

# Measuring & Mitigating Electric Vehicle Adoption Barriers

by

Tommy Carpenter

A thesis  
presented to the University of Waterloo  
in fulfillment of the  
thesis requirement for the degree of  
Doctor of Philosophy  
in  
Computer Science

Waterloo, Ontario, Canada, 2015

© Tommy Carpenter 2015

### **Author's Declaration**

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

## Abstract

Transitioning our cars to run on renewable sources of energy is crucial to addressing concerns over energy security and climate change. *Electric vehicles* (EVs), vehicles that are fully or partially powered by batteries charged from the electrical grid, allow for such a transition. Specifically, if hydro, solar, and wind generation continues to be integrated into the global power system, we can power an EV-based transportation network cleanly and sustainably.

To this end, major car manufacturers are now producing and marketing EVs. Unfortunately, at the time of this writing, drivers are slow to adopt EVs due to a number of concerns. The two greatest concerns are *range anxiety*—the fear of being stranded without power and the fear that necessary charging infrastructure does not exist—and the unknown return on investment of EVs over their lifetime.

This thesis presents computational approaches for *measuring* and *mitigating* EV adoption barriers. Towards measuring the barriers to adoption, we build a sentiment analysis system for programmatically mining detailed perceptions towards EVs from ownership forums. In addition, we design the most comprehensive *electric bike* trial to date, which allows us to study several aspects of electric vehicles, including range anxiety, at a much lower cost. Towards mitigation, we develop algorithms for managing a network of gasoline vehicles to be used by EV owners when a planned trip exceeds the range of their EV. Further, we design a model for taxi companies to compute whether it is profitable to transition a fraction of their fleet to EVs.

To summarize our findings, we find that sentiments towards EVs are very positive, especially regarding performance and maintenance, but there are concerns over range anxiety and the higher initial price of EVs. There is a delicate balance between these two adoption barriers. Larger batteries cost more, so alleviating range anxiety with larger batteries leads to pricier vehicles. Conversely, EVs with low range capabilities can also induce costs, because drivers and fleets that own EVs may have to often acquire (or own as an additional vehicle) a gasoline vehicle to fully meet their mobility demands. As a result, EVs are best suited for drivers and fleets that are able to make long-term return on investment calculations, and whose mobility patterns do not include many very long trips. Fleets can greatly reduce their operating costs by adopting EVs because they have the capital to make upfront investments that are profitable long-term. We show that even under conservative assumptions about revenue loss due to battery depletion, EVs are already profitable (the company saves more than enough money to recoup all initial investments) for a large taxi company in San Francisco. Similarly, EVs can be profitable for two-car families (those who already have a gasoline car) and for those who can easily acquire a gasoline vehicle when needed, hence our work on sizing networks of gasoline-vehicle pools for EV owners. Finally, we find that not only are electric bikes and EVs operationally similar, the sentiments towards the two technologies are as well. Advancements made in the battery sector, especially those that reduce costs or weight, are likely to accelerate sales in both markets.

The results presented in this thesis, as well as in prior work, suggest that EVs are suitable for many drivers and will hence serve a role in our eventual transition away from fossil fuels.

# Acknowledgements

First and foremost, I would like to thank Prof S. Keshav, my advisor, for this dissertation would not exist without his guidance. I started working with Prof Keshav in August 2010, and in just shy of five years, he has shaped the way I think about research, engineering, the world, and my place in it. Many methods in this thesis are a result of his vast knowledge of systems and networking, and his endless supply of red pens. Many positive changes/events in my life are a result of his teachings. I am very fortunate to have been supervised by him.

I thank my PhD committee: Prof Hussein Mouftah, Prof Lukasz Golab, Prof Johnny Wong, and Prof Yuying Li for all of their help and suggestions on this thesis. It is greatly improved due to their guidance.

My wife, Lisa Patel Carpenter, I cannot imagine life—past, present, or future—without. She has provided infinite love and support throughout my life and journey. Additionally, she has meticulously read my work and provided countless clarifications and fixes.

I thank my parents for their support and motivation. Each has helped me greatly in a variety of ways.

Next, I thank the co-authors and collaborators on all projects presented in this thesis. In order by chapter, they are Prof Lukasz Golab and Sohail Javaad Syed (Chapter 4), Prof Wong (Chapters 5 and 6), Pavlo Kystmey (Chapter 6), Dr. Andrew Curtis (Chapter 7), Costin Ograda-Bratu, Rayman Preet, Milad Khaki, and Tobias Schroeder (Chapter 8).

Chapters 5 and 6 would not have been possible without data from Community Carshare. Similarly, much of my knowledge of queueing theory, which is essential to these chapters, is inherited from Prof Wong.

I thank Dr. Andrew Curtis and Dr. Earl Oliver, the two senior PhD students under Prof Keshav's supervision at the time of my arrival, for their help early in my graduate career.

Jason Baek, Prof Augustin Chaintreau, Prof Catherine Rosenberg, Parsiad Azimzadeh, and Anthony Schmieder all made contributions to the work presented in this thesis.

To all members of the Tetherless Computing/ISS4E research lab present from August 2010—January 2015: thank you for making DC3562 a second home.

Finally, I would not have applied to Waterloo if it were not for Prof Frédéric Latour, my undergraduate mathematics professor, a Waterloo alumni. I also thank all professors in CCSU's Computer Science department for pushing me as an undergraduate student.



# Table of Contents

<b>List of Tables</b>	<b>ix</b>
<b>List of Figures</b>	<b>xi</b>
<b>List of Symbols</b>	<b>xv</b>
<b>List of Acronyms</b>	<b>xix</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Technology Overview</b>	<b>4</b>
2.1 Electric Vehicles . . . . .	4
2.1.1 Battery Electric Vehicles . . . . .	4
2.1.2 Plug In Hybrid Electric Vehicles . . . . .	5
2.1.3 Hybrid Electric Vehicles . . . . .	5
2.1.4 Price Comparisons . . . . .	5
2.2 Electric Bikes . . . . .	6
2.3 Batteries . . . . .	7
2.4 Charging ("Refeuling") . . . . .	8
2.5 Glossary . . . . .	8
<b>3 Related Work</b>	<b>10</b>
3.1 Gauging Driver Perceptions . . . . .	10
3.1.1 Field Trials . . . . .	11
3.1.2 Surveys . . . . .	14
3.2 Range Anxiety . . . . .	19
3.3 Calculating the Return On Investment . . . . .	21
3.4 Using GPS to Infer Mobility Demands . . . . .	26
<b>4 Mining Electric Vehicle Opinions</b>	<b>29</b>
4.1 Synopsis . . . . .	29
4.2 Terminology . . . . .	30
4.3 System Architecture . . . . .	30
4.4 Phase 1: Data Collection and Preprocessing . . . . .	32
4.5 Phase 2: Building The Feature Space . . . . .	32

4.6	Phase 3: Chunking	32
4.6.1	Part of Speech Tagging	33
4.6.2	NLTK Chunking	33
4.6.3	Advantages Of Chunking	35
4.7	Phase 4: Sentiment Querying	36
4.7.1	Handling Context-Dependent Opinions	36
4.7.2	Sentiment Querying Pipeline	37
4.8	Phase 5: Results	38
4.8.1	Evaluation Methodology	38
4.8.2	Distribution of Sentiments	40
4.8.3	Performance	42
4.9	Prior Feature-Based Opinion Mining Methodologies	44
4.9.1	Feature-Based Opinion Mining	45
4.9.2	Word and Sentence Polarity Classification	46
4.10	Conclusions & Future Work	47
<b>5</b>	<b>A Vehicle Pool Model For Reducing Range Anxiety</b>	<b>49</b>
5.1	Synopsis	49
5.2	Sizing A Pool Subject To A QoS Target	50
5.2.1	Binomial Based Sizing	51
5.2.2	Sizing Via The Engset Loss Model	52
5.2.3	Dataset-based sizing	58
5.2.4	Optimizing The Pool Size	58
5.3	Evaluation Methodology	59
5.3.1	Dataset	60
5.3.2	Performance Metrics	60
5.3.3	Bootstrapping	61
5.3.4	Parameter Values	62
5.3.5	Performance Metric Computation	62
5.4	Evaluation Results	63
5.4.1	Arrivals are Non-Markovian and Non-Stationary	63
5.4.2	Performance	65
5.5	Conclusions & Future Work	69
<b>6</b>	<b>A Carsharing Model For Reducing Range Anxiety</b>	<b>70</b>
6.1	Synopsis	70
6.2	Problem Formulation	71
6.2.1	System Architecture	71
6.2.2	Assumptions	72
6.2.3	Methodology Overview	72
6.3	Pool Sizing Methods	73
6.3.1	Engset Loss Model	73
6.3.2	Additive Decrease Multiplicative Increase (ADMI)	75
6.3.3	Offline Optimal Baseline	76

6.4	Fleet Management	76
6.5	Evaluation Methodology	78
6.5.1	Bootstrapping and Confidence Intervals	78
6.5.2	Performance Metrics	78
6.5.3	Performance Metric Computation	79
6.5.4	Parameter Prediction	79
6.6	Results	80
6.6.1	Single Pool Dynamic Sizing	81
6.6.2	Mutli-Pool Dynamic Sizing	83
6.7	Modeling Removals	86
6.8	Prior Fleet Sizing and Management Models	89
6.9	Conclusions And Future Work	91
6.9.1	Supporting One Way Trips	91
6.9.2	Subscribers Borrowing From Multiple Locations	92
6.9.3	Reservations	92
6.9.4	Budget Limitations	93
6.9.5	Multi Period Future Optimization	93
6.9.6	Redirecting Blocked Arrivals	93
<b>7</b>	<b>A Profitability Model for Fleet Transition</b>	<b>94</b>
7.1	Synopsis	94
7.2	Taxi Operation Overview	95
7.3	Data-Oriented Process to Estimating ROI	97
7.3.1	Model Assumptions	97
7.3.2	Inputs	98
7.3.3	Model Outputs	99
7.3.4	Estimating Charge Levels	99
7.3.5	Using The Model To Infer Costs	102
7.4	Switching Station Optimization Algorithm	105
7.5	San Francisco Case Study	106
7.5.1	Dataset and Preprocessing	107
7.5.2	Clustering Locations	108
7.5.3	Assumptions For BEV/PHEV Revenue Analysis	108
7.5.4	Case Study EV Scenarios	109
7.5.5	Existing Taxi Revenue	110
7.5.6	Revenue Analysis for BEVs	110
7.5.7	Revenue Analysis for PHEVs	114
7.5.8	PHEV vs. BEV Comparison	115
7.5.9	ICEVs With Increased Efficiency, HEVs	115
7.5.10	Sensitivity Analysis	116
7.6	Conclusions and Future Work	118
<b>8</b>	<b>Designing A Trial To Explore the Barriers to Electric Bike Adoption</b>	<b>120</b>
8.1	Motivating Research Questions	121

8.1.1	Coordinating Usage Data With Preconceived Barriers . . . . .	122
8.1.2	Questions Relevant to EV Researchers/Manufacturers . . . . .	123
8.1.3	HCI Questions . . . . .	124
8.2	Field Trial Design . . . . .	124
8.2.1	The eBikes . . . . .	124
8.2.2	The Sensor Kit . . . . .	125
8.2.3	Data Transmission and Analysis . . . . .	125
8.2.4	Participant Selection . . . . .	127
8.2.5	Incentivizing Participation & Care Of Bikes . . . . .	128
8.3	Mapping Battery Voltage to SOC . . . . .	128
8.3.1	Estimating SOC Based On Voltage . . . . .	129
8.3.2	Voltage Filtering . . . . .	130
8.3.3	Biking Interpolation . . . . .	132
8.4	Trip Detection . . . . .	133
8.5	Survey Analysis . . . . .	135
8.5.1	Common Phrases In Survey . . . . .	135
8.5.2	Sentiments In Survey . . . . .	137
8.6	WeBike Data . . . . .	139
8.7	Prior Bicycle Telematic Systems . . . . .	144
8.8	Conclusions . . . . .	145
<b>9</b>	<b>Conclusions</b>	<b>147</b>
9.1	Summary Of Contributions . . . . .	147
9.2	Directions for Future Work . . . . .	150
9.3	Concluding Remarks . . . . .	152
	<b>References</b>	<b>153</b>
	<b>Appendix A Sentiment Analysis Chunking Grammar</b>	<b>173</b>
	<b>Appendix B Metric Over Time Figures For Chapter 6</b>	<b>174</b>
	<b>Appendix C Proof Of Convergence For Algorithm 1</b>	<b>178</b>
C.1	Blocking Probability As A Fixed Point . . . . .	178
C.2	Existence and Uniqueness of a Fixed Point . . . . .	179
C.3	Fixed Point Iteration . . . . .	179
	<b>Appendix D WeBike Field Trial Participation Survey</b>	<b>181</b>
	<b>Appendix E WeBike To EV Mapping</b>	<b>193</b>

# List of Tables

4.1	Examples of solvable sentiment analysis errors. "FSD" denotes "feature specific dictionary" . . . . .	44
4.2	Examples of sentiment analysis errors that are "better left unsolved". . . . .	44
4.3	Examples of ambiguous sentiment analysis classifications. . . . .	44
6.1	Comparison of two ELM and two ADMI sizing configurations (corresponds to Figure 6.4a). $P$ denotes the pool number. ADMI does not size for a QoS target. . . . .	82
6.2	Sensitivity of ELM to the QoS parameter (see Figure 6.4b). $P$ denotes the pool number. . . . .	83
6.3	Comparison of four sizing configurations (corresponding to Figure B.2a). $P$ denotes the pool number. ADMI does not size for a QoS target. . . . .	84
6.4	Sensitivity analysis to QoS target (corresponding to Figure B.2b). $P$ denotes pool number. . . . .	84
6.5	Comparison of metrics across the two pools (corresponding to Figure B.3a). $P$ denotes the pool number. ADMI does not size for a QoS target. . . . .	84
6.6	The averages of each metric achieved over all periods across pools and for two different ADMI configurations (also see Figure B.1). $P$ denotes the pool number. ADMI does not size for a QoS target. $R$ is a boolean indicating whether removals are implemented. . . . .	88
6.7	The averages of each metric achieved over all periods for pool 0 and four different ELM configurations with and without removals (also see Figure B.3b). $P$ denotes the pool number. $R$ is a boolean indicating whether removals are implemented. Note the offline optimal baseline always removes cars when beneficial with respect to performance metrics. . . . .	88
7.1	Table of variables in the Bayesian network. O = Observable, H = Helper, Hd = Hidden, D = Discrete, C = Continuous . . . . .	102
7.2	Nissan Leaf Specifications [Nis14b] . . . . .	111
7.3	Percentage of fares lost in different BEV scenarios. BSS: battery switching stations. . . . .	111

7.4	Example showing how $\tau$ affects the number of battery swaps performed and the remaining SOC of swapped batteries at the airport switching station for one taxi. $q_b(\tau)$ (Eq(7.15) is computed assuming the 24kWh can charge in four hours from depleted using level 2 charging, i.e., $B_{Fx} = 4/24$ . With respect to the last column, the operator cannot buy a fraction of a battery, but the fractional batteries needed in the system per taxi per station are summed to compute $c_{ext}$ .	113
7.5	Chevrolet Volt Specifications [Che14a]	116
8.1	Sensor hardware and functionality. †: this is not the exhaustive list of sensors in the Galaxy S III, which, for example, also includes a barometer for sensing atmospheric pressure and a sensor for measuring whether there is a magnetic field being applied to the X,Y,Z axes.	126

# List of Figures

2.1	A Tesla Model S charging at a Level 3 PEVSE pillar ("Supercharger"). Image source: Tesla [ <a href="#">Tes14c</a> ]. . . . .	9
4.1	Sentiment analysis system architecture . . . . .	31
4.2	Scrapy[ <a href="#">Scr14</a> ] webcrawling process . . . . .	31
4.3	Pipeline for querying the sentiment of $(F, o)$ pairs . . . . .	37
4.4	Sentiments of all Leaf (top) and Volt (bottom) sentences in the ground truth corpus. There are three bars for each feature. The first shows the polarity distribution of sentences the system classified as positive, negative or neutral, even if for the incorrect choice of feature $(s^+, s^-, s^N)$ . The second shows the distribution of those the system correctly classified for the correct feature $(c^+, c^-, c^N)$ . Finally, the last shows the distribution of ground truth sentences $(t^+, t^-, t^N)$ . . . . .	41
4.5	Sentiments for all Tesla Model S sentences in the feature corpus. Because there is no GTC, we only show one bar, the polarity distribution of sentences the system classified $(s^+, s^-, s^N)$ . . . . .	41
4.6	Performance of classifying all Leaf (top) and Volt (bottom) sentences in the GTC. . . . .	43
5.1	Example arrival time distribution over one day to the carshare pool . . . . .	54
5.2	Figure 5.1 zoomed on [8,9] . . . . .	54
5.3	Subscribers active in the carshare over time . . . . .	61
5.4	The black distribution shows the probability of a specific number of vehicles in use. The grey distribution shows the CDF. . . . .	63
5.5	Interarrival time (top) and service time distribution (bottom) for all trips, observed vs. exponential using MLE. The grey region shows where the observed and predicted curves overlap. The black region shows observed arrivals, and the white region shows arrivals as predicted by the exponential distribution. . . . .	64
5.6	Interarrival time (top) and service time distribution (bottom) for trips >80km, observed vs. exponential using MLE . . . . .	64

5.7	Top: Arrival hour vs. frequency for all trips in our dataset. Bottom: daily number of arrivals for the last 300 days ( 10 months) in our dataset. The x-axis shows the day number starting from the first day in our dataset in March 2005; the region shown is approximately all 2013 data. . . . .	65
5.8	Comparison of achieved performance metrics for the full dataset, $\epsilon = .05$ . ELM <sub>S</sub> is abbreviated "S" and ELM <sub>D</sub> is abbreviated "D". . . . .	67
5.9	Comparison of achieved performance metrics for the full dataset, $\epsilon = .1$ . ELM <sub>D</sub> (2160) is drawn under ELM <sub>D</sub> (4368) because they achieve the same performance. ELM <sub>S</sub> is abbreviated "S" and ELM <sub>D</sub> is abbreviated "D". . . . .	67
5.10	Comparison of achieved performance metrics for trips over 80km only, $\epsilon = .05$ . ELM <sub>S</sub> is abbreviated "S" and ELM <sub>D</sub> is abbreviated "D". . . . .	68
6.1	New carshare architecture snapshot during period $p$ . . . . .	74
6.2	Joint Optimization of fleet size and vehicle movement, for a single period. . . . .	77
6.3	ADMI equivalence illustration . . . . .	79
6.5	Joint optimization program with systematic removals . . . . .	87
7.1	Bayesian network to simulate ICEV taxis as EVs by estimating their battery SOC . . . . .	101
7.2	Switching station placement optimization . . . . .	107
7.3	Points represent taxi mobility data. The grid shows the reduced coordinate space. We used a denser grid in downtown San Francisco due to the large number of data points in this region. . . . .	108
7.4	Distribution over pick-up and drop-off locations as distribution. Black shows the probability (percentage) of a pickup, red shows the probability of a dropoff, but the two distributions are nearly identical. The airport and Union square locations are annotated with "A" and "U" respectively. For YCSF, annotated with "Y", we used an identifier of 1. . . . .	112
7.5	Threshold $\tau$ vs. $r_B(\tau)$ as a percentage of $r_E$ (left Y-axis), $r_B(\tau)$ (right Y-axis) for varying gasoline prices, electricity fixed at \$.22/kWh. . . . .	112
7.6	Threshold $\tau$ vs. $r_B(\tau)$ as a percentage of $r_E$ (left Y-axis), $r_B(\tau)$ (right Y-axis) for varying gasoline prices, gas fixed at \$1.08/liter. . . . .	112
7.7	Gas and electricity prices vs. $r_B(\tau)$ using three switching stations. Figure does not take into account cost of switching stations, but takes into account the cost of extra batteries $c_{exbt}$ . . . . .	114
7.8	Gas and electricity prices vs. $r_P$ as a function of $r_E$ , $r_P$ . No roadside charging. All costs taken into account, including the cost of additional PHEVs $c_{expv}$ . . . . .	114
7.9	Region where BEVs are profitable after amortizing three switching stations. Electricity fixed at \$.22/kWh. The red line shows the current gas price in San Francisco (\$1.08/liter). . . . .	114
7.10	Left: Average U.S. price per gallon of gas since 1971 (2011 dollars) [Zfa10]. Right: Average electricity price given by three major utilities in California since 1980 [Tom07] . . . . .	116



7.11	Return on investment, per ICEV, in purchasing higher efficiency ICEVs. Gasoline fixed at \$1.08/l. The X-axis shows the assumed cost of acquiring higher efficiency ICEVs. The Y-axis shows the fuel efficiency, where 10.58km/l corresponds to 25mpg and 25.4km/l corresponds to 60mpg. The central point shows the case assumed in all prior results—an efficiency of 10.58—so the relative return is 0. The red arrow shows the return ( $\approx$ \$28000) of the 2014 Toyota Prius. . . . .	117
7.12	Distribution of fare trip lengths (the bar for 50 represents all trips over 50km)	118
8.1	Neural Network for predicting EV usage & adoption based on survey input parameters. . . . .	123
8.2	The eProdigy Whistler with our sensor kit attached. . . . .	126
8.3	eBike sensor kit assembly. Left: connected sensors, middle: sensors enclosed in weather proof case, right: battery + sensors + case assembly shown on bike. . . . .	127
8.4	eBike data pipeline . . . . .	128
8.5	Battery model received from battery manufacturer. The “mode” annotations are added by us and are not part of the original figure. The downward-sloping curves show the <i>already discharged</i> mAh (X-axis) vs. the battery voltage for six different battery temperatures. The capacity of each cell is $\approx$ 200mAh, and there are eight cells in each eBike battery. . . . .	131
8.6	Top: our transcription of the original battery model obtained by the battery manufacturer (Figure 8.5, scaled by 10 cells because there are 10 LiON cells in each battery. The curves are inverted to show the SOC on the Y-axis, because the SOC is the unknown variable we are modeling, and voltage (known) is shown on the X-axis. Bottom: our SOC estimation models from Eq(8.1). . . . .	131
8.7	Top: voltage readings over a subset of one day for one user. The blue curve shows the raw data, and the red curve shows the readings after a low-pass filter is run. Bottom: our SOC estimation using linear interpolation during periods of biking. The slightly decreasing SOC observed during non-biking periods represents the power consumed by the phone and sensor kit. . . . .	132
8.8	Frequencies of all single-word terms used at least 10 times. Red terms were classified as negative, green as positive, and blue as neutral. . . . .	136
8.9	Frequencies of all bigrams (double-word terms) used at least 4 times. Red terms were classified as negative, green as positive, and blue as neutral. . . . .	136
8.10	Frequencies of all single-word terms not classified as neutral used at least 5 times. Red terms were classified as negative, green as positive. . . . .	137
8.11	Frequencies of all bigrams not classified as neutral used at least 3 times. Red terms were classified as negative, green as positive. . . . .	137
8.12	Sentiments found in all sentences in the eBike survey corpus. “Usage Conditions” refers to sentiments found in the context of weather or seasons, e.g., “cannot be ridden in Winter”. . . . .	138

8.13	Trip length distribution for 24 participants from 9/7/2014—12/20/2014. .	139
8.14	Cumulative distance traveled by 24 participants from 9/7/2014—12/20/2014. . . . . .	140
8.15	The number of observed charging events v.s., the SOC of the eBike battery at the start of charging. The black curve shows how many charging events began when the SOC was in that range, and the grey curve shows the cumulation. . . . .	142
8.16	Distribution of charging events for 24 participants from 9/7/2014—12/20/2014.	142
8.17	This figure shows the percentage of the eBike battery that was depleted during all trips taken by the 24 participants from 9/7/2014—12/20/2014. Each circle represents a trip. The X-axis shows the length of the trip. The annotations show the percentage of the battery was depleted on the trip. The Y-axis shows the range of the eBike assuming all of the user’s trips required the same amount of battery per kilometer as the trip shown. For example, if a 5km trip requires 25% of the battery, the Y-axis will show an empirical range of 20km. . . . .	143
8.18	WeBike user portal interface . . . . .	145
8.19	reproduced from Dozza and Fernandez [DF14b]. "IMU": inertial measure- ment unit. . . . .	146
B.1	Top left: ADMI(1,.5,30,1) without removals, top right: ADMI(1,.5,30,1) with removals implemented, bottom left: ADMI(1,.5,15,1) without removals, bottom right: ADMI(1,.5,15,1) with removals implemented. These are plot- ted together with the offline optimal baseline for a 95% QoS target. The averages of each metric achieved over all periods are summarized in Table 6.6. . . . .	175

# List of Symbols

---

Chapter 4 Notation Reference Table

---

$o$	indexes an opinion phrase
$F$	indexes a feature
$\mathcal{P} = \{P_1, P_2, \dots\}$	set of products the user wishes to mine reviews for
$\mathcal{F}^p = \{\vec{\mathbf{F}}_1^p, \vec{\mathbf{F}}_2^p, \dots\}$	feature space of $p \in \mathcal{P}$
$\vec{\mathbf{F}}_i^p$	vector of synonyms for feature $i$
$\mathcal{O}$	set of all recognized opinion phrases
$(F, o)$	feature-opinion pair
$N/ + / -$	Neutral/positive/negative
$s_F^+, s_F^-$	# sentences the system classifies as $+, -$ about $F$
$t_F^+, t_F^-$	# sentences we manually classify as $+, -$ about $F$
$c_F^+, c_F^-$	# sentences the system correctly classifies as $+/ -$ about $F$

---

Chapter 5 Notation Reference Table

---

$s$	indexes a subscriber
$S$	the number of pool subscribers
$l$	number of bootstrap iterations
$\epsilon$	QoS target
$K$	the length of the pools busy periods
$p(\cdot)$	denotes "probability of ( $\cdot$ )"
$p(a)$	daily Binomial probability each subscriber will make a request (§5.2.1)
$p(b m)$	blocking probability given a pool size of $m$
$m$	the size of the pool
$m_{max}$	maximum size of the pool
$\rho(s, B)$	probability $s$ arrives given the pool is in a busy period
$1/\lambda_s$	$s$ 's mean think time (MTT)
$\lambda_s$	$s$ 's mean think rate (MTR)
$1/\lambda_s$	$s$ 's mean service time (MST)
$\lambda_s$	$s$ 's mean service rate (MSR)
$\psi$	"offered traffic" parameter
$\rho_s$	$\lambda_s/\mu_s$ , a measure of $s$ 's activeness
$1/\lambda_B$	MTT of all subscribers during busy periods
$1/\mu$	MST of all subscribers

$\rho$	the weighted mean of all $\rho_s$ given by $\lambda_B/\mu$
$B$	the pool's busiest period(s)
$A_B$	the arrival rate of subscribers to the pool during $B$
$c_s$	$s$ 's mean cycle time
$1/c_B$	average cycle time of subscribers during busy periods
$\omega_s$	$1/c_s$ ; $s$ 's cycle rate
$n_B$	number of subscribers that arrive during busy periods
$\mathcal{K}/ \mathcal{K} $	the set/number of pools potential busy periods (§5.2.2)
$\sigma$	desired iteration accuracy of Algorithm 1
$D(\cdot)$	legend notation; $\cdot$ = length of historical busiest period $ELM_D$ sizes for
$S(\cdot)$	legend notation; $\cdot$ is the set of daily chunks $ELM_S$ uses

---

Chapter 6 Notation Reference Table

---

$s$	indexes a subscriber
$j$	indexes a pool
$p$	indexes time periods
$l$	number of bootstrap iterations
$\tilde{\cdot}$	prediction of parameter/variable $\cdot$
$\epsilon$	QoS target
$K$	the length of the pools busy periods
$\mathcal{J}$	set of all pools
$\mathcal{S}_j^p$	denotes the set of members assigned to $j$ during $p$
$S_j^p$	$ \mathcal{S}_j^p $ , the number of members of pool $j$
$m_j^p$	size of $j$ during $p$
$\tilde{m}_{optj}^p$	calculated minimum size of $j$ during $p$ needed to meet QoS
$\gamma_j^p$	blocking probability of $j$ during $p$ .
$1/\lambda_j^p$	mean think time during busy periods of $j$ during $p$
$1/\mu_j^p$	mean service time of $j$ during $p$
$V_{j,k}^p$	cost to move a car between pools $j$ and $k$ at (the start of) $p$
$A_j^p$	cost to buy a car for pool $j$ at $p$
$R_j^p$	profit from selling a car from pool $j$ at $p$
$v_{j,k}^p$	number of cars to move from $j \rightarrow k$ at $p$
$a_j^p$	number of cars to buy for $j$ at $p$
$r_j^p$	number of cars to sell from $j$ at $p$
$U$	the utilization performance metric
$A$	the availability performance metric
$M2V$	the M2V performance metric
$M$	figure notation to indicate the pool size

$K$	length of ELM busy period
$VPS_j$	ADMI virtual pool size variable for pool $j$
$ADMI_{in}$	initial guess of pool size for ADMI
$\alpha, \beta$	ADMI dynamic decrease and increase parameters
$y$	# of served requests before decrease triggered for ADMI
$u_j^{t \in P}$	# members of pool $j$ assigned a car at time instant $t$ in period $p$

---

Chapter 7 Notation Reference Table

---

$i$	indexes a taxi
$f$	indexes a fare
$b$	indexes a battery switching station
$t_k$	indexes timesteps ( $t_k$ is a fixed point in time)
$w$	indexes a shift
$T$	total number of taxis the company has
$\mathcal{T}_i$	set of all fares completed by taxi $i$
$\mathcal{W}_i$	set of all shifts completed by taxi $i$
$L$	the replacement rate (every $L$ years) for all ICEVs, BEVs, PHEVs, and batteries
$\mathcal{P}(X)$	parents of node $X$ in Bayesian network
$\Delta_B(\tau)$	total ROI in transitioning $x$ ICEVs to BEVs
$\Delta_P$	total ROI in transitioning $x$ ICEVs to PHEVs
$r_B(\tau)$	average ROI per BEV over $L$
$r_P$	average ROI per PHEV over $L$
$x$	no. of vehicles the company transitions to EVs
$r_E$	average revenue generated by the company's existing ICEV taxis
$r_L(\tau)$	average revenue lost per taxi from missed fares
$r_T^i(w, \tau)$	revenue generated during shift $w$ of ICEV taxi $i$
$r_B^i(w, \tau)$	revenue generated during shift $w$ of BEV taxi $i$ before battery depletion
$s_B(\tau)$	average fuel savings per taxi from using BEVs instead of ICEVs
$s_P$	average fuel savings per taxi from using PHEVs instead of ICEVs
$r_{FARE}(f)$	revenue generated from fare $f$
$V_E$	efficiency of an ICEV (mpg)
$B_E$	efficiency of the BEV (km per kWh)
$P_E$	efficiency of the PHEV (km per kWh)
$P_G$	efficiency of the PHEV after battery depletion (mpg)
$p_f^i$	time the company's taxi $i$ is parked at lights while on fair $f$
$d_S^i$	total distance driven by the company's ICEV taxi $i$
$d_B^i(\tau)$	total distance driven by BEV taxi $i$ before depletion over all shifts

$d_E^i$	total distance driven by PHEV taxi $i$ on electricity
$d_G^i$	total distance driven by PHEV taxi $i$ on petroleum
$n_{BSS}$	the number of switching stations needed
$q_j(\tau)$	average no. of batteries needed at $b$ for each taxi (fractional)
$\lambda_j(\tau)$	average no. of battery switches at $b$ per taxi per day
$l_j(\tau)$	average remaining charge level (kWh) of batteries switched at $b$
$\lambda_H$	rate at which PHEVs return to headquarters (PHEVs/day)
$B_{gx}$	energy gained per second (kWh/second) assuming level $x$ charging
$B_{Fx}$	days (fractional) to charge a fully depleted battery w/ level $x$ charging
$\zeta$	the capacity of each battery (kWh)
$R_{Px}$	PHEV charging rate (days/PHEV) assuming level $x$ charging
$D$	discharge rate of the EV (kWh/km)
$C_g$	cost of gas per liter
$C_e$	price of electricity per kWh
$C_{bat}$	the cost of one battery (dollars)
$C_{phev}$	full cost of a PHEV (dollars)
$c_{exbt}(\tau)$	average cost of BEV batteries needed per taxi
$c_{expv}$	cost of extra PHEVs, per taxi, so drivers start shifts with full batteries
$c_{BSS}(\tau)$	cost of all battery switching stations needed
$C_{bev}$	cost(BEV - ICEV); incremental BEV cost
$C_{phev}$	cost(PHEV - ICEV); incremental PHEV cost
$\varphi(i, w)$	<i>opportunity cost</i> of $i$ during $w$
$n_{BSS}$	number of switching stations to locate
$\mathcal{Y}$	set of possible switching station locations
$\text{cost}(l)$	price of placing a station at $l \in \mathcal{Y}$
$\Upsilon(\cdot)$	indicates that location $\cdot$ has a switching station

# List of Acronyms

<i>Abbr</i>	<i>Description</i>
ADMI	additive decrease multiplicative increase sizing method (Chapter 6)
ADS	arrival departure stream (Chapter 5)
AER	all electric range
BEV <sub>x</sub>	battery electric vehicle with x km of range
BSS	battery switching station (Chapter 7)
BIN	binomial sizing method (Chapter 5)
CDOP	context dependent opinion phrase (Chapter 4)
CFG	context free grammar (Chapter 4)
DCLGN	dynamic conditional linear Gaussian network (Chapter 7)
ELM <sub>S/D</sub>	Engset Loss Model, survey/dataset variant (Chapters 5)
ELM	Engset Loss Model (Chapters 6)
EPRI	Electric Power Research Institute
EREOI	energy returned on energy invested
EV	electric vehicle (inclusive of BEV + PHEV)
FBOM	feature based opinion mining (Chapter 4)
FSD	feature specific dictionary (Chapter 4)
GPS	global positioning system
GTC	ground truth corpus (Chapter 4)
HCI	human computer interaction (Chapter 8)
HEV	hybrid electric vehicle
ICE(V)	internal combustion engine (vehicle)
IEA	International Energy Agency
LDV	light duty vehicle
LiON	Lithium-ion battery
M2V	member-to-vehicle ratio (Chapters 5, 6)
MLE	maximum likelihood estimation (Chapter 5)
MPQA	multi-perspective question answering opinion corpus (Chapter 4)
MSRP	manufacturer suggested retail price
MSR	mean service rate (Chapters 5, 6)

MST	mean service time (Chapters 5, 6)
MTR	mean think rate (Chapters 5, 6)
MTT	mean think time (Chapters 5, 6)
NAOP	non-adjective opinion phrase (Chapter 4)
NLTK	Natural Language Toolkit (Chapter 4)
OPEC	Organization of the Petroleum Exporting Countries
OPT	optimal baseline sizing method (Chapter 6)
(P)EVSE	(public) electric vehicle supply equipment
PHEV <sub>x</sub>	PHEV with $x$ km of AER
PPE	percent prediction error (Chapter 6)
QoS	quality of service (Chapters 5, 6)
RA	range anxiety
ROI	return on investment
SOC	battery state of charge
TCO	total cost of ownership
TCP	transmission congestion protocol (Chapter 6)
VKT	vehicle kilometers traveled
WTP	willingness to pay
YCSF	Yellow Cab San Francisco (Chapter 7)



# Chapter 1

## Introduction

If oil demand outpaces available supply, the effects could be devastating to the global economic, transportation, and food systems [HBW05, Llo10, Bun10, Ele09, NPK<sup>+</sup>11]. As evidence, ten out of the past eleven major recessions, including the “Great Recession” of 2008-2009, followed or coincided with a sharp increase in oil prices [Ham11, Ele09, U.S14c], and our modern food supply and agricultural systems—everything from farming equipment to food transport to fertilizer and pesticide production—is dependent upon oil [NPK<sup>+</sup>11].

According to OPEC, there are  $\approx 40$  years left of reserves remaining at current consumption levels [OPE12, Int12]. However, *light sweet crude* is nearing depletion [GHOC11, HBW05]. Remaining crude is harder to find and *heavier*, i.e., it requires much more refining. Consequently, the *energy returned on energy invested*<sup>1</sup> (EROEI) of finding oil has decreased exponentially from 1200 in 1919 to just 5 currently [GHOC11]. Moreover, many remaining reserves are located in oil sands and in offshore deposits, both of which are environmentally dangerous to extract in addition to having low EROEIs. Tar sands, for example, have an EROEI of only 1.6 [Inm13]. The 2010 Deepwater Horizon offshore drilling disaster is now cited as the *worst* human-induced environmental disaster in U.S. history [Tim10, RFI10], and the 2013 Bakken oil shale spill is one of the largest on-shore oil spills in U.S. history [Glo13]. practice. Finally, oil supply is not a *robust* system. More than 80% of all remaining proven reserves are controlled by OPEC [OPE12] so an embargo would threaten global supplies, which has historical precedence: the 1973 oil crisis [U.S13b].

To remove our dependence on oil and avert such problems, we should begin by transitioning light duty vehicles (LDVs), which include passenger cars, to use an alternative fuel source<sup>2</sup>. The question then becomes, *what should we fuel our cars with?* There is strong evidence that the answer is *grid generated electricity*. That is, we should transition our internal combustion engine vehicles (ICEVs) to *electric vehicles* (EVs) instead of cars powered by natural gas, hydrogen, or ethanol:

---

<sup>1</sup>The EROEI gives the number of energy units returned for every unit invested.

<sup>2</sup> LDVs dominate world oil consumption. As a case study, the U.S. LDV fleet alone requires >8 million barrels of petroleum per day, constituting 60% of U.S. transportation energy use and 42% of all U.S. petroleum use [U.S13d, U.S13c]. As a comparison, air travel represents only 9% of transportation energy use and 2.6% of all oil use [U.S13d, U.S13c].

- It is possible to convert ICEVs to run on natural gas, but natural gas has a much lower EREOI than other grid generated electricity sources like coal. Even clean renewables such as hydro and wind have EREOIs of 40+ and 20 respectively, compared to just 7-10 for natural gas [Inm13, Hei09]. Moreover, the EREOI of natural gas is steadily declining because natural gas is often found alongside oil [GHOC11, Hei09]. Finally, natural gas is a finite<sup>3</sup> fossil fuel. Even if conditions become perfect (increased discoveries, better searching and production techniques, large increase in EREOI, etc) to transition LDVs to natural gas, it will only represent an interim transition. Transitioning our cars from one fossil fuel to another, especially one already used for critical applications such as home heating, simply delays our problem, and does not reduce CO2 emissions either.
- EVs are currently 6-10 times more efficient than hydrogen powered vehicles [Mac09]. Hydrogen suffers from multiple problems [Mac09, Hei09]. First, it is only a storage medium like a battery; energy must be converted to and from hydrogen to use it as a transportation fuel, and this double-conversion process is inefficient. Second, hydrogen has low energy density even when stored at very high pressures, requiring vehicles to have impractically large storage tanks. Finally, hydrogen leaks and cannot be stored in vehicles for long durations of time without dissipating.
- Ethanol, a fuel produced from corn and other crops that many propose as an alternative to gasoline, is the easiest to dismiss: it currently has a *negative* EREOI; more energy is invested than returned [Inm13]. Moreover, ethanol is produced on land that can otherwise be used for agriculture and livestock.
- EVs are energy efficient. To compare to our current cars, 75% of gasoline in an ICE is wasted as heat and less than 25% of input fuel is converted to useful energy [Mac09]. This remains true after a century of developments. However, EVs are 60-90% efficient in converting input energy<sup>4</sup> to useful energy [Mac09]. As an example, Tesla claims the Model S is 88% efficient [Tes14a].
- EVs can be powered using sustainable and renewable sources of energy including solar, wind, and hydro. See, for example, the E-carsharing Sylt project offering EV rentals that are charged using 100% renewables [E-W14]. Moreover, it is more efficient to store renewable sources directly using batteries than to convert them to and from hydrogen.

Our hypothesis is thus: *the majority of LDV ICEVs in our transportation system will eventually be replaced with EVs*. This is also the view of several major auto manufacturers who are now mass-producing EVs.

Transitioning from century-old ICEVs to EVs represents a large change for drivers. EV sales

---

<sup>3</sup>Currently the primary battery technology for EVs is Lithium Ion (LiON) batteries, which are 3% composed of lithium, a finite metal [Mac09]. The lithium in used batteries is 100% recyclable although the recycling process is currently unprofitable [Kum13]. Moreover, the concept of storing grid generated electricity in a battery, i.e., an electric vehicle, is independent of the battery technology—future EVs may use a different battery technology to store electricity than LiON. Battery chemistries currently under research and development, e.g., phosphate and manganese based batteries, have little or no valuable metals [Kum13].

<sup>4</sup>By input energy we mean “from the charging outlet”, but there are also transmission and distribution losses that sum to about 7% between the power plant and the charging outlet [U.S12].

are steadily increasing [Ele14] but are still compose a negligible amount (1.6% in August 2014 [Arg14]) of total LDV sales. From the comprehensive literature review discussed in §3, we conclude the three most prominent adoption barriers are as follows.

1. Pure electric vehicles, or “BEVs”<sup>5</sup> must be recharged between trips and have limited range on a single charge, compared to ICEVs which can be refueled in minutes at widely deployed stations. Some BEV owners and prospective buyers consequently have *range anxiety*, the term given to the fear of depleting their batteries before reaching their destination. This term also refers to the fear that, for the foreseeable future, BEVs will be unusable for long trips.
2. EVs have a higher upfront purchase price than ICEVs. However, the cost per km driven using electricity is significantly lower than the cost per km driven on gasoline, so over the lifetime of the vehicle drivers can recoup some or all of their initial investment through fuel savings. Most drivers cannot or do not compute this long term return on investment (ROI) when purchasing a vehicle, thus are wary to pay the higher initial price of EVs despite possible long term savings. Moreover, the resale price is not known.
3. The lack of detailed knowledge of how early adopters of EVs perceive their vehicles is also a barrier to wider adoption. For example, early adopters’ perceptions can help manufacturers build improved EVs that more drivers are likely to adopt, and can help marketers target EVs appropriately (i.e., EVs are not well suited for all mobility patterns or geographical regions). Unfortunately, field trials of sufficient size and length are very expensive to conduct, so opinions towards EVs have mostly been elicited from non-EV drivers through mass surveys.

Other barriers have also been cited in the literature, but the above are the three most cited barriers to adoption, so this thesis focuses on three problems. Specifically, this thesis is organized as follows:

- We overview EVs and their corresponding technologies in Chapter 2.
- We present a survey of related work in this area in Chapter 3.
- To determine sentiments towards EVs, we present a system that mines perceptions from online ownership forums in Chapter 4.
- Towards alleviating range anxiety, Chapters 5 and 6 present work on sizing networks of ICEVs to be used by BEV drivers when a planned trip exceeds their BEV range
- We present an ROI model for taxi fleets, a viable market segment for EV adoption, to determine when it is profitable to transition to EVs in Chapter 7.
- We design the most comprehensive *electric bike* trial to date in Chapter 8. This is in effort to explore the barriers to adopting this form of EV in North America<sup>6</sup>.
- We summarize our contributions and conclude in Chapter 9.

---

<sup>5</sup>Chapter §2 gives an overview of the three types of commercially available EVs.

<sup>6</sup>Electric bikes, two wheeled EVs as described in §2.2, are widely adopted in countries like China where car ownership (EVs and ICEVs) are prohibitively expensive, but are seldom adopted in North America.

# Chapter 2

## Technology Overview

This chapter presents an overview of the vehicle technologies explored in this thesis. In §2.1, we describe the three major types of EVs. We overview electric bikes in §2.2, as they are the subject of Chapter §8. An overview of battery technologies is given in §2.3. We describe the “refueling” process of EVs, *charging*, in §2.4. Finally, a glossary of domain specific terms is given in §2.5. For further information on these technologies, please refer to the IEA [IEA13], McKinsey [MHN+12], and Young et al. [YVWS12].

### 2.1 Electric Vehicles

Most of this thesis focuses on measuring and mitigating the adoption barriers to *electric vehicles* (EVs). Here we describe t/he three types of EVs available as of 2014. Throughout this thesis, we use the term “ICEV” when discussing standard gasoline cars powered by internal combustion engines (ICEs), “EV” when the statement applies to both BEVs (§2.1.1) and PHEVs (§2.1.2), and “PHEV”/“BEV” when the statement applies only to the specific type of EV.

#### 2.1.1 Battery Electric Vehicles

Battery Electric Vehicles (BEVs) are fully powered by batteries and do not have an ICE. We largely focus on the adoption of BEVs because these vehicles use no petroleum. Because BEVs do not have an ICE, they have a limited range on a single charge. When the battery is depleted, the vehicle cannot be used until the battery is recharged. Fully recharging takes on the order of hours and is dependent on the battery size and the specifications of the BEV and charging equipment (see §2.4).

The two best selling BEVs at the time of writing are the Nissan Leaf and the Tesla Model S. In 2013, Nissan sold 22,610 Leafs and Tesla sold 22,450 Model Ss [Ohn14, Voe14]. The Leaf and the Model S have ranges of 160km (24kWh battery) and 480km (84kWh battery), and MSRPs (after tax incentives) of \$21,500 and \$63,570 [Nis14a, Tes14c] respectively.

The Nissan Leaf debuted much earlier than the Tesla so it is used as the “standard” BEV, primarily with respect to range, throughout this thesis. The Model S began selling after most of the research in this thesis was published.

### 2.1.2 Plug In Hybrid Electric Vehicles

Plug-in Hybrid Electric Vehicles (PHEVs) have a grid-chargeable battery and an ICE. PHEVs are powered from the battery for first portion of a trip without using the ICE. The standard notation to describe a PHEV is “PHEVxxm”, where xx refers to the distance in kilometers (km) the PHEV is expected to drive using only the battery, known as the “all electric range” (AER). Once the battery is depleted, the ICE is used for the remainder of the trip and the battery may continue to power electronics onboard the vehicle<sup>1</sup>. PHEVs do not have range limitations due to the ICE.

Currently, the best selling PHEV is the Chevrolet Volt. The Volt has an AER of 64km (16kWh battery) and has an MSRP of \$26,685 after tax incentives [Che14b]. In 2013, Chevrolet sold 23,094 Volts [Voe14].

### 2.1.3 Hybrid Electric Vehicles

Hybrid Electric Vehicles (HEVs) are ICEVs fitted with a small onboard battery and electric motor. The battery is charged by capturing energy from regenerative braking and powers the electric motor which assists with onboard electronics and acceleration. HEVs are completely dependent on petroleum and are not charged from the grid, though they have higher fuel efficiencies than ICEVs. Therefore, this technology is not studied in this thesis. To date, the best selling HEV is the Toyota Prius [Toy14a], which currently (2014) has an efficiency of 51mpg and an MSRP of \$24,000.

### 2.1.4 Price Comparisons

Here we briefly compare the price of the aforementioned EVs to each other and to comparable ICEVs. We make these comparisons because, as discussed in our literature review in §3, the relatively high initial cost of EVs remains one of their biggest adoption barriers.

First, because a BEV’s battery constitutes a large fraction of the vehicle price, and the cost of batteries increases linearly with size (see §2.3), the price of BEVs tends to increase linearly with range. This can be seen, for example, when comparing the Model S with the Leaf; the Model S has 3x the range and costs 3x as much. Their identical sales thus far

---

<sup>1</sup>Some PHEVs operate in “blended operation” mode, where the battery and ICEV are always both used in contrast to using the ICE only upon depletion. However, most mass market PHEVs like the Volt operate as described above, so we only consider PHEVs of this type for simplicity.

suggests there are markets for both higher range EVs for a higher price and lower range EVs for a lower price.

Second, because PHEVs have both an ICE and a battery, they cost more than BEVs with comparable battery sizes. The Volt, while only having a 16 kWh battery, costs \$5,000 more than the Leaf, which has a 24 kWh battery. The premium is due to the ICE and the need for a more complicated drivetrain that integrates both the ICE and battery.

To compare the Leaf and Volt to similar non-luxury ICEVs, a 2014 Toyota Corolla is \$16,800 MSRP [Toy14b] and a 2014 Honda Civic is \$18,390 MSRP [Hon14]. Thus, the Leaf has a  $\approx$ \$3-5k premium over comparable ICEVs and the Volt has an  $\approx$ \$8-10k premium. While the Model S costs significantly more than the Leaf due to its larger battery and range, it is marketed as a luxury sports car and contains many vehicle features found only in luxury ICEVs. The Model S is comparable to the Audi A7 (\$64,500) [Aud14] and the Mercedes-Benz E Class (\$58,000) [Mer14].

Finally, many sources predict that EV prices will fall due to decreasing battery prices (see §2.3) and future competition. At the time of writing, the EV market is still in its infancy; currently, the Leaf, Volt, and Model S are the only three PHEVs/BEVs with non-trivial sales. However, many manufacturers are soon releasing EVs hoping to gain a share of this upcoming market—three times as many more models are coming to market in late 2014 and 2015 alone than are available to date [Ing14]. The next releases will occur in late 2014 when BMW begins selling three EVs [BMW14].

## 2.2 Electric Bikes

Chapter §8 focuses on barriers to *electric bicycle* ("eBike") adoption. EBikes are bicycles extended with an on board LiON battery. When desired, the rider can turn on various levels of *electric assistance*, provided the battery is not depleted, for easier travel. The major benefit of eBikes v.s. their conventional counterparts is easier traveling over long distances and in hilly areas. EBikes are in the scope of this thesis for three reasons:

1. They are a form of petroleum-free electric transit, and have no emissions.
2. They are widely adopted in some regions like China where car-ownership, including EV ownership, is prohibitively expensive, but are much less adopted in other regions like North America. However, little is known about *why* there is such a disparity in sales.
3. They are essentially two-wheeled PHEVs<sup>2</sup>. When the cyclist is using partial or full electric assistance, it is analogous to a PHEV prior to battery depletion, and when the cyclist must pedal because the battery is depleted, it is analogous to a PHEV engaging the ICE. In addition, eBikes use the same battery technology as EVs and are charged identically.

---

<sup>2</sup>A minor distinction is that a PHEV engages the ICE usually only once after the battery is depleted, whereas an eBike user can alternate, if they wish, between pedaling and electric assist.

EBikes typically provide about 50km of range using a  $\approx$ .350kwh battery and take about five hours to charge. The eBike studied in Chapter §8 is discussed in detail in §8.2.1 and is shown in Figure 8.2.

## 2.3 Batteries

Currently, the two major battery technologies used in EVs are nickel metal hydride (NiMH) and Lithium-ion (LiON). Many HEVs, like the Prius, use NiMH batteries because they predate BEVs and PHEVs like the Leaf, Tesla, and Volt. All current PHEVs/BEVs use LiON due to its increased energy density v.s. NiMH batteries. For a detailed discussion of battery technologies and their operation, see Young et al. [YVWS12].

Two related research areas (beyond the scope of this thesis) regarding batteries are *increasing energy density* and *decreasing cost*. Building EVs comparable to ICEVs has been challenging because gasoline has very high energy density, 3,000 Wh/kg, compared to current battery technologies such as LiON (whose energy density is only 120 Wh/kg). Lower energy density means large batteries are needed to build EVs with significant range adding to weight and cost. Fortunately, the price of LiON batteries has fallen fast from \$1000/kWh in 2008 to only \$450/kWh in 2013, and prices are projected to continue falling. Tesla recent announced their "Gigafactory" to produce LiON batteries at a massive scale [Tes14b], hoping to drive prices below \$200/kWh. McKinsey projects this will happen by 2020 [MHN<sup>+</sup>12], not only because of demand from the EV industry, but also because LiON is the primary battery technology used in ubiquitous computing devices such as cell phones, laptops, and tablets [MHN<sup>+</sup>12]. To quote, "*Many advances in battery technologies are likely to be achieved first in consumer-electronics applications where manufacturing volumes and fierce competition facilitate price reductions before making their way into the automotive industry.*"

Future battery technologies may have much higher energy densities than current LiON batteries. Two technologies under development are Lithium-air (Li-A) and Lithium-sulphur (Li-S) batteries [VN14]. IBM's "Battery 500 Project" has the goal of "improving battery energy density tenfold" using Li-A batteries [IBM12]. To date, IBM researchers have demonstrated the fundamental chemistry of the charge-and-recharge process for Li-A batteries. In 2012, the Department of Energy granted the US Center for Energy Storage Research \$120M for Li-S research. The stated goal is to make cells "five times more energy dense, and five times cheaper, in just five years". When Li-S was first proposed over 40 years ago, the batteries had very short lifespans, only 100 charging cycles, but modern Li-S cells can now maintain half of their capacity after 1500 cycles, a performance on par with LiON cells [VN14]. However, these Li-A and Li-S battery technologies are still in the research and development stage and are not currently available to consumers.



## 2.4 Charging (“Refeuling”)

BEVs and PHEVs are “refueled” by plugging them into a charging unit with a connection to the electrical grid, a process known as *battery charging*. These units transfer power from the grid into the battery at a *charging rate* dependent on the type of charging unit, the specifications of the EV, and the availability of power in that location. This process takes under an hour to many hours depending on the charging rate. There are three standardized levels of charging:

1. Level 1 charging, commonly referred to as “slow charging” or “overnight charging”, is the slowest form and uses a 110V AC connection found in any standard electrical outlet in North America. A typical charge time from depleted-to-full for a BEV160 like the Leaf is 8-12 hours and much longer for BEVs with large batteries like the Model S.
2. Level 2 charging uses a 220V AC connection, which is standard in Europe and also present in most North American homes for large appliances. Typically, EVs can be charged in half the time using level 2 charging compared to level 1 charging.
3. Level 3, commonly referred to as “quick charging” or “DC fast charging”, uses a  $\geq 480V$  DC connection. Level 3 chargers have very high charging rates. For example, the Tesla Superchargers (shown in Figure 2.1) charge up to a rate of 120 kW, which is capable of replenishing half of the 84kWh battery in 30 minutes [Tes14c]. This is equivalent to fully charging a Leaf in 15 minutes.

The other option to refuel EVs is *battery switching*. Here, a battery switching station physically changes the batteries in an EV. A user comes to the station with a nearly-depleted battery and the battery is replaced with a fully charged battery. This process happens in a matter of minutes. The depleted batteries are then charged at the switching station. This type of refueling is similar to petroleum vehicles, where the vehicle is refueled in minutes instead of hours. The first major manufacturer of switching station infrastructure, Better Place, is recently (2013) defunct [Cha14a]. However, other EV manufacturers, such as Tesla [Tes14c], have an interest in battery switching.

## 2.5 Glossary

The following terms associated with EVs are used throughout this thesis:

- *Battery charging*: refueling an EV by plugging the vehicle into a charging outlet.
- *Capacity*: the amount of energy (usually stated in kWh) that can be stored in its battery.
- *Discharge rate*: the rate at which the EV consumes energy in kwh/km.
- *Kilowatt (kW) & Kilowatt-Hour (kWh)*: a kW is a unit of power equal to 1000 watts. A kWh, the basic unit of energy in which EV batteries are described, is the energy consumed by a device drawing a kW of power for one hour (3,600 joules).





Figure 2.1: A Tesla Model S charging at a Level 3 PEVSE pillar ("Supercharger"). Image source: Tesla [Tes14c].

- *PHEVxx*: refers to a PHEV that can be driven for xx km (on average) on battery before the ICE engages.
- *(Public) Electric Vehicle Supply Equipment ((P)EVSE)*: commonly referred to as "charging stations" or "charging points", these stations deliver electrical energy from an electricity source to charge an EV's batteries. "Public" is used when referring to a charging station outside of the home. EVSE can be a standard electrical outlet, but most EV manufacturers sell dedicated PEVSE pillars that include additional functionality, e.g., allowing the charging rate to be controlled by a smart phone. Figure 2.1 shows an example of a PEVSE pillar (Tesla's Supercharger).
- *Range*: the expected distance it can travel given normal driving conditions when fully fueled (using petroleum for ICEVs, electricity for BEVs, and both for PHEVs). PHEV manufacturers additionally state the range that can be driven using electric power before the ICE is used.
- *Regenerative Braking*: a process that stores some of the energy lost in the braking process back into the EV battery.
- *State Of Charge (SOC)*: the remaining charge level of an EV battery, analogous to the remaining fuel level in an ICEV.

# Chapter 3

## Related Work

This chapter presents related work in the following four areas:

1. Gauging drivers' perceptions towards EVs and eBikes (§3.1)
2. Range anxiety (§3.2)
3. Determining the ROI for drivers and fleets transitioning to EVs (§3.3)
4. Using GPS to infer mobility demands (§3.4)

The first three directly correspond to the three areas explored in this thesis as discussed in Chapter 1. In addition, several of our methodologies depend on measuring drivers' mobility patterns from GPS traces, so related work in this area is also presented.

To complement the rest of this section, we refer the reader to the Electrification Coalition whitepaper [Ele09] and Boulanger et al. [BCMW11] for excellent overviews of the EV ecosystem. These reports overview EVs and discuss the barriers to adoption, consumer education, the adoption of charging standards, PEVSE deployment, battery technology, and the interaction between EVs and the electrical grid.

### 3.1 Gauging Driver Perceptions

Understanding consumer perceptions is critical for manufacturers that want to build EVs that drivers want to adopt. Researchers and manufacturers have thus far used two primary tools to identify drivers' perceptions towards EVs. First, many have conducted field trials where participants were given EVs to drive and were periodically interviewed about their experiences. Others have interviewed drivers as to their perceptions towards EVs. These papers are described in §3.1.1 and §3.1.2 respectively. Morton et al. [MSA11] who summarizes other techniques used in prior work to gauge drivers' perceptions including rational choice theory, social interaction theory, prospect theory, and theory of planned behavior.

### 3.1.1 Field Trials

#### Electric Vehicle Field Trials

In this section we discuss several EV field trials. In these trials, participants were supplied with EVs and monitored. Monitoring consisted of drivers recording their trip information in travel diaries, surveys, and interviews throughout the trials. In some cases, vehicles were also fitted with GPS data loggers that recorded location and charging information. The primary purpose of conducting field trials is to determine how perceptions towards EVs change as drivers become more experienced with the technology.

While field trials are useful for drawing conclusions about drivers' experiences with and perceptions towards EVs, they are subject to at least three limitations:

1. Field trials are expensive because multiple EVs must be purchased or leased for the trial. It is therefore expensive, especially for academic researchers, to conduct field trials with a large number of participants for significant durations.
2. During shorter trials, drivers may not have time to adjust to driving BEVs or have time to derive well-informed conclusions. As a consequence of (1), conclusions drawn from field trials are usually drawn from a small number of still-inexperienced drivers, hence they may not be widely applicable.
3. Some drivers stated they changed their normal driving habits during the trials to fully explore and "push" the vehicles' capabilities, thus the results may not indicate whether EVs are suitable for their "normal" driving behavior. This behavior is similar to the *Hawthorne effect*, which states subjects in an experiment often alter their behavior for the duration of an experiment [Mac07].

To summarize, it is very expensive to conduct large (in terms of number of participants) and long (duration) field trials, but conclusions from smaller and shorter trials may not be valid. Here we present the parameters (participants, location, length of duration, data collection methods, vehicle specifications) of the most significant (in terms of size and duration) field trials that have taken place. We first review the details of each trial, then present the common conclusions drawn from these trials.

Bunce et al. [BHB14] present results from the Technology Strategy Board's UK based Nationwide Ultra Low Carbon Vehicle Demonstrator Programme. 135 participants completed questionnaires and were interviewed prior to the trial and after driving the EV for three months. Some drivers continued driving the vehicle for an additional nine months. Most drivers were charged a few hundred pounds per month to lease the EV, in contrast to other trials where the vehicles were supplied to participants. The nine types of BEVs evaluated had different vehicle specifications. The vehicles were fitted with GPS loggers that also record charging metrics, but the exact data collected from the loggers is not discussed. Everett et al. [EBH<sup>+</sup>11] present preliminary results as part of the same program.

Franke et al. report on the "MINI-E trial" in Germany [FK13]. The trial consisted of two sub-trials where 40 participants drove BEVs for six months and were interviewed prior to

the trial, three months into the trial, and again when the trial was complete. The BEVs were MINI E Coopers with a 28kWh battery pack and a range of 160km. Charging habits were recorded by the drivers through travel diaries, and the cars were fitted with GPS data loggers to record location, speed, parking habits, and time of use. The authors find a large disparity between drivers' range preferences and their range requirements. Pre-trial and post-trial surveys of the participants about their range preferences indicated that *a priori*, drivers expected ranges to be much higher than their average needs, but these range expectations decreased throughout the trial as EV experience increased.

Turrentine et al. [TGLW11] report results from the US portion of the same "MINI-E trial". 235 drivers participated (using the same BEVs as the prior study) for one year, and 54 of these drivers were interviewed further for the report. They used travel diaries and interviews to collect information. Drivers in this study had to lease the (reasonably expensive) vehicles to participate; most of the drivers had higher than average incomes.

Graham-Rowe et al. [GRGA<sup>+</sup>12] present results from a one-week UK PHEV & BEV trial. 40 drivers participated, 20 of whom drove a BEV and 20 of whom drove a PHEV. The BEVs were Mitsubishi iMEVs and Citroen C1s that had ranges of 120km and 144km respectively. The PHEVs were Toyota Pruis PHEV48s<sup>1</sup>. Participants were interviewed at the conclusion of the study. The researchers were primarily interested in participants' views on EV cost, their confidence in the technology, whether they felt EVs were still in their infancy, and whether they were willing to adapt their behavior to purchase an EV.

Cenex [Car10] held a six month BEV field trial. They gave four BEVs with an average range of 72km to 10 different fleet organizations in the UK. The BEVs were driven for one to four weeks by each organization, with a total of 195 participating drivers. The vehicles were equipped with GPS data loggers that recorded extensive charging and location information. Drivers were questioned prior to and after their test trial. This study was uniquely focused on whether company fleets present an early adoption market for BEVs. Kurani [Kur09] interviewed 34 drivers following a four week PHEV trial as part of the same program.

The trials suggest that perceptions towards EVs generally improve with EV experience. Here we summarize the most important conclusions from these trials:

- Drivers enjoy the advantages EVs have over ICEVs, e.g., increased acceleration, lower carbon emissions, lower noise levels, and the need to brake less often<sup>2</sup>. The performance and handling of EVs are perceived as advantages during post-trial interviews.
- Drivers find the charging process easy but express concern over plug design and standardization. Drivers want the ability to charge their EV at all EVSE stations without concerns of plug compatibility.
- Drivers overestimate the amount they will use public charging infrastructure (PEVSE). During pre-trial interviews, many drivers stated they would not consider buying an EV until PEVSE is widely available, but during post trial interviews, many drivers indicated

---

<sup>1</sup>The first Toyota Prius was an HEV, as discussed in §2, but a PHEV version was released later.

<sup>2</sup>EVs with *regenerative braking*, a mechanism that stores some of the energy usually lost in the braking process back into the EV battery, do not require much use of brakes at city driving speeds.

that they did not need to use PEVSE and concluded the relative lack of PEVSE (compared to ICEV gas stations) is a “perceptual problem”.

- Drivers expressed in post-trial interviews that they prefer the recharging process over the ICEV refueling process. Drivers like the ease of overnight charging.
- The time it takes to charge an EV battery did not impact the mobility of most drivers.
- After the trials, many drivers expressed a willingness to pay a slightly larger upfront premium ( $\approx$ \$2,000 more than a comparable ICEV) for the EV because the driving experience exceeded their expectations. This is a large indication that perceptions towards EVs improves with experience and also indicates that inexperienced drivers may have misconceptions.
- There is not a clear consensus as to whether range anxiety (RA) increases or decreases as drivers experience driving BEVs. For some drivers, fears subside after using BEVs for their normal routines without problems, but for others, seeing the remaining charge level decrease induces anxiety even when they reach their destinations without problems. However, it is clear that drivers have less range anxiety after actually experiencing EV use. In pre-trial interviews, drivers’ vastly overestimate their range needs, but then later acknowledge this in post-trial interviews. Post-trial, many of the participants state they would consider buying a BEV160 as a secondary vehicle and a BEV240 as a primary vehicle. These range expectations are far lower than those reported in surveys of ICEV drivers with no EV experience (as discussed in §3.1.2).

## Electric Bike Field Trials

Gehlert et al. [GKS<sup>+</sup>12] present preliminary results from the 2012/2013 German Pedelec Naturalistic Cycling study. There were 90 participants in the four week measurement period; 30 cyclists and 60 eBike users<sup>3</sup>. All bicycles and eBikes were equipped with a sensor kit that measured participants’ GPS locations (using the SM Modellbau GPS-Logger), traffic/road/weather conditions using a continuous video feed (using the ACME FlyCamOne eco V2 camera), and their horizontal movements with a wheel sensor, a speed sensor, and an acceleration sensor (using the SM Modellbau Unilog2). Questionnaires assessing their experiences were given to participants before and after the observation period. For one of the four weeks, participants also recorded their activities in a travel diary to correlate with data collected. At this time, an English-language post-trial report has not been published.

Dozza et al. [DWM13] present preliminary results from a Swedish eBike field trial called e-BikeSAFE. At the time of the report, 20 participants had driven monitored eBikes for two weeks each. All participants kept a trip diary, took two questionnaires (one before and one after data collection), and underwent a post-trial interview. The purpose of the trial

---

<sup>3</sup>In Germany, eBikes with a motor larger than 250W are classified as “two wheeled vehicles” and require a license and insurance to operate. Out of the 60 “eBikes”, 50 were “pedelecs”, eBikes with a 250W motor (most similar to the eBikes discussed in the other studies), and 10 were of the larger variety similar to North American “mopeds”.

was to study when the eBike were in dangerous situations, and was based on a prior study called BikeSAFE which studied the same for bicyclists. The authors survey of potential participants covers 1) demographic data 2) their cycling behavior and usage of their bicycles 3) their opinion towards eBikes, and 4) whether they already have experience with eBikes. Their sensor kit is detailed in [DF14a] and summarized in §8.7. Interestingly, their sensor kit includes a push button mounted on the handlebar to record user-triggered alerts. This button allows researchers to see which user-defined safety alerts can be detected using the raw sensor data, and when combined with GPS, to geographically plot "danger zones". During the trial, participants were asked to press it during all dangerous situations (e.g, near accidents, closer than comfortable proximity to cars, etc). At the time of the preliminary report, the participants had taken 332 trips covering 1549 km over 114 hours. 63 dangerous situations were reported by the participants including six where the cyclist crashed, corresponding to a critical situation on 18% of trips on average. At this time, a follow up report has not yet been published.

Paefgen et al. [PM10] show eBike usage metrics, e.g., average and maximum velocity, trip lengths and distributions, and trip routes, from a small trial. Location and acceleration data is obtained for four months from 17 eBikes rigged with the Picotrack GPS kit, an accelerometer, and an internal Li-Ion battery for power. Data is presented for only two of the 17 eBikes, so the paper is quite limited in scope. It is worth noting, because we use the same computation, that the authors compute the distance between two successive GPS coordinates using the Haversine Distance formula<sup>4</sup>, a trigonometric formula (based on the Law of Haversines) for calculating the shortest distance between two coordinates on an ellipsoid.

### 3.1.2 Surveys

#### Surveys on EV Perceptions

Researchers and manufacturers have also interviewed drivers about their perceptions towards EVs. Most of these surveys are of drivers with little or no experience with EVs, which gives a different perspective than those interviewed during field trials. A benefit of surveys is that thousands of drivers can be interviewed at little or no cost using online tools. Below we describe the surveys that have taken place and any unique insights; insights common to the surveys are summarized afterwards.

Lebeau et al. [LML<sup>+</sup>13] interviewed 1196 Belgian respondents on perceptions towards BEVs. The authors claim that BEVs are ideal for Belgian mobility based on various travel statistics (short trips being common, flat roads, etc). The respondents rated a list (given by the surveyors) of several potential advantages and disadvantages on a numeric scale, thus the

---

<sup>4</sup>[http://www.codecadex.com/wiki/Calculate\\_distance\\_between\\_two\\_points\\_on\\_a\\_globe](http://www.codecadex.com/wiki/Calculate_distance_between_two_points_on_a_globe)



results may not be comparable to other open-ended surveys discussed in this section<sup>5</sup>.

Tal and Nicholas [TN13] compare data from two data sources in California, which at the time of writing, represented over 42% of the US EV market. First, data was obtained from drivers who applied for the California tax rebates given to EV buyers based on the battery in their vehicle. Second, the authors used the 2012 CALTRANS travel survey, which was given by various government partners to 9600 households who bought a new 2012 or later vehicle. The two sources thus contain data from EV and non EV new car buyers. This allows the authors to compare socio-economic variables such as age, income, location, commute distance, and education, between PHEV, BEV, hybrid, and ICEV owners.

Krause et al. [KCLG13] surveyed 2302 drivers in 21 large US cities and conclude that drivers are uneducated about the technology. More than 66% of respondents failed to answer basic questions about EVs and 75% undervalued or were unaware of EV benefits. For example, the majority of respondents believed EVs cost between 10 and 50% *more to maintain* than ICEVs<sup>6</sup> and more than 70% underestimate fuel cost savings. Very few (5%) knew about local incentives to purchase EVs in their region. Only 30% of respondents stated they could identify a PEV if they saw one.

Egbue and Long [EL12] surveyed drivers who are more likely to be early adopters than the general population about barriers to adoption. Instead of randomly surveying the population, data was collected from students and faculty (n=481) in a technological university that specializes in science programs. Roughly half of the respondents indicated they had some experience with alternative fuel vehicles, which is far greater than the average population. Thus, while this may not provide an accurate depiction of EVs from the general public, it gives a different perspective from those more familiar with the technology. Interest in EVs was found to be much higher than reported in other surveys (likely due to population bias), with over 80% showing some interest in EVs. The authors provide the full responses of the survey for further analysis.

Deloitte [Del11] surveyed 13,000 respondents in 17 countries about BEVs (but not PHEVs). Respondents were surveyed as to their intent to purchase BEVs, as well as perceived BEV selling points and adoption barriers. The survey took place from 2010-2011. The majority of respondents were willing to "consider purchase", but many wanted vehicle specifications that exceeded what was available in 2010/2011. To quote, "*when consumers' actual expectations for range, charge time, and purchase price (in every country around the world included in this study) are compared to the actual market offerings available today, no more than 2 to 4 percent of the population in any country would have their expectations met today based on a data analysis of all 13,000 individual responses to the survey.*" Consumers in all regions stated their interest in EVs would increase with higher gas/petrol prices, e.g., at a U.S.

---

<sup>5</sup> Asking respondents to rate a list of disadvantages is not equivalent to asking what the disadvantages are. In the former case, the respondents may be presented with items they would not have thought of in an open ended survey. Moreover, they may have perceptions not on the given list that would have been revealed in an open-ended survey.

<sup>6</sup>One of the major selling points of EVs is that they have hundreds less moving parts requiring lubricants and fluids (no engine) leading to less mechanical failures and maintenance.

fuel price of \$5/g (a 37% increase from prices at the time of writing) the percentage of respondents interested in EVs rose to 78%.

Musti and Kockelman [MK11] used a 2009 web based survey of Austin, Texas residents to simulate the adoption of EVs in Austin over a 25 year period. The survey asked residents to choose between 12 vehicles in different scenarios. The scenarios were based on different vehicle specifications, price assumptions, and monetized values for the adverse affects caused by ICEVs such as global warming and health impacts due to carbon emissions. The authors used the survey results to estimate likely adoption in the coming years in the given scenarios, and to determine the most significant factors to EV adoption. Using the computed adoption rates and the the 2001 National Household Travel Survey, they also estimate the percentage of vehicle-km displaced as a result under the different scenarios.

Peters et al. [PD14] interviewed 969 drivers in Germany<sup>7</sup>. The authors tried to place respondents into four distinct customer groups: current owners, and those with concrete intent to purchase, some interest in EVs, and no interest in EVs. The authors followed the *diffusion of innovations* model proposed by Rogers [Rog03] to determine what variables (e.g., demographic variables) the groups can be parameterized by.

The City Of New York [Cit10] interviewed 1,400 drivers throughout NYC as to their EV adoption plans. The goal was to determine what percentage of residents were interested in adopting EVs and where they were located, both to see whether distribution and transmission infrastructure needed to be updated, and to place well-utilized PEVSE. This study gives a unique perspective on EV adoption because driving patterns and car ownership in NYC differs vastly from the rest of the US. The NYC transit system carries more passengers than the five next largest transit systems in the U.S. combined, so only 44% of NYC households own a car compared to 90% nationally as a result.

The EPRI (Electric Power Research Institute) [EPR10] surveyed 900 Southern California residents as to their concerns over EV adoption. Interestingly, they survey many HEV owners in addition to ICEV owners and are able to compare and contrast the responses from these two classes of drivers. HEV owners seem much more apt to purchase a PHEV/BEV than ICEV owners, which shows customer experience and education about the technology is essential to long term adoption.

Deloitte [Del10] interviewed major automotive executives, clean technology startups, auto dealers, energy companies, and 2,000 vehicle owners. The study focuses on drivers' knowledge of EV technology, how important brand name is to drivers, RA, charging opportunities and infrastructure, and drivers' perception towards the cost of EV purchase and ownership.

Ernst & Young [Ern10] interviewed 4,000 drivers in the US, China, Japan, and Europe about their knowledge of EVs and their perceived adoption barriers. The study compares drivers' responses across different regions of the world and focuses on barriers such as EV cost, EV range, lack of public charging infrastructure, and drivers' misconceptions about EVs.

Here we summarize the most important conclusions from these surveys:

---

<sup>7</sup>The survey data for this publication was collected in 2010, so vehicle models and perceptions are not as of 2014 (year of publication)



- The studies overwhelmingly conclude that range anxiety is the largest barrier to EV adoption. Surveys indicate a large gap between range requirements and expectations, often referred to as the “range paradox”. In Lebeau et al. [LML+13], only 35% of respondents stated they traveled more than 40km per day, but 90% of respondents said a range lower than 200km was unacceptable and 68% said a range lower than 300km unacceptable. In Egbue and Long [EL12], 87% of respondents indicated they travel fewer than 64km per day on average, but only 32% of respondents were interested in BEVs with a battery range between 0—160km, while 23% preferred ranges between 160km—320km and the remaining 45% preferred more than 320km of range. The average range desired was 345km—four times the average daily distance driven by participants. Deloitte [Del11] finds that 80% of respondents drive less than 80 kilometers per day but 50% of respondents demand more than 160km of range and 30% demand more than 400km.
- The second most cited barrier was the initial cost of EVs. Most drivers demand that EVs be cost comparable with ICEVs before they consider purchase. The majority of respondents in each survey stated they would not be willing to pay an upfront premium for EVs relative to an ICEV, regardless of whether the lifetime TCO is lower.
- Early EV adopters are well educated and have higher than average incomes. The common hypothesis for this phenomenon is that because many drivers cannot afford the high initial price of EVs (relative to ICEVs), adopters tend to have high paying jobs, which tends to be correlated with holding high degrees.
- The top three perceived advantages to converting from ICEVs to EVs are: 1) the low fuel and maintenance cost per km, 2) the ability to charge at home and make few “fueling stops”, and 3) supporting sustainability.
- Drivers tend to assume EVs “should” operate similar to the ICEVs and are discouraged upon learning the differences, e.g., longer recharging vs. near-instant refueling. This effect is related to prior work showing people are not good at evaluating new products that are “psychologically distant” from those they are accustomed [PD14, SG11, GRGA+12].
- Drivers are currently unaware of many attributes of EVs including range levels, recharging times, noise levels, safety ratings, and performance characteristics and are unaware of the differences between EVs and ICEVs. This can lead to a feedback loop in which consumer interest further decreases—misconceptions lead to lower interest in EVs, which in turn affects their answers in marketing and opinion surveys on intent to purchase [KCLG13]. If manufacturers and governments believe consumer interest in EVs is low, they may be less likely to invest capital in the space. The lack of investments can lead to stunted technology maturity, which further reduces consumer interest etc. Simply put, few investments are made when consumer interest is low, but consumer interest may be artificially low due to misperceptions.

## Surveys on eBike Perceptions

An et al. [ACX+13] surveyed 470 respondents about their eBike usage in Shanghai. Shanghai is one of the most eBike-incentivized cities due to the inner city ban of motorcycles and the

very high cost of car ownership due to government regulation [ACX+13]. The authors find most eBikes in Shanghai are used for commuting (42.7% of trips) and shopping (36.5%). A third of all eBike commute trips are longer than 10km, and the average eBike commuter trip time is 27.3 minutes, suggesting that eBikes are not only used for short trips. Among the two thirds of respondents with shorter (<10km) commutes, many use their eBikes as an alternative to bus transit. Similar to [Eng12] but in contrast to [WHS14], 56% of the respondents transitioned to eBikes from cycling and 33% transitioned from bussing; only 2% of respondents transitioned to eBikes from owning a private car. The top three reasons respondents purchased an eBike were 1) punctuality (arriving to destination on time despite traffic congestion), 2) reduced commute time (v.s. cycling and public transit), and 3) economic reasons (cheaper than car ownership).

Hiselius and Svenssona [WHS14] surveyed 321 eBike users in Sweden whose emails were obtained from a Swedish eBike retailer. Nearly all respondents owned or had access to a car, and the majority of users transitioned from using cars to e-bikes. This contrasts other studies where most eBike users were transitioning a portion of their cycling trips [ACX+13, Eng12]. The top three most cited reasons for purchasing an eBike were 1) increased bicycle mobility in the rain or wind (cited by 58% of respondents), 2) eBikes are more environmentally friendly than cars (again 58%), and 3) to reduce their commute time v.s. their bicycle (42%).

Popovich et al. [PGS+14] interviewed 27 Sacramento eBike users about their experiences with and sentiments towards their eBikes. The interviews were held as open-ended conversations, in contrast to other works where people fill out a survey or questionnaire. Interestingly the authors find that the eBike ownership households have high incomes, a correlation also true for EV owners as discussed. Most respondents purchased eBikes 1) because someone in their social circle recommended buying an eBike (this was also the top reason also reported by Engelmoer [Eng12]), 2) due to their greater speed and acceleration v.s. bicycles, and 3) because they require less exertion than bicycles. Many users, in contrast to Engelmoer [Eng12] but in line with Hiselius [WHS14], replaced a portion of their driving, with some users getting rid of their car altogether. The most common negative sentiments reported by users were 1) theft concerns, 2) safety concerns, 3) unwieldiness<sup>8</sup>, and 4) range anxiety. Regarding range anxiety, though eBikes can be pedaled upon battery depletion, the extra weight of the eBikes v.s. bicycle makes pedaling more difficult. The authors conclude that eBikes allow those with some interest in cycling but are unable/unwilling to commute longer distances or in hilly regions to transition some of their car usage to eBike usage.

Engelmoer [Eng12] surveyed eBike users and regional mobility experts in the Netherlands, primarily to see whether eBikes can reduce the environmental impact of commuter traffic in Dutch cities. The author states the notion that "eBikes are for the elderly" is still somewhat prevalent but changing rapidly in Dutch cities—eBikes now appeal to a wider population as a cheap, clean and flexible alternative to car ownership. However, the author finds that most Dutch eBike users are offsetting cycling trips, instead of car trips. This is unsurprising, however, because 27% of all mobility in the Netherlands is completed by bicycle, the highest

---

<sup>8</sup>Due to the increased weight of eBikes, users could not transport their eBikes easily, e.g., on stairs.

percentage in the world, and nearly 50% of trips under 8km are made by bicycle. The high rate of cycling presents both an incentive and disincentive for eBike use in the Netherlands: eBikes may attract those who are no longer able to cycle or those with longer commutes, but may repel those who are capable and prefer bicycles [Eng12, WHS14]. The author finds that the top three factors with a statistical relation to eBike ownership are 1) having friends or family that also own an eBike (41%), 2) older age (16%), and 3) possessing longer than average commute distances (16%). Several hypothetical eBike penetration scenarios are simulated to determine the impacts on the environment assuming eBikes become widely adopted. The simulation input parameters include: the percentage of trips taken by eBike/car/bicycle, energy use per vehicle, emissions per vehicle, emissions of local power generation, assumed average commute distance, and the number of people that commute daily. The authors find that there would (in theory) be a reduction in total energy use and most emissions with higher eBike penetration levels, but certain types of emissions (e.g., SO<sub>2</sub>) would increase, as those types of emissions are a byproduct of electricity generation but not petrol consumption.

Cherry et al. [CWJ11, LCY+13] discuss the operation of an eBike sharing system, the cycleUshare pilot at the University of Tennessee, but much of the report focuses on surveys of the eBike share participants. The surveyed participants' sentiments match the other surveys discussed here, e.g., most (77%) respondents either agreed or strongly agreed that e-bikes are more attractive than regular bicycles because they are better for longer trips and hilly areas ("terrain barriers"). Interestingly, 64% of respondents disagreed that regular bicycles are more attractive because battery range is not an issue—most saw *any* electric assistance from the eBikes as a benefit over bicycles, as long as the eBike can be pedaled normally after battery depletion.

## 3.2 Range Anxiety

Range anxiety (RA) is the most cited barrier to EV adoption. The surveys and trials in §3.1, as well as the GPS studies in §3.4, indicate that although current BEV160s range suffice for nearly all of drivers' mobility needs, drivers still prefer higher range. This has been cited in the literature as the "range paradox" or "range discrepancy". Strong examples include Ernst & Young's interviews [Ern10] of 4,000 drivers where only 40% believed a BEV160 would suffice but only 2% said they often drove more than 160km, and the survey by Egbue et al. [EL12] where respondents preferred a range of four times their average vehicle use. This suggests that drivers desire a "range buffer", battery capacity beyond what is normally required. RA can unnecessarily lower BEV utility—Neubauer and Wood [NW14] discuss that drivers may opt to use another vehicle even if they believe the trip might induce RA, some of which could be accommodated by the BEV. To complement this section, Nilsson [Nil11] gives an excellent overview of RA, its implications, and potential solutions.

So far, manufacturers and researchers have tried to reduce RA using one of three methods.

The first is the approach taken by Tesla: build BEVs with much larger batteries. See §2.1.1 for the range comparisons.

Second, Many players, including manufacturers, governments, and researchers, propose that deploying public charging stations (PEVSE) is a solution to alleviating RA. Drivers report they would feel more comfortable driving EVs knowing they can charge away from home if needed [BCMw11, Ele09, Del10, EPR10, MSA11, GRGA+12, Nil11, Adv11, WU13, Bak11, NW14]. However, data from PEVSE networks suggests that drivers infrequently charge in public:

- The EV Project [U.S13a] (2010-2013) was the largest (to date) deployment of EVSE, having deployed over 9,000 charging stations in major US cities, including 6,000 in participants' residences (for free in exchange for drive cycle and charging data). Through December 2012, the project logged over 96 million km and 1.6 million charging events from 7,376 participating Leaf and Volt owners [ECO13]. Their results report only  $\approx 15\%$  of charging occurs outside the home.
- ChargePoint America [Cha14b] (2011-present) is a program sponsored by Coulomb Technologies to provide EVSE to nine major-population regions in the United States. This program has collected charging data from their 4,217 stations. They reported in March 2013 that only 34% of charging occurs outside the home [Cha13], however, they have deployed more public stations than residential stations [Cha14b].
- The Cabled UK project [ARU12] (2009-2012) measured 340 vehicles for a year each, and reported that 15% of charging occurred away from the drivers' homes [Lov11]. Assuming these studies are a reliable indicator of actual charging demand, we conclude public charging accounts for only 15-20% of charging demand.

Regardless of whether PEVSE has thus-far been well utilized, we do expect PEVSE to play a role in alleviating range anxiety in the long-term. First, drivers have stated in numerous surveys that having the *option* to charge away from home makes them feel safer, so range anxiety should decrease as PEVSE becomes more ubiquitous. This is expected to happen, because many key players have incentives to deploy PEVSE as demand grows. For example, businesses with centralized parking such as shopping centers and restaurants can use on-site PEVSE to encourage EV owners to spend extra time on premises. Moreover, as of January 2015, at least one EV manufacturer, Tesla, is deploying worldwide fast charging stations [Tes15]. These L3 stations allow drivers to travel 170 miles (272km) after just one hour of charging. If stations with such high charging speeds continue to be deployed, range anxiety should largely decrease.

In this thesis, we do not explore PEVSE deployment because there is a considerable body of work on this subject already; see Liselotte [Lis12], Wirgesa et al. [WLK12], Neubauer and Wood [NW14], Bakker [Bak11], Efthymiou et al. [EATM12], and The Schatz Energy Research Center [Sch13] for various PEVSE deployment models.

It may be many years before PEVSE is widely deployed. In addition, PEVSE may still not serve drivers that need to take very long trips, and surveys reveal that drivers do consider this "worst-case usage" when making purchasing decisions. For these reasons, BEV owners

may still desire occasional access to an ICEV. If BEV owners can cheaply and easily acquire a vehicle on days their mobility demands exceed their BEV range, they may have less concerns about range limitations and range anxiety. To this end, in chapters §5 and §6, we design algorithms for sizing and managing a carshare for EV owners—a network of ICEVs pools to be used by BEV owners when they need to make long trips. This is referred to as a *vehicle access model*. There is one related work in this space, King et al. [KGWS13], which was published independently and simultaneously with our first work. Our model is more comprehensive than their work in several ways, but we defer this discussion to §6.8 after we introduce necessary queueing terminology and discuss papers focusing on the sizing of similar systems.

We end this section by summarizing a survey on measuring drivers' *willingness to pay* (WTP) for more range. Drivers' WTP for additional range in a BEV is a tradeoff between the two largest adoption barriers, RA and upfront costs. Dimitropoulos et al. [DRvO13] gives a meta-analysis of 33 prior studies investigating consumer preferences for EVs and other alternative-fuel vehicles. The authors also summarize a related metric, *compensating variation*, which describes the monetary compensation required to restore one's utility after experiencing a change (here, the change from "unlimited" ICEV range to limited BEV range). The compensating variation for a change in driving range from  $R$  to  $R'$  is the integral of WTP over  $[R, R']$ . The authors use several statistical techniques to resolve statistical biases between the 33 studies. The authors find the utility of driving range is nonlinear; the WTP for additional range is higher when the base range is lower, and decreases as the base range increases e.g., drivers will pay more to extend a BEV80's range than a BEV300's range. Of the 33 studies analyzed, the mean WTP with a base of 160km is \$41/km<sup>9</sup> and the median is \$35, suggesting the existence of high WTP values. The median compensating variation from 160km to 560km (the latter is used as the "range for an ICEV") was \$13100, suggesting that a BEV160 would need to be \$13,000 *cheaper* than an ICEV because their "range" changes from 560km to 160km. The author notes the actual value is probably lower because changing from an ICEV to a BEV leads to fuel and maintenance savings, and this does not include tax benefits such as the US credit of \$7,500 for BEVs. Finally, the WTP will vary among different individuals based on their income, mobility demands, and whether they have access to an ICEV.

### 3.3 Calculating the Return On Investment

The higher initial cost of EVs compared to ICEVs is the second most cited adoption barrier. While drivers perceive EVs as having lower operating and fueling costs than ICEVs, most do not or cannot calculate their expected return on investment (ROI) when making purchasing decisions [BHB14, EPR10, GRGA<sup>+</sup>12, GM12, Pru10, Del10, Nat10, BST09, Aec09, MSA11, BMW11]. To quote Morton et al. [MSA11]: *"the majority of consumers do not*

---

<sup>9</sup>these figures are converted from dollars per mile to dollars per kilometer to be unit-consistent with the rest of this thesis.

*have even the fundamental building blocks to be able to make detailed payback calculations. Moreover, many drivers did not even know what electricity costs and had little idea of how much electricity was required to propel an EV (it is widely concluded however that the price of gasoline will have significant influence on EV adoption in the coming years). Drivers believe the overall operating cost of an EV is 'probably lower' than an ICEV, but because they cannot easily calculate to what extent this is the case, they express deep concern about the high initial cost of EVs".*

Given this, drivers need tools to compute their ROI and payback period when purchasing an EV. However, building such a tool is difficult because the ROI and payback period for EVs is heavily dependent upon the drivers' mobility patterns. As stated earlier, EVs will be more profitable for drivers whose mobility consists of many short trips than drivers who make longer trips. In the former case, more distance can be traveled under a PHEV's AER or a BEV's range, thus recouping more of the initial cost through fuel savings<sup>10</sup>. Furthermore, for BEVs, if a driver regularly makes trips exceeding the BEV range, they may need to pay for access to another vehicle, decreasing the ROI of the BEV. External factors such as climate (air conditioning and heating can heavily affect range) and regional incentives also heavily affect EV utility. Consequently, the ROI for two different drivers can differ significantly even if they have the same average vehicle demand (e.g., average kilometers driven per week/year).

The following papers present ROI models for drivers' and fleets transitioning to EVs. While we study the ROI for fleets transitioning to EVs in §7 and do not explore the ROI for individuals in this thesis, we present work in both spaces both for completeness and to illustrate many of the complexities in building ROI models. Due to changing vehicle and battery prices, we present the most recent papers first.

The EPRI [EPR13] gives an extensive analysis of the ROI and payback period for PHEVs vs. hybrids and conventional ICEVs. The report stresses that profitability models must incorporate a large number of parameters to tailor the results to drivers with different mobility patterns, charging availability, and budgets. The two vehicles considered are the Volt (PHEV64) and the Leaf (BEV160). The model assumes that only at-home charging is available to all drivers in the baseline scenario, but other charging opportunities such as workplace charging are discussed in a sensitivity analysis. The analysis uses the Puget Sound dataset (also used by Neubauer and Wood [NW14], Neubauer et al. [NBW12], and Wu et al. [WDL14]) to evaluate different drive cycles. This dataset contains GPS measurements of 759,000 trips from 415 vehicles measured over 3 to 18 months from 2004-2006, and is available to researchers [Pug08]. However, the authors note this dataset under-represents high and low mileage drivers, thus using it to compute PHEV payback periods may present an inaccurate picture. Trips in the dataset that cannot be completed by the Leaf are assumed

---

<sup>10</sup>This barrier to adoption will be harder to overcome in the U.S. The average gasoline tax in the US is 47cents/gallon, whereas in the UK, for example, the tax is equivalent to \$3.28 per gallon—almost 20 times larger [Ele09]. The total price of gas is consequently much lower in the U.S. than in much of the world, so less of the initial cost is recovered through fuel savings; EVs are much more cost competitive in countries with higher gasoline prices and taxes.



to be "replacement miles", that is, the driver must pay for access to another vehicle on those days (though this cost is assumed to be \$0 in the following results, it is perturbed in a sensitivity analysis). A problem with this paper is that although the authors discuss several input parameters such as maintenance costs, the actual mathematical model is not presented—only results are given<sup>11</sup>. The authors' results are as follows. First, with 2013 incentives and prices, both the Volt and the Leaf are cost competitive for most drivers over their lifetime, though the payback may happen only near the lifetime of the vehicle. The Volt is cost competitive:  $\approx$ \$700 cheaper over the lifetime vs. an equivalent ICEV. The variation in ROI across drivers is significant but most of the variance is associated with positive ROIs—the risk of significantly high negative return (relative to an ICEV) is low. The authors find the Leaf is much more cost competitive than the Volt<sup>12</sup>. However, they also find that while the Leaf is thousands of dollars cheaper over the vehicle lifetime vs. an ICEV, the ROI has a wide variation for different drivers. This suggests BEVs require more careful consideration by drivers making purchase decisions. A sensitivity analysis suggests that changes in gasoline prices will have a significant impact on the ROI and payback period, as expected, but incentives and rebates have an even larger impact. The authors conclude by discussing how the results change if the EV is financed instead of purchased upfront and if model parameters are altered.

Wu et al. [WDL14] measure whether various PHEVs are profitable for drivers by using the Puget Sound dataset [Pug08]. The paper is similar to our (earlier) work on taxi profitability; the authors emulate PHEVs using the ICEVs in the dataset and use similar modeling equations. The authors estimate the operating costs of various PHEV AERs under different PEVSE availability and gasoline price scenarios. The authors do not consider BEVs or different electricity prices (which is fixed at 12c/kWh). The authors conclude that the incremental battery cost of larger AER PHEVs is difficult to justify without government subsidies or gasoline prices higher than \$4/gallon (which it is near at the time of this writing). If gasoline costs less than \$4/gallon, even with a battery cost of \$200/kWh (well below the current price), only PHEV16s are profitable. Below \$3/gallon, HEVs or ICEVs are best. As expected, however, the authors find the price of gas significantly affects the results, e.g., when gas increases from \$4/gallon to \$5/gallon, the number of PHEV users who benefit from a larger battery increases significantly under all charging scenarios.

Al-Alawi et al. [AAB13] give a comprehensive ROI and payback model for PHEVs. The authors examine four PHEV classes: compact and mid sized passenger cars (LDVs), and mid to large light trucks. The price of all EV parts not found in ICEVs excluding the battery are derived from EPRI estimates, and the authors assume a battery price of  $\approx$ \$300/kWh, lower than most current estimates. The authors assume that the distance driven for LDVs and LTVs is fixed at 19,300km and 24,000km in the first year and declines at  $\approx$ 3% per year afterwards. The fixed first year usage is derived from 2005 U.S. national averages. The authors further assume that the percentage of kms driven on electricity is a fixed

---

<sup>11</sup>In contrast, in Chapter §7 we give our full mathematical profitability model.

<sup>12</sup>At the time of the authors' writing, the Leaf was significantly cheaper than the Volt at \$29,022 vs. \$35,200 [EPR13]. The same difference exists today, as discussed in §2.1.1, but both are now priced lower.

function of their total VKT. That is, they do not consider different drive cycles where the annual VKT is composed of many long trips, some long trips, or many short trips, etc, which could significantly change their results. The authors further estimate values for fuel prices, electricity prices, resale values, vehicle life, maintenance costs, insurance costs, and registration costs from various sources. Their results suggest the payback period of PHEVs ranges from 14-20 years for compact LDVs, 7 to 10 years for mid-sized LDVs, 6 to 7 years for medium LTVs, and 3 to 5 years for large LTVs. This intuitively suggests that due to the relatively poor mileage of larger vehicles, the same sized battery will lead to higher fuel savings if installed in a SUV vs. a compact car; the savings are greater when "lower efficiency kms" are displaced with electricity. Thus the payback period is directly proportional to the mileage of the ICEV counterpart. The authors further note that including maintenance costs and the salvage value of vehicles in the payback analysis (most prior papers do not) significantly decreases the PHEV payback period.

Neubauer et al. [NBW12] use the National Renewable Energy Laboratory's Battery Ownership Model [OABM+10] to examine the sensitivity of BEV economics to drive patterns, vehicle specifications, charge strategies, and battery replacement rates using two different methods for accounting for trips that exceed the BEV range. The model outputs the cost ratio of an EV to a standard ICEV. The model requires many inputs, including drive cycle data, the EV and ICEV's vehicle specifications, financing structure (how the vehicles are purchased) and tax incentives. The authors assume trips that cannot be completed with the BEV are handled by another household vehicle (the "low cost scenario") or via a car rental or carshare (the "high cost scenario"). In the former scenario, the authors sum fuel and maintenance costs from having to use an ICEV with the EV cost of ownership, and in the latter, rental costs are also included. The authors apply a battery degradation model so that more trips are unable to be completed over time. IEA forecasts are used for electricity and gas prices. The authors simulate several parameters for the inputs such as vehicle range and charging availability/capabilities. The authors simulate drive cycles using the Puget Sound dataset [Pug08] (also used by the EPRI study [EPR13] above). Under their drive cycle data, a BEV160 requires using another vehicle through either of the two usage scenarios (household ICEV or rental),  $\approx 40$  days per year on average. Given this relatively large number, the authors find a significant difference in the ROI depending on whether the low or high cost scenario is employed. The authors find that BEVs are cost effective, based on the simulated inputs, for  $\approx 20\%$  of the drive cycles simulated, suggesting there exists a substantial market for BEVs. Due to the model's sensitivity to drive patterns, detailed knowledge of drivers' individual or household driving patterns is required to make cost-optimal decisions.

Garcia and Miguel present a systems dynamics model to predict the adoption of EVs in Spain through 2020 [GM12]. The model, which is designed to predict long term trends, models interesting dynamics of EV adoption, e.g., as EV adoption increases, carbon emissions decrease, giving the government an empirical reason to increase government subsidies for EVs, further encouraging EV adoption. The model can be used to identify feedback loops in which a variable influences itself over time, and thus predict how EV adoption may evolve as a function of these feedback loops. In the report, the authors focus on how the evolution of gasoline prices, electricity prices, battery prices, maintenance costs and



government subsidies will affect the cost of EV ownership—however only a small subset of the model is presented. The authors conclude that in the most pessimistic case, about 750k EVs may be adopted by 2020, which does not meet the Spanish government goals, but with the introduction of government subsidies and with reductions in future battery prices, government goals may be met. The authors also predict government goals will not be met even with reductions in battery and vehicle prices without government subsidies.

The National Research Council [Nat10] evaluated several pricing scenarios to predict the year PHEV16s and PHEV64s will be profitable for consumers. Their model is similar to the others described here—it is not tailored to different drivers, makes assumptions based on national averages, and considers only a few parameters such as gasoline and electricity prices. The authors conclude that because of expected price of lithium-ion batteries through 2030, it is unlikely for PHEVs to be cost comparable with ICEVs when gas is under \$1.03/liter (\$4.00/gal) (currently, gas prices are at this point but are highly volatile—they must stay above this point to remain profitable [Gas]).

Aecom [Aec09] forecasts when various types of EVs will be profitable for Australian drivers. The authors build a choice model which predicts when drivers will purchase EVs given vehicle prices, range, cost, and the availability of PEVSE. The model estimates the financial equivalent of each these factors to predict whether customers will purchase EVs, e.g., “an increase in range from 100km to 200km is worth  $\approx$ \$3,000 to drivers”. The model does not calculate the TCO of EVs based on mobility patterns or driver data. The model does consider a wide range of parameters including the cost of installing PEVSE given various EV penetration levels and the monetary benefit of reducing CO<sub>2</sub> emissions, but the derivations for many of the assumed values are not stated. The authors conclude all types of EVs will become profitable near 2030 due to falling vehicle prices and savings incurred by reductions in greenhouse gas emissions and air pollution. At the time the report was written (2009), the lifetime savings incurred by purchasing EVs compared to ICEVs was not enough to offset the higher initial cost of EVs, and the authors projected this would hold for many years. They predict a transition to HEVs in the next 5-10 years, PHEVs in the next 5-20 years, and BEVs in 20+ years. However, gasoline price predictions used in the report turned out to be incorrect. The report predicted the price of petroleum per barrel will be less than \$80USD through 2040, but it has passed \$100USD several times since the report [Blo12] was published. The authors give a sensitivity analysis and note deviations from their predicted gasoline prices largely affect their profitability analysis. As a result, EVs may be profitable in Australia sooner than predicted. Moreover, the authors state EVs will be more expensive (by roughly \$10k) in Australia compared to other nations due to import costs and a lack of local vehicle manufacturing, and caution the results may not generalize.

Becker et al. [BST09] study how the adoption of EVs through 2030 can be profitable for the overall US economy and individual drivers if batteries were leased instead of purchased. To model customer adoption, the authors amortize PEVSE and battery costs through driver *pay per mile* contracts envisioned to be run by operators of networked public charging facilities. This idea stems from the *network externality model* [LM94], with the idea being that as more drivers purchase BEVs, the relative cost per driver of public charging infrastructure and

battery switching facilities decreases while its utilization increases, and therefore, all costs relating to batteries and charging infrastructure should be amortized over many drivers. Batteries would be leased on a per distance basis and exchanged at switching facilities rather than paid for by individual drivers, reducing the initial cost of BEVs and removing drivers' concerns over limited battery life. Under this payment model, the authors predict that between 2012 and 2017, BEVs are likely to be cost comparable to ICEVs. Averages from the 2001 NHTS survey are used to predict the likely market segment for BEVs—the authors estimate BEVs will account for 64% of U.S. light-vehicle sales and 24% of the U.S. light-vehicle fleet by 2030. Following their predictions for the adoption rate of BEVs, they estimate trade deficit reductions due to decreased petroleum imports, economic growth via the creation of battery manufacturing jobs, and reductions in health care costs due to lower overall emissions from the transportation sector. The results for each of these macroscopic issues are positive. BEV adoption would lead to massive decreases in the trade deficit, the number of jobs created as a result of battery manufacturing would outnumber those lost in the ICEV service industries, and health care cost reductions would be significant. Unfortunately, while companies such as (the now defunct) Better Place support this battery leasing model [Bet12], currently most vehicle manufacturers do not. It is possible that this model presents a long term solution to the higher initial cost of EVs.

### 3.4 Using GPS to Infer Mobility Demands

Whether EV adoption is profitable for a particular driver is heavily dependent upon their mobility demands. For example, they should not adopt a BEV (as the only household vehicle) if most of their trips are longer than the BEV's range, because they'll often need to pay for access to another vehicle. The same holds for PHEV adoption; if they often travel under the PHEVs AER, they will recoup more of their initial investment through fuel savings, but if they often need to engage the ICE, a PHEV will be more expensive than comparable ICEVs. Thus measuring existing ICEV mobility patterns can estimate the degree of range anxiety, estimate how often BEV owners will need access to an additional vehicle, and help estimate whether EVs will be profitable for a given driver or fleet operator. Moreover, for BEVs specifically, the percentage of trips that cannot be completed using a BEV cannot be observed with a BEV; existing ICEVs must be measured for manufacturers to build BEVs with specifications suitable for the majority of existing travel. Measuring existing mobility also helps plan PEVSE deployments.

The seminal work in this space is that of Pearre et al. [PKGE11]. The authors use a 2004 mobility dataset to derive conclusions about required EV range. The dataset used is one of the most extensive available—484 Atlanta-based vehicles were measured for one year by GPS data loggers. Their goal was to answer whether BEVs are compatible with drivers' current travel patterns. More specifically, they serve to find  $x, y, z$  in the question: *given a BEV $x$  ( $x$  km of range),  $y\%$  of drivers would find BEVs satisfy their driving patterns on all but  $z$  days of the year.* The authors find BEV160s would suffice the needs of 9% of drivers with zero driving modifications, 17% of drivers on all but 2 days a year with no

mid-day charging, and 32% of drivers on all but six times a year. They also find greater than 80% of vehicles are parked at any given time, including weekday rush hours—drivers have ample time to charge their vehicles given the existence of infrastructure. These results show currently available BEVs could be adopted by a significant portion of the population with little modification given properly located PEVSE.

Kahn and Kockelman [KK12] give a similar analysis using a year of GPS data from 255 Seattle households. The data was collected during 2004-2006 and includes 269,357 trips. The authors conclude that mobility patterns in Seattle are EV friendly, at least relative to those in Atlanta studied by Pearre [PKGE11]. The results suggest that BEV160s would suffice for 50% of single vehicle households if the households have access to another vehicle (e.g., a rental car) four days per year. The same 160km BEV would cover 90% of households 95% of the year (all but 19 days). Alternatively, A PHEV64 would allow Seattle households, which travel 42km per day on average, to electrify 80% of their travel. If gas is 3.50\$/gallon and electricity is 11.2c/kWh, a PHEV such as the Volt would save the average household \$535 per year, thus at those price levels there is still a significant payback period due to the initial premium. A BEV160 like the Leaf would save the average household \$780 per year if they could arrange free transit on days their BEV would not suffice. See §3.3 for another work using the same dataset for inferring EV profitability (Wu et. al [WDL14]).

The Swedish government [Kar13] undertook a large GPS study to analyze whether EVs are suitable for the Swedish population. Participants were recruited randomly from the Swedish motor-vehicle register. From 2010-2012, over 700 cars installed GPS equipment that was sent by mail. The resulting database contains 714 cars with data, 528 cars with >30 days of data, and 450 with >50 days of data. Questionnaires were also given to the participants. The results show that current BEV160s would suffice for the majority of Swedish drivers. Of the 134,425 recorded trips (totaling 1.3M kms), only  $\approx 3\%$  of trips were over 160km and these 3% of trips accounted for 14% of total kms traveled. Moreover, only 25% of all stopping occasions were more than 6 hours but these stops accounted for more than 80% of all parking duration, meaning that at least 80% of the time any car was parked, it could have charged fully from empty on a L2 (or L3) charger.

Slater et al. [SDT+09] agree that high initial cost and range limitations remain the two most significant barriers to EV adoption in their UK based study. The authors used data from the UK Department for Transport to gather data about vehicle use in the UK. The patterns show UK mobility is dominated by many short trips, far less than the range of currently available BEVs. For example, two-thirds of all commuting (to work) trips are less than 16km, suggesting BEVs are well suitable for average daily usage. Moreover, a BEV with a range of even 80km (half of the range of the Leaf) is suitable for 50% of all UK vehicle-km assuming only at-home charging, with the remaining distance undertaken by a small number of high mileage individuals. While long trips beyond available BEV range are rare, drivers place high disutility on limited range while making purchasing decisions. Interestingly, the authors note most EV owners in the UK own multiple vehicles and are able to use their ICEVs for long trips to combat this limitation, which also shows early EV adopters are typically from higher income households. To increase the adoption rate of EVs, the authors

conclude fast EV charging (the ability to fully charge the vehicle within an hour) may alleviate many driver concerns. Interestingly, while city-wide PEVSE is attractive to drivers, from an analysis of mobility patterns, the authors conclude PEVSE may be underutilized with the proliferation of home and workplace charging. Finding the minimum density of public recharging infrastructure sufficient for customer perception, as more infrastructure would be wasteful and underutilized, is a difficult but interesting problem

Gonder et al. [GMST07] analyze a dataset of 227 vehicles monitored for one day each in 2007 by GPS data loggers in St Louis. They simulate whether different EV models can fulfill the drivers' patterns, and if so, the estimated fuel reductions if the drivers adopted those EVs. They find HEVs would reduce fuel consumption by about 29% relative to ICEVs, PHEV32s about 55%, and PHEV64s about 66%. The GPS traces indicate drivers accelerate and decelerate very aggressively (which reduces vehicle efficiency). They also show that HEVs and PHEVs (in 2007) were roughly 33% and 40% cheaper to fuel than ICEVs respectively. The fuel savings are likely greater now—since this study the US average cost of electricity has risen by 9% [U.S14b] but the cost of gasoline has increased 40% [Gas].

Plotz [Plo14] gives the statistical theory to compute the mean number of days per year that a users driving demand will exceed some range threshold (i.e., the number of days they will need access to another vehicle) based on their past driving behavior. This work is theoretical and intended for use by other researchers with data; the author does not conclude how often BEVs will not suffice for users (e.g., like Pearre [PKGE11]) but instead provides statistical tools to calculate such probabilities given data as input. The author also computes the standard error of this mean given the sample data. The methodology requires that distribution of daily vehicle kilometers traveled for the user is log-normally distributed. The author numerically validates their estimators using a case study of the "average" German driver, and shows that the standard error of the estimates decreases (as expected) with more observations.

We end by noting that many authors rely on statistics from national travel surveys such as the periodic US National Household Travel Survey [US 13]. In these surveys, households respond to questionnaires about their *average* travel habits. While these are useful for determining average mobility statistics, e.g., the average length of a US car trip, they do not give any finer-grained statistics of interest to EV researchers such as trip length distributions, distributions of parking durations, and temporal distributions of parking events. These can only be obtained using trip-level information, which the above papers attempt to measure.

## Chapter 4

# Mining and Classifying Electric Vehicle Owners' Opinions

**Publication reference:** *T. Carpenter, L. Golab, S. J. Syed. Is The Grass Greener? Mining Electric Vehicle Opinions. e-Energy 2014.*

---

### 4.1 Synopsis

As discussed in §3, while field trials and surveys have been the two primary methods to gauge the barriers to EV adoption, both methods have limitations. Because EVs are expensive, field trials are usually limited in duration and in the number of participants, and surveys usually receive a large number of respondents with little to no experience with EVs. We thus propose an additional method to elicit perceptions: *online sentiment analysis*. We build an open-source system that programmatically mines EV owners' sentiments from online forums [Mis14b, Mis14c, Mis14a]. These forums contain detailed product discussions authored by owners. Analyzing their perceptions is beneficial for at least three sets of users: prospective buyers (what features do owners like and dislike?), marketers (what well-reviewed features should be advertised?) and manufacturers (what features should be improved?).

Our contributions in this chapter are twofold:

1. We augment prior methods of gauging the barriers to EV adoption with a system that programmatically mines perceptions from EV ownership forums. Our system produces a high-level product overview with the ability to drill down into opinionative sentences about specific features of interest. The system not only mines perceptions found through expensive field trials, but also opinions that are only observable after longer periods of time, e.g., those towards battery degradation, that are not found during field trials.

2. We augment prior sentiment mining techniques with new methods for parsing sentence fragments containing opinions about features (§4.6) and handling context dependent opinions (§4.7.1).

We find our system significantly reduces the amount of text the user must read to determine owners' opinions. In our case evaluation, we found that 90% of the sentences in the forums did not contain any feature synonyms, and moreover, 70% of the sentences that did contain features were neutral (did not express a sentiment). Thus, given the feature set of interest, the system reduced the text by 97% to the 3% of relevant text.

We evaluate the system's performance using a corpus of 830,000 sentences, 8,000 of which we manually labelled for ground truth. The system extracts and classifies opinions with a precision and recall of  $\approx 60\%$ , which is on par or better than the previous opinion mining systems discussed in §4.9.1.

We have open-sourced our system [Car] unlike prior sentiment mining systems that were commercialized.

## 4.2 Terminology

In this chapter, we use the following notation:

- All usages of *dictionary* refer to the data structure mapping keys to values, *not* indexes of English word definitions. We use the domain terminology *lexicon* for the latter.
- $\mathcal{P} = \{P_1, P_2, \dots\}$  represents the *set of products* the user wishes to mine reviews for.
- *Feature* refers to a product feature or attribute.
- $\mathcal{F}^p = \{\vec{F}_1^p, \vec{F}_2^p, \dots\}$  represents the *feature space* of  $p \in \mathcal{P}$ , where  $\vec{F}_i^p$  is a vector of synonyms for feature  $i$ . For example, if feature 1 of  $p$  is *fuel economy*,  $\vec{F}_1^p$  may be  $\langle \text{fuel economy, efficiency, gas mileage, fuel efficiency, mpg..} \rangle$
- $\mathcal{O}$  represents the set of opinion phrases recognized. Until §4.7.2 this definition is sufficient; we formally define  $\mathcal{O}$  in Eq(4.1) after several other definitions are introduced.
- We refer to a sentence fragment in which opinion  $o$  refers to feature  $F$  as a  $(F, o)$  pair.
- We abbreviate "neutral" with  $N$ , positive with  $+$ , and negative with  $-$ .

## 4.3 System Architecture

Because customers usually desire certain specifications when shopping for expensive products like vehicles<sup>1</sup>, we build a *feature-based opinion mining* (FBOM) system that extracts and classifies opinionative statements about specific *product features* the user is interested in (e.g., "battery capacity" and "safety"). There are five main phases in FBOM:

---

<sup>1</sup>E.g., one buyer may seek performance while another may look for top safety ratings.

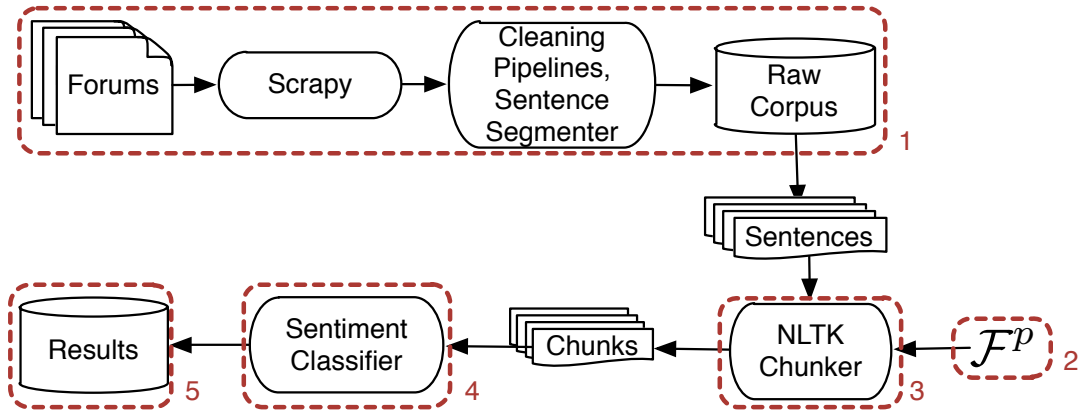


Figure 4.1: Sentiment analysis system architecture

1. Building a text corpus to be mined
2. Defining or mining the product and product features of interest
3. Extracting sentence fragments from the corpus containing opinions about those features; these fragments denoted as  $(F, o)$  pairs.
4. Classifying each  $(F, o)$  pair as  $\{+, N, -\}$
5. Aggregating results and computing various statistics

Our system is depicted in Figure 4.1. The red boxes labeled 1—5 correspond to the five aforementioned phases. First, forums are first crawled, cleaned, and split into individual sentences. These steps are collectively labeled as step 1 (see §4.4). The feature space is then defined or mined (§4.5). Sentences are then mined for  $(F, o)$  pairs using a process known as *chunking* (§4.6). Chunks, fragments containing  $(F, o)$  pairs, are then classified as having a positive, negative, or neutral sentiment (§4.7). Finally, results are output (§4.8).

We note that in the remainder of this chapter, we define and refer to several lists of words and dictionaries. Rather than reproduce these lengthy data structures here, some of which contain thousands of words, we refer the reader to the codebase [Car].

Prior work on building feature based opinion mining (FBOM) systems is discussed in §4.9.1.




Figure 4.2: Scrapy[Scr14] webcrawling process



## 4.4 Phase 1: Data Collection and Preprocessing

We use the Python Scrapy package [Scr14] to collect EV reviews from the Web. Scrapy is a system in which the user writes *spiders* containing two sets of regular expressions (regex). URLs matching any expression in the first set are *content pages*, and are sent to a *parsing pipeline*. URLs matching any expression in the second set are *linking pages* which hold links to other (content & linking) pages. The crawling process is depicted in Figure 4.2. The parsing pipeline processes the page text, fragments the text into sentences, and inserts the sentences into a database we denote as the *raw corpus*. For processing, we first remove all HTML tags and links. We then convert all characters to lowercase. Next, we iterate through a large list of common typos [Wik14a] and fix misspellings. Finally, we iterate through a list of contractions [Wik14b] and expand them. This is done because words such as *not* are *valence shifters* which change the sentiment of opinion phrases as discussed in §4.7.

## 4.5 Phase 2: Building The Feature Space

We build the feature space  $\mathcal{F}^p$  for each product  $p \in \mathcal{P}$  as follows:

1. We first manually create a *seed* set of features and define a few synonyms for each.
2. We then use *word frequencies* to generate more candidate features from the raw corpus. If a noun in the raw corpus has a high frequency, it may be a feature synonym. We manually filter these results, because not all common nouns are features; e.g., *road* is quite common in our EV review database.
3. We next use NLTK's *collocation functionality*, which produces bigrams and trigrams with high-scoring point-wise mutual information—sets of two and three words that often occur together. Multiple-noun features like *battery capacity* and non-adjective opinion phrases like *warranty issues* are found this way.
4. Finally, we use NLTK's *concordance* functionality. Concordance shows the words surrounding each usage of the target word, e.g., `concordance("battery", k)` prints each occurrence of *battery* with the  $k/2$  surrounding words on both sides. Manually reviewing concordance helps further identify multi-word features. This process can be repeated indefinitely as time permits.

## 4.6 Phase 3: Chunking

In the following two sections, we describe our sentence parsing methodology, *chunking*, which works by grouping part of speech (pos) tags with regular expressions (regex) [Fri02]. To the best of our knowledge, we are the first work to apply this natural language processing



technique, built for finding syntactic constituents in English sentences such as noun phrases<sup>2</sup>, to sentiment mining. We also describe its advantages over prior work.

### 4.6.1 Part of Speech Tagging

Chunking groups sequences of words based on their part of speech (pos), so the first phase of chunking is to label each word in each sentence with its *pos tag*. We first filter out a list of *stopwords*<sup>3</sup> from each sentence. We then tag each word with its pos using the Stanford Tagger [Sta14] for this task, which has an interface for NLTK [NLT14]. This produces a list of tuples of the form [(*word*<sub>1</sub>, < *pos* >), (*word*<sub>2</sub>, < *pos* >).] for each sentence.

We then iterate over each tagged sentence and replace the tags of select words with special tags. The usage of these five special tags is discussed in the following section.

- Sentences are searched for all synonyms for all features, and all matches' pos tags are replaced with <*feature*>.
- We replace the pos tags of *implicit features*, words that are both feature synonyms and opinion phrases, with <*implicitfeat*>. For example, *noisy* is both a synonym of *sound* (which we later grouped into Miscellaneous) and also a negative opinion.
- We search for words contained in a list of *non-adjective opinion phrases*<sup>4</sup> (NAOPs), and replace all matches' tags with <*naop*>.
- We replace the pos tags of *feature changers*, words that when used near one feature indicate the sentence is actually referring to another<sup>5</sup> with the <*featchange*> tag.
- We replace the pos tags of *valence shifters* [KI06] like *not* with the <*vs*> tag.

### 4.6.2 NLTK Chunking

English is not a regular language [CFL10], thus arbitrary English sentences cannot be parsed using regular expressions. Fortunately, most of the sentence constructs people use can. We recognize English text using a context free grammar, following common practice. We use the Python NLTK (*Natural Language Toolkit* [NLT13]) package to parse sentences using *chunking* [BKL09], which makes use of regex to group word sequences with particular parts of speech (pos) together.

A context free grammar (CFG) is a set of production rules of the form  $A \rightarrow B$ , where this denotes A can be replaced with B in any "string" in the language. In parsing natural language, CFGs state "replace instances of B with the higher-level notion of A". For instance, the rule

---

<sup>2</sup><http://www.nltk.org/api/nltk.chunk.html>

<sup>3</sup>These are the "extra" words in a sentence, which help with sentence structure and flow, but are not integral to sentence meaning, and are usually not indexed in information retrieval systems [Ras09].

<sup>4</sup>such as *disgrace* and *problem*

<sup>5</sup>As examples, if *costs* is used near the feature *car*, the sentence is likely referring to the *price* feature of the vehicle, and if the word *handles* is used near *car*, the sentence is probably referring to the *car's performance*.

$\langle\text{verb-phrase}\rangle \rightarrow \langle\text{subject}\rangle\langle\text{verb}\rangle$  replaces the tags  $\langle\text{subject}\rangle\langle\text{verb}\rangle$  with the higher level concept  $\langle\text{verb-phrase}\rangle$ . Chunking is simply an *extended CFG* (E-CFG), that is, a CFG in which the right hand side of production rules can be regex, instead of lower level concepts or terminal strings. While E-CFGs provide no functional benefit over traditional CFGs because they describe exactly the same set of languages [AGW00]—an infinite number of CFG production rules may be needed to express the same rule of an E-CFG [AGW00]. The regex operators  $\{+, *\}$  provide *compactness*—a way to specify an infinite number of patterns that greatly condense the set of needed rules.

We then define a NLTK *chunking grammar*, a series of regex *rules* executed on these tagged sentences that combines tuples into *chunks*, corresponding to higher level concepts. The rules are executed in-order and are non-overlapping; that is, words consumed during one chunking will not be part of another chunk. The chunking grammar can include as many rules as desired. The most specific rules should be defined first in the grammar since rules are executed in order, and rules with the most flexibility should come last. Each expression attempts to match a sequence of tags. The standard regex tokens  $\{*, ., +, ?\}$  can be used to capture arbitrarily long groups of tags, and allow for optional parts of speech. For example, the rule *X*:

$$X: \{\langle\text{det}\rangle^6 ?\langle\text{noun}\rangle\langle\text{verb}\rangle + \langle\text{adverb}\rangle * \langle\text{adjective}\rangle +\}$$

chunks both *the (noun) is really superb* and *my (noun) has been reliable*.

Recall that five special tags are inserted during the pos tagging phase:  $\langle\text{feature}\rangle$ ,  $\langle\text{implicitfeat}\rangle$ ,  $\langle\text{naop}\rangle$ , and  $\langle\text{featchange}\rangle$ . These tags are used while chunking as follows:

- The  $\langle\text{feature}\rangle$  tag is included in every rule; the goal is to mine  $(F, o)$  pairs.
- The  $\langle\text{implicitfeat}\rangle$  tag indicates the word should be parsed as a feature synonym and an opinion. For example, we define *noisy* as an implicit synonym for the feature *sound*, and when tagging a sentence, we tag *noisy* with  $\langle\text{implicitfeat}\rangle$ . During the chunking phase, the tuple  $(\text{noisy}, \langle\text{implicitfeat}\rangle)$  is chunked as the  $(F, o)$  pair  $(\text{sound}, \text{noisy})$ , where *sound* is found using an inverted dictionary<sup>7</sup> that maps synonyms to features.
- The  $\langle\text{naop}\rangle$  tag is used in the chunking grammar wherever  $\langle\text{adjective}\rangle$  is used.
- Each chunk is post-processed to see if it contains a  $\langle\text{featchange}\rangle$  tag. If so, the  $\langle\text{featchange}\rangle$  word is checked in a dictionary of *feature changers* for the  $\langle\text{feature}\rangle$  word. If the  $\langle\text{featchange}\rangle$  word is a dictionary key, its value replaces the original feature. For example, we insert *handles* into the *car* feature changer dictionary with the value *performance*. The fragment *car handles great*, tagged:

$$(\text{car}, \langle\text{FEATURE}\rangle), (\text{handles} \langle\text{featchange}\rangle), (\text{great} \langle\text{opinion}\rangle)$$

<sup>6</sup> $\langle\text{det}\rangle$  refers to a *determiner* such as *this* or *my*

<sup>7</sup>If a dictionary has the property that no two keys have the same value, the dictionary can be *inverted*, such that given a value, it can be queried to produce the corresponding key. In this context, the required property means that no word is a synonym for two different features.

will be post-processed, prior to querying the chunk, and changed to:

$(performance, \langle FEATURE \rangle), (handles \langle featchange \rangle), (great \langle opinion \rangle)$

leading to the  $(F, o)$  pair  $(performance, great)$ .

- If the  $\langle vs \rangle$  tag is found in a chunk, the chunk's sentiment is inverted after querying.

With a well-crafted grammar only a few rules are needed; all results presented in §4.8.1 come from the four rules given in §A.

### 4.6.3 Advantages Of Chunking

In this section, we provide intuition as to why chunking works better than methods used in prior work for mining  $(F, o)$  pairs.

Prior work uses one of two methods to classify  $(F, o)$  pairs. The first is to compute a scoring function using the sentiment of adjectives in the sentence and their proximity to features or products. These methods are misleading if used on a sentence-wide basis because sentence structure plays a vital role in the meaning of sentences. Consider two sentences:

$s_0$  : *it does not have good (feature)*

$s_1$  : *the (feature1) is not good, but its (feature2) is excellent*

Sentences like  $s_0$  are problematic for sentence-wide scoring methods because the feature is close to the positive opinion *good*, but the sentiment is negative. On the other hand, including rules such as "invert the sentiment if a valence shifter is found in the sentence" would incorrectly classify  $s_1$  as negative for product2 . The solution is to apply the valence inversion rule to only the first part of  $s_1$ . We do so using the  $\langle vs \rangle$  tag, for example:

$r_0$  :  $\{ \langle verb \rangle \langle vs \rangle \langle "have/has" \rangle \langle opinion \rangle \langle feature \rangle \}$

$r_1$  :  $\{ \langle feature \rangle \langle verb \rangle \langle vs \rangle ? \langle opinion \rangle \}$

Using this inversion rule,  $s_0$  triggers  $r_0$  which correctly classifies it as negative. Moreover,  $s_1$ , chunked as  $[the (r_1), but its (r_1)]$  fails to trigger  $r_0$  because the sentence structure does not match, but triggers  $r_1$  for the first chunk which correctly classifies it as negative and triggers  $r_1$  again which correctly classifies the second chunk as positive.

Moreover, scoring functions can be used within chunks because they may contain multiple, sometimes conflicting opinions, e.g., *the interior is stylish but too small*. We use the function:

$$\sum_{o \in \text{chunk}} sentiment(F, o) * \left( \frac{1}{d(F, o)} \right)$$

Where  $sentiment(F, o)$  is determined using the methods discussed in the next section and  $d(F, o)$  is the distance (number of words) the opinion is from the feature in the chunk. This scoring function would classify the above chunk as positive. Another scoring function may

assign all opinion phrases in a chunk equal weight, e.g., assign each of *stylish* and *too small* a weight of .5 and classify the chunk as neutral.

The second method used to mine  $(F, o)$  pairs in prior work is to parse sentences according to syntactical templates, which are similar to chunking but do not use the regex operators  $\{+, *, ?\}$  [KIM<sup>+</sup>04, PE05, ZJZ06]. For example, a syntactical template may be:

$$\{\text{"The"} \langle \text{feature} \rangle \text{"is"} \langle \text{opinion} \rangle \}$$

which captures, for example, *the (feature) is (adjective)*, but would fail to capture *the (feature) is really (adjective)*. With chunking, we can capture both using optional tags:

$$\{\langle \text{feature} \rangle \langle \text{verb} \rangle \langle \text{adverb} \rangle * \langle \text{opinion} \rangle \}$$

To further demonstrate the expressiveness of optional tags, rules with optional valence shifters and "filler" words can capture many sentence constructs in a single rule, e.g,

$$\{\langle \text{valenceshift} \rangle ? \langle \text{opinion} \rangle \langle \text{"with"} \rangle \langle \text{det} \rangle ? \langle \text{article} \rangle ? \langle \text{feature} \rangle \}$$

matches all of "*no problems with the <feature>*", "*problems with my <feature>*", and "*no issues with my <feature>*".

Thus, chunking has advantages over both scoring based and template based methods.

## 4.7 Phase 4: Sentiment Querying

Here we describe how we classify  $(F, o)$  pairs. We first introduce two methods for handling context-dependent opinions (§4.7.1) and then discuss our classification algorithm (§4.7.2).

### 4.7.1 Handling Context-Dependent Opinions

Some opinion phrases are context-dependent, for example *cheap quality* vs. *cheap price*, and pose a challenge for opinion mining systems that use lexicons or WordNet [Pri14] to determine the sentiment of a word. Before discussing our sentiment querying algorithm, we introduce two concepts to handle context-dependent opinion phrases (CDOPs).

First, we define a *feature specific dictionary* (FSD),  $S_{f \in \mathcal{F}_p}^{p \in \mathcal{P}}[o \in \mathcal{O}]$ , for each  $\{\text{product, feature}\}$  pair  $(p, F)$ . These dictionaries are small and only contain phrases that have a sentiment when referring to  $(p, F)$  that is different than the sentiment it holds when used in other contexts. They produce a sentiment when queried with an opinion phrase  $o$  if  $o$  is defined to be context-dependent for that feature, or "unknown" otherwise:

$$\begin{aligned} S[\text{Car, Quality}](\text{cheap}) &\rightarrow - \\ S[\text{Car, Price}](\text{cheap}) &\rightarrow + \\ S[\text{Car, Performance}](\text{cheap}) &\rightarrow \text{Key Error} \end{aligned}$$




Figure 4.3: Pipeline for querying the sentiment of  $(F, o)$  pairs

We currently build these dictionaries manually by classifying a set of chunks, classifying them using the system, and inspecting chunks containing the same opinion phrase that were classified correctly by the system in some cases and incorrectly in others. The intuition is that if  $o$  is used in two chunks referring to  $f_1$  and  $f_2$ , if the system classifies  $(f_1, o)$  correctly but  $(f_2, o)$  incorrectly, it may indicate that  $o$  is a CDOP. Building these dictionaries automatically is left as future work, for example, using the work of Orimaye et al. [OAEG11].

Second, for some features, “more or less is always better”, e.g., *performance* and *price* respectively. We refer to such features as *oriented features*. We maintain whether each feature is positively, negatively, or non oriented, because *Intensity modifiers* like *low* and *high* (a list can be found in Paradis [Par97]) change their context when referring to such features. For example, note the orientations of the following  $(F, o)$  pairs:

(Range, low)  $\rightarrow -$ , (Range, high)  $\rightarrow +$ , (Maintenance, low)  $\rightarrow +$ , (Gear, low/high)  $\rightarrow N$

The methodology for querying oriented features is discussed in the next section.

## 4.7.2 Sentiment Querying Pipeline

We now describe how we classify whether opinion  $o$  expresses a positive, negative, or neutral sentiment about feature  $F$  in mined  $(F, o)$  pairs. Prior algorithms for opinion phrase classification are discussed in §4.9.2. To determine the sentiment of a  $(F, o)$  pair, we check the following set of *ordered* rules, summarized in Figure 4.3. The rules are checked from first (being the most specific) to last (being the default rule of simply returning  $N$ ), and when a rule triggers, a sentiment is returned and the rest of the rules are ignored.

1. If  $o$  is an intensity modifier and  $F$  is an oriented feature, we use the following sub-rules:

$$\begin{array}{ll}
 (o : +, F : +) \rightarrow \text{return } +, & (o : +, F : -) \rightarrow \text{return } - \\
 (o : -, F : +) \rightarrow \text{return } -, & (o : -, F : -) \rightarrow \text{return } +
 \end{array}$$

2. If the feature is found in the FSD  $S_F^P$ , the corresponding label is returned.
3. If querying the *default sentiment lexicon* returns a sentiment, the sentiment is returned. Our default sentiment lexicon is built from the MPQA Opinion Corpus [Wil, Wil08] and contains over 6,800 opinion phrases.
4. Return neutral ( $N$ ).

There is one exception to this pipeline. Phrases such as *wondering*, *curious*, *as long as*, *can*, *may*, and *will* are common in chunks that contain an opinion phrase but do not express an opinion, rather, they ask a question or talk about a hypothetical scenario. If any word in a chunk is found in a list of these phrases, the chunk is classified as neutral.

We now formally define  $\mathcal{O}$  (§4.2) as:

$$\{\text{NAOPs}\} \cup \{\text{intensity modifiers}\} \cup \{\text{MPQA}\} \cup \{S_{f \in \mathcal{F}^p}^{\text{PEP}}.keys \forall (p, F)\} \quad (4.1)$$

## 4.8 Phase 5: Results

Our system gives prospective EV owners a high-level product overview with the ability to drill down into opinionative sentences about features of interest. The sentences reveal sentiments found during field trials, but also some that were only realized after the owners possessed their vehicles for a significant amount of time. Moreover, the system significantly reduces the amount of text the user must read to determine owners’ opinions because it filters out sentences that do not contain features of interest or opinions.

We discuss corpora generation and our performance metrics in §4.8.1. In §4.8.2 we support the above claims. We discuss precision, recall, and classification errors in §4.8.3.

### 4.8.1 Evaluation Methodology

Here we discuss our evaluation methodology. We discuss our corpora generation in §4.8.1, define our evaluation metrics in §4.8.1, and define terms used in our results in §4.8.1.

#### Corpora Generation

For our evaluation, we crawled the owner discussion forums for the three best selling (see §2.1.1) EVs—the Nissan Leaf [Mis14c], the Chevrolet Volt [Mis14b], and the Tesla Model S [Mis14a]. After mining the forums, our raw corpus contained 107,293 Volt sentences, 220,906 Leaf sentences, and 500,668 Model S sentences. We then filtered out all sentences containing no synonyms in  $\mathcal{F}^{\text{Volt}}$ ,  $\mathcal{F}^{\text{Leaf}}$ , and  $\mathcal{F}^{\text{Tesla}}$ , leaving 10,519 Volt, 19,799 Leaf, and 73,228 Tesla Model S sentences, which we collectively denote as the *feature corpus*.

We demonstrate our system’s operation in two ways:

- For the Leaf and Volt, we manually labeled—read and annotated with our opinion of the sentences’ sentiments—a small fraction ( $\approx 25\%$ ) of the feature corpus. We denote this set of labeled sentences as the *ground truth corpus* (GTC). The GTC contains 2,566 Volt and 5,514 Leaf sentences containing at least one feature. In §4.8.1 we show the precision, recall, and the distribution of sentiments for the GTC. No results are presented for the Leaf and Volt for the unlabeled portion of the feature corpus.

- For the Tesla Model S, we demonstrate the sentiment extraction capability on a large corpus which we cannot manually label: we show just the distribution of sentiments for the entire feature corpus (73,228 sentences).

## Precision and Recall

Let

- $s_F^+, s_F^-$  be the number of sentences the **system** classifies as +,− about feature  $F$ .
- $t_F^+, t_F^-$  be the number of sentences **we** manually classify as +,− about feature  $F$ .
- $c_F^+, c_F^-$ , where  $c$  stands for “correct”, denote the number of sentences the system classifies correctly as +/− about  $F$  that **we also classify** as +,− about  $F$ . Intuitively,  $c_F^+ \leq s_F^+$  is the number of sentences that we agree are positive about  $F$ . We stress *about  $F$*  because sometimes the system classifies a sentence’s sentiment correctly but for the wrong feature (the sentence is actually referring to a different feature).

We compute two metrics, *opinion precision* and *opinion recall*, for each feature  $F$ :

$$\text{precision}(f) = \frac{c_F^+ + c_F^-}{s_F^+ + s_F^-} \quad \text{recall}(f) = \frac{c_F^+ + c_F^-}{t_F^+ + t_F^-}$$

Chunks misclassified as (+/−) lead to lowered precision, and (+/−) chunks misclassified as neutral lead to lowered recall.

Due to our methodology of filtering out all sentences which contain no feature synonyms, our measure of recall is not “true recall”; we overestimate recall. This is because some sentences in the raw corpus may refer to a feature implicitly or using a rare synonym. However, because most sentences do not contain features, simply selecting a percentage of the (massive) raw corpus to classify for ground truth, without first filtering the sentences by known feature synonyms, may lead to a ground truth corpus containing very few sentences containing features. We thus believe only labeling sentences containing known features is a reasonable approximation.

## Definitions

Here we define some features of electric vehicles used in our figures:

- *General* refers to any opinion referring to the *car* itself and not a specific feature, such as *this car is amazing*.
- Current EV batteries lose capacity over time as they are repeatedly charged and discharged, and if they are subjected to extreme temperatures [BCM11]. *Degradation* refers to the effect of charging and climate on a batteries capacity and life.
- We denote anything related to heating and cooling, including features such as heated seats and pre-warming (warming the EV while it is still plugged in at home), as *HVAC*.



- *Carwings* and *Onstar* are products included with the Leaf and Volt respectively that provide various feedback, charging, and safety services to drivers.
- *MiscFeats* refers to a mix of other features including power steering and navigation.

## 4.8.2 Distribution of Sentiments

As discussed in the prior section, we show the distribution of sentiments found in the ground truth corpus (GTC) for the Leaf and Volt, and for the entire feature corpus for the Tesla Model S. The distributions of sentiments found in the GTC for the Leaf and Volt are shown in Figure 4.4. The distribution of sentiments found in the feature corpus for the Tesla Model S is shown in Figure 4.5<sup>8</sup>. Examining these figures gives a high-level view of opinions about the various product features. If the user wishes to learn more about each feature, they can then read the corresponding classified sentences. Some sentiments found in this manner match sentiments found during expensive field trials:

- As discussed in §3, *price* and *range anxiety* are commonly cited as the two largest EV adoption barriers. This is supported by all three ownership forums, though for the Leaf and Volt the number of statements about these features in the GTC is small (for price, this problem is further discussed in §4.8.3). For the Leaf, which has only 100 miles of range, the sentiments about range and range anxiety are mostly negative, and sentiments towards range anxiety for the Volt are positive (no range limitations). The Tesla Model S has three times the range of the Leaf so most sentiments regarding anxiety are positive. For price, the majority of sentiments for all three products is negative.
- *Maintenance* is also commonly cited as a major selling point of BEVs—the absence of an ICE means fewer moving parts that can fail and less fluids to change. From Figures 4.4 and 4.5, we see that sentiments towards maintenance for the Leaf and Tesla (BEVs with no ICE) are overwhelmingly positive. Interestingly, the Volt also has positive reviews for maintenance.
- Field trials often conclude that participants prefer the EV charging process over the ICEV refueling process, because many drivers charge their EVs overnight or while at work. We also find that sentiments towards charging are positive. Reading the classified sentences reveals that many drivers receive free charging at work, are not concerned about the lack of public charging stations, and live in areas with time-of-use electricity pricing so they are able to charge overnight cheaply.
- As a final example, the mostly positive sentiments for the *general* category reveals some early adopter bias. Even when drivers express concern about the set of features they dislike, many end their discussions with comments like “...but I love my Leaf/Volt/Tesla”.

---

<sup>8</sup>The numbers above the bars in the Figures 4.4 and 4.5 show the number of sentences classified for that feature. These numbers do not sum to the exactly the size of the GTC for the Volt and Leaf (2,566 Volt and 5,514 Leaf sentences) or the size of the feature corpus for the Model S (73,228 sentences). This is because while all of these sentences contain features, they do not all contain opinion phrases or even adjectives—these (small number of) sentences are discarded during the classification stage and were not labeled.



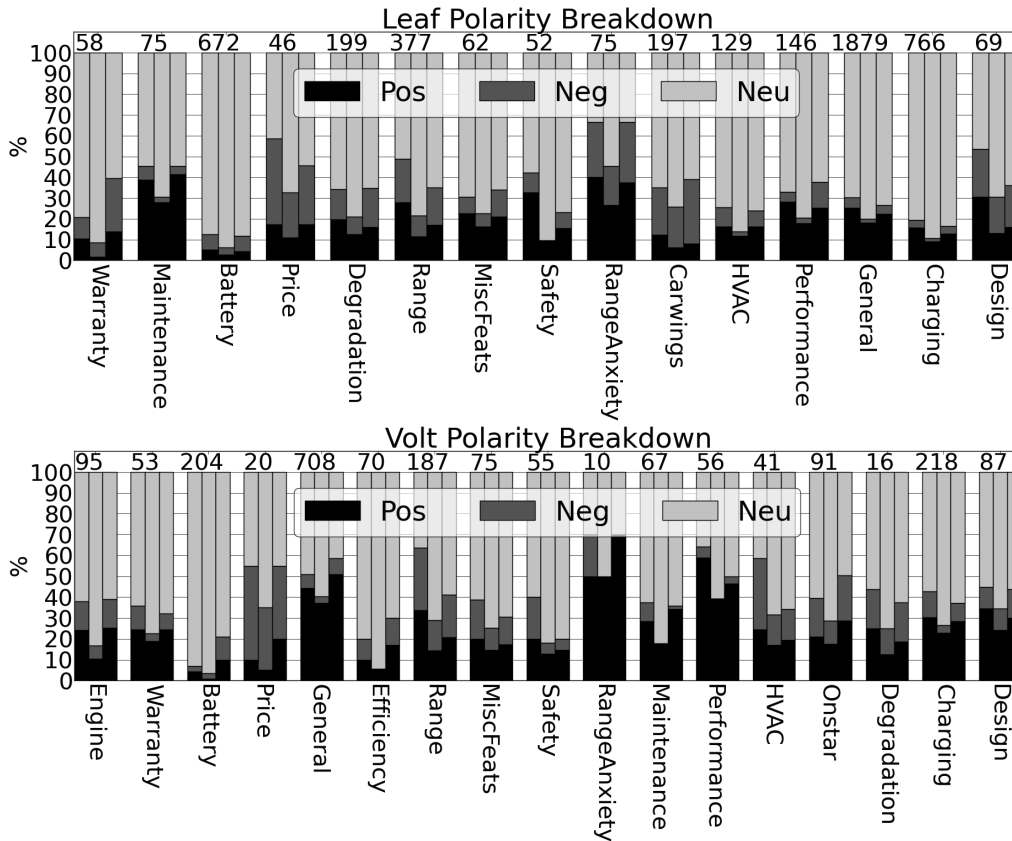


Figure 4.4: Sentiments of all Leaf (top) and Volt (bottom) sentences in the ground truth corpus. There are three bars for each feature. The first shows the polarity distribution of sentences the system classified as positive, negative or neutral, even if for the incorrect choice of feature ( $s^+$ ,  $s^-$ ,  $s^N$ ). The second shows the distribution of those the system correctly classified for the correct feature ( $c^+$ ,  $c^-$ ,  $c^N$ ). Finally, the last shows the distribution of ground truth sentences ( $t^+$ ,  $t^-$ ,  $t^N$ ).

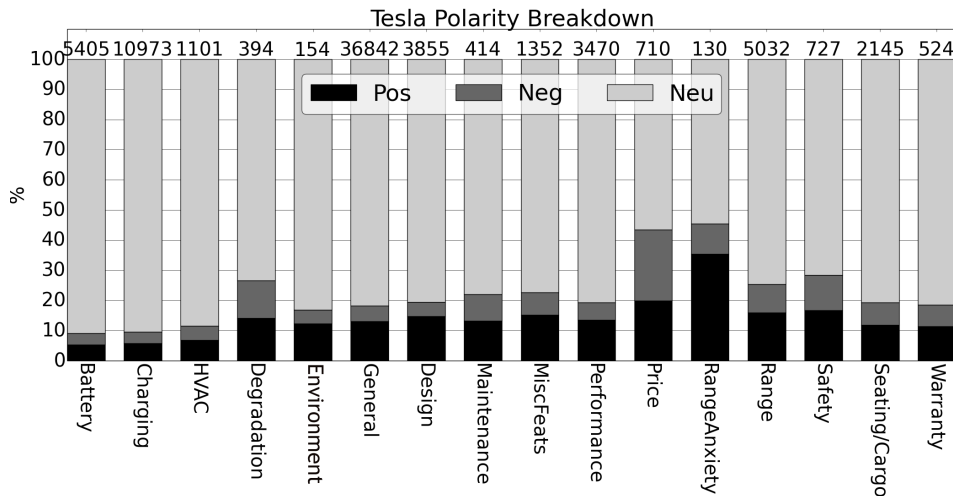


Figure 4.5: Sentiments for all Tesla Model S sentences in the feature corpus. Because there is no GTC, we only show one bar, the polarity distribution of sentences the system classified ( $s^+$ ,  $s^-$ ,  $s^N$ ).

In addition, we are also able to derive insights that were only perceived after the owners had their vehicles for a significant duration of time. Such insights are not possible to elicit from field trials. The most important example is owners' experience with *battery degradation*, the effect of repeated charging and climate on battery capacity over time. While it is known that climate and charging cycles affect battery life, it is unknown to what extent this is the case [BCMw11]. We find that some owners are experiencing non-trivial battery degradation and about 50% of sentiments towards degradation are negative<sup>9</sup>. Manufacturers can use these sentiments in conjunction with the owners' location (if available on the forum) to derive conclusions about the effect of climate on battery capacity, e.g., we find that owners in hot regions such as Arizona and southern California post more often about degradation. Closely related are sentiments on *warranty*, which are mostly negative, because some owners have filed for battery replacements through battery warranties against capacity loss. Reading classified statements from warranty and maintenance reveals other vehicle problems of interest to manufacturers, such as the replacement and maintenance rate of various parts. These problems may not appear in shorter duration field trials.

Figures 4.4 and 4.5 also demonstrate the text reduction capability of our system. As discussed in §4.8.1, out of the  $\approx 830,000$  sentences in the raw corpus, only  $\approx 100,000$  contained a feature synonym of interest, specifically, a feature synonym in  $\mathcal{F}^{Volt}$ ,  $\mathcal{F}^{Tesla}$ , and  $\mathcal{F}^{Leaf}$  collectively. Moreover, we see in Figures 4.4 and 4.5 that  $\approx 70\%$  of sentences containing a feature of interest are neutral. Thus, given the set of defined features in this case study, the system reduced the space of text by 97% to the 3% of relevant text.

### 4.8.3 Performance

Figure 4.6 show the precision and recall of our system on the GTC. The performance is on par with or better than prior FBOM systems discussed in §4.9.1. To interpret these results, we now present several types of errors we encountered because they give insight into the complexities of opinion mining. Some of these errors show how the system performance can be improved by continuously fine tuning the system as time permits, while others show that perfect performance is impossible due to inherent English ambiguity.

Some features are hard to mine or classify opinions for. For example, our system performs poorly on classifying opinions related to *safety*. We find the word *issue* is heavily overloaded but used often—in some instances it is used synonymously with *hazard*, such as *that is a safety issue!*, and in other cases it is used synonymously with *feature*, e.g., *grounded charging is a safety issue*. Our system also performs poorly on classifying *warranty* opinions for similar reasons. For example, it is difficult to tell programmatically when the phrase *not covered* is used as a negative or neutral sentiment. Sometimes posters state facts with this phrase, e.g., *the windshield is not covered in your warranty*, and other times to express frustration, such as *the repair was not covered by my warranty*. Moreover, notice that

---

<sup>9</sup>Positive sentences about negative features usually include valence shifters, e.g., "I have not had any battery degradation".




Figure 4.6: Performance of classifying all Leaf (top) and Volt (bottom) sentences in the GTC.

*price*, even though it is cited as a major adoption barrier, has very few comments. We find many posts comment on the price of something other than the price of the vehicle, such as the price of charging and electricity. After experiencing a high rate of classification errors regarding this feature, we tuned our system to only classify a chunk as referring to *price* if the chunk contained both a synonym of *price* and a synonym of *General* such as *car*, *Volt*, *Leaf*, etc. This led to a tradeoff: from Figures 5 and 6 we see our system performs well with respect to precision and recall for *price*, but the number of classified sentences to draw conclusions from is small. Conversely, our system classifies *General* opinions well for both products, for which there are many. This is because sentences expressing general sentiments are often clear and brief, e.g., *I love my Leaf!* and *this is an excellent car*.

Table 4.1 shows examples of errors we fixed by updating the feature-specific sentiment dictionaries, updating the default sentiment lexicon, or other adding other domain knowledge (DK) as shown. Others errors are "better left unsolved"—tuning the system to correct these creates larger problems elsewhere. Sometimes a feature synonym is used in two different contexts, e.g, "ice" can refer to an engine (Internal Combustion Engine) or the weather condition. We set the system according to which usage is most common, but errors will occur when the word is used in the less common context. Others represent chunking errors where changing the chunking grammar to fix the error caused more problems in other sentences because the offending sentence structure is uncommon. Finally, some words are context-dependent even within the context of one feature. We accept such errors, shown

in Table 4.2, as necessary due to the ambiguity and complex structure of English. Finally, some ambiguous sentences can be classified differently by different human readers. Some correspond to a parameter which sounds positive to some, such as “100 miles per gallon”, but for which it is hard to impose a strict cutoff for which all human readers agree, e.g., “all mileages over X are positive”. Others are ambiguous sentences that could be classified as either positive or negative. Examples of such errors are shown in Table 4.3.

Sentence	Class	Tr	Problem	Solution
<i>The standard warranty is more than enough.</i>	<i>N(Warranty)</i>	+	<i>more than enough</i> not recognized as an opinion	add it to the default sentiment dictionary
<i>My Carwings sometimes hangs.</i>	<i>N(Carwings)</i>	-	<i>hangs</i> not recognized	add it as - to the Carwings FSD
<i>This car is a blast to drive!</i>	-( <i>General</i> )	+	<i>blast</i> is context dependent	add it as + to the General FSD
<i>The dealer showed me how to perform the recommended maintenance</i>	+( <i>Maintenance</i> )	<i>N</i>	<i>recommended</i> is context dependent	add it as <i>N</i> to the Maintenance FSD
<i>I love my Audi, it is a great car.</i>	+( <i>General</i> )	<i>N</i>	Subject error	classify chunks w/ other popular car names as <i>N</i>
<i>I had a low battery warning.</i>	-( <i>Battery</i> )	<i>N</i>	<i>warning</i> is context dependent	classify chunks w/ this phrase as <i>N</i>
<i>It has a 5 star safety rating.</i>	<i>N(Safety)</i>	+	<i>5 star</i> not recognized	add it as + to the Safety FSD
<i>The hot weather kills my range.</i>	<i>N(Range)</i>	-	<i>kills</i> , a verb, not recognized as an opinion	add it as a non-adjective opinion to the Range FSD.

Table 4.1: Examples of solvable sentiment analysis errors. “FSD” denotes “feature specific dictionary”

Sentence	Class	Tr	Problem
<i>"It was charge free."</i>	+( <i>Charging</i> )	<i>N</i>	poster is not referring to charging but rather the price of something. However, sometimes posters talk about free charging.
<i>The Volt is the best car ever and I will never go back to a crude ICE</i>	-( <i>General</i> )	+	ICE most often refers to the Volt’s ICE, but sometimes to non hybrid vehicles in general.
<i>My new radio has far less degradation near mountains.</i>	+( <i>Degradation</i> )	<i>N</i>	here the poster is referring to radio signal degradation; static near mountains, power lines, or tunnels.
<i>It felt like the engine was on.</i>	+( <i>Engine</i> )	<i>N</i>	<i>like</i> is a tough word. Even within the context of one feature, it can be used as a comparator or as a positive sentiment (more common).
<i>I love everything about this car, with the exception of the exterior.</i>	<i>N(Exterior)</i>	-	Chunking error. We find most opinions do not “cross over” commas, hence commas are not included in chunking rules. Here the valence shifter “exception” loses the opinion to negate—love, since it is not in the chunk.

Table 4.2: Examples of sentiment analysis errors that are “better left unsolved”.

Sentence	Class	Tr	Problem
<i>The Leaf handles 99% of my annual driving."</i>	<i>N(General)</i>	?	Should this be +? What’s the +/- Cutoff?
<i>I get around 80 mpg."</i>	<i>N(Efficiency)</i>	?	Should this be + for Efficiency? Cutoff?
<i>The Leaf is very easy to push.</i>	+( <i>General</i> )	?	This might mean the poster’s car died, or something related to performance/handling?
<i>Level 2 charging is very efficient.</i>	+( <i>Charging</i> )	?	Should this be +? The poster may either be happy with their charging experience or just stating a fact.
<i>I never worry about my mileage.</i>	+( <i>Efficiency</i> )	?	Is the poster stating a fact or that they have plenty of range?
<i>I am sad my Volt is in for maintenance.</i>	-( <i>Maintenance</i> )	?	is the poster disappointed the car requires maintenance, or are they stating they miss driving their car?

Table 4.3: Examples of ambiguous sentiment analysis classifications.

## 4.9 Prior Feature-Based Opinion Mining Methodologies

This section discussed related work on feature-based opinion mining §4.9.1 and sentiment classification §4.9.2.

### 4.9.1 Feature-Based Opinion Mining

Hu and Lui define the concept of feature-based opinion mining [HL04a] and introduce the first FBOM system, Opinion Observer [LHC05]. Their system first builds  $\mathcal{F}^p$ , then finds and summarizes positive and negative opinions corresponding to each feature. The authors use a lexicon to determine the sentiment polarity of adjectives in sentences containing product features, then classify sentences based on the number of positive and negative words in a sentence.

Hu and Lui [HL04b] expand on the product feature identification phase. The authors present an association rule mining process to build  $\mathcal{F}^p$ . The mining process finds noun phrases (e.g., *digital camera*) that are likely to be product features. Pruning rules are used to trim the set of mined product features. Lui, Hu, and Cheng [LHC05] further update their system to use supervised learning for detecting *implicit features*, e.g., *fast* refers to the feature *performance*. Finally, Ding, Lui, and Yu [DLY09] update their system with a better sentiment classifier. For each feature  $F$  in a sentence, the authors compute a scoring function based on all adjectives in the sentence and their distance from  $F$ . Hence, if there are two features, the adjectives closest to each will influence their scores the most, but all adjectives have a non-zero contribution to the score of all features. This improvement better classifies  $(F, o)$  pairs than simply averaging the classification of all adjectives.

Scaffidi et al. [SBC+07] build a system called "Red Opal" which allows users to search for products based on the ratings of specific product features. Products are ranked feature-wise based on numerical review ratings, like those found on Amazon, rather than opinion words in the reviews. While the system achieves good results when numerical reviews are available (which they are typically on online retailers), This system cannot mine forums, article comments, or other text.

Popescu and Etzioni [PE05] present OPINE, a system that uses different algorithms for building  $\mathcal{F}^p$  and mining/classifying  $(F, o)$  pairs. Their feature selector uses pointwise mutual information between potential features and *metonymy*<sup>10</sup> *discriminators*, such as *of object*, *object has*, *object comes with*. These phrases can distinguish whether a sentence is opinionative or factual, leading to a higher precision for detecting  $(F, o)$  pairs. To mine  $(F, o)$  pairs, the authors use *syntactical templates* such as  $\langle \text{feature} \rangle$  *is*  $\langle \text{value} \rangle$ —if a sentence matches this pattern,  $(\text{feature}, \text{value})$  is mined as an  $(F, o)$  pair. These syntactical templates motivated our use of *chunking* to parse  $(F, o)$  pairs. Zhang et al [ZNC10] use a graph mining approach to rank several products according to various product features. The authors divide opinionative sentences into two sets, those that express opinions on just one product (subjective), and those that compare two or more products (comparative). Products are treated as nodes in a "feature graph". Subjective sentences and their classification are used to weigh nodes, while comparative sentences and their classification are used to weigh edges between the two products being compared. Then a PageRank algorithm is used to rank the set of nodes according to the feature. In the future, when many EV models are sold and EV sales increase, this work may help compare several EV models. on 4.6.

---

<sup>10</sup> noun: *substitution of the name of an attribute for that of the thing meant, e.g., "suit" for "business"*

This resolves problems with Hu’s [HL04a] system because not all opinions in a sentence are associated with every feature in the sentence. They use statistics and classifier-based methods for classifying word sentiments, as opposed to Hu’s lexicon based approach.

Jin et al. [JHS09] discuss Opinion Miner, which trains a hidden Markov model to find  $(F, o)$  pairs and simultaneously classify them; the only work we know of that merges these steps. The model is trained using linguistic constructs, syntactical templates, and word sentiments. The model learns to mine constructs such as “negative opinion about [feature]”, instead of first finding  $(F, o)$  pairs and then separately classifying each pair. The authors manually tag certain constructs, then use synonyms, antonyms, linguistic constructs, and other bootstrapping techniques to grow the set of training examples.

Others have improved upon these systems. Kobayashi et al. [KIM+04] suggest a domain-knowledge-driven feature and opinion phrase selection process, instead of the general association mining techniques. Their work is concerned only with mining product features and opinion phrases, and does not discuss sentiment classification. They introduce an iterative algorithm that generates candidate features and opinion phrases, and manually select those that are valid. In each iteration, more candidates are selected based on the prior iteration, and the process is repeated until an iteration goes by where the human selects no candidates. Zhuang et al. [ZJZ06] present a case study of Hu’s system using movie reviews. The authors first identify features such as special effects, acting, plot, directors, cast. The authors then identify opinions discussing those features and classify the opinions as  $\{+, -\}$ . Finally, the authors summarize the results. The authors incorporate domain knowledge using supervised learning into feature and opinion keyword mining. For example, they have several movie fans manually tag reviews for features, feature opinions, and cast members. To parse pairs, the authors use syntactical templates like Etzioni [PE05] and Kobayashi [KIM+04].

Our system builds upon these systems as follows. Like Zhuang et al. [ZJZ06] and Kobayashi et al. [KIM+04], we incorporate domain knowledge (DK) into our mining process, specifically in the chunking and querying phases. For feature mining, a part of the chunking phase, we use Hu’s association mining technique [HL04b] and then manually prune and collapse the feature set like Kobayashi et al. [KIM+04]. We note this is feasible because there are fewer than one hundred common features one may talk about within the context of a car. For parsing sentences, our use of chunking is similar to using syntactical templates like Kobayashi, Popescu, and Zhuang [KIM+04, PE05, ZJZ06], but more powerful (see §4.6.3). During our sentiment classification stage, we make use of an open-source sentiment dictionary, the MPQA Opinion Corpus [Wil, Wil08]. Like Ding et al. [DLY09], we handle context opinions, but introduce two new methods for this in §4.7.1.

## 4.9.2 Word and Sentence Polarity Classification

Here we discuss the four prior methods of opinion phrase classification:

1. *Lexicon methods* [DLY09, ES06a, KH04, HL04b, HL04a, LHC05, ES06b], which we use in our system, start with a small set of classified *seed words*. These sets are then grown

using synonyms from WordNet [Pri14] or other glosses.

2. *Semantic methods*, e.g., [KI06, HM12, ZJZ06] classify the sentiments of words and sentences based on lexicons and the semantic rules of the English language. For example, two adjectives joined by *and* are likely to share the same polarity, e.g., *sunny* and *beautiful*—if we know the polarity of one in the conjunction, we can infer the other.
3. *Distance methods*, e.g., [TAV06, Tur02, GA05] measure the polarity of a given word based on the *distance* of that word from a set of positive and negative seed words. Distance is normally computed via WordNet or by analyzing the co-occurrence of words in a large corpus, with an example function being  $d(\text{word}, \text{good}) - d(\text{word}, \text{bad})$  where  $d(x, y)$  gives the distance (computed via WordNet) from word  $x$  to word  $y$ . Another common distance measure is *pointwise mutual information* [TC03].
4. *Classification methods*, e.g., [WWH04, PE05, JHS09] treat the problem of determining the polarity of opinion phrases as a machine learning problem. Rather than learning the sentiment of individual words, a classifier is trained to classify sentences directly. These authors manually label sentences, train a polarity classifier based on this labeled training data, and then classify sentences in the unlabeled data using the trained classifier.

Some opinion phrases are context-dependent and pose a challenge for opinion mining systems. Ding et al. [DLY09] give an algorithm for querying the sentiment of such phrases. They first attempt to query all opinion phrases in a sentence using a lexicon approach. They then use an algorithm which considers syntactical constructs and the sentiments in neighboring sentences to classify unclassified phrases per these lexicons. Two other works have also studied classifying context-dependent adjectives [TP10, WW11].

## 4.10 Conclusions & Future Work

Understanding EV owners' experiences with and perceptions towards EVs is helpful for manufacturers to build later-generation models more aligned with drivers' mobility preferences and requirements. We build an opinion mining system that classifies opinions found in EV ownership forums. Our system helps the user obtain a high-level overview of opinions on various product features, and greatly reduces the space of text the user is required to read to extract opinions. These opinions are useful to prospective buyers (what features do owners like and dislike?), marketers (what features should be advertised?) and manufacturers (what features should be improved?). To build this system, we combine prior opinion mining systems with several new optimizations and our EV domain knowledge. We furthermore open sourced our system for extension by other researchers, as we find most prior opinion mining systems are unavailable.

While our results are on par with prior opinion mining systems, we emphasize two of our contributions. First, prior FBOM systems have focused on products where lots of numerical reviews *are* available. For example, Hu and Lui [HL04a] focus on digital camera reviews, Kobayashi et al. [KIM<sup>+</sup>04] focus on movie reviews, and Scaffidi et al. [SBC<sup>+</sup>07] focus on



any product with numerical reviews. While these papers mine opinions on specific features, a useful addition for consumers, reviews with easily-parseable schema exist for these products (e.g., Amazon and IMDB). We present work on mining opinions in free text that does not follow any schema. Second, our system is open source [Car]. The package includes our domain knowledge input and all manually built lists and dictionaries. Future researchers or users that wish to extend the system have a comprehensive starting point. In contrast, the three closest prior systems are either closed source or unavailable—Opinion Miner is proprietary [Ope], the Opinion Observer system turned into the proprietary OpinionEQ [Opi, LH], and the OPINE system [PE05] was never released by the authors.

We end with several avenues for extending and improving our work:

1. An interesting future avenue would be to see how sentiments change over time by analyzing sentiments periodically, e.g., monthly. It is possible to see the change in an individual's opinion over time if they post using a user name. For owners that post anonymously, we can examine how collective sentiments change. Field trials have attempted to determine how perceptions change over time, for example, by interviewing participants both before and after the trial. However, because individual participants in trials are often only given a vehicle for a few weeks or months, their opinions may not change as much as they would over a period of years. Thus, tracking owners' opinions over time online can reveal insights about how perceptions change with vehicle experience.
2. Similarly to (1), we currently weight all data found in the forums equally. An alternative methodology would be to weight recently mined sentiments more than older sentiments.
3. We do not currently perform pronoun resolution. A poster may explicitly mention a feature in one sentence and then write several more opinionative sentences about the same feature using pronouns easily resolved by a human reader. Our system only categorizes opinionative sentences where an explicit or implicit feature is found. Pronoun resolution, while difficult, may be used to infer the features being discussed.
4. We aggregate data for each individual product in a separate database. When mining for product  $p$ , we assume all  $(F, o)$  pairs found in the database for  $p$  refer to  $p$ . This produces erroneous results when a poster discusses another product in their post. Future work could detect the product being discussed from the context.
5. We do not perform spam or malicious text detection. We treat all sentences equally even though some may contain sentences injected by malicious sources, such as drivers who oppose a particular brand, or advertisements posted by spamming bots.
6. Our methods for building the FSDs and lists of oriented adjectives are simple; we use pre-built lists and manually add others. This part of our system can be improved using recent work in the field of classifying context-dependent adjectives [TP10, WW11].
7. We do not distinguish between chunks referring to one product and chunks comparing two products as proposed by Zhang et al [ZNC10]. Modifying the chunking grammar to include comparative templates may reduce classification errors.



# Chapter 5

## A Vehicle Pool Model For Reducing Range Anxiety

**Publication reference:** *T. Carpenter, S. Keshav, J. Wong. Sizing Finite Population Vehicle Pools. IEEE Transactions on Intelligent Transportation Systems, Volume PP, Issue 99, 2014.*

---

### 5.1 Synopsis

One way of alleviating range anxiety is to give BEV owners access to an ICEV when they need to make a trip longer than their BEV range. In this chapter, we focus on sizing ICEV *pools*, which are a number of ICEVs stored at the same location to be used on demand by BEV owners. This pool can be at existing dealerships (i.e., the pool is formed by some of the dealers's unsold ICEVs) as BMW recently proposed<sup>1</sup> [Ing13], or operated as a community or government run program to facilitate BEV adoption. Regardless of the business model, our goal is to minimize the number of vehicles in the pool while still meeting a desired percentage of subscriber requests.

In this chapter, we analyze the *static* sizing of a *single* vehicle pool. That is, we assume the pool is sized only once for all future demand (given information about previous and expected future demand), and we assume that all ICEVs are stored at a single location. In the next chapter, we generalize both of these assumptions, i.e., we extend the methods presented here to *periodically* size a *network* of pools.

The main contributions of this chapter are:

---

<sup>1</sup>While this concept was contemporaneously and independently proposed by EV manufacturers, they did not address *sizing* such pools.

1. We study three techniques to size ICE pools to meet a desired *quality of service* (QoS) target (§5.2), defined as the fraction of customer requests that are met. The three methods assume different types of data (or no data) regarding customer demand patterns are available.
2. We apply busy period sizing to finite population queueing systems, which, to the best of our knowledge, has not been done in prior work. Busy period sizing allows us to apply queueing methods, which require stationary arrivals, to a non-stationary system by examining the system during the period where the arrival traffic is most stationary. In prior work, this has only been done with infinite population queueing systems.
3. We propose a weighted-average methodology to model a population of heterogenous users as a population of "average" users. This allows us to apply queueing methods that require the population be composed to *i.i.d* users to a system where users are not *i.i.d*.

We numerically evaluate the performance of the three sizing methods using eight years of data from an Ontario-based car share (§5.3 and §5.4). We use three performance metrics: *availability* (percentage of requests served), *utilization* (the percentage of time vehicles in the pool are used), and the *member-to-vehicle ratio* (the size of the user population relative to the size of the pool). We find our sizing methods, when properly configured, size a pool within 1—3% of desired QoS (availability) targets.

## 5.2 Sizing A Pool Subject To A QoS Target

The precise statement of our problem is as follows: we wish to determine the least number of ICEVs that should be made available to a finite set of *subscribers* so that the probability of an unmet request is smaller than a target value of  $\epsilon$ . The parameters  $\epsilon$  and  $1 - \epsilon$  are interchangeably referred to as the *quality of service* target, e.g., a target of  $\epsilon = .05$  means that more than 95% of vehicle requests should be met. We envision subscribers to be BEV owners that pay for access to the pooling service. This sizing is difficult because demand for ICEVs is non-stationary. Vehicles are most commonly used during the day, and moreover, there are holidays during which the demand for vehicles is higher than during the rest of the year. We use a dataset of carshare use to show in §5.4.1 that demand for a real-world carshare is non-stationary. *Busy period sizing*, that is, sizing a system according to the period of highest demand, is often used to size non-stationary infinite population systems in telephone networks [Ive09]. However, we are not aware of any prior work which uses busy period sizing to size *finite* population systems with non-stationary demand. While the motivation for our work was to size a pool of ICEVs to be used by BEV owners, this work can be used to size any finite population, non-stationary queueing system according to a desired QoS target. For example, this work can be used to size a private parking lot.

We present three sizing techniques based on different levels of information about subscribers' demand patterns. We discuss a Binomial-based sizing method in §5.2.1. This approach only requires knowing the average number of times per year the average subscriber will use the

pool. We present a queueing theoretic algorithm in §5.2.2. This approach requires knowing, on average, how often subscribers need vehicles and how long they need them for. Finally, we study a different queueing approach in 5.2.3. This approach requires the pool operator to track the arrival times of subscribers.

Throughout this chapter we use the term *loss probability* to denote the probability that a subscriber arrives to an empty pool (i.e., is 'blocked'), given the pool has size  $m$  vehicles, and denote it by  $p(b|m)$ . Our goal is to size the pool such that  $p(b|m) < \epsilon$ .

We note that some carshares permit advance *reservations* allowing their subscribers to pre-reserve a vehicle for a predetermined period of time. Our sizing methods can be used in this case as well, by sizing the pool such that the probability a reservation cannot be accommodated is less than  $\epsilon$ .

## 5.2.1 Binomial Based Sizing

Our first technique is based on the Binomial distribution, which can be used to conservatively size a pool with only a minimal amount of information about future usage. Consider a pool on a given day. We want to find the pool size  $m$  such that  $p(b|m) < \epsilon$ . Let

- $p(a)$  be the probability a subscriber arrives on a given day. In this conservative sizing method, we assume all subscribers are independent.
- $S$  be the total number of subscribers
- $m$  be the number of vehicles in the pool
- $\epsilon$  be the QoS

Assume subscribers always arrive to the pool at the same hour  $h$  and always return their vehicle before  $h$  the following day. These two assumptions imply that the same vehicle cannot be reused multiple times a day and no vehicle is used for multiple days in a row. The probability exactly  $k$  subscribers arrive to the pool on the same day (during  $h$ ) is given by:

$$p(k) = \binom{S}{k} p(a)^k (1 - p(a))^{S-k}$$

If exactly  $k$  subscribers arrive,  $p(b|m, k)$  is given by:

$$p(b|m, k) = \begin{cases} 0 & \text{if } k \leq m \\ (k - m)/k & \text{otherwise} \end{cases}$$

Thus, given  $S$  subscribers, the probability a random subscriber finds an empty pool is given by marginalizing out  $k$ :

$$p(b|m) = \sum_{k=0}^S p(b|m, k) p(k) \tag{5.1}$$

Eq(5.1) gives  $p(b|m)$  for a fixed  $m$ ; the algorithm presented in §5.2.4 determines the smallest value of  $m$  that meets the QOS target.

## 5.2.2 Sizing Via The Engset Loss Model

Given additional data regarding subscriber workload, pools can also be sized using the *Engset Loss Model* (ELM) [Kle75a, Ive09, Tij03], also known as the  $G_I/G_I/m/m/S$  queue. In this model, we treat arriving subscribers as jobs arriving to a queueing system. We further treat vehicles in the pool as parallel servers that serve incoming jobs. Each job (subscriber) that arrives receives dedicated service from one server (takes a vehicle) until the job is processed (the subscriber brings back the vehicle). We assume subscribers will not wait for a vehicle when the pool is empty, thus there is no buffer in the system. Finally, there are only a finite number of sources from which jobs can be generated—the pool subscribers—leading to a finite population model. We thus model our system as a  $G/G/m/m/S$  system.

When a subscriber is borrowing a vehicle, we say they are in the *service state* and remain in this state on average for their *mean service time* (*MST*), the average duration they need a vehicle for. After finishing service, the subscriber enters the *thinking state* and waits on average their *mean think time* (*MTT*), their average duration between completing service and their next request.

We first give the formula for the blocking probability of this queue. Then, we describe our methodology for estimating the input parameters. *If*:

1. All subscribers' think times are *i.i.d* from an arbitrary distribution  $G_{tnk}$  with mean  $1/\lambda_B$
2. All subscribers' service times are *i.i.d* from an arbitrary distribution  $G_{ser}$  with mean  $1/\mu$

Then the probability all  $m$  vehicles are being used is given by [Kle75a, Ive09, Tij03]:

$$p(b|m) = \frac{\binom{S}{m} \rho^m}{\sum_{i=0}^m \binom{S}{i} \rho^i}, \quad \rho = \frac{1/\mu}{1/\lambda_B} = \frac{\lambda_B}{\mu} \quad (5.2)$$

$$= \frac{\binom{S}{m} \psi^m (1 - \psi)^{S-m}}{\sum_{i=0}^m \binom{S}{i} \psi^i (1 - \psi)^{S-i}}, \quad \psi = \frac{\rho}{1 + \rho} \quad (5.3)$$

Eq(5.2) is the *Engset distribution*. It is equivalent to Eq(5.3), which is a *truncated binomial distribution* [Tij03, Ive09]. The variable  $\psi$  is known as "offered traffic" and is a measure of how busy the queueing system is (closer to one is busier). Unfortunately, both Eq(5.2) and Eq(5.3) become numerically unstable as  $S$  and  $m$  grow because  $\binom{S}{m}$  becomes too large to compute. Instead, we calculate  $p(b|m)$  using an algebraically equivalent but stable recursive formulation [Ive09]:

$$\rho(b|m) = \begin{cases} \frac{\rho \cdot (S-m+1) \cdot \rho(b|m-1)}{m + \rho \cdot (S-m+1) \cdot \rho(b|m-1)} & \text{if } m > 0 \\ 1 & \text{if } m = 0. \end{cases} \quad (5.4)$$

$$\rho = \frac{1/\mu}{1/\lambda_B} \quad (5.5)$$

Note that the number of subscribers  $S$  is an input parameter. Moreover, these formulae give the blocking probability as a function of a given pool size  $m$ , so  $m$  is also an input parameter (an algorithm is given in §5.2.4 to find  $m$  such that  $\rho(b|m) < \epsilon$ ). Hence, the only non-input parameter is  $\rho$ , which is a function of the MST  $1/\mu$  and the busy period MTT  $1/\lambda_B$ .

There are two problems with directly applying this model:

1. We have heterogeneous subscribers who may have different MTTs and MSTs, so the assumption that subscribers' MTTs and MSTs are *i.i.d.* may not hold.
2. Even if they are *i.i.d.*, we expect the pool to have *busy periods* such that the aggregate stream of arrivals to the pool is non-stationary, and the ELM is only applicable to stationary queueing systems.

We describe how we deal with these problems in the following sections.

## Modeling the Average Subscriber

To deal with the problem that MTTs and MSTs, may not be *i.i.d.*, we model the heterogeneous population of subscribers as a homogeneous population of *average subscribers*. We assume the following data can be collected from each subscriber  $s = 1..S$ :

- $s$ 's *mean think time* (MTT)  $1/\lambda_s$  corresponding to their *mean think rate* (MTR)  $\lambda_s$ .
- $s$ 's *mean service time* (MST)  $1/\mu_s$  corresponding to their *mean service rate* (MSR)  $\mu_s$ .
- $\rho(s, B)$ , the probability  $s$  arrives given the pool is in a busy period.

This data can be collected from subscribers using customer surveys of the form "How many times per year will you need a vehicle?", "How long are you likely to need it for?", and "Will you arrive during any of *these* busy times?"

## Assumptions

We make seven assumptions:

1. A weighted average of all subscribers' MTTs is a good estimate of the mean of the arrival distribution  $1/\lambda_B$ . In the ELM model, only the mean of  $G_{tnk}$  (not the distribution of think times) is needed to compute  $\rho(b|m)$ , so our assumption that subscribers are identically distributed according to the think-time distribution  $G_{tnk}$  simply means this weighted average is a good estimate of  $1/\lambda_B$ .




Figure 5.1: Example arrival time distribution over one day to the carshare pool




Figure 5.2: Figure 5.1 zoomed on [8,9]

2. Because each subscriber has their own mobility patterns and vehicle needs, subscribers' think times are independent. When combined with the prior assumption, think times are *i.i.d* according to  $G_{tnk}$ .
3. Similarly, a weighted average of subscribers' MSTs is a good estimate of its mean  $1/\mu$ .
4. The duration for which subscribers need their vehicles and the times they return their vehicles are independent. Thus, service times are *i.i.d* according to  $G_{ser}$ .
5. Subscribers arrive at the pool at most once during each *busy period*. Since the busy period is a subset of each day, or a day/weekend during a year, we assume a subscriber will not rent more than once per busy period.
6. The loss probability  $p(b|m)$  is the same for all subscribers and is given by steady state results for the ELM given later in this section. This naturally follows from our methodology of modeling the average subscriber.
7. The MTT during the busiest period,  $1/\lambda_B$ , *always holds*. This assumption is discussed further in the following section.

## Non-Stationarity & Busy Period Sizing

We now deal with the issue that the aggregate stream of arrivals to a pool may be non-stationary. Suppose the probability subscriber  $s$  arrives to a pool at a particular time on a day is shown in Figure 5.1. If many subscribers followed the same distribution, clearly [8am,9am] would be a busy period for the pool compared to the rest of the day. Notice that within this hour, shown in the boxed region and enlarged in Figure 5.2, the distribution is approximately uniform—subscribers who arrive during this busy period are equally probable to arrive at any time. *Busy period modeling* allows us to exploit this [Rio51, Flo01] by assuming the pool is *always* in a busy period so that arrivals appear stationary. We do this by calculating the MTT  $1/\lambda_B$  by examining the system only during busy periods, as explained in the following section.

There are two busy period sizing methodologies in common use. First, one can size the system according to the historical single busiest period of a defined length. The electrical grid, for example, is often sized according to the highest electrical demand during any one hour ever recorded. A dataset is needed to use this approach, see §5.2.3. Alternatively,

one can size the system according to the average demand during recurring busy periods. For example, a teletraffic network may be sized according to the call rate during all 5-6pm periods. In this section, we estimate the MTT during the average busy period, rather than the historical single busiest period.

Within this sizing methodology, there are two further options. We can size the pool according to a daily busiest period of  $K$  hours or a yearly busiest period of  $K$  contiguous days. For both options, the probability subscriber  $s$  arrives given the pool is in a busy period  $B$ ,  $p(s, B)$ , is derived in the following section.

### Deriving the ELM Parameter $\rho$

We now derive the parameter  $\rho$  in Eq(5.5). We use the following notation:

- $A_B$  denotes the arrival rate of pool subscribers during  $B$ , one of the pools busy periods.
- $c_s = 1/\lambda_s + (1 - p(b|m))1/\mu_s$  denotes subscriber  $s$ 's *mean cycle time*, the mean time between  $s$ 's requests since on average they think for time  $1/\lambda_s$  and then receive service for time  $1/\mu_s$ , at which point they begin thinking again. With probability  $p(b|m)$   $s$  is blocked and does not receive service and immediately goes back to thinking.
- $\omega_s = 1/c_s$  represent  $s$ 's *cycle rate*, the "rate" at which they arrive, which represents how active a subscriber  $s$  is. It is also the probability that  $s$  arrives on any one day of the cycle.
- $n_B$  denotes the number of subscribers that arrive to the pool (including those blocked) during busy periods.
- $K$  is the length of busy periods, given as input.

Recall from the assumptions discussed in §5.2.2 that we assume a weighted average of all subscribers' MTTs is a good estimate of the mean of the arrival distribution  $1/\lambda_B$ , and similarly that a weighted average of subscribers' MSTs is a good estimate of its mean  $1/\mu$ . Our methodology is to use  $\omega_s$  as the weight for each subscriber  $s$ . The intuition is to give lower weight to data from subscribers who rarely use the service. Note that while  $\omega_s$  includes  $1/\mu_s$ , typically  $1/\lambda_s \gg 1/\mu_s$ , thus we are not discounting users with long service times. Future work could include studying other weighting functions (see §5.5).

We start by computing the MST  $1/\mu$  as a weighted average of subscribers' MSTs:

$$\frac{1}{\mu} = \frac{\sum_{i=1}^S \omega_i \frac{1}{\mu_i}}{\sum_{i=1}^S \omega_i} \quad (5.6)$$

We now derive the parameter  $1/\lambda_B$ . Consider the probability subscriber  $s$  arrives given the pool is in a busy period,  $p(s, B)$ . We first examine the case where there is a recurrent daily busy period<sup>2</sup>. We assume the pool owner breaks up a day into a set of arbitrary-length

---

<sup>2</sup>Which period of the day is busiest can be computed through customer surveys.

chunks, which we denote  $\mathcal{K}$ . We define the vector  $\mathbf{t}_s$  to be the probability that a given arrival by subscriber  $s$  occurs during each chunk of the day. We stress that  $\mathbf{t}_s$  does *not* give the probability subscriber  $s$  arrives during each period of each day. That is:

$$\mathbf{t}_s = p(\text{current time} \in \text{chunk}_1 | s \text{ arrives}), \dots, (\text{current time} \in \text{chunk}_{|\mathcal{K}|} | s \text{ arrives}) \quad (5.7)$$

$$\neq p(s \text{ arrives} | \text{current time} \in \text{chunk}_1), \dots, (s \text{ arrives} | \text{current time} \in \text{chunk}_{|\mathcal{K}|}) \quad (5.8)$$

This can be derived from a subscriber survey of the form "on a day you need a vehicle, how likely is it that you arrive during [chunk 1], ..., [chunk  $|\mathcal{K}|$ ]".

For the case where we are sizing the pool for a daily busy period, we assume the probability  $s$  arrives on any given day is uniform and given by  $\omega_s$ . We estimate the busiest period by finding the weighted average probability a subscriber makes a request during each period, and then finding the maximum of these probabilities over all periods:

$$B = \arg \max_{k \in \mathcal{K}} \left( \frac{\sum_{i=1}^S (\omega_i \mathbf{t}_i[k])}{\sum_{j=1}^S \omega_j} \right)$$

Then  $p(s, B)$ , the probability that  $s$  arrives *given* the pool is in a busy period, is given by:

$$p(s, B) = \omega_s * \mathbf{t}_s[B] \quad (5.9)$$

If we instead wish to size the pool according to the busiest contiguous block of days over the year, we assume the pool owner knows the period they wish to size for, e.g., "Mother's Day weekend". We assume the pool owner can determine this by asking its subscribers the question "out of the  $X$  times you expect to need a vehicle per year, how many times  $Y$  do you expect to be during [this busy period]" <sup>3</sup>? Then  $p(s, B) = Y/X$ .

Next, we derive  $n_B$ , the number of users expected to arrive to the pool during any given busy period. We weight an arrival by each subscriber  $s$  by  $p(s, B)$ :

$$n_B = \sum_{i=1}^S p(s, B) \omega_i \quad (5.10)$$

We use the following rationale to derive the mean think time,  $1/\lambda_B$ : *the effective arrival rate of subscribers to the pool during busy periods is equal to the number of subscribers that arrive during the busy period divided by the length of the busy period.* That is,

$$A_B = \frac{n_B}{K} \quad (5.11)$$

Suppose we know the MTT of users during the busy period,  $1/\lambda_B$ .  $A_B$  is also given by

$$A_B = S \left( \frac{1}{c_B} \right) = S \left( \frac{1}{\frac{1}{\lambda_B} + (1 - p(b|m)) \frac{1}{\mu}} \right) \quad (5.12)$$

---

<sup>3</sup>This can be generalized to the case where the pool implementor has knowledge that one of a few busy days or weekends will be the busiest, but needs to survey customers to determine which is actually the busiest.



That is, if  $1/c_B$  represents the *average cycle time of each individual subscriber during busy periods*, then we expect  $S/c_B$  to arrive each time unit.

Finally, by combining Equations (5.11) and (5.12), we get:

$$\frac{S}{\frac{1}{\lambda_B} + (1 - \rho(b|m))\frac{1}{\mu}} = \frac{n_B}{K} \quad (5.13)$$

And consequently:

$$\frac{1}{\lambda_B} = \frac{SK}{n_B} - (1 - \rho(b|m))\frac{1}{\mu} \quad (5.14)$$

Finally, the parameter  $\rho$  in Eq(5.5) is given by:

$$\rho = \frac{1/\mu}{1/\lambda_B} = \frac{Eq(5.6)}{Eq(5.14)} \quad (5.15)$$

### Blocking Probability Iteration Algorithm

Note that Eq(5.15) finds  $\rho$  by assuming a value for  $\rho(b|m)$  in Eq(5.14), and Eq(5.4) finds  $\rho(b|m)$  assuming a value for  $\rho$ . Thus, we iterate the two equations to find the values of these two parameters.

The iteration algorithm we use to compute  $\rho(b|m)$  for a fixed  $m$  to a desired accuracy  $\sigma$  is given in Algorithm 1. We stop when  $\rho(b|m)$  changes between two iterations by less than  $\sigma$ , our desired accuracy. This iteration gives  $\rho(b|m)$  for a fixed  $m$ ; an algorithm is given in §5.2.4 to optimize  $m$ .

---

#### Algorithm 1 Compute $\rho(b|m)$

---

Let  $\star^{(k)}$  represent the value of  $\star$  during the  $k$ th iteration

Inputs:  $S, m, \sigma$

- 1: set  $k = 0, \rho(b|m)^{(k)} = 0.5$
  - 2: **do**
  - 3:   Compute  $\rho^{(k)}$  using  $\rho(b|m)^{(k)}$  and Eq(5.15)
  - 4:   Compute  $\rho(b|m)^{(k+1)}$  using  $\rho^{(k)}$  and Eq(5.4)
  - 5:    $k++$
  - 6: **while** ( $|\rho(b|m)^{(k+1)} - \rho(b|m)^{(k)}| > \sigma$ )
- 

The pool operator can always guarantee convergence of this algorithm; proofs are deferred to §C. In rare cases, the exact criteria of which are given in §C.3, the operator may need to increase the sizing parameter  $K$ . The problem is that Eq(5.14) becomes *negative* when  $K$  is very small, which is invalid because the blocking probabilities given by Eq(5.2) and Eq(5.4) are not defined for negative values of  $1/\lambda_B$  ("negative arrival rates").

### 5.2.3 Dataset-based sizing

In this approach, the input to the sizing algorithm is a trace of past requests to an existing pool that contains the times at which each request is made, and the duration of each request. With this dataset, we can directly measure  $A_B$  and  $1/\mu$  and subsequently solve for  $1/\lambda_B$ . Suppose the granularity of the dataset is  $H$  hours, that is, suppose it is possible to measure the arrival rate during any period of  $H$  from the dataset. Let  $K = qH$ , where  $q$  is an integer (i.e.,  $K$  must be a multiple of  $H$ ) represent the length of the historical busiest period the operator wishes to size the pool for. Let  $A_i$  represent the arrival rate during each segment of length  $K$ . Suppose the dataset is divided into  $n$  segments. We use a sliding window approach to calculate the arrival rate during the busiest segment in the dataset as:

$$A_B = \max \left( A_0 = \frac{\sum_{i=0}^{q-1} H_i}{K}, \dots, A_j = \frac{\sum_{i=j}^{q+j-1} H_i}{K}, \dots \right) \quad (5.16)$$

It is straightforward to compute  $1/\mu$  from the dataset; we average the duration of all requests. Then, we compute  $1/\lambda_B$  from Eq(5.12) because  $A_B$  is known. Eq(5.12) is iterated with Eq(5.4) to obtain a pairing of  $p(b|m)$ ,  $1/\lambda_B$  for a given  $m$ .

We note that if the goal is to obtain the most conservative sizing by finding the highest MTR observed during a period (or conversely, the lowest MTT observed), the dataset should be examined at the finest granularity possible. That is, finding the MTR using chunks of a larger size will never be higher than the MTR found using chunks of a smaller size. *Proof.* Let  $T$  be an arbitrary period of time. Suppose we divide  $T$  into  $c$  equally sized smaller periods. Denote the MTR during period  $i$  as  $\lambda_i$ . We have:

$$\lambda_T = \frac{\sum_{i=1}^c \lambda_i}{c} \leq \frac{c(\max_{i=1..c} \lambda_i)}{c} \leq \max_{i=1..c} \lambda_i \quad \square$$

As with the previous approach, this iteration computes  $p(b|m)$  for a fixed  $m$ .

### 5.2.4 Optimizing The Pool Size

Sections §5.2.1, §5.2.2, and §5.2.3 have presented different methods of computing the loss probability  $p(b|m)$  for a fixed number of vehicles  $m$  depending on what data is available. We use Program 1 to size pools given these three sizing methodologies. Objective (5.20) minimizes the number of vehicles ( $m$ ) in the pool. Constraint (5.21) ensures the QoS target is met<sup>4</sup>, but Constraint (5.22) states  $m$  cannot exceed the maximum number of vehicles that can be stored in the pool. Thus, this program is infeasible if  $m_{max}$  is too small given the number of subscribers and their usage patterns.

Program 1, despite being an integer program, is  $\mathcal{O}(TS \log m^*)$ , where  $T$  is the time it takes to compute  $p(b|m)$  according to which of the three methods is being used, and  $m^*$

<sup>4</sup>In prediction; we can never guarantee the QoS is met because the training set may not be perfectly indicative of demand seen in the validation set, even with many bootstrapping iterations.

---

### Program 1: Pool Sizing Integer Program

---

Inputs:	$m_{max}$	(maximum pool size)	(5.17)
	$\epsilon$	(desired QoS)	(5.18)
Decision Variables:	$m \in \mathbb{N}$	(pool size)	(5.19)
Objective:	$\min m$		(5.20)
Subject To:	$\rho(b m) < \epsilon$		(5.21)
	$m \leq m_{max}$		(5.22)

---

#### Algorithm 2 :Solution to Program 1

---

```

1:  $m_L, m_H \leftarrow 1$ 
2: while  $\rho(b|m_H) > \epsilon$  // Phase 1: Bracketing
3:    $m_L = m_H$ 
4:    $m_H = 2 * m_H$ 
5: while 1 // Phase 2: Binary Search
6:    $m_M \leftarrow \frac{m_L + m_H}{2}$ 
7:   if  $\rho(b|m_M) \leq \epsilon$  then  $m_H \leftarrow m_M$  else  $m_L \leftarrow m_M$ 
8:   if  $m_H - m_L \leq 1$ 
9:     if  $m_H \leq k$  then return  $m_H$  else return "Infeasible"

```

---

is the optimal solution, using Algorithm 2 as follows. First, during the *bracketing phase*, we repeatedly double the pool size  $m_H$ , until  $\rho(b|m_H) \leq \epsilon$ . We also maintain the previous  $m$  value before this threshold is reached, denoted  $m_L$ . Thus, we find an interval  $(m_L, m_H]$  that is smaller than  $[1, S]$  in which  $m^*$  lies. This process is logarithmic in  $m^*$  because we grow  $m_H$  at an exponential rate, so during the bracketing phase, in the worst case we compute  $\rho(b|m) \log m^*$  times. When this procedure terminates,  $m_L < m^* \leq m_H$ . We then use binary search on the interval  $(m_L, m_H]$  to find  $m^*$ <sup>5</sup>. Binary search is logarithmic in the width of the interval, which is at most  $m^* - 1$ . The worst case is when  $m_H = m^* - 1$  and is doubled, giving us  $m_H - m_L = (2m^* - 2) - (m^* - 1) = m^* - 1$ . Therefore, we compute  $\rho(b|m)$  another  $\log m^*$  times during the search phase, and algorithm 2 is in  $\mathcal{O}(TS \log m^*)$ .

## 5.3 Evaluation Methodology

In this section, we describe our evaluation methodology. Henceforth, we abbreviate each of our sizing methodologies as follows: we refer to the binomial based method (§5.2.1) as BIN, the Engset Loss Model based on subscriber surveys (§5.2.2) as ELM<sub>S</sub>, and the Engset

---

<sup>5</sup> $\rho(b|m)$  is monotonically decreasing with increasing  $m$ , so we can create an array  $[\rho(b|m_L), \dots, \rho(b|m_H)]$  which is naturally sorted in decreasing order. We can then use binary search to find the minimum index such that  $\rho(b|m) < \epsilon$ . We do not need to compute all the values of this array however: we only need to work with the boundaries and midpoint at each step.

Loss Model which uses the dataset (§5.2.3) as  $ELM_D$ .

We use a dataset of carshare use as a proxy for demand for this BEV pooling service. Specifically, we evaluate our methodologies using two sets of experiments. In the first set, we assume demand for the pool exactly matches the demand for a single location carshare. In the second set, we consider only the demand for trips longer than 80km, assuming shorter trips were accommodated by the subscribers' personal BEVs.

This methodology has both disadvantages and advantages. For the former, carshares are generally priced to discourage against long trips in terms of time and distance. For example, if the pricing of an ICEV pool for BEV owners was the same as the carshare used for our evaluation, it may be cheaper for subscribers to use rental cars for long trips. Second, 80km is less than the range of today's BEVs, so these trips may not be a perfect proxy for demand by BEV owners. Unfortunately, there are not enough trips in our dataset over 160km to obtain statistically significant results. However, while this carshare reservation dataset is not ideal for studying demand primarily composed of long trips, it does allow us to evaluate our methods on real world non-stationary demand. The alternative methodology would be to use prior measurements of how often BEV owners are likely to need an ICEV, for example, based on the works of Pearre et al. [PKGE11] or other works discussed in §3.4. While this would correctly capture the demand of individual BEV owners, it would not capture the difficulty that much of our work attempts to overcome—non-stationary demand as observed by the pool. We show in this section that demand for the carshare used in our evaluation is very non-stationary—it has periods with demand five times higher than the average demand.

### 5.3.1 Dataset

We use data from a local car share, Community Carshare [Com14], to evaluate our sizing methodologies. The dataset is composed of all trips made between January 2005 and October 2013. There are 51,223 records in the dataset, each detailing the start time and end time of the trip, as well as the distance driven in km. Figure 5.3 shows the number of "active" subscribers over time; where we say a subscriber is active at time  $t$  if they made at least one reservation in the year prior to  $t$ . The number of subscribers steadily rose over the eight years from 50 in 2005 to over 700 in October 2013.

### 5.3.2 Performance Metrics

To evaluate our methodology, we use the same three metrics that car shares use to evaluate their operation:

- *Availability*: the percentage of served requests. This is our QoS metric. Data for availability is limited, but Shaheen suggests that carshares in the past aimed for an availability of 95% [SS98].




Figure 5.3: Subscribers active in the carshare over time

- *Member-to-vehicle ratio (M2V)*: the number of subscribers for each vehicle in the pool. The M2V metric for carshares has steadily, but not monotonically increased from 15-20 in 1998 to nearly 50 in 2012 [SS98, SSW04, SCC09, You13, Bro04].
- *Utilization*: the percentage of time vehicles are used. Carshares generally have utilizations in the range of 20-40% [SM07, LB06, Bro09, BHT11]. Rental operations aim for nearly double this because they have a much larger number of potential renters (they are “infinite population” systems) [Rol11, Inv, Ten10, Bro12, Rul13, Bro09].

The last two metrics are inversely proportional to the first—they decrease with larger pool sizes but availability increases with larger pool sizes. Both existing car shares and rental companies have traditionally aimed to maximize utilization at the expense of availability, because profit is directly proportional to utilization. For the BEV pool application however, availability may be essential—the service may only alleviate range anxiety if subscribers almost always receive a vehicle when they need one.

### 5.3.3 Bootstrapping

We use *random bootstrapping with replacement* [Efr79] and *cross validation* to evaluate our sizing methodologies as follows. We perform  $I$  bootstrap iterations. In each iteration, half of the dataset is selected at random. This half, known as the training set, is used to compute parameters needed for the sizing methodologies. The other half, known as the test set, is used to evaluate the performance metrics (M2V, utilization, availability) for each sizing methodology. After each iteration, the entire dataset is sampled again in the same fashion (hence with replacement). Thus, we obtain a total of  $I$  samples of the three performance metrics. Bootstrapping treats each performance metric as a sampling statistic. We therefore obtain a total of  $I$  values for each of these sampling statistics and thus are able to compute confidence intervals.

---

**Algorithm 3** : Replay Arrival Departure Stream

---

```
1: // out array tracks the # of vehicles out of the pool
2: blocked_count, out[0] = 0
3: for i in range [0, len(ADS))
4:   if ADS[i] == "Arr" and out[i - 1] < m // not blocked
5:     out[i] = out[i - 1] + 1
6:   else if ADS[i] == "Arr" // blocked
7:     out[i] = m, blocked_count ++
8:   else out[i] = max(0, out[i - 1] - 1) // departure
```

---

### 5.3.4 Parameter Values

The parameters  $\epsilon$ ,  $K$ ,  $|\mathcal{K}|$  (the QoS target, the number of daily chunks for  $\text{ELM}_S$ , and the length of each chunk) are given as input to our methods. The rest of the parameters are computed from the training set as follows. During each sampling iteration, for all methods, we compute  $S$ , the number of subscribers, and we compute  $m$  such that  $p(b|m) < \epsilon$ . For BIN we need to compute  $p(a)$ , the mean probability a subscriber needs a vehicle on a given day. Using the terminology defined in §5.2.2, the probability subscriber  $s$  needs a vehicle on a given day is  $\omega_s$ , which represents " $s$  arrives once every  $c_s$  days", or alternatively, "the probability  $s$  arrives on a given day is  $\omega_s$ ", hence, we compute  $p(a) = \sum_{s=1}^S \omega_s / S$ . For  $\text{ELM}_S$ , each subscriber's MST and MTT are computed from the database. Then  $n_B$ ,  $1/\mu$ , and  $1/\lambda_B$  are computed using the methodology in §5.2.2. That is, we pretend that customers respond to surveys perfectly accurately, giving answers that we compute from the dataset. For  $\text{ELM}_D$ , we first compute  $A_B$  using a sliding window of size  $K$ , then,  $1/\mu$  and  $1/\lambda_B$ .

### 5.3.5 Performance Metric Computation

We now describe how we compute the performance metrics. During each iteration, the test set is used to compute a stream of arrivals and departures, referred to as an arrival departure stream (ADS). We first form two lists containing the sorted times of all arrivals (start of vehicle use) and departures (end of vehicle use) seen in the test set. These two lists are then merged into one stream. When there is a tie—one subscriber arrives while another is returning a vehicle—we conservatively assume the arrival comes first and cannot be served by the incoming vehicle.

For BIN,  $\text{ELM}_S$ , and  $\text{ELM}_D$ , the pool size  $m$  is computed during the training phase. During each iteration, we replay the ADS for these three methodologies using Algorithm 3. Availability is computed as  $1 - \text{blocked}/\text{len}(\text{ADS})$ . Let  $\text{out}[i]$  denote the number of vehicles in use at time  $i$ , and let  $\Delta_i$  denote the amount of time between arrivals  $i - 1$  and  $i$ . We compute utilization as:

$$\frac{\sum_i \text{out}[i] * \Delta[i]}{m * (\text{out}[\text{end of dataset}] - \text{out}[0])}$$

The numerator is the Riemann sum giving the total number of “utilized vehicle-hours” and the denominator is the maximum number of vehicle-hours possible assuming every vehicle was always in use, so the fraction gives the percentage of time each vehicle is used on average. Finally, we annotate each trip record in the dataset with the number of active subscribers at that time. This allows us to compute the average number of pool subscribers  $S$ . We then compute M2V as  $S/m$ .

## 5.4 Evaluation Results

We now discuss our evaluation results. We first use statistical tests to show the carshare used for our evaluation has non-Markovian, non-stationary arrivals in §5.4.1. We then show the performance metrics achieved by several sizing configurations in §5.4.2.

### 5.4.1 Arrivals are Non-Markovian and Non-Stationary

We demonstrate that the demand patterns of our carshare are non-Markovian and non-stationary. Figure 5.7 shows the number of arrivals as a function of the time of day, and the number of arrivals over the last 300 days in our dataset. They clearly show demand both within the average day and over days of the year are non-stationary. Figure 5.4 shows probability distribution over the number of vehicles in use, which demonstrates that the pool experiences rare but large peaks in demand (another demonstration of non-stationarity). We see that the average number of vehicles in use at any time is approximately three, but at times the number out is much larger, e.g., the peak of nearly 30 is 10 times larger.

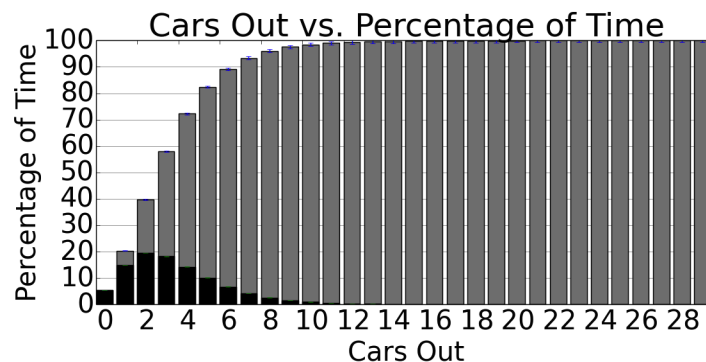


Figure 5.4: The black distribution shows the probability of a specific number of vehicles in use. The grey distribution shows the CDF.

We further tested a null hypothesis *that the carshare subscribers’ interarrival times are exponentially distributed*. The parameter of this distribution is obtained by using maximum likelihood estimation [Was04]. We tested the same hypothesis for the service time distribution. Figure 5.5 shows the distribution of interarrival times split in half hour bins (0-30min, 30-60min, etc.), and service times using one hour bins. The null hypotheses that the interarrival time and service time distributions are exponentially distributed were both rejected

at the 99% confidence level using  $\chi^2$  tests. This represents strong evidence against the null hypotheses. We also tested the same results by only including trips greater than 80km assuming that all trips under 80km. The null hypotheses were again rejected at 99% confidence level. The distributions for this case are shown in Figure 5.6. Hence, there is strong statistical evidence that a Markovian model is inappropriate.

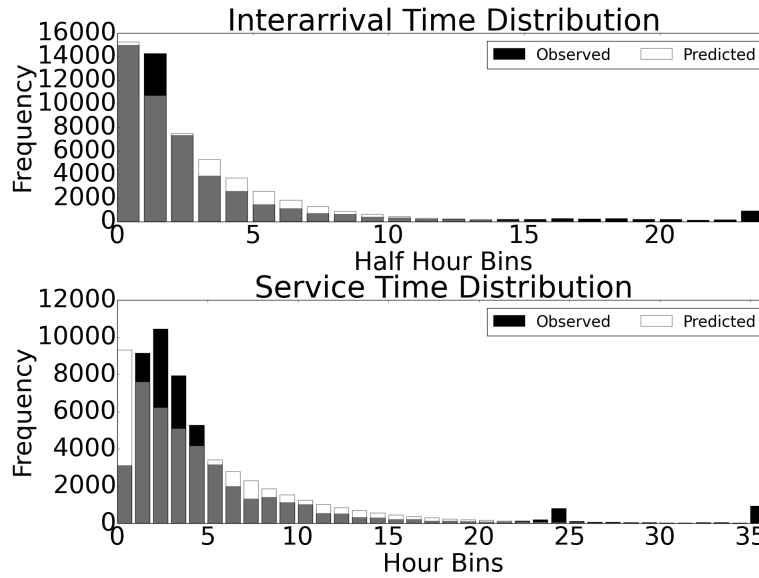


Figure 5.5: Interarrival time (top) and service time distribution (bottom) for all trips, observed vs. exponential using MLE. The grey region shows where the observed and predicted curves overlap. The black region shows observed arrivals, and the white region shows arrivals as predicted by the exponential distribution.




Figure 5.6: Interarrival time (top) and service time distribution (bottom) for trips >80km, observed vs. exponential using MLE



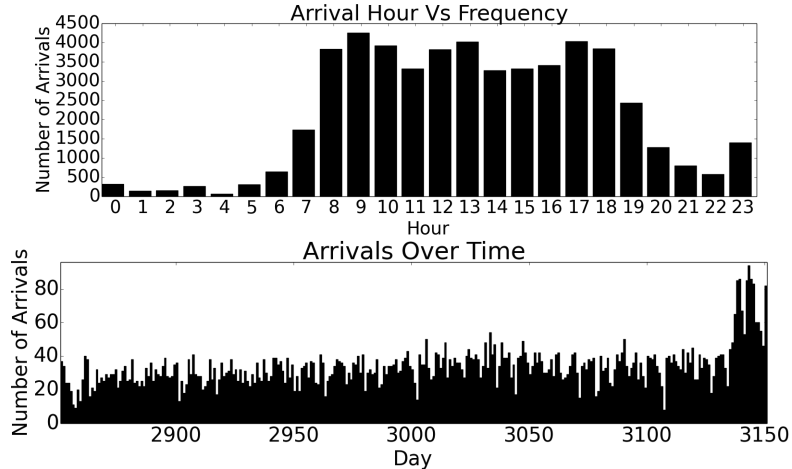


Figure 5.7: Top: Arrival hour vs. frequency for all trips in our dataset. Bottom: daily number of arrivals for the last 300 days (10 months) in our dataset. The x-axis shows the day number starting from the first day in our dataset in March 2005; the region shown is approximately all 2013 data.

## 5.4.2 Performance

Here we present the performance achieved by our sizing methods. We first discuss the reason for variability in the performance metrics (leading to the inclusion of error bars). As discussed in §5.3.3, we sample the M2V, utilization, and availability distributions  $l$  times to compute confidence intervals; for these results all confidence intervals are normal 99% intervals over  $l = 100$  iterations. In each iteration, the subscribers seen in the test set differs, which causes the  $m$  as determined by the sizing methods to vary. Moreover,  $S$  varies with each test set. These two factors contribute to the variance in M2V, shown as horizontal error bars. During each iteration, the ADS also differs, which leads to different computations of availability and utilization because the Riemann sums change. The computations of utilization and availability also vary as the pool size varies, reflected in the vertical error bars.

We graph each metric’s performance as a point in the 3D metric space formed by M2V, availability, and utilization. Figure 5.8 shows performance metrics for BIN and several configurations of  $ELM_S$  and  $ELM_D$ . The notation  $D(\cdot)$  indicates  $\cdot$  is the length of the historical single busiest period the pool is sized for using  $ELM_D$ . For  $ELM_S$ ,  $S(\cdot)$  indicates that  $\cdot$  is the set of daily chunks used; each number gives the start time of a new chunk. For example,  $S(0, 8, 18)$  indicates three daily chunks: [12am-7:59am], [8am-5:59pm], [6pm-11:59pm]. The QoS  $\epsilon = .05$ , corresponding to a busy period availability target of 95%. We make several observations about the results:

1. With respect to a 95% availability target,  $ELM_D(4368)$ , which sizes for the busiest 6 months observed in the training dataset over the eight years, and  $ELM_S(0,9,17)$ , which sizes for the busy period of 9am—5pm, work best.  $ELM_D(4368)$  slightly oversizes, obtaining  $A \approx 95\%$ , and  $ELM_S(0,9,17)$  slightly undersizes, obtaining  $A \approx 94\%$ . These are within one percent of the QoS target, and have a utilization at the high end of those reported in the literature.
2. BIN is too conservative. Giving out customer surveys (to size using  $ELM_S$ ) or record-

ing pool demand over time (for using  $ELM_D$ ) leads to large multiplexing benefits. BIN achieves a M2V of only five (one car per every five subscribers). Even  $ELM_D(8)$ , which conservatively sizes for the busiest eight hours seen in the dataset, outperforms BIN.

3.  $ELM_S(0)$  sizes the pool assuming that there is only one daily chunk starting at midnight that lasts 24 hours, thus the pool is sized according to the average day<sup>6</sup>. We see that sizing this carshare for the typical day does not work well. This is a result of the pool demand being non-stationary. Sizing for the average day leads to very low availability during busy periods, and consequently lower than desired overall availability. We hypothesize  $ELM_S(0)$  would work well for more stationary pools where demand does not fluctuate as much.  $ELM_S(0,12)$  exhibits the same problem, but to a lesser extent.
4. As proven in §5.2.3, sizing using  $ELM_D(X)$  v.s.  $ELM_D(Y)$  where  $X < Y$  ensures that  $m|X \geq m|Y$ . That is, availability is non-increasing as the busy period size increases. Figure 5.8 shows that availability monotonically decreases as the size of the historical busy period increases, but we later show a result where increasing the busy period size leads to the same sizing.
5. Utilization is linearly inversely proportional to  $m$ ; as  $m$  decreases and M2V increases, utilization increases linearly. However, a standard queueing result is that availability is *not* a linear function of the number of servers (cars in this case). As  $m$  decreases and M2V increases, availability decreases at an increasing rate. This can be seen with the points BIN,  $ELM_D(2160)$ , and  $ELM_S(0)$ ; when the M2V increases from 5 to 35, there is a 1.5% decrease in availability, but when M2V increases from 35 to 65, there is a 10% decrease in availability. The inverse of this result is intuitive; as more cars are added to the system, there are diminishing returns with respect to increasing availability.

As a sensitivity analysis to the parameter  $\epsilon$ , we show the same results for a QoS target of 90% availability in Figure 5.9. The trends from Figure 5.8 and the ordering of methods by their achieved metrics are preserved, except that  $ELM_D(4368)$  achieves the same performance as  $ELM_D(2160)$ .

---

<sup>6</sup>the probability each subscriber arrives during any daily busy period is the same as their probability of arriving on any day. Specifically,  $\mathbf{t}_s[B]$  in Eq(5.9) is 1 (the whole day is the busy period), hence each subscriber  $s$ 's busy hour weight  $p(s, B) = \omega_s$ , their probability of arrival on any day.




Figure 5.8: Comparison of achieved performance metrics for the full dataset,  $\epsilon = .05$ .  $ELM_S$  is abbreviated "S" and  $ELM_D$  is abbreviated "D".

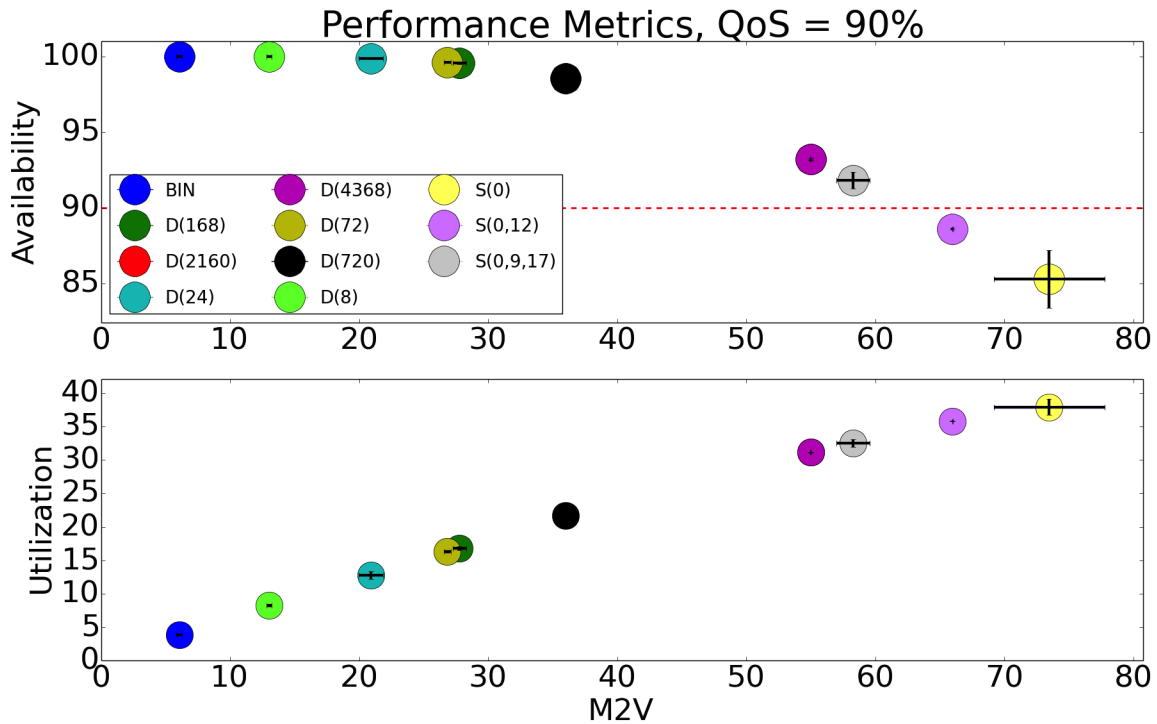


Figure 5.9: Comparison of achieved performance metrics for the full dataset,  $\epsilon = .1$ .  $ELM_D(2160)$  is drawn under  $ELM_D(4368)$  because they achieve the same performance.  $ELM_S$  is abbreviated "S" and  $ELM_D$  is abbreviated "D".

Figure 5.10 shows results after removing all trips under 80km from the database. Only a small portion of the dataset (hence, demand) (14.5%) remains after shorter trips are removed. Some parameter configurations are removed as they are too conservative for this case and comparable to BIN, e.g.,  $ELM_D(8)$ . Some of the trends are different than Figure 5.8 but these changes are intuitive. For a given pool size of  $m$  there is a much higher overall availability, compared to the previous case, because overall demand is lower, the gaps between demand are necessarily larger, and the number of vehicles out at any time is lower. When multiple "long trip" users arrive on the same day, the probability their arrivals and departures will be staggered such that they are able to share the same vehicle is lower (longer trips require more usage time). Because there is less multiplexing,  $ELM_S$  and  $ELM_D$  pool sizes are larger for the same parameter configurations. Because this leads to higher availability as discussed, M2V and utilization are lower in all ELM configurations shown.

As discussed, the methodology of removing all trips under 80km from this carshare, which is priced against longer trips, may not accurately simulate demand for a pool of ICEVs for BEV owners because actual demand may include many more longer trips. With more demand, we hypothesize the results to be similar to Figure 5.8. Obtaining another dataset would allow us to better study this usage case.

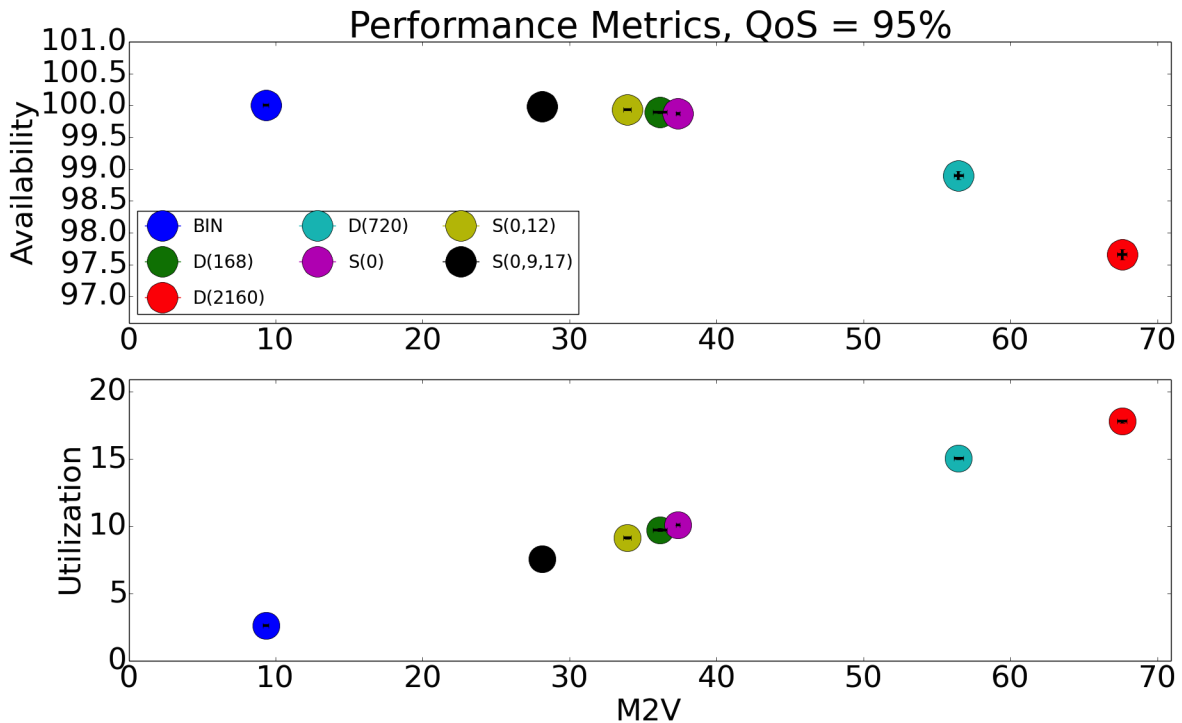


Figure 5.10: Comparison of achieved performance metrics for trips over 80km only,  $\epsilon = .05$ .  $ELM_S$  is abbreviated "S" and  $ELM_D$  is abbreviated "D".

## 5.5 Conclusions & Future Work

We study sizing finite population vehicle pools with non-stationary demand. Vehicle demand varies with time of day and time of year, and moreover many pool applications have a finite population, hence there is a need for such sizing methodologies. We propose three such methods to size to meet a QoS  $\epsilon$ . We further show that each of our methods has advantages and disadvantages with respect to various performance metrics.

Regarding our main application of interest, we note there are several advantages for potential BEV buyers and BEV dealerships of offering this pooling service. First, integrating this service into dealerships significantly reduces subscribers' *transactional cost*. Dealerships could collect all subscribers' information at the time of purchase so subscribers need not fill out paperwork each time they obtain an ICV. Moreover, because the subscription is offered by the dealership, subscribers would not have to compare the price of several rental vehicle agencies to determine the cheapest option. Second, dealerships can internally compute the cost of this service and amortize this cost into the price of their BEVs. The service can then be sold as "free" to potential customers, which would appeal from a marketing perspective. Finally, dealerships already maintain a pool of vehicles for customers awaiting repairs for their vehicles, thus our approach does not impose any radical changes to current practices.

We suggest three avenues to extend our work:

- Let  $\Gamma_s : \mathbb{R} \rightarrow \mathbb{R}$  be a function that denotes subscriber  $s$ 's disutility derived from a specific value of  $p(b|m)$ , the QoS of the pool. Subscribers' disutility of driving to an empty pool may increase nonlinearly, perhaps exponentially, as  $p(b|m)$  increases. Consider modifying Program 1 by removing Constraint (5.21) and modifying Objective (5.20) as

$$\arg \min_m \left( \sum_{s=1}^S \Gamma_s(p(b|m)) \right)$$

This formulation minimizes subscriber *risk* instead of forcing an explicit QoS  $\epsilon$ .

- We have not taken pricing into account, however, several authors have studied the economics of similar finite population systems, including parking lots [AdP04, Jan10, VNR12, LS10] and car share subscriptions [HG10, SCR06, MbMS<sup>+</sup>05, SM07].
- In §5.2.2, we weight each subscriber  $s$  by their cycle rate  $\omega_s$  when calculating  $1/\lambda_B$  and  $1/\mu$ . Future work could evaluate other weighting functions.

# Chapter 6

## A Multi-period Multi-pool Carsharing Model For Reducing Range Anxiety

---

### 6.1 Synopsis

This chapter is a generalization of the previous chapter. In the previous chapter, we statically sized a single vehicle pool for use by a finite population of users ("subscribers"). We now focus on dynamic fleet management—the repeated addition, removal, and movement of vehicles—of a multi-pool network. This fleet management is repeated on a periodic basis, e.g., every two weeks. We study the generalization to the multi period (dynamic), multi pool case for two reasons:

1. Instead of using averages of subscribers' information as done in static sizing, our methods now respond to changes in demand.
2. The model presented in the prior chapter is only applicable to single pool scenarios, e.g., in the case that an EV dealership allows their customers to use their ICEVs with the purchase of an EV. However, a multi pool network is more convenient for customers (more options to choose from), so it may be desirable to operate a multi location system independent of dealerships (e.g., as a carshare).

Although motivated by range anxiety, our work is general, in that it deals with all multi-pool vehicle systems. A network of ICEV pools for BEV owners is isomorphic to a carshare, so we use the term carshare henceforth. Prior methods for managing vehicle sharing systems make simplifying assumptions about demand patterns, including infinite population<sup>1</sup>, Markovian

---

<sup>1</sup>In infinite population models, demand requests can be generated from a potentially "infinite" or very large population of users. In finite population models, requests can only be generated from a known set of sources, in our case the pools' subscribers.

interarrival and service times, and a stationary arrival process. The application of their results is thus restricted to carshares with these demand patterns.

In the proposed system, we expect BEV owners to infrequently need access to a vehicle for a long trip, a different demand distribution than traditional carshares which are sized for more frequent trips of shorter duration. We therefore study the management of a carshare for *arbitrary* demand profiles, so that it is widely applicable.

Our main contribution in this chapter is as follows. We propose two sizing methods that, in conjunction with a repeated optimization program, dynamically manage a carshare fleet to meet any desired vehicle availability target. In contrast to prior work, our methods place no restrictions on vehicle demand patterns and are designed to serve non-stationary demand. Although motivated by range anxiety, our work is applicable to all multi-pool carshares.

We simulate our sizing methods using eight years of data from a local carshare. We show our methods perform well for any given vehicle availability target on many different workloads generated from the data.

We formulate our problem in §6.2. Two sizing methods are discussed in §6.3. An optimization program is formulated in §6.4 to compute the addition/movement of vehicles to/within the system. We describe our evaluation methodology and results in §6.5 and §6.6 respectively. We extend our model to allow for the systematic removal of vehicles, and attempt to simulate the usage case of an ICEV pool for BEV owners in §6.7. We conclude and discuss extensions to our model in §6.9.

## 6.2 Problem Formulation

Here we define our system architecture, state our assumptions, and overview our methodology.

### 6.2.1 System Architecture

The carshare under consideration operates as follows. The carshare has a set of pools of vehicles at locations  $\mathcal{J}$ . Subscribers of the carshare arrive to their preferred pool when they need a vehicle, and if one is available, they use it and return the vehicle to the same location afterwards<sup>2</sup>. The fraction of requests that are *blocked*—unserved because no vehicles are available—is referred to as the *blocking probability*. Let  $\gamma_j$  be the blocking probability at pool  $j \in \mathcal{J}$ . The objective of the carshare is to ensure  $\gamma_j < \epsilon \forall j$ , where  $1 - \epsilon$  is our *quality of service (QoS) target*. Ensuring  $\gamma_j < \epsilon \forall j$  requires vehicles to be periodically added to or moved between locations in the carshare. If demand grows at some locations but decreases at others, cars may be moved between locations. Additionally, if demand grows network wide, cars can be added to the carshare. The addition and movement of vehicles

---

<sup>2</sup>We discuss extending our model to allow subscribers to return their vehicle to different locations in §6.9.1

is henceforth referred to as *fleet management*. We assume fleet management takes place at periodic intervals. We denote the number of cars to move from pool  $j$  to  $k$  at the start of period  $p$  as  $v_{j,k}^p$ , and denote the number of cars to buy for  $j$  at  $p$  as  $a_j^p$ . Our goal is, for every period  $p$ , to correctly set  $v_{j,k}^p \forall (j, k) \in (\mathcal{J} \times \mathcal{J})$  and  $a_j^p \forall j \in \mathcal{J}$  to ensure  $\gamma_j^p < \epsilon \forall j$ .

Sizing and managing a carshare such that a systemwide QoS is met is challenging because demand patterns, such as the average time between requests, the average length of requests, etc, can be arbitrary and are constantly changing as subscribers join and leave. Moreover, for our application of interest, a carshare for EV owners, we expect that subscribers must be guaranteed a vehicle when they need it with high probability for the program to be successful. We are unaware of any prior work that is able to size a carshare for arbitrary demand patterns to meet a systemwide QoS target. Prior models on sizing and managing fleets/carshares are discussed in §6.8.

## 6.2.2 Assumptions

Our work makes the following assumptions:

- The pools are located in a relatively small, city-sized region. We are not considering a nation-wide system where it may be impractical to move vehicles between pools.
- Each carshare subscriber prefers one pool (pool selection is discussed in §6.5.1) and all of their trips begin and end at that pool. Our model can be extended to remove these assumptions as shown in §6.9.1 and §6.9.2.
- The size of each pool is set at the beginning of each period  $p$  and does not change until  $p+1$ . The length of each period can be made arbitrarily small, thus the fleet management can be repeated as frequently as desired.
- We initially assume that removals and replacements of cars are handled exogenously by the carshare operator because the decision of when to sell vehicles is complex. We remove this assumption in §6.7, i.e., we present fleet management that allows the addition, movement, and removal of vehicles, but also show it may not be preferable to do so.
- Subscribers who are blocked leave the pool immediately—there is no queueing.
- In our case study, we assume that it is cheaper to move cars between pools than to purchase new vehicles. However, the cost to move and purchase vehicles are given as input to our optimization program, so this assumption can be removed if desired.

## 6.2.3 Methodology Overview

Here we present our formulation of this problem. We consider a single period myopic optimization framework, that is, sizing decisions are made by considering only one future time period—this decision is further discussed in §6.9.5. Under this framework, at the start of each period  $p$ , we execute four phases:



1. The minimum size necessary to meet the desired QoS for period  $p$  at each pool is predicted (§6.3).
2. The fleet management is computed (§6.4).
3. The carshare then operates for the current period  $p$ . Afterwards, various performance metrics are computed.
4. Finally, data needed for step (1) for the following period are measured.

## 6.3 Pool Sizing Methods

In this section we propose two sizing methods for sizing each pool at the start of period  $p$ . The first method (§6.3.1) is an extension of  $ELM_D$ , the dataset-based variant of the Engset Loss Model we proposed in the prior chapter (see §5.2.3). The second method (§6.3.2) is inspired by the TCP *transmission control protocol* used in the Internet.

Note that, because sizing is computed at the start of period  $p$  based on past demand and expected future demand, we can never ensure the QoS will be met for  $p$  because future demand can be arbitrary. Nevertheless, we find our solution is quite robust, in that, despite not knowing the future demand exactly, the QoS target is nearly always met.

We do not consider the survey variant of the ELM,  $ELM_S$ , which we used in Chapter §5 for static sizing. While surveys can be distributed to customers infrequently at little inconvenience (the prior chapter's static sizing requires just one survey), we hypothesize subscribers would not want to take a survey at the beginning of every period  $p$  if periods are relatively short. Shorter periods are preferable for best results; long periods result in "lag times" between observed changes in demand and pool resizing because resizing is only done at the beginning of each period. For example, we use a period duration of two weeks.

### 6.3.1 Engset Loss Model

Here we extend the  $ELM_D$  method proposed in §5.2.3 to the multi pool case. We assume the reader is familiar with the derivations in §5.2.2 and §5.2.3 as they are not reproduced here. To extend ELM to size a multiple location carshare, we make the following changes:

- The population of pools may change over time, so  $S$ , the set of subscribers, is indexed by the pool number  $j$  and the period  $p$  as  $S_j^p$ .
- We size each pool for its' busiest period. We calculate the mean think time (MTT) of  $j$  during  $p$ , denoted  $1/\lambda_j^p$ , by considering only the busiest  $K$  consecutive hours of  $j$  during  $p$ .  $K$  itself is given as input and is assumed to be the same for all pools. The set of parameters  $1/\lambda_j^p \forall j \in \mathcal{J}$  replaces our old notation for the busy period MTT of a single pool,  $1/\lambda_B$ . We assume the busy period MTTs  $1/\lambda_j^p \forall j \in \mathcal{J}$ , hold throughout  $p$ .

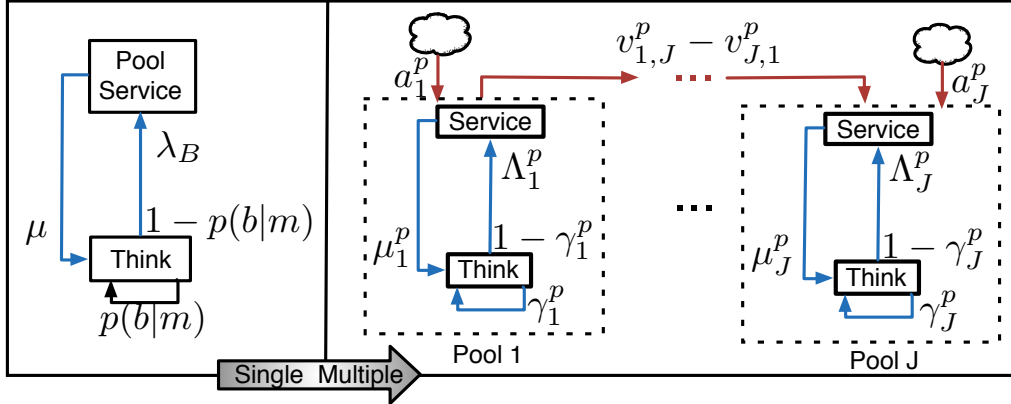


Figure 6.1: New carshare architecture snapshot during period  $p$

- We compute the mean service time (MST) of each pool separately. We compute the MST of  $j$  during  $p$  as the average of all request durations from  $j$  during  $p$ . The set of parameters  $1/\mu_j^p \forall j \in \mathcal{J}$  replaces our old notation for the MST of a single pool,  $1/\mu$ . We assume the MSTs  $1/\mu_j^p \forall j \in \mathcal{J}$ , hold throughout  $p$ .
- The parameters needed for  $j$  during  $p$ ,  $S_j^p$ ,  $1/\lambda_j^p$ , and  $1/\mu_j^p$ , are unknown at the start of  $p$ . Subscribers may join and leave during  $p$  and can have arbitrary demand patterns during  $p$ . However, we predict these parameters for  $p$  based on their past observed values. We describe various predictors in §6.5.4. We use the notation  $\tilde{\cdot}$  to denote a prediction of  $\cdot$ .
- The blocking probability is computed for each pool. Let the predicted values for  $j$  during  $p$  be  $\tilde{S}_j^p$ ,  $1/\tilde{\lambda}_j^p$ , and  $1/\tilde{\mu}_j^p$ . With a pool size  $m_j^p$ , the predicted blocking probability of  $j$  during  $p$  is given by (extended from Eq 5.3 and Eq 5.4):

$$\tilde{\gamma}_j^p(m_j^p) = \begin{cases} \frac{\psi_j^p}{m_j^p + \psi_j^p} & \text{if } m_j^p > 0 \\ 1 & \text{if } m_j^p = 0. \end{cases} \quad (6.1)$$

$$\psi_j^p = \left( \frac{1/\tilde{\mu}_j^p}{1/\tilde{\lambda}_j^p} \right) \cdot \left( \tilde{S}_j^p - m_j^p + 1 \right) \cdot \left( \tilde{\gamma} (m_j^p - 1) \right) \quad (6.2)$$

The set of predicted blocking probabilities for  $p$  at the start of  $p$ ,  $\tilde{\gamma}_j^p(m_j^p) \forall j \in \mathcal{J}$ , and the set of observed blocking probabilities for  $p$  observed after  $p$ ,  $\gamma_j^p(m_j^p) \forall j \in \mathcal{J}$ , replace our old notation of the observed blocking probabilities of a single pool  $p(b|m)$ .

- Algorithm 2 is again used to compute the predicted minimum size of  $j$  during  $p$ , denoted  $\tilde{m}_{opt_j}^p$ , such that  $\tilde{\gamma}_j^p < \epsilon \forall j$ . The set of pool sizes for  $p$  which are predicted to satisfy the QoS at each pool,  $\tilde{m}_{opt_j}^p \forall j \in \mathcal{J}$ , replaces our old notation of  $m^*$ .

We denote this method henceforth as ELM(K), where  $K$  is the busy period length in hours. The  $D$  subscript is dropped because there is no longer ambiguity between two ELM methods.

With this terminology defined, Figure 6.1 shows the model we consider in the prior chapter on the left and the queueing model we consider in this chapter on the right.

### 6.3.2 Additive Decrease Multiplicative Increase (ADMI)

Here we present a sizing method that does not rely on historical data and does not make assumptions about subscriber behavior. This method differs from ELM in that it does not attempt to guarantee a QoS. The method is motivated by the algorithm used for TCP congestion control [KR09, CJ89, YL00] in the Internet. In TCP congestion control, a sender begins with a small *window size*, defined as the number of packets it can send without receiving an acknowledgement. The sender increases its window size by  $\alpha \in \mathbb{N}$  repeatedly ("additive increase") until it is too large and a packet loss is observed due to network congestion. At this point, their window size is multiplied by a fraction  $\beta \in (0, 1)$  ("multiplicative decrease") and the process repeats<sup>3</sup>. This results in a sawtooth curve that converges to the optimal size if  $\alpha$  and  $\beta$  obey<sup>4</sup>[YL00, CJ89]:

$$0 < \alpha < \text{pool size}, \alpha \in \mathbb{N} \quad (6.3)$$

$$0 < \beta < 1 \quad (6.4)$$

This algorithm is an *additive increase multiplicative decrease* algorithm. For pool sizing, we take the opposite approach and propose an *additive decrease multiplicative increase* algorithm, which we abbreviate by ADMI. This algorithm works as follows. We maintain a "virtual pool size" variable, denoted by  $VPS_j$ .  $VPS_j$  is originally set to some initial guess, denoted  $ADMI_{in}$ , and changes as follows. Whenever there are  $y$  unblocked arrivals in a row, the virtual pool size is decreased by an additive constant  $\alpha \in [1, VPS_j)$ , that is,  $VPS_j = VPS_j - \alpha$ . Whenever an arrival is blocked, the counter for  $y$  is set to zero, and the pool size is increased by a factor  $\beta$ , that is,  $VPS_j = \frac{VPS_j}{\beta}$ . This algorithm decreases the pool size slowly over time when it is too large, and reacts aggressively to blocked arrivals. The parameters  $y$  and  $\alpha$  controls how aggressive the algorithm is in keeping a low pool size, and  $\beta$  controls how aggressively the pool size increases upon a blocked arrival. As long as  $\alpha$  is bounded to be in  $[1, VPS_j)$  and  $\beta \in (0, 1)$ , the size of the pool will oscillate around the optimal pool size even if this changes over time due to non-stationary arrivals [YL00, CJ89].

We now describe the relationship between the tracking variable  $VPS_j$  and the predicted optimal pool size  $\tilde{m}_{opt_j}^p$ . Recall that we only set the size of each pool at the start of each period. We use the heuristic:

$$\tilde{m}_{opt_j}^p = \text{avg}(VPS_j^{[p-1, p)})$$

The average value of  $VPS_j$  over some window is likely to be close to the optimal pool size in that window. This heuristic thus "lags" by assuming what was likely optimal for period  $p - 1$  will be optimal for period  $p$ , which works well when pool demand does not change too rapidly. Several other heuristics could be implemented, for example we could set  $\tilde{m}_{opt_j}^p$  to the 70th percentile of  $VPS_j^{[p-1, p)}$  to be more conservative. Note that  $y$  controls the tradeoff between the pool size and the pool availability. If  $y$  is small, the virtual pool size is

<sup>3</sup>In TCP Reno congestion control,  $\alpha = 1, \beta = .5$ .

<sup>4</sup>This is true if all senders use the same values for  $\alpha$  and  $\beta$  [YL00, CJ89].

reduced quickly, which will lead to a lower average of  $VPS_j$  and consequently a lower  $m_j^p$ . If  $y$  is large, the virtual pool is left in an "oversized" state for longer, which will increase the average of  $VPS_j$  and consequently  $m_j^p$ .

Although ADMI does not provide probabilistic bounds on QoS, we show in §6.6 that our heuristic works well in practice.

### 6.3.3 Offline Optimal Baseline

We also compute the optimal offline sizing to compare our metrics with. After each period  $p$ , for all pools  $j$ , we compute the optimal number of vehicles that should have been stored in  $j$  during  $p$ ,  $m_j^p$ , as:

$$\min m_j^p \text{ subject to } \gamma_j \leq \epsilon \quad \forall j$$

The M2V and utilization metrics improve with a smaller pool size, so the minimum size of each pool such that the QoS is met is optimal with respect to all performance metrics. In practice, it is impossible to achieve these values because this requires perfect knowledge of demand and that all excess cars be removed at the beginning of every period. These values hence serve as a benchmark to compare other methods. We note this baseline method does not guarantee that the achieved availability exactly matches the QoS target, i.e., we cannot ensure  $\gamma_j = \epsilon$ . Specifically, for some periods  $p$ , there is no number of vehicles such that the QoS is exactly met;  $m_{opt_j}^p$  over satisfies the QoS but  $m_{opt_j}^p - 1$  under satisfies the QoS.

## 6.4 Fleet Management

Demand for carshare pools may change over time. If demand increases at some pools and decreases at others, it may be possible to move vehicles between pools to satisfy the QoS rather than purchasing new vehicles. We formulate a single period optimization program that computes the movements and addition of vehicles to a set of pools such that the expected QoS is met. Let:

- $V_{j,k}^p$  be the cost to move a car between pools  $j$  and  $k$  (recall  $v_{j,k}^p$  denotes the number of cars to move from  $j$  to  $k$  during  $p$ )
- $A_j^p$  be the cost to buy a car for pool  $j$  at  $p$  (recall  $a_j^p$  denotes the number of cars to buy for  $j$  during  $p$ )

Figure 6.2 shows our formulation of this problem. The optimization computes the optimal fleet management for each period  $p$ , that is, it outputs the number of cars to move between each pair of pools, and the number to purchase for each pool. We first explain the pre-computation of  $\tilde{m}_{opt_j}^p$ . Recall from §6.3.1 and §6.3.2 that  $\tilde{m}_{opt_j}^p$  is the predicted minimum size of  $j$  during  $p$  such that  $\tilde{\gamma}_j^p < \epsilon$ . All cars in excess of  $\tilde{m}_{opt_j}^p$  can be moved to other pools, and if  $m_j^{p-1} < \tilde{m}_{opt_j}^p$ , this deficit must be resolved by moving one or more vehicles from other pools or adding purchased cars to  $j$ . Thus from  $\tilde{m}_{opt_j}^p$  we compute the deficit or

## Single Period Fleet Optimization Program

Inputs	$\forall j, k : V_{j,k}^p, \forall j : A_j^p$	(6.5)
Pre-computation	$\forall j : \tilde{m}_{opt_j}^p = \begin{cases} \text{ELM} : \min m \text{ s.t. } \tilde{\gamma}_j^p < \epsilon \\ \text{ADMI} : \text{avg}(\text{VPS}_j^{[p-1,p]}) \end{cases}$	(6.6)
Decision Variables	$\forall j, k : v_{j,k}^p \in \mathbb{N}^0$	(6.7)
	$\forall j : a_j^p \in \mathbb{N}^0$	(6.8)
Objective	$\min \sum_{j=1}^{ \mathcal{J} } \sum_{k=1}^{ \mathcal{J} } v_{j,k}^p V_{j,k}^p + \sum_{j=1}^{ \mathcal{J} } a_j^p A_j^p$	(6.9)
Subject To	$\forall j : m_j^p = m_j^{p-1} + a_j^p + \sum_{k=1}^{ \mathcal{J} } v_{k,j}^p - \sum_{k=1}^{ \mathcal{J} } v_{j,k}^p \geq \tilde{m}_{opt_j}^p \dagger$	(6.10)
	$\forall j : m_j^p \leq m_{max_j}^p$	(6.11)

†:The new pool size  $m_j^p$  is equal to the previous size ( $m_j^{p-1}$ ) plus all cars added to  $j$  and moved to  $j$ , minus all cars moved from  $j$ . Moreover this should be at least  $\tilde{m}_{opt_j}^p$ , the predicted minimum size of  $j$  during  $p$  such that the QoS is met.

Figure 6.2: Joint Optimization of fleet size and vehicle movement, for a single period.

excess at  $j$ , if any. However,  $\gamma_j^p$  is a nonlinear function of the input parameters so adding the constraint  $\tilde{\gamma}_j^p < \epsilon$  to the optimization formulation would make it an integer *non-linear* program which require computationally expensive heuristics [Das13, Ree93]. Fortunately,  $\tilde{m}_{opt_j}^p$  is not dependent upon  $\tilde{m}_{opt_k}^p$  for  $j \neq k$ ; it is dependent only upon the behavior and demand of the subscribers of pool  $j$ . That is, the *sizing requirements* of each pool are independent and can be precomputed.

Objective (6.9) ensures cars are moved to satisfy the QoS when possible. Cars are added to pool  $k$  if moving a car from any other pool  $j$  to  $k$  would result in the expected QoS not being met at  $j$ . Constraint (6.10) states that  $m_j^p$  be at least  $\tilde{m}_{opt_j}^p$ , that is, the predicted QoS be met (statistically for ELM, heuristically for ADMI). Constraint (6.19) states  $m_j^p$  cannot exceed the maximum size of pool  $j$  during  $p$ ,  $m_{max_j}^p$ . Combined with Constraint (6.10), this program is infeasible if the bound  $m_{max_j}^p$  is too small given the number of subscribers and their usage patterns. Specifically, the predicted minimum number of vehicles required to meet the QoS target  $\tilde{m}_{opt_j}^p$  may be greater than  $m_{max_j}^p$ . In this case, the pool operator must decide whether to lower the QoS target or increase the maximum pool size.

This is an integer program that is computationally difficult [GJ79], but because the objectives and constraints are linear, it can be solved for moderate problem sizes using linear relaxation and branch and bound [Das13, Ree93] by optimization packages such as Gurobi [Gur14] and CPLEX [IBM14].

For the special case where  $V_{j,k}^p$  is the same for all pairs  $\{j, k\}$ , fleet management is simpler. In this case, we would:

1. Compute the excess or deficit at each pool,  $m_j^{p-1} - \tilde{m}_{opt_j}^p$
2. Move from the pools with excess to the pools with deficits
3. Add cars to pools with deficits, if applicable

## 6.5 Evaluation Methodology

In this section, we describe our evaluation methodology. We use the same dataset for our evaluation as described in §5.3.1.

### 6.5.1 Bootstrapping and Confidence Intervals

To evaluate our sizing and optimization methods on a multi-pool carshare, we split the dataset from Chapter 5 “virtually” among multiple pools using *bootstrapping with replacement* [Efr79]. We perform  $I$  bootstrapping iterations. Recall that we assume subscribers have a preference for a single pool, e.g., the one closest to their home or work, and make all requests from that pool. Within each bootstrap iteration  $i$ , we assign each<sup>5</sup> unique subscriber ID to a random pool. When running the sizing algorithm during iteration  $i$ , we assume subscribers make all vehicle requests to the pool assigned to them during  $i$ . Thus, one bootstrap iteration represents a possible permutation of subscriber preferences. The more iterations that are performed, the more confidence we have that our results are not dependent on a particular assignment. Thus, bootstrapping gives us a sensitivity analysis to subscriber assignment to pools.

With this process, we obtain  $I$  observations of each performance metric—availability, utilization, and M2V—during each period  $p$  for both sizing methods. By the central limit theorem, observations of the performance metrics are expected to be normally distributed when  $I$  is large [Kes12]. In our evaluation, we perform  $I = 100$  iterations and compute the 95% confidence intervals (CI) for every (method, metric, period) triple.

### 6.5.2 Performance Metrics

We use the same three metrics used in the prior chapter (see §5.3.2). To summarize:

- *Member-to-vehicle ratio (M2V)*: the number subscribers per vehicle, typically 30-50.
- *Utilization (U)*: the percentage of time each car is used, typically 20-40%.
- *Availability (A)*: the percentage of served requests at each pool, typically 90-95%.

---

<sup>5</sup>All users are re-assigned every round, hence *Bootstrapping with replacement*.

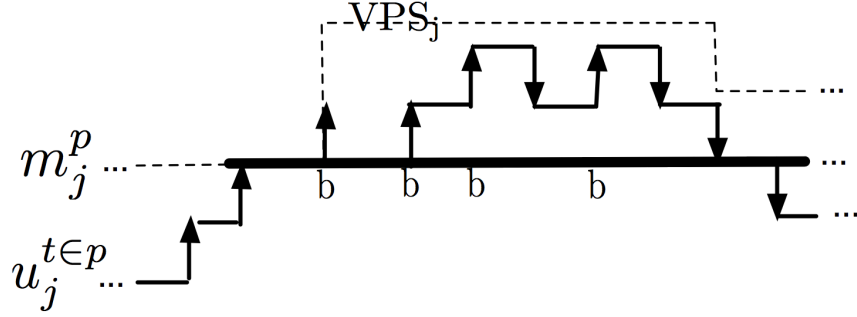


Figure 6.3: ADMI equivalence illustration

### 6.5.3 Performance Metric Computation

During each period, at time  $t$ , we track the number of users assigned to cars,  $u_j^t$ , and the number of vehicles available,  $c_j^t$ . If an arrival finds  $c_j^t = 0$ , i.e., the pool is empty, that arrival is blocked, and  $u_j^t$  is not incremented, with the exception of ADMI as explained below. When simulating ADMI, we allow the number of subscribers of pool  $j$  assigned a car at time instant  $t$  in period  $p$ , denoted  $u_j^{t \in p}$ , to exceed the pool size  $m_j^p$ , but not the virtual pool size  $VPS_j$ . We do this to correctly simulate ADMI—recall that  $VPS_j$  increases to  $VPS_j / \beta$  when  $u_j^{t \in p}$  exceeds  $VPS_j$ , and this condition would never be met if  $u_j^{t \in p} \leq m_j^p$ . Counting arrivals as blocked whenever  $u_j^{t \in p} > m_j^p$  leads to the same metrics as if we restricted  $u_j^{t \in p}$  to be below  $m_j^p$ . For intuition, see Figure 6.3. It is evident that for each event “arrival to the pool when  $u_j^{t \in p} \geq m_j^p$ ” there is a corresponding event “arrival of a request to the pool when  $u_j^{t \in p} = m_j^p$ ”. Since all such arrivals are blocked in both cases, and the event space is the same, availability can be computed using either case.

Denote the length of each period as  $L$ . We compute metrics for period  $p$  as follows:

- $A$  for both ELM and ADMI as  $1 - \gamma_j^p$ , which is the percentage of arrivals within  $p$  that are not blocked.
- $M2V$  for both ELM and ADMI as  $S_j^p / m_j^p$ .
- $U$  for ELM as the Riemann sum of  $u_j^{t \in p}$ , divided by the total area possible,  $m_j^p L$ . We also compute and show availability and utilization within the ELM busy periods.
- $U$  for ADMI as the Riemann sum formed by  $\min(u_j^{t \in p}, m_j^p) \forall t \in p$ , divided by  $m_j^p L$ .

### 6.5.4 Parameter Prediction

Recall that at the start of period  $p$ , three parameters,  $1/\lambda_j^p, 1/\mu_j^p, S_j^p$  (the MTT, MST, and the number of subscribers) are predicted for the ELM method. This prediction can use any values measured during periods  $[0, p - 1]$ . We evaluate two different prediction approaches:

1. Several *moving average* predictors using a moving average of the previous  $H$  values:

$$1/\tilde{\lambda}_j^p = \frac{\sum_{i=p-H}^{p-1} 1/\lambda_j^i}{H}, \dots$$

2. Several Matlab [Mat14] *neural network* predictors. Some networks were configured to use *all* previous values, and others used only some (e.g., 20 previous values).

We define the percent prediction error (PPE) for a metric  $x$  during  $p$  and its prediction as:  $\left| \frac{x^p - \tilde{x}^p}{x^p} \right| \times 100$ . Surprisingly, we found that the PPE did not differ significantly between the two prediction methods. Method (1) gives an average (over all periods and all time series) PPE of  $\approx 20\%$  and method (2) gives an average error of  $\approx 15\%$ . We use method (1) using an empirically derived  $H = 3$  hereafter, because while method (2) gives slightly better predictions, it requires much more computational time than method (1).

We are aware that many more advanced prediction methods exist; we intend to evaluate more predictors as future work.

## 6.6 Results

In this section we present the results of our simulations. Our goal is to show that our sizing methods often meet the desired QoS target despite non-Markovian, non-stationary demand. We generalize from the single pool static sizing case presented in the prior chapter to the single pool dynamic sizing case in §6.6.1. In §6.6.2, we further generalize to the multi-pool dynamic sizing case.

Results shown are for 15-day periods. We use the following notation:

- For ADMI, we use "ADMI( $\alpha$ ,  $\beta$ ,  $\gamma$ , ADMI $_{in}$ ), $j$ " to represent the four ADMI parameters discussed in §6.3.2 (the linear decrease, the multiplicative increase, the number of unblocked arrivals before a decrease is executed, and the initial pool size), and the pool number.
- For ELM, we use "ELM( $K$ ), $j$ " where  $K$  is the busy period length in hours and  $j$  is the pool number.
- We show performance metrics for ELM both overall and during busy periods. In figure legends, the subscript  $B$  indicates "within the busy period".
- The legend notation  $M$  represents the pool size  $m_j^p$ .
- Dashed lines labeled  $op$  represent the offline optimal baseline. Recall from §6.3.3 that the baseline may exceed requirements due to quantization.
- When subscribers are re-assigned to pools during the bootstrapping process, the demand patterns for each pool change. We refer to the demand for any pool, given its current assignment of subscribers, as a *subscriber demand profile*.

Note that the  $M2V$  metric is not a percentage and thus not on the same scale as the other metrics. Further recall that  $op$  always gives the maximum utilization, maximum  $M2V$ , and minimum availability possible such that the target QoS is still met. Whenever the utilization  $U$  or  $M2V$  increases beyond  $op$ , the availability  $A$  is below the QoS target due to an error in demand prediction. When  $A$  is above  $op$ , excess cars reduce the other performance metrics.



## 6.6.1 Single Pool Dynamic Sizing

Here we show results for the dynamic sizing of a single pool. These results are for comparison with those in §5.4.2 and specifically Figure 5.8. Note that the confidence intervals are of size zero for this case because there is no bootstrapping (i.e., subscribers are not randomly assigned to pools because there is only one); no randomness is involved in this simulation.

Table 6.1 shows the average metrics achieved over all periods for two ELM and two ADMI sizing configurations. The ELM QoS target is 95%. ADMI does not size to meet a QoS target, but the configurations shown are plotted together with the offline optimal baseline for a 95% QoS target for comparison. The performance metrics for each period are shown in Figure 6.4a. Table 6.2 shows the average metrics achieved over all periods for the same two ELM configurations for QoS targets of 99% and 90% as a sensitivity analysis to the QoS parameter  $\epsilon$ . The performance metrics for each period are shown in Figure 6.4a.

We first discuss the large improvement in the performance metrics compared to the static sizing of the same pool. Under static sizing, the utilization was between 25—30% for methods that achieved a QoS within 1% of the target (refer to Figure 5.8), but when dynamically sizing the same pool, we achieve utilizations of 35—45%. The reason is as follows. During static sizing, bootstrapping is used to repeatedly split the dataset for training and testing—half of the rows are selected at random for each purpose. Each record includes the number of active subscribers at the time the trip record began—one point on the curve shown in Figure 5.3. During each iteration, half of the points on this curve are sampled yielding an approximation of the number of pool subscribers  $S$ , denoted  $\tilde{S}$ , of  $\tilde{S} \approx 330$ . This is not a good approximation except when  $S \approx 330 = \tilde{S}$  during mid 2011.  $\tilde{S}$  overestimates  $S$  prior to this period and underestimates  $S$  afterwards, but when performing static sizing this  $\tilde{S}$  is assumed constant. Thus, static sizing ignores an important feature of this carshare, which is that it is growing over time. The same problem applies to subscriber demand patterns; for the static sizing case, all demand patterns visible in the training dataset are averaged (using weights) to represent the “typical subscriber” (of which  $\tilde{S}$  are assumed to be present). However, because subscribers with different demands join and leave, this average, though computed using weights representing each subscriber’s “activeness”, may not be indicative of demand patterns observed in the test dataset. When performing dynamic sizing, during  $p$ , the predictions  $\tilde{S}_j^p$  (and  $\tilde{\lambda}_j^p, \tilde{\mu}_j^p$ ) are computed by observing only the prior few periods under the moving average prediction methodology, so they better resemble  $S_j^p$  (and  $\tilde{\lambda}_j^p, \tilde{\mu}_j^p$ ).

We make several other observations:

- All four configurations perform well with respect to the 95% QoS target (with the exception of the last few periods, which we discuss next). Recall that for ELM, the sizing for period  $p$  is performed based on prior observed data. It cannot, of course, guarantee the QoS is met because future demand can be arbitrary. Nonetheless, with few exceptions, the ELM methods achieve the QoS target, showing that both of our proposed approaches are robust. We also see that ADMI achieves availability between 90—95%, despite not sizing for a specific QoS; ADMI performs very well considering it

solely operates based on when and how often arrivals are blocked. ADMI is slightly less robust than ELM, but requires significantly fewer system parameters for its operation.

- We clearly see that utilization ( $U$ ) and  $M2V$  are inversely proportional to availability ( $A$ ) and the QoS target from Tables 6.1 and 6.2. For example, a QoS of 90% (bottom) leads to much higher vehicle utilization, about 15% on average, compared to a QoS of 99% (top curve), and a higher  $M2V$ . A QoS of 95% has intermediate performance, as expected. This emphasizes that there is no way to simultaneously maximize the three performance metrics; a smaller number of vehicles is better for the carshare operator in terms of costs and infrastructure utilization, but worse for customers as they will have a lower QoS guarantee.
- Near period 175, availability falls noticeably below the QoS target, as seen in Figures 6.4a and 6.4b. This is due to the onset of highest demand ever in the system (shown in Figure 5.7) near the end of our dataset in October 2013. This is a consequence of changing pool sizes only at periodic intervals: a *lag time* is introduced between observed increases in demand and increased pool sizes to meet that demand. ADMI rebounds faster because the ELM sizing method uses moving average predictors, and it takes several periods for increases in demand to significantly affect the moving average.
- The two ADMI configurations in Figure Table 6.1 (corresponding to Figure 6.4a) are very different.  $ADMI(2, .5, 5, 1)$  is aggressive: it doubles the virtual pool size  $VPS_j$  upon a blocked arrival, decreases the pool size after just five unblocked arrivals, and decreases it by two each time.  $ADMI(1, .75, 10, 1)$  is much less reactive: it increases  $VPS_j$  by 25%, waits for 10 unblocked arrivals before decreasing the pool size, and decreases by only one. However, they have almost identical performance—the only noticeable difference is that the more aggressive ADMI maintains a slightly lower availability because it sheds (virtual) vehicles more frequently leading to lower periodic averages of  $VPS_j$ . We find ADMI is, for the most part, insensitive to its configuration parameters.
- The  $M2V$  metric steadily but not monotonically increases over time in all of our simulations. This parallels what has happened in practice (as discussed in §5.3.2). The observed utilization is also at the upper end of the ranges reported by existing carshares.

	QoS	P	$\bar{A}$	$\bar{M}$	$\bar{M2V}$	$\bar{U}$
OPT-95	95	0	96.94	7.61	27.13	42.68
ELM(96)	95	0	97.53	8.48	24.47	38.37
ELM(144)	95	0	96.52	8.05	25.96	39.87
ADMI(1,0.75,10,1)	n/a	0	98.21	8.60	23.81	38.14
ADMI(2,0.5,5,1)	n/a	0	97.55	8.41	24.33	37.41

Table 6.1: Comparison of two ELM and two ADMI sizing configurations (corresponds to Figure 6.4a).  $P$  denotes the pool number. ADMI does not size for a QoS target.

	QoS	P	$\bar{A}$	$\bar{M}$	$M\bar{2}V$	$\bar{U}$
OPT-99	99	0	99.61	9.11	22.64	36.85
ELM(96)	99	0	99.25	10.5	19.84	31.17
ELM(144)	99	0	98.63	9.47	22.16	34.58
OPT-90	90	0	92.81	6.63	31.14	47.06
ELM(96)	90	0	96.47	8.04	26.02	39.88
ELM(144)	90	0	94.27	7.12	28.88	44.10

Table 6.2: Sensitivity of ELM to the QoS parameter (see Figure 6.4b).  $P$  denotes the pool number.

## 6.6.2 Mutli-Pool Dynamic Sizing

We have five dimensions of data: 100 bootstrap iterations,  $\approx 190$  periods (eight years of data), three performance metrics, two sizing methods, and  $J$  pools. For ease of exposition, we show representative results for only two pools and note that the trends discussed did not change with more pools. Unlike the prior section, the figures showing the performance metrics over time are deferred to Appendix B, as most conclusions are drawn directly from the summarizing tables.

Table 6.3 shows the average metrics achieved over all periods for four different sizing configurations. The ELM QoS target is 95%. ADMI does not size to meet a QoS target, but the configurations shown are plotted together with the offline optimal baseline for a 95% QoS target for comparison. The performance metrics for each period are shown in Appendix B in Figure B.2a. Table 6.4 shows the average metrics achieved over all periods for the same two ELM configurations (also for pool 0) for QoS targets of 99% and 90% as a sensitivity analysis to the QoS parameter  $\epsilon$ . The performance metrics for each period are shown in Appendix B in Figure B.2b. Finally, Table 6.5 shows the average metrics achieved over all periods for both carshare pools (pool 0 and pool 1) to examine the effect of varying subscriber demand profiles. The performance metrics for each period are shown in Appendix B in Figure B.3a.

We find that the overall results from the single pool dynamic sizing case presented in §6.6.1 continue to hold here. The ELM methods work well with respect to their given QoS target, with few exceptions due to prediction errors and unexpected demand, as shown in Figures B.2a, B.2b, and B.3a. ADMI also achieves an average QoS between 90—95% despite not sizing for a QoS target. We find the same insensitivity of ADMI to its sizing parameters as we did previously; the two configurations shown in Figure B.2a are very different yet achieve near identical performance. We also see a decrease in performance near period 175 due to the onset of very high demand.

We make two additional observations. First, the performance metrics for each pool in a multi pool system are worse, on average, than metrics achieved with a single pool. This is a standard result in queueing theory: a shared resource is more efficient than the same resource strictly partitioned amongst sets of users [HK88]. Second, the differences in achieved metrics for the same sizing method between both carshare pools, as shown in Table 6.5 (corresponding to Figure B.3a) are negligible. Each pool has 100 bootstrap iterations of

randomly assigned subscribers each with their own demand profiles, yet the methods achieve nearly the same metrics on each iteration for each pool. Moreover, the 95% confidence intervals in each graph for each period, computed using 100 bootstrap iterations, are very small—less than  $\pm 3\%$  of the mean. This is supporting evidence that our methods work are insensitive to subscriber demand distributions.

	QoS	P	$\bar{A}$	$\bar{M}$	$M\bar{2}V$	$\bar{U}$
OPT	95	0	97.07	5.02	20.46	33.34
ELM(96)	95	0	97.69	5.58	18.98	30.04
ELM(168)	95	0	96.70	5.17	20.49	32.03
ADMI(1,0.9,20,1)	n/a	0	97.26	5.21	19.81	31.85
ADMI(2,0.5,5,1)	n/a	0	95.09	4.67	22.35	33.35





Table 6.3: Comparison of four sizing configurations (corresponding to Figure B.2a).  $P$  denotes the pool number. ADMI does not size for a QoS target.

	QoS	P	$\bar{A}$	$\bar{M}$	$M\bar{2}V$	$\bar{U}$
OPT-99	99	0	99.49	6.35	16.16	27.33
ELM(96)	99	0	99.30	7.11	15.07	24.71
ELM(168)	99	0	99.05	6.56	16.09	26.54
OPT-90	90	0	94.14	4.37	23.48	36.90
ELM(96)	90	0	96.04	5.01	21.11	32.93
ELM(168)	90	0	94.29	4.59	22.95	35.41





Table 6.4: Sensitivity analysis to QoS target (corresponding to Figure B.2b).  $P$  denotes pool number.

	QoS	P	$\bar{A}$	$\bar{M}$	$M\bar{2}V$	$\bar{U}$
OPT	95	0	97.07	5.02	20.46	33.34
ELM(216)	95	0	96.42	5.08	20.82	32.58
ADMI(1,0.75,10,1)	n/a	0	97.05	5.21	19.89	31.73
OPT	95	1	97.04	5.0	20.57	33.15
ELM(216)	95	1	96.40	5.07	20.88	32.29
ADMI(1,0.75,10,1)	n/a	1	97.02	5.20	19.97	32.17

Table 6.5: Comparison of metrics across the two pools (corresponding to Figure B.3a).  $P$  denotes the pool number. ADMI does not size for a QoS target.



(a) Performance metrics for two ELM configurations and two ADM configurations, single pool dynamic sizing case. These are plotted together with the offline optimal baseline for a 95% QoS target. The averages of each metric achieved over all periods are summarized in Table 6.1.



(b) Performance metrics for ELM(96) and ELM(144), the same two ELM configurations shown in 6.4a, for QoS targets of 99% (left) and 90% (right). The averages of each metric achieved over all periods are summarized in Table 6.2.

## 6.7 Modeling Removals

Thus far, we have assumed that fleet management consists of adding vehicles to the network and moving vehicles between pools in the network. However, vehicles added during periods with high demand may not be needed forever. If the carshare experiences an extended period with lower demand, the operator may wish to decrease the size of the system. We have so far assumed that *removals*, the selling of vehicles, would be handled exogenously by the carshare operator. In this section, we:

1. extend our optimization model to allow vehicles to be removed each time period
2. show a demand scenario where implementing removals has performance benefits
3. discuss why implementing automatic removals is problematic in practice and why we propose that removals should be handled on a case-by-case basis by the carshare operator

### Implementing Removals

We first describe how we extend our model to allow car removals. The new optimization program, modified from Figure 6.2, is shown in Figure 6.5. We introduce variables  $r_j^p$  and  $R_j^p$ , indicating the number of cars to remove from pool  $j$  during time period  $p$ , and the salvage value from doing so. The objective is modified to include an additional term,  $-\sum_j r_j^p R_j^p$  (the total profit incurred from selling), and the constraint is modified to include a  $-r_j^p$  term (the new pool size is reduced by the number of vehicles removed).

This myopic optimization program removes cars when *optimal with respect to performance metrics* in a single time period—cars are removed from  $j$  at the beginning of  $p$  if the QoS for  $p$  given  $m_j^{p-1}$  is predicted to be over-met. This implementation allows us to compute an *upper bound* on the performance metrics gains we can achieve by implementing removals, assuming there was no overhead involved in buying and selling cars. At the end of this section, we discuss another approach that optimizes for long term costs rather than performance, but show it is infeasible to use in practice.

### Performance With Pool Removals

In §5.4.2 we showed results for the static sizing of a single pool assuming all trips under 80km are handed by the subscribers' personal vehicles. Once trips under 80km are removed, very little demand ( $\approx 10\%$ ) remains, so the gaps between demand are larger. After periods with increased demand, the sizing methods react by increasing the pool size, but without removals the pool size remains unnecessarily high through periods of low demand.

Table 6.7 (corresponding to Figure B.3b in Appendix B) shows results for the dynamic sizing of a single pool (pool 0), considering only trips over 80km, for two ELM configurations both with or without removals. Table 6.6 (corresponding to Figure B.1 in Appendix B) shows the same for two ADMI configurations. We make several observations:



## Fleet Optimization Program With Removals

Inputs	$\forall j, k : V_{j,k}^p$	(6.12)
	$\forall j : A_j^p, R_j^p$	(6.13)
Pre-computation	$\forall j : \tilde{m}_{opt_j}^p = \begin{cases} \text{ELM} : \min m \text{ s.t. } \tilde{\gamma}_j^p < \epsilon \\ \text{ADMI} : \text{avg}(\text{VPS}_j^{[p-1,p]}) \end{cases}$	(6.14)
Decision Variables	$\forall j, k : v_{j,k}^p \in \mathbb{N}^0$	(6.15)
	$\forall j : a_j^p, r_j^p \in \mathbb{N}^0$	(6.16)
Objective	$\min \sum_{j=1}^J \sum_{k=1}^J v_{j,k}^p V_{j,k}^p + \sum_{j=1}^J a_j^p A_j^p - \sum_{j=1}^J r_j^p R_j^p$	(6.17)
Subject To	$\forall j : m_j^p = m_j^{p-1} + a_j^p - r_j^p + \sum_{k=1}^J v_{k,j}^p - \sum_{k=1}^J v_{j,k}^p \geq \tilde{m}_{opt_j}^p$	(6.18)
	$\forall j : m_j^p \leq m_{max_j}^p$	(6.19)

*Figure 6.5: Joint optimization program with systematic removals*

- We see from Figure B.1 that  $Mop$  fluctuates much more than when dynamically sizing a pool. This suggests that not allowing  $M$  to increase *and decrease* in response to fluctuations in demand decreases  $M2V$  and  $A$ . As evidence, without removals, the four conservative methods give a pool size  $M$  that is always greater than  $Mop$ .
- For ELM(8), a very conservative sizing method, nearly 100% availability is maintained *with and without* removals. However, with removals implemented, the  $M$  curve is a tighter upper envelope so the other performance metrics are improved. This is an example where removals can improve the other performance metrics without affecting availability.
- For ELM(72), with removals implemented, the average performance metrics nearly match the optimal baseline.
- Due to the reactive nature of the sizing methods, the pool size is increased and decreased after increases and decreases in demand. This lag can be observed two ELM and two ADMI configurations with removals implemented. Consequently, when the optimal pool size oscillates rapidly, the lag causes the pool size to be suboptimal during some periods.
- Though removals can lead to decreased availability during some periods, it also leads to large performance gains. For example, in Table 6.7 the two ADMI configurations with removals implemented see much higher utilization and  $M2V$  metrics in exchange for a few periods with low availability.
- Even with removals implemented, the performance metrics are lower for the truncated demand case where all trips under 80km are assumed to be handled by the subscribers' personal vehicles. When periods with high demand are far apart, the moving average

predictors may “reset” to low demand during the interim periods. Moreover, with less demand, ADMI has fewer arrivals to react to; ADMI works best when it can determine the pool size based on a large number of arrivals.

	QoS	P	R	$\bar{A}$	$\bar{M}$	$\bar{M}^2V$	$\bar{U}$
OPT	95	0	0	98.06	4.92	45.34	31.20
ADMI(1,0.5,15,1)	n/a	0	0	99.81	12.08	17.08	13.46
ADMI(1,0.5,30,1)	n/a	0	0	99.81	11.26	18.32	14.28
OPT	95	0	1	98.06	4.92	45.34	31.20
ADMI(1,0.5,15,1)	n/a	0	1	98.33	7.5	31.04	22.79
ADMI(1,0.5,30,1)	n/a	0	1	99.52	9.04	25.34	18.09

Table 6.6: The averages of each metric achieved over all periods across pools and for two different ADMI configurations (also see Figure B.1).  $P$  denotes the pool number. ADMI does not size for a QoS target.  $R$  is a boolean indicating whether removals are implemented.

	QoS	P	R	$\bar{A}$	$\bar{M}$	$\bar{M}^2V$	$\bar{U}$
OPT	95	0	both	98.06	4.92	45.34	31.20
ELM(8)	95	0	0	99.81	18.12	11.66	8.86
ELM(8)	95	0	1	99.81	11.847	18.64	13.62
ELM(72)	95	0	0	99.36	7.86	26.62	20.34
ELM(72)	95	0	1	96.43	5.33	40.83	28.41

Table 6.7: The averages of each metric achieved over all periods for pool 0 and four different ELM configurations with and without removals (also see Figure B.3b).  $P$  denotes the pool number.  $R$  is a boolean indicating whether removals are implemented. Note the offline optimal baseline always removes cars when beneficial with respect to performance metrics.

## Implementing Systematic Removals In Practice

The myopic optimization program shown in Figure 6.5 removes vehicles from the system whenever it is optimal with respect to performance metrics. With myopic optimization, when demand fluctuates, some vehicles may be sold during low demand and repurchased later during high demand. However, assuming vehicles depreciate in value, this may be suboptimal over time with respect to cost. Let  $A_j^p - R_j^p$  denote the *rebuying penalty*, which intuitively represents the cost to buy a vehicle after selling one too hastily (due to myopic optimization). To implement removals while optimizing for cost instead of performance, we could model the *the holding cost*  $c_h$ , the cost to store each vehicle for one period, and sell vehicles only when the rebuying penalty is cheaper than the holding cost over several periods. The problem is that to make an optimal decision w.r.t. cost at the start of any period  $p$ , the operator needs to know *when* cars that are removed will be needed in the future. The operator cannot optimize for cost in practice unless demand is perfectly known far ahead of time:



To see this, consider a single pool with holding cost  $c_h$ , and let the profit from selling a car be  $R_j$ . Let  $X$  be the largest integer such that  $A_j - R_j \geq c_h X$ ;  $X$  represents the maximum number of periods for which it is cheaper to hold an unused car than to sell it re-buy it later. Suppose  $m_j^{p-1} = 2$  and  $\tilde{m}_{opt_j}^p = 1$ . Should the carshare sell the 2nd vehicle at the start of period  $p$ ? It is optimal to keep the car if and only if the next period at which  $\tilde{m}_{opt_j} \geq 2$  is no more than  $X$  periods away, and beneficial to sell it otherwise. That is, to make a decision optimal with respect to cost at the start of period  $p$ , demand must be known for all periods  $[p, \dots, p + X]$ .

In conclusion, it is difficult to implement the optimization in Figure 6.5 in practice. Thus, we propose that the complex problem of when to sell vehicles should be handled by the carshare operator on a case-by-case basis, rather than on the basis of an optimization program. Our formulation in Figure 6.2 advises when cars should be added to or moved within the system, and whether any cars are "excess" at any time, but it does not automatically remove the excess vehicles. We propose that this optimization be used to inform the operator when there are excess vehicles in the system so they can judge whether to sell the vehicle.

## 6.8 Prior Fleet Sizing and Management Models

Prior work in fleet management makes one or more of the following modeling assumptions:

1. the number of users who can create demand in the system (i.e., request a vehicle from some location) is "infinite" and not restricted to a finite group of subscribers
2. the distribution of interarrival times and service times at locations are Poisson, i.e., the system is Markovian
3. the arrival processes to locations is stationary, i.e., does not vary with time of day or day of year.

These three assumptions, made for modeling simplicity, are often inaccurate. To demonstrate this, we show that a local Waterloo carshare has non-Markovian arrivals and departures, has non-stationary demand, and has a finite population of subscribers in §5. Other authors also made similar observations. In a study of a company car pool, Dondeti et al. [DM09] find that not only are the arrival and service distributions non-Markovian, but are not drawn from "any standard distribution". The arrival process for another business pool was found to be close to Markovian, but not the service process [WF79]. With respect to stationary demand, many carshares have higher pricing rates during peak periods [SM07], i.e., they have pricing that reflects (and attempts to balance) their non-stationary demand. Finally, infinite population systems are suitable for large carshares with thousands of subscribers, but not smaller carshares with a small subscriber base (like the one we study)<sup>6</sup>. Carshares often store their fleet across multiple locations to provide better access for their geographically dispersed subscribers. Because subscribers often borrow vehicles from the

---

<sup>6</sup>We note that this is the least unrealistic assumption made by prior work, because infinite population models approximate finite population models as the size of the population grows [Kle75a].

same pools consistently, e.g., the one closest to their home or work, often each location can be modeled as a smaller finite population system.

To the best of our knowledge, §6 presents the first work to manage a multi-pool vehicle sharing system without making any of the three aforementioned assumptions (infinite population, Markovian arrival and departure processes, and stationary demand). There are several prior models that make some or all of these assumptions.

The three closest related models are by George et al. [GX11], Hampshire et al. [HG10], and King et al. [KGWS13]. George et al. [GX11] study sizing a network of car rental pools. The arrival process is assumed to be Markovian, and service times at rental stations are assumed to be Markovian under a virtual service model, thus rental pools are modeled as Markovian queues. Non-stationary demand is not considered. The authors minimize costs by assuming there is a financial penalty for meeting any given QoS level under 100% and a cost to buy vehicles. Hence it may be cheaper to maintain a low QoS if the price of vehicles is high. Hampshire and Gaites [HG10] focus on the sizing and profitability of a peer-to-peer vehicle sharing service. However, their sizing algorithm makes two assumptions that do not apply to our problem. First, the authors assume that the arrival of subscribers to car shares is stationary. It is not, especially near major holidays—some car shares even have higher rates during peak periods [SM07]. Thus, much of our work deals with sizing according to busy periods. Second, the authors assume the population of subscribers is infinite. We instead focus on smaller pools that cannot be approximated with an infinite population model. King et al. [KGWS13] present a vehicle access model to reduce range anxiety and present sizing methods to manage the size of the ICEV fleet. However, all service times are assumed to be deterministic—one day each. The sizing of the system is based on stationary demand (the average arrival rate). Finally, the authors assume that all load is shiftable and that requests queue, i.e., if a user demands a vehicle for day  $k$  and one is not available until day  $k + 3$ , the user will use the vehicle then. However, not all vehicle demand is shiftable by one or more days, especially demand for long trips because they tend to occur on weekends and holidays. See [PT08, Par77, LT09, KKN03, Koc97, SHvH13] for similar queueing models that make the same assumptions.

Two other papers have used different (non-queueing) tools to study the fleet management problem. Barth [BT99] use an agent-based modeling approach for the sizing and management of a vehicle sharing system. An origin-destination (OD) matrix of carshare trips is created using online travel surveys to build a Markov process that generates arrivals. That is, the arrival process is assumed to be Markovian and stationary. Service times are assumed to be known and deterministic between each OD pair. For the simulation, OD trip origins are generated randomly according to this Markov process, and the duration of each trip is chosen based on the origin and destination. The authors propose pool sizing and vehicle relocation algorithms based on trips simulated using the Markov process and the OD matrix. Dondeti et al. [DM09] use a similar approach to Barth, but trips are generated using six years of pool data. The authors find that the arrival and service distributions do not follow any standard distribution. Instead of employing generic queueing models, the authors simulate trips from the empirical distributions directly by discretizing trip length into half

hour buckets and determining the probability a random trip is in each. They assume all trip requests that cannot be fulfilled from the pool are served from a rental organization for a price dependent upon the length of the trip request. Their goal is to determine the optimal size of the pool given the tradeoff between the purchase and holding cost of pool vehicle and the cost to acquire rentals for unserved demand.

These fleet management models are related to, but differ from, models for the joint sizing and routing (SR) of a vehicle fleet [Lap92, PJO95]. These two sets of models have different objectives. SR models optimize the size and routes of transport fleets serving customer shipping demands over a multi-period horizon to minimize cost. For each period, vehicles must be assigned to demands and routes, and vehicles must be moved to each node based on anticipated demand. However, SR models allow demand to be *backlogged*, that is, unserved demand (due to a lack of vehicles at a particular time) can be served later for a penalty. In the seminal SR model developed by Beaujon [BT91], *all* demand is backlogged.

This is not always possible in managing fleets such as carshares. For example, if no carshare vehicles are available, and a subscriber needs to make a trip, they may not be able to postpone the trip to another time. For example, a subscriber may want a vehicle for a holiday weekend trip. Due to backlogging, in some SR models all demand may eventually be met, i.e., there is an "eventual QoS" of 100%. In other models, unmet demand is dropped if it is cheaper to do so. Thus SR models minimize cost and do not ensure a given QoS is met within any specific time period. Queueing models are not used for sizing. Instead, SR models use multi-period optimization models to minimize cost.

## 6.9 Conclusions And Future Work

We propose that a carshare for EV owners may help reduce range anxiety by giving owners occasional access to an ICEV. Subscribership to such a program can be incentivized or even subsidized by EV dealerships. We hypothesize that a high quality of service (vehicle availability) is essential to the success of the program. Contrary to existing fleet management models, our model can be used to size vehicle sharing systems for arbitrary demand patterns to ensure a high QoS. We solve the dynamic sizing problem using repeated myopic optimization. At periodic intervals, we compute the minimum size of each carshare pool necessary to ensure a systemwide QoS target is met. We adapt two sizing methods we proposed in the previous chapter for this purpose. Cars are then moved between or added to pools to satisfy the QoS. We show that our fleet management methodology performs well with respect to optimal baselines for a variety of subscriber demand profiles.

We now discuss several ways our model can be extended.

### 6.9.1 Supporting One Way Trips

We do not currently model one way trips (i.e., cars that are rented from one pool and returned to a different pool) for two reasons. First, there already exist several methods

to *rebalance* sharing systems—methodologies to optimally move a fixed number of vehicles between locations after a series of one way trips to resolve pool imbalances [CMR12, MW05, MHN10, FSS13, UMW07]<sup>7</sup>. We can include these methods to allow for one way trips as follows. Sizing is still done at the start of each period using our methodology, and our fleet management is performed immediately after sizing to compute whether cars need to be added to the system and whether cars should be moved due to demand fluctuations. These procedures ensure each pool is sufficiently sized for the coming period, that is,  $m_j^p \geq \tilde{m}_{opt_j}^p \forall j$ . To implement one way trips, any of the cited re-balancing algorithms can be employed daily (or weekly, etc.) to ensure  $m_j^p \geq \tilde{m}_{opt_j}^p \forall j$ . Since vehicles do not leave the system between periods, it is always possible to rebalance.

Second, at the time of writing, the carshare used in our evaluation does not allow for one way trips [Com14]. The carshare has many pools spread over a wide geographical region to serve a wide population, and many of these pools are limited to just a few spaces because of parking costs and availability. Allowing one way carsharing creates logistical problems, e.g., what should happen when a subscriber brings a vehicle back to a pool which is already full? Subscribers may have to visit several locations to return vehicles due to limited parking availability.

## 6.9.2 Subscribers Borrowing From Multiple Locations

Our model assumes that subscribers always borrow cars from a particular pool, e.g., the one closest to their home or work, and do not borrow vehicles from different pools at different points in time. Recent work has shown that there is value in exploiting communications technologies to allow users to choose vehicles from more than one pool [CJA14].

However, our model can be easily adapted to handle this. For the ELM method, the population of each pool would now be the total number of subscribers,  $|\mathcal{S}|$ , as subscribers can rent from any pool. Furthermore, the mean think time and mean service time ( $1/\lambda_j, 1/\mu_j$ ) of pool  $j$  should be calculated as a weighted mean of *all* subscribers' think and service time distributions, with the weight being *how often they use pool  $j$* . ADMI needs no modifications. If the dataset logs the start and end location of each trip, the simulation methodology can be modified to use this information. We did not perform these extensions because our dataset does not provide this information.

## 6.9.3 Reservations

A carshare operator may wish to implement a *reservation system* where subscribers call ahead to reserve a vehicle. We do not model this type of system because we are using a historical dataset and cannot “reschedule” past requests, and because many existing carshares size purely based on predicted demand [Mat04]. We plan to consider this in future work.

---

<sup>7</sup>This list is not exhaustive: a large body of work exists on this topic

## 6.9.4 Budget Limitations

Our current fleet management formulation (Program 6.2) minimizes cost subject to a hard constraint on QoS. However, this formulation may ignore cost constraints—meeting the desired QoS systemwide may not be feasible given budgetary constraints. Future work could consider an alternative formulation that instead maximizes the QoS subject to fixed cost constraint. A possible objective may be to maximize the average availability across pools, e.g.,  $\max \sum_{j=1}^{|\mathcal{J}|} \tilde{\gamma}_j^p$ , where the expected blocking probability  $\tilde{\gamma}_j^p$  is a function of the pool size  $m_j^p$ , and cost is also a function of  $m_j^p$ . However, we caution that this requires computing  $\tilde{\gamma}_j^p$  *within* the optimization program over the space of pool sizes  $m_j^p$ , as  $m_j^p$  would then be a decision variable, and the computation of  $\tilde{\gamma}_j^p$  (Eq(6.1)) is recursive and its analytical equivalence (Eq(5.2)) is nonlinear. In contrast, in our formulation,  $m_j^p | \tilde{\gamma}_j^p < \epsilon$  is pre-computed and given as input to the optimization program.

Alternatively, the carshare operator can compute the cost of maintaining a few acceptable QoS targets (e.g., 80%, 85%,...) using Program 6.2 and then make a decision based on cost constraints.

## 6.9.5 Multi Period Future Optimization

We show in §6.7 that without reliable prediction of future demand, the operator cannot optimize over a multi period horizon. Moreover, we observe that prediction errors of 20% per period can compound, i.e., at the start of period  $p$ , the prediction error for demand during  $p$  is  $\approx 20\%$ , but the PPE for period  $p + 1$  can be 40%. This is an inherent problem caused by unrestricted demand—subscribers can join and leave at any time and may change their demand patterns significantly without notice. Thus, unless demand is regulated by the carshare provider, perhaps by placing a restriction on the usage by subscribers, this problem will be difficult to overcome in practice. Nonetheless, as future work we plan on evaluating methods for optimizing multiple periods using a receding horizon optimization framework into the future.

## 6.9.6 Redirecting Blocked Arrivals

Another avenue for extending our work is to model the multi pool network as a queueing network, where if a subscriber is blocked (does not receive a vehicle) at one pool, they try to receive a vehicle from a different pool instead of returning to their think state. This requires alternating the MTT  $1/\lambda_j$  and the MST  $1/\mu_j$  at each pool  $j$  according to the additional demand from those who came from other pools after being blocked.

# Chapter 7

## The Return On Investment for Taxi Companies Transitioning to Electric Vehicles

**Publication reference:** *T. Carpenter, A. R. Curtis, S. Keshav. The Return On Investment for Taxi Companies Transitioning to Electric Vehicles: A Case Study in San Francisco. Transportation Vol 40, No. 4, July 2013.*

---

### 7.1 Synopsis

Fleets are a fruitful adoption market for EVs for a number of reasons. Fleet operators:

- are able to assign vehicles based on demand, e.g., a fleet can dispatch an EV or ICEV depending on the range requirements of planned trips
- are rational, in that they are comfortable making purchase decisions based on long term ROI calculations
- avail substantial discounts due to volume purchases
- often have access to centralized parking that is suited to EV charging

We focus on taxi fleets, a popular type of vehicle fleet in cities. Taxi operators can displace significant amounts of petroleum use by transitioning to EVs<sup>1</sup>. However, they will only transition if it is profitable to do so. In this chapter, we study whether taxi companies can simultaneously save petroleum and money by transitioning to electric vehicles.

---

<sup>1</sup>Based on our data, for triple-shift taxis (those driven 24 hours a day) we estimate savings of approximately 15,000 liters each year per taxi, and 10,000 liters for double-shift (16 hours a day) taxis

As discussed in §3.3, several authors have studied the ROI for individuals transitioning to EVs. We cannot use the mobility patterns of individuals to derive conclusions about the ROI implications of fleets switching to EVs, because fleet mobility patterns are very different. For example, we find that the taxis in our dataset are parked only 12% of the time, whereas private vehicles are parked 80—90% of the time (see §3.4).

In this chapter, our main contributions are:

1. We present the first work on computing the return on investment (ROI) for a taxi fleet transitioning to EVs. Our data-oriented approach evaluates different infrastructure scenarios, including battery switching and roadside charging, and all EV types: BEVs, PHEVs, and HEVs. For each scenario, we quantify the ROI and the investment payback period, and extrapolate the analysis to a wide array of electricity and petroleum prices.
2. As part of (1), we present a graphical model for using data collected from the company's ICEV taxis to estimate the SOC of a fleet of electric taxis over time.
3. We formulate the problem of locating battery switching stations that serve the taxi fleet as a revenue-maximizing optimization problem.

We do not make any assumptions about the vehicles' mobility patterns; instead, we use a data set of GPS coordinates of the company's existing ICEVs to derive conclusions. Our process requires an input dataset containing GPS and passenger fare information for a taxi company's existing ICEV taxis.

As a case study, we analyze the adoption of PHEVs, BEVs, and HEVs by a taxi company with over 530 vehicles, Yellow Cab San Francisco. Our study shows that both PHEVs and HEVs have a positive ROI as of late 2014 in San Francisco, and if battery switching stations are available<sup>2</sup>, BEVs are profitable as well. Using our algorithm, we find only three battery switching stations are needed for Yellow Cab San Francisco for BEVs to be profitable. Furthermore, gasoline prices in San Francisco are 5.4% higher than the rest of the United States, but electricity prices are 75% higher, so taxi companies with similar practices and mobility patterns in other cities are likely to profit more than YCSF by transitioning to EVs.

This chapter is organized as follows. An overview of taxi operations is given in §7.2. We present our ROI analysis methodology in §7.3. In §7.4 we give an algorithm to locate necessary BEV infrastructure. The results from our case study are given in §7.5. Conclusions and future work are discussed in §7.6.

## 7.2 Taxi Operation Overview

We briefly describe the operation of a taxi company as it relates to our work and how its operating practices may change if it transitions its fleet to EVs.

---

<sup>2</sup>As discussed in §2, the first manufacturer of switching station infrastructure, Better Place, is recently (2013) defunct. However, other EV manufacturers, such as Tesla [Tes14c], have an interest in battery switching. We continue to use Better Place's estimates for switching infrastructure pricing for lack of an alternative.



Our model applies to taxi companies in *fare-regulated* taxi markets where the local government or local Taxi Commission controls the pricing structure of taxis within the region. Moreover, we assume that employee-drivers operate company-owned vehicles in *shifts* that typically last eight hours. At the end of the shift, drivers return the vehicle to the company premises where the vehicles are re-fueled and handed over to another driver for another shift. Drivers only refuel the vehicles while they are not carrying passengers.

We use the term *fare* to refer to a contract between a driver and a passenger to transport the passenger to a desired destination for some price. Fares may be pre-arranged by calling the company to schedule a pickup, or they can be arranged on-the-fly by signaling taxis as they drive past. Each taxi company has its own pricing model, that is, how it charges for fares. It is usually a function of time, distance, and other charges. A driver's goal is to complete as many fares as possible during their shift, as this is the sole source of revenue for the company.

To maximize the likelihood of fares, taxi drivers may continuously drive around looking for passengers or may wait at busy locations such as airports and city centers. This behavior is not fuel efficient, but the revenue from additional fares usually compensates for the cost of wasted fuel.

We now note how this existing operation may change if the taxi company were to convert their fleet to BEVs or PHEVs, where BEVs are allowed to switch their batteries at a switching station. Such a change can impact the frequency of refueling, the potential introduction of refueling delays between shifts, and driver behavior between fares.

- *BEVs*. The range of a typical BEV is about a third of an ICEV. Thus, BEVs must be "refueled" about three times more often than ICEVs. Consequently, either drivers must refuel more often between fares or turn down more fares. Note that installing battery switching stations allows BEVs to be refueled as quickly as ICEVs. Therefore, there is no additional delay at the company premises between shifts.

Many equations in the following sections are dependent on a variable  $\tau$ , which represents the battery charge threshold below which taxi drivers switch their battery if they are at a location with a switching station. We use the notation  $\cdot(\tau)$  to represent a variable's value assuming the switching threshold is  $\tau$ . We assume that a BEV driver switches their battery whenever the SOC is less than  $\tau$  and the driver is at a location with a switching station. That is, we assume drivers never modify their trajectories to switch batteries. We discuss computing the optimal value of  $\tau$  in §7.3.5.

- *PHEVs*. PHEVs do not have to be refueled more often than ICEVs. However, the primary gain from switching to PHEVs is to reduce fuel costs by driving the taxis primarily using the battery. This reduction is possible only if taxis rarely switch to their ICE mode, which requires their batteries to be fully charged after the end of a shift. PHEV batteries cannot be switched in today's models. This introduces a delay between shifts while the vehicles are being charged. To avoid this delay, the taxi company could purchase additional PHEVs to ensure vehicle availability for the next shift. This issue is discussed in detail in §7.3.5.



- In both cases, a driver's practice of opportunistically attracting fares by driving around would be affected. Drivers need to trade off the benefit from additional fares for the cost of battery depletion.

## 7.3 Data-Oriented Process to Estimating ROI

Our goal is to calculate the company's ROI in transitioning a certain fraction of their taxi fleet to EVs. To do this, we describe an analytical technique that allows us to use measurements of the company's existing ICEVs to study whether they should adopt EVs. This section is outlined as follows. In §7.3.1 we discuss the assumptions we make in our model. In §7.3.2, we describe the required inputs necessary to use our model. In §7.3.3 we describe the model outputs. We describe how we use data collected from the company's ICEVs to model EVs in §7.3.4. In §7.3.5 we discuss how the model is used to infer the company's ROI as a result of transitioning a portion of their fleet to EVs.

### 7.3.1 Model Assumptions

In our model, we assume that:

- if a taxi depletes its battery during a shift, all revenue the taxi would have generated during the remainder of that shift is lost. This is a conservative assumption given that in practice a taxi could potentially drive to a switching station, but this would change the trajectory of the vehicle. Since we are using records of past trips, we would not know whether the fares after this trip to a switching station would be valid.
- the company replaces every ICEV, BEV, and PHEV, after  $L$  years of use, regardless of the level of usage. The lifetime  $L$  is given as input. For simplicity, we also assume that extra batteries in the system (for BEVs) are replaced after  $L$  years of use. Future work could examine different replacement rates for different infrastructure.
- the taxi system analyzed as input operates in a fare-regulated taxi market. We cannot estimate the revenue derived from taxis or the relative costs of owning an ICEV compared to an EV in an unregulated market, as taxi drivers could make arbitrary changes to their fare structure in response to changes in gasoline and electricity prices, making our projections inaccurate. From surveying a representative sample of cities in North America however, we conclude most taxi markets are regulated by the local Taxi Commission or local government.
- the company does not make a large change to their fleet size so these values over the next  $L$  years will resemble the values from the dataset. Our model is data oriented—it relies on measured data rather than a survey of taxi operations. Any dataset will have values "built in" because the values are measured from an already existing system, such as the average number of fares completed per taxi, the average utilization of the company's taxis, and the average revenue brought in per taxi.

- the price of gasoline and electricity do not go below some lower bounds during  $L$ . That is, our model computes a worst-case ROI—it assumes the price of gasoline and electricity are at least as high as the input parameters. Removing this assumption would require us to introduce stochastic processes for the price of gasoline and electricity, greatly complicating the model.
- EVs are introduced into the fleet at the start of the analysis period. Our model computes the ROI for the taxi company over  $L$  years assuming a fraction of their fleet is replaced with EVs at the *start of  $L$* . It is straightforward to extend this analysis to the case where EVs are introduced in batches, as each ROI analysis is independent, and the total ROI would be the sum of each batch.
- batteries charge at a constant rate and the battery capacity does not change over its lifetime.
- taxis do not deviate from the routes in the data set, that is, we do not modify the dataset in any way. Taxis are only allowed to switch batteries if they are at a location with a switching station and do not have a fare. In practice, taxis would monitor their battery levels and may drive to a switching facility in between fares.
- maintenance costs for a fleet of EVs is the same as a fleet of ICEVs. This is a conservative assumption because EVs, with fewer moving parts, require less maintenance than ICEVs. We further assume that driver salaries and dispatch expenses remain fixed whether the company uses EVs or ICEVs. We do not model these three parameters.
- taxis do not charge their batteries with the air conditioning (AC) on. We model the effect of AC usage on the battery capacity, but make this assumption so that at any time the taxi is either consuming or gaining energy, but not both.

These assumptions essentially state the taxi market remains relatively constant over  $L$  years. Further limitations are discussed in are discussed in §7.6.

### 7.3.2 Inputs

Our process for determining the changes in revenue for the company as a result of switching to EVs requires the following inputs:

1. *Mobility Data*. A critical input to our taxi model is mobility data from the existing ICEV fleet. We require the periodic collection, from each taxi, of its geographical location and fare status, for a period of several weeks. This could be obtained by collecting a log file for each ICEV, where each record of the log file has a time stamp, the GPS location of the ICEV, and whether there is a paying passenger currently in the vehicle. Commercial systems that record this type of GPS data are readily available today [Geo14].

We organize the input dataset into a set of *shift files*, where a shift file represents data for one drivers working shift as defined in §7.2. We require a set of shift files for each driver and each taxi.

2. *Reduced Coordinate Space.* A second input to our model is a reduced geographic coordinate space that minimizes model dimensionality without overly affecting its correctness. We overlay the taxi company’s geographical operating region with a set of points in the reduced space and we map GPS data to its closest grid coordinate using Euclidean distances.
3. *Fare Pricing Model.* Every taxi company has their own pricing function they use to charge for fares, and this needs to be given as input. We use  $r_{\text{FARE}}(f)$  be the cost of fare  $f$ .
4. *Operating Costs.* We require gasoline, electricity, and vehicle prices in the taxi company’s region of operation.
5. *Vehicle Specifications.* We require specification of EV parameters such as battery size, range, and charging rates.
6. *Vehicle Replacement Rate.* We require the lifetime  $L$  of the company’s vehicles. Some taxi commissions require vehicles be replaced based on their age in years, while others require taxis be replaced after having driven a certain distance. Because this is a data-oriented model, for the latter case  $L$  can be estimated from the mean distance driven by the taxis each year.

### 7.3.3 Model Outputs

The process produces the following outputs:

1. The company’s ROI over  $L$  years, based on the fraction of the fleet transitioned BEVs or PHEVs, and the average lifetime  $L$  of the company’s vehicles.
2. Assuming a PHEV transition, the number of additional vehicles that must be purchased so that each driver can begin their shift with a fully-charged vehicle.
3. Assuming a BEV transition, the number of extra batteries the company must purchase and the number and location of needed battery switching stations<sup>3</sup>.

### 7.3.4 Estimating Charge Levels

We simulate the taxi company’s ICEV taxis as though they were EVs by estimating their SOC over time. Specifically, we develop a Bayesian network to estimate the charge level of a taxi at any time given the time-series of GPS coordinates from the corresponding ICEV. For the remainder of this chapter, we assume the reader is familiar with basic Bayesian networks; for an introduction, see Koller and Friedman [KF09]. We first present some necessary definitions:

- The set of parents  $\mathcal{P}(X)$  of a node  $X$  in a directed acyclic graph  $G(V, E)$  is defined as  $\{Y | (Y, X) \in E, Y \neq X\}$ .

---

<sup>3</sup>In the absence of switching stations, our PHEV analysis holds for BEVs that are charged.

- Nodes can be *observed* (we can directly observe or compute their values) or *hidden* (we estimate their value because we cannot observe their values). As shown later in this section, the only hidden node in our Bayesian network is the SOC node.
- Nodes can be either *discrete* or *continuous*; discrete variables can only take values from a countable set of values, such as the integers, whereas continuous variables can be any real number.
- A *Bayesian network* is a directed acyclic graph that defines the relationship  $p(X|\mathcal{P}(X))$  between every node and its parents. The probability of any variable  $X$  is dependent only upon its parents and is independent of all other nodes in the network.

We use a *dynamic conditional linear Gaussian network* (DCLGN) [KF09] to infer taxis' SOC. *Dynamic* refers to the ability to track hidden variables that change over time by inferring their value at discrete<sup>4</sup> *timeslices*. A *timeslice* is an instantaneous point in time and the  $k$ th timeslice is denoted  $t_k$ . These timeslices are spaced by a *timestep* which can be constant or variable. We use a variable timestep by using one timestep for every GPS measurement (which are not equally spaced). *Conditional linear Gaussian* refers to the assumption that the probability distribution of variable  $X$  is a Gaussian whose mean is a linear function of its parents' means:

$$p(X|\mathcal{P}(X)) = \mathcal{N}(\kappa_0 + \kappa_1 E[p_1] + \dots + \kappa_P E[p_P]; \sigma^2)$$

where  $\mathcal{P}(X) = \{p_1, \dots, p_P\}$ . We note that the variance  $\sigma^2$  is often assumed to be independent of  $\mathcal{P}(X)$ , and in many cases, it is assumed to be zero [KF09]. In this case only the expectations of each variable in the network are considered. In this chapter, we assume  $\sigma^2 = 0$  for all nodes in the network. This assumption can be removed as part of future work; see §7.6.

For further simplification, let  $\mathcal{P}(X, t_k)$  to denote  $X$ 's parents at time  $t_k$ . For computational simplicity, we assume variables follow the *Markov assumption*:

$$p(X, t_k) | [\mathcal{P}(X, t_0), \dots, \mathcal{P}(X, t_k)] = p(X, t_k) | [\mathcal{P}(X, t_{k-1}), \mathcal{P}(X, t_k)]$$

With respect to SOC, this means the SOC level of the taxi *SOC* at time  $t_k$  depends on the other variables in the network only at times  $t_{k-1}$  and  $t_k$ . This is a reasonable assumption because the SOC of a taxi at time  $t_k$  can be computed based on the SOC at time  $t_{k-1}$  and the change in SOC between  $t_{k-1}$  and  $t_k$ .

Figure 7.1 shows the model that is used to estimate the battery charge levels over time. We introduce *helper variables* that reduce the number of parents of variables we query. Table 7.1 shows the variables in our network, and whether they are observed, helper variables, hidden, discrete or continuous. The dotted arrows in Figure 7.1 represent variables that have an effect on the next timeslice called *persistence edges*. The solid lines represent *inter-time edges* that do not affect variables at the next timeslice.

<sup>4</sup>It is important not to confuse the relationship between continuous variables and discrete timeslices. Continuous variables in dynamic networks are real-valued but are observed at discrete timeslices. For example *time* is a continuous variable in our network even though we observe this variable at discrete timeslices.

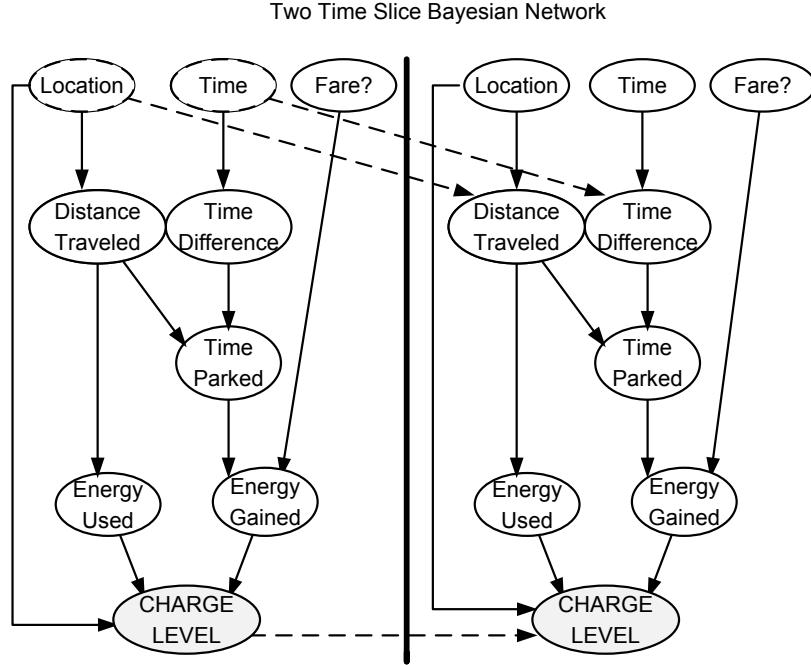


Figure 7.1: Bayesian network to simulate ICEV taxis as EVs by estimating their battery SOC

We explain the three most important variables:

- *Energy Used*. This variable represents the energy the taxi  $i$  consumes between two timesteps. Let  $D$  be the discharge rate of the EV (kWh/km), and  $e$  represent the most significant energy usage of an EV other than propelling the vehicle: air conditioning (AC). Then,

$$u(i, t_{k-1}, t_k) = d(i, t_{k-1}, t_k)D + e \quad (7.1)$$

The US National Renewable Energy Laboratory states “Air conditioning loads can reduce EV range and HEV fuel economy by nearly 40% depending on the size of air conditioner and driving cycle”. We discuss our AC assumptions in §7.5, and note that AC usage by taxi companies is different depending on climate in their regions.

- *Energy Gained*. This variable is only used when studying the effect of roadside charging on battery charge level. If a taxi is parked between two timesteps, we assume the taxi could have been charged during this time. Let  $g_x(i, t_{k-1}, t_k)$  be the energy gained by taxi  $i$  between two timesteps assuming level  $x$  charging (kWh), and  $B_{gx}$  is the amount of energy gained per second (kWh/second) assuming level  $x$  charging (this depends on the BEV or PHEV model). We assume drivers never charge with passengers in the vehicle:

$$g_1(i, t_{k-1}, t_k) = fare(i, t_k)p(i, t_{k-1}, t_k) \cdot B_{g1} \quad (7.2)$$

$$g_2(i, t_{k-1}, t_k) = fare(i, t_k)p(i, t_{k-1}, t_k) \cdot B_{g2} \quad (7.3)$$

- *Charge Level*. Charge level is the hidden node our network is designed to query. This variable has 4 parents: the charge level from the previous state, the current location, the energy used, and the energy gained. The edge between the two variables *charge level*

Variables			
Name	Symbol in Eq.	Meaning	Type
Location	$Loc(i, t_k)$	Taxi $i$ 's location at $t_k$	O, D
Fare	$fare(i, t_k)$ (0/1)	Whether $i$ has a passenger at $t_k$	O, D
Time		Value of time	O, C
Distance Traveled	$d(i, t_{k-1}, t_k)$	Km traveled by $i$ between last two timesteps $t_{k-1}, t_k$	H, C
Time Difference		Time elapsed between last two timesteps $t_{k-1}, t_k$	H, C
Time Parked	$p(i, t_{k-1}, t_k)$	Seconds $i$ parked between last two timesteps $t_{k-1}, t_k$	H, C
Energy Used	$u(i, t_{k-1}, t_k)$	Energy $i$ used between last two timesteps $t_{k-1}, t_k$	H, C
Energy Gained	$g_x(i, t_{k-1}, t_k)$	Energy $i$ gained between last two timesteps $t_{k-1}, t_k$	H, C
Charge Level	$SOC(i, t_k)$	$i$ 's charge level at $t_k$	Hd, C

Table 7.1: Table of variables in the Bayesian network.  $O$  = Observeable,  $H$  = Helper,  $Hd$  = Hidden,  $D$  = Discrete,  $C$  = Continuous

and *location* is because the *charge level* is dependent upon *location* because of battery switching (in the case of BEVs). We model  $E[SOC(i, t_k)]$ , the expected charge level at timestep  $t_k$ , for all non-battery switching scenarios as:

$$E[SOC(i, t_k)] = E[SOC(i, t_{k-1})] - u(i, t_{k-1}, t_k) + g_x(i, t_{k-1}, k) \quad (7.4)$$

When studying BEVs with switching stations, we model  $E[SOC(i, t_k)]$  as:

$$E[SOC(i, t_k)] = \begin{cases} \text{full if } (\Upsilon(Loc(i, t_k)) = 1) \wedge (E[SOC(i, t_k)] < \tau) \wedge (fare(i, t_k) = 0) \\ E[SOC(t_{k-1})] - u(i, t_{k-1}, t_k) + g_x(i, t_{k-1}, k) \text{ otherwise} \end{cases} \quad (7.5)$$

where  $\Upsilon(Loc)$  indicates that the location has a switching station; the optimization problem discussed in §7.4 decides this variable.

Note that one of  $u(i, t_{k-1}, t_k)$ ,  $g_x(i, t_{k-1}, t_k)$  will always be zero—either the taxi is parked and does not use energy or the taxi travels and uses energy.

### 7.3.5 Using The Model To Infer Costs

We now describe the process to use this Bayesian model to determine company's ROI in transitioning their fleet to EVs.

## Notation

In the following equations, we use  $i$  to index a specific taxi,  $w$  to index a specific shift,  $f$  to index a specific fare,  $b$  to index a specific battery switching station, and  $T$  to represent the total number of taxis the company has (in the dataset).

## Methodology

For BEVs and PHEVs respectively, our process is to compute:

$$r_B(\tau) = r_E - r_L(\tau) + s_B(\tau) - C_{bev} - C_{exbt}(\tau) \quad (7.6)$$

$$r_P = r_E + s_P - C_{phev} - C_{expv} \quad (7.7)$$

$$\Delta_B(\tau) = r_B(\tau) \cdot X - C_{BSS}(\tau) \quad (7.8)$$

$$\Delta_P = r_P \cdot X \quad (7.9)$$

where the terms are defined in Table 1. We now explain how we compute these costs.

## Determining Existing Taxi Revenue

We compute the company's existing revenue  $r_E$  using the fare data and the company's pricing model. Let  $\mathcal{T}_i$  be the set of all fares completed by taxi  $i$ . Using the input  $r_{FARE}$  from §7.3.2, the pricing function the company uses for a fare,

$$r_E = \frac{\sum_{i=1}^T (\sum_{f \in \mathcal{T}_i} r_{FARE}(f))}{T} \quad (7.10)$$

## Revenue Loss from Lost Fares

We now show how to compute the revenue loss due to transitioning to BEVs,  $r_L(\tau)$  (PHEVs do not have revenue losses as they use the ICE after battery depletion). We assume that if a taxi depletes its battery during a shift, all revenue the taxi would have generated during the remainder of that shift is lost. This upper bounds revenue losses because we assume drivers can only switch batteries if they are in a location with a switching station and they do not modify their paths to drive to a switching station. As a result of this worst-case restriction, the drivers may deplete their battery on their shift under our model. Let  $\mathcal{W}_i$  be the set of all shifts completed by taxi  $i$ . We determine the revenue loss by:

$$r_L(\tau) = \frac{\sum_{i=1}^T (\sum_{w \in \mathcal{W}_i} r_T^i(w, \tau) - r_B^i(w, \tau))}{T} \quad (7.11)$$

## Fuel Cost Reduction

Fuel savings are computed as follows:

$$s_B(\tau) = \frac{\sum_{i=1}^T \left( C_g \left( \frac{d_S^i}{V_E} \right) - C_e \left( \frac{d_B^i}{B_E} \right) \right)}{T} \quad (7.12)$$

$$s_P = \frac{\sum_{i=1}^T \left( C_g \left( \frac{d_S^i}{V_E} \right) - \left( C_e \left( \frac{d_E^i}{P_E} \right) + C_g \left( \frac{d_G^i}{P_G} \right) \right) \right)}{T} \quad (7.13)$$

where the variables are defined in Table 1. Variables  $d_B^i(\tau)$ ,  $d_E^i$ , and  $d_G^i$  come from the Bayesian network. We start the analysis with the first datapoint of the first shift for each taxi and assume the charge level of the vehicle is full. At each datapoint (GPS reading), we update the total distance driven by the taxi so far, and query the Bayesian network for the charge level of the vehicle. Assuming we are analyzing PHEVs, if the charge level ever reaches zero, then  $d_E^i$  is the distance driven to that point and  $d_G^i$  is the distance driven throughout the remainder of the shift. If we are analyzing BEVs, if the charge level reaches zero,  $d_B^i(\tau)$  is the distance driven to that point (then Equation (7.11) must be used to compute the revenue losses).

## Switching Station Infrastructure

Battery switching allows drivers to have a fully charged battery within minutes. This mitigates the range limitations of BEVs, assuming there are enough switching stations to service the taxi fleet. Switching stations have a large upfront cost, estimated to be \$500,000 by Better Place<sup>5</sup>, a manufacturer of EV switching infrastructure [Yar09, Gal09]. This does not take into account the cost of real estate. To provide an adequate coverage area, the fleet may need to be served by several switching stations spread across a city. Given the expense of switching stations, we want to find the minimal number and optimal location of stations to supply the fleet *without wasting money on buying unnecessary stations*. This problem can be stated as an optimization problem: given a set of taxis and the mobility data, find the optimal location(s) for switching stations such that the taxi company's profits are maximized. We discuss this problem in §7.4.

## Battery and Extra Vehicle Costs

This section presents the computation of the cost of batteries ( $c_{exbt}(\tau)$ ) and additional PHEVs ( $c_{expv}$ ). Taxis start each shift with a fully charged battery. This requires purchasing extra batteries to be kept at each switching station (BEVs) or storing extra PHEVs at the headquarters (PHEVs).

---

<sup>5</sup>We note that since the completion of this work, Better Place has dissolved. Other companies may design and manufacture switching infrastructure in the future.



Using Little’s law [Kle75b],  $c_{exbt}(\tau)$  and  $c_{expv}$  can be computed on a per taxi basis:

$$c_{exbt}(\tau) = \sum_{b=1}^{n_{BSS}} q_b(\tau) \cdot C_{bat} \quad (7.14)$$

$$q_b(\tau) = \lambda_b(\tau) \cdot B_{Fx} \left( \frac{\zeta - l_b(\tau)}{\zeta} \right) \quad (7.15)$$

$$c_{expv} = \lambda_H \cdot R_{px} \cdot C_{phev} \quad (7.16)$$

where the terms are defined in Table 1. Equation (7.14) multiplies the number of batteries needed at switching station  $i$  (per taxi) by the cost of each battery, and sums over all needed switching stations.

Note that  $l_b(\tau)$  and  $\lambda_b(\tau)$  are proportional to  $\tau$ . If  $\tau$  increases, batteries are switched with higher remaining capacity and take less time to charge. As  $\tau$  decreases, batteries are switched with lower remaining capacity and take more time to charge.

We assume additional PHEVs are kept only at the headquarters and drivers only switch PHEVs at the end of their shifts. We are not considering storing and charging PHEVs at the BEV battery switching stations. This is because it is less expensive to store batteries than vehicles—batteries can be stacked and stored in the same building but vehicles require expensive real estate for parking.

### Optimal Switching Threshold

The optimal value of  $\tau$  is unknown. We numerically evaluate  $\Delta_B(\tau)$  for each value of  $\tau$  in the set  $\{10\%, 20\%, \dots, 100\%\}$  and choose the value of  $\tau$  that maximizes  $\Delta_B(\tau)$ .

## 7.4 Switching Station Optimization Algorithm

In this section, we describe the optimization program we use to find optimal locations for BEV switching stations. Prior work [CCK13] shows that the problem is NP-hard, which implies that it is unlikely to be able to be solved by an algorithm that runs in polynomial-time, but the problem can be solved when the number of stations to place is small. Prior algorithms for locating fueling stations are divided into two classes: those assuming fueling stations are *discretionary facilities* (drivers only stop at them if they are en route to other destinations) [BLF90, BLF92] and those assuming drivers would deviate from their paths to a refueling station. Because we do not modify the trajectories of the taxis seen in the dataset, we treat refueling facilities as discretionary. Prior algorithms for locating discretionary alternative fueling stations are *flow interception models*—they place facilities along roads with high vehicle flows [SPV11, BLF90, BLF92, KL05]. This heuristic works well for siting public facilities because they intercept the most amount of traffic, again assuming drivers will not deviate from their paths. However, we propose a different algorithm for siting private

facilities because we are interested in maximizing *profit*. Placing facilities to maximize the number of intercepted flows does not necessarily maximize profit. Because we assume taxi drivers do not stop while serving a fare to switch their batteries, stations located where taxis park and wait without fares will be more utilized than stations alongside popular fare routes. We assume the taxi company’s objective is to maximize their overall revenue and introduce an optimization framework based on the discretized locations of the taxis and their charge levels.

We now formally describe the switching station location problem. Let  $\varphi(i, w)$  be the *opportunity cost* of  $i$  during  $w$ . If BEV  $i$  can complete all fares during  $w$ ,  $\varphi(i, w) = 0$ , otherwise,  $\varphi(i, w)$  is the sum of revenue lost from fares after  $i$ ’s charge level is too low to complete a fare (further fares during  $w$  are assumed lost). Our optimization problem, Program 7.2, is: *given the set of taxis’ historical mobility patterns, and a desired number of stations to place, find the optimal location(s) for switching stations such that the taxi company’s profits would have been maximized*. Note that this formulation takes the number of stations  $n_{BSS}$  as input. To optimize  $n_{BSS}$ , we search over  $n_{BSS} = 0, 1, 2, \dots$ , each time computing the ROI after considering the infrastructure pay back costs. When  $n_{BSS}$  is too small, adding additional stations increases the ROI, but at some point the diminishing returns of adding more stations does not outweigh the additional station costs. The inflection point is the optimal value of  $n_{BSS}$ .

When  $\mathcal{Y}$  contains only a few locations or  $n_{BSS}$  is small (as we found in our case study), we can optimally compute the switching station locations using brute force. This brute force approach may not be feasible over larger areas with more locations. In this case, it is possible to use heuristic algorithms to find a solution, though these heuristics cannot guarantee the optimality of their solution. Algorithms such as simulated annealing, tabu search, and hill climbing are general optimization methods, and could be used to find approximate solutions to the switching station location problem [RN03, GL97, KGV83].

## 7.5 San Francisco Case Study

We applied our process to a data set collected by Yellow Cab San Francisco (YCSF) as part of the Cabspotting project [The08, Cra09]. San Francisco is a regulated taxi market. The San Francisco Municipal Transportation Agency sets the fare price for all taxi companies in the region. In addition to controlling fares, San Francisco is a *medallion-based market*—the agency also controls when and how many new taxis can be integrated into the region. When this agency votes to allow new taxis, medallions are sold to companies allowing them to operate one new taxi; these medallions were (as of October 2012) sold for \$300,000 [Bay12]. Incidentally, in 2011 two medallions were issued specifically to operate battery switched electric taxis [CBS11].

This section is laid out as follows. In §7.5.1 we discuss our dataset and our preprocessing of the data. In §7.5.2 we discuss clustering GPS coordinates into a finite set of locations. We discuss assumptions specific to our case study in §7.5.3. In §7.5.4 we detail what EV

## Switching Station Placement Optimization

---

Inputs	$n_{BSS}$ : number of stations to place <span style="float: right;">(7.17)</span>
	$\mathcal{Y}$ : set of potential locations where a station can be placed <span style="float: right;">(7.18)</span>
	$\forall l \in \mathcal{Y}, \text{cost}(l)$ : price of placing a station at $l \in \mathcal{Y}$ <span style="float: right;">(7.19)</span>
	$\forall i, \mathcal{W}_i$ : set of all shifts completed by taxi $i$ <span style="float: right;">(7.20)</span>
	$\forall i, \forall w \in \mathcal{W}_i, w = \{t_{start}, \dots, t_k, \dots, t_{end}\}$ : start and end times of each shift <span style="float: right;">(7.21)</span>
	$\forall i, \forall w \in \mathcal{W}_i, \forall t_k \in w, \text{Loc}(i, t_k), u(i, t_{k-1}, t_k), \text{fare}(i, t_k), g_x(i, t_{k-1}, k)$ : (Table 7.1) <span style="float: right;">(7.22)</span>
Decision	$\forall l \in \mathcal{Y}, \Upsilon(l)$ : whether to place station at $l$ <span style="float: right;">(7.23)</span>
Objective	$\min_{\Upsilon(l)} \sum_i \sum_{w \in \mathcal{W}_i} \varphi(i, w)$ <span style="float: right;">(7.24)</span>
Sub. To	$\varphi(i, w) = \sum_{f \in w[\Omega(i, w), t_{end}]} r_{\text{FARE}}(f)$ <span style="float: right;">(7.25)</span>
	$\Omega(i, w) = \arg \min_{t_k \in w} E[\text{SOC}(i, t_k)] = 0$ † <span style="float: right;">(7.26)</span>
	$E[\text{SOC}(i, t_k)] = \begin{cases} \text{full if } (\Upsilon(\text{Loc}(i, t_k)) = 1) \wedge (E[\text{SOC}(i, t_k)] < \tau) \wedge (\text{fare}(i, t_k) = 0) \\ E[\text{SOC}(i, t_{k-1})] - u(i, t_{k-1}, t_k) + g_x(i, t_{k-1}, t_k) & \text{otherwise} \end{cases}$ <span style="float: right;">(7.27)</span>
	$\forall l \in \mathcal{Y}, \Upsilon(l) \in \{0, 1\}$ <span style="float: right;">(7.28)</span>
	$\sum_{l \in \mathcal{Y}} \Upsilon(l) \leq n_{BSS}$ <span style="float: right;">(7.29)</span>

† If the battery is depleted during shift  $w$ ,  $\Omega$  returns the time this first occurs,  $t_k$ . We then set  $\varphi(i, w)$  as the sum all fares that would have been completed starting with  $t_k$  until for the remainder of shift  $w$ .

---

*Figure 7.2: Switching station placement optimization*

adoption scenarios we examine. §7.5.5 gives the fare price structure ( $r_{\text{FARE}}$ ) for YCSF. The BEV and PHEV revenue analyses are discussed in §7.5.6 and §7.5.7 respectively. These two revenue analyses are compared in §7.5.8. Finally, we discuss how relevant results from YCSF would be to other taxi companies in §7.5.10.

### 7.5.1 Dataset and Preprocessing

The dataset includes the following information for 536 YCSF taxis during May 17, 2008 – June 10, 2008. Each measurement includes:

- Latitude and longitude to 5 decimal places
- Whether a paying passenger is inside the vehicle
- The current time of the data point

The average timestep between each data point is 60-90 seconds. As discussed in §7.3.2, we split the data from each taxi into shifts.

GPS devices sometimes report erroneous data, so we preprocessed the dataset to remove inconsistencies. For example, in some cases a taxi's position would be incorrectly reported between two correct readings. We noticed data from a taxi was either >99.5% correct or very erroneous due to a faulty GPS device in that taxi. For the taxis that only had few erroneous points we simply removed those points, whereas the taxis with many problems were simply discarded and excluded from all results. We discarded all data from seven out of 536 taxis.

## 7.5.2 Clustering Locations

We clustered the GPS coordinates in our data using the reduced coordinate space displayed in Figure 2. The clustering locations are spaced 4km apart with the exception of downtown San Francisco. For the downtown area, we used a denser grid (1km x 1km) because of the higher density of data within this region. After collapsing each GPS datapoint into its closest grid point, the GPS data was discarded.




Figure 7.3: Points represent taxi mobility data. The grid shows the reduced coordinate space. We used a denser grid in downtown San Francisco due to the large number of data points in this region.

## 7.5.3 Assumptions For BEV/PHEV Revenue Analysis

Here we state the assumptions we made while performing the revenue analysis.

- EVs may use up to 40% of their battery for AC while it is on [FR00]. We assume that ACs doubled in efficiency since 2000 and taxi drivers in San Francisco use AC 50% of the time—hence we assume 10% of the battery is used for AC.
- We assume Leaf batteries (purchased and stored at switching stations) cost \$450/kWh as discussed in §2.3. Therefore, with a 24kWh battery,  $C_{bat} = \$11,000$  in Eq(7.14).
- We assume the companies ICEV taxis have an efficiency of 25mpg, which is the national average efficiency of all new cars sold in the U.S. as of August 2014 [Rea14].
- Our data set indicates that each taxi is driven 112,000–177,000kms per year. Based on this figure, we assume the company replaces each ICEV, BEV, PHEV, and *battery* after four years of use ( $L = 4$ ). (This does not include replacing parts over the four years). Companies do not disclose their vehicle replacement rates, which makes estimating this figure difficult. However, the The Taxi and Limousine Commission of New York City states “Cars brought into service as taxicabs must be brand new vehicles and generally must be replaced five years after being placed into service. [Sch06].
- We assume the company spends \$17,000 for a new ICEV when replacing an old ICEV, using the average price of the non-luxury ICEVs discussed in §2.1.4. While taxi fleets probably receive discounts for buying vehicles in bulk, we do not apply discounts to EVs either—we simply assume the MSRP rate for both.
- At the time of writing, gas in San Francisco costs  $\approx \$1.08/\text{liter}$  and electricity costs  $\$.22/\text{kWh}$  [U.S14a]. These values are used as a basis for our ROI analysis but we show our results for a wide range of electricity and gas prices.
- We assume that switching infrastructure lasts 15 years.

#### 7.5.4 Case Study EV Scenarios

We studied ten different scenarios, as follows:

- (1-2) BEVs with Level 1 and 2 roadside charging only
- (3-4) BEVs with Level 1 and 2 roadside charging and battery switching
- (5-6) PHEVs with Level 1 and 2 roadside charging only
- (7-8) PHEVs with Level 1 and 2 roadside charging and PHEV switching at YCSF headquarters
- (9) BEVs with only battery switching
- (10) PHEVs with PHEV switching at YCSF headquarters only
- (11) ICEVs with increased efficiency (e.g., HEVs)

However, we found that scenarios with roadside charging (1—8) did not allow for much charging at all because the taxis were rarely parked. Taxis in our case study were parked only 12% of the time and were constantly driving the other 88% (we note this may be atypical, but this is indicated by our dataset). Consequently, even when we assume level 2 roadside

charging is available *everywhere* in San Francisco (an extremely unrealistic assumption), the results changed by less than 15% for both PHEVs and BEVs. Thus, for BEVs we only show results for scenario 9, and for PHEVs we show results for only scenario 10. In §7.5.9, we show the return on investment in purchasing ICEVs with higher fuel efficiencies.

### 7.5.5 Existing Taxi Revenue

As discussed in §7.3.2, each company has their own fare pricing model. For YCSF,  $r_{\text{FARE}}(f)$  is given on their website [Yel11]:

$$r_{\text{FARE}}(f) = 3.10 + .45 \cdot (p + (d - .2)) + 2A \quad (7.30)$$

where  $d$  is the distance of the trip in miles,  $p$  is the time the taxi was parked (at traffic lights) during the fare, and  $A$  is one if the passengers' destination was the airport and zero otherwise (the company charges an airport surcharge fee).

### 7.5.6 Revenue Analysis for BEVs

We now compute the change in revenue by using BEVs and a switching threshold of  $\tau$ ,  $\Delta_B(\tau)$  given current prices and vehicle specifications using Equations 7.8 and 7.9. First, we show how we compute the revenue losses,  $r_E$ , in §7.5.5. The cost of the BEV we study,  $C_{\text{bev}}$ , is derived in §7.5.6. In §7.5.6 we derive the cost of the battery switching stations,  $c_{\text{BSS}}(\tau)$ , and show its relationship to  $r_L(\tau)$ . In §7.5.6 we measure the relationship between the threshold  $\tau$ ,  $r_L(\tau)$  and the cost of extra batteries needed,  $c_{\text{exbt}}(\tau)$ . Roadside charging is briefly discussed in §7.5.4. We incorporate the revenue loss  $r_L(\tau)$ , the fuel savings  $s_B(\tau)$ , and  $c_{\text{exbt}}(\tau)$  in §7.5.6 which also computes the overall return on investment  $\Delta_B(\tau)$ .

#### Nissan Leaf Specifications

For our BEV experiments, we study the Nissan Leaf. The Leaf does not have a switchable battery, but vehicles with similar attributes with switchable batteries may be sold eventually. As discussed in §2.1.1, the price of the Leaf at the time of writing (2014) is \$21,500. so under our assumption that the company replaces ICEVs for \$17,000,  $C_{\text{bev}} = \$4,500$ .

Even though manufacturers list the full capacity of a battery, the full capacity is not actually used—the battery is not fully charged or discharged to preserve the life of the battery [BCMW11]. However, manufactures list the expected range based on the *usable* portion of the battery—the figure we are interested in. Table 7.2 gives the values of the constants needed for our revenue analysis for the Leaf. These figures were derived from the specifications given on their website [Nis14b].

<i>Constant</i>	<i>Value for Nissan Leaf</i>
$D$ in Eq. 7.1 (kWh/km)	0.15
$B_{g1}$ in Equation 7.2 (kWh/s)	.00033
$B_{g2}$ in Equation 7.2 (kWh/s)	.001
$B_{f1}$ in Eq. 7.15 (days/battery)	.83
$B_{f2}$ in Eq. 7.15 (days/battery)	.29
$\zeta$ in Eq. 7.15 (kWh)	24
$C_{bev}$ in Table 1	\$4,500

Table 7.2: Nissan Leaf Specifications [Nis14b]

### Switching Station Location and Distribution Over Locations

We now calculate the cost of battery switching stations,  $c_{BSS}(\tau)$ , and show its relationship to the revenue losses incurred,  $r_L(\tau)$ . We find the locations of switching stations by applying the algorithm presented in §7.4. We find that three switching stations are optimal, so  $c_{BSS} = \$1,500,000$ .

The relationship between the revenue loss  $r_L(\tau)$  and the cost of battery switching stations  $c_{BSS}(\tau)$  is shown in Table 7.3. Without battery switching, even if we assume charging infrastructure is available everywhere (i.e., whenever a taxi is stopped, its battery charges while it is parked), a third of all fares are lost. However, with additional switching stations at the San Francisco airport and Yellow Cab headquarters, only 3% of fares are lost. Adding additional stations to these three has negligible impact on  $r_L(\tau)$  but greatly drives up  $c_{BSS}(\tau)$ ; three stations represents the optimal value for YCSF.

We find the distribution over all locations where fares began and ended to gain intuition as to why three stations are adequate. Figure 7.4 shows this distribution. We find approximately 90% of all pick-ups and drop-offs occur in only 20% of the locations. This explains why a small number of switching station locations suffice; switching stations near these locations will be heavily used.

No Charging or Switching	41.5%
L2 Roadside charging only	37 %
Union Square BSS (no charging)	15%
YC, Union Square, Airport BSS	3%

Table 7.3: Percentage of fares lost in different BEV scenarios. BSS: battery switching stations.

### Switching Threshold Analysis

We now analyze the ROI  $\Delta_B(\tau)$  as a function of  $\tau$  as discussed in §7.3.5. Figures 7.5 and 7.6 show  $r_B(\tau)$  vs.  $\tau$  for a fixed electricity price (\$.22/kWh) and a fixed gas price (\$1.08/l).




Figure 7.4: Distribution over pick-up and drop-off locations as distribution. Black shows the probability (percentage) of a pickup, red shows the probability of a dropoff, but the two distributions are nearly identical. The airport and Union square locations are annotated with "A" and "U" respectively. For YCSF, annotated with "Y", we used an identifier of 1.

For our case study, we find that the revenue per BEV monotonically increases as a function of  $\tau$ . In the absence of charging (switching only), then, BEVs should switch whenever they are at a switching station regardless of their SOC if a battery is available. This is because their SOC monotonically decreases throughout their shift so future switching opportunities necessarily occur at lower  $\tau$  values leading to lower revenue gain or further revenue loss.




Figure 7.5: Threshold  $\tau$  vs.  $r_B(\tau)$  as a percentage of  $r_E$  (left Y-axis),  $r_B(\tau)$  (right Y-axis) for varying gasoline prices, electricity fixed at \$.22/kWh.




Figure 7.6: Threshold  $\tau$  vs.  $r_B(\tau)$  as a percentage of  $r_E$  (left Y-axis),  $r_B(\tau)$  (right Y-axis) for varying gasoline prices, gas fixed at \$1.08/liter.

Increasing the threshold increases the taxis' average charge levels, because their switched batteries have a higher remaining SOC. This increases the fuel savings  $s_B(\tau)$  and decreases the revenue loss  $r_L(\tau)$ . Moreover, the cost of extra batteries  $c_{exbt}$  needed to meet this additional switching load grows slowly due to a combination of Little's Law, given in Eq(7.15), and multiplexing. Specifically, Little's law states that the number of batteries in the system is given by the rate at which batteries "enter the system" (are swapped and need to be charged) multiplied by the time batteries "remain in the system" (are required to charge).



$\tau$	Swaps	remaining battery SOC after swap at $b$ : $l_b(\tau)$	rate taxi swaps at $b$ : $\lambda_b(\tau)$	time required to charge battery: $B_{Fx} \left( \frac{\zeta - l_b(\tau)}{\zeta} \right)$	extra batteries needed: $q_b(\tau)$
10	6	1638.95	0.25722	0.15528	0.03994
20	15	3663.16	0.64305	0.14122	0.09081
30	21	6077.7	0.90027	0.12446	0.11204
40	27	8181.24	1.15749	0.10985	0.12715
50	32	10610.1	1.37184	0.09298	0.12756
60	41	12607.5	1.75767	0.07911	0.13905
70	50	14864.5	2.14350	0.06344	0.13598
80	76	17316.1	3.25813	0.04641	0.15122
90	131	19952.3	5.61599	0.02810	0.15785

Table 7.4: Example showing how  $\tau$  affects the number of battery swaps performed and the remaining SOC of swapped batteries at the airport switching station for one taxi.  $q_b(\tau)$  (Eq(7.15)) is computed assuming the 24kWh can charge in four hours from depleted using level 2 charging, i.e.,  $B_{Fx} = 4/24$ . With respect to the last column, the operator cannot buy a fraction of a battery, but the fractional batteries needed in the system per taxi per station are summed to compute  $c_{ext}$ .

As  $\tau$  increases, the rate at which batteries are swapped,  $\lambda_b(\tau)$  in Eq(7.15), increases. However, when batteries are swapped more often, the remaining SOC of swapped batteries increases, so the time required to charge them,  $B_{Fx} \left( \frac{\zeta - l_b(\tau)}{\zeta} \right)$  in Eq(7.15), decreases. The number of extra batteries needed  $q_b(\tau)$ , the product of these inversely-changing rates, increases slowly. For intuition, Table 7.4 shows how these parameters change with  $\tau$ , for a randomly selected taxi at the airport station as  $\tau$  increases. In summary, the cost of the extra batteries is recouped with lower operating costs and fewer lost fares.

## Overall BEV Transition Cost

Figure 7.7 shows the cost to transition each individual ICEV to a BEV ( $r_B$ ) for a wide array of gas and electricity prices, *without* taking into account the cost of switching stations. For any gasoline-electricity pair (X and Y axes) for which the value of  $r_B$  is positive (Z-axis), the company can begin paying back the switching station costs. We now compute  $\Delta_B(\tau)$  as in Eq(7.8). We are interested in the "break even" point where  $x$  BEVs can be operated for the exact cost that  $x$  ICEVs can, including all switching station costs. Under our assumption that switching infrastructure lasts 15 years, we amortize the \$1.5M cost of the three stations accordingly, yielding a cost of \$100k/year and \$400,000 over the 4 year replacement period (that  $r_B(\tau)$  is computed over). Thus, we compute:

$$\Delta_B(\tau) = r_B(\tau) \cdot x - 400000 = 0$$

Figure 7.9 shows the when  $\Delta_B(\tau) = 0$  for a fixed electricity price of \$.22/kWh. For a given point on the line, if gas prices rise, the company accrues profits. At current prices, the company needs to operate at least 30 BEVs to be profitable.




Figure 7.7: Gas and electricity prices vs.  $r_B(\tau)$  using three switching stations. Figure does not take into account cost of switching stations, but takes into account the cost of extra batteries  $c_{exbt}$ .




Figure 7.8: Gas and electricity prices vs.  $r_P$  as a function of  $r_E$ ,  $r_P$ . No roadside charging. All costs taken into account, including the cost of additional PHEVs  $c_{expv}$ .




Figure 7.9: Region where BEVs are profitable after amortizing three switching stations. Electricity fixed at \$.22/kWh. The red line shows the current gas price in San Francisco (\$1.08/liter).

### 7.5.7 Revenue Analysis for PHEVs

We now compute the costs of switching to PHEVs instead of BEVs. We study the Chevrolet Volt. As discussed in §2.1.2, the price of the Volt at the time of writing is \$26,685, so we use this as the cost of extra PHEVs needed at the headquarters,  $P_C$  in Eq(7.16). Under our assumption that the company replaces ICEVs for \$17,000, the incremental cost  $C_{phev} = \$9,685$ . Table 7.5 gives the values of the constants needed for our revenue analysis

for the Volt. These are derived from the specifications given on their website [Che14a].

Figure 7.8 shows the ROI per PHEV,  $r_P$ , without any roadside charging. This figure includes the cost of additional PHEVs needed,  $c_{expv}$  in Eq(7.16). At current prices, PHEVs are less expensive to operate than ICEVs. Moreover, there are no infrastructural investments, so the cost of transitioning  $x$  ICEVs to PHEVs,  $\Delta_P$ , is simply given by  $r_P x$ . The company can therefore switch to PHEVs on a per vehicle basis.

### 7.5.8 PHEV vs. BEV Comparison

PHEVs and BEVs are both profitable in our analysis under a range of gasoline and electricity prices. However, it is difficult to project which EV will be more profitable for taxi companies in coming years because many possible trends can affect their relative profitability:

- BEVs have larger batteries than PHEVs. Battery prices are projected to continue falling, as discussed in §2.3, and if this happens, the price of BEVs will fall faster than PHEVs (which still have ICEs). However, if this projection turns out incorrect, and battery prices increase, for example due to a Lithium shortage, this trend may reverse.
- Figure 7.10 shows gasoline prices per gallon (one gallon = 3.78 liters) since 1971 adjusted for inflation [Zfa10] and the average price of electricity in California (adjusted for inflation) since 1980 [Tom07]<sup>6</sup>. These historical trends suggest that BEVs will be more profitable as fuel prices rise, which is displayed by their higher potential-ROI shown in Figures 7.7 and 7.8. However, a sharp (unprecedented) decline in gasoline prices would significantly reduce or eliminate the profitability of BEVs.
- In our analysis we assume that switching stations cost \$500,000, which was the price estimated by Better Place at the time the journal publication of this chapter [CCK13] was written. Unfortunately, in 2013 Better Place declared bankruptcy [Cha14a]. While other EV manufactures have stated intention to produce swappable BEVs, e.g., Tesla [Tes14c], if this does not occur, BEVs may be unprofitable for taxi companies.

### 7.5.9 ICEVs With Increased Efficiency, HEVs

Finally, we show the return on investment in simply purchasing ICEVs with higher fuel efficiency. Recall that for all results thus far, we assumed the company's ICEV taxis have an efficiency of 25mpg, which is the national average efficiency of all new cars sold in the U.S. as of August 2014 [Rea14]. Figure 7.11 shows the change in revenue per taxi assuming the company purchases ICEVs with higher fuel efficiencies under different cost assumptions. Recall from §2.1.3 that HEVs are ICEVs with a small, nonchargeable battery that increases

---

<sup>6</sup>Note that Figure 7.10 does not show a price of \$.22/kWh electricity (the figure used throughout this chapter) because electricity prices in San Francisco are nearly double than in the rest of California—see the U.S. Bureau [U.S14a] for a comparison—and we could not find a long history of electricity prices in San Francisco alone.

<i>Constant</i>	<i>Value for Chevrolet Volt</i>
$D$ in Eq. 7.1 (kWh/km)	0.25
$B_{g1}$ in Equation 7.2 (kWh/s)	.0004
$B_{g2}$ in Equation 7.2 (kWh/s)	.0011
$R_{p1}$ in Eq. 7.15 (days/PHEV)	.45
$R_{p2}$ in Eq. 7.15 (days/PHEV)	.16
$\zeta$ in Eq. 7.15 (kWh)	16
$C_{phev}$ in Table 1	\$9,685

Table 7.5: Chevrolet Volt Specifications [Che14a]




Figure 7.10: Left: Average U.S. price per gallon of gas since 1971 (2011 dollars) [Zfa10]. Right: Average electricity price given by three major utilities in California since 1980 [Tom07]

the fuel efficiency of the vehicle, thus HEVs are analogous to high-efficiency ICEVs. At the time of writing, the best selling HEV, the Toyota Prius, has an efficiency of 51mpg and an MSRP of \$24,000. This corresponds to an efficiency of 21.58km/l and an incremental cost of \$7,000. The return on investment for a vehicle with these specifications is shown at the red arrow, which corresponds to an ROI of \$28000. Thus, at current prices, it is more profitable to invest in HEVs than BEVs or PHEVs. Note, however, that HEVs are still completely petroleum dependent, thus they represent a profitable but non-sustainable (though to a lesser extent) alternative to standard ICEVs. Individual taxi companies should decide whether to invest in BEVs, PHEVs, or HEVs, which are all profitable at today's prices, based on their relative merits (profitability v.s. sustainability).

### 7.5.10 Sensitivity Analysis

We now analyze whether our case study results would generalize to taxi companies in different cities. Although we cannot draw definite conclusions without re-running our study for a taxi company in a different city, we attempt to answer this question here.

1. *Average Trip Length and City Density.* Switching infrastructure is expensive, so it will initially be sparsely deployed, in contrast with current petroleum infrastructure. Con-




Figure 7.11: Return on investment, per ICEV, in purchasing higher efficiency ICEVs. Gasoline fixed at \$1.08/l. The X-axis shows the assumed cost of acquiring higher efficiency ICEVs. The Y-axis shows the fuel efficiency, where 10.58km/l corresponds to 25mpg and 25.4km/l corresponds to 60mpg. The central point shows the case assumed in all prior results—an efficiency of 10.58—so the relative return is 0. The red arrow shows the return ( $\approx$ \$28000) of the 2014 Toyota Prius.

sequently, the geography of a city affects the feasibility of BEVs. Large cities with widespread points of interest are less suitable than dense cities with concentrated points of interest. One way we can measure this for a given city is to determine the distribution over fare trip lengths. We can use the distribution of how far people commonly travel as a heuristic to estimate how many switching stations will be needed. Figure 7.12 shows the distribution of trip lengths for all fares the YCSF taxis completed during the study period. From the cumulative density function of this distribution, we find 85% of all fares are less than 10 km, which is why few stations are needed in San Francisco. This figure shows a two-peaked distribution. From the probability mass function, we find 8% of the fares are between 20 and 30 km (roughly 7% of all trips are to the San Francisco International Airport, which is 24 km by highway from Union Square in downtown San Francisco). The average trip length can also help us determine whether PHEVs or BEVs are better suited for a region. For a PHEV<sub>xxm</sub>, it does not matter how many trips are completed before battery depletion, because the financial benefit comes from the fuel savings on the first xx km. BEVs are range limited, however, and completing a large number of short trips (before battery depletion) is more profitable than a short number of long trips, due to the initial charge to each passenger that requests a fare. Therefore, PHEVs are likely better suited for cities with many long trips, whereas BEVs will be more

profitable in cities with a large number of short trips, like San Francisco.

2. *Distribution over locations.* Closely related to the distribution over trip lengths is the distribution over locations discussed in §7.5.6. We found roughly 90% of YCSF fares start or end at fewer than 20% of the grid locations. If this distribution was less concentrated, the average trip length shown in Figure 7.12 may have increased. In our case study, trips in the downtown area within 3 km of Union Square accounted for more than half of all trips. This is highly conducive to centralized switching station placement. Taxis in larger cities may find they have to travel a greater distance out of their way to refuel.
3. *Gas and electricity prices.* Although gas prices in San Francisco (\$1.08/liter) are higher than the rest of the United States, they are lower compared to the rest of the world, e.g., gas prices across Europe are significantly higher. If the mobility patterns of taxis in these regions are similar to those in San Francisco, transitioning to EVs would be even more profitable. We also note that electricity prices in San Francisco are twice the United States national average, while gas prices are not [Uni10].
4. *Temperature and weather.* Reference [SB08] shows that at cold temperatures ( $< 0^{\circ}\text{C}$ ), over 10% of the energy in a battery is lost compared to at  $21^{\circ}\text{C}$ . It rarely snows or drops below freezing in San Francisco, even during the winter months, but taxi companies in cities with colder climates should expect worse performance. Furthermore, passengers in cities with extreme weather temperatures require more heating and cooling, which further drains the battery.




Figure 7.12: Distribution of fare trip lengths (the bar for 50 represents all trips over 50km)

## 7.6 Conclusions and Future Work

In this chapter, we proposed a process to determine the ROI for a taxi corporation transitioning to electric vehicles. We first model taxi mobility and then used the model to compute the economic costs of the transition. The model can be configured with a wide array of

input parameters, including the type of vehicle to be tested, electricity and gasoline prices, and roadside charging/battery switching infrastructure assumptions. We then used our approach to analyze a fleet of over 500 taxis in San Francisco. We found that PHEVs, BEVs and HEVs are all currently profitable.

To obtain realistic results for our case study, we used only commercially available vehicles and their manufacturer specifications. However, predicting the outcome of a major transition prior to it occurring is an error-prone process. We now discuss some avenues for future work.

1. In §7.3.4, we assume that the variance of all nodes in our Gaussian network is zero. Future work can include extending this to a stochastic model where  $\sigma^2$  is calculated based on the co-variances between variables in the network.
2. We have not accounted for vehicle maintenance costs. Maintenance costs for a fleet of EVs is thought to be lower than ICEVs [BCMW11], but we are not aware of any quantitative analysis comparing the two. A maintenance cost analysis for a large fleet of PHEVs/BEVs would greatly improve our cost model.
3. Our switching station optimization assumes that the locations can charge any number of batteries and can be placed anywhere in the city. In reality, distribution network limitations may place some restrictions on switching station placement and battery charging; areas with a fully utilized distribution network may not be able to accommodate the new load.
4. We have not considered real estate prices for switching stations, other than the cost of the stations themselves. We should account for the cost of acquiring space to build the switching station. We also have not modeled swapping station maintenance costs.
5. Obtaining a second data set from a different city would provide a better foundation for the sensitivity analysis section.
6. Batteries do not charge at a constant rate as we have assumed. A better assumption would be to use a two phase linear approximation; have a higher charge rate while the state of charge (SOC) is less than 80%, and a lower rate when the SOC is above 80% [Nis14b].
7. Queueing delays at switching stations due to multiple taxis attempting to switch their batteries in parallel should be modeled.



# Chapter 8

## Designing A Trial To Explore the Barriers to Electric Bike Adoption

The field trial described in this chapter was co-engineered by Costin Ograda-Bratu, Rayman Preet, Milad Khaki, Tobias Schroeder, and S. Keshav. The trial components that were mainly engineered by these co-authors, and not by the thesis author, have been removed from this thesis.

---

### Synopsis

There has recently been a large uptake of electric bicycles (eBikes) in many densely populated cities where vehicle travel is inconvenient (due to traffic) or costly due to government regulations (usually in effort to reduce vehicle emissions [Par06]). For example, Chinese manufacturers estimate there are now 200 million eBikes in use in China [Tim13], representing a 40% increase from 2010 estimates of 120 million [Goo10]. Likewise, in the EU, eBike sales have increased tenfold in the last decade [Bak13]. However, their adoption currently varies greatly by region—in contrast, of the  $\approx 100$  million bikes sold in the US in the last decade [Nat13], only 1% (1 million) were electric [Bak13]. The large disparity in sales between regions shows that, at least in North America, there are *barriers to eBike adoption*. This chapter focuses on the discovery and study of such barriers.

In §3.1, we discussed that EV field trials have thus far been problematic because EVs are expensive. When trial conductors (e.g., University researchers) must lease expensive vehicles, the vehicles are typically shared among participants for a short duration each. This limits both the number of participants and the duration for which participants can evaluate their vehicles. This in turn leads to less reliable conclusions; because perceptions may change with experience, it is difficult to predict whether the participants would have had the same



sentiments if given the EVs for longer, and moreover, it is clearly better to obtain opinions from larger populations.

However, these problems are mitigated when the vehicles being leased or purchased are relatively inexpensive. We therefore are conducting a field trial with 31 eBikes. The eBikes used in our field trial cost approximately \$1200 (after bulk discounts) and each telemetry kit costs \$800, allowing us to conduct a trial with 31 bikes for \$65,000. For comparison, for \$65,000, only three EVs could be purchased or five EVs could be leased at \$350/mo for the three years<sup>1</sup>.

Given the lack of eBike sales in North America, and that problems with field trails are mitigated by using less expensive vehicles, we *use a field trial methodology to study eBike usage in a North American context*. In this chapter, we present our design of the most comprehensive eBike field trial to date—a three year trial of eBikes in the Waterloo region—named “WeBike”. Specifically, we discuss:

1. the motivating research questions for designing and operating an eBike field trial (§8.1)
2. the design of WeBike, including the design of custom eBike sensor kits (§8.2)
3. how we estimate the eBikes’ state of charge (SOC) levels, a necessary first step in answering many research questions (§8.3)
4. our algorithm for detecting cycling trips from GPS traces (§8.4)
5. preliminary insights obtained from the participant selection survey (§8.5) and the data collected from the eBikes (§8.6)

WeBike uses a fully automated data collection and transmission process; it does not rely on travel diaries or any other user action other than keeping the eBikes charged, in contrast to many prior EV and eBike field trials (see §3.1.1) requiring the user to collect or upload data. As discussed throughout this chapter, fully automating the secure collection and transmission of data requires custom hardware, software, and solutions to many technical problems.

Currently, our trial consists of 31 bikes. However, these 31 bikes are being tested as a proof of concept for a much larger 500-eBike trial.

## 8.1 Motivating Research Questions

In this section, we discuss the primary motivations for conducting the WeBike trial. The three-year trial has only been under operation for a few months at the time of writing, and the questions presented below will not be answered until a significant amount of data has been collected. However, we pose these questions to demonstrate why and how we designed WeBike to understand the adoptability of eBikes. The remainder of this chapter is about

---

<sup>1</sup>We give the eBikes to our participants to keep after three years as an incentive, but because we have purchased the eBikes, we could have conducted the trial “indefinitely” (until hardware failure).

how we designed WeBike to be able to study the questions posed here, which primarily lie in three spaces: coordinating usage data with preconceived barriers (§8.1.1), measuring parameters of interest to EV researchers and manufacturers (§8.1.2), and human-computer interaction (§8.1.3).

### 8.1.1 Coordinating Usage Data With Preconceived Barriers

One of the main goals of WeBike is to study whether the participants' eBike usage differs from their preconceived notions. We measured participants' preconceived notions using a survey (see §8.2.4) that was given prior to participant selection. By combining our pre-trial survey with the sensor data, we hope to answer the following:

- Do people use eBikes more or less than they thought they would? Did users use their eBikes during conditions where they thought they would not, and visa versa? These questions are important for separating the *perceptual barriers* to eBike adoption, i.e., barriers similar to range anxiety that prevent adoption in some cases where adoption is possible, and the *actual barriers* to eBike adoption, i.e., the scenarios where eBikes do not fit a given individual's mobility needs.
- One can envision three potential target populations for eBikes: 1) those transitioning car trips to eBike trips, 2) those transitioning public transit trips to eBike trips, and 3) those transitioning from bicycles hoping to travel further or in more hilly areas. Are there noticeable differences in the usage of eBikes across these populations? E.g., do users that transitioned their trips from car trips tend to be shorter than those who transitioned from cycling?
- On a similar note, can we cluster eBike users into different usage categories? Do some use their eBikes primarily for commuting or shorter trips while others use the eBikes for longer trips (possible for health/exercise), or do users have similar usage patterns?
- Finally, can we use the survey data collected from participants to predict their actual usage? This question is posed by a co-author of our participant selection survey (§8.2.4) and the co-creator of a neural-network model of transportation choice [WSNdH14]. The model, depicted in Figure 8.4, poses transportation choice as a constraint-satisfaction problem, where different means of transportation facilitate multiple, sometimes competing objectives. Decision-making under the model is partly deliberative, consisting of beliefs about the degree to which certain models of transportation are compatible with the different goals, and and partly emotional, implemented using an emotional coherence algorithm [Tha06]. We can use parameters from our participant selection survey (specifically item (2) in §8.2.4) to parameterize the model to predict the potential adoption of eBikes. If these parameters correlate with the actual usage of these bikes, this would help validate the model to predict actual transportation behavior.

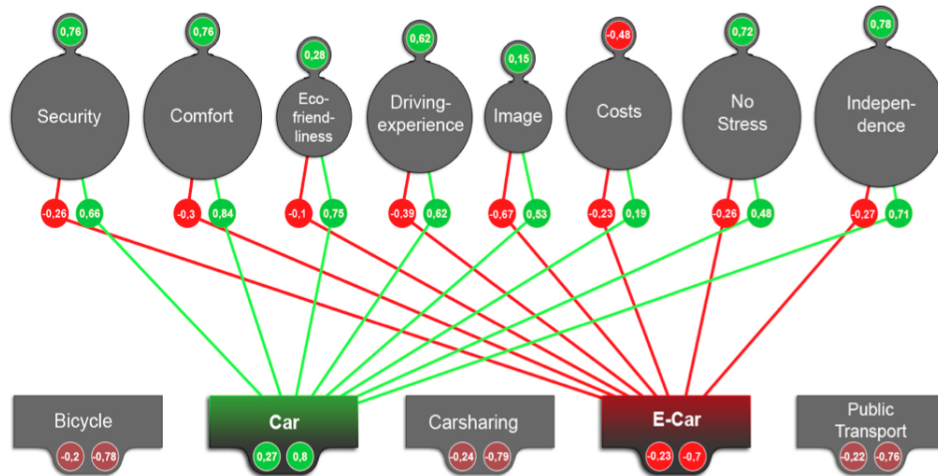


Figure 8.1: Neural Network for predicting EV usage & adoption based on survey input parameters.

### 8.1.2 Questions Relevant to EV Researchers/Manufacturers

The second main objective of WeBike is to use eBikes as a “Wind tunnel”—an inexpensive way to draw some conclusions about EVs. As discussed in §2.2, eBikes are isomorphic to PHEVs and use the same battery technology. Given that eBikes are  $\approx 1/10$ th the price, we can therefore study some aspects of EVs at low cost. Some examples are given below:

#### Range anxiety:

- At what SOC do people normally start charging?
- Do participants deplete or nearly deplete their eBike batteries? I.e., are they pushing their range limits?
- Do we observe the “range paradox” reported in many EV field trials and surveys (§3.1)? I.e., is there a mismatch between the users’ range needs and their perceptions or charging behavior?

#### Parking habits/charging:

- What is the temporal distribution of parking/charging events?
- Where are the “hotspots” for eBikes in a University city? Can this help us site PEVSE?
- What are the observed charging losses? That is, for each kWh of energy that is drawn from a charging outlet, what percentage is stored into the eBike battery?

#### LiION battery properties:

- What are the effects of different drive cycles and different levels of electric assistance on battery life/range?
- How does battery life/range depend on temperature?
- How much does the battery capacity degrade over its lifetime?

The WeBike data that can be used to study these questions (among others) is given in the Mind Map in Appendix §E. The data necessary to study all of these questions is collected during WeBike.

### 8.1.3 HCI Questions

There are many questions regarding what information to display (if any) to eBike users while they are biking:

- Should the information be displayed audibly or visually? For example, we could mount a small screen on the handlebars and stream information to it, or we can use a speaker to display information audibly.
- What information would be most appealing to users? Should they be health focused (e.g., calories burned), environment focused (e.g., CO<sub>2</sub> offset v.s. a car), or logistically focused (e.g., traffic conditions or route planning).
- Regarding audible information, how often should we refresh? I.e., should the information be given at fixed time or fixed distance intervals? (Visually, the screen can be refreshed frequently.)
- Should we have any user input similar to e-bikeSAFE<sup>2</sup>?
- What is the effect, if any, of displaying this information to users on their behavior? As examples, if we start displaying health information to users after a period of time without doing so, will they use the electric assistance less afterwards? If we alert the user that they are entering a region with heavy car traffic, do users reroute?

## 8.2 Field Trial Design

Here we discuss several components of WeBike, specifically, the eBikes (§8.2.1), our sensor kits (§8.2.2), the data pipeline (§8.2.3), our participant selection process (§8.2.4), and participant incentivization (§8.2.5). Prior work on conducting eBike trials is described in §3.1.1.

### 8.2.1 The eBikes

The eBikes selected for WeBike are 2014 eProdigy Whistlers [ePr14], shown in Figure 8.2. The 2014 Whistler model can provide electric assistance of up to 32kmph. The rider can select from five different levels of electric assistance, allowing the user can pedal as much or as little as they desire (provided the battery is not depleted). The LiON battery provides up to 40km of electric range, depending on terrain. With respect to degradation, the manufacturer claims the batteries retain “80% capacity after 1000+ cycles”, but we plan to evaluate this claim with future data.

The sensor kit is attached to the battery, and the battery is removable, as shown in Figure 8.3. This allows participants to park their eBike and bring their battery assembly with them

---

<sup>2</sup>See Dozza [DWM13] in §8.7.

to charge at any standard outlet (e.g., inside their home or office). Using a standard 110V AC outlet, the battery takes about five hours to charge<sup>3</sup>.

## 8.2.2 The Sensor Kit

Here we describe the sensor kit we engineered that is mounted on each eBike. To the best of our knowledge, we are the first to build a telematics system specifically for eBikes, though several researchers have designed telematics systems for standard bicycles (see §8.7).

At the minimum, to collect the data we need in order to answer our motivating research questions, our WeBike sensor kit has to be able to:

1. store and upload all data collected from each of our sensors (see see §8.2.3)
2. record GPS, as it is used to compute speed, distance traveled, and location
3. measure charging times and charging rates, which are necessary for all charging-based research questions in §8.1.2 and Appendix §E.
4. measure battery voltage and battery temperature, so that SOC can be inferred (see §8.3), which is necessary for all research questions relating to battery life and range

Upon studying these requirements, we realized that modern smartphones represent an elegant solution to (1) and (2). Modern smartphone can store many gigabytes of data, can transmit data either over a cellular network and/or wifi, and come with a plethora of internal sensors—15 for the phone we use—including GPS. Moreover, modern smartphones can act as a host device—we can use it to read data coming from other sensing devices, which we need because the smartphone cannot handle (3) and (4). For these tasks, we use a collection of Phidget sensors<sup>4</sup> and Digikey temperature sensors<sup>5</sup>. Collectively, this led to the sensor kit shown in Table 8.1. The Phidget sensors are connected to an interface kit, which acts as a bus. This interface is then connected to the Android smartphone as a slave device. The Android smartphone hosts the Phidget bus and reads and stores data from all Phidget sensors through the interface kit. The Android smartphone, Phidget sensors, and all required cables are packaged into a weatherproof box as shown in Figure 8.3. Full assembly guides are available online [ISS14a].

## 8.2.3 Data Transmission and Analysis

Figure 8.4 shows our data pipeline, specifically, how hardware and software enable the collection, transmission, and processing of data. The green box represents the client (eBike) side and the red box represents the server side. Blue nodes represent raw data, orange

<sup>3</sup>A 110V AC outlet, which is capable of charging much larger batteries in the same amount of time (see “level 1 charging” in §2.4). This rate is reduced by the eBike battery controller to prevent battery problems.

<sup>4</sup><http://www.phidgets.com/>

<sup>5</sup><http://www.digikey.ca/product-search/en?x=0&y=0&lang=en&site=ca&KeyWords=lm335zns%2Fnopb>




Figure 8.2: The eProdigy Whistler with our sensor kit attached.

nodes represent processing of the data, grey nodes represent hardware, and the yellow node represents the transmission of data from client to server via wifi. Black arrows represent flow. Currently, data only flows from the client side to the server side, but the single dashed arrow labeled "f" for "future" represents the flow of data *back* to the eBikes (recall §8.1.3).

On the client side, data is collected from the sensing kit (currently 10 times per minute) and stored in a file. Every hour, the file is named according to the current date and time, and stored in the kit for transmission. When the eBike user brings the bike to campus, the kit connects to the Eduroam campus-wide wireless network and performs an rsync that uploads all data files that have not previously been uploaded to our servers. For security, the data is sent through an SSH tunnel. The client side software is open source and publicly available [ISS14b]. After transmission, on the server side, the data is unzipped and inserted into a database. Analytics are then run on the server.

<i>Description</i>	<i>Data</i>
Galaxy S III (Android) <sup>†</sup>	Time, GPS, Speed Acceleration (including gravity) along X,Y,Z axes Linear acceleration (excluding gravity) along X,Y,Z axes Gyroscope for rotation along X,Y,Z axes
Phidget Voltage Sensor	battery voltage for inferring battery SOC
Phidget Current Transducer	charging rate
Digikey temperature sensor	temperature inside sensor box battery temperature

Table 8.1: Sensor hardware and functionality. †: this is not the exhaustive list of sensors in the Galaxy S III, which, for example, also includes a barometer for sensing atmospheric pressure and a sensor for measuring whether there is a magnetic field being applied to the X,Y,Z axes.





Figure 8.3: eBike sensor kit assembly. Left: connected sensors, middle: sensors enclosed in weather proof case, right: battery + sensors + case assembly shown on bike.

## 8.2.4 Participant Selection

To select participants for our field trial, we sent a survey to students, staff, and faculty in several UW departments. The survey<sup>6</sup>, included in Appendix §D, had four sections:

1. Prior knowledge & attitudes about eBikes (questions 2–5), e.g., “please provide up to 20 thoughts on eBikes”. These questions are used to determine the respondents’ perceptions towards eBikes. The responses are analyzed in §8.5.
2. Transport-related motives/emotions (questions 8—27), e.g., “how important is safety when choosing your mode of transport?”. These questions retrieve the necessary inputs to the neural-network discussed in §8.1.1.
3. Usage information (questions 6,7,28—30), e.g., “how many kilometers will you ride daily in summer?”. These responses were used to select respondents for the trial who were likely to produce the most data as described below.
4. Background information (Questions 35—38), e.g., “what is your gender/profession?”. These responses were used to balance for gender and profession as described below.

We received 165 responses to this survey. About half (84) of the respondents were interested in the survey but not in participating in the field trial. From the remaining 81 that were interested in the field trial, 25 participants were selected by our collaborator, Tobias Schroeder, to prevent selection bias, and six eBikes were given to our research staff, totaling 31 eBikes.

<sup>6</sup>This survey is very similar to the one given by Dozza et al. [DWM13].

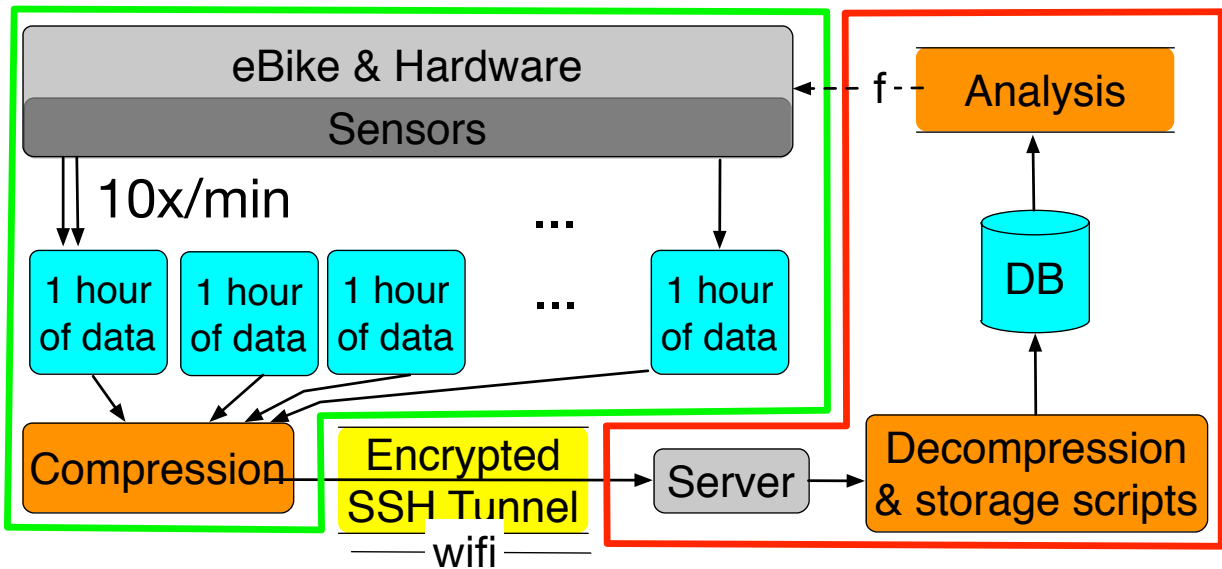


Figure 8.4: eBike data pipeline

### 8.2.5 Incentivizing Participation & Care Of Bikes

Two problems are common to operating most field trials. First, the trial operator must incentivize participants to join the study. Often, there is some work required of users beyond using the technology, and the work required is sometimes a deterrent to potential participants. In many EV field trials, for example, participants were required to keep a travel diary of their vehicle usage. In our trial, participants are required to 1) keep the eBike batteries charged at all times, because if the eBike battery depletes, our data collection and theft detection mechanisms fail, and 2) bring the bike to campus at least once per week so that the eBike can upload its data to our servers. Second, the operator must incentivize the participants to take care of the technology. With WeBike, this means keeping the eBikes and batteries out of the rain when possible and ensuring the eBikes are properly locked at all times to deter theft attempts.

We attempt to solve these two problems at the same time—by giving the eBikes to the participants after three years of data collection. The 2014 retail price of the bikes is \$2,000, so participants are incentivized to participate and keep the bike protected (from rain and theft) because failing to do so would degrade the value of the eBike they are inheriting.

## 8.3 Mapping Battery Voltage to SOC

In this section, we discuss how we estimate the eBikes' battery SOC at all times from voltage measurements. The battery SOC is a necessary input for answering many of the questions posed in §8.1.2. Unfortunately, the batteries equipped on the eBikes do not have SOC



sensors that we can read programmatically<sup>7</sup>. Therefore, we develop a process to estimate the SOC of the eBike batteries.

### 8.3.1 Estimating SOC Based On Voltage

The SOC of a LiON battery can be estimated precisely if the battery current and voltage is available [TMLK11, SPV13]. Unfortunately, we cannot measure the discharge current of the battery, because the wire connecting the battery to the motor is inaccessible. Given this, we estimate the eBikes' battery SOCs given their voltage  $V$  and their temperature  $T$ , parameters we do measure, but we note that this method can be more inaccurate than estimating the SOC based on voltage and current [Bat14].

First, we start with the battery operation model we received from the battery manufacturer. The model, shown in the downward sloping lines in Figure 8.5, shows the discharged capacity (battery capacity - SOC) v.s., voltage for six different battery temperatures. The figure corresponds to a single LiON cell; there are 10 cells in each eBike battery. Next, we manually transcribed the voltage level at every 40th mAh level, leading to 50 measurements of (mAh, voltage) for each<sup>8</sup> curve. We then multiplied each of these measurements by 10 to obtain a model for the 10-cell battery. Finally, we scale each voltage by  $22/32$ ; this is because although the maximum voltage of the 10-cell battery is  $\approx 40V$  (Y-axis of Figure 8.5 multiplied by 10), our battery voltage sensor can only read up to  $28V$ <sup>9</sup>. Our final transcription of the battery model (Figure 8.5) is shown at the upper part of Figure 8.6.

Our goal is to estimate the SOC of a battery given its voltage and temperature ( $V, T$ ). From Figure 8.5, we see that each battery has three "modes of operation":

1. Mode 1, where the battery voltage increases or decreases quickly
2. Mode 2, where the battery is operating stably and the SOC decreases nearly linearly with the battery voltage
3. Mode 3, where the battery voltage decreases rapidly as the battery nears its *depth of discharge*, the point at which the usable portion of the battery has been depleted

Given this observation, our approach is to fit a linear model to each operation mode for each temperature curve, that is, to use a *three line model*. However, notice from Figure 8.5 and the top of Figure 8.6, that for the  $-20C$  and  $-10C$  temperature curves, the SOC is not a *function* of voltage—the same voltage reading corresponds to multiple SOCs. For these curves, the SOC is ambiguous for many ( $V, T$ ) pairs. In light of this, we use a single linear model for these two curves. Specifically, we model each SOC curve as:

---

<sup>7</sup>On each battery, there is a button that displays the SOC to the nearest 20% using a LED display when manually pushed, but this value cannot be polled programmatically.

<sup>8</sup>We did not transcribe the  $60C$  curve as it is unlikely for our eBike batteries to reach this temperature.

<sup>9</sup>Thus, we want a voltage of  $40V$ , corresponding to a full battery, to correspond to the highest voltage our sensor can read ( $28V$ ), and  $40(22/32) = 27.5$

$$SOC(V, -20) = Vm_{-20} + b_{-20} \quad (8.1)$$

$$SOC(V, -10) = Vm_{-10} + b_{-10} \quad (8.2)$$

$$SOC(V, 0) = \begin{cases} Vm_0^1 + b_0^1 & V \geq 24.86 \\ Vm_0^2 + b_0^2 & 24.86 > V \geq 21.58 \\ Vm_0^3 + b_0^3 & V < 21.58 \end{cases} \quad (8.3)$$

$$SOC(V, 23) = \begin{cases} Vm_{23}^1 + b_{23}^1 & V \geq 27.12 \\ Vm_{23}^2 + b_{23}^2 & 27.12 > V \geq 22.13 \\ Vm_{23}^3 + b_{23}^3 & V < 22.13 \end{cases} \quad (8.4)$$

$$SOC(V, 45) = \begin{cases} Vm_{45}^1 + b_{45}^1 & V \geq 27.77 \\ Vm_{45}^2 + b_{45}^2 & 27.77 > V \geq 22.41 \\ Vm_{45}^3 + b_{45}^3 & V < 22.41 \end{cases} \quad (8.5)$$

where the  $m$  and  $b$  linear-fit parameters are chosen to minimize the least squared error on the transcribed data (top of Figure 8.6) for that temperature curve. When training the linear models for the -20C and -10C curves, we ignore the data corresponding to Modes 1 and 3 because 1) these extreme values heavily skew the model slope and 2) most of the time the battery is in Mode 2 where the SOC is a linear function of voltage. The top of The model resulting from Eqs(8.1)—(8.5) is shown at the bottom of Figure 8.6.

### 8.3.2 Voltage Filtering

The voltage sensor is noisy. When using voltage to estimate SOC, the noisy voltage measurements can lead to incorrect SOC fluctuations. To resolve this, we perform a low pass filter on voltage readings prior to estimating the battery SOC. Let  $V_1, \dots, V_n$  represent a time series of successive battery voltage readings. An example series is shown in blue in the top of Figure 8.7. We remove noise from these readings by creating a new series  $\hat{V}_1, \dots, \hat{V}_n$ :

$$\begin{aligned} \hat{V}_i &= \alpha \hat{V}_{i-1} + (1 - \alpha)V_i \\ \hat{V}_1 &= V_1, \quad \alpha \in [0, 1] \end{aligned}$$

The smoothed data is shown in red in the top of Figure 8.7 for a value of  $\alpha = .95$ . This time series is then used for SOC estimation using the aforementioned model. The estimation using our linear model is shown at the bottom of Figure 8.7—in the next section, we discuss why the SOC appears linear in our estimation during periods of biking.




Figure 8.5: Battery model received from battery manufacturer. The “mode” annotations are added by us and are not part of the original figure. The downward-sloping curves show the already discharged mAh (X-axis) vs. the battery voltage for six different battery temperatures. The capacity of each cell is  $\approx 200\text{mAh}$ , and there are eight cells in each eBike battery.

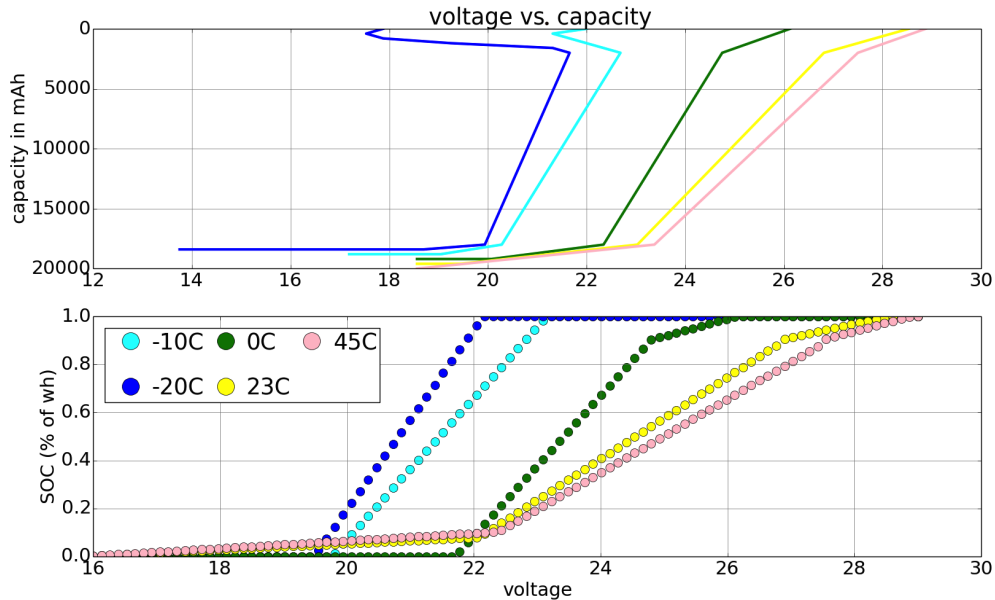


Figure 8.6: Top: our transcription of the original battery model obtained by the battery manufacturer (Figure 8.5, scaled by 10 cells because there are 10 LiON cells in each battery). The curves are inverted to show the SOC on the Y-axis, because the SOC is the unknown variable we are modeling, and voltage (known) is shown on the X-axis. Bottom: our SOC estimation models from Eq(8.1).

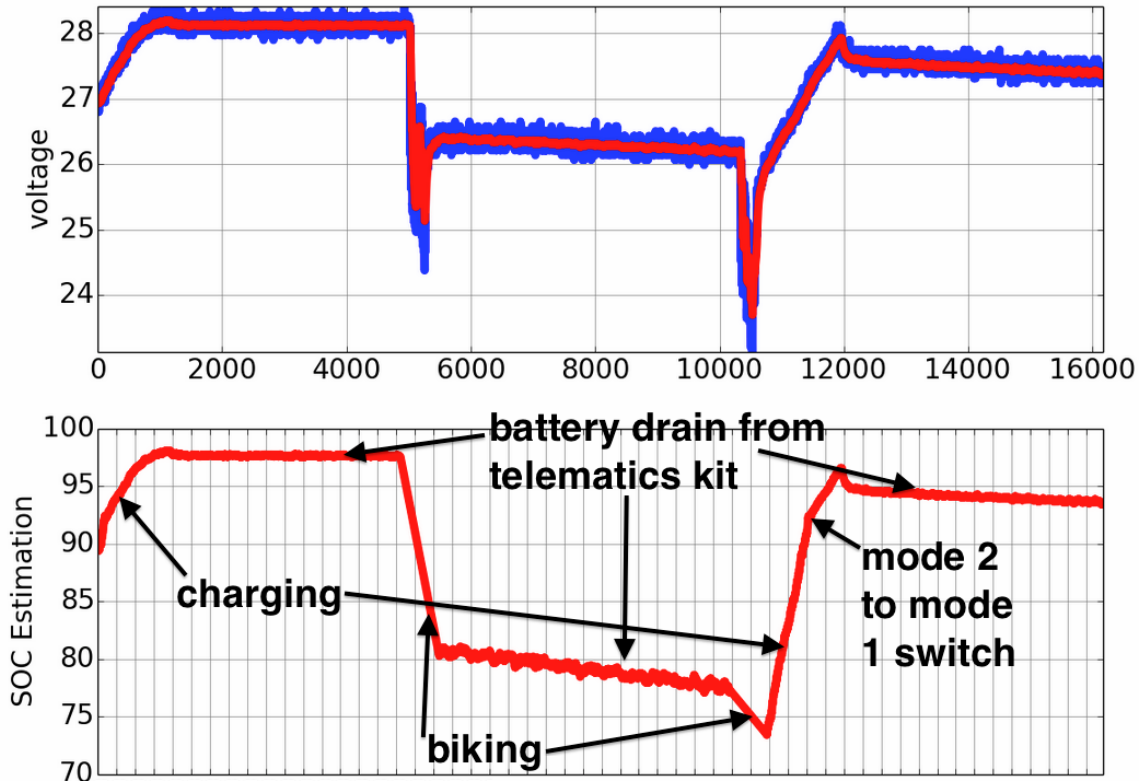


Figure 8.7: Top: voltage readings over a subset of one day for one user. The blue curve shows the raw data, and the red curve shows the readings after a low-pass filter is run. Bottom: our SOC estimation using linear interpolation during periods of biking. The slightly decreasing SOC observed during non-biking periods represents the power consumed by the phone and sensor kit.

### 8.3.3 Biking Interpolation

Finally, we discuss one last complication when estimating the SOC based on voltage. Voltage based SOC estimation cannot be performed while there is a load on the battery [Bat14]. In our case, the voltage-based SOC estimation is inaccurate during periods of biking. When the eBike battery is in use to provide electric assistance of power  $P = IV$ , the battery current and voltage may fluctuate as long as their product remains constant. The fluctuations are based on the drive cycle, e.g., when a user is cycling up a hill, more current is applied, leading to a sharp drop in voltage. This can be observed between 7:19—8:12AM in Figure 8.7, which corresponds to a period the user was cycling. The voltage drops sharply (due to hills) during their commute. To resolve this, we compute the SOC of each battery using the aforementioned linear model while the bike is at rest, and linearly interpolate the SOC during periods of biking. Specifically, if a user is biking from  $t_1$  to  $t_2$ , we compute the SOC at  $t_1$  and  $t_2$  and linearly interpolate the SOC during the interval. This is exemplified at the bottom of Figure 8.7.

## 8.4 Trip Detection

Here we overview our trip detection algorithm. Given a time series  $G_1, \dots, G_n$  of successive GPS coordinates, the algorithm returns a set of trips  $\mathcal{T}$ , where each trip is characterized by its start time, end time, and distance. GPS data is always ordered before processing, but it may be lossy—there may be intermittent periods of arbitrary length missing from the data. Our algorithm is shown in Algorithm 4 and makes the following assumptions:

1. Only records five or more seconds apart are considered; if  $G_k$  is processed by the algorithm, the next record processed is the first non-missing record starting with  $G_{k+5sec}$ .
2. Distance is computed using Haversine distance<sup>10</sup>. This is a standard algorithm for computing the distance between two coordinates on a 3-dimensional ellipsoid surface.
3. If a user bikes from  $t_a \rightarrow t_b$ , and from  $t_c \rightarrow t_d$ , where  $d > c > b > a$ , these are counted as two separate trips if  $t_b$  and  $t_c$  are more than five minutes apart. Alternatively stated, a biking trip containing a period without biking of longer than 5 minutes is treated as two separate trips.
4. If a user bikes from  $t_a \rightarrow t_d$ , and this time interval contains five or more minutes of missing data from time  $b \rightarrow c$ ,  $t_a \rightarrow t_b$  and  $t_c \rightarrow t_d$  are treated as two separate trips. Alternatively stated, a biking trip that is missing five or more contiguous minutes is treated as two trips, one ending right before the missing data, and one starting right after the missing data.
5. If a user bikes less than 15 meters in 5 or more seconds, it is considered to be GPS error and the accumulated distance is not included as part of any trip.
6. If a user bikes at a speed less than 1kmph or greater than 80kmph, averaged over a duration of 5 or more seconds, it is considered to be GPS error and the accumulated distance is not included as part of any trip.
7. Trips less than 1km are discarded as GPS error. Despite the above error checks, we find that a series of GPS errors can still register as a trip of several hundred meters.

In the algorithm,

- Lines 1—6 initialize variables
- Lines 7—8 compute the time and distance delta (per (2) above) between the current data row being processed and the last row that was processed
- Lines 9—11 enforce (4) above.
- Line 12 enforces (1) above
- Line 13 enforces (5) and (6) above
- Lines 14—17 start or update a trip after biking is detected
- Lines 18—25 enforce (3) above
- Lines 29—36 enforce (7) above

---

<sup>10</sup>The exact implementation we use is given here: <http://stackoverflow.com/questions/4913349/haversine-formula-in-python-bearing-and-distance-between-two-gps-points>

---

**Algorithm 4** Trip Detection algorithm

---

```
1: tripStartTimes = [], tripEndTimes = [], tripDists = [] ▷  $\mathcal{T}$ 
2: ResetTripCounters()
3: for l in data: do ▷ l is each data row.
4:   if lastRow == "": then ▷ first row?
5:     lastRow = l, lastRowTime = l[0] ▷ l[0] is time
6:   end if
7:   TIMELAPSE = (l[0]-lastRowTime).seconds()
8:   d = haversine(l[3],lastRow[3], l[4], lastRow[4]) ▷ l[3],l[4] are GPS lat,long
9:   if TIMELAPSE >= 300 and tripHasStarted == 1: then
10:    endTrip() ▷ end prior trip if >10 minute gap in data
11:    lastRow = l, lastRowTime = l[0]
12:   else if TIMELAPSE ≥ 5: then ▷ only consider rows 5 seconds apart
▷ Assume < 15m, <1kmph, or >80kmph are GPS errors
13:     if d > 15 and 1000 < d/(TIMELAPSE/3600) < 80000: then
14:       if tripHasStarted == 0 then ▷ start trip, or add to existing trip.
15:         tripHasStarted = 1, tripStart = l[0]
16:       end if
17:       tripDist += d, tripEnd = l[0] ▷ Update trip distance, trip end
18:     else ▷ did not move
19:       if tripHasStarted == 1: then
20:         secondsSinceLastSignificantMovement += TIMELAPSE
21:         if secondsSinceLastSignificantMovement ≥ 300: then
22:           endTrip(): ▷ end trip if did not move in last 5 mins
23:         end if
24:       end if
25:     end if
26:     lastRow = l, lastRowTime = l[0] ▷ Don't do this if TIMELAPSE < 5
27:   end if
28: end for

29: procedure endTrip():
30:   if 1 < tripDist < 1000: then ▷ Filter out trips <1km
31:     tripStartTimes.append(tripStart)
32:     tripEndTimes.append(tripEnd)
33:     tripEndTimes.append(tripDist)
34:   end if
35:   ResetTripCounters()
36: end procedure

37: procedure ResetTripCounters():
38:   tripHasStarted = 0, tripStartTime = 0, tripEndTime = 0, tripDist = 0
39:   secondsSinceLastSignificantMovement = 0
40: end procedure
```

---

## 8.5 Survey Analysis

In this section, we analyze the responses received from the participant selection survey. This is in effort to gauge perceptions towards eBikes prior to our field trial.

Out of the 38 questions in the survey (Appendix §D), 6 included text responses. Three questions (5, 31, and 34) were open ended. Three other questions (3,4, and 6) included a comment field where respondents further explained their answers. Collectively, we refer to all text written by the respondents for these six questions as the *survey corpus*. The majority of the corpus is composed of the responses to Q5, which asked the respondents to list up to random 20 thoughts on eBikes.

In this section, we visualize the comments in the survey using two methods: classifying the sentiments of the most common phrases used by the survey respondents, and mining opinions about eBike product features using the system described in Chapter 4.

### 8.5.1 Common Phrases In Survey

Figure 8.8, 8.9, 8.10, and 8.11 show the sentiments of commonly used phrases in the survey corpus. Figure 8.8 and 8.10 show frequencies of *unigrams* (single word terms), and Figures 8.9 and 8.11 show the frequencies of *bigrams* (double word phrases). To produce these figures:

- The sentiments were classified using the the MPQA Opinion Corpus [Wil, Wil08], the same sentiment dictionary used in Chapter 4.
- A list of NLTK *stopwords*<sup>11</sup> (mostly inclusive of conjunctions) were removed prior to analysis for all four figures. Consequently bigrams such as “better for the environment” appear as “better environment”.
- To produce Figures 8.8 and 8.10 which show frequencies of unigrams, additional nouns that were frequent but do not give insight on their own, such as “road” and “bike”, were added to the list of stopwords. These words are not excluded from the bigram analysis so that phrases such as “faster than bikes” are mined.
- Different forms and synonyms of the same word are merged, e.g., the frequencies of “dangerous”, “danger”, and “unsafe” are summed and shown for “dangerous”.

---

<sup>11</sup><http://www.nltk.org/book/ch02.html>




Figure 8.8: Frequencies of all single-word terms used at least 10 times. Red terms were classified as negative, green as positive, and blue as neutral.




Figure 8.9: Frequencies of all bigrams (double-word terms) used at least 4 times. Red terms were classified as negative, green as positive, and blue as neutral.



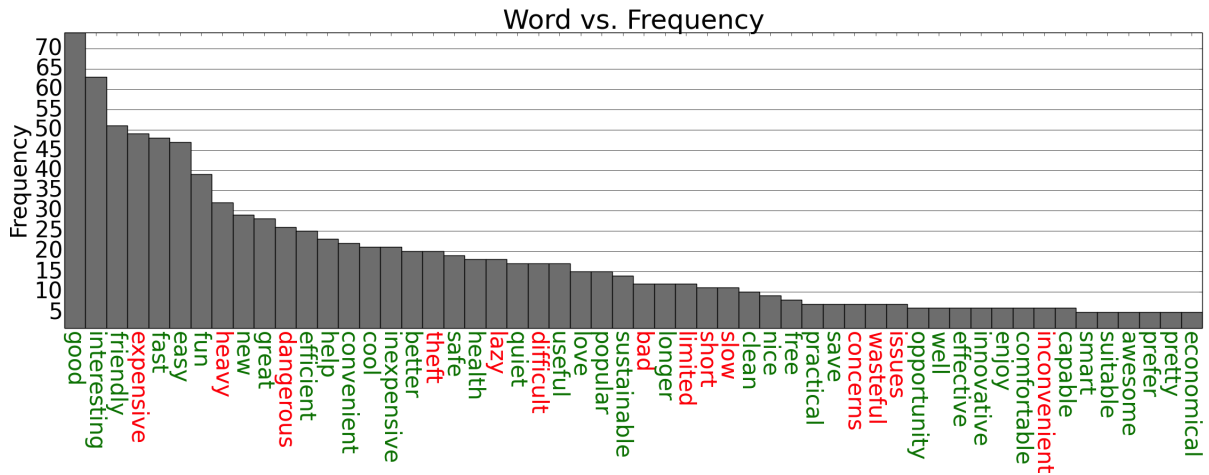


Figure 8.10: Frequencies of all single-word terms not classified as neutral used at least 5 times. Red terms were classified as negative, green as positive.




Figure 8.11: Frequencies of all bigrams not classified as neutral used at least 3 times. Red terms were classified as negative, green as positive.

### 8.5.2 Sentiments In Survey

We also set up the sentiment system presented in Chapter 4 to mine the survey corpus for opinions about eBike features. The sentiments are shown in Figure 8.12. Note that only opinions containing both a feature and an opinion are mined; the number of opinions below will be less than the number of positive and negative phrases used in the corpus, which were shown in the prior section. Here we state our conclusions:

- (N) Most respondents were unfamiliar with or had no opinions regarding the battery, charging properties, design, or maintenance of the eBikes.
- (+) Interestingly, respondents compared eBikes to public transit when talking about "convenience", but not to cars. This could be a bias of the sample (University students




Figure 8.12: Sentiments found in all sentences in the eBike survey corpus. "Usage Conditions" refers to sentiments found in the context of weather or seasons, e.g., "cannot be ridden in Winter".

and faculty), as many do not own cars. Respondents find that eBikes are convenient when compared to public transit because they can be ridden without waiting. This is also shown in Figure 8.11 with 22 occurrences of "convenient" but only five of "inconvenient"<sup>12</sup>.

- (+) eBikes are seen as environmentally friendly and high performance (fast relative to cycling). Moreover, the majority of general<sup>13</sup> sentiments towards eBikes are positive.
- (–) The majority of sentiments regarding range and price are negative. This exactly matches the sentiments towards EVs found in many surveys and field trials (see §3.1).
- (–) Some respondents feel eBikes are for those who are "lazy", "cheaters", or "unhealthy". Many of these respondents were cyclists.
- (–) As a form of range anxiety, even though eBikes can be pedaled after battery depletion, respondents have the perception that the extra weight (v.s. standard bicycles) make it infeasible to do so.

<sup>12</sup>These five were not used in conjunction with "eBikes". The phrases "(in)convenient" were set up as feature changers for "General" (see §4.6.1), so opinions were only mined for the "convenience" feature if the chunk also contained "eBike".

<sup>13</sup>The large number of comments for the "General" feature are due to the many questions "eBikes are...".

## 8.6 WeBike Data

The questions posed in §8.1 will be fully answered only after the completion of this thesis. However, we present some preliminary results that are available. At the time of writing, the WeBike participants have had functioning eBikes for two to three months. When participants first received their eBikes in August 2014, we encountered numerous hardware and software problems that led to a lack of data early in the trial. Most of the WeBike participants received eBikes with functioning hardware and software in October and November. Thus, the results presented here are limited to a few months, but we note that much of our work was in developing the algorithms used to present these results (which can be updated as biking continues).

Our main focus is on the questions posed in §8.1.2, as they are the most relevant questions to the rest of this thesis. We study six of the questions presented in §8.1. All Figures presented that show trip lengths or distance traveled use Algorithm 4 to detect trips.

**1) Do some use their eBikes primarily for commuting or shorter trips while others use the eBikes for longer trips (possible for health/exercise), or do users have similar usage patterns?**

Figure 8.13 shows the trip length distribution of all eBike trips made in the same date range for the same set of eBikes. Figure 8.14 shows the distance traveled by 24 participants from 9/7/2014—12/20/2014. We see from these figures that, thus far, the users have mostly used the eBikes for shorter trips, e.g., commuting. The average trip length is 3—4km, the commute distance to campus for most participants. There are about 20 continuous trips (recall that a trip is considered two trips if it contains a period of no biking that is longer than five minutes) of more than 11km in this time interval. As expected, we also see that some users bike more than others. Out of the 24 participants shown, five surpassed 50km, four surpassed 100km, and one user has biked more than 450 km.




Figure 8.13: Trip length distribution for 24 participants from 9/7/2014—12/20/2014.




Figure 8.14: Cumulative distance traveled by 24 participants from 9/7/2014—12/20/2014.

## 2) At what SOC do people normally start charging?

Figure 8.15 shows the number of charging events started at each of 10 SOC ranges. We observe that users charge very conservatively. More than half of all charging events began when the battery had a SOC of 50% or more. About 25% of all charging events started when the SOC was already >90%, i.e., users “top off” their batteries frequently. Only 10% of charging events occurred when the eBike had less than 30% of battery left. The mean SOC range at which users start charging is 60—70%.

## 3) Do participants deplete or nearly deplete their eBike batteries? I.e., are they pushing their range limits?

Figure 8.17 shows the “empirical range” of the eBikes; this figure is explained in greater detail in its caption. From this figure, we see that nearly all user trips deplete less than 25% of the battery. In addition, from Figure 8.13, we see that most trips are well under the advertised range of 40km. Thus, we conclude that users are not yet pushing their range limits. We note that if a user fully depletes their battery, and does not charge it for long enough that the smartphone battery also depletes, the charging event will be observed when the eBike battery is around 10% and the phone reboots. Thus, this happened at most eight times (Figure 8.15).

**4) Do we observe the “range paradox” reported in many EV field trials and surveys? I.e., is there a mismatch between the users’ range needs and their perceptions or behavior?**

As discussed, nearly all user trips deplete less than 25% of the battery (Figure 8.17) and most trips are under 4km (Figure 8.13). This contradicts the charging characteristics observed in Figure 8.15, i.e., people are charging their batteries frequently despite the fact that they are not yet close to range limitations. In addition, this contradicts the sentiments towards range and weight observed in Figure 8.12, because the users do not need more range for their current mobility patterns and are not at risk of having to pedal the heavier eBike (vs. conventional bicycles) after depletion, yet the majority of sentiments for both of these features are negative. Thus, we conclude that there is a disparity between range needs and behavior/perceptions.

**5) LiON battery properties: What are the effects of different drive cycles and different levels of electric assistance on battery life/range?**

From Figure 8.17, we see that the eBike range is variable, and that the range estimates given by the eBike manufacturer (“up to” 40km, see §8.2.1) are fairly accurate. The average range achieved by WeBike users, thus far, is  $\approx 30$ km. Range varies based on three major factors: the amount of electric assistance used, the number of hills on the route, and the temperature. We observe a few instances with higher than estimated range (40—70km), which corresponds to the user applying minimal electric assistance, and a few instances where the battery was depleted almost immediately (0—10km), corresponding to full electric assistance, possibly on trips with steep hills. With respect to temperature, we note that the results shown here correspond to fall (September through December) temperatures, so results may change later when winter and summer biking is also observed.

**6) Parking habits/charging: What is the temporal distribution of charging events?**

Figure 8.16 shows the temporal distribution of all charging events for the 24 participants from 9/7/2014—12/20/2014. Thus far, there does not appear to be a significant bias towards charging at any particular time. Some users are charging their eBikes during working/schooling hours, as indicated by the mid-day charging—recall that the battery is detachable and can be charged indoors. There is also a minor peak in charging at 5pm, representing after-work charging. Still, the temporal distribution of charging events is quite evenly distributed. These preliminary charging patterns do not correlate with EV charging patterns, where many owners charge their vehicles overnight.




Figure 8.15: The number of observed charging events v.s., the SOC of the eBike battery at the start of charging. The black curve shows how many charging events began when the SOC was in that range, and the grey curve shows the cumulation.




Figure 8.16: Distribution of charging events for 24 participants from 9/7/2014—12/20/2014.

**User Portal** To conclude this section, we note that we have created a website for users to log in<sup>14</sup> and view their biking data. At the time of writing, users can generate six analyses:

1. View their trips (start time, end time, distance) on a given day
2. View their cumulative distance traveled per day (Figure 8.13 for just their eBike) over any given time period
3. Plot their battery life over time on a given day (identical to Figure 8.7).
4. Plot their trip length distribution over a given time period (Figure 8.13 for their eBike)
5. Plot their speed over time on a given day
6. Plot their data on Google Maps on a given day

A screenshot of the user interface, at the time of writing, is shown in Figure 8.18. The code for all analyses in this section, as well the website, is open source [Car14].

<sup>14</sup><http://blizzard.cs.uwaterloo.ca/webike>




Figure 8.17: This figure shows the percentage of the eBike battery that was depleted during all trips taken by the 24 participants from 9/7/2014—12/20/2014. Each circle represents a trip. The X-axis shows the length of the trip. The annotations show the percentage of the battery was depleted on the trip. The Y-axis shows the range of the eBike assuming all of the user’s trips required the same amount of battery per kilometer as the trip shown. For example, if a 5km trip requires 25% of the battery, the Y-axis will show an empirical range of 20km.

## 8.7 Prior Bicycle Telematic Systems

To the best of our knowledge, there are no prior works on building sensor kits specifically to measure eBikes. The sensing kit for the German pedelec study [GKS<sup>+</sup>12] (see related work in §3.1.1) did not include any eBike specific sensors, and the same is true for Paefgen et al. [PM10]. The kits used for e-bikeSAFE [DWM13, DF14b], as summarized below, was nearly the same kit used for their bikeSAFE project; it included only one additional eBike specific sensor to measure the level of electric assistance used at all times. The two prior bicycle measuring frameworks that served as a starting point for our sensor kit are described below.

Dozza and Fernandez design a safety-focused sensor kit for their bikeSAFE project [DF14b], which was also used in their e-bikeSAFE pilot [DWM13]. The only eBike-specific sensor added for the e-bikeSAFE project is a current sensor which measures the level of electric assistance used at all times. Their bikeSAFE kit collects the data shown in Figure 8.19 (this table does not include the electric assistance sensor). The logger, the Phidgets SBC2 microcomputer, acquires, processes, and saves all sensor data and is housed in a waterproof case (Pelicase 1150). The logger is turned on/off automatically when the user starts/stops pedaling. The sensors and logging unit are powered by a separate onboard battery (not the eBike battery). A push button is also mounted on the handlebars that allows the user to trigger an alert so the time of the event can later be correlated with data obtained at that time. On the server side, MATLAB is used to linearly interpolate missing data and assess the data quality of each log file. Files with much missing data and files containing a low number of trustable samples are discarded.

Eisenman et al. [EML<sup>+</sup>09] present BikeNet, an advanced sensor framework for bicycles—there are no eBike specific sensors. The sensor hardware includes:

- the Moteiv Tmote Invent platform—contains a two axis accelerometer (for acceleration, angle of incline, lateral tilt), a thermistor, a photodiode, a microphone, and capabilities to interface with all other sensors (similar to the Phidget platform)
- Invent push button—records the timestamp whenever the user presses<sup>15</sup>
- the Nokia N80 phone—contains a camera, microphone, and GSM (for data upload)
- a magnet-triggered reed relay—records the angular velocity of the wheel and pedal, forward speed, distance traveled
- Honeywell HMC1052L dual axis magneto-inductive sensor—measures direction and deviation with respect to the Earth's magnetic field, metal detection (for car detection)
- Garmin Etrex GPS unit—records location and time
- Telaire 7001 CO2/Temperature Monitor—measures carbon dioxide levels
- ArcherKit Biofeedback Monitor—measures stress level of cyclist

---

<sup>15</sup>used similarly to how the bikeSAFE push button allowed the user to record dangerous situations to see whether the event was also detected in the sensor data



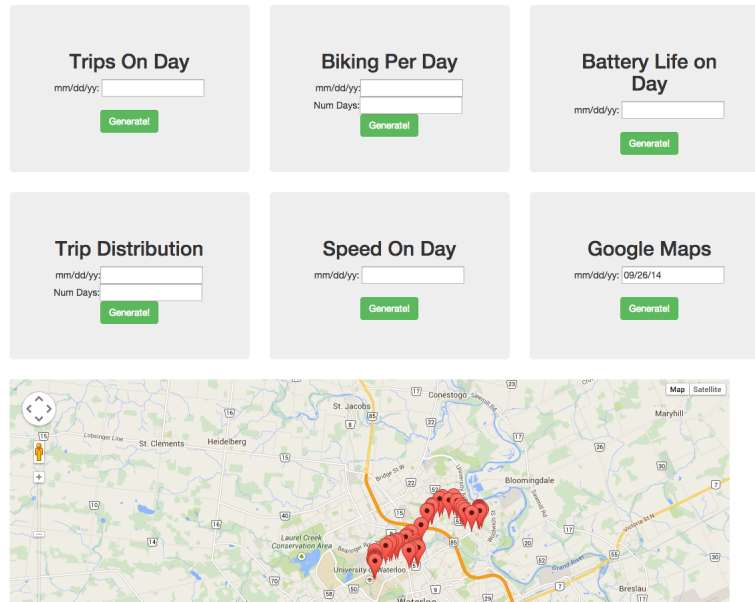


Figure 8.18: WeBike user portal interface

- OtterBox 1600 Case—protects the above hardware from inclement weather

To validate events detected in the sensor readings, cameras were also mounted on the cyclists helmet so that sensor readings could be correlated with video footage.

## 8.8 Conclusions

In this chapter, we have presented our design of WeBike, a field trial for studying electric bike (eBike) usage in North America. WeBike is designed to study why eBikes are very well adopted in some regions (e.g., 200 million in China) and seldom in others (e.g., less than 1 million in the U.S.). Moreover, we hope to use eBikes as a “wind tunnel”—a way to study parameters of interest to EV researchers and manufacturers at low cost. Our participant selection survey reveals that many view eBikes similarly to EVs, and that both share a set of common adoption barriers, including range anxiety. Our preliminary analysis of the WeBike data supports this comparison.

In the beginning of this chapter, we introduced many research questions that we hope to answer throughout WeBike. Future work will entail answering these questions. However, we conclude with a final suggestion for future work. An interesting future avenue would be to adapt our carshare sizing methodologies presented in Chapters 5 and 6 to size *eBike sharing systems*. This may not be a straight forward adaptation as Cherry et al. [CWJ11, LCY<sup>+</sup>13] discuss. Their discussion of operating<sup>16</sup> of an eBike sharing system show there are many differences between operating a carshare and an eBike sharing system. First, while eBike users can pedal when the battery is depleted, fatigue may prevent users from returning their

<sup>16</sup>not sizing; the fleet size was fixed at 20 eBikes based on budget constraints and was not based on demand

Sensor Type	Sensor Name	Data Provided	Resolution	Sample frequency
Camera	GoPro Hero, Hero2	Video frames	1920x1080 pixel	30 fps
Accelerometer	Phidgets IMU 1056	3D Acceleration	22 $\mu\text{g}$	100 Hz
Compass		3D Directional vector	400 $\mu\text{G}$	100 Hz
Gyroscope		3D Angular rate	0.02°/s	100 Hz
GPS	Phidgets GPS 1040	Latitude and longitude	Minimum circular error 2.5 m)	10 Hz
		Heading	(Varies with circular error)	10 Hz
		Velocity	100 mm/s	10 Hz
Brake force sensor	Flexiforce resistive force sensor	Pressure	0.01 N	100 Hz

Figure 8.19: reproduced from Dozza and Fernandez [DF14b]. "IMU": inertial measurement unit.

eBikes or even reaching their destination. Moreover, when an eBike is returned, it has to be charged before it can be used by the next subscriber, and this charging time depends on both the previous subscriber (e.g., weight) and their mobility patterns (e.g., distance driven, percentage of electric assist). Therefore, range and charging time constraints must be taken into account during both the sizing and rebalancing phases. Additionally, share locations can be equipped with battery swapping stations because eBike batteries are lightweight and removable. To reduce the recharging time, additional batteries could be kept at each location to instantly bring a returned eBike back into service. The rebalancing problem then entails the balance of both eBikes and spare batteries.

# Chapter 9

## Conclusions

This chapter concludes the thesis. In §9.1, we summarize the contributions presented in this thesis. We summarize some avenues for future work in §9.2. Finally, we present concluding remarks in §9.3.

### 9.1 Summary Of Contributions

As discussed throughout this thesis, and validated by the literature survey in Chapter 3, range anxiety and the higher initial price of EVs compared to ICEVs are the two primary barriers preventing widespread EV adoption. We have presented computational methods towards alleviating these concerns. Here, we summarize our contributions.

In Chapter 4, we present an *online sentiment analysis* system that mines EV owners' perceptions from online forums. Such perceptions have historically been obtained through expensive field trials or targeted surveys. The system mines opinions about specific vehicle features such as "range" and "safety". It outputs a high-level product overview—a breakdown of sentiments for each feature—with the ability to drill down into desired opinionative sentences. These feature-wise opinions can help multiple types of users: they help prospective buyers determine whether a given EV matches their needs, they help marketers better advertise EVs (by promoting the best-reviewed features), and they help manufacturers improve their vehicles so that later-generation models are more aligned with drivers' preferences. Note that some perceptions, for example those relating to battery degradation after years of charging, are not captured during field trials, but our system is able to do so.

To build our system, we extended previous review mining systems with several new optimizations and EV domain knowledge. In addition, the system is general and customizable so new vehicles and product types can be quickly added; we tuned the same system to classify sentiments towards eBikes. We have open-sourced our system [Car] to allow for future extensions.

In Chapter 5, we propose that giving BEV owners access to an ICEV, which can be used on days their BEV does not have sufficient range, can help alleviate range anxiety. For a

BEV owner to feel as though all of their mobility demands are met through a combination of their BEV and the pooling service, the ICEV access should have a very low transactional cost. This "ICEV pool" can be implemented in one of many ways: it can be composed of a BEV dealership's unsold ICEVs [Ing13]<sup>1</sup>, run as a for-profit-service similar to a carshare, or as a community or government run non-profit to catalyze EV adoption. In this chapter, we propose three algorithms to statically size such a pool of ICEVs given the demand patterns of its collective BEV "subscribers". Our goal is to find the minimum number of vehicles that can be stored in the pool such that a desired percentage of subscriber requests, the QoS target, are met. Sizing vehicle pools to ensure a desired QoS target is difficult because demand is highly non-stationary—there are periods with very high demand both within a single day and over the course of a year—but a high QoS guarantee is vital to effectively reduce range anxiety. Using eight years of data collected from an Ontario carshare, we show that our algorithms, when correctly parameterized, can achieve within 1–3% of the QoS target.

In addition to having a low transactional cost and a high QoS, another aspect that should be considered is the convenience of pickup—minimizing the transit time and distance for subscribers to retrieve their ICEVs. In Chapter 6, we extend the methods in Chapter 5, which statically size a single ICEV pool, to *periodically* size a *network* of pool locations. A multi pool network, if sited appropriately, is more convenient for subscribers because the average time required to retrieve a vehicle is reduced with more access locations. In this chapter, we propose two algorithms to compute the addition, removal, and movement of vehicles within such a network, which we refer to as *fleet management*. The methods adapt the size of each location in response to changing membership or demand at fixed time intervals known as *periods*, the length of which are given as input (we use two weeks). Our methodology ensures that vehicles are "rebalanced" instead of purchased when possible—if demand increases at location  $L_1$  but decreases at  $L_2$ , we attempt to relocate a vehicle from  $L_2$  to  $L_1$ . Prior methods for fleet management make simplifying but incorrect assumptions about demand patterns, e.g., the arrival process is often assumed to be Markovian and stationary. In contrast, our methods work with *arbitrary* demand patterns. We show using the same eight-year dataset that our methods perform, on average, within 1–3% of the QoS target across all network locations, despite non-Markovian and non-stationary demand.

Fleets are a likely adoption market for EVs. Fleet operators have the initial capital to fund long-term investments, can dispatch multiple types of vehicles (EV/ICEV) to serve different demands, and often have centralized parking where charging infrastructure can be sited. In light of this, in Chapter 7, we present a data-based ROI model for taxi fleets, which are large consumers of petroleum in cities worldwide, to compute whether it is profitable to invest in EVs. We do not make assumptions about the taxis' mobility. Instead, we use the GPS coordinates of the company's existing taxis to model their mobility patterns. The GPS time series are input into a Bayesian network that estimates each taxi's battery SOC over

---

<sup>1</sup>This approach has several benefits; it reduces transactional costs because the dealership has the subscriber's information from their EV purchase, and moreover dealerships can compute the cost of operating this service and amortize the cost into the EV purchase price so the EV and ICEV access are sold as a single, convenient package

each shift. Several ROI metrics are computed based on when taxis deplete their batteries (according to the SOC estimation) and must either begin using gasoline (for a PHEV) or return to the headquarters (for a BEV). We calculate the revenue losses or gains, as well as the payback period (if applicable), considering many different infrastructure and pricing scenarios. We use a dataset containing over 20 million GPS readings from 530 taxis in San Francisco for our evaluation. We find that, as of 2014, HEVs, PHEVs, and BEVs are all profitable for the company, to different extents, and each have merits. To our knowledge, this work is the first of its kind; work prior to ours focused on ROI analyses for individual drivers (see §3.3) who have different mobility patterns and requirements than taxis.

Finally, in Chapter 8, we discuss a different type of electric vehicle—electric bikes (eBikes). EBikes have been heavily adopted in some regions, e.g., China which has over 200 million in use [Tim13], but seldom adopted in others, e.g., the US where there are less than 1 million are in use [Bak13]. The large disparity in sales could be due to differing attitudes towards bikes/eBikes, differences in transportation networks, differences in weather patterns, or something else altogether—this issue has not been well studied. Towards this end, we designed *WeBike*, a Waterloo based eBike trial. After surveying 160 respondents about their perceptions towards eBikes in this North American context (where eBikes are not well adopted), we chose 31 respondents to participate in a three year trial. Much of our work thus far went into designing the telematics strategy for *WeBike*. Each eBike is equipped with over 15 different sensors, most importantly GPS and battery sensors. Our telematics solution only requires that participants bring their eBikes to campus periodically and to keep their eBike batteries from fully depleting (requiring a charge every four to five days)—everything else, including the transmission of data, is fully automated, in contrast to prior EV field trials (see §3.1.1) requiring the user to collect or upload data. Additionally, we have developed algorithms to process the raw data, such as GPS and battery voltage readings, into higher level concepts like biking trips and battery SOC. We show that eBikes and EVs have almost identical perceptions by using the sentiment analysis system from Chapter 4 to analyze the survey responses. We further show how eBikes can be used to study several parameters pertinent to EV researchers, including range anxiety, parking and charging habits, and LiON battery properties.

## 9.2 Directions for Future Work

Here we summarize some interesting extensions to our work.

The sentiment analysis system presented in Chapter 4 could be extended in two ways. First, our system currently only examines the sentiments towards a vehicle at a given snapshot in time. Future work can include expanding the system to study how sentiments change over time. For example, how do opinions about specific features change when the vehicle is updated in each model year? Do opinions about range anxiety change as owners possess their vehicles for longer? We are currently not capturing these temporal trends. Second, the enhancements discussed in §4.10 would help increase the performance of the system. These improvements include pronoun resolution, which will help capture many more opinions<sup>2</sup>, detecting the product being discussed from the context rather than forum/page titles, detecting comparisons of two products, and better context-dependent opinion handling.

Because Chapter 6 extends the mathematical foundations presented in Chapter 5, we discuss ways to extend these two chapters simultaneously. First, we have only analyzed one QoS metric—availability. In both chapters<sup>3</sup>, we consider the system requirements met if more than  $1 - \epsilon$  percentage of requests are served. However, other QoS targets or objective functions could be explored. For example, subscribers may be willing to receive a vehicle less often if the service is significantly cheaper. Depending on the demand patterns, the pool size may have to increase greatly to meet slightly higher QoS targets. For example, Table 6.4 shows that the pool size needs to increase by 50% to meet a meet QoS of 99% vs. 90%. Consequently, always maintaining a high QoS target may lead to larger pool sizes and high subscription fees. An alternative may be to consider both the service availability and the service operating costs in the objective function. Second, we have not studied the problem of *locating* the pool(s). If the pooling service is offered by an EV dealership, naturally the pool is formed by their unsold ICEVs. However, if one or more pools are to be operated by another entity separate from the dealership, the location of the pool(s) is important. This becomes a *facility location* problem with many considerations—the cost of real estate, the distance of the pools to subscribers, the distance between pools which affects fleet management costs, whether potential pool locations are easily resizable (e.g., can new spaces be added if the pool needs to grow), etc. These considerations can lead to computationally difficult problem formulations, as many facility location problems are NP-Hard or NP-Complete. Finally, several ways to improve our fleet management algorithms are given in §6.9. While we show how to extend our mathematical model to allow for one way trips, we have not evaluated this model extension and it would be interesting to see how often vehicle rebalancing must occur for the QoS targets to be maintained if one way trips are allowed. Similarly, we show how our model can be extended to allow members to obtain their vehicles from different locations instead of always going to the same pool, but it would be interesting to see how the pool sizes change under this scenario.

---

<sup>2</sup>For example, sometimes an entire paragraph is written about a feature, but only the first sentence of the paragraph explicitly mentions the feature. Currently, the other related opinions are lost.

<sup>3</sup>In Chapter 6, we have the *secondary* objective of minimizing fleet management costs.

We note that for sufficient evaluation of these scenarios, a different dataset which shows the demand at several pool locations is desirable. Lastly, our fleet management currently employs repeated myopic optimization; each time period, we optimize only for the next coming period. There are several challenges that we have yet to overcome in implementing multi-period optimization as discussed in §6.9.

While we attempted to make the ROI model presented in Chapter 7 comprehensive, there are several avenues for extending it. First, obtaining a dataset of EV taxi usage (to be used in combination with our dataset of ICEV taxi usage) would greatly improve our evaluation. By predicting the SOC using our model, then comparing the results to data obtained from actual electric taxis, we can calculate the *variance* at each node in the network, which we currently ignore. This would allow us to have confidence intervals in our results. Second, our assumption that all fares are lost for the remainder of the shift upon BEV battery depletion may be too conservative. If the taxi company has additional ICEVs at the headquarters, it may only take on the order of minutes to bring the BEV back (just prior to depletion) and begin using an ICEV. While this process may limit the driver's ability to complete fares during this switching period, if depletion occurs early enough in the shift, they may still be able to complete fares afterwards. We note, however, that modeling this switching behavior 1) requires altering the taxi trajectories seen in the dataset and 2) removes the certainty of conservative results. Similarly, our approach can currently only be used by taxi companies whose vehicles are brought back to a common location after each driver's shift. Future work could generalize the process to taxi companies with multiple vehicle repositories. Finally, additional extensions, including modeling of vehicle maintenance costs, real estate prices, and queueing delays at switching stations, are discussed in §7.6.

The WeBike trial discussed in Chapter 8 has many promising avenues of future work. First, we are planning to release additional participant surveys every 6—12 months, including one early next year. These surveys will help study how perceptions towards eBikes change with continued eBike use. Next, many of the questions posed in §8.1 will be addressed as more data becomes available. There are two major research avenues we wish to explore in this context. Primarily, we hope to coordinate the WeBike usage data with the perceptions obtained through the initial and subsequent participant surveys; many specific questions are given in §8.1. Examining how the actual eBike usage correlates to those individuals' perceptions will allow us to study the differences between perception and reality, which should help us understand why eBikes are not well adopted in some regions. Finally, one can extend our fleet management algorithms to eBike sharing systems. With respect to rebalancing, this is an easier problem, because many eBikes can be moved at once between locations using a truck; the operator does not have to pay a driver to move each individual vehicle between locations. However, there are challenges that do not exist in the ICEV pool context, for example, the eBikes and their batteries need to be charged prior to use, so sizing must take into account the time needed to charge the batteries, which in turn depends on a variety of factors including the users' mobility patterns.

## 9.3 Concluding Remarks

We believe that transitioning our transportation network from petroleum based fuels to cleaner, more sustainable fuels is necessary. As sweet crude becomes more difficult to find, we have turned to unconventional sources of oil, including tar sands and offshore drilling. Recovering petroleum from these sources is not only inefficient, but has led to several large-scale environmental disasters. Simultaneously, petroleum-based transportation is a large contributor of CO<sub>2</sub> emissions, and many believe that if CO<sub>2</sub> emissions are not reduced greatly and rapidly, climate change will soon pass an irreversible “tipping point”.

Over the next several decades, our transportation network may be composed of a number of alternative fuels, including electricity, hydrogen, and natural gas. Currently, the best alternative seems to be electricity. In November 2014, global EV sales surpassed 600,000 vehicles, which still pales in comparison to ICEV sales, but far exceeds all of the other alternatives. Thus, we have worked towards facilitating the adoption of this technology.

The transition to EVs presents a large change for drivers. Re-fueling is no longer instant; batteries take on the order of hours to charge, though some “fast charging” technologies seem promising. Consequently, range is no longer infinite; drivers must plan routes and charging stops to reach some destinations. Moreover, while gasoline is a far more expensive fuel than electricity, it is paid for over time, whereas EV batteries must be paid for upfront when purchasing an EV. Many drivers find it difficult to compute when the fuel savings will outweigh this higher initial cost.

We have worked towards solving these problems. We have built a system that helps drivers determine whether a particular EV is right for them and for manufacturers to build vehicles better aligned with drivers’ preferences, proposed algorithms to size and manage a network of ICEVs to be used by range-limited EV owners, built an ROI analysis for taxi companies—large consumers of petroleum in cities across the world—to compute their ROI in transitioning to EVs, and designed a comprehensive field trial of electric bikes that can be used to study EVs at low cost. Hopefully these contributions, when combined with others’ research, will help us move towards a cleaner, more sustainable future.



# References

- [AAB13] Baha M Al-Alawi and Thomas H Bradley. Total cost of ownership, payback, and consumer preference modeling of plug-in hybrid electric vehicles. *Applied Energy*, 103(C), March 2013.
- [ACX<sup>+</sup>13] Kang An, Xiaohong Chen, Feifei Xin, Bin Lin, and Longyu Wei. Travel Characteristics of E-bike Users: Survey and Analysis in Shanghai. *Procedia - Social and Behavioral Sciences*, 96, November 2013.
- [AdP04] Simon P Anderson and André de Palma. The economics of pricing parking. *Journal of Urban Economics*, 55(1), January 2004.
- [Adv11] Advanced Energy. Community Planning Guide For Plug-in Electric Vehicles. *Advanced Energy Report*, 2011.
- [Aec09] Aecom. Economic Viability of Electric Vehicles. *Aecom Report*, September 2009.
- [AGW00] Jürgen Albert, Dora Giammarresi, and Derick Wood. Normal form algorithms for extended context-free grammars. *Theoretical Computer Science*, 267, 2000.
- [Arg14] Argonne National Laboratory. Light duty electric drive vehicles monthly sales updates, 2014. Acc. Sept 30, 2014. [http://www.transportation.anl.gov/technology\\_analysis/edrive\\_vehicle\\_monthly\\_sales.html](http://www.transportation.anl.gov/technology_analysis/edrive_vehicle_monthly_sales.html).
- [ARU12] ARUP. Coventry and Birmingham Low Emission Demonstrators (CABLED) Project, 2012. <http://www.arup.com/Projects/CABLED.aspx>.
- [Aud14] Audi. 2014 Audi A7 vs Tesla Model S Comparison, 2014. Acc. Apr 4th 2014. <http://tinyurl.com/l2e628h>.
- [Bak11] Jorrit J Bakker. Contesting Range Anxiety: The Role of EV Charging Infrastructure in the Transportation Transition. Master's thesis, Eindhoven University Of Technology, May 2011.
- [Bak13] Brandon Baker. Electric Bike Sales Soar Worldwide, Nov 2013. Acc. July 20th 2014. <http://ecowatch.com/2013/11/06/electric-bike-sales-soar-worldwide/>.

- [Bat14] Battery University. How to measure state-of-charge, December 2014. Accessed Jan 10, 2015. [http://batteryuniversity.com/learn/article/how\\_to\\_measure\\_state\\_of\\_charge](http://batteryuniversity.com/learn/article/how_to_measure_state_of_charge).
- [Bay12] Bay City News. SFMTA Approves Controversial Taxi Medallion Plan, August 2012. Acc. October 20th 2012 <http://sanfrancisco.cbslocal.com/2011/09/02/111118/>.
- [BCMWW11] Albert G Boulanger, Andrew C Chu, Suzanne Maxx, and David L Waltz. Vehicle Electrification: Status and Issues. *Proceedings of the IEEE*, 99(6), June 2011.
- [Bet12] Better Place. Battery Switch Stations, 2012. Acc. May 15th 2012. <http://www.betterplace.com/the-solution-switch-stations>.
- [BHB14] Louise Bunce, Margaret Harris, and Mark Burgess. Charge up then charge out? Drivers' perceptions and experiences of electric vehicles in the UK. *Transportation Research Part A*, 59(C), January 2014.
- [BHT11] Matthew Barth, Jing Han, and Michael Todd. Performance Evaluation Of A Multi-Station Shared Vehicle System. *IEEE Intelligent Transportation Systems Conference*, 2011.
- [BKL09] Steven Bird, Ewan Klein, and Edward Loper. *Natural Language Processing with Python. Chapter 7: NLTK: Extracting Information from Text*. O'Reilly Media, Inc., 2009.
- [BLF90] Oded Berman, Richard Larson, and Nikoletta Fouska. Optimal Location of Discretionary Service Facilities. *Operations Research*, November 1990.
- [BLF92] Oded Berman, Richard Larson, and Nikoletta Fouska. Optimal Location of Discretionary Service Facilities . *Transportation Science*, 26, August 1992.
- [Blo12] Bloomberg. Energy & Oil Prices, 2012. Acc. Apr 20th 2012. <http://www.bloomberg.com/energy/>.
- [BMW14] BMW USA. BMW i, 2014. Acc. Aug 4th 2014. <http://www.bmwusa.com/standard/content/vehicles/2014/bmwi/default.aspx?mobileoverride=true#i3>.
- [Bro04] David Brook. Carsharing – Start Up Issues and New Operational Models. *Transportation Research Board*, January 2004.
- [Bro09] Chris Brown. How to run a successful car-sharing operation, Oct 2009. Acc. Apr 12th 2013, <http://www.autorentalnews.com/article/story/2009/09/how-to-run-a-successful-carsharing-operation/page/2.aspx>.

- [Bro12] Chris Brown. Car rental: A tough year for pricing?, Feb 2012. Acc. Apr 10th 2013, <http://www.businessfleet.com/blog/auto-focus/story/2012/02/car-rental-a-tough-year-for-pricing.aspx>.
- [BST09] Thomas A Becker, Ikhtlaq Sidhu, and Burghardt Tenderich. Electric Vehicles in the United States: A New Model with Forecasts to 2030. *Center for Entrepreneurship & Technology*, August 2009.
- [BT91] George J Beaujon and Mark A Turnquist. A Model for Fleet Sizing and Vehicle Allocation. *Transportation Science*, 25(1), Feb 1991.
- [BT99] Matthew Barth and Michael Todd. Simulation model performance analysis of a multiple station shared vehicle system. *Transportation Research Part C*, 7, Oct 1999.
- [Buc02] Buchholz, S. Penn treebank ii tag set. source: Memory-based grammatical relation finding, 2002. Acc. Sept 1 2014. <http://www.clips.ua.ac.be/pages/mbsp-tags>.
- [Bun10] Bundeswehr Transformation Center's Future Analysis Division. