs are effective for classification problems where the decision boundary between classes is non-linear. Using kernel functions, SVM can capture complex relationships, making it suitable for predicting ASE effectiveness categories.

   Strengths: Works well in high-dimensional spaces and is robust to overfitting, especially in smaller datasets [52].

5. **K-Nearest Neighbors (KNN)**

   Why Chosen: KNN was included for its simplicity and ability to capture local patterns in the data. Since the effectiveness of ASE can vary based on similar characteristics (e.g., road geometry or traffic volume), KNN's instance-based learning approach can identify these localized patterns effectively.

Strengths: Simple to implement, does not require a training phase, and can handle non-linear relationships naturally [53].

This diverse selection ensures robust and reliable predictions of ASE effectiveness, aligning with the study's goals to provide actionable insights into road segment characteristics.

However, as displayed in Figure12, the number of locations in each cluster is not distributed evenly, and only 13 locations belong to the highly effective cluster, which is called a class imbalance. Class imbalance occurs when some classes are underrepresented compared to others. This can lead to biased models that perform poorly on the minority class. To address the class imbalance in our dataset, we used the SMOTE and Tomek method, which combines two techniques:

- **SMOTE (Synthetic Minority Over-sampling Technique)**: This method generates synthetic examples for the minority class by interpolating between existing examples. This increases the number of minority class examples and helps the model learn better from them.

- **Tomek Links**: This under-sampling technique removes examples of the majority class that are close to examples of the minority class, helping to create clearer class boundaries and reduce overlap.

Combining these methods helps balance the class distribution by increasing the minority class and cleaning the majority class. To further address class imbalance and demonstrate the reliability of the models, we tested our machine learning models on five-fold cross-validation. Cross-validation is a technique for evaluating machine learning models by training several models on subsets of the available input data and evaluating them on the complementary subset of the data. In this work, we used 5-fold cross-validation, where the dataset is divided into 5 subsets, and the model is trained and evaluated 5 times, each time using a different subset as the evaluation set and the remaining subsets as the training set. This process is defined as follows:

$$\text{CV Accuracy} = \frac{1}{k}\sum_{i=1}^{k} \text{Accuracy}_i \tag{1}$$

where k is the number of folds, and Accuracy_i is the accuracy of the model on the i-th fold.

To evaluate the performance of the trained models, we used the following metrics:

1. **Accuracy**: The ratio of correctly predicted instances to the total instances. It gives a general measure of model performance but can be misleading if classes are imbalanced. The formula for accuracy is:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{2}$$

   where TP is the number of true positives, TN is the number of true negatives, FP is the number of false positives, and FN is the number of false negatives.

2. **Classification Report**: Provides detailed metrics for each class, including:

   - **Precision**: The ratio of true positive predictions to the total predicted positives. It indicates how many of the predicted positive instances are actually positive.

$$\text{Precision} = \frac{TP}{TP + FP} \tag{3}$$

   - **Recall**: The ratio of true positive predictions to the actual positives. It measures how many of the actual positive instances the model correctly identified.

$$\text{Recall} = \frac{TP}{TP+FN} \tag{4}$$

- **F1-Score**: The harmonic mean of precision and recall. It provides a single metric that balances both concerns.

$$\text{F1-Score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{5}$$

- **Support**: The number of true instances for each class in the test data.

3. **Confusion Matrix**: A table used to describe the performance of a classification model by comparing the actual and predicted classes. It shows the counts of true positives, true negatives, false positives, and false negatives, providing a comprehensive view of the model's performance. The structure of a confusion matrix is:

	Predicted Positive	Predicted Negative
Actual Positive	TP	FN
Actual Negative	FP	TN

(6)

3.3.3 Results

In this section, we present the results of our analysis, detailing the performance of various machine learning models in predicting the effectiveness of ASE locations. Our evaluation considered multiple classifiers, which were assessed using cross-validation with SMOTETomek for handling class imbalance. The key metrics used to gauge model performance were mean accuracy. The results for each model are summarized in Table 4.

Table 4: Performance of Machine Learning Models with SMOTETomek

Model	Mean Accuracy $\pm$ Std
LightGBM	0.66 ± 0.04
XGBoost	0.67 ± 0.05
Logistic Regression	0.47 ± 0.04
Support Vector Machine (SVM)	0.48 ± 0.06
K-Nearest Neighbors (KNN)	0.61 ± 0.02
Gradient Boosting	0.67 ± 0.02
RandomForest	$\mathbf{0.68 \pm 0.04}$

As shown in Table 4 the RandomForest Classifier emerged as the best-performing model, achieving 0.68 accuracy with a standard deviation of 0.04. This indicates that the RandomForest model consistently performed well across different folds of the cross-validation process, making it a reliable choice for predicting the effectiveness of ASE locations. The higher accuracy and lower standard deviation of the RandomForest model suggest that it has better generalization capabilities and is less likely to overfit compared to other models. Based on the cross-validation results, we selected the RandomForest Classifier as the best model. We then performed a final evaluation using cross-validation to confirm its performance on the original dataset. The classification report provides detailed metrics for each class, including precision, recall, and F1-score, as shown in Table 5.

Table 5: Classification Report for RandomForest Model

Class	**Precision**	**Recall**	**F1-Score**
Highly Effective	0.65	0.60	0.62
Moderately Effective	0.70	0.68	0.69
Not Effective	0.70	0.75	0.72

The classification report in Table 5 provides detailed insights into the model's performance across different classes. Precision, recall, and F1-score are critical metrics for understanding the effectiveness of the RandomForest model in correctly classifying ASE locations as 'Highly Effective,' 'Moderately Effective,' or 'Not Effective.'

3.3.4 Applying predictive Model on Vaughan Data

In the new experiment we applied our best model in previous step to the new test data from the City of Vaughan's data. The new dataset, which includes road and traffic metrics similar to those used in the Toronto dataset, was used to predict the effectiveness of ASE locations in Vaughan. The distribution of these predicted categories is illustrated in Figure 19.

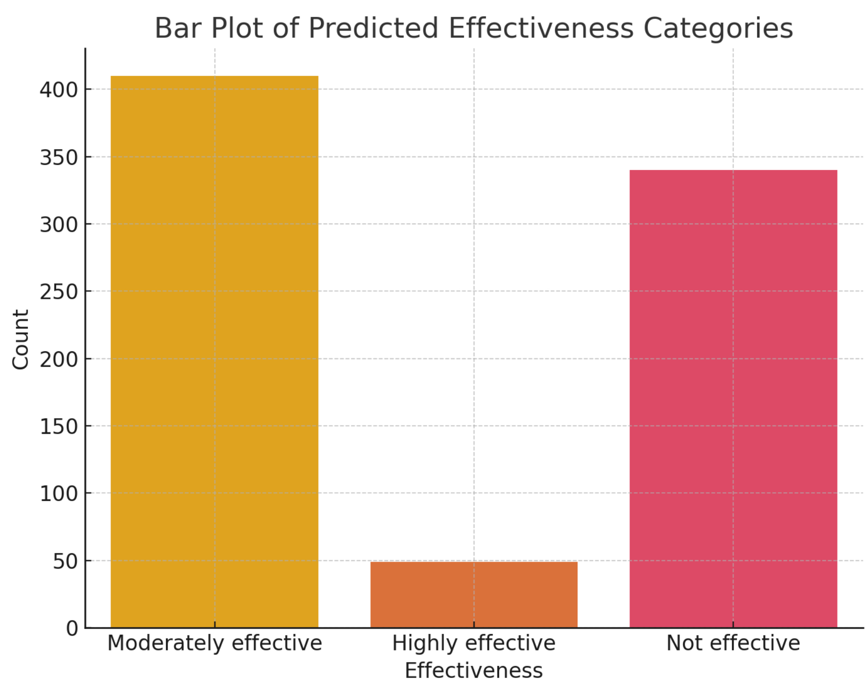

Figure 19: Bar Plot of Predicted Effectiveness Categories for City of Vaughan

These results indicate a similar pattern to our findings in the City of Toronto, with the majority of locations falling into the 'Moderately Effective' category. The presence of a significant number of 'Not Effective' locations highlights areas where enforcement strategies may need to be revisited to improve their impact. Conversely, the 'Highly Effective' category, although smaller in count, suggests that certain enforcement locations and strategies are performing exceptionally well.

4 Optimization Model

A clear and structured methodology is necessary to optimize the placement of speed cameras across multiple locations to minimize traffic violations. This section describes the approach to model and solve this complex decision-making problem under uncertainty. We use Markov Decision Processes (MDP), ideal for environments where outcomes are partly random and partly controlled by the decision-maker. In this model, randomness comes from changes in traffic violation rates at different locations, depending on whether a speed camera is present. The main goal is to reduce the overall number of traffic violations over a set period across various locations. With a limited number of speed cameras and a specific timeframe for their use, the challenge is to allocate these cameras effectively to lower violations. This section briefly reviews the literature on speed camera scheduling and then details how MDP is applied to solve this problem.

4.1 Model Framework

As mentioned previously, although it is possible to predict the probable number of violations in Number ommunity safety zone during enforcement, finding the optimal scheduling plan to maximize safety benefits remains challenging.

A comprehensive model is needed that considers the impact of enforcement over time in reducing collision risks . As previously mentioned, speed violations in this study represent safety concerns. We have a long-term T-epoch plan with limited resources (available cameras) and many sites (N community safety zones) to cover. We aim to maximize safety benefits during the enforcement at one site and across the entire planning period and subsequent planning cycles. Additionally, we lack the total network violation data; however, we could estimate the daily probability of violations for each enforcement month from predictions made earlier. This challenge is somewhat analogous to distributing emergency facilities across a network, where the optimal placement of facilities is crucial to minimize travel time during service demands.

Although our 'facilities' are fixed speed cameras, the objective is to position them in locations where they can cover the most speed violations across the network. This strategy aims to maximize the impact of ASE in reducing violations not only during a specific epoch, defined here as a unit of time equivalent to one month, but also throughout the entire planning period consisting of T epochs. In this context, T represents the total number of months over which the planning and optimization of camera placements occur. Oded Berman and Odeni [54] used Markovian processes to locate mobile servers on a network, exploring optimal server placements by assuming that travel time for each link could take a finite number of values. Consequently, the network could be in one of a finite number of states at any given time, dynamically transitioning from one state to another based on changes in travel times for links. Our approach similarly seeks to dynamically optimize the placement of speed cameras to best respond to fluctuating violation probabilities across the network. A Markov process, or Markov chain, is a fundamental concept in stochastic processes that models systems changing states over time with the property that the future state depends only on the present state and not on how the system arrived there. This property is known as the Markov property.

Markov processes can be classified into discrete-time and continuous-time Markov chains based on the timing of state transitions. In a discrete-time Markov chain, transitions between states occur at a set of discrete time points, while in a continuous-time Markov chain, transitions can occur at any time. The mathematical representation of a Markov process involves a set of states and the probabilities of transitioning from one state to another, known as transition probabilities. These probabilities are typically represented in a matrix form known as a transition matrix. Applications of Markov processes are vast and include areas such as queuing theory, population dynamics, economics, genetics, etc. They are instrumental in modelling systems where it is reasonable to assume that future behaviour only depends on the current state. In a Markov process (or Markov chain), one set of transition probabilities describes the likelihood of moving from one state to another. This set is typically represented as a matrix where each entry $p_{s,r}$ indicates the probability of transitioning from state s to state r in a one-time step. In a simple Markov process, these probabilities are fixed and depend only on the current state and not the decision-maker's actions. As a result, the Markov Decision Process (MDP) extends the idea of a Markov process by incorporating actions and rewards.

An MDP provides a mathematical framework for modelling decision-making where each decision influences immediate rewards and subsequent situations and decisions. An MDP is defined by several key components: 1- States, 2- Actions, 3-Transision Probabilities 4-Rewards and policies. the definition of each of these components in MDP frame work and in this problem are described in following sections.

4.1.1 States

A state represents each possible situation in the environment. The set of all possible states is called the state space. In practical applications, a state could represent anything from the current amount of money a company has to the robot's position in a room. In the ASE implementation problem, states play a crucial role in our analysis, specifically focusing on speed violations as a traffic safety measure. As stated earlier, the primary goal is to reduce the number of speed violence safety overall. Therefore, a set of states, X_1 through X_M, is defined, corresponding to levels of speed violations at various locations. These states are discrete and range from 1 (representing the highest number of violations to very low numbers of violations). Based on the predicted speed violations for each location, the corresponding state is assigned from this spectrum of M states. This classification allows us to systematically assess and respond to varying safety concerns across different locations.

4.1.2 Actions

For each state, multiple actions can typically be taken. The set of all possible actions available in a state is called the action space. Actions represent the choices available to the decision-maker, such as moving left or right in a maze, buying, selling, or holding stock in financial applications.In the MDP framework of this model, the actions determine the enforcement status of community safety zones at specific time epochs. Specifically, the action variable a_k(t) represents whether to enforce a safety zone i at time epoch t. This binary variable is set to 1 to indicate enforcement action is taken, and it is 0 otherwise. This simple yet effective binary designation allows us to model the decision-making process regarding the deployment of speed enforcement measures systematically and clearly.

4.1.3 Transition Probabilities

These are probabilities that define the dynamics of the environment. For each state and action, a transition probability determines the likelihood of transitioning to any other state. This component captures the uncertainty of decision-making, as it is not always sure what the exact outcome of an action will be due to the stochastic nature of the environment. In this model, transition probabilities are crucial for determining how the state of each location in terms of speed violations $\{X_1, \ldots, X_M\}$ changes from one epoch to the next, depending on the enforcement actions taken. These probabilities quantify how likely it is for a location i to transition from one violation state to another by the end of epoch t, given the action at the beginning of epoch t, denoted as a_k(t), Enforce or Non-Enforce. The transition probability matrix varies based on the action taken. If a_k (t) =1 (enforcement action is taken), the transition probability matrix is $P_{E,k}$; if a_k (t) =0 (no enforcement action is taken), the matrix is $P_{N,k}$. Both matrices are $M \times M$ in size, where each element $[P_{E,k}]_{sr}$ represents the probability of transitioning to state X_r at the end of epoch t if the initial state was X_s at the start of the epoch, under enforcement. Similarly, $P_{N,k}$ represents the transition probabilities during non-enforcement periods for location k.

4.1.4 Markovian Rewards

A reward function assigns a numerical reward for each action taken in a given state. The reward typically reflects the desirability of an outcome, guiding the decision-maker's actions toward long-term goals. For instance, in a navigation problem, a positive reward might be given for moving closer to a target, while a negative reward might be received for hitting an obstacle. In this problem, what is typically referred to as 'rewards' are conceptualized as penalties, aiming to minimize these penalties to maximize the overall benefit. The penalties in this model primarily depend on each location's state rather than the specific actions taken. Consequently, our 'rewards' are represented by a column vector $\mathbf{v}$, where each element of it $\mathbf{v}_s$ corresponds to the number of speed violations associated with that state$s$. This penalty value is constant for each state across all times, locations, and actions, reflecting the number of speed violations define each state.

4.1.5 Markovian Policy

A policy is a strategy that specifies the action to be taken in each state. It is a fundamental concept in MDPs, directly affecting the decision-making process. Most MDP scenarios aim to find an optimal policy that maximizes the expected return over time, often computed as the sum of discounted rewards received. The policy for addressing this problem will be defined in subsequent steps, following the mathematical notation and framework established in the earlier sections.

In the next section, we will present the mathematical model, considering the Markov Decision Process (MDP) components of our problem discussed in this section.

4.2 Mathematical Model Notations

We investigate the problem of optimizing the number of speed violations in various locations, which we divide into M potential violation states. The number of speed violations characterizes these states, and we represent each state by a violation number. There are N potential locations and P cameras for monitoring. Rewards amounts are represented by an $M \times one$ column vector matrix $\mathbf{v}$, where the r-th row corresponds to the number of violations for r. The probability of being at state r at time t for a given location k is denoted by $P_k(t)$. We introduce $[P_{E,k}]_{sr}$ and $[P_{N,k}]_{sr}$ as the probabilities of transitioning from state s to state r at location k if enforcement is present or absent, respectively. These probabilities are organized into $M \times M$ transition matrices $P_{E,k}$ and $P_{N,k}$ where the sum of each row is always 1, with enforcement at location k during epoch t determined by the binary variable a_k(t).

New: We defined P_k (t) to be the state probability vector for location k at the end of time t, which determines the probability of being in each state for location k at the beginning of time epoch t+1, this vector is solely dependent on the P_k (t-1) which is the probability state vector at the end of time t-1 equal to the probability state at the beginning of time epoch t and the action which is taken during time epoch t. which can be written as following:

$$P_k(t) = P_k(t-1) \cdot \left(P_{E,k}^{a_k(t)} . P_{N,k}^{(1-a_k(t))} \right) \tag{7}$$

based on the equation (7), we can determine the probability state vector of location k at the beginning of time t+1 to be equal P_k(t) and the probability state vector of location k at the end of time epoch t+1 is as following:

$$P_k(t+1) = P_k(t) \cdot \left(P_{E,k}^{a_k(t+1)} . P_{N,k}^{(1-a_k(t+1))} \right) \tag{8}$$

$$P_k(t+1) = P_k(t-1) \cdot \left(P_{E,k}^{a_k(t)} . P_{N,k}^{(1-a_k(t))} \right) \cdot \left(P_{E,k}^{a_k(t+1)} . P_{N,k}^{(1-a_k(t+1))} \right) \tag{9}$$

$$P_k(t) = P_k(t-1) \cdot \left(P_{E,k}^{a_k(t)} . P_{N,k}^{(1-a_k(t))} \right) \tag{10}$$

$$P_k(T) = P_k(0) \cdot \prod_{t=1}^{T} \left(P_{E,k}^{a_k(t)} . P_{N,k}^{(1-a_k(t))} \right) \tag{11}$$

$$V_k(T) = P_k(0) \cdot \mathbf{V} = P_k(0) \cdot \left[\prod_{t=1}^{T} \left(P_{E,k}^{a_k(t)} . P_{N,k}^{(1-a_k(t))} \right) \right] \cdot \mathbf{V} \tag{12}$$

The objective is to minimize violations and camera relocation costs over T time epochs by optimally allocating our P facilities to N locations, considering the maximum budget constraint B. We also consider cost K for relocating each camera and cost L for the constant cost of the relocation process regardless of frequency. Where $y_k(t)$ is a binary variable that determines whether the enforcement status of location k differs from the previous time epoch $t-1$.

$$y_k(t) = |a_k(t) - a_k(t-1)| \tag{13}$$

The objective function is

$$\begin{aligned}
\min \quad & \sum_{t=1}^{T}\sum_{k=1}^{N} P_k(0) \cdot \left[\prod_{s=1}^{t}\left(P_{E,k}^{a_k(s)} . P_{N,k}^{(1-a_k(s))}\right)\right] \cdot \mathbf{V} && && (14)\\
\text{s.t.} \quad & \sum_{k=1}^{N} a_k(t) = p && \forall t && (15)\\
& \sum_{t=1}^{T} w(t) + \frac{\sum_{k=1}^{N} y_k(t)}{2} \leq B && \forall t && (16)\\
& \sum_{k=1}^{N} y_k(t) \leq M_{Large} w(t) && \forall t && (17)
\end{aligned}$$

The constraint ensures the total number of active camerasNumbery given time is P. $W(t)$ is also a binary variable that defines whether or not the system experienced any relocation at time t.

Solving nonlinear optimization problems with nonlinear constraints presents significant challenges, primarily due to the non-commutativity of matrices involved. To manage this complexity, we have developed a schedule within a finite time frame consisting of c epochs. We assume the system reaches a steady state, employing a heuristic approach where the schedule is repeated indefinitely.

A heuristic is a problem-solving strategy that uses approximations or rules of thumb to find practical solutions, especially when exact methods are computationally infeasible. Unlike exact optimization methods, heuristics prioritize efficiency and manageability over guaranteeing mathematically optimal solutions, making them particularly suitable for large-scale, complex systems like ours. In this context, the heuristic simplifies the ASE scheduling framework by assuming a steady-state cyclic schedule, where enforcement and non-enforcement periods repeat consistently across epochs.

The heuristic approach offers several advantages. It reduces computational overhead by structuring the problem into repetitive cycles, enabling efficient implementation without solving the full non-linear optimization problem. By ensuring feasibility within the constraints—such as the number of active cameras and relocation limits—it provides a practical solution that balances accuracy with real-world applicability. Although heuristics do not guarantee optimality, they are effective for deriving good-enough solutions within a reasonable timeframe, which is essential in operational contexts. This methodology forms the foundation of our cyclic policy in the Markov Decision Process (MDP) framework, ensuring a stable state vector for each location as enforcement cycles progress.

In this framework, enforcement for each location commences at the start of the cycle and is potentially followed by a non-enforcement period. By consistently ordering enforcement and non-enforcement across the c-epoch cycles for each location, we significantly simplify the problem's complexity. Initially, the state probability for location k is denoted as $P_k(0)$. After undergoing χ_k epochs of enforcement followed by $c - \chi_k$ epochs of non-enforcement, the state probability $P_k(c)$ adheres to equation 18. Consequently, the state vector at the beginning of the subsequent cycle equals the result of equation 18. Considering the schedule will be repeated in the second cycle, with χ_k epochs of enforcement followed by $c - \chi_k$ epochs of non-enforcement, at the end of the second cycle, the probability state vector is calculated using equation 19.

Assuming $P_{E,k}^{\chi_k} \cdot P_{N,k}^{(c-\chi_k)} = A$, and establishing that matrix A qualifies as a transition probability matrix (refer to Appendix A), it follows that after h cycles, the state vector will stabilize at $P_k(h{\cdot}c)$. Given that matrix A possesses a steady state vector, there exists a vector $\hat{P}_k$ such that $\hat{P}_k \cdot A = \hat{P}_k$

. This vector, $\hat{P}_k$, represents the steady state vector for location k under a heuristic cyclic scheme with c epochs.

From these assumptions, it is deduced that at the start of the h-th cycle, in the long term, the state probability vector stabilizes at $\hat{P}_k$, and similarly at the beginning of cycle $h+1$ (or the end of cycle h), the probability vector remains $\hat{P}_k$, as detailed in equation 21. This methodology forms the cornerstone of our policy within the Markov Decision Process (MDP) framework.

$$P_k(c) = P_k(0).P_{E,k}^{\chi_k}.P_{N,k}^{c-\chi_k} = P_k(0).\mathbf{A}_k \tag{18}$$

$$P_k(2c) = P_k(c).P_{E,k}^{\chi_k}.P_{N,k}^{c-\chi_k} = P_k(c).(P_{E,k}^{\chi_k}.P_{N,k}^{c-\chi_k})^2 \tag{19}$$

$$P_k(h.c) = P_k(0).\mathbf{A}_k^h \tag{20}$$

$$P_k((h+1).c) = P_k(h.c) \tag{21}$$

The number of speed violations during one cycle in stable condition, for cycle h and after, at given location k, is calculated using equation 22. Considering the Distributive Property of matrixes, we can simplify the equation more and use equation 23:

$$V_k = \sum_{t=1}^{\chi_k} \hat{P}_k.P_{E,k}^t.\mathbf{V} + \sum_{t=1}^{c-\chi_k} \hat{P}_k.P_{E,k}^{\chi_k}.P_{N,k}^t.\mathbf{V} \tag{22}$$

$$V_k = \hat{P}_k \left[\sum_{t=1}^{\chi_k} P_{E,k}^t + P_{E,k}^{\chi_k} \sum_{t=1}^{c-\chi_k} P_{N,k}^t \right].\mathbf{V} \tag{23}$$

Consequently, the objective function for this heuristic problem will be equal to the objective function to . It is noteworthy that constraints will be updated. First of all, ensuring that at most only one-time location will be enforced during one cycle, we assumed that the relocation and maintenance costs are low enough compared to the budget to be neglected. Also, constraint 15 that made sure to consider the limitations of available cameras is transformed to constraint 26 to make sure total enforcement periods in the network won't exceed the total enforcement periods that the network could have in a c-epoch duration. Moreover, constraint 25 represents the enforcement duration for each location for a c-epoch scheduling plan.

$$\min \quad \frac{T}{c}\sum_{k=1}^{N}\hat{P}_k\left[\sum_{t=1}^{\chi_k}P^t_{E,k}+P^{\chi_k}_{E,k}\sum_{t=1}^{c-\chi_k}P^t_{N,k}\right].\mathbf{V} \tag{24}$$

$$\text{s.t.} \quad \chi_k \leq c \qquad \forall k \tag{25}$$

$$\sum_{k=1}^{N}\chi_k \leq cp \tag{26}$$

4.3 Model evaluation

4.3.1 Scenario definition

To assess the effectiveness of our model, we utilized Matlab's `GA` function. GAs are inherently suited for nonlinear problems as they do not require gradient information, which can be particularly advantageous in scenarios where the objective function is complex or not differentiable. The linear constraints are efficiently handled through a penalty approach integrated within the GA, ensuring that solutions conform to the constraints without compromising the algorithm's ability to explore the solution space extensively.

Our initial step involved simplifying the analysis to a scenario with just two locations and one surveillance camera. In this setup, we considered two distinct states of traffic violations—high and low—resulting in each area being represented by a state probability vector as a 1x2 matrix. The matrices for enforcement and non-enforcement states are configured as 2x2 triangular matrices, reflecting transitions between these violation states. For the enforcement scenarios, we used an upper triangular matrix (`PE,k`), indicating that a transition from a low to a high violation state under enforcement is deemed impossible, thus assigning a zero probability. Conversely, the non-enforcement state matrix (`PN,k`) is a lower triangular matrix based on the premise that shifting from a high to a low violation state without enforcement is equally unlikely.

The analysis involved pairwise comparisons across three distinct location types, each characterized by unique `PN,k` or `PE,k` matrices. The baseline location Type 1 has `PE,k` matrix values of `[.3 .7; 0 1]` and a `PN,k` matrix of `[1 0; .6 .4]`, setting a foundation for comparison. We then explored how slight adjustments in the `PE,k` matrix elements could affect decisions on camera placements. Although location Type 2 shares the same `PN,k` as Type 1, it displayed different propensities for transition during enforcement periods. Specifically, location Type 2, with `PE,2 = [.2 .8; 0 1]`, showed a higher likelihood of transitioning from a high to a low violation state than location Type 1. Location Type 3, mirroring location Type 1 in terms of its `PE,k` matrix, differentiates itself through a unique `PN,k` matrix of `[1 0; .8 .2]`. This variation in the `PN,k` matrix suggests that location Type 3, despite similar enforcement dynamics as Type 1, experiences a substantially different behaviour under non-enforcement conditions, potentially affecting the overall efficacy of the non-enforcement measures.

The core of our evaluation focused on how these different transition matrices influence the allocation of enforcement time across location types. By conducting pairwise comparisons—starting with each location type against the baseline Type 1—we aimed to determine the optimal enforcement time distribution. This was achieved by incrementally adjusting the enforcement time for

Type 1 and observing the total number of violations acrNumbere network. The precision of our optimization model was then validated by comparing the ideal pair of enforcement times, derived from each pairwise analysis, against the enforcement times calculated by the `fmincon` model. This methodology not only highlights the nuanced approach required to optimize enforcement strategies but also demonstrates the capability of our model to adapt and identify efficient allocations of enforcement resources based on the probabilistic behaviour of traffic violations across different locations.

