How To Get Over Traffic: Optimizing High-Traffic Areas Using Bridges

Bachelor's Project Thesis

Darie Petcu, s3990044, Petcu@student.rug.nl,
Supervisors: Prof. Dr. L.C. Verbrugge

Abstract: Traffic optimization research has shown little diversity in recent years, with traffic lights seemingly being the only topic of interest. This project shifts the focus towards other means of controlling traffic (sign priority, right-of-way, adding overpasses), covering scenarios where traffic lights might not be suitable or necessary. An Agent-Based Model was employed to simulate the traffic. Two intersections governed by sign priority were analyzed for possible improvements, by reassigning the priority road and by adding an overpass. Traffic flow was measured before and after these changes were made to determine the efficacy of the modifications. The aim of this project is to act as a proof of concept, establishing whether these simulations can be used for assessing real-life infrastructure projects. Real-world data was used to simulate two areas containing busy intersections in the city of Groningen and investigate optimization prospects. The results showed minor improvements for one of the modelled intersections with crowded traffic conditions. However, the other intersection experienced significant deterioration of car flow with these changes, regardless of traffic conditions. The contradictory nature of these results has been attributed to the confounding variables that differentiate between the intersections. Further investigation is required to isolate the effectiveness of the simulated modifications from the local particularities of each intersection.

1 Introduction

In today's world, the number of people living in large cities keeps growing, and with them the road networks are becoming increasingly complex. Inhabitants of urban settlements are faced with the challenges of commuting on most days of their lives, be it within one neighbourhood or across town. Local administrations are thus faced with the task of enabling their citizens to efficiently move around town, with reasons ranging from environmental to time saving.

The need to optimize traffic

However, the means of fulfilling this are somewhat limited. Streets are generally located between constraints, such as buildings, that prevent the free moving and reshaping of road networks. As they can seldom be entirely modified, the task at hand comes down to selecting the optimal configuration of intersection types in an area to allow as many traffic participants as possible to pass through. These intersection types are roundabouts or standard intersections, the latter of which can also be governed by various types of priority rules: traffic signs, traffic lights, right-of-way.

Alternatively, some intersections can be further optimized by the addition of multi-level crossings, such as overpasses or underground tunnels. This measure is usually justified in busy areas, where it enables one axis of transit along the intersection to completely bypass the bottleneck effect of the intersection. This has beneficial effects not only for people travelling along the overpass, but also for the other traffic participants from that intersection, which becomes overall less crowded.

The need to use agent-based simulations

The aforementioned complexity of modern traffic poses a serious question of what kind of infrastructure to implement for optimization. If decisions concerning road ensembles in the past were fairly straight-forward, present-day traffic is more diverse and complex than ever before. This means that the solutions civil engineers must find become dependent on multiple variables, including the number of cars coming and going in each direction, the effects of rush hour on the traffic flow, and overall safety.

Variation of all these factors cannot be simultaneously taken into account accurately via mental deliberation, so making use of simulations is a reliable and reproducible way to model traffic and analyze the up- and downsides of the alternatives proposed by civil engineers. Alternatives to simulations also include equation-based models or experiments; however, simulations have the advantage of a visual output. Using an agent-based model (ABM) allows one to closely observe the emergent behaviour of the traffic, as well as predict the traffic flow in the real world upon implementation of the simulated alterations to the roads (Bazghandi, 2012).

Another trait of traffic that can be naturally investigated via ABM is people's subjectivity when it comes to route choice, traffic behaviour, and human error. All of these factors can be modelled as part of the car agents. The emergent behaviour that arises, such as gridlocks in intersections where cars are simultaneously turning left from multiple directions, is therefore a closer representation of the real-world counterpart, while also facilitating a hasty implementation of the structure (Doniec et al., 2008).

How do simulations apply to the real world?

All of the aforementioned traits of traffic can be analyzed through ABM, but one has to make sure that the conclusions drawn from the analysis can truly be extrapolated to the real world. ABM metatheory is concerned with the way models are relatable to each other and to the real world, and the extent to which this is possible. Any simulation of real-world data needs to tackle the challenge of validation and verification to the real world. Otherwise, the simulation is merely a digital environment with no grounds in the real world and no means to apply conclusions drawn from the simulation outside of it. The assumptions upon which the model is built, the level of abstraction that the model applies to some real-world features, and what means are available to include real-world data into the simulation are all aspects that must be discussed when researching an ABM. Omitting these crucial aspects from a research paper associated with an ABM inevitably results in its conclusions being deemed inapplicable to the real-world counterpart of what the ABM emulates.

ABM verification and validation against the real world is deeply elaborated in Gräbner (2018). This paper brings in epistemological questions such as "If model verification and validation are needed, *what kind of* verification and validation is adequate for the model at hand? ". Among the more straightforward ways to validate a simulation is to incorporate real-world data into it. F. Zhang et al. (2005) applied this strategy and successfully modeled a valid single-lane traffic simulation. The current research was also conducted with incorporated real-world data, namely traffic volume data provided by the Gemeente Groningen (2021).

Another aspect to establish about the ABM is determining which parts of the modelled real-world scenario act simultaneously and in the same way, qualifying them to behave as a homogeneous agent. One could argue that, for e.g., computational reasons, groups of cars with the same destination could also be implemented as one agent. This is called macroscopic modelling, and Haman et al. (2017) discussed the concept of holon, an agent comprised of agents. They argued that traffic could be modelled on multiple levels, with microscopic models sacrificing computational power while macroscopic models sacrifice detail capturing. Macroscopic models are therefore limited in urban areas, where individual cars change lanes at different times and often have different destinations.

It is for these reasons that the current paper only made use of microscopic modeling, by designing each traffic participant as an individual agent. The model was implemented as dynamic, meaning that moving traffic was implemented and analyzed. Static traffic models are merely concerned with representing journeys on study areas (Haman et al., 2017), which does not overlap with what the current research aims to achieve.

What to investigate

Recent traffic optimization research has mostly been accomplished using ABM. One, if not the most, popular topic is controlling traffic lights using Reinforcement Learning. Research such as that of Arel et al. (2010) investigated preventing gridlocks in the centre of intersections, using synced traffic lights to ensure a steady flow of traffic over a long road segment, or adapting the traffic light waiting times based on the demand from each side of the intersection. However, the literature on reassigning the priority road is quite sparse, much like that on the necessity for building additional road structures.

The need to use ABM in traffic research, along with the replicability and reproducibility of it, the multitude of variables to take into account when attempting to optimize traffic, the unpredictability of the effect of some infrastructural improvement, and the amount of detail such research can yield have all been discussed in this section. It is for these reasons that the current research looked into the issue of optimizing traffic on a neighbourhood/area level which contains an agglomerated intersection, and how to make that intersection less crowded. Car flow was modelled over the desired area, and the targeted intersection was adjusted to observe the outcome.

These adjustments came in two forms: mild and strong. A *mild adjustment* consisted of keeping the same road structure, but altering the priority road and observing the effect on traffic flow. A *strong adjustment*, on the other hand, came down to selecting two ends of the intersection and connecting them with an overpass. This allowed cars going from one of these two roads to the other to bypass the traffic inside the intersection. For readability reasons, the terms “overpass” and “bridge” are used interchangeably in this paper.

The reasoning behind these two adjustments stemmed from the idea that if an intersection is busy, then something must be wrong with it and a solution must be found. It might be the case that a mild adjustment would suffice to help clear up the intersection. In that case, a substantial financial advantage is created, because the construction of an overpass or some other complex expensive solution can be avoided. Otherwise, a strong adjustment could fix the traffic problem, and the simulation can help find the best configuration for this (i.e., which roads to link via an overpass).

To further assess the efficacy of the adjustments and their effects on each intersection, a second experiment was performed, focused on overloading each intersection. Having in mind the reasoning that real-world traffic might not be busy enough to observe significant alterations in car flow from the modifications, massively crowding the intersections was targeted.

Given this terminology, the current article investigated the extent to which areas can be optimized via application of these mild or strong adjustments in one intersection. As such, the following research question was formulated: **Can car flow in a neighbourhood be improved using ABM by optimizing one intersection?** The answer was pursued by measuring the car flow in the modeled area with and without the different configurations (both mild and strong adjustments), then checking to see if there is a significant increase. The following hypotheses were formulated:

- H_0: **No attempted configuration (priority changes or bridges) can single-handedly significantly modify car flow;**
- H_1: **Car flow in a neighbourhood can be significantly modified by changing the priority road in an intersection;**
- H_2: **Car flow in a neighbourhood can be significantly modified by building a bridge over an intersection.**

2 Methods

This section discusses each aspect of building the simulation and running experiments using the dataset from the real world.

2.1 Traffic Flow Data

As mentioned in Section 1, the data that was used in the experiments discussed in the following subsections was provided by the Gemeente Groningen (2021). Their Groningen Open Data portal website offers various datasets with information about the city. The only available traffic volume data contained figures for several important intersections in the city of Groningen. The information is structured in rows labeled with each intersection, and columns for every year between 2010-2020. The numbers are expressed in daily car count averaged over each year.

The year 2020 was excluded as outlier from the data, due to the pandemic’s influence on traffic figures.

The yearly evolution of the traffic should not be viewed as parallel measurements of the traffic, but instead as a progression through time. Since the traffic is influenced by real-world events such as new buildings appearing on some streets or population fluctuations, its development must be seen as an irreversible process. As such, it would not be correct to average all the yearly figures. Instead, the counts for each intersection in the year 2019, the most recent year without lockdowns, were used.

The data was not rough, given that it had already been averaged over all days of the year. It only required minimal processing to be used in the experiments. Since each experiment only runs for a limited amount of time and not a full day, the figures in the dataset were first divided to obtain hourly averages. The experiments were run for multiple traffic scenarios, namely high and low traffic. To capture the increase in traffic at rush hour, a distribution on hourly traffic factors (Regehr et al., 2015) was taken. This distribution illustrates what percentage of the total daily count is in traffic during which hours, enabling the accurate division of the available data into hours.

The hourly car counts were assumed to remain constant throughout the hour. Therefore, they were divided by a factor of 30 to obtain the average number of cars in a 2-minute interval with the mean from the dataset. The higher obtained figures were used for rush hour simulations, whereas the lower one from during the rest of the day were used for low traffic simulations.

For each of the two intersections used in the experiments, surrounding areas were selected as spawn locations for the vehicles in the simulation. The count for these adjacent intersections determined the frequency of vehicles coming from that direction, and the vehicles were given paths that crossed through the intersection of interest. The paths for vehicles from each spawn location had multiple destinations, determining the vehicles to follow different lanes and trajectories in the intersection that was investigated.

2.2 Building the Simulation

A framework specialized for simulating traffic was used, CityFlow (H. Zhang et al., 2019). Despite SUMO (Lopez et al., 2018) being state-of-the-art for traffic simulation, CityFlow uses multi-threading and is thus faster. Additionally, it is more recent than SUMO, and would not have had time to establish itself as a viable alternative. It also supports reinforcement learning using Gym, which offers the perfect opportunity to extend the current research in the future.

In order to simulate the intersections of concern, several steps were taken. Their overview is listed below, and each step is detailed in the following subsections:

- converting the maps to the format CityFlow requires;
- cleaning the map files of structures irrelevant for the current research, namely pedestrian paths and crosswalks, bicycle lanes, and traffic lights;
- editing the maps to create the changed priority roads and the bridges that were compared to the current intersections;
- adding vehicles to the simulation in the key points that were discussed in Section 2.1.

2.2.1 Creating the Maps

The experiments described in this paper used two different intersections. The intersections are located where the following roads meet:

1. Eikenlaan - Kastanjelaan;
2. Paterswoldseweg - Laan Corpus den Hoorn;

The general area surrounding the intersections was modeled. This allows a natural flow of vehicles approaching the intersection, instead of abruptly spawning the vehicles right next to the intersection. The two intersections were selected based on two criteria: they tend to become crowded at high-traffic hours, and they are in areas of different road dimensions. To clarify, intersection 1 is in a suburban area with many small streets surrounding it, whereas intersection 2 is in a highly circulated area, with large streets dominating the landscape. This diversity allows the findings to be applicable to multiple scenarios, and diminishes the suspicion that they stem from confounding variables, such as particularities of the investigated locations.

After selecting the areas to be modeled, the OpenStreetMap website was used to download the corresponding maps. They were converted from the `.osm` to the SUMO-native `.net.xml` file format. SUMO's NetEdit and NetConvert tools for automatic and manual editing of map files were subsequently used for cleaning the files. In this process, the pedestrian and bicycle-only paths were removed, as well as any crosswalks and traffic lights. These are all elements that were not used in the experiments, so they could be safely deleted.

Roads were also inspected to ensure that none were corrupted during the download and conversion process. In this case, corruption could mean any distortion to the shape of the roads, any inconsistencies in how lanes diverge or converge, or any unusable intersections due to, e.g., priority conflicts. Once each file had been inspected, it was converted to CityFlow format using the CityFlow-provided converter script (available on the CityFlow Documentation).

2.2.2 Editing the Intersections

For each intersection, three files were used: one with the current, real-world layout of the intersection, one with the same layout but a different priority road (i.e., the mild adjustment mentioned in Section 1), and one with the modified layout including the bridge(s) (i.e., the strong adjustment mentioned in Section 1). The aforementioned NetEdit SUMO tool was used to this extent, which is an application with a graphical interface that allows editing SUMO map files. Screenshots of this interface, featuring all used intersection configurations, can be found in the Appendix, Images A.1, A.2. After the editing process, the maps were converted into CityFlow format. The process to create each structure is described below.

Changing the Priority: NetEdit currently does not support selecting an intersection and changing the priority road directly, but only inspecting priority order. However, a workaround was used to make up for this logistical problem. The road segments leading into the intersection were deleted, and new ones were created in the same location. Creating new segments allows manually choosing which one has priority, so this pipeline was used for creating the maps with a changed priority.

In Intersection 1, the priority road that was researched as an alternative was the South - West axis, i.e., Kastanjelaan - Eikenlaan (Elzenlaan direction). The choice of which of the two halves of Eikenlaan to include in the priority road was made based on traffic volume: in the dataset, the west part is always busier. In Intersection 2, the priority road was simply swapped, by making the north - south axis prioritary instead of the east - west one.

Creating the Bridges: In the two intersections, the simulated bridge was created on predefined conditions: it should diminish waiting time for those vehicles that were previously forced to yield to another road. The bridges were therefore built starting from one of the secondary roads in the intersection. This criterion was applied to the intersections to select the locations of the bridge ends, and this process is described in detail for each intersection in the next paragraphs.

Intersection 2 contains a four-way intersection, so the adopted bridges go forward above the priority road, allowing vehicles to cross the intersection without yielding to perpendicular traffic. Two bridges were built, one on each side (i.e., one for the North - South direction, and one for the opposite).

Intersection 1, however, is a 3-way T-shaped intersection, with priority along the straight road. The least prioritary trajectory in this intersection is to come from the secondary road and turn left, because this scenario requires the vehicle to yield to most vehicles in the intersection. There are no other vehicles that must, in turn, yield to this trajectory. This direction was therefore the first one that was optimized with a bridge.

The second bridge was, once again, connected to the secondary road. This time, it was linked to the right side of it, which leads away from Intersection 1. There are three directions from which vehicles could turn onto this road: either turning left or right from the priority road, or performing a U-turn starting on the other side of the secondary road. Out of these three trajectories, the one that turns left from the priority road was selected for building the second bridge. One reason was the logistical issues that would arise from attempting to build a road over or under this intersection on the aforementioned right or U-turn trajectories, but also the futility of using bridges for right- and U-turns instead of slip lanes. Having said that, the main reason for not selecting this configuration was that it protects the vehicles travelling straight on the priority road from any left-turn interruption. Both possible left turns in this intersection have thus been rerouted to the bridges.

2.2.3 Populating the Maps with Vehicles

CityFlow requires the following information for spawning vehicles:

- A start and end location for their trajectory;
- A frequency at which vehicles are spawned at the start location;
- Particularities of the vehicle.

The particularities of the vehicle include information such as length, width, acceleration, distance maintained from surrounding vehicles, speed etc. These can be used to implement variety into the traffic participants and their behaviours, for a more accurate model of the real world. The scope of the current research does not include making use of these customizable particularities, but they are a good extension point for future research.

Some of these values were changed from default, but experiments were not run for multiple different values. `minGap` was changed from 2.5m to 15m, `maxPosAcc` was changed from $2\frac{m}{s}$ to $1.5\frac{m}{s}$, and `usualPosAcc` was changed from $2\frac{m}{s}$ to $1.0\frac{m}{s}$.

These values were changed because, during preliminary experiments, the vehicles were not reacting to the intersection changes regardless of how intense traffic was. In oher words, the measured traffic flow was identical in all configurations, regardless of its intensity. Vehicles were slowed down through these measures, and the likelihood of traffic jams was increased by raising the distance kept between two vehicles. The effect of this measure was that groups of vehicles were less compact, and thus overall traffic was less efficient.

The data discussed in Section 2.1 was consulted to determine spawn locations for vehicles: the locations where data was gathered were used as the spawn locations around all investigated intersections. Practically speaking, these were the only locations on the map where car counts from the real world were known. Table 2.1 shows the spawn locations used for each map. All maps therefore spanned across the intersection of interest, with the selected spawn locations at or close to the edges of the maps.

Spawn Location	Intersection 1	Intersection 2
North	None	Paterswoldseweg 180
East	Eikenlaan 2	Laan Corpus Den Hoorn 200
South	Kastanjelaan 2	Paterswoldseweg 810
West	Eikenlaan 280	Laan Corpus Den Hoorn 100

Table 2.1: Spawn Locations For Each Intersection

2.3 Measuring Traffic Flow

The information of interest from the experiments was the number of vehicles to pass through the intersection during an experiment run (120 seconds). This was an arbitrarily chosen duration; any duration would have worked, provided it is not absurdly short, because CityFlow is a deterministic simulation and yields the same results over any number of simulation runs. This was used to quantify the effects of each attempted intersection configuration. An intersection is more efficient when more cars are able to pass through it in a certain period of time. In a real-life scenario, this would be measured by counting the number of unique cars to pass through the intersection. However, the simulation operationalization is limited by CityFlow's Data Access API.

The only way to measure the aforementioned number is through a function that returns a dictionary with all vehicles in the simulation, grouped by lanes. This dictionary represents the state of the simulation in the current time step, which means that the function needs to be called at each iteration. In order to count the number of unique vehicles passing through the intersection throughout an entire experiment run, the aforementioned function was used to monitor all entrances and exits to the intersection. The closest road segments to each intersection were selected. These are referred to as "detection zones" for the rest of this paper.

Each detection zone can have one or multiple lanes. Due to not allowing duplicates, Python sets were used to count unique vehicles in the intersection. One set was used for each detection zone, to keep track of which vehicles passed through that segment. At the end of a simulation run, the sizes of all these sets were collected. This meant that the results communicated the total number of cars to enter the intersection from each direction, as well as the total number of cars to exit the intersection in each direction.

2.4 Running the experiment

The particularities of running and collecting results for each experiment are discussed in the following paragraphs. Key information required for their replicability is also covered.

2.4.1 Experiment 1

The first experiment was the main focus of the project. It relied on the real-world data and looked into how the investigated configurations would impact traffic in real life. For each experiment run, multiple particularities were first selected: which of the intersections to simulate, which configuration to test for the selected intersection, and the time of day (i.e., the heaviness of traffic). The former two determined which of the available files should be used for the map in the simulation, and the latter determined the frequency of spawned cars coming from each direction.

Once all these independent variables had been selected, the simulation commenced. The vehicles were first given enough time to reach and populate the measured intersection. In real-life scenarios, there are always cars inside the intersection or approaching it, whereas cars in the simulations are initially only present at the spawn locations and must first make their way to the intersection. Given this contrast between real life and simulation, the aforementioned initial travel time was required. Regardless of the intersection, vehicles were given 120 seconds to travel from their spawn locations to the intersection of interest.

After these 120 seconds, it was assumed (and also confirmed via the graphical user interface provided by CityFlow) that the vehicles had indeed populated the intersection and the measurements could begin. Employing the pipeline described in Section 2.3, the simulation ran for 120 more seconds during which the vehicles passing through the simulated intersection were counted. The obtained counts for each direction were then collected.

2.4.2 Experiment 2

To attain the desired "intersection overloading" effect required in Experiment 2, the number of cars in the intersection was increased. This was achieved by lowering the spawn interval, which is the frequency of spawning cars at a given location. The simulation has shown convergent results when the spawn interval is decreased. This means spawning more cars than a certain value would have no additional effect upon the results. In other words, this was the maximum number of cars that could possibly cross the intersection. The least frequent spawn interval where the results converge was considered the "convergence threshold". This convergence threshold was found for each intersection by experimenting with various values. The same spawn interval was used at all spawn locations, in contrast to how Experiment 1 spawned vehicles.

When the results would converge in the standard intersection configuration, the discovered convergence threshold was stored and subsequently used on all the configurations for that intersection. As described in the previous paragraph, the experiment was run with this convergence threshold as the spawn interval to see whether the configurations would modify the results, and in what way. The outcomes were also cross-checked against Experiment 1, to separate the effects of the real-world data from the specifics of the intersections and configurations.

The convergence threshold used for Intersection 1 was 2 seconds, and that for Intersection 2 was 2.5 seconds. Another change from Experiment 1 was that the waiting time before starting to count the vehicles was increased to 220 seconds. This measure was taken to ensure the intersections were as crowded as they can be when the vehicle count started. The previous 120 seconds were enough to populate the intersection, but user-interface-supported inspection revealed, in the overloading scenario, cars would continue to pile up past this simulation step, instead of exhibiting the desired non-evolving behaviour. The results of this second experiment can be observed in Tables 3.3 and 3.4, and can also be visualized in the Appendix (Figures A.5 and A.6).

3 Results

Once the simulation had finalized, the dependent variable was made available for inspection. This was in the form of car counts going into and out of each intersection's roads, which were labelled using cardinal points (North, East, South, West). The total number of vehicles entering and exiting the intersection was also computed, by summing the counts for all cardinal points. The obtained numbers can be observed in Tables 3.1 and 3.2. It is worth mentioning that, despite the standard deviations appearing in the plots for both experiments, they do not further improve data interpretation. The bar plots already provide the necessary visual cues to understand the volume difference between different traffic flows, so the standard deviations only reiterate this information.

3.1 Experiment 1

As can be observed in Image A.3, **Intersection 1** does not draw any noticeable benefits from the alternative priority: despite morning rush traffic allowing one extra car to pass through, during the clear hour it is on par with the standard configuration, and the afternoon rush hour sees one less vehicle passing through the intersection. The bridge appears to improve both morning and afternoon rush hours, increasing car flow by 8 and 17 vehicles respectively. Clear hour, however, is encumbered by the bridge and decreases overall traffic flow by 3 vehicles.

In **Intersection 2**, however, the results do not feature the same stagnations or slight improvements. Image A.4 shows that, in the clear hour, the alternative priority is roughly equivalent to the standard configuration. Two less vehicles can enter the intersection, but those that manage to do so pass faster, with one extra vehicle crossing the intersection during the experimental time frame. Nevertheless, both rush hours show traffic volume depreciation in the investigated configurations, compared to the standard one. The alternative priority decreases the total car count by 5 vehicles in the morning and 9 in the afternoon, this decrease being evenly distributed between the entering and exiting vehicles. The bridge configuration continues and accentuates this trend, by exhibiting a total of 26 less cars in the morning and 35 less cars in the afternoon to pass through the intersection. These effects are, once again, visible both in entering and exiting vehicle counts.

3.2 Experiment 2

The results of Experiment 2 are available in Tables 3.3 and 3.4. As Figures A.5 and A.6 show, the trends from the first experiment are not changed. There is still a slight increase with the changed priority in Intersection 1, but the slight decrease in Intersection 2 is less visible than in Experiment 1. The bridge is proved, yet again, to be a valuable improvement for Intersection 1: 84 more vehicles are able to pass through the intersection on account of the bridge. This is a massive difference for Intersection 1, given that some cases from Experiment 1 see less than 84 cars pass through in total.

Intersection 2 also shows results consistent with Experiment 1: the bridge is still a major flaw in this case. In total, 33 less cars are able to pass through the overloaded intersection compared to the standard configuration.

4 Discussion

This section covers a deeper analysis of the results and a critical interpretation. A comprehensive overview of ways to improve the current research is also included.

4.1 Intersection-Specific Analysis

Before discussing the results, an overview of Braess's paradox is due, being a recurring topic in the next paragraphs. Braess (1968) discusses how adding additional roads in a road network could in fact decrease the overall traffic flow for the road network. Other contexts, such as biology or power grids, also sometimes feature this paradox. In colloquial terms, the explanation is that too many drivers opt for the quicker route, and end up decreasing the overall travel time for the traffic participants. In more formal terms, the Nash Equilibrium for each individual driver does not match with the optimum configuration for the intersection, which means that the individual best outcome does not arise from the same choices as those required to reach the collective best outcome.

Table 3.1: Eikenlaan: Number of Vehicles Entering and Exiting the Intersection, by Traffic and Road

Traffic Hour	Clear Hour			Morning Rush			Afternoon Rush		
Adjustment	Standard	Mild	Strong	Standard	Mild	Strong	Standard	Mild	Strong
East Entrance	14	14	11	22	22	17	27	27	21
East Exit	12	11	11	15	15	19	16	15	22
South Entrance	5	6	6	9	9	9	12	12	12
South Exit	13	13	13	18	19	21	20	20	26
West Entrance	17	17	5	20	20	26	20	20	31
West Exit	9	9	18	12	13	12	15	15	15
Total In	36	37	25	51	51	52	59	59	64
Total Out	34	33	42	45	47	52	51	50	63

Table 3.2: Paterswoldseweg: Number of Vehicles Entering and Exiting the Intersection, by Traffic and Road

Traffic	Clear Hour			Morning Rush			Afternoon Rush		
Adjustment	Standard	Mild	Strong	Standard	Mild	Strong	Standard	Mild	Strong
North Entrance	22	22	16	31	31	20	31	31	20
North Exit	28	29	26	37	37	35	44	44	39
East Entrance	28	27	25	42	40	45	50	42	45
East Exit	24	24	21	27	27	21	31	31	27
South Entrance	17	16	12	27	27	19	30	30	22
South Exit	22	22	20	30	29	28	29	30	28
West Entrance	21	21	21	21	21	21	21	21	21
West Exit	22	22	21	30	28	30	33	31	32
Total In	88	86	74	121	119	105	132	124	108
Total Out	96	97	88	124	121	114	137	136	126

Table 3.3: Eikenlaan: Number of Vehicles Entering and Exiting the Intersection in Experiment 2, by Road

Adjustment	Standard				Mild				Strong			
Road	East	South	West	TOTAL	East	South	West	TOTAL	East	South	West	TOTAL
Entering	28	18	18	64	33	24	20	77	31	30	40	101
Exiting	17	19	19	55	25	20	22	67	37	35	30	102

Table 3.4: Paterswoldseweg: Number of Vehicles Entering and Exiting the Intersection in Experiment 2, by Road

Adjustment	Standard					Mild					Strong				
Road	North	East	South	West	Total	North	East	South	West	Total	North	East	South	West	Total
Entering	31	49	36	20	136	32	45	36	20	133	20	53	24	20	117
Exiting	46	32	30	32	140	46	33	30	32	141	34	33	28	31	126

4.2 Experiment 1

A closer look at the **Intersection 1** in and out counts for the clear hour reveals that only the number of vehicles entering the intersection is decreased, and there are in fact more vehicles exiting the intersection in the bridge configuration than in the other two. A form of Braess's paradox is therefore present, with a possible explanation being that too many vehicles rush to cross the intersection via the bridge and prevent additional vehicles from entering the intersection in due time.

Section 3 discussed a decrease in the number of vehicles in the bridge scenario for **Intersection 2**. Although these premises would recommend Intersection 2 results for Braess's paradox, much like Intersection 1, this time the explanation is much simpler. This result can be attributed to the situation vehicles have to deal with after crossing the bridge, when they must rejoin regular traffic. They now yield to vehicles coming from the ground-level lanes, which are the rest of the vehicles that have turned into the same road where the bridge ends.

Before building the bridge, vehicles on the path optimized by the bridge had to first yield to vehicles coming from their left, then to those coming from their right, and could then cross the full intersection and continue driving. After building the bridge, these vehicles now pass over the entire intersection, but then have to yield to the summation of vehicles to turn onto this road from all directions, and at the same time. They must now wait for everyone that has turned into this new road, without breaking down this task into multiple steps.

4.3 Experiment 2

To take a deeper dive into the results of the second experiment, it is interesting to take a closer look at road-specific numbers. Intersection 1 shows a relatively similar proportion of vehicles travelling on each road in the standard and changed priority case. However, the bridge manages to increase the vehicle counts across the South and West axes, raising them to be on par with the East axis. The latter is also increased compared to the other two scenarios, but its tally does not stand above the others anymore with the bridges. Given that both bridges were connected to the South road, this result does not come as a surprise. Furthermore, the West road is still the priority road, but its vehicles do not have to coordinate as much with the yielding vehicles when the latter would join the main road.

Intersection 2 also shows some interesting developments over the individual roads. Despite the standard and changed priority configurations not showing much of a change, this is not the case for the bridge. This sees an increase over the East road, but a larger decrease over the North and South roads. As with Intersection 1, the priority road (East-West) visibly benefits from not facing as many vehicles inside the intersection anymore.

However, the bridges seem to pose a bottleneck for the North and South roads. As stated in Section 4.2, one cause could be that they must yield when exiting the bridge. On top of that, perhaps the bridges should have featured multiple lanes instead of a single lane per direction. This could have allowed more vehicles to cross over the intersection, and Experiment 2 evidently contains enough vehicles to fill multiple lanes on the bridges.

4.4 Intersections Cross-Comparison

Everything put together, the alternative priority does not show significant improvements in any intersection regardless of the time of day, but the bridges sometimes improve and sometimes impair traffic flow. The contradictory nature of these results can mean one of two things. The experiment has either led to inconclusive findings, or the confounding variables specific to each intersection are so compelling as to cause opposite results between the investigated intersections.

These confounding variables are hardly insignificant. With diversity in mind, the intersections selected for this research vary in a number of ways. Intersection 1 only has one lane per direction of travel, whereas Intersection 2 has roads with as many as 3 lanes on one side of the road. Traffic volume is also a major contrast point: Intersection 1 has 10421 cars coming from its busiest road over the course of one day, whereas Intersection 2 has 19798, almost twice as many. Even the number of roads entering the intersection is different, with Intersection 1 linking three roads while Intersection 2 connects four.

A deeper dive into the configurations investigated in the discussed experiments allows for a more specific explanation of the results. While the bridge itself provides a hasty alternative to waiting for an opening to turn left in Intersection 1, it appears to be more a hindrance than a help in the case of Intersection 2. The decreased number of vehicles to enter the intersection in both bridge scenarios (for all hours in Intersection 2 and clear hour in Intersection 1) suggests an increased waiting time before the intersection. The cause of this appears to be the bridge across both intersections, thus confirming the aforementioned theory that too many vehicles choose the bridge route and cannot then hastily merge into traffic upon crossing the bridge. This creates a queue of vehicles on the bridge that extends to the other side of the intersection and prevents additional vehicles from entering it.

These affirmations are *confirmed by the results of the second experiment*, which proves that the findings remain valid even during drastically crowded scenarios.

4.5 Future research

The current research was performed with several assumptions in mind. These assumptions allowed for a straight-forward implementation of the traffic model, but steps can be taken to improve the accuracy and real-world applicability of the findings.

The incremental improvement that can be brought to these experiments is solving the conflict between the results of the two intersections. This contrast appears across both experiments. Simulating a third intersection is a straight-forward way to accomplish this, because then a majority will be created, either proving or disproving overpasses' ability to optimize traffic. The current project was intended to include three intersections. However, due to technical issues related to NetConvert, this was eventually dropped.

The designated third intersection was where the streets Europaweg and Sontweg meet. This intersection also had the advantage of possibly validating the results, because the dataset also contained the number of vehicles exiting this intersection towards its southern road. As discussed in Section 1, validation is a very important aspect of ABM, and that constitutes a big advantage of the aforementioned intersection.

As for the assumptions discussed in the beginning of this section, the first simplification of this project that can be tackled is the traffic data used in the simulations. This data was collected over long periods of time, more than 10 years being available in the dataset. External short-term data, or manually collected data at the desired intersections over short intervals, would have been more suitable and could be collected to extend this project in the future. On top of that, the traffic distribution that was used to approximate hourly traffic was taken from a paper based on United States data, whereas the experiments run in this project took place in Europe. Collecting short-interval data for future research would remove the need for the hourly traffic distribution, but research specific to Europe would have been more adequate for this project's purposes regardless of that.

Despite these complaints brought to the broad dataset, they do not mean that the dataset itself should not be used. In fact, the yearly progression that can be observed in the dataset can come in handy for future research. Combined with new infrastructure, recent constructions, and similar information, this data can be used to predict how traffic volume will evolve in future years. Even without these external factors, the yearly counts can still be used to project next years' car counts, and further investigate how the changes to the investigated intersections would affect the future. After all, the assessed infrastructure would be built for the future, so the best data to test it on would be car counts in the years following its construction.

As for the mild and strong adjustments featured in the experiments, their limitations are no secret. The alternative priority could have been tested between multiple different roads, and the same goes for the bridges. More maps can be created in future research, containing bridges between any two or more roads. This can also be extended by including other types of intersections, such as roundabouts or traffic lights. A more complete project would ideally run through all possible configurations, with all types of intersections.

Furthermore, a more futuristic project could see the implementation of other types of traffic managing. For example, Dresner & Stone (2008) implemented a mechanism that coordinates traffic over an intersection with continuously changing rules. They state autonomous, AI-guided vehicles could be taken advantage of by coordinating them in a more efficient manner than current intersection types allow. Their mechanism is able to mimic current intersection types, but also other intersection behaviours that rely on a "detailed communication protocol".

Another assumption that laid the foundation of this project is that the number of cars coming from each direction does not change after the intersection's configuration changes. This is, evidently, not always the case. One could easily see that changing the main road in an intersection would cause drivers to opt for different paths on their route. In the case of building a bridge, this is all the more relevant, as vehicle numbers would definitely change upon such a significant infrastructure project's finalization.

The second experiment could also be extended: the convergence threshold discussed in this paper must surely contain key information about the intersection. After all, a lower threshold (i.e., more cars are required to converge to a maximum) entails an intersection can hold more vehicles, which is ultimately the goal of optimizing intersections. Future research could replicate this experiment and take a closer look at this information.

On the topic of trajectory choice, another weakness of this project comes to light: in the bridge scenarios, an additional way of crossing the intersection in a specific trajectory is effectively added. Vehicles may still opt to cross the intersection on ground level, despite following the bridge's direction. This would occur at least due to drivers changing their minds regarding which path to take after passing the entrance to the bridge. This happens everyday, with people often missing their highway exits or in other similar situations. A more realistic behaviour should have been modeled by making the agents sometimes opt to cross through the intersection despite taking the bridge's trajectory, and this point of improvement could be addressed in future research.

Finally, the robustness of the bridges modeled in this project should be criticized. The situations between the different maps are not directly comparable, because one map has a bridge between a primary and secondary road, and the other one has a bridge between two secondary roads. Moreover, one map sees bridges optimizing left-hand turns, whereas the other targets the straight trajectory in both directions of the yielding road.

This matter, however, was done purposefully. The aim of the project was to cover multiple scenarios of traffic optimization, and analyze the situations in each of them. The aforementioned diversity in map choice was therefore also applied to the simulated bridges, in order to separate the effects of the investigated changes as much as possible from the other particularities that were not of concern.

5 Conclusions

Given the inconclusive nature of the mild adjustments in all investigated intersections, H_1 *is rejected*. However, the bridge has shown minor traffic improvements in a few isolated cases, and major traffic deteriorations in the other cases. As such, H_2 *is accepted and* H_0 *rejected*. A notable aspect is that, while H_2 is accepted, the modifications it refers to consist of traffic deterioration, not improvements. The results have shown car flow can be significantly decreased by building a bridge.

Future iterations of these experiments should include different types of intersections, as well as more precise real-world data. Moreover, a consensus on the effects of the investigated changes, regardless of which intersection they are applied to, must be sought out. This can be attained by experimenting on a larger number of intersections, effectively isolating specifics of each intersection from the obtained results.

References

Arel, I., Liu, C., Urbanik, T., & Kohls, A. G. (2010). Reinforcement learning-based multi-agent system for network traffic signal control. *IET Intelligent Transport Systems*, *4*(2), 128–135.

Bazghandi, A. (2012). Techniques, advantages and problems of agent based modeling for traffic simulation. *International Journal of Computer Science Issues (IJCSI)*, *9*(1), 115.

Braess, D. (1968). Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, *12*(1), 258–268.

Doniec, A., Mandiau, R., Piechowiak, S., & Espié, S. (2008). A behavioral multi-agent model for road traffic simulation. *Engineering Applications of Artificial Intelligence*, *21*(8), 1443–1454.

Dresner, K., & Stone, P. (2008). A multiagent approach to autonomous intersection management. *Journal of artificial intelligence research*, *31*, 591–656.

Gemeente Groningen. (2021). *Counter figures car traffic.* Retrieved from `https://groningen.dataplatform.nl/#/data/c1b9271b-e3c7-4254-b9dd-691178bf20cb`

Gräbner, C. (2018). How to relate models to reality? An epistemological framework for the validation and verification of computational models. *Journal of Artificial Societies and Social Simulation*, *21*(3).

Haman, I. T., Kamla, V. C., Galland, S., & Kamgang, J. C. (2017). Towards a multilevel agent-based model for traffic simulation. *Procedia Computer Science*, *109*, 887–892.

Lopez, P. A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., … Wießner, E. (2018). Microscopic traffic simulation using SUMO. In *2018 21st International Conference on Intelligent Transportation Systems (itsc)* (pp. 2575–2582).

Regehr, J. D., Montufar, J., & Hernández-Vega, H. (2015). Traffic pattern groups based on hourly traffic variations in urban areas. *Journal of Transportation of the Institute of Transportation Engineers*, *7*(1), 1–16.

Zhang, F., Li, J., & Zhao, Q. (2005). Single-lane traffic simulation with multi-agent system. In *Proceedings. 2005 IEEE Intelligent Transportation Systems, 2005.* (pp. 56–60).

Zhang, H., Feng, S., Liu, C., Ding, Y., Zhu, Y., Zhou, Z., … Li, Z. (2019). Cityflow: A multi-agent reinforcement learning environment for large scale city traffic scenario. In *The World Wide Web Conference* (pp. 3620–3624).

A Appendix

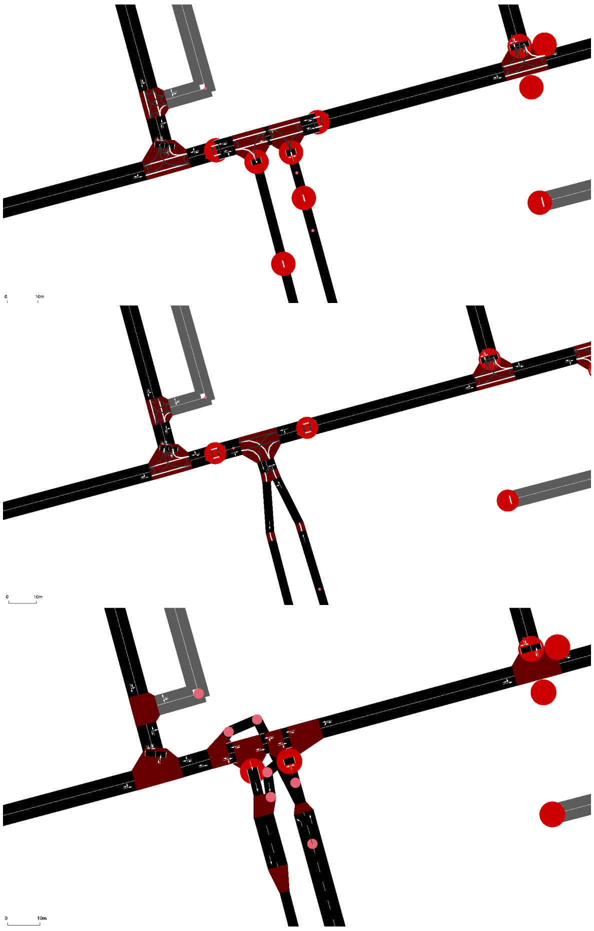

Figure A.1: Eikenlaan Intersection Configurations. From top: standard, alternative priority, bridge.

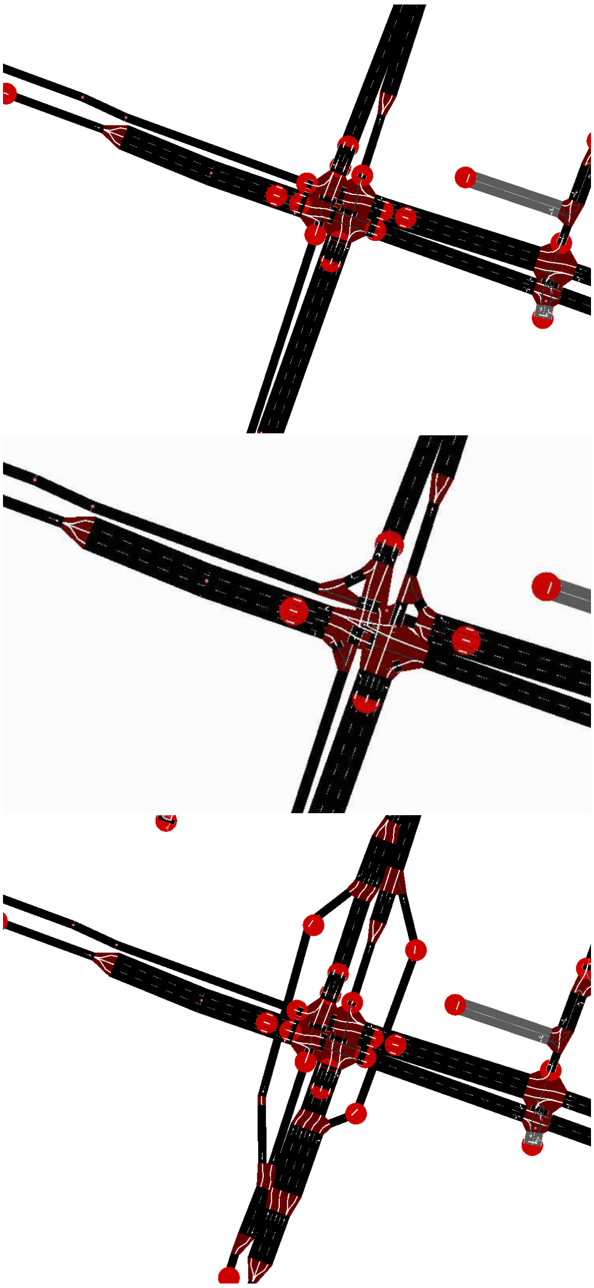

Figure A.2: Paterswoldseweg Intersection Configurations. From top: standard, alternative priority, bridge.

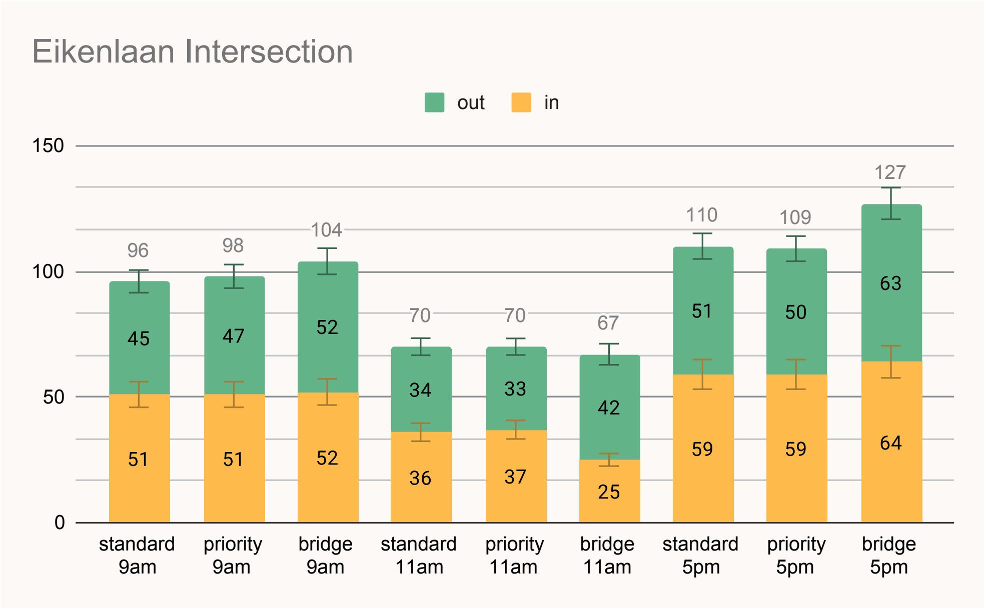

Figure A.3: Total vehicles entering and exiting Intersection 1

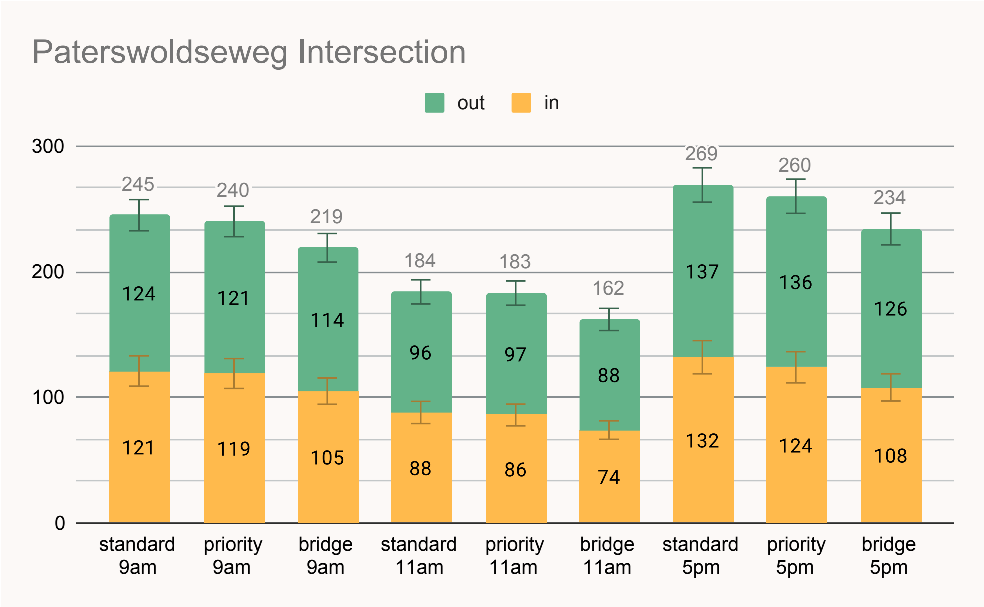

Figure A.4: Total vehicles entering and exiting Intersection 2

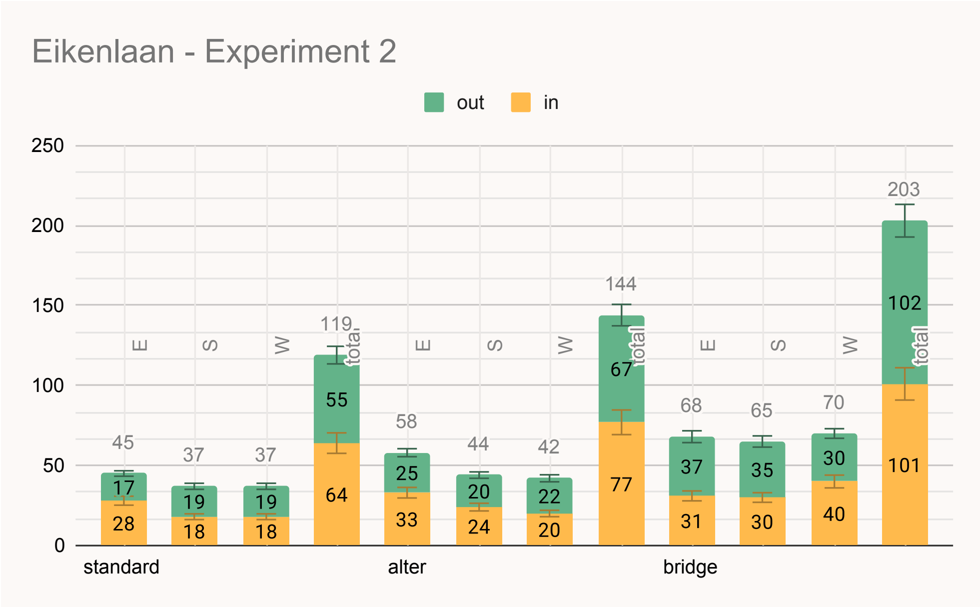

Figure A.5: Number of vehicles entering and exiting Intersection 1 in Experiment 2, by Road and Total

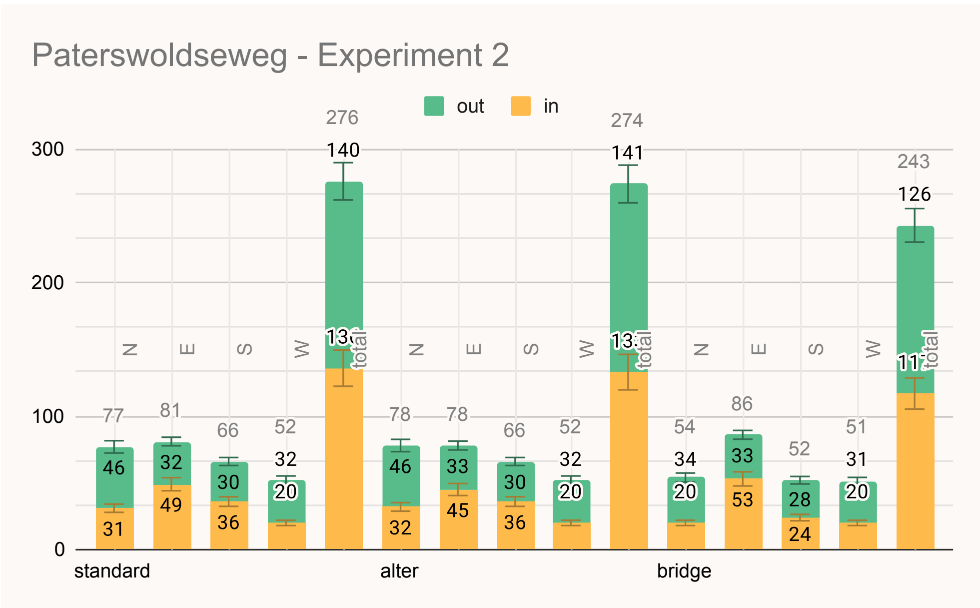

Figure A.6: Number of vehicles entering and exiting Intersection 2 in Experiment 2, by Road and Total