Location Type 1 serves as our foundational benchmark for comparison. This type and Type 2 share the same matrix for non-enforcement periods, albeit with slight variations in their transition matrices during enforcement. Specifically, location Type 2 is inclined towards transitioning to lower violation states more effectively than location Type 1, indicating a better response to enforcement measures. On the other hand, while location Type 3 shares the same enforcement transition matrix as location Type 1, it exhibits differences in its non-enforcement transition behaviours. Location Type 3 displays a lower tendency to stay in a low-violation state and a higher probability of reverting to high-violation states without enforcement. This delineation underscores the nuanced dynamics of violation state transitions under varying enforcement conditions.

4.3.2 Different Location Types

Our analysis methodically compared each location type against the baseline (Type 1) across several pairs: 11, 12, and 13. In each of these scenarios, we considered only two locations, one from each type, over a total enforcement cycle of nine months, with a single camera allocated for monitoring. The enforcement initiative commences with location Type 1, a factor determined to have no significant impact on the study's outcome, as will be detailed subsequently. We explored ten timing scenarios for each pair of locations, aiming to quantify the total number of violations at Numbero locations across varying enforcement durations—ranging from 0 to 9 months for location Type 1. This approach allowed us to ascertain the optimal distribution of enforcement time between location Type 1 and its counterpart in each pair. The objective was to delineate an approximate optimal enforcement period that effectively balances the reduction of violations across different location types relative to our baseline (Type 1).

The following section will examine these comparisons more closely, providing a detailed breakdown of the enforcement time scenarios for each pair of locations. This analysis aims to clarify how enforcement resources can be strategically allocated to maximize the effectiveness of ASE measures in reducing traffic violations on different road segments.

To conduct a comparative analysis between location types 1 and 2, we explored ten potential time distribution scenarios, such as [0,9], [1,8], ..., and [9,0]. In these time scenarios, the first element of each array represents the enforcement time at location Type 1. In contrast, the second element denotes the enforcement time at the paired location, which could be Type 2, 3, or an additional instance of Type 1. Given the constraint of a single camera shared between two locations, the sum of enforcement times is fixed at nine months (c=9), ensuring exclusive camera presence at one location at any given moment. For example, if the camera is initially installed at location Type 1 at the start of the cycle (t=0) and later relocated to location Type 2 at time x , the remaining time y is calculated as = c - x, with the camera relocation time assumed to be negligible.

To determine the total number of violations for[x,y] time scenario We rigorously calculated the number of speed violations the network will record in one cycle for each time-location pair scenario. Figure 20 displays these results for all three location scenarios, with distinct colours assigned to each scenario to facilitate easy differentiation. As observed, the combination of location types

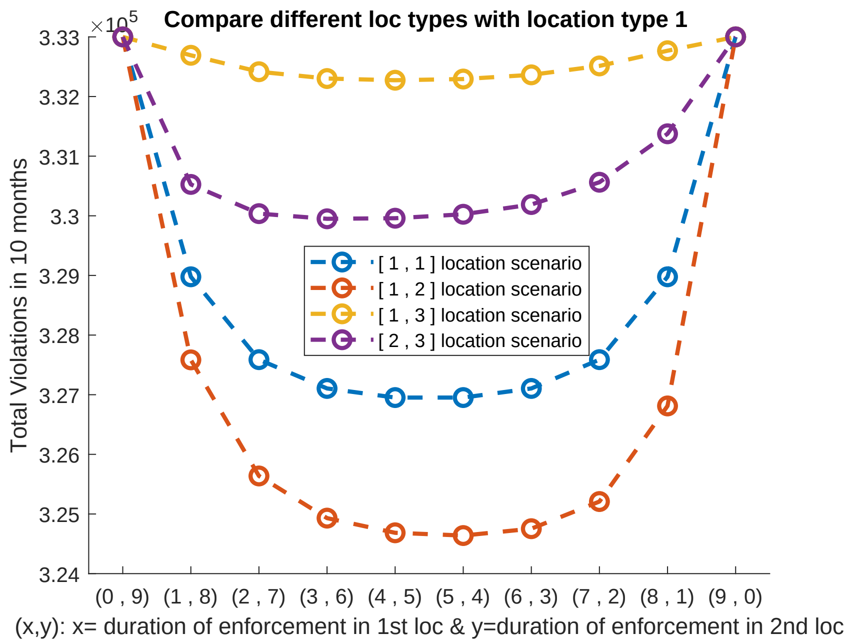

Figure 20: High and Low violation state probability for location types 1 and 2

1 and 2 consistently shows the lowest number of speed violations across all time scenarios. In contrast, the combination of types 1 and 3 exhibits the highest number of violations in the time scenario. This outcome aligns with expectations, given that location Type 2 demonstrates a stronger propensity to transition to lower violation states during enforcement and non-enforcement periods than location Type 3.

Figure 20 reveals that the three scenario lines are somewhat symmetrical. The scenario pairing location types 1 and 1 is perfectly symmetrical, indicating a balanced enforcement strategy. The scenario involving types 1 and 2 tends towards more enforcement towards location Type 1, which may initially seem counterintuitive. Commonly, one might expect a greater enforcement focus on location Type 2 due to its tendency to transition to a lower violation state when enforced. However, delving deeper into the dynamics of the problem, it becomes evident that location Type 1 requires prolonged enforcement to mitigate its slower transition towards low violation states effectively. As a result, the optimal enforcement strategy achieves a balance—intensifying enforcement on location Type 2 to leverage its responsive nature to enforcement measures while bolstering enforcement for location Type 1 to offset its sluggish transition dynamics. Consequently, the time scenario [5,4] registers the lowest number of violations, numbering it as this pair's most effective enforcement configuration. In scenarios involving location types 1 and 3, there is a noticeable trend towards allocating more enforcement time to location Type 3 and less to location Type 1. This strategy stems from the model's preference to minimize non-enforcement periods for location Type 3 due to its tendency to revert to high violation states. Conversely, location Type 1, which exhibits a lesser inclination to maintain high violation states during non-enforcement periods, receives less enforcement time. This differential enforcement strategy highlights the nuanced understanding required to optimize resource allocation effectively based on the distinct behavioural patterns of each location type. The model prefers assigning approximately equal enforcement periods for scenarios where both locations are of Type 1. This symmetrical enforcement approach indicates the model's strategy for evenly distributing enforcement efforts when dealing with homogeneous location types. Subsequent sections will delve deeper into how the model allocates the camera in such scenarios, providing a granular understanding of the enforcement preferences and their implications for traffic violation mitigation. When comparing location types 2 and 3, our analysis explored the ten timing scenarios to understand their impact on violation rates. The plot reveals that location Type 2 requires more consistent enforcement due to its high transition rate to low violation states under enforcement. In contrast, location Type 3's higher propensity to revert to high violation states necessitates focused enforcement during critical periods. The strategy here is to balance enforcement by allocating sufficient time to location Type 2 to maintain its low violation state while concentrating efforts on location Type 3 to prevent rapid escalation in violation rates. This nuanced approach underscores the need for dynamic enforcement allocation based on the specific characteristics and behaviors of each location type. The results from this scenario further reinforce the importance of adaptive enforcement strategies tailored to the unique needs of different location types to maximize the effectiveness of ASE measures.

4.3.3 Violation distribution over one cycle

Now we can give more details on how we calculated the number of violations for each location each time for any location and time scenario. We first calculated the violations for each location using 23, which requires each location's state transition probabilities $\hat{P}_k$, $P_{E,k}$, $P_{N,k}$, the cycle length (c), and the enforcement periods x and y. Finding P_k for each location involves solving P_k . $(I\text{-}{P_{E,k}}^{\chi_k}.{P_{N,k}}^{c-\chi_k}) = 0$ with Matlab, which could introduce challenges since our model assumes that a location's enforcement phase precedes its non-enforcement phase within a cycle. However, for our scenarios, the second location is enforced in the latter part of the cycle. Addressing this requires understanding that our model is heuristically solved, assuming enforcement cycles have been ongoing indefinitely. The enforcement for each location begins when its state probability vector matches P_k for that location, denoted as time t=t*. The state probability updates over time based on the elapsed time since t*, with the state probability vector calculated differently depending on whether t-t* is within the enforcement period. For locations 1 and 2, we plotted expected violations over time, considering endless successive cycles. Since the P_k for locations is independent, we can align location Type 2's cycle with location Type 1's without losing generality. Thus, we shift location Type 2's violation curve right by location Type 1's enforcement duration, as shown in Figure 21a. This figure illustrates the predicted violations by summing each location's violations over time, with the total violations equal to the area under the green curve, a sum of the areas under the blue (location Type 1 violations) and red (location Type 2 violations) curves. The shift in location

Type 2's curve does not alter the total violation count but may affect the distribution of predicted violations at each point.

Through this analysis, particularly with the scenarios for locations 1 and 2 illustrated in Figure 20 , we observed that the [5,4] time scenario results in the lowest total number of violations. Figures 22b and 21c depict the changing probability of being in each state for both locations over the cycle for location types 1 and 2, respectively. The left axes represent the probability of being in the high-violation state, while the right axes show the low-violation state probability. To enhance visualization, the areas under each curve are color-coded, with the high-state area shaded distinctly from the low-state area.

Figures 22b and 21c illustrate the probabilities of being in high and low violation states over time for location types 1 and 2, respectively. In location Type 1, the probability of high violation states decreases as enforcement begins at time 0. This downward trend persists until Epoch 5, when enforcement mechanisms are implemented at this location. However, the rate of decrease becomes more gradual as time progresses. Once enforcement is removed from location Type 1, the probability of high violation states starts to rise again, although the increase is less steep as the cycle advances. Conversely, for location Type 2, the probability of high violation states increases at the beginning of the cycle from its lowest probability in a high violation state to the highest probability of being in a low violation state. This trend is expected, considering that the cycles repeat and the enforcement mechanism had been active at location Type 2 at the end of the previous cycle. The probability rises and flattens at its peak around epoch 5, indicating a steady-state condition. At this time, the enforcement mechanism is installed at location Type 2, leading to a subsequent decrease in the probability of high violation states. As the cycle concludes, location Type 1 approaches its steady state, and the enforcement mechanism is rotated back, initiating the next cycle. The same

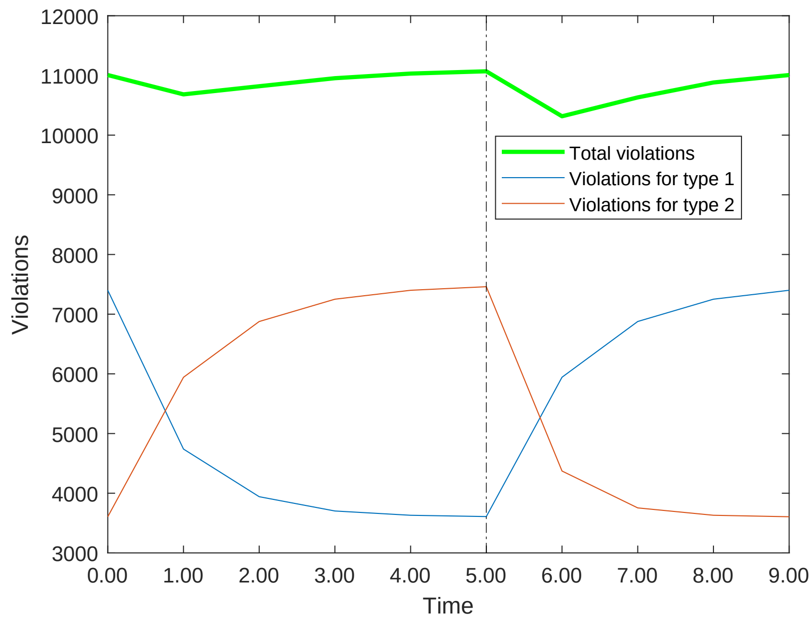

(a) Speed violation distribution over time for each location and for the whole network

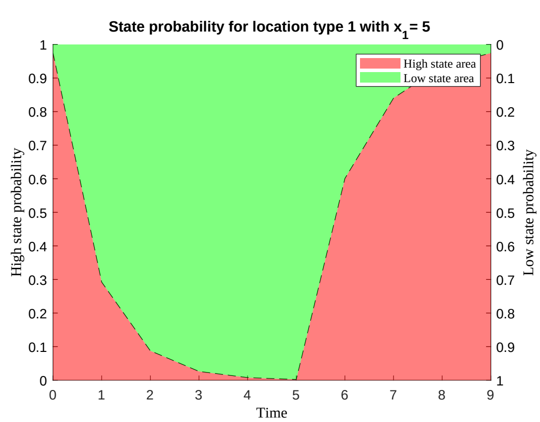

(b) probability state vector for location Type 1

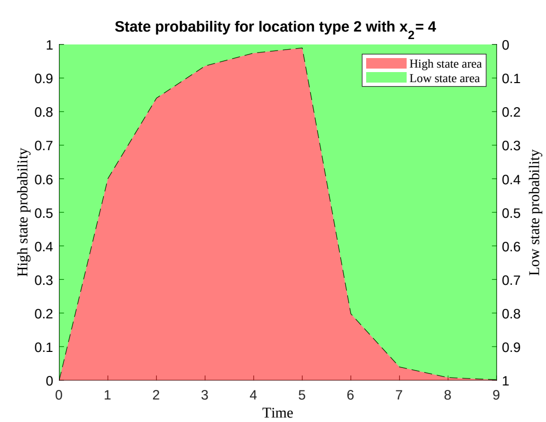

(c) probability state vector for location Type 2

Figure 21: The performance of the network for optimum distribution of camera between location types 1 and 2 during one cycle

pattern is observed when comparing location types 1 and 3 and location Type 1 against itself across different cycles. The consistent behavior across these comparisons highlights the predictable nature of enforcement's impact on violation probabilities. This cyclical enforcement strategy effectively manages the probability of violations by periodically redistributing enforcement resources, thereby maintaining a balance in compliance rates across various locations. These findings underscore the dynamic impact of enforcement mechanisms on violation probabilities, revealing distinct patterns for different locations and phases within the enforcement cycle.

Figure 24a illustrates a scenario involving three locations, each representing a different type of location. The enforcement strategy begins at location Type 1, where the enforcement mechanism is active for three epochs, then moves to location Type 2 for two epochs, and finally to location Type 3 for four epochs. The figure displays the number of speed violations at each location over a complete cycle, with each location's data represented by a distinct color. The total number of violations over time is shown in green. The blue curve, representing location Type 1, shows a decrease in violations at the start of the cycle, reaching a minimum at the end of the third epoch. Following the removal of enforcement at this location, the number of speed violations increases, eventually peaking as the enforcement is relocated. Simultaneously, the red curve, representing location Type 2, starts with a relatively high number of violations, displaying a slight positive slope close to zero until epoch three. At this point, enforcement is applied, resulting in a reduction in violations over the next two epochs. Once the enforcement is moved to location Type 3, the number of violations at location Type 2 begins to rise again. For location Type 3, represented by the yellow curve, the number of violations is relatively low at the beginning of the cycle. This pattern can be attributed to the repetitive enforcement schedule, where the camera was present at location Type 3 at the end of the previous cycle. After the removal of the camera, the violations initially increased

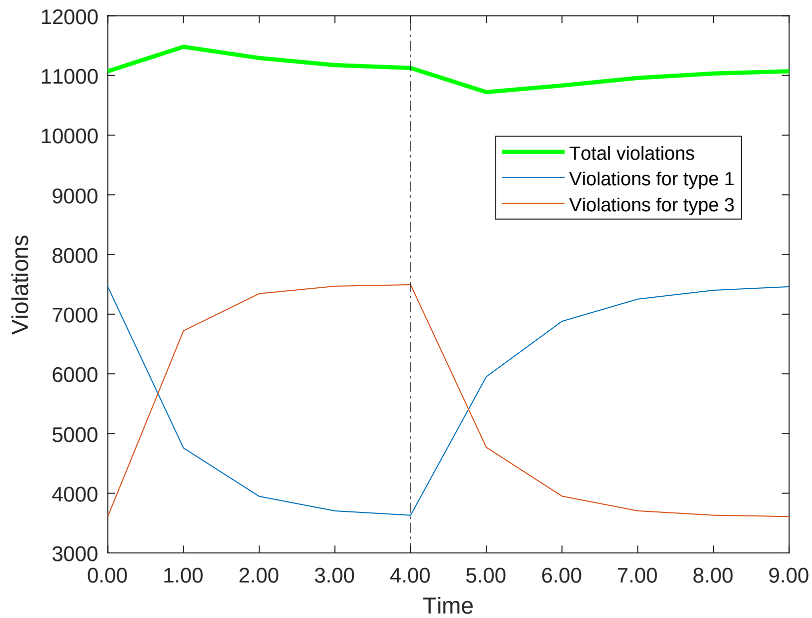

(a) Speed violation distribution over time for each location and for whole network

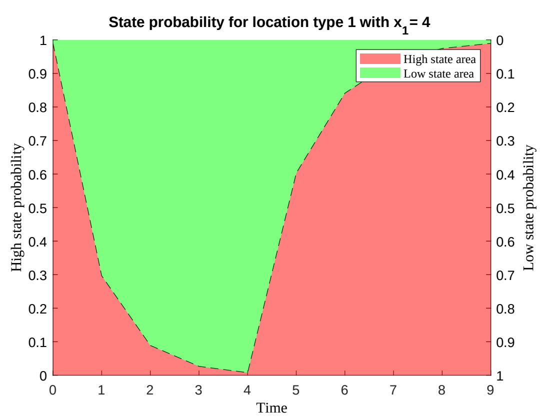

(b) probability state vector for location Type 1

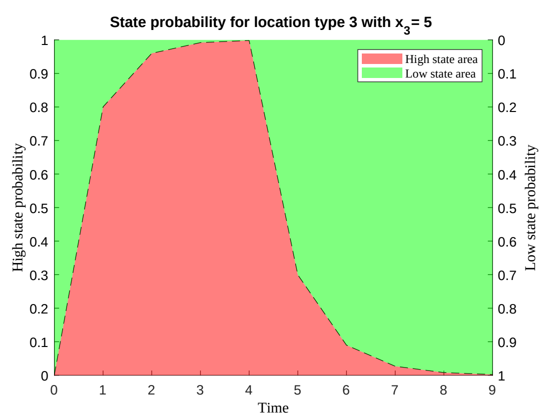

(c) probability state vector for location Type 3

Figure 22: The performance of the network for optimum distribution of camera between location types 1 and 3 during one cycle

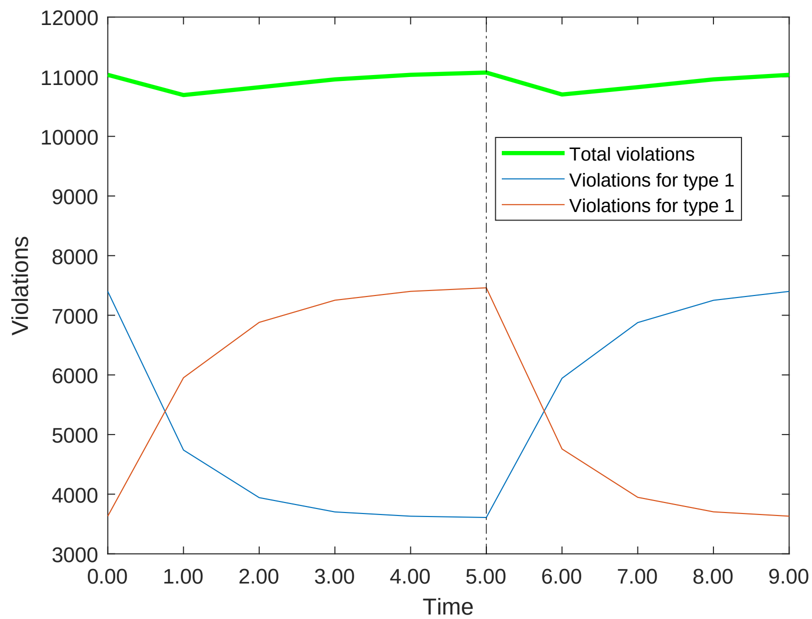

(a) Speed violation distribution over time for each location and for whole network

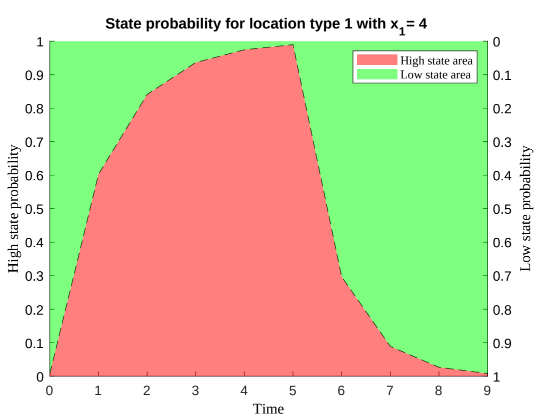

(b) probability state vector first location with $x_1 = 5$

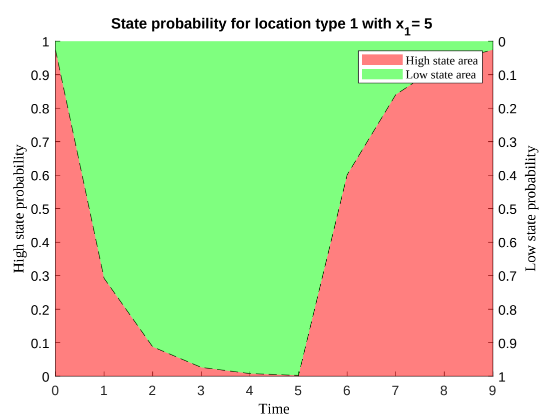

(c) probability state vector for second location with $x_1 = 4$

Figure 23: The performance of the network for optimum distribution of camera between location types 1 and 1 during one cycle

until epoch five. The rate of increase then flattens, and from epoch five onwards, the number of violations starts to decrease as the camera is reinstalled at this location, continuing to decline until the end of the cycle.

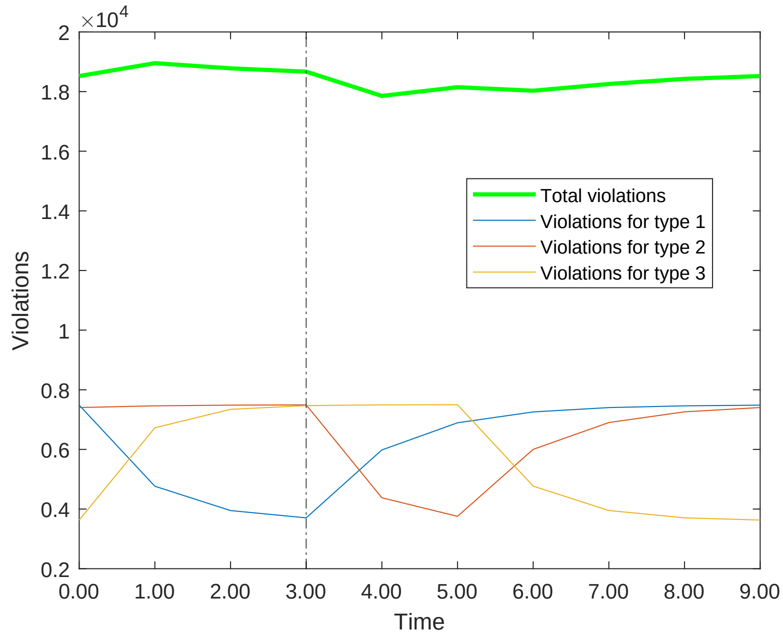

(a) Speed violation distribution over time for each location and for whole network

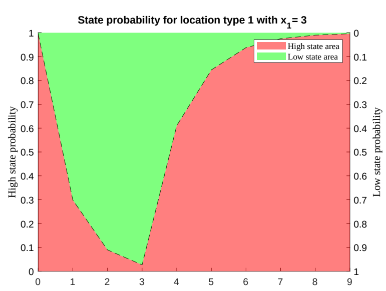

(b) probability state vector for location Type 1

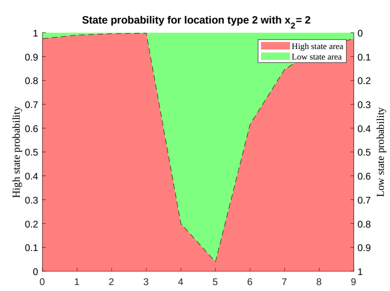

(c) probability state vector for location Type 2

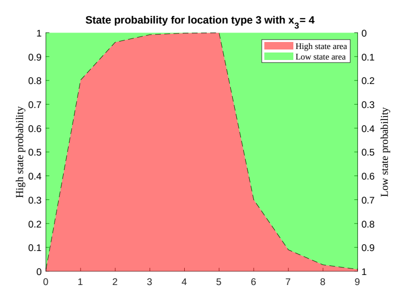

(d) probability state vector for location Type 3

Figure 24: The performance of the network for optimum distribution of camera between location types 1 2 3 during one cycle

In Figure 24b, the high violation state probability for location Type 1 decreases at the beginning of the cycle, continuing until time epoch 3. This decrease coincides with the initial enforcement period at this location. However, once the enforcement camera is uninstalled, the probability of high violations starts to increase, eventually stabilizing at a steady-state level by the end of the cycle. Figure 24c illustrates the scenario for location Type 2. Initially, there is a high probability of being in the high violation state, which gradually increases until epoch 3, reaching its peak. This increase is attributable to the lack of enforcement at this location during the early stages of the cycle. As soon as the enforcement camera is installed at time epoch 3, a marked decrease in high violation state probability is observed. This decrease continues until the camera is removed at epoch 5, after which the probability begins to rise again. For location Type 3, as shown in Figure 24d, the high violation state probability is low at the beginning of the cycle. This initial low probability aligns with the enforcement camera being present at the end of the previous cycle. As time progresses and the camera is absent, the probability of high violations increases, though the rate of increase slows, eventually flattening. At time epoch 5, the high violation state probability reaches its maximum and then decreases once the camera is reinstalled at this location, continuing to decline until the end of the cycle. The following section aims to gain a deeper understanding of how the objective function changes with varying values of c, M, and N. We assume $r_i = \frac{c \cdot M}{N}$ as the average enforcement period for each location in the network. While the model may not distribute the enforcement period equally due to certain preferences discussed in Section 20, this ratio provides a useful metric for comparing different scenarios.

4.3.4 Density of location types in Network

This section aims to evaluate the model's performance across different densities of location types 1 and 2. We focus specifically on these location types due to the limitations of our available real-life data, which lack detailed information regarding the enforcement period. Consequently, we assume the same non-enforcement matrix for all location types, allowing us to isolate and analyze the effects of location Type 1 and 2 densities, with different enforcement probability matrixes and the same non-enforcement matrixes, on the model's outcomes.

In the first set of scenarios, we consider a constant planning period of $c = 10$ epochs, with different ratios of $\frac{M}{N}$. Figures 25a, 25c, and 25e display three curves corresponding to networks with different M and N but the same M/N in each plot. Figure 25a illustrates the scenario with a constant planning period of $c = 10$ epochs and an M/N ratio of 0.1. The figure presents three curves corresponding to networks with 10 possible sites and 1 camera, 20 road segments and 2 cameras, and 100 road segments and 10 cameras. We observe a general decreasing trend in the total number of speed violations as the density of Type two locations increases. This trend is consistent across all networks, regardless of the M/N ratio. However, networks with higher values of N exhibit a greater total number of violations, which is logical given that more road segments provide a greater potential for violations.

To normalize the impact of varying numbers of road segments on the total number of violations, we introduce a parameter α_i for each scenario i. This parameter is defined such that $\alpha_i \cdot N_i = 100$. For example, in the first scenario with $N = 10$, $\alpha_1 = 10$, while in the second scenario with $N = 20$, $\alpha_2 = 5$. By multiplying α_i with the violations in each network, we obtain a normalized comparison of the results. The normalized plot, shown in Figure 25b, reveals that the trend in violations relative to the density of location types remains nearly identical across different networks. This consistency suggests that the primary factor influencing the number of violations is the M/N ratio, rather than the absolute number of cameras or road segments. Figure 25c,25d and 25eand 25f further support these findings with $M/N = 0.2$ and $M/N = 0.5$, respectively, each showing similar trends in violation rates across networks of different sizes. The parallel trends in these figures indicate that the enforcement strategy's effectiveness is proportional to the M/N ratio, independent of the absolute size of the network.

In the next step, we aimed to understand the impact of increasing the number of cameras within a single network. Figure 26a presents the total number of violations for a network with 100 road segments (N=100), showing data for configurations with 1, 5, 10, 20, and 50 cameras. This analysis provides insights into how the trend in violations might change if the number of available cameras increases in the future.A general decrease in the total number of network violations is observable as the distribution of Type 2 locations increases. Moreover, as the number of cameras increases, there is a corresponding decrease in total violations across the network. This trend suggests that enhancing enforcement coverage by adding more cameras effectively reduces the number of speed violations. The decline in violations appears consistent and parallel across the various configurations with the same number of road segments, regardless of the number of cameras. This consistent pattern highlights the significant role that camera coverage plays in managing and reducing violations. As more cameras are deployed, they likely create a stronger deterrent effect, leading to fewer violations. This effect is observable across different network sizes, as similar results were noted for networks with 10 and 20 road segments displayed in Figures 26b and 26c, as previously discussed.

In all previous scenarios, the duration of the cycle was assumed to be fixed and equal to 10 epochs. As the final step of this section, we tested how the model performs with different types and densities for various cycle durations c in a 10-road segment network with 1 camera. For $c = 5, 10$, and 20, the speed violations were calculated, as shown in Figure 27a. Consistent with previous results, there is a decreasing trend in the number of speed violations as the density of location Type 2 increases in the network, regardless of the cycle duration. To accurately evaluate the model's performance across different c values, it was necessary to normalize the results. We defined β_i, which assumes

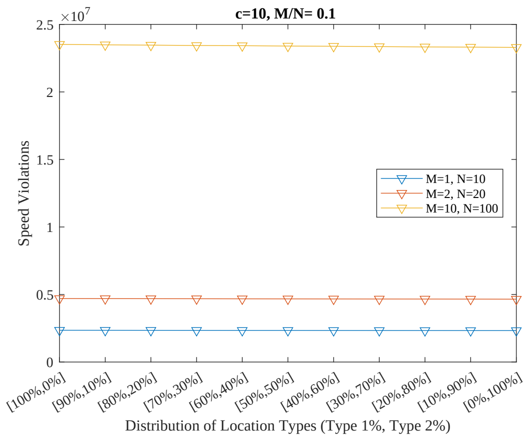

(a) Parallel trend for M/N=0.1

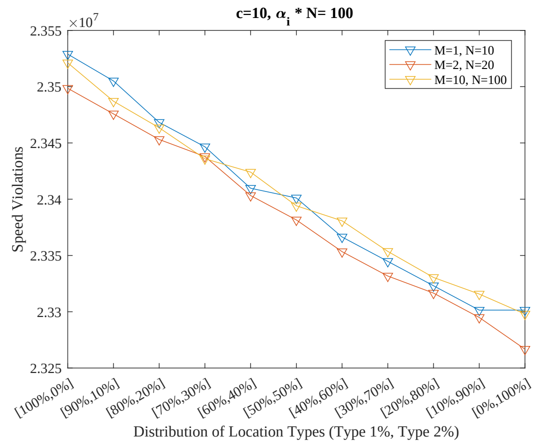

(b) While $\alpha_i \cdot N = 100$

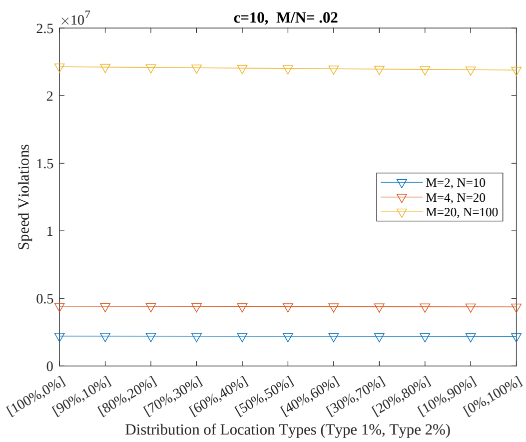

(c) Parallel trend for M/N=0.2

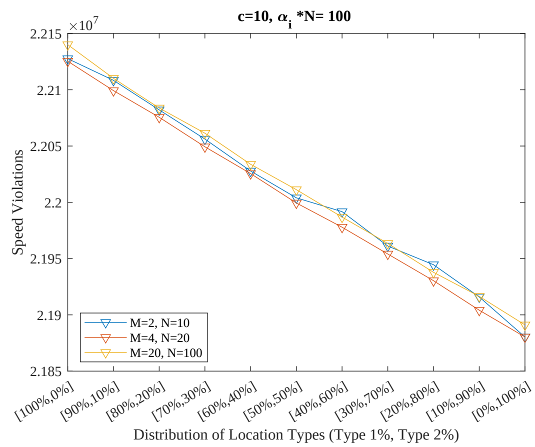

(d) While $\alpha_i \cdot N = 100$

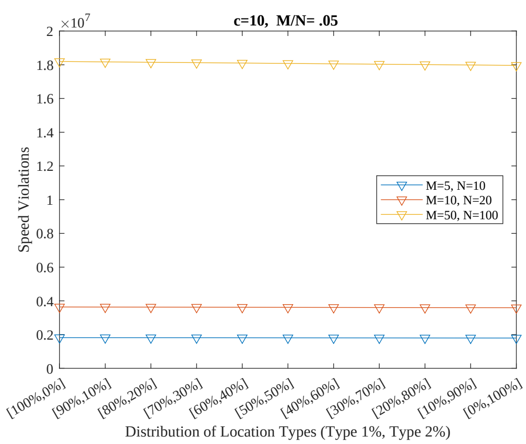

(e) Parallel trend for M/N=0.5

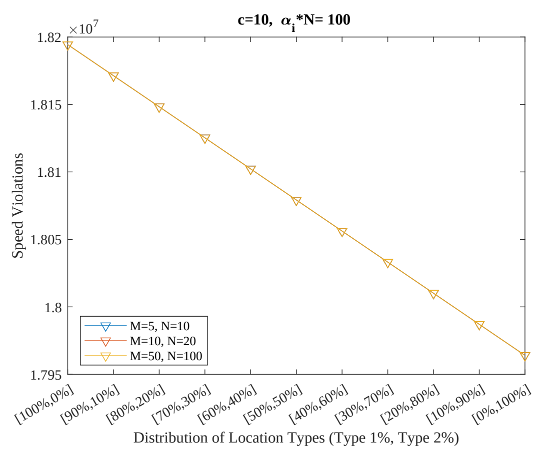

(f) While $\alpha_i \cdot N = 100$

Figure 25: Comparison of networks with 100, 20, and 10 site locations for different M/N ratios and $\alpha_i \cdot N = 100$.

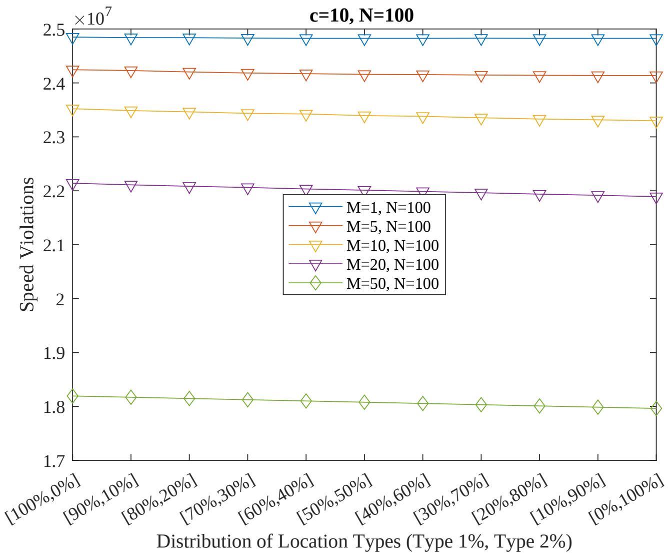

(a) Compare calculated speed violation for a 100-road-segment network for different numbers of cameras for different location type distributions

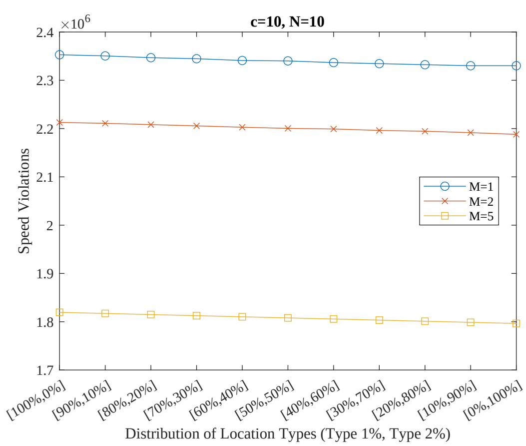

(b) Compare calculated speed violation for a 10-road-segment network for different numbers of cameras for different location type distributions

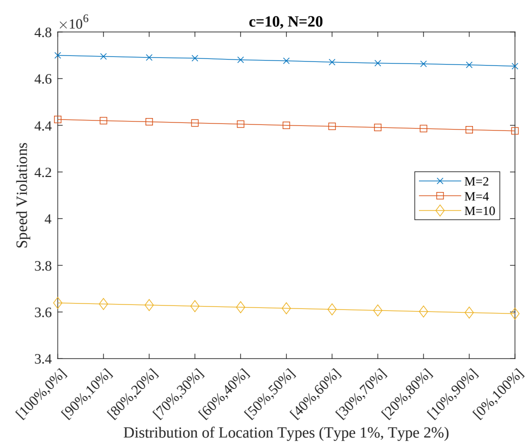

(c) Compare calculated speed violation for a 20-road-segment network for different numbers of cameras for different location type distributions

Figure 26: How the model optimizes different networks with different numbers of resources

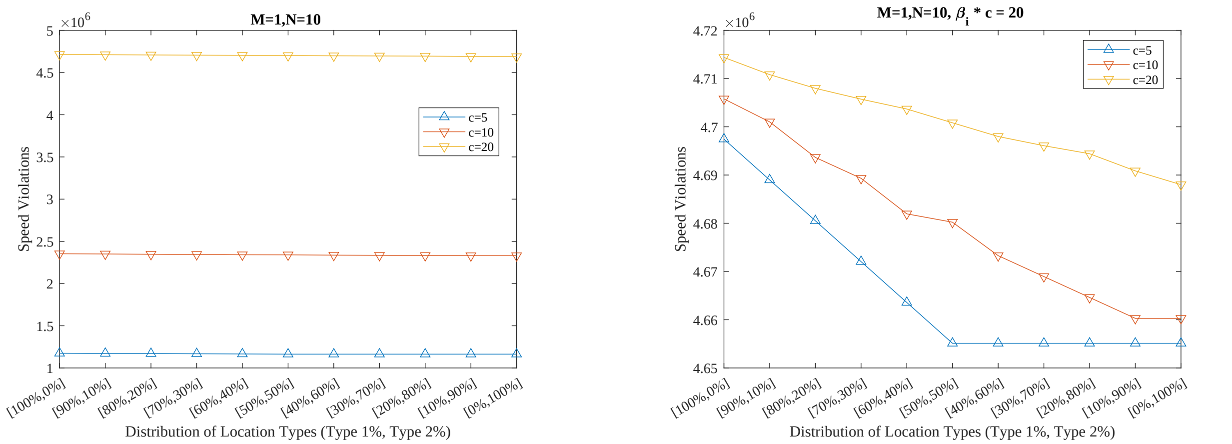

(a) Compare calculated speed violation of the model for a 10 road segment network with 1 camera, for different cycle duration

(b) Normalized speed violations for different c durations

Figure 27: Evaluate speed violation of the network for varied cycle duration and different network location type combinations

that $\beta_i \cdot c_i = 20$. As a result, for the first scenario, $\beta_1 = 4$; for the second, $\beta_2 = 2$; and for the third, $\beta_3 = 1$. The violations in each scenario were then multiplied by their respective β values. The normalized results are presented in Figure 27b. Although the normalized violations are closer in value, the data still shows a higher number of violations for longer cycle durations. This suggests that the duration of the cycle has a significant impact on the model's performance, and further investigation into the effects of cycle duration will be necessary in the next section.

4.3.5 Analysis of Scenarios with different c duration

In this section, the performance of the model under varying cycle durations (c) and different numbers of cameras (M) is evaluated. The previous analysis concluded that as the density of Type 2 locations in the network increases, the number of speed violations decreases, displaying an almost linear relationship between these factors. For simplicity and to reduce computational complexity, this section assumes that all networks consist entirely of Type 2 road segments. As illustrated in Figure 28, we examined speed violations for networks with varying cycle durations, specifically

$c = 1, 2, \ldots, 10, 15, 20$, and different numbers of cameras. It is observed that increasing the number of cameras generally leads to a reduction in speed violations, a finding consistent with earlier research. However, due to the varying study periods across different c values, there is an observable increase in the total number of violations as c increases. This trend necessitates a normalization to isolate the effect of c itself on the results. To address this, normalization factor, γ_i, was introduced for each cycle duration, ensuring that $\gamma_i \cdot c = 10$. Consequently, for scenarios with $c = 1$, the speed violations are multiplied by 10; for $c = 2$, by 5; and so forth, with $c = 20$ being multiplied by 0.5. This normalization allows us to focus solely on the influence of c on the network's speed violations, removing the confounding effect of different study periods. Figure 29 depicts the normalized data, revealing that each camera scenario exhibits a distinct minimum point. Specifically, for $M = 50$, the optimal c is 2; for $M = 20$, the optimal c is 4; and for $M = 10$, the optimal c is 8. This pattern indicates that as the number of cameras increases, the model favors shorter cycle durations, suggesting that more frequent rotations are beneficial in higher surveillance environments.

To evaluate whether the M/N ratio influences the optimal c duration, the model was tested on two additional networks with 10 and 20 road segments. The results, displayed in Figures 30a and 30, demonstrate a similar trend: as the M/N ratio increases, there is a higher tendency for the network to relocate the cameras more frequently. Additionally, the results indicate that networks with similar M/N ratios tend to have nearly the same optimal c duration values, highlighting a consistent relationship between camera density and the frequency of relocations. The findings from this study highlight the importance of both the number of cameras and the duration of their deployment cycles in optimizing traffic enforcement strategies. The model's performance indicates that with more cameras, shorter cycle durations are more effective in reducing speed violations.

The findings from this study highlight the importance of both the number of cameras and the duration of their deployment cycles in optimizing traffic enforcement strategies. The model's performance indicates that with more cameras, shorter cycle durations are more effective in reducing speed violations. This is likely because more frequent relocations prevent drivers from adapting to the presence of enforcement devices, thus maintaining the deterrent effect. The normalization technique employed here proved effective in isolating the impact of cycle duration, providing clearer insights into its role in traffic enforcement strategies. This is likely because more frequent relocations prevent drivers from adapting to the presence of enforcement devices, thus maintaining the deterrent effect. The normalization technique used here proved effective in isolating the impact of cycle duration, providing clearer insights into its role in traffic enforcement strategies. These results offer valuable guidance for traffic authorities and policymakers in designing more effective Automated Speed Enforcement (ASE) programs. By optimizing the deployment and rotation schedule of cameras, speed violations can be minimized, thereby enhancing road safety. Further research could explore the impact of mixed road segment types and varying enforcement technologies to refine these strategies further.

4.4 Implement model on real data

To validate our MDP heuristic optimization model, we used city of toronto ticketing data. As previously mentioned, the components of the Markov Decision Process (MDP)—states, transition probabilities, rewards, and penalties—must be clearly defined. This section details each component, using data from the City of Toronto's Automated Speed Enforcement (ASE) ticketing records.

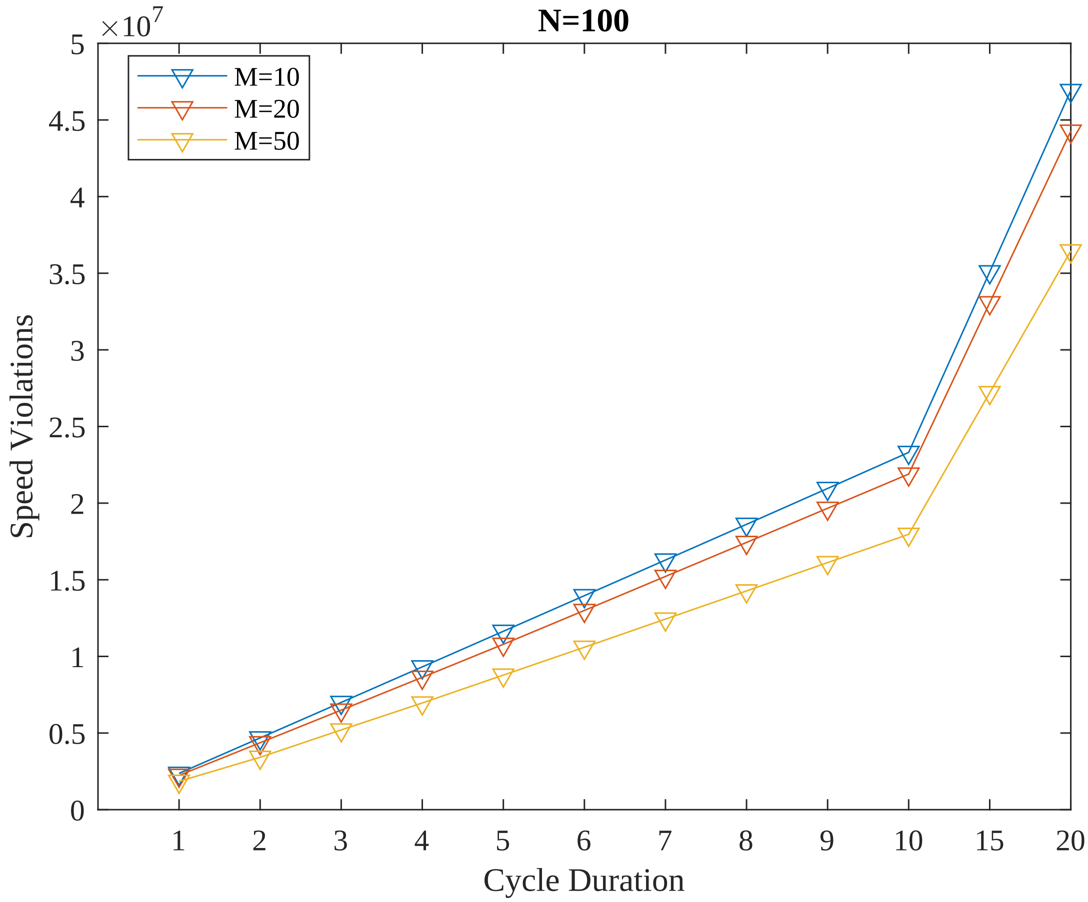

Figure 28: Speed violations across a 100-road segment network for different cycle durations and varying numbers of cameras (M)

4.4.1 Time

While not a direct component of the MDP, defining the time epochs is crucial for accurate calculations. In this study, each epoch corresponds to one month. This time unit is chosen for two main reasons. Firstly, the available ASE ticketing data reports the number of tickets issued per month, preventing us from analyzing fluctuations within shorter periods. Secondly, it is more practical to consider camera relocations on a monthly basis due to the financial costs and logistical challenges faced by municipalities, including legal and staffing considerations.

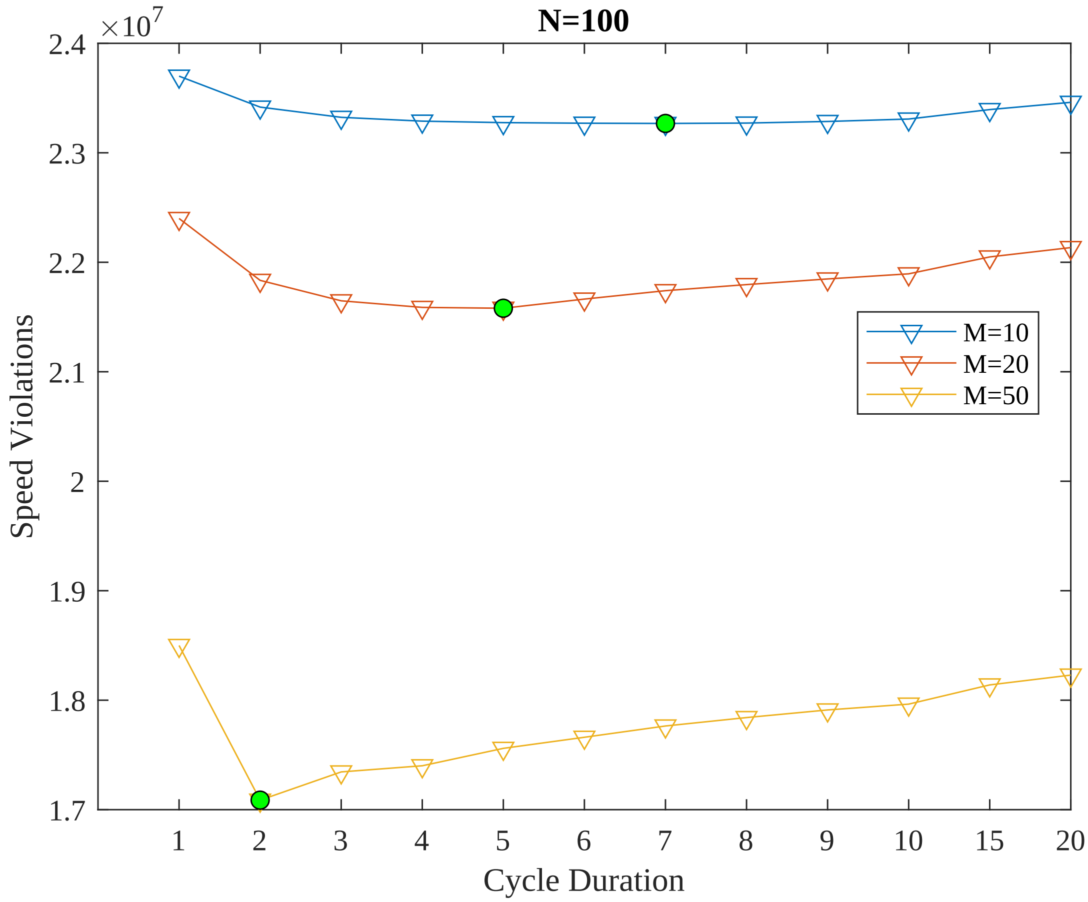

Figure 29: Normalized speed violations across a 100-road segment network, showing optimal cycle duration for different numbers of cameras(M).

4.4.2 States

Defining the states is fundamental to the model, as all other components depend on the initial state definitions. In this context, states represent the number of speed violations recorded at a specific location during one time epoch, or month. Although the number of speed violations can range from 0 to a much higher number, defining each possible number as a distinct state would complicate calculations. Therefore, we define ranges of speed violations, with each range representing a state. For example, if we define two states—'low violation' and 'high violation'—where 'low violation' includes 0 to 50 violations per day and 'high violation' includes 51 and above, then a location with

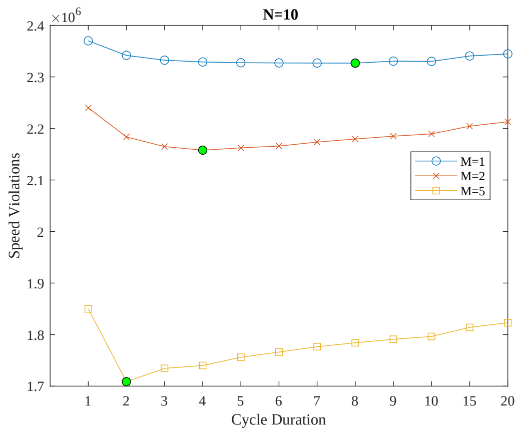

(a) Network with 10 road segments.

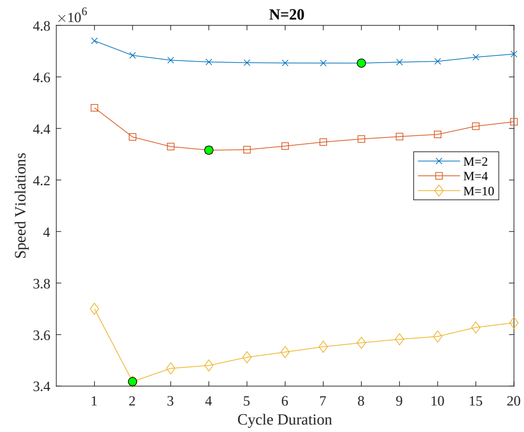

(b) Network with 20 road segments.

Figure 30: Impact of varying M/N ratios on optimal cycle duration

5 violations in a month would fall into the 'low violation' state. To establish the number of states and their corresponding ranges, we use data from the preparation section. Specifically, we compiled an Excel file listing the average number of speed violations per day (denoted as Y) for each month at each location. The Y column contains all the potential speed violation counts observed over monthly intervals. Figure 31, a histogram of the Y values, illustrates the frequency of violations per month, revealing a right-skewed distribution with a long tail. To determine the number of states (i.e., the number of ranges), we employed the k-means clustering algorithm, which has also been used previously to categorize road segments based on ASE effectiveness. K-means is a straightforward method for partitioning daily violation data into distinct groups (e.g., low, medium, high violations). To find the optimal number of clusters for K-means clustering, in other words the number of states, an elbow method implemented on the 'Y', to plot the total within-cluster sum of square (WCSS) distances for different numbers of clusters and identify the point where the WCSS begins to decrease more slowly (the "elbow"). This indicates that adding more clusters does not significantly improve the fit. Based on the elbow plot shown in Figure 32, the "elbow"—the point where the Within-Cluster Sum of Squares (WCSS) begins to level off—occurs around 3 to 4 clus-

ters. This indicates that the optimal number of clusters for the k-means algorithm is likely between 3 and 4. While using 4 clusters would allow for a more detailed model by differentiating states with subtle variations, we opted for 3 clusters to simplify the problem. Figure 33 displays the histogram of the Y values clustered into three groups using the k-means algorithm.

This clustering effectively segments the data into three distinct categories, making it easier to interpret and manage the ASE data in the context of the MDP framework.

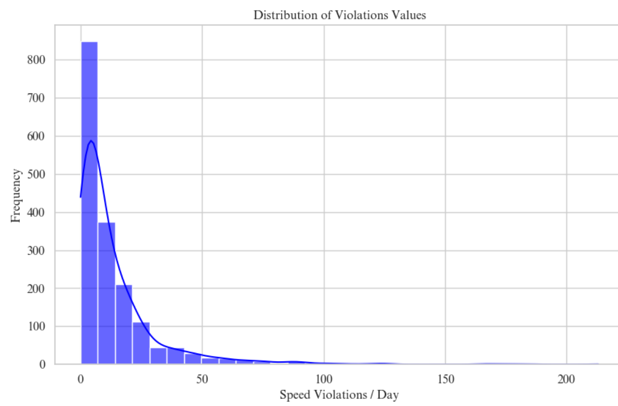

Figure 31: Histogram of ASE ticketing in Toronto

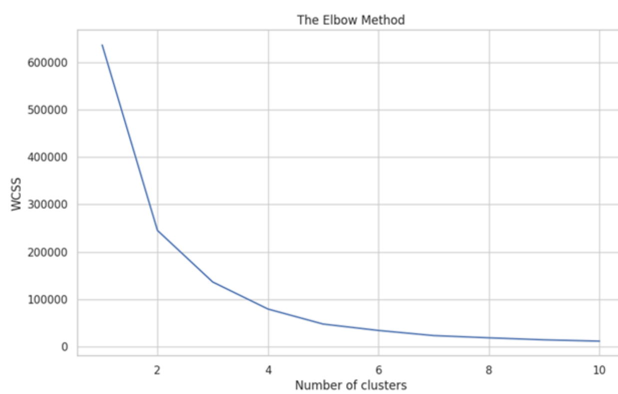

Figure 32: Elbow method implemented on ASE Ticketing

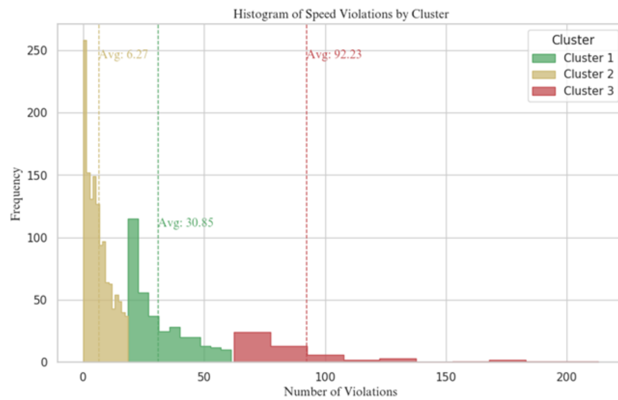

Figure 33: Histogram of ASE ticketing in Toronto

4.5 Defining States and Transition Probabilities

After establishing a framework for defining the states, we need to determine the state of each location at any given time during enforcement. Ideally, data on the number of speed violations during non-enforcement periods would be available; however, since this data is not provided, we focus on the enforcement period and make assumptions for the non-enforcement period. The next step involves calculating the transition probability matrices for enforcement and non-enforcement actions within our network. Since data for the non-enforcement period is not available, we assume a constant non-enforcement transition probability regardless of the location type. This approach allows the model to primarily rely on enforcement period transition probabilities for optimization. Based on the location types defined in the data preparation section through k-means clustering, we assume that locations within the same cluster have identical transition matrices during enforcement. Consequently, we calculate three transition matrices, one for each location type. With three states, each matrix will be 3x3, requiring the calculation of probabilities for transitioning from each state to all other possible states. If we denote the low violation state as 0, medium violation state as 1, and high violation state as 2, the transition probabilities $p(0,0)$, $p(0,1)$, ..., $p(2,2)$ need to be calculated. To determine these probabilities, we examine the frequency of transitions from one state to another for each cluster type. For each location i at time t, we prepare a tuple (s,r), where s is the state of location i at time $t-1$ and r is the state at time t. Two main considerations were taken into account when preparing these tuples: first, tuples could not be defined for rows where $t=1$ represents the first month of enforcement. Second, nine locations were enforced over more than one period, meaning they were initially enforced, then non-enforced, and re-enforced after some time. For these rows, we ensured that the previous state for a location was based on the same enforcement period. Taking these factors into account, we ended up with 1372 tuples in total. After clustering these rows based on their location type, we calculated the density of each tuple for each location type. Figure Y (replace with the actual figure number) shows the frequency of each tuple for each location type. For each type, the transition probability $p(s,r)$ is calculated using the formula:

$$p(s,r) = \frac{f(s,r)}{\sum_{h=0}^{2} f(s,h)}$$

where $f(s,r)$ is the frequency of transitions from state s to state r. The transition probabilities matrix $[P_{E,k}]_{s,r}$ was calculated, and tables 6 to 8, presents the final transition probabilities for each possible transition based on location type.

Metric	(0, 0)	(0, 1)	(0, 2)	(1, 0)	(1, 1)	(1, 2)	(2, 0)	(2, 1)	(2, 2)
Frequency	42	22	6	31	331	0	1	1	13
Distribution (%)	60	31.43	8.57	8.56	91.44	0	6.67	6.67	86.67

Table 6: Transition Probabilities for Non-effectivity ASE location

Metric	(0, 0)	(0, 1)	(0, 2)	(1, 0)	(1, 1)	(1, 2)	(2, 0)	(2, 1)	(2, 2)
Frequency	88	81	1	41	662	0	1	1	0
Distribution (%)	51.76	47.65	0.59	5.83	94.17	0	50	50	0

Table 7: Transition Probabilities for Medium-Effectivity ASE locations

Metric	(0, 0)	(0, 1)	(0, 2)	(1, 0)	(1, 1)	(1, 2)	(2, 0)	(2, 1)	(2, 2)
Frequency	17	9	4	2	7	0	8	1	2
Distribution (%)	56.67	30	13.33	22.22	77.78	0	72.73	9.09	18.18

Table 8: Transition Probabilities for Highly-Effective ASE locations

$$P_{E,\text{Highly-Effective}} = \begin{pmatrix} 0.1818 & 0.0909 & 0.7273 \\ 0 & 0.7778 & 0.2222 \\ 0.1333 & 0.3 & 0.5667 \end{pmatrix}$$

$$P_{E,\text{Non-Effective}} = \begin{pmatrix} 0.8667 & 0.0667 & 0.0667 \\ 0 & 0.9144 & 0.0856 \\ 0.0857 & 0.3143 & 0.6 \end{pmatrix}$$

$$P_{E,\text{Medium-Effective}} = \begin{pmatrix} 0 & 0.5 & 0.5 \\ 0 & 0.9417 & 0.0583 \\ 0.0059 & 0.4765 & 0.5176 \end{pmatrix}$$

$$P_N = \begin{pmatrix} 1 & 0 & 0 \\ 0.6 & 0.4 & 0 \\ 0.5 & 0.3 & 0.2 \end{pmatrix}$$

4.5.1 Results

The model was implemented to prepare scheduling for the City of Vaughan, focusing on the 100 road segments deemed the highest priority by city planners. Using the prediction model described in the data prediction section, we categorized these segments: 30 segments were classified as non-effective ASE locations, 10 as highly effective, and 60 as medium-effective. Given that the City of Vaughan has only 10 cameras, we evaluated different camera deployment durations to find the optimal scheduling plan, as shown in the Figure 34

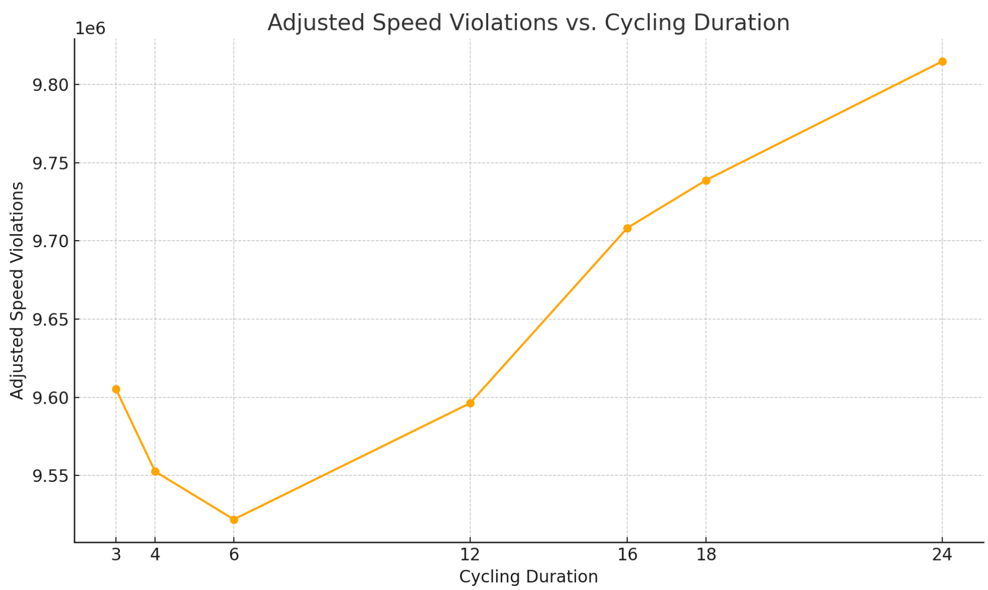

Figure 34: Comparison of the number of speed violations for different cycling periods after normalizing the scheduling period for 12 months

As we can see, the optimal cycling duration for these locations is determined to be 6 months. Delving into the scheduling details, we observe that out of the total 60 months of enforcement available (calculated as $c \times M = 6 \times 10 = 60$ months), 8 months are allocated to highly effective locations, 39 months to medium-effective locations, and 13 months to non-effective locations. While the allocation to medium-effective locations appears significantly higher, comprising the largest share of the enforcement period, and even the non-effective locations receive more enforcement time than the highly effective ones, this distribution must be contextualized by considering the average enforcement period per location type. Specifically, the average enforcement duration per location type is 2.3 months for highly effective locations, 0.65 months for medium-effective locations, and 0.43 months for non-effective locations. This indicates a strategic allocation pattern where the model prioritizes longer enforcement periods for highly effective locations. The tendency to allocate more total enforcement time to medium-effective and non-effective locations reflects the larger number of these locations rather than an equal enforcement emphasis, ensuring that highly effective locations receive proportionally more enforcement time on average. This approach optimizes the use of limited resources, such as cameras, by focusing on areas where the enforcement is predicted to have the most significant impact.

5 City of Vaughan ASE

The amendment of the School Act to the Highway Traffic Act of Ontario, which introduced ASE systems in school and community safety zones, has led many cities in Ontario to adopt this technology. Additionally, numerous other cities are planning future implementations. Among these, the City of Vaughan has expressed intentions to deploy an ASE system.

This tool comprises various functionalities, from showcasing the different road characteristics and network features within the City of Vaughan to identifying community safety zones. Furthermore, it facilitates the creation of an optimal ASE schedule that endeavors to minimize the potential for violations within the designated ASE scenario. A critical aspect of this tool is its ability to allow users to simulate different planning scenarios. This approach involves optimizing the length of enforcement schedules, the number of cameras deployed, the frequency of camera rotations, and the prioritization of objectives within the enforcement strategy—whether the focus is on maximizing revenue from issued tickets or enhancing safety by reducing violations—according to user preferences. Additionally, the tool provides the capability to generate detailed reports in PDF format, which include both the enforcement schedules and analyses of community safety zones.

In this chapter, we will detail the comprehensive process undertaken from the initial data preparation for the City of Vaughan to the development of the ASE tool. This involved an in-depth examination of Vaughan's transportation planning initiative, termed "Move Smart," from which various data sets were extracted and analyzed. A pragmatic model for ASE enforcement was then formulated, followed by a review of the city's latest strategies for defining community safety zones. Ultimately, both the model and the safety zones were integrated into the Mapbox tool, creating a robust, data-driven solution for ASE enforcement planning in Vaughan. This chapter aims to outline each step of this process, demonstrating how theoretical planning and data analysis were translated into a practical tool that addresses the nuanced requirements of ASE system implementation.

5.1 City of Vaughan Move Smart(optional read)

The City of Vaughan has launched the MOVE Smart initiative, focusing on improving transportation safety and efficiency. This initiative prioritizes safety with strategies to reduce traffic collisions and ensure all road users' well-being. The core safety goals of the MOVE Smart strategy include:

1. Data-Driven Safety Enhancements

   Vaughan employs a data-driven approach to pinpoint high-risk areas and traffic patterns, using this insight to tailor safety improvements. This method helps identify where and how to make roads safer for drivers, cyclists, and pedestrians.

2. Focus on Vulnerable Road Users

   A significant part of Vaughan's strategy is protecting vulnerable road users like pedestrians and cyclists. Efforts include creating pedestrian zones, bike lanes, and accessible transit options to make the city's transportation network inclusive and safe for everyone.

3. Safety Culture Promotion

   Beyond physical infrastructure, Vaughan aims to cultivate a safety culture. Through educational campaigns and community involvement, the city encourages everyone to contribute to a safer environment, emphasizing shared responsibility among road users.

4. Leveraging Technology

The strategy also explores using cutting-edge technologies, such as smart traffic signals and ASE. These tools are seen as vital in managing traffic flow, improving response to collisions, and enforcing laws more effectively.

5. Continuous Monitoring for Improvement

   An ongoing evaluation of safety measures ensures that Vaughan's strategies remain effective and responsive to new challenges. This commitment to monitoring and refining approaches guarantees long-term enhancements in road safety

5.2 Data Analysis

The foundation of our analysis in the internship lay in navigating the complexities of municipal data management, notably through our engagement with Traffic Engineering Software (TES). Despite its designed purpose for municipal data oversight, challenges were encountered, particularly with data completeness and accuracy. For instance, collision data was notably constrained to a specific timeframe, and discrepancies were observed in the Average Annual Daily Traffic (AADT) predictions, including the reporting of negative values for certain segments.

In response to these limitations, our analysis was underpinned by five key data sources, carefully selected for their relevance and potential to inform the development of ASE strategies within the City of Vaughan. These sources encompassed a VTC-mapping layer for traffic visualization, cloud-based speed radar records for empirical speed data, historical and current lists of community safety zones, and a comprehensive record of traffic signage. Each source was instrumental in painting a holistic picture of the traffic and safety landscape, serving as critical inputs for our subsequent modeling and tool development efforts. each of these resources and their regarding analysis are

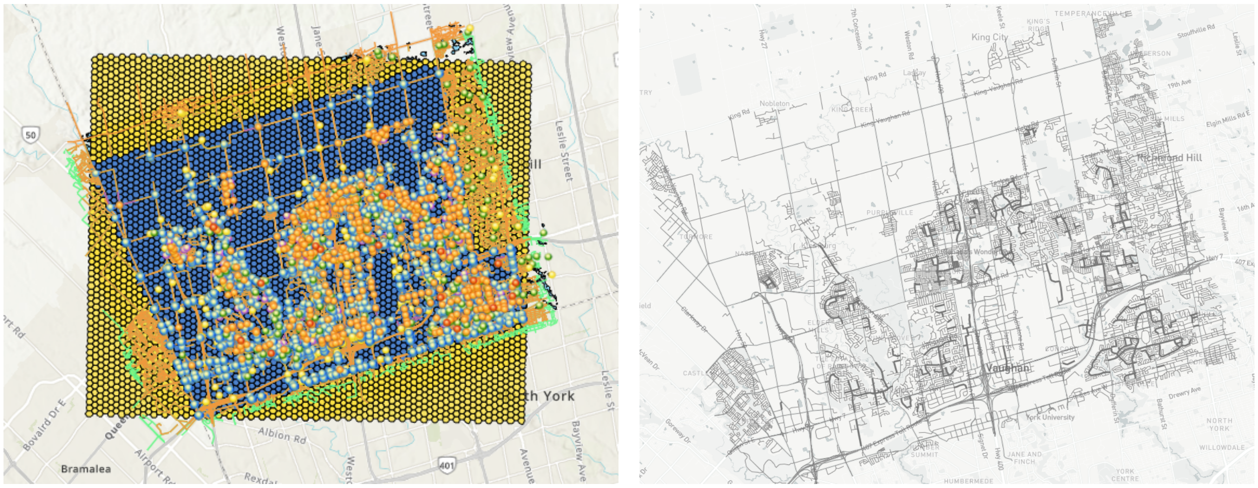

(a) Layers retrived from VTC package (b) City of Vaughan road network retrived from Python

Figure 35: Visualization of the availble data from City of Vaughan

mentioned in the following.

5.2.1 VTC Mapping Layer Data

In the course of our study, we employed a variety of data types derived from the VTC mapping layer displayed in Figure 35a to enrich our understanding of traffic dynamics and safety within the City of Vaughan. The data types included serious injury collisions, the number of lanes, speed limits, intersection volumes, and ward borders. The comprehensive data package offered multiple layers for analysis; however, the dataset documenting collisions with fatal injuries was initially available only in Excel format, lacking spatial coordinates in a conventional degree measure.

To address this challenge, we embarked on a process to geocode the provided X and Y coordinates, which were not initially presented in degree measures. The intersection volume data, downloaded as a CSV file, underwent a similar geocoding process in ArcGIS Pro. We selected the NAD_1983_UTM_Zone_17N projected coordinate system for our data, ensuring precise geolocation within the City of Vaughan. Each data layer was subsequently saved as a separate hosted feature layer within the GIS environment.

Further, we integrated and analyzed the speed limit and number of lanes data in conjunction with a city network extracted from OpenStreetMap. The extraction was facilitated by the Python package osmnx, treating the city's road network as a graph with road segments as edges and intersections as nodes. It is important to note our analytical focus was primarily on the edges (road segments), while the attributes of intersections (nodes) were not extensively explored as seen in Figure 35b.

The road network data offered detailed attributes for each segment, such as road type (e.g., highway, residential, arterial), directionality (two-direction or one-direction), and toll status. However, this information was not uniformly available for all road segments within the city network, resulting in some data gaps. Consequently, the data related to speed limits and the number of lanes were utilized as reference metrics. Through a spatial join process in ArcGIS Pro, we consolidated these metrics with the broader city network data into a single, comprehensive shapefile.

The data preparation and analysis process underpinned our subsequent development of a tool for ASE enforcement planning. By integrating diverse data layers and ensuring precise geolocation, we laid a solid foundation for a data-driven approach that enhances traffic safety and management strategies within the City of Vaughan.

5.2.2 Cloud-Based Speed Radar Data

The cloud database was instrumental in providing access to speed radar data, among other types of traffic-related information. Specifically, the ”Watch Your Speed” (WYS) program in the City of Toronto, which employs radars to record and display the speed of passing vehicles, served as a primary data source. These radars, positioned alongside roads, are not intended for enforcement purposes but rather aim to increase driver awareness of their current speeds by displaying it on a digital board adjacent to a speed limit sign. This allows drivers to immediately compare their speed with the legal limit.

The strategy for deploying these radars in the City of Vaughan primarily hinges on community requests. Residents with concerns about violating the speed limit in their neighborhoods can contact the city to request the installation of a speed radar. Consequently, the placement and duration of these radars are not dictated by a uniform safety strategy but are instead variable, dependent on community feedback. This approach, while responsive to immediate community concerns, presents a challenge for systematic data analysis due to the ad hoc nature of radar locations and the limited number of devices that are moved based on public requests.

The cloud platform that houses the radar data proved to be user-unfriendly, complicating the extraction process. The data for each radar installation period was presented in a two-column table, with one column listing various speed ranges (from 0 to over 160 km/h) and the other detailing the number of vehicles detected within each speed range(Screenshot of the platform is displayed in Figure 5.2.2). Each location's data had to be downloaded separately as an Excel file and then aggregated for comprehensive analysis. Further complicating the process was the dispersed listing of locations with WYS experiences, requiring a manual search and export operation for each site based on the year of installation.

Given these challenges, our analysis initially focused solely on the speed radar data from 2020, the most recent year available at the time of our study. We compiled data from 155 radar locations into a single Excel file, documenting for each site the start and end dates of radar installation and the recorded speed counts across the specified ranges. The duration of each radar's operation was calculated from these dates, and the locations were geocoded using ArcGIS Pro to produce a shapefile mapping all 2020 WYS installations.

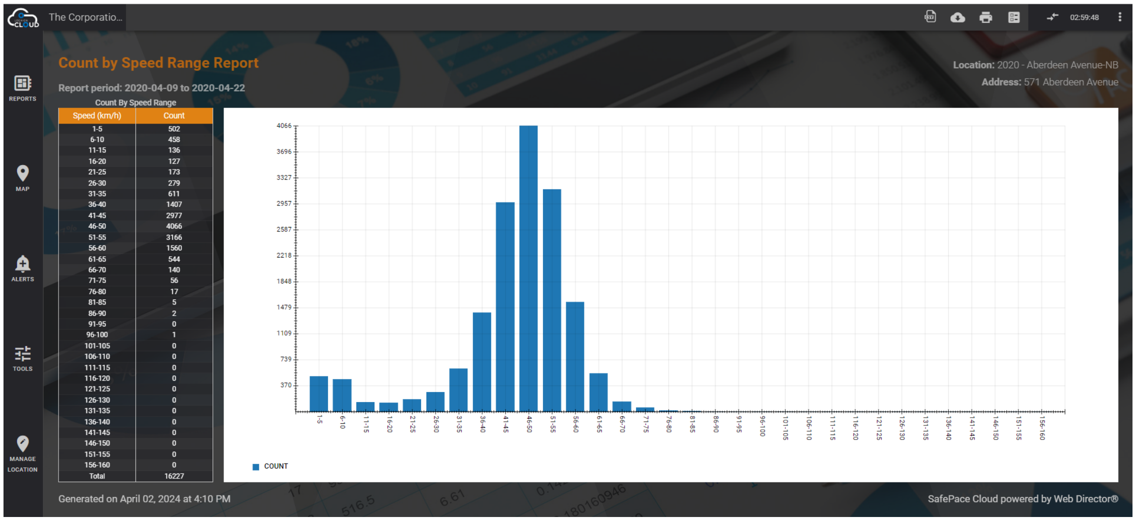

Figure 36: Screen shot from the Speed Radar Cloud Platform

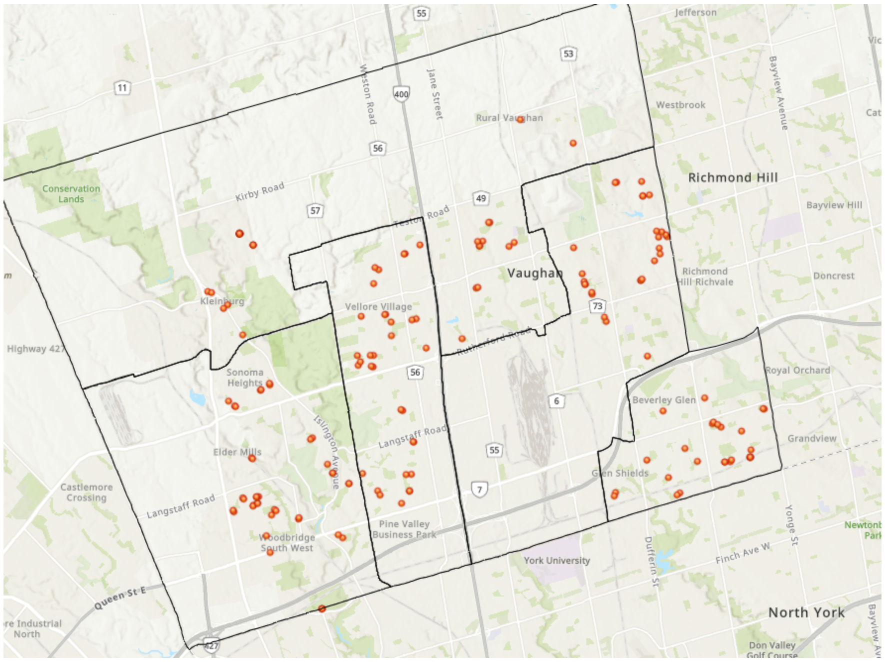

Figure 37: Location of all Speed Radars installed in 2020, after geocoding.

To investigate the issue of speed violations further, we performed a spatial join in ArcGIS Pro between the Speed Radar shapefile and a shapefile retrieved from the VTC layering file, which contained speed limit data for each location. This analysis enabled us to quantify the number of speed violations recorded at each radar location during its operational period.

The initial strategy for optimizing ASE scheduling involved leveraging speed radar data to estimate potential speed violations across different road segments. However, a significant obstacle emerged: none of the speed radar installations were situated within community safety zones as defined at the time. To overcome this limitation, we explored machine learning techniques and training models on existing speed radar data to predict the likelihood of violations based on various road characteristics, such as speed limits, lane numbers, and the proximity of serious injury collisions.

Data preparation for this machine learning analysis entailed augmenting the speed radar dataset with additional variables, including lane count, speed limit, traffic volume, collision history within a two-kilometre radius, and distance to the nearest school. These attributes were sequentially merged with the radar data using spatial joins in ArcGIS Pro.

Our initial attempt to model the relationship between these features and daily speed violation counts used a simple linear regression. However, the limited size of our dataset (155 locations) suggested that a continuous regression model might not be the most appropriate approach. We examined the distribution of speed violations per day across all radar locations to categorize the locations based on their violations. As depicted in Figure 38, the histogram shows a skewed distribution, with the majority of locations experiencing fewer than 100 violations per day. The 50th and 85th percentiles, marked by dashed lines, correspond to approximately 47.52 and 379.48 violations per day, respectively. These thresholds were critical in categorizing the radar locations into low, medium, and high-risk groups.

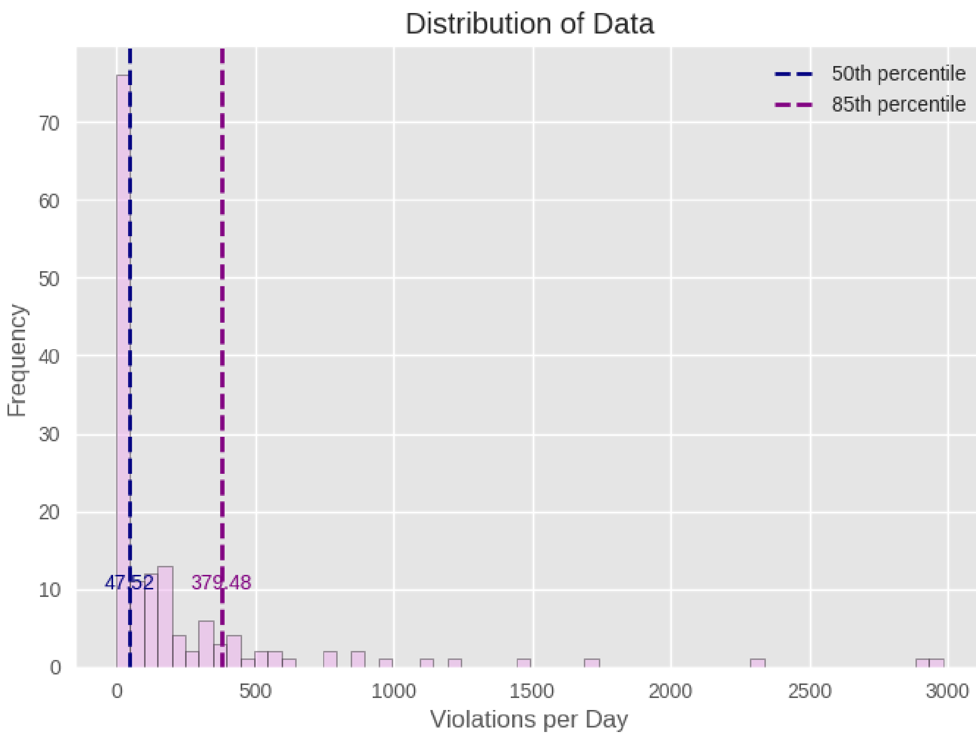

Figure 38: Distribution of Speed Violations per Day

The histogram highlights that most radar locations fall into the low-risk category, with a significant drop in the frequency of high violation observations. This distribution pattern underscores the need for targeted enforcement in high-risk areas to maximize the impact of ASE initiatives. By focusing resources on locations with higher predicted violations, road safety will be enhanced.

First of all, the radar locations were categorized into three risk groups based on daily violation counts: low (50 or fewer violations per day), medium (between 50 and 380 violations), and high (over 380 violations). Each location was assigned to its respective group, facilitating the use of supervised machine learning models such as Random Forest, Logistic Regression, SVM, K-Nearest Neighbors, Decision Tree, and MLP classifiers. After that, a feature importance analysis was conducted using several classifiers, including Logistic Regression, Decision Tree Classification, Random Forest Classification, and XGBoost Classification. The results, summarized in Table 1, highlight the relative importance of various features in predicting speed violations.

Table 9: Feature Importance Across Different Models

Model	Feature	Importance
LogisticRegression	Number of Collision within 1 km Radius	-0.01
LogisticRegression	Number of Lanes	0.00
LogisticRegression	SPEED_LIMI	0.03
LogisticRegression	Volume	-8.94
LogisticRegression	Distance to Closest School	4.63
DecisionTreeClassifier	Number of Collision within 1 km Radius	0.38
DecisionTreeClassifier	Number of Lanes	0.01
DecisionTreeClassifier	SPEED_LIMI	0.03
DecisionTreeClassifier	Volume	0.19
DecisionTreeClassifier	Distance to Closest School	0.39
RandomForestClassifier	Number of Collision within 1 km Radius	0.32
RandomForestClassifier	Number of Lanes	0.06
RandomForestClassifier	SPEED_LIMI	0.07
RandomForestClassifier	Volume	0.21
RandomForestClassifier	Distance to Closest School	0.33
XGBClassifier	Number of Collision within 1 km Radius	0.02
XGBClassifier	Number of Lanes	0.14
XGBClassifier	SPEED_LIMI	0.29
XGBClassifier	Volume	0.20
XGBClassifier	Distance to Closest School	0.17

Based on this analysis, ”Volume” emerged as a highly significant predictor, especially in the Logistic Regression model with an importance score of -8.94. This negative importance suggests an inverse relationship between traffic volume and speed violations, which warrants further investigation. Similarly, ’Distance to Closest School’ was also a significant feature, with importance scores of 4.63 in Logistic Regression and 0.39 in the Decision Tree Classifier, indicating that proximity to schools is a crucial factor in predicting speed violations. Other important features included ”Number of Collision within 1 km Radius” and ”Number of Lanes”. The ”Number of Collisions within 1 km Radius” showed varying degrees of importance across different models, with the Decision Tree Classifier assigning it the highest importance score of 0.38. This underscores the relevance of collision history in anticipating speed violations. The ”Number of Lanes’ feature, although showing lower importance scores than other features, still provided valuable insights, particularly in the XGBClassifier, with an importance score of 0.14. These insights enabled us to refine our predictive models, ensuring they account for the most impactful road characteristics. To further analyze the performance of our supervised machine learning models, we compared the accuracy and precision of each classifier. The following bar chart in Figure 40 illustrates the accuracy and precision scores for Logistic Regression, Decision Tree Classifier, Random Forest Classifier, MLP Classifier, SVC, XGBClassifier, K-Neighbors Classifier, and GaussianNB. From Figure 40 and Table 10, we can observe that the models exhibit varying levels of accuracy and precision. Logistic Regression and K-Neighbors Classifier achieved the highest accuracy scores of 0.70, indicating a strong overall performance in correctly predicting speed violations. K-Neighbors Classifier also demonstrated the highest precision score of 0.63, which suggests that it is particularly effective at correctly identifying true positives and minimizing false positives. The random forest and MLP classifiers were followed closely, with balanced accuracy and precision scores of around 0.65-0.68 for accuracy and 0.61-0.62 for precision, showcasing their robustness in handling the dataset. SVC and XGBClassifier also performed reasonably well, with accuracy scores of 0.69 and 0.67, respectively, though their precision scores were slightly lower, indicating a moderate rate of false positives. Decision Tree Classifier and GaussianNB had the lowest precision scores, at 0.55 and 0.57, respectively, suggesting a higher incidence of false positives compared to other models. Overall, the K-Neighbors Classifier emerged as the most balanced model with high scores in both accuracy and precision, making it a strong candidate for deployment in predicting speed violations. The insights gained from this analysis will inform the strategic placement of speed radars and the scheduling of ASE initiatives, ultimately enhancing road safety.

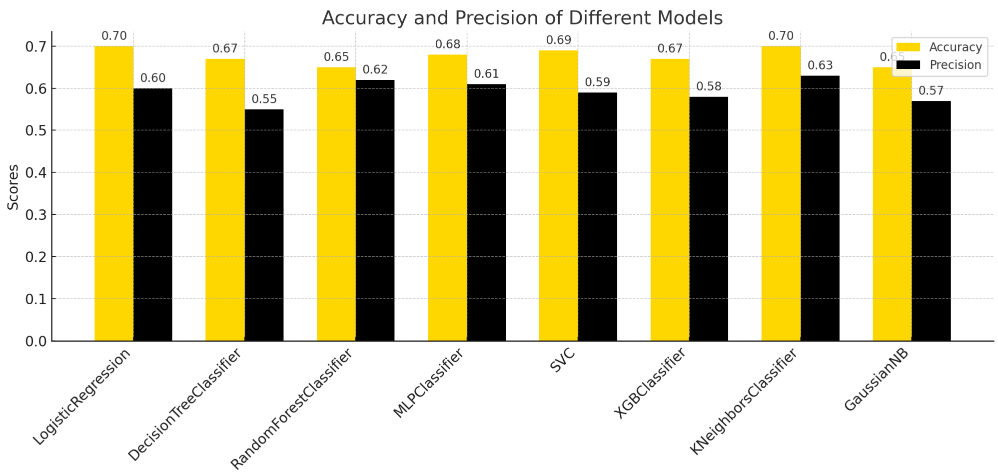

Figure 39: Accuracy and Precision of Different Models

Table 10: Accuracy and Precision Scores of Different Models

Model	**Accuracy**	**Precision**
LogisticRegression	0.70	0.60
DecisionTreeClassifier	0.67	0.55
RandomForestClassifier	0.65	0.62
MLPClassifier	0.68	0.61
SVC	0.69	0.59
XGBClassifier	0.67	0.58
KNeighborsClassifier	0.70	0.63
GaussianNB	0.65	0.57

5.2.3 Act by Law Data

In this section of our analysis, we focused on integrating act-by-law data pertaining to community school zones within the City of Vaughan. In PDF format, the provided data outlined specific road segments identified for potential implementation of ASE systems or other forms of automated enforcement until September 2023. Each segment was detailed with start and end intersection addresses, facilitating their eligibility assessment for ASE deployment. These addresses were first digitized into an Excel spreadsheet to streamline the geocoding process. Subsequently, the geocoding process transformed these textual addresses into geographic coordinates, which allowed for the precise mapping of each road segment. In ArcGIS Pro, these mapped segments were manually traced with sketch lines to visually represent the designated areas between the start and end points as specified in the act by law document.

Following the tracing process, all individual sketch lines were consolidated into a singular layer, which was then saved and hosted as a feature layer within our geographic information system (GIS) platform. This systematic approach facilitated the detailed visualization and analysis of areas earmarked for ASE implementation, underlining the spatial extent and distribution of these zones within the urban fabric of Vaughan.

In parallel, the City of Vaughan was engaged in efforts to refine and expand the scope of these act by law designated locations by developing a more comprehensive definition for community safety zones. This endeavor was part of a broader initiative to enhance road safety, particularly in areas frequented by vulnerable populations such as schoolchildren. The consultancy firm CiMA+ was commissioned to lead this project, and as part of their preliminary work, some potential areas were delineated and prepared as shapefiles for initial model testing.

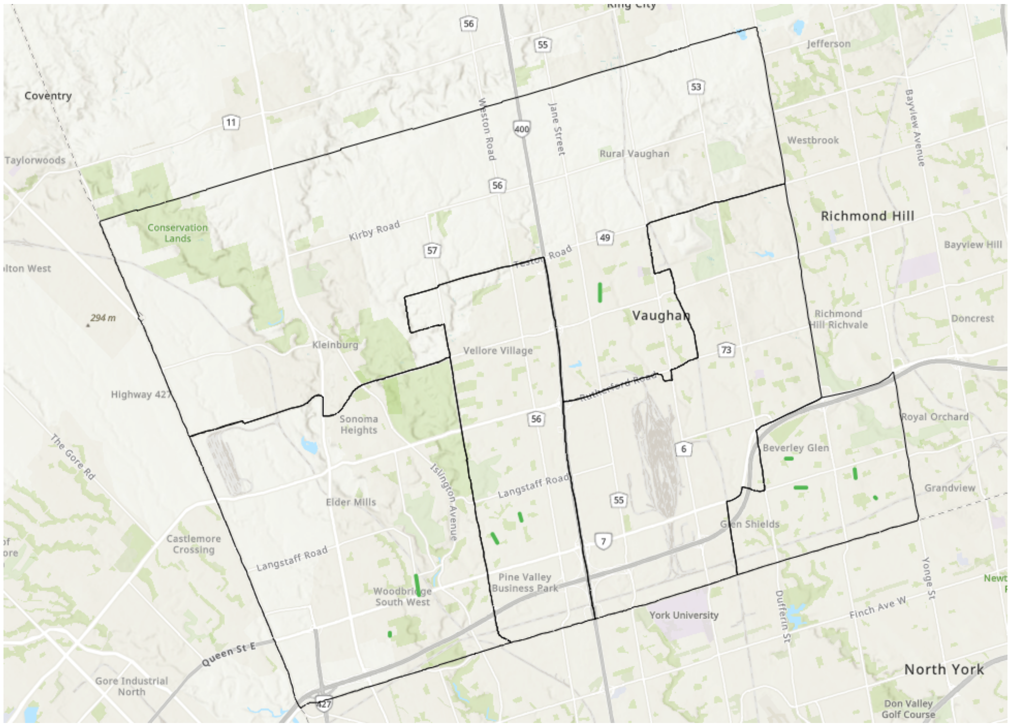

Figure 40: Accuracy and Precision of Different Models

5.2.4 Traffic Sign Data

In this study, a comprehensive dataset containing information on all traffic signs within the City of Vaughan was pivotal in identifying potential locations for community safety zones. This dataset, cataloged in an Excel file, included details such as sign types, locations (latitude and longitude), and additional attributes pertinent to our analysis. Among the various types of signs, those indicating school zones and areas with a maximum speed limit of 40 km/h were particularly interesting, as they indicate zones with high pedestrian activity and heightened safety considerations.

The initial step involved geocoding the signs based on the latitude and longitude coordinates they provided. Of the 19,038 signs listed in the dataset, we successfully geocoded 17,510. Although the URL for the photo of each sign was removed from our dataset, a unique identifier column was added. This identifier was intended to facilitate the linkage of each sign to photographs in a separate

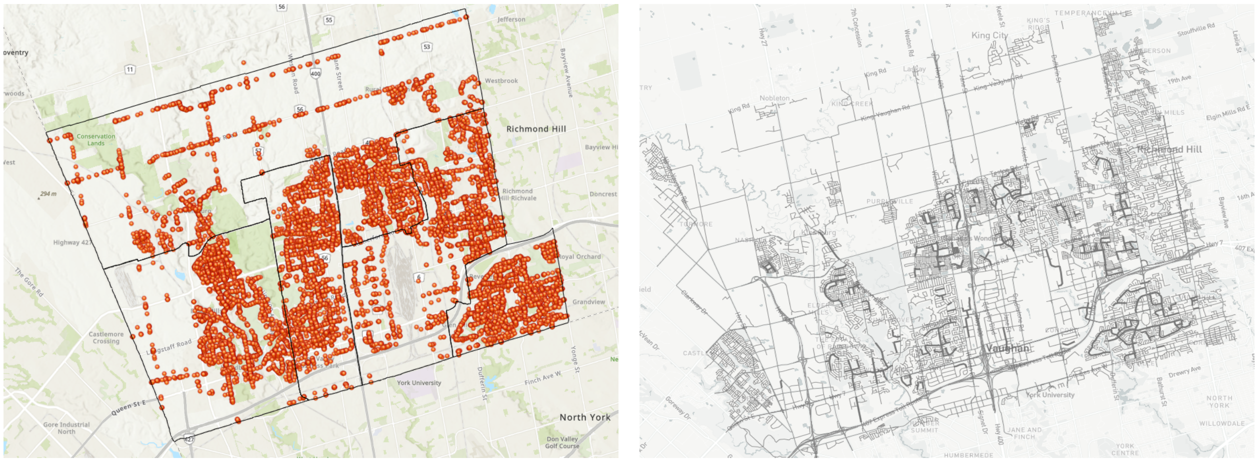

(a) Location of all sign throughout the City of Vaughan after Geocoding the Excel file

(b) Location of 'SCHOOL ZONE MAXIMUM SPEED' signs and possible school zone

Figure 41: City of Vaughan Sign data

Excel file for potential future research.

Subsequently, the geocoded sign locations were imported into ArcGIS Online for further analysis as in Figure 41a. We applied a filter to isolate signs labeled as 'SCHOOL ZONE MAXIMUM SPEED' (in Yellow and White), deeming these locations as potential candidates for community safety zones. To enhance the accuracy of our identification process, we also incorporated a feature layer of school zone locations from the VTC map layer into our ArcGIS Online map view. This allowed for a direct comparison and verification of potential safety zones against officially designated school zones.

However, a challenge arose in determining the exact road segments to classify as potential community safety zones, especially when a sign was placed at only one intersection near a school. To address this, we conducted manual searches for each location on Google Maps, utilizing Street View to understand the surrounding school area better. This meticulous process enabled us to delineate specific road segments that aligned with the maximum speed school zones, culminating in creating a shapefile representing these identified segments in Figure 41b.

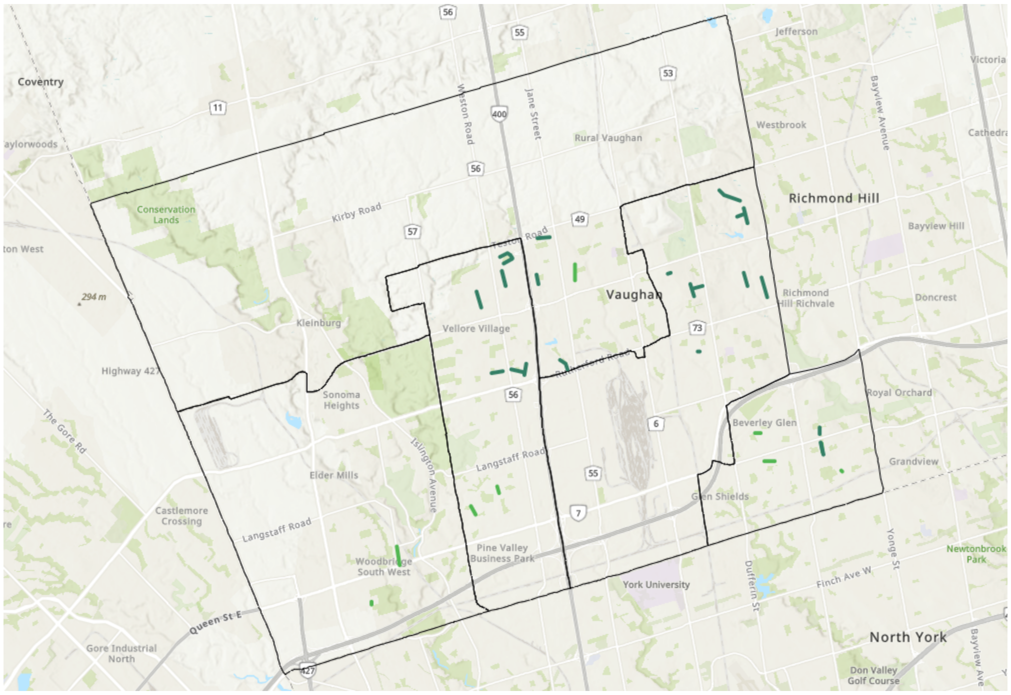

Figure 42: Suggested Community safety zones, based on both Act-by law data, and sign data

It is noteworthy that Although this shapefile was not directly utilized in subsequent phases of the study, the integration of this shape file with the Act by law shape file, which is shown in 42 proved invaluable for CIMA+ consultancy and the City of Vaughan. It facilitated a more informed perspective on potential school zones, thereby contributing to the city's ongoing efforts to enhance road safety through the strategic designation of community safety zones.

5.3 Optimization Model

Implementing ASE systems aims to mitigate the risk and severity of traffic collisions by reducing the incidence of speed violations. While the intrinsic goal of such systems is enhancing road safety, the financial aspect of operating these systems, mainly the potential revenue from issuing violating the speed limit tickets, cannot be overlooked. Although the City of Vaughan emphasizes safety as its paramount concern, developing a model that can also account for revenue generation is prudent, enabling city planners to consider its impact if necessary.

5.3.1 Objective Function and Constraints

To construct an effective optimization model, we must first delineate the objectives and constraints of the ASE system. Speed camera studies typically assess effectiveness through metrics such as average speed reduction, collision rates, and the frequency of speed violations. For this study, we focus on minimizing speed violations as the primary safety metric, aiming to maximize the number of speed violations prevented across the network during the planning period while also considering the revenue generated from tickets as a secondary objective.

Given that ASE cameras are to be periodically relocated, the cost associated with such relocation must also be factored into the model. Our dual objective is thus to maximize both safety outcomes and ticket revenue. By assigning a zero weight to ticket revenue, the model will exclusively prioritize safety, defined here as the reduction in speed violations detectable by ASE during enforcement periods.

Leveraging ASE ticketing data from the City of Toronto, we observed a general decline in ticket issuance overtime at most sites, suggesting a deterrent effect on violating the speed limit behaviour. This trend is incorporated into our model, assuming a linear decrease in ticket issuance post-enforcement, with violations returning to baseline levels upon camera removal.

5.3.2 Objective Function

The model employs a binary variable x_{ij} to indicate whether a community safety zone i is enforced at time j, with $x_{ij} = 1$ signifying enforcement and $x_{ij} = 0$ otherwise. Given a limited number of cameras M, the sum of x_{ij} across all zones at any given time j must not exceed M. Furthermore, if a zone i is enforced at time j, it cannot be enforced again until after k cycles, imposing a constraint on the sequence of enforcement across cycles.

The City of Vaughan has expressed a desire for equitable camera distribution across wards, introducing a constraint *perWard* to ensure that the sum of x_{ij} for each ward at any time j does not exceed a predefined limit.The optimization problem is formulated as follows:

$$\max \sum_{i=1}^{n} \sum_{j=1}^{k} R_{ij} \cdot x_{ij} \tag{27}$$

Where R_{ij} represents the combined metric of safety value and revenue for enforcing a community safety zone i at time j,R_{ij} is actually retrieved from the machine learning model that predicts the number of speed violations after t month of enforcement described in the data analysis section.

It is assumed that the minimum violation rate for any location cannot fall below four tickets per day, a threshold based on empirical data from Toronto, to ensure that the model does not predict unrealistically low violation numbers.

By determining the optimal enforcement schedule x_{ij} for each community safety zone across different cycles, this model seeks to balance safety enhancements with potential revenue generation, subject to the constraints of camera availability, enforcement cycle timing, and ward-based distribution requirements.

5.3.3 Constraints

Given the objective function designed to maximize both safety and revenue from the ASE system, we must consider several constraints that govern the feasibility and efficacy of the ASE scheduling plan.

1. Constraint 1: Camera Availability

   To ensure that the deployment does not exceed the available number of cameras M, we apply the following constraint for each time period j:

$$\sum_{i=1}^{n} x_{ij} \leq M, \quad \forall j \tag{28}$$

   This constraint ensures that the total number of community safety zones enforced at any time does not exceed the total number of available ASE cameras.

2. Constraint 2: Enforcement Cycle

   Considering that each location i can only be enforced once within a certain number of cycles k to allow for camera rotation, we introduce a constraint to manage the enforcement frequency for each zone:

$$\sum_{l=j}^{j+k-1} x_{il} \leq 1, \quad \forall i, \forall j \tag{29}$$

   This constraint ensures that if a zone is enforced at time j, it cannot be enforced again until after k cycles have elapsed.

3. Constraint 3: Equitable Distribution Across Wards

   To accommodate the City of Vaughan's requirement for equitable camera distribution, we define a parameter *perWard* that limits the number of cameras per ward at any given time:

$$\sum_{i \in W} x_{ij} \leq perWard_W, \quad \forall j, \forall W \tag{30}$$

where W represents each ward, and $perWard_W$ is the maximum number of cameras allocated to ward W at any time j.

4. Constraint 4: Minimum Violation Threshold

   Lastly, we set a minimum daily violation rate to maintain realism in the model by ensuring that the predicted number of violations does not fall below a certain threshold. This reflects the understanding that a certain level of violations is likely to persist regardless of enforcement, modelled as follows:

$$R_{ij} \geq 4, \quad \forall i, \forall j \tag{31}$$

   Where R_{ij} is the expected number of violations (or tickets) for zone i at time j, and the minimum threshold is set to 4 tickets per day based on empirical data.

Incorporating these constraints into our optimization model allows for a realistic and practical approach to scheduling ASE deployments within the City of Vaughan. The model balances the primary objective of enhancing road safety with the secondary consideration of revenue generation, all while adhering to logistical limitations and policy requirements. Through this mathematical formulation, city planners can devise an ASE scheduling plan that optimally utilizes available resources to achieve the greatest possible impact on road safety. However, the completion of community safety zone definitions is a prerequisite for finalizing the ASE scheduling for the City of Vaughan.

5.4 Community safety Zones definition

In Ontario, the concept of Community Safety Zones plays a crucial role in enhancing road safety, particularly in areas frequented by vulnerable road users such as school children, pedestrians, and cyclists. Defined under Section 214.1 of the Highway Traffic Act (HTA), these zones empower municipalities to designate specific portions of highways under their jurisdiction as areas where public safety is of paramount concern. The primary legislative intent behind establishing Community Safety Zones is to deter dangerous driving behaviours by imposing enhanced penalties for traffic violations within these designated areas. Designating a Community Safety Zone involves enacting a municipal by-law, which specifies the highway segment to be designated and the effective times of the designation. Once established, these zones are marked with clear signage to inform drivers of their entry into an area with significantly higher stakes for traffic violations. The HTA mandates that fines for offences such as violating the speed limit, reckless driving, and other hazardous behaviours are doubled within Community Safety Zones, aiming to foster a safer environment for all road users. The importance of Community Safety Zones extends beyond the mere legal framework and imposition of penalties. These zones are a testament to the commitment of municipalities and law enforcement agencies toward reducing traffic-related collisions, injuries, and fatalities. By targeting areas with a high concentration of vulnerable road users or a history of traffic collisions, Community Safety Zones act as a deterrent against dangerous driving behaviours, thereby enhancing the community's overall safety. Moreover, establishing Community Safety Zones catalyzes community awareness and engagement regarding road safety issues. It encourages residents to advocate for safer driving behaviours and participate in road safety initiatives, fostering a culture of mutual respect and care among all road users. This collective effort is pivotal in creating safer road environments, especially in areas where children and other vulnerable populations are present. The City of Vaughan, with the help of CiMA +, developed a comprehensive two-step warrant system designed to effectively justify the designation of Community Safety Zones (CSZs) in Ontario. This system introduces a nuanced approach to enhancing road safety, particularly for vulnerable road users in high-risk areas. The first step, Warrant 1, focuses on identifying Designated Areas of Special Concern. The second step, Warrant 2, introduces a Safety Warrant that assesses and quantifies potential safety risks, guiding the decision-making process for implementing CSZs.

- **Warrant1: Designated Area of Special Concern**

  Warrant 1 targets specific locales where the convergence of vulnerable road users and vehicular traffic poses significant safety concerns. These areas are often characterized by their propensity to attract individuals who are more susceptible to traffic-related hazards. Key land uses satisfying Warrant 1 criteria include School Zones, Trail Access Points, Retirement Housing, Community Centers, Parks, and Places of Worship. Notably, School Zones are deemed inherently warranting for CSZ designation, exempting them from further validation under Warrant 2. Additionally, when schools and parks are proximate, they should be treated as a unified candidate site for potential CSZ application. The evaluation extends to other boundary roads surrounding the candidate location to ascertain their eligibility for CSZ status.

- **Warrant 2: Safety Warrant**

  Following the confirmation of a Designated Area of Special Concern via Warrant 1, the

Safety Warrant (Warrant 2) assesses the feasibility and potential impact of implementing a CSZ to mitigate safety risks. This assessment leverages a risk scoring matrix that evaluates each candidate location against a series of risk factors. These include Vehicle Volume, indicating the higher potential for conflicts due to increased vehicle presence; Number of Lanes, with multi-lane roads presenting longer crossing distances and heightened exposure to vehicle interactions; Length of Sidewalk, offering a physical buffer between pedestrians and vehicular lanes; Truck Volume, noting the compounded risk due to trucks' maneuverability challenges and collision severity; Bus Stops, acting as indicators of pedestrian activity and crossing frequency; Intersections and Entrances, increasing conflict points; Operating Speeds, where higher speeds correlate with more severe collision outcomes; and Collision History, providing a historical perspective on the location's safety performance. Each risk factor is assigned a score—'high' (3), 'moderate' (2), or 'low' (1)—to quantify the safety risk present and prioritize CSZ designation accordingly. This scoring system facilitates a data-driven approach to identifying and mitigating road safety hazards in Ontario's communities. The proposed two-step warrant system offers a structured and evidence-based framework for enhancing the safety of Ontario's roads, especially for the most vulnerable users. By systematically identifying areas of special concern and quantifying safety risks, this policy framework aims to support the strategic implementation of Community Safety Zones, ultimately contributing to safer community environments and reducing traffic-related injuries and fatalities.

The Excel file containing all road segments within warrant one was provided alongside the shapefile of the final community safety zones as determined by this policy visualized in Figure 43. Unfortunately, access to the shapefile for the initial warrant was restricted, preventing further analysis. Consequently, our focus shifted to the final dataset made available to us.

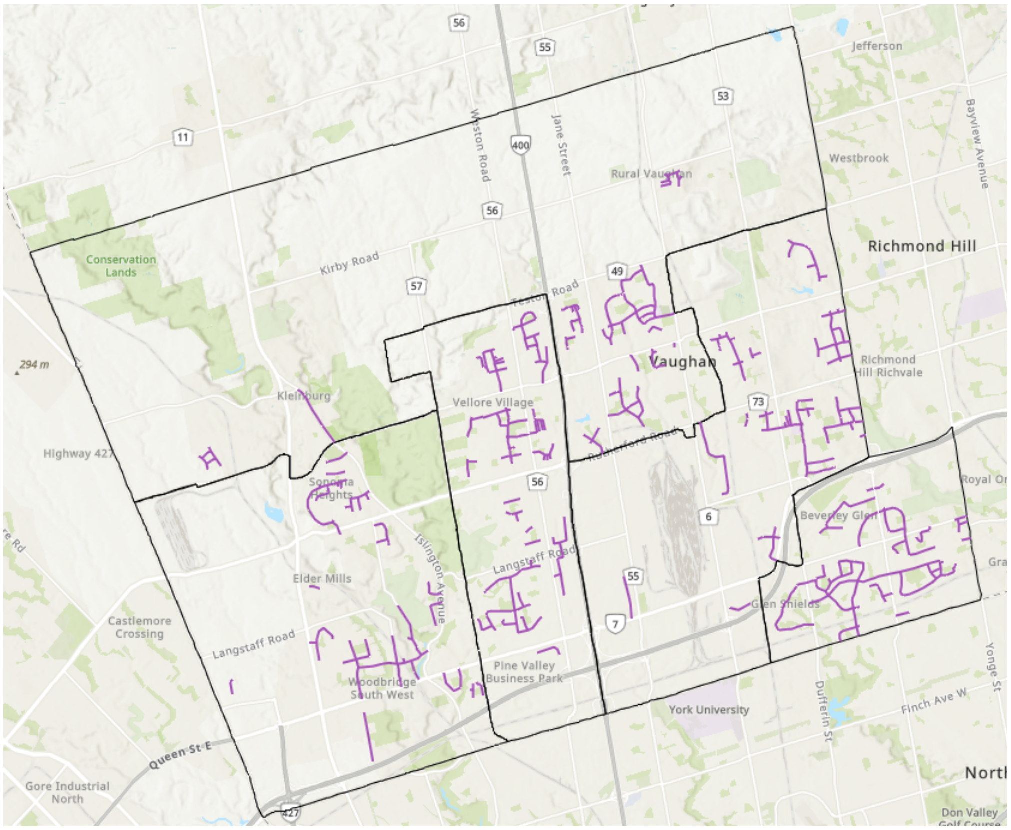

Figure 43: Final Community Safety Zones shapefile based on the 2023 Policy of city of Vaughan

5.5 MapBox JS Tool

Mapbox JS is a sophisticated, interactive mapping library that harnesses the full potential of JavaScript, offering an extensive suite of tools for the creation and management of dynamic maps. Predicated on the solid framework provided by Mapbox's customizable mapping solutions, this library enables seamless integration of geographical data into both web and mobile applications. It empowers developers to craft map experiences that are both highly customized and interactive, allowing for the manipulation of map features, adjustment of visual elements, and real-time interaction with geographic information. Given its versatile API, Mapbox JS supports a myriad of mapping functionalities, including layering diverse map styles, incorporating custom data sources, and facilitating user interactions. Consequently, Mapbox JS emerges as the preferred choice among developers seeking to incorporate sophisticated mapping features into their applications, ensuring precision, scalability, and a rich user experience.

Within the scope of this project, we harnessed Mapbox JS's capabilities to develop an innovative ASE tool designed to consolidate our analytical and optimization models into a singular, integrated platform. This tool is aimed at municipal city planners, offering them not merely a means to visualize current traffic and safety data but also a dynamic mechanism for engaging in the planning process. It facilitates the exploration of various scenarios, allowing planners to devise and refine ASE schedules with greater flexibility and foresight.

The ensuing sections will detail the various components of the Mapbox-based ASE tool, elucidated with accompanying screenshots to enhance comprehension and illustrate the tool's functionality. Each segment of the tool has been meticulously developed to cater to the specific needs of city planners, enabling a comprehensive understanding and evaluation of potential ASE deployments. Through this tool, planners are equipped with the means to dynamically interact with data, assess the impact of different enforcement strategies, and ultimately make informed decisions that align with safety and community objectives.

5.5.1 View of Tool

Figure 44 provides an introductory view of the ASE planning tool. Upon launching the tool, the user is presented with a detailed map of the City of Vaughan, delineating its intricate road network and the defined boundaries of its political wards. Strategically positioned across the map are a

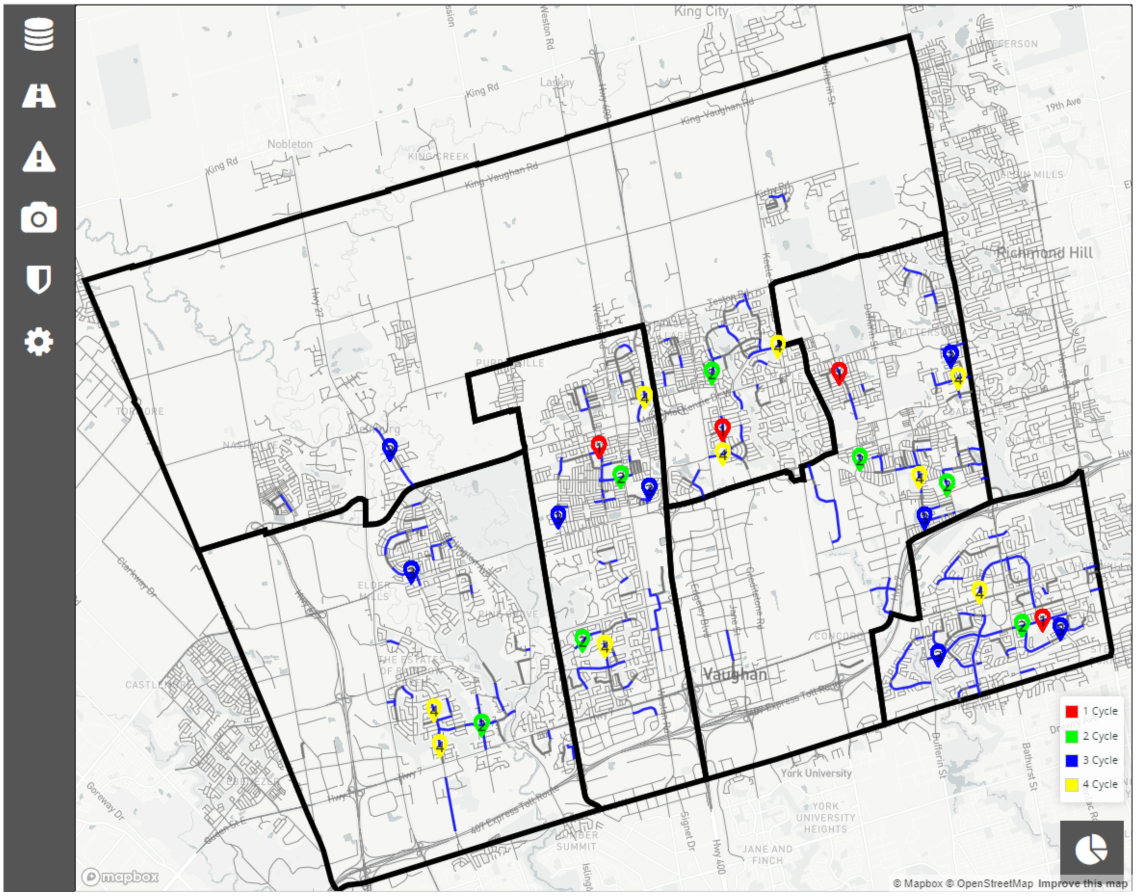

Figure 44: Introductory view of the ASE planning tool.

series of markers, each distinguished by a unique colour—red, green, blue, and yellow—adorned with the numerals 1 through 4, respectively. These markers represent the deployment locations of ASE cameras for the tool's predefined default scenario, detailed in subsequent sections.

The colour-coded markers correlate to distinct enforcement cycles within the tool's operational framework, indicating that the default scenario is designed to encompass four distinct cycles. A closer inspection reveals that each colour marker appears no more than twice within a single ward, suggesting that the default scenario allocates a maximum of two ASE cameras per ward. With ten markers for each colour present on the map, it can be inferred that the scenario anticipates the utilization of a total of 40 ASE deployment points across the city.

An interactive legend is situated on the right side of the user interface, offering the functionality to filter the visualization of markers according to their assigned cycle. Figures 45 and 46 demonstrate this filtering capability, showcasing the selective display of cycles one, and cycles two and four, respectively.

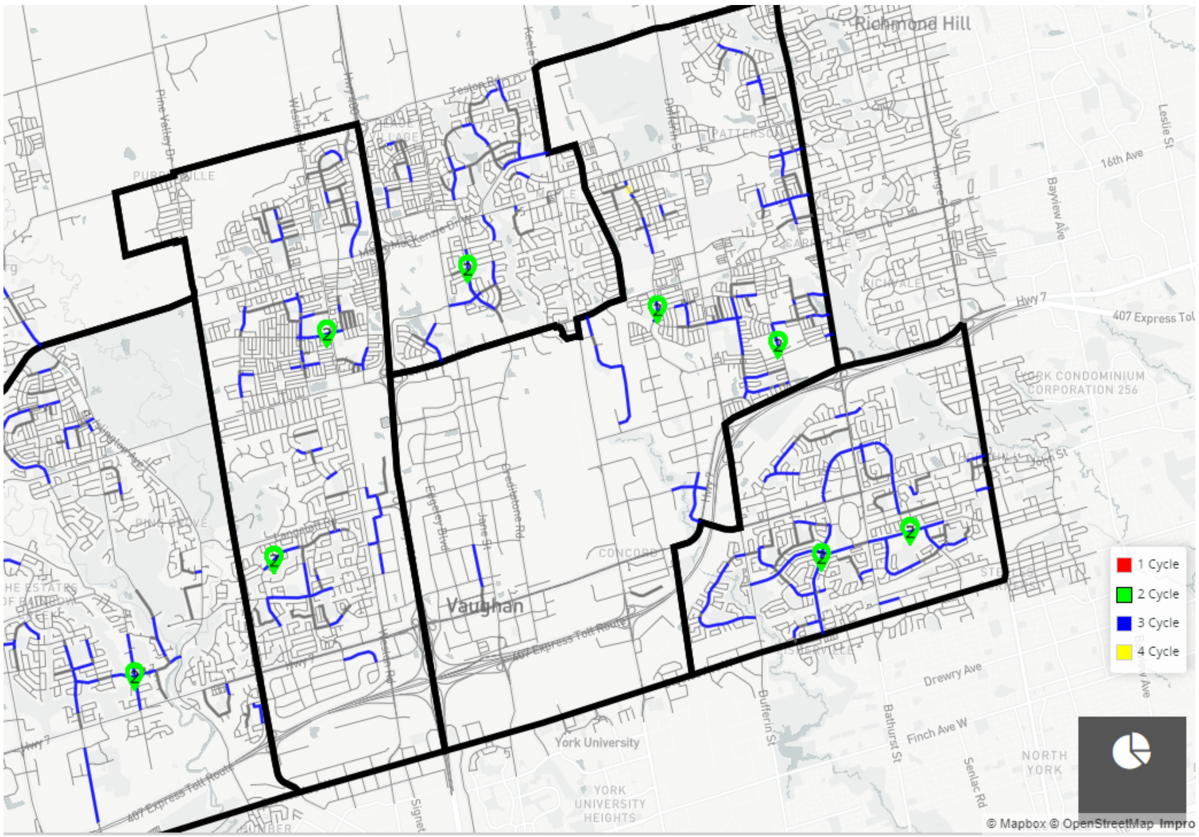

Figure 45: Filter displaying cycle one.

Vaughan's road network is depicted in a subdued gray palette, with certain segments accentuated in bold gray or blue. These prominent segments, regardless of their hue, represent the designated community safety zones as finalized by the City of Vaughan. Those highlighted in blue have achieved a score of 14 or higher within the scoring system of Warrant 2, a metric indicative of higher traffic safety concerns.

By zooming into the map, one can observe that the ASE camera markers are strategically positioned in proximity to the blue road segments. This visual cue highlights the model's emphasis on

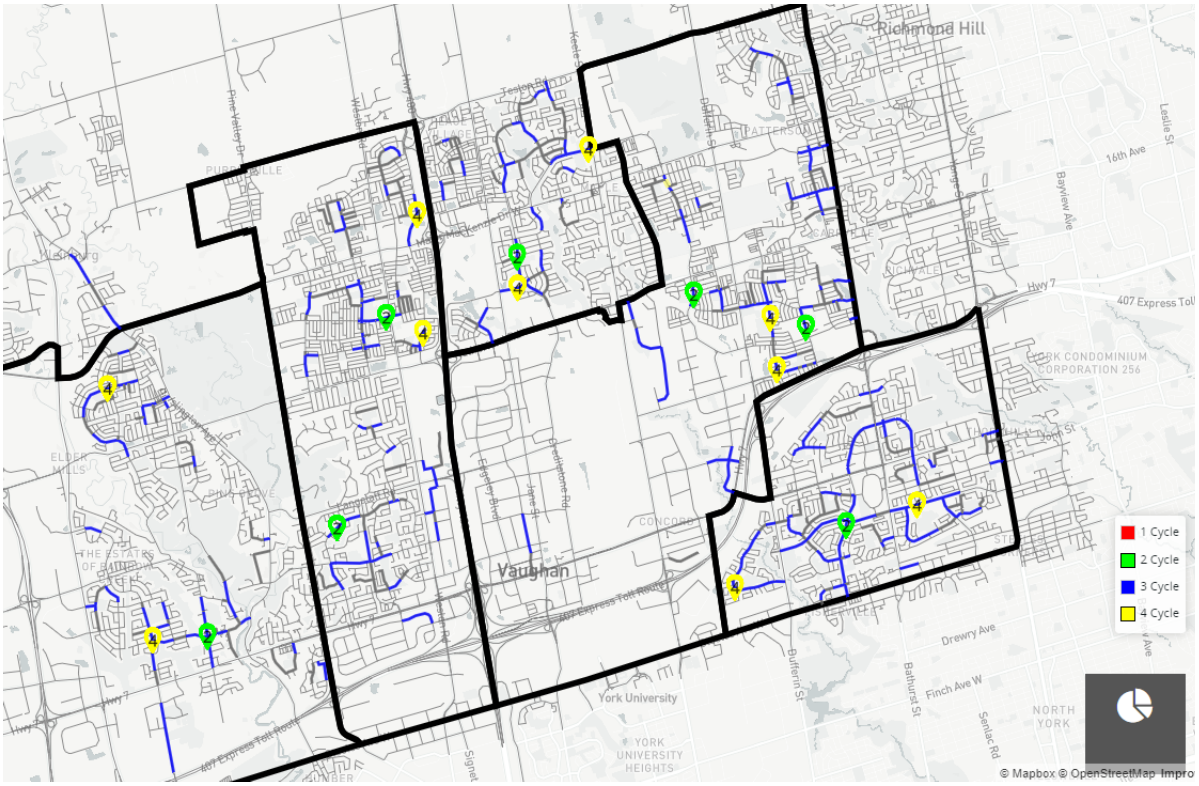

Figure 46: Filter displaying cycles two and four.

prioritizing road segments that have satisfied the criteria of Warrant 2, aligning the deployment of ASE cameras with areas identified as having elevated safety requirements.

5.5.2 Left Panel Interface and Functionalities

The interface features a versatile vertical panel on the left-hand side, comprising six distinct sections: Filters, Road Properties, WYS, Optimization, Community Safety Zone, and Export Panel. Each section allows for interactive engagement with the map, facilitating simultaneous visualization of the resultant changes. The functionality of each panel is elucidated below:

Filtering Panel: Illustrated in Figure 47, this panel contains six toggleable checkboxes enabling users to customize the map layers that are visible. The road network and ward boundaries are selected by default to provide a fundamental framework of the city's geography. The "Speed Radars"

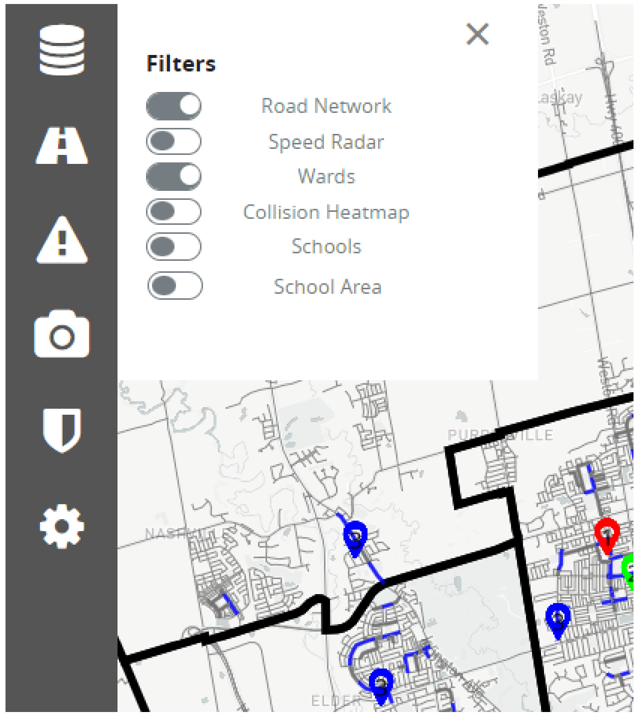

Figure 47: Filtering options in the ASE planning tool.

option pertains to the location data of speed radars discussed earlier. Activation of the ”Heatmap” filter reveals a heatmap overlay, representing the distribution of serious injury collisions, thereby enhancing the user’s comprehension of high-risk areas within the city’s road network. The ”Schools” filter marks the positions of educational institutions. At the same time, the ”School Area” demarcates the vicinity within a specified radius from the center of each school, aiding in identifying zones of heightened pedestrian activity. Figure 48 demonstrates the interface when the ”Collision Heatmap” is enabled and the ”Road Network” is deselected.

Road Properties Section: This panel is designed to assist traffic engineers and city planners in estimating the impact of altering the number of lanes or the speed limit on the community safety zone score and the final schedule of the optimization model. By clicking on any road segment, users can access detailed information, including the total number of collisions, the number of cycling

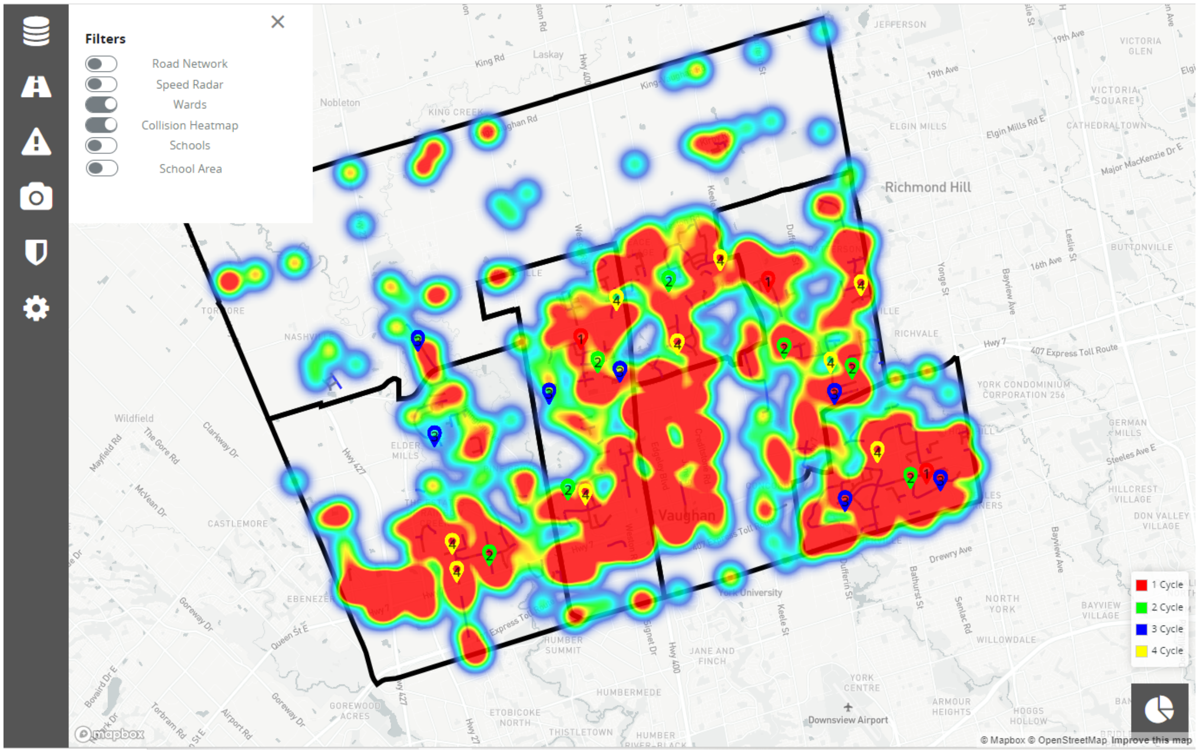

Figure 48: Interactive filtering demonstrated in the ASE planning tool.

and pedestrian collisions, and the number of schools within a 1 km radius. When a road segment is selected, it is highlighted in red for easy identification.

The panel allows users to modify the number of lanes and the speed limit directly. There is also an option to add a toll, which is particularly relevant for freeway segments. After making the desired changes, users can click the ”Update” button. This action applies the new speed limit and number of lanes to the road segment, updating the community safety zone score and the final schedule of the optimization model accordingly.

By providing an interactive and user-friendly interface, this panel facilitates a granular analysis of potential changes, supporting more informed and effective traffic management decisions. The real-time feedback loop ensures that users can quickly assess the consequences of their modifications and make data-driven decisions to enhance road safety and optimize traffic flow.

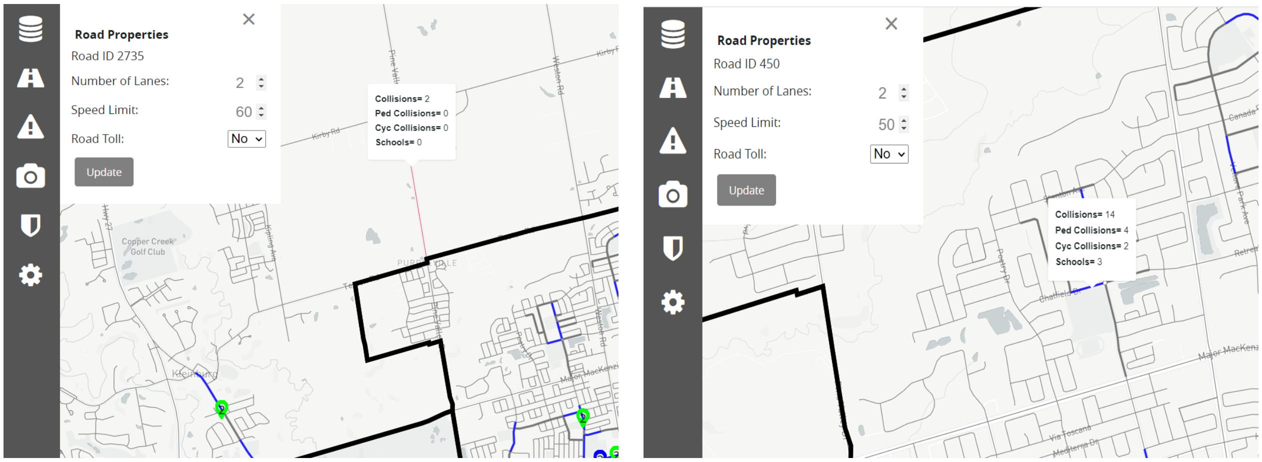

(a) Random road segment (b) Random Community Safety Zone

Figure 49: Updating road properties

WYS Panel/Speed Radar Panel: Currently, the WYS panel, also known as the Speed Radar panel, facilitates the statistical analysis of speed radar data through interactive filtering mechanisms. As depicted in Figure 50, the panel presents an option to toggle the visibility of Speed Radar (SR) statistics, which is deactivated by default. Located at the lower end of the panel are two adjustable sliders; the first slider allows users to filter WYS locations by the average number of violations per day, while the second slider governs the display of locations based on the duration of enforcement. Upon activation of SR statistics, as shown in Figure 52, visual markers—circles centred on SR locations—become evident on the map. It is important to note that for clearer visual interpretation, the Speed Radar layer is also activated within the filtering panel. The absence of a circle at a specific SR location indicates that the selected parameters for either the duration of installation or the average violations do not meet the specified thresholds. The diameter of each circle is proportionally related to the duration of enforcement, with larger circles emerging as the corresponding slider is adjusted to the right. Furthermore, the circle colour coding indicates the violation groupings as defined in the speed radar data analysis section. Locations categorized within the low violation group are represented by yellow circles, medium violations by blue, and

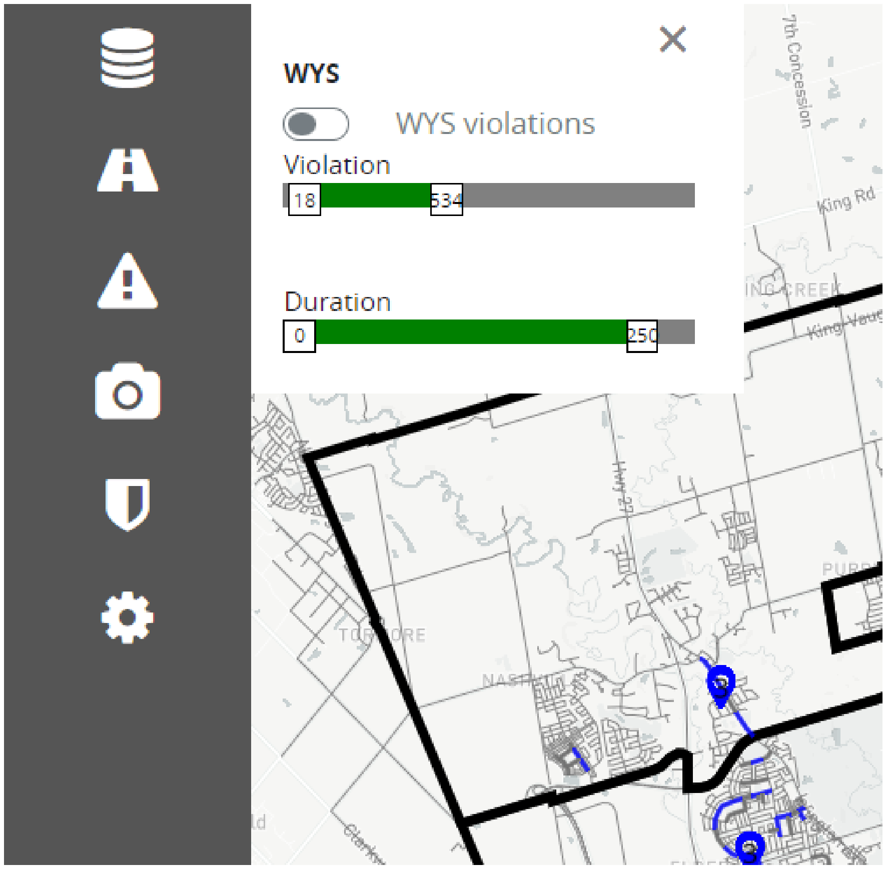

Figure 50: The WYS/Speed Radar Panel interface with statistical visibility options.

high violations by red. Thus, manipulating the violation slider to exclude figures above 380 results in the disappearance of red circles, denoting high violation locations.

While the WYS panel's current implementation offers a foundational set of statistical filtering tools for analyzing speed radar data, this component of the tool is slated for further development. The planned enhancements will aim to refine the user experience, expand the analytical capabilities, and introduce additional layers of interactive data visualization to support more nuanced and comprehensive traffic safety assessments.

Optimization Panel The Optimization Panel is a pivotal feature of the ASE planning tool, designed with the foresight of enabling city planners to meticulously craft ASE enforcement scenarios

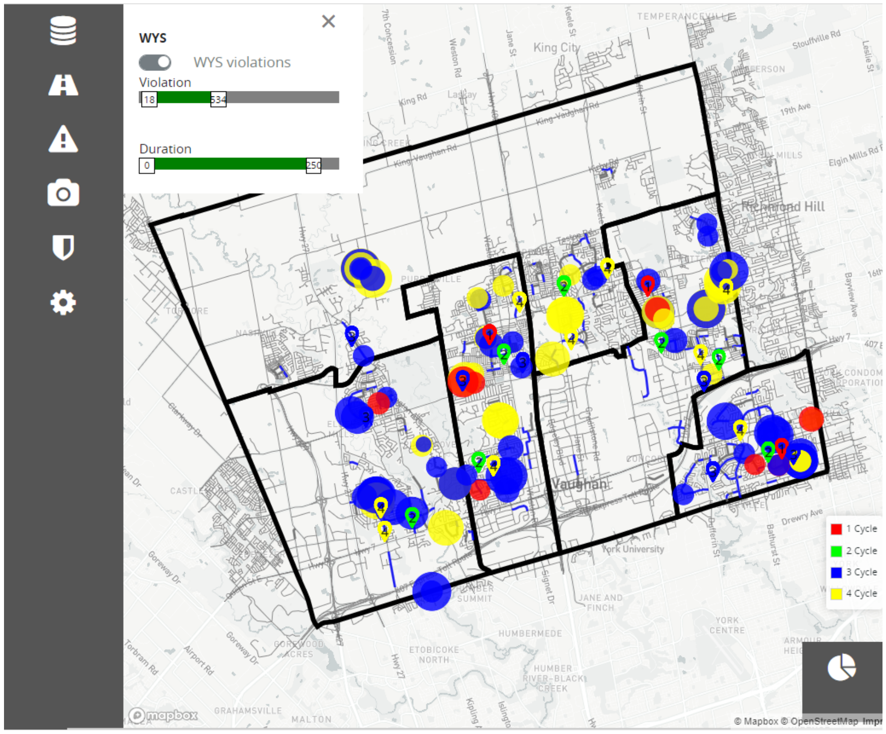

Figure 51: Activation of SR statistics displaying violation circles.

Figure 52: Filtered view highlighting specific SR data

that align with the unique traffic safety objectives of the City of Vaughan. Illustrated in Figure 53, this panel introduces a dual-objective function mechanism at the outset, facilitating a strategic selection between violation reduction and score maximization as primary enforcement goals.

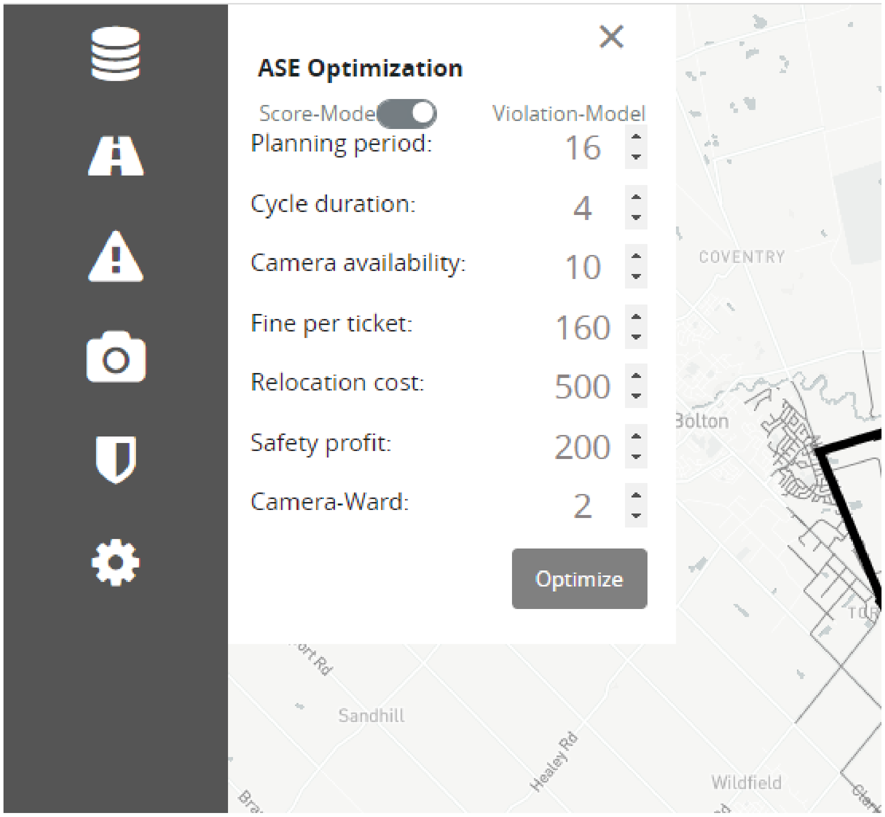

Figure 53: Initial view of the Optimization Panel with objective function toggle.

City planners can toggle between two meticulously defined objective functions within the Optimization Panel's interface. These functions are foundational to formulating ASE enforcement strategies that are responsive and tailored to the City of Vaughan's traffic safety goals. *Violation Reduction Model (Default)*: The primary objective of the default model is to minimize the occurrence of traffic violations across the city, thereby enhancing road safety, or in other words, maximizing the total number of speed violations that will be prevented. Mathematically, this objective is represented as the minimization of the expected number of violations across all road segments and enforcement cycles:

$$\max \sum_{i=1}^{n} \sum_{j=1}^{k} R_{ij} \cdot x_{ij} \tag{32}$$

Here, R_{ij} denotes the predicted number of violations for road segment i during cycle j, and x_{ij} is a binary decision variable indicating whether road segment i is enforced during cycle j (1 if enforced, 0 otherwise).

Score Maximization Model: As an alternative to the violation-centric approach, the score model shifts focus towards maximizing the overall safety score of enforced road segments within the community safety zones. This model's objective function aims to select road segments that collectively achieve the highest safety score, as defined by the criteria set in Warrant 2:

$$\max \sum_{i=1}^{n} \sum_{j=1}^{k} S_{i} \cdot x_{ij} \tag{33}$$

In this equation, S_i represents the safety score assigned to road segment i based on its evaluated safety criteria, and x_{ij} maintains its role as the binary decision variable.

The introduction of these two objective functions within the Optimization Panel underscores the ASE planning tool's versatility. By enabling planners to select between focusing on violation reduction or score maximization, the tool caters to a spectrum of strategic enforcement priorities, ranging from direct safety enhancements to broader, score-based safety improvements.

Following this initial selection, planners are prompted to input seven key parameters:

- Total Planning Period: The overall duration of the enforcement strategy, measured in months.

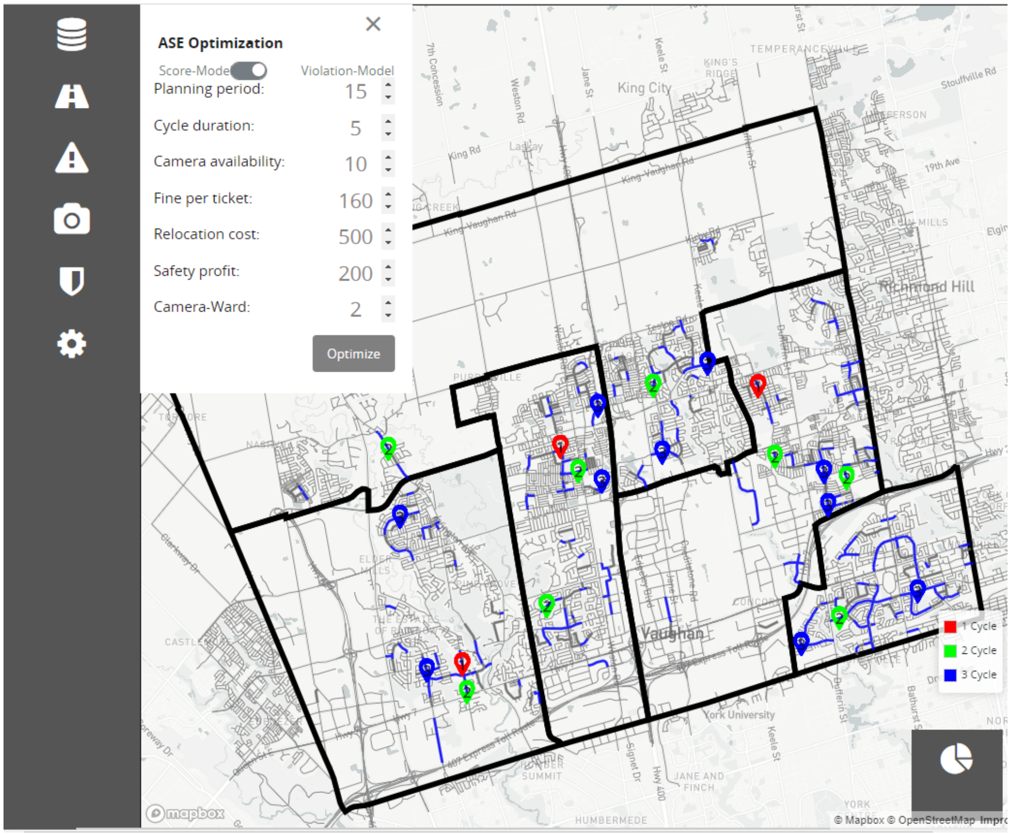

Figure 54: Visualization of a new enforcement scenario developed through the Optimization Panel.

- Cycle Duration: The frequency at which cameras are relocated, defining each enforcement cycle's length.
- Number of Cameras: Specifies the total cameras available for deployment.

The subsequent inputs comprise a weighting system influencing the model's decision-making process:

- Average Revenue per Ticket: This input allows planners to account for the financial aspect of traffic enforcement. Entering "0" disregards revenue considerations entirely.
- Relocation Cost per Camera: Despite the constant nature of this cost, it influences financial reporting without affecting the optimization outcome due to the fixed cycle number.

- Safety Value: Quantifies the importance of preventing violations, with higher values indicating a preference for safety over revenue maximization.

The final input, Camera Distribution per Ward, ensures cameras are allocated equitably across wards. A default setting of ”2” cameras per ward, with a total of ”10” available cameras, indicates a balanced distribution approach.

Executing the optimization with these parameters adjusts the ASE deployment plan, as evidenced in Figure 54. This scenario, specifying a 15-month planning period with 5-month cycles, demonstrates the model's responsiveness. The repositioning of markers and the disappearance of markers from the fourth cycle reflect the tool's capacity to adapt to revised strategic inputs, showcasing the potential for three cycles within the new planning framework.

School Zones Panel: Refining Community Safety Zone Strategies The School Zones Panel within the ASE planning tool serves a dual purpose, aligning with the City of Vaughan's strategic objectives for traffic safety and policy-making. This panel facilitates a dynamic approach to evaluating and adjusting community safety zone (CSZ) parameters, directly influencing future policy development and the efficacy of current safety measures.

1. Policy-making Support: Primarily, the panel assists city planners in revising scoring system thresholds for road segments under Warrant 1, laying the groundwork for future policy adjustments. This capability ensures that policy-making remains responsive to evolving traffic safety needs and data insights.

2. Enhancement of Current CSZ Effectiveness: Additionally, it enables users to modify the impact of individual metrics on the final score for existing CSZs, offering a granular level

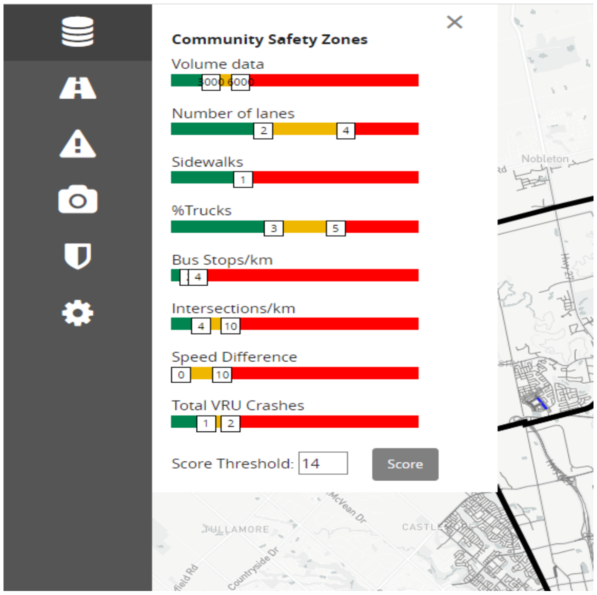

(a) Interactive sliders in the School Zones Panel.

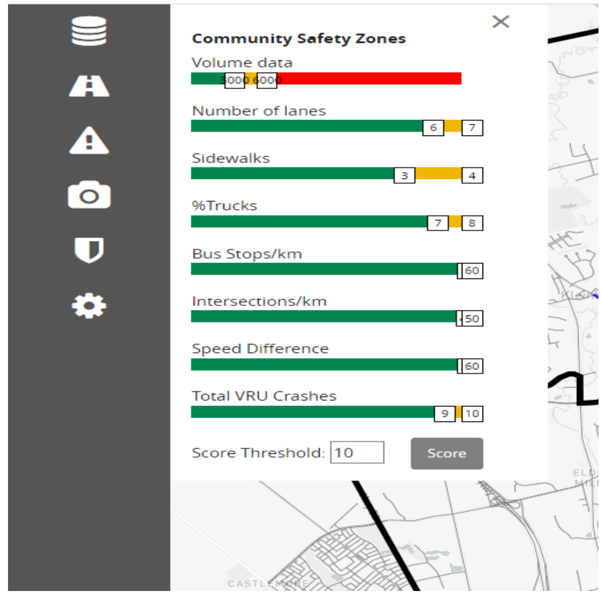

(b) Sliders for a vehicle volume-focused scenario.

Figure 55: Visualization of the School Zones Panel's flexibility in adjusting to new safety scenarios.

of control over safety prioritization. This feature is particularly valuable for fine-tuning the criteria that define CSZs, ensuring that enforcement strategies are both targeted and effective.

As depicted in Figure 55a, the panel presents eight sliders corresponding to the eight metrics outlined in Warrant 2 of the City of Vaughan's community safety zone policy. Traditionally, thresholds for scoring segments within Warrant 1 have been static. However, the tool transitions this method from static to dynamic, empowering city planners to adjust metric thresholds as needed, thereby adapting the scoring system to future safety and policy requirements.

Each slider represents a metric, with the position of the slider's thumb determining the threshold level. Metrics falling within specific ranges are scored accordingly: a green range scores 1, yellow scores 2, and red scores 3. This scoring mechanism is pivotal in identifying road segments eligible for ASE implementation, prioritizing those that meet the requisite score based on adjusted thresholds.Notably, the tool incorporates a provision for segments within finalized CSZs that may not meet Warrant 2 criteria yet are integral to maintaining zone continuity. To address this, a user-input field at the panel's base allows for the specification of a minimum total score for ASE candidacy, ensuring a coherent and logical approach to camera placement. Upon activating the ”calculate” button, the tool processes the user-defined thresholds for each metric, aggregating scores for comprehensive evaluation. Road segments achieving a score equal to or exceeding the user-defined minimum are flagged as potential ASE locations, highlighted in bold blue on the map for easy identification. This nuanced approach enables planners to weight metrics according to strategic safety goals. For example, emphasizing vehicle volume allows for adjusting non-volume metric thresholds to ensure they consistently score 1, thereby focusing on volume data. This methodical adjustment of thresholds ensures that road segments of particular concern, based on specific metrics like volume, are accurately identified for potential ASE deployment. Figures 55a and 55b visually illustrate the slider adjustments for a scenario focused on vehicle volume and the resulting identification of new ASE candidate locations. This comparative analysis with the original scenario (as seen in Figure 44) highlights the tool's flexibility and its profound impact on strategic planning for community safety enhancements.

This detailed exploration of the School Zones Panel underscores its significance in the ASE planning toolkit. It offers a sophisticated means for city planners to refine, evaluate, and optimize community safety zone strategies within the City of Vaughan.

Export Panel: Facilitating Comprehensive Data Analysis and Reporting The Export Panel, delineated in Figure 57, represents a crucial component of the ASE planning tool, specifically engineered to meet the nuanced reporting needs of city planners. This panel is meticulously designed to allow for the exportation of analytical results and scenario evaluations in various formats, accommodating diverse planning and documentation requirements.

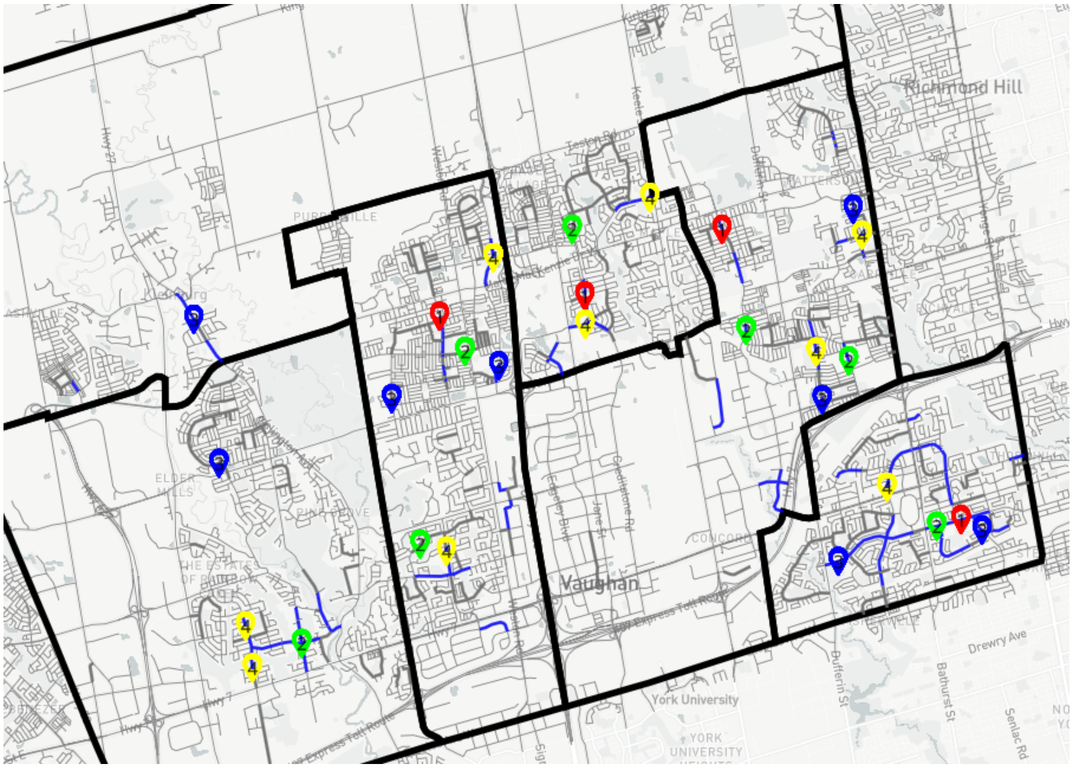

Figure 56: Identification of new potential ASE candidates based on the defined CSZ scenario, illustrating the tool's dynamic adjustment capability

Central to its functionality is the capacity to generate and export comprehensive reports, detailed analyses, and scenario comparisons in formats that include PDF reports, CSV files, and GeoJSON shapefiles. This versatility ensures that planners have the flexibility to choose the format that best suits their needs, whether for in-depth analysis, sharing with stakeholders, or integration with other planning tools.

A standout feature of the Export Panel is its ability to facilitate comparative studies of up to three distinct scenarios within a single report. This feature empowers planners to not only document individual scenario outcomes but also to juxtapose these scenarios side by side, providing a clear, consolidated view of potential strategies and their implications. Planners are prompted to name and save their scenarios, streamlining the process of selecting and compiling the desired report format.

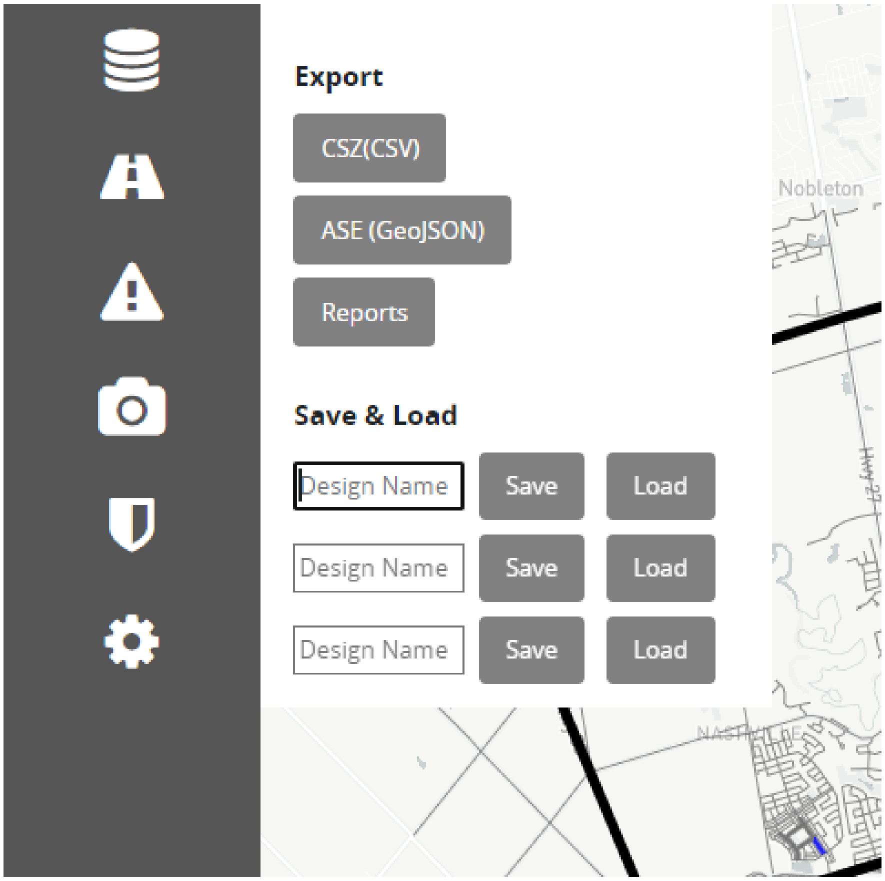

Figure 57: The Export Panel, showcasing the interface for exporting scenario analyses and comparisons.

By enabling the efficient exportation of tailored reports and data files, the Export Panel significantly enhances the planning process. It not only aids in the archival and documentation of strategic planning efforts but also supports informed decision-making by visualizing the potential outcomes and comparative merits of different ASE deployment scenarios. This functionality underscores the panel's value as an indispensable resource for city planners dedicated to optimizing traffic safety measures within the City of Vaughan.

5.5.3 Bottom plots

Having completed the description of the left panel, we now shift our focus to the bottom panel of the tool. This section is a vital component of our interactive mapping interface, designed to provide comprehensive analysis and facilitate informed decision-making processes. Accessible through a click on the "analysis" icon beneath the legend on the right side of the screen, the bottom panel unfolds into four distinct subsections: the ASE Optimization Panel, Community Safety Zone Panel, Ward Analysis, and Property Histogram Panel displayed on Figure 58. Each panel plays a crucial role in the analytical and final scheduling phase, offering a unique set of tools and data visualizations to enhance user interaction and understanding.

ASE Optimization Panel The ASE Optimization Panel is a critical feature designed to display the outcomes of ASE scheduling, updated to reflect the most recent optimization processes. As depicted in Figure 58, this panel provides a visual and interactive summary of the new optimization scenarios we have elaborated upon. It meticulously outlines the operational details for each ASE cycle, including the number of cameras deployed, the anticipated reduction in speed violations due to the ASE enforcement, and the number of tickets that will be issued. Furthermore, it specifies the locations of ASE camera installations across different wards. This panel is not just informational but interactive. It includes a dynamic table that links directly to the map. By selecting any location detail within the table, users can prompt the map to zoom in on the corresponding road segment, which is simultaneously highlighted. Through this integration of detailed data presentation and interactive mapping capabilities, the ASE Optimization Panel is a powerful instrument for understanding and evaluating the effectiveness of speed enforcement strategies within urban environments.

The Community Safety Zone (CSZ) bottom Panel The Community Safety Zone (CSZ) bottom Panel is critical in analyzing and visualizing urban safety metrics, especially concerning road segments within designated safety zones. This panel meticulously displays the outcomes of pre-

Figure 58: The ASE Optimization bottom Panel

defined scenarios for identifying road segments that qualify under the criteria set for community safety zones. Such identification is dynamically updated in response to user interactions, specifically the activation of the 'Calculate' button within the main CSZ panel interface. This user-driven update mechanism ensures that the analysis reflects the most current data and user-defined parameters, enhancing the tool's applicability and relevance in urban planning and safety assessments. The road segments displayed within the CSZ Panel as seen in Figure 59a are meticulously ranked according to their total safety scores, ranging from the highest to the lowest. This ranking facilitates a priority-based approach to safety interventions, allowing urban planners and safety officials to identify and address the most critical areas first. For each road segment, the panel provides a comprehensive summary that includes its designation as a community safety zone, the overall safety score, and an indicator of proximity to schools. This information is crucial for stakeholders interested in implementing targeted safety measures in areas where children and school communities are most affected. Additionally, the CSZ Panel employs an intuitive colour-coding scheme to visually communicate the performance of each road segment across eight critical safety metrics. These colours—green, yellow, and red—correspond to scores of 1, 2, and 3, respectively, based on predefined thresholds. This scheme is aligned with the logic of interactive sliders used elsewhere in the tool, ensuring a consistent and user-friendly experience. Green indicates segments that meet or exceed safety standards; yellow signifies moderate concern, and red alerts users to areas requiring immediate attention. As it is displayed in Figure 59a and 59b, the colours of the metrics vary as the

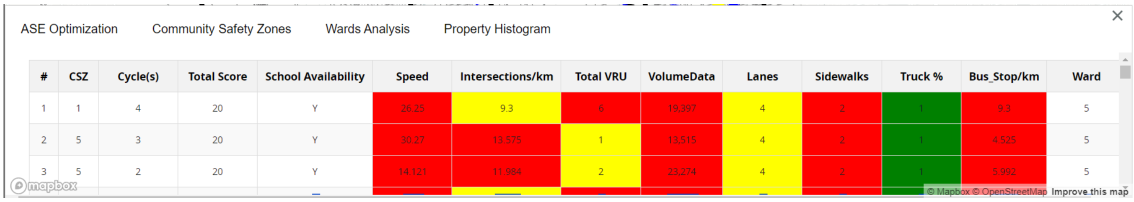

ASE Optimization Community Safety Zones Wards Analysis Property Histogram

#	CSZ	Cycle(s)	Total Score	School Availability	Speed	Intersections/km	Total VRU	VolumeData	Lanes	Sidewalks	Truck %	Bus_Stop/km	Ward
1	1	4	20	Y	26.25	9.3	6	19,397	4	2	1	9.3	5
2	5	3	20	Y	30.27	13.575	1	13,515	4	2	1	4.525	5
3	5	2	20	Y	14.121	11.984	2	23,274	4	2	1	5.992	5

(a) CSZ bottom panel for the default scenario

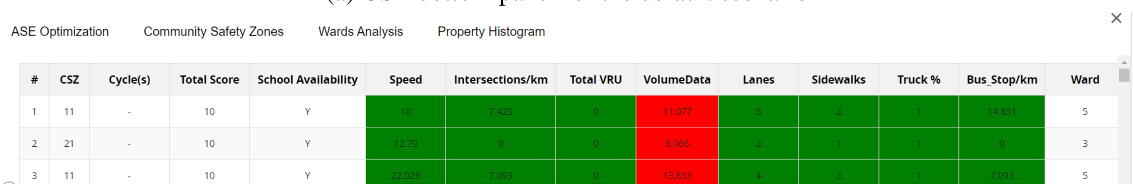

ASE Optimization Community Safety Zones Wards Analysis Property Histogram

#	CSZ	Cycle(s)	Total Score	School Availability	Speed	Intersections/km	Total VRU	VolumeData	Lanes	Sidewalks	Truck %	Bus_Stop/km	Ward
1	11	-	10	Y	10	7.425	0	11,077	5	2	1	14.851	5
2	21	-	10	Y	12.79	0	0	6,066	2	1	1	0	3
3	11	-	10	Y	22.029	7.093	0	13,853	4	2	1	7.093	5

(b) CSZ bottom panel for the Volume concern scenario

Figure 59: Dynamic change of CSZ bottom panel after changing metric thresholds in main CSZ panel

thresholds of the sliders changed from default scenario in Figure 55a to volume focused scenario in Figure 55b. The CSZ Panel's interactivity extends beyond visual analytics. By clicking on any row within the table, users can instantaneously zoom in on the corresponding road segment on the map, with the segment highlighted for easy identification. This feature enhances the tool's navigability and allows users to correlate tabulated data with real-world geography. Such integration between tabular data and geospatial visualization underscores the panel's role in facilitating a deeper understanding of urban safety dynamics, enabling stakeholders to make informed decisions based on comprehensive analyses and interactive data exploration.

Ward Analysis The Ward Analysis section is ingeniously designed to facilitate a comparative analysis of safety improvements and potential revenue generation across different wards. This comparison is pivotal for strategic planning and prioritization in deploying ASE systems. As illustrated in Figure 60, this section visually represents its analyses through two pie charts and five bar charts each for one ward. The first pie chart delineates the predicted number of speed violations that could be averted thanks to the proactive measures suggested by the ASE optimization section. Conversely, the second pie chart forecasts the potential revenue through the predicted number of

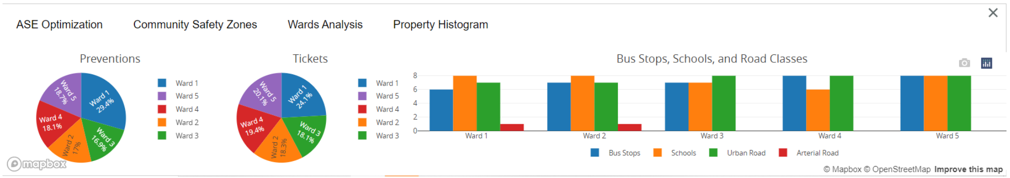

Figure 60: Ward analysis bottom Panel

tickets to be issued post-ASE deployment. These visualizations provide a straightforward, at-a-glance understanding of the impact of ASE systems, aiding in the equitable distribution of traffic safety measures and resource allocation among the wards. This section will expand to encompass broader and more detailed visualizations and comparative analyses between wards using bar charts. These enhancements will facilitate deeper insights into safety and revenue aspects and provide statistics regarding the total number of urban road segments and arterial roads determined for ASE installation during the planning period.

Histogram Property Moving to the Histogram Property section, this part of the tool is meticulously crafted to enhance the comprehension of the distribution of metric values used in the Community safety zone panel. Figure 61 features a histogram focusing on the volume metric as a case study. The x-axis of the histogram represents the volume values, while the y-axis indicates the frequency of these values across all analyzed road segments. The colour coding of the bars enriches the visual data representation, correlating each bar's colour with the score attributed to the respective volume value. This colour-coded approach not only simplifies the understanding of data distribution but also aligns with the tool's overarching design philosophy, ensuring consistency in user experience. By providing a visual breakdown of metric distributions, this section empowers users to identify patterns, outliers, and critical areas of interest within the dataset, thereby facilitating a more granular and informed analysis of community safety zones. Figures 62a and 62b show

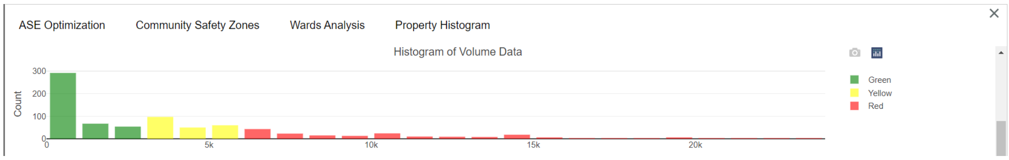

Figure 61: Color coded histogram of volume metric

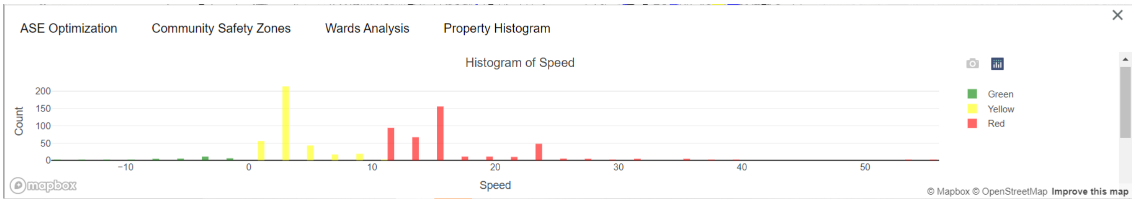

(a) Histogram of speed in default scenario

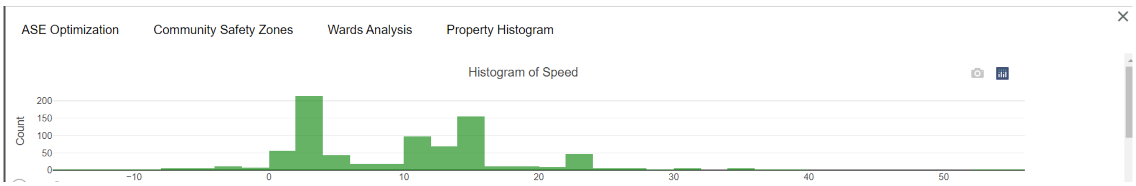

(b) Histogram of speed in volume focus scenario

Figure 62: Dynamic change of histogram colours with the change in CSZ thresholds.

how the colours change in the speed histogram between the default scenario and the volume focus scenario.

The ASE tool is a big step in making our cities safer and helping manage traffic better. It uses maps and data to show us where roads can be safer and how violating the speed limit can be reduced. As we plan to add more features, like better ways to compare different areas and make the tool more interactive, it will become even more useful. This tool is really about helping city planners and people who make laws to understand what's happening on the streets and make better decisions. It's a practical way to use technology to tackle real-world problems, aiming to make our communities safer for everyone.

6 Conclusion

This study explored a novel approach to optimizing Automated Speed Enforcement (ASE) scheduling in urban areas to enhance road safety. In North America, municipalities use ASE systems as a key strategy to reduce traffic violations and enhance safety, especially in school zones and community safety zones. Despite their proven effectiveness, challenges persist in optimizing the placement and scheduling of ASE cameras due to limited resources and the dynamic nature of driver behavior. Our research developed and validated a model based on the Markov Decision Process (MDP) to address these challenges. This model considers various factors influencing ASE effectiveness, including traffic patterns, violation rates, and camera placement strategies. It aims to minimize the total number of speed violations over a planning period by strategically deploying and relocating ASE cameras across different locations.

To consolidate these findings and provide a comprehensive understanding of our approach, we systematically defined the necessary metrics, evaluated ASE effectiveness across multiple sites, predicted future outcomes using advanced machine learning models, and optimized scheduling through the Markov Decision Process (MDP) framework.The key components of this study are:

- **Safety Metric Definition:** We began by defining a robust safety metric to evaluate ASE effectiveness. Speed violations were selected as the primary metric due to their immediacy and practicality in indicating the impact of ASE systems in urban areas. This metric aligns with the need to measure effectiveness over shorter observation periods, especially when cameras are periodically relocated.

- **Evaluation of ASE Effectiveness:** We conducted a detailed evaluation of ASE effectiveness at various sites using linear regression analysis to determine the trend of enforcement impact over time. By assessing the slope of regression lines for each site, we categorized locations into groups based on their effectiveness levels. Further clustering using the K-means algorithm refined these categories into ”Not Effective,” ”Moderately Effective,” and ”Highly Effective” groups, enhancing the precision of our predictive models. This method provided a clearer understanding of the effectiveness patterns and allowed for more targeted feature engineering to improve the accuracy of future ASE deployments.

- **Prediction of ASE Effectiveness:** This study integrated comprehensive traffic data, including road characteristics and traffic volume, speed limits, road types, number of lanes, and proximity to schools with each ASE site location to analyze factors influencing ticketing patterns. This enriched dataset was used to train multiple machine learning models, with the Random Forest Classifier achieving the highest performance, with a mean accuracy of 0.68, validated using precision, recall, and F1-score metrics. The model's robustness was further ensured by addressing class imbalance with the SMOTETomek method and employing cross-validation techniques. Applying this model to new data from the City of Vaughan confirmed its broader applicability, highlighting areas where enforcement strategies could be improved and reinforcing the value of data-driven approaches in enhancing traffic safety and speed enforcement.

- **Optimization of ASE Scheduling:** The study demonstrated the application of a Markov Decision Process (MDP) framework to optimize ASE scheduling. By simulating various deployment scenarios, we found that the model tends to allocate longer enforcement periods to highly effective locations while strategically distributing the remaining resources among moderately and non-effective locations. This approach balances maximizing safety benefits and maintaining unpredictability in enforcement to prevent driver adaptation.

6.1 Future of Work

For future work, a primary focus should be on obtaining more precise data that captures the full distribution of vehicle speeds, rather than solely documenting instances of speed violations. This approach would provide a more holistic understanding of speed behavior across different road segments and conditions. Additionally, implementing hidden cameras to collect speed data before and after enforcement periods could offer valuable insights into changes in speed patterns over time. It is important to emphasize that these hidden cameras would be used exclusively for data collection purposes and not for issuing any tickets.

The data gathered from hidden cameras could facilitate the development of accurate probability state matrices for periods without enforcement, allowing for a more nuanced analysis of speed behavior. Further research could also explore enhancing the objective function by considering the spillover effects of enforcement at one location on traffic behavior in nearby areas. For example, enforcing speed limits on a particular road segment may cause drivers to reroute to alternative paths, altering traffic volumes and potentially impacting road safety conditions in adjacent areas. Incorporating these dynamic interactions would lead to more effective and strategically targeted enforcement measures.

Furthermore, gathering more comprehensive and detailed data would improve the predictive models used to assess ASE effectiveness, enabling more accurate and precise predictions. By refining these models with richer datasets, future studies can better understand the factors influencing enforcement success, optimize resource allocation, and ultimately enhance road safety outcomes.

Appendices

A Appendix 1

Proofs

Exponentiation of Transition Probability Matrices

Theorem 1 *If P is a transition probability matrix, then P^k is also a transition probability matrix for any positive integer k.*

The proof proceeds by induction on k.

Base Case ($k = 1$): Trivial, as $P^1 = P$ is a transition probability matrix by definition.

Inductive Step: Assume P^k is a transition probability matrix for some $k \geq 1$. We need to show P^{k+1} is also a transition probability matrix.

Consider $P^{k+1} = P^k \cdot P$. Each element of P^{k+1} is given by:

$$(P^{k+1})_{ij} = \sum_{r=1}^{n} (P^k)_{ir} P_{rj}$$

Since P^k is row stochastic, we have:

$$\sum_{r=1}^{n} (P^k)_{ir} = 1 \quad \text{and} \quad (P^k)_{ir} \geq 0$$

Each $P_{rj} \geq 0$ and $\sum_{j=1}^{n} P_{rj} = 1$. Therefore, we have:

$$\sum_{j=1}^{n} (P^{k+1})_{ij} = \sum_{j=1}^{n} \sum_{r=1}^{n} (P^k)_{ir} P_{rj} = \sum_{r=1}^{n} (P^k)_{ir} \sum_{j=1}^{n} P_{rj} = \sum_{r=1}^{n} (P^k)_{ir} = 1$$

Also, $(P^{k+1})_{ij} \geq 0$ because it is a sum of non-negative terms.

By induction, P^k is a transition probability matrix for all k.

Multiplication of Transition Probability Matrices

Theorem 2 *If A and B are transition probability matrices, then AB is also a transition probability matrix.*

Consider the matrices A and B, each with non-negative entries where each row sums to one.

Each element of AB is given by:

$$(AB)_{ij} = \sum_{k=1}^{n} A_{ik} B_{kj}$$

Since A is row stochastic, $A_{ik} \geq 0$ and $\sum_{k=1}^{n} A_{ik} = 1$ for all i. Since B is row stochastic, $B_{kj} \geq 0$ and $\sum_{j=1}^{n} B_{kj} = 1$ for all k.

Therefore, we have:

$$\sum_{j=1}^{n} (AB)_{ij} = \sum_{j=1}^{n} \sum_{k=1}^{n} A_{ik} B_{kj} = \sum_{k=1}^{n} A_{ik} \sum_{j=1}^{n} B_{kj} = \sum_{k=1}^{n} A_{ik} = 1$$

Also, $(AB)_{ij} \geq 0$ as it is a sum of non-negative products.

B Appendix 2

B.1 Random Forest Classifier

Random Forest is an ensemble learning method that builds multiple decision trees during training and merges their results to improve accuracy and control overfitting.

- **Training**: Each tree in the forest is trained on a random subset of the training data using bootstrap sampling. The algorithm splits the data at each node based on the feature that provides the best split according to a criterion like Gini impurity or entropy. The Gini impurity is calculated as follows:

$$\text{Gini impurity} = 1 - \sum_{i=1}^{n} p_i^2 \tag{34}$$

  where p_i is the probability of an instance being classified into a particular class.

- **Prediction**: For classification, each tree in the forest casts a vote for the class, and the class with the most votes becomes the model's prediction.

B.2 XGBoost Classifier

XGBoost (Extreme Gradient Boosting) is a highly efficient and scalable implementation of gradient-boosted decision trees. It builds trees sequentially, where each tree tries to correct the errors of the previous ones.

- **Training**: The algorithm minimizes a loss function (like logistic loss for classification) and adds trees to the model to reduce this loss. Regularization techniques are used to prevent overfitting. The logistic loss function is given by:

$$\text{Logistic Loss} = -\sum_{i=1}^{N} [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)] \tag{35}$$

where y_i is the actual class label, p_i is the predicted probability, and N is the number of instances.

- **Prediction**: The predictions from all trees are summed to provide the final prediction, with more recent trees having a greater influence.

B.3 Logistic Regression

Logistic Regression is a statistical model that uses a logistic function to model a binary dependent variable. It estimates the probability that a given input point belongs to a certain class.

- **Training**: The algorithm finds the best-fitting model to describe the relationship between the independent variables and the dependent binary outcome by maximizing the likelihood of the observed data. The logistic function is defined as:

$$\sigma(z) = \frac{1}{1 + e^{-z}} \tag{36}$$

  where $z = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_n x_n$, with $\beta_0, \beta_1, \ldots, \beta_n$ being the coefficients and $x_1, x_2, \ldots, x_n$ being the input features.

- **Prediction**: The logistic function outputs a probability between 0 and 1, which is then thresholded to decide the class.

B.4 Support Vector Machine (SVM)

SVM is a supervised learning model that analyzes data for classification and regression analysis by finding the hyperplane that best separates the classes.

- **Training**: SVM finds the optimal hyperplane that maximizes the margin between the nearest points of the classes, called support vectors. The hinge loss function is used for training:

$$\text{Hinge Loss} = \sum_{i=1}^{N} \max(0, 1 - y_i \cdot (w \cdot x_i + b)) \tag{37}$$

  where y_i is the actual class label, x_i is the input feature vector, w is the weight vector, b is the bias term, and N is the number of instances.

- **Prediction**: The model predicts the class of a new input based on which side of the hyperplane the point falls.

B.5 K-Nearest Neighbors (KNN)

KNN is a non-parametric method used for classification and regression. It predicts the class of a data point based on the classes of its k nearest neighbors.

- **Training**: KNN does not have a training phase per se, as it stores the entire training dataset.

- **Prediction**: For a new data point, the algorithm finds the k nearest points in the training data and assigns the class that is most common among them. The Euclidean distance is commonly used to measure the closeness of points:

$$\text{Distance} = \sqrt{\sum_{i=1}^{n} (x_i - y_i)^2} \tag{38}$$

where x_i and y_i are the feature values of the instances.

B.6 Gradient Boosting Classifier

Gradient Boosting is an ensemble technique that builds models sequentially, each trying to correct the errors of the previous models.

- **Training**: The algorithm starts with an initial model (like a decision tree) and adds new models that predict the residuals (errors) of the existing model. The loss function for training can be mean squared error:

$$\text{Loss} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \tag{39}$$

  where y_i is the actual value, $\hat{y}_i$ is the predicted value, and N is the number of instances.

- **Prediction**: The predictions from all models are combined to produce the final output, typically using a weighted sum.

B.7 LightGBM Classifier

LightGBM (Light Gradient Boosting Machine) is a gradient boosting framework that uses tree-based learning algorithms. It is designed to be efficient and scalable.

- **Training**: Similar to XGBoost, LightGBM builds trees sequentially and optimizes a loss function. It uses histogram-based algorithms and leaf-wise tree growth to speed up training and reduce memory usage.

- **Prediction**: Predictions are made by summing the outputs of all trees, with each tree contributing to the final decision.

Bibliography

[1] W. H. Organization, "Global status report on road safety 2015," 2015.

[2] M.-Å. Belin, P. Tillgren, and E. Vedung, "Vision zero–a road safety policy innovation," *International Journal of Injury Control and Safety Promotion*, pp. 1–9, 2012.

[3] S. Bowen, "Moral philosophy," 2005.

[4] M.-Å. Belin, *Public Road Safety Policy Change and its Implementation–Vision Zero a road safety policy innovation*. PhD thesis, Karolinska Institutet, 2005.

[5] S. Government, "Bill 1996/97:137: Vision zero and the traffic safety work of the future," tech. rep., 1997.

[6] W. Haddon Jr, "Changing approaches to epidemiology, prevention, and amelioration of trauma: transition to approaches etiologically rather than descriptively based," *American Journal of Public Health and the Nation's Health*, vol. 58, no. 10, pp. 1943–1955, 1968.

[7] ASE Ontario, "Ase ontario," 2024. Accessed: 2024-07-27.

[8] P. D. Thomas and L. E. Decina, "A review of the effectiveness of speed enforcement cameras in reducing road traffic collisions," *Traffic Injury Prevention*, vol. 9, no. 1, pp. 37–43, 2008.

[9] L. E. Decina and P. D. Thomas, "Automated enforcement of traffic signals: A literature review," *Transportation Research Record*, vol. 2019, pp. 10–15, 2007.

[10] H. W. Mountain, L.J. and M. Maher, "The relationship between speed and accidents," *Accident Analysis & Prevention*, vol. 37, no. 4, pp. 731–741, 2005.

[11] C. A. Lave and L. B. Lave, "Did the 65 mph speed limit save lives?," *Accident Analysis & Prevention*, vol. 31, no. 1, pp. 1–7, 1999.

[12] J. Stuster and Z. Coffman, "Synthesis of safety research related to speed and speed limits," *Federal Highway Administration*, 1998.

[13] G. T. F. Skszek, Stephen and W. W. Recker, "An empirical assessment of the relationship between speed and roadway safety," *Transportation Research Record*, vol. 1897, pp. 36–43, 2004.

[14] Y.-J. Kweon and K. M. Kockelman, "Safety impacts of speed limit changes: use of empirical bayes method," *Transportation Research Record*, vol. 1908, pp. 45–53, 2005.

[15] W. Snyder, "Speed dispersion and its effect on highway fatalities: Re-evaluating lave's model," *Accident Analysis & Prevention*, vol. 21, no. 1, pp. 1–5, 1989.

[16] N. J. Garber and R. M. Gadiraju, "Traffic and highway safety: An analysis of speed limits," *Virginia Transportation Research Council*, 1989.

[17] N. J. Garber and A. A. Ehrhart, "Effectiveness of speed monitoring displays in increasing speed limit compliance in highway work zones," *Transportation Research Record*, vol. 1723, pp. 7–15, 2000.

[18] D. T. Levy and P. Asch, "The lack of speed coordination implies greater risk: A re-evaluation of lave's model," *Journal of Safety Research*, vol. 20, no. 2, pp. 53–63, 1989.

[19] R. A. Retting, C. M. Farmer, and A. T. McCartt, "Evaluation of automated speed enforcement in montgomery county, maryland," *Traffic injury prevention*, vol. 9, no. 5, pp. 440–445, 2008.

[20] K. Shin, S. P. Washington, and I. van Schalkwyk, “Evaluation of the scottsdale loop 101 automated speed enforcement demonstration program,” *Accident Analysis & Prevention*, vol. 41, no. 3, pp. 393–403, 2009.

[21] A. Montella, B. Persaud, M. D’Apuzzo, and L. L. Imbriani, “Safety evaluation of automated section speed enforcement system,” *Transportation research record*, vol. 2281, no. 1, pp. 16–25, 2012.

[22] D. A. Quistberg, L. L. Thompson, J. Curtin, F. P. Rivara, and B. E. Ebel, “Impact of automated photo enforcement of vehicle speed in school zones: interrupted time series analysis,” *Injury prevention*, vol. 25, no. 5, pp. 400–406, 2019.

[23] C. Wang, C. Xu, and P. Fan, “Effects of traffic enforcement cameras on macro-level traffic safety: A spatial modeling analysis considering interactions with roadway and land use characteristics,” *Accident Analysis & Prevention*, vol. 144, p. 105659, 2020.

[24] C. Wilson, C. Willis, J. K. Hendrikz, R. Le Brocque, and N. Bellamy, “Speed cameras for the prevention of road traffic injuries and deaths,” *Cochrane database of systematic reviews*, no. 11, 2010.

[25] H. Li, D. J. Graham, and A. Majumdar, “The impacts of speed cameras on road accidents: An application of propensity score matching methods,” *Accident Analysis & Prevention*, vol. 60, pp. 148–157, 2013.

[26] L. Carnis, “Automated speed enforcement: What the french experience can teach us,” *Journal of Transportation Safety & Security*, vol. 3, no. 1, pp. 15–26, 2011.

[27] K. Jeong-Gyu, “Changes of speed and safety by automated speed enforcement systems,” *IATSS research*, vol. 26, no. 2, pp. 38–44, 2002.

[28] L. J. Thomas, R. Srinivasan, L. E. Decina, and L. Staplin, “Safety effects of automated speed enforcement programs: critical review of international literature,” *Transportation Research Record*, vol. 2078, no. 1, pp. 117–126, 2008.

[29] A. Gains, M. Nordstrom, J. SHREWSBURY, L. MOUNTAIN, and M. MAHER, “The national safety camera programme-four-year evaluation report-december 2005,” 2005.

[30] C. J. Rodier, S. A. Shaheen, and E. Cavanagh, “Automated speed enforcement in the us: a review of the literature on benefits and barriers to implementation,” in *Transportation Research Board 87th Annual Meeting. CD-ROM. Washington, DC*, 2007.

[31] A. Montella, L. L. Imbriani, V. Marzano, and F. Mauriello, “Effects on speed and safety of point-to-point speed enforcement systems: Evaluation on the urban motorway a56 tangenziale di napoli,” *Accident analysis & prevention*, vol. 75, pp. 164–178, 2015.

[32] R. A. Retting and C. M. Farmer, “Evaluation of speed camera enforcement in the district of columbia,” *Transportation Research Record*, vol. 1830, no. 1, pp. 34–37, 2003.

[33] R. Tay, “Speed cameras improving safety or raising revenue?,” *Journal of Transport Economics and Policy (JTEP)*, vol. 44, no. 2, pp. 247–257, 2010.

[34] H. Li, Y. Zhang, and G. Ren, “A causal analysis of time-varying speed camera safety effects based on the propensity score method,” *Journal of safety research*, vol. 75, pp. 119–127, 2020.

[35] J.-P. Moon and J. E. Hummer, "Speed enforcement cameras in charlotte, north carolina: Estimation of longer-term safety effects," *Transportation research record*, vol. 2182, no. 1, pp. 31–39, 2010.

[36] P. Pilkington and S. Kinra, "Effectiveness of speed cameras in preventing road traffic collisions and related casualties: systematic review," *Bmj*, vol. 330, no. 7487, pp. 331–334, 2005.

[37] L. Mountain, W. Hirst, and M. Maher, "Are speed enforcement cameras more effective than other speed management measures?: The impact of speed management schemes on 30 mph roads," *Accident Analysis & Prevention*, vol. 37, no. 4, pp. 742–754, 2005.

[38] H. Li and D. J. Graham, "Heterogeneous treatment effects of speed cameras on road safety," *Accident Analysis & Prevention*, vol. 97, pp. 153–161, 2016.

[39] A. M. Novoa, K. Pérez, E. Santamariña-Rubio, M. Mar´i-Dell'Olmo, and A. Tob´ias, "Effectiveness of speed enforcement through fixed speed cameras: a time series study," *Injury prevention*, vol. 16, no. 1, pp. 12–16, 2010.

[40] E. De Pauw, S. Daniels, T. Brijs, E. Hermans, and G. Wets, "An evaluation of the traffic safety effect of fixed speed cameras," *Safety science*, vol. 62, pp. 168–174, 2014.

[41] F. Wegman and C. Goldenbeld, "Speed management: enforcement and new technologies," 2006.

[42] R. Factor, N. Haviv, and G. Keren, "Enforcement and behavior: the effects of suspending enforcement through automatic speed cameras," *Journal of experimental criminology*, vol. 19, no. 3, pp. 743–759, 2023.

[43] A. M. Kim, X. Wang, K. El-Basyouny, and Q. Fu, "Operating a mobile photo radar enforcement program: A framework for site selection, resource allocation, scheduling, and evaluation," *Journal of Transportation Safety & Security*, vol. 8, no. 3, pp. 257–280, 2016.

[44] A. Talebpour, H. S. Mahmassani, and A. Elfar, "A nonlinear integer programming model for the optimization of police speed camera enforcement locations," *Transportation Research Part C: Emerging Technologies*, vol. 58, pp. 106–123, 2014.

[45] J. Ambros, J. Elgner, R. Turek, and V. Valentová, "Where and when do drivers speed? a feasibility study of using probe vehicle data for speed enforcement planning," *Accident Analysis & Prevention*, vol. 105, pp. 128–137, 2017.

[46] R. Gang and Z. Zhuping, "Traffic safety forecasting method by particle swarm optimization and support vector machine," *Expert Systems with Applications*, vol. 38, no. 9, pp. 10983–10990, 2011.

[47] P. Rogerson, M. Cameron, and A. Graham, "Effectiveness of speed cameras in reducing accidents in victoria, australia," *Accident Analysis & Prevention*, vol. 26, no. 4, pp. 543–550, 1994.

[48] A. Delaney, H. Ward, M. Cameron, and A. F. Williams, "Controversies and speed cameras: Lessons learnt internationally," *Journal of Safety Research*, vol. 36, no. 4, pp. 283–287, 2005.

[49] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[50] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794, ACM, 2016.

[51] D. W. Hosmer, S. Lemeshow, and R. X. Sturdivant, *Applied logistic regression*. Wiley, 2013.

[52] C. Cortes and V. Vapnik, “Support-vector networks,” *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[53] T. M. Cover and P. E. Hart, “Nearest neighbor pattern classification,” *IEEE Transactions on Information Theory*, vol. 13, no. 1, pp. 21–27, 1967.

[54] O. Berman and A. R. Odoni, “Locating mobile servers on a network with markovian processes,” 1978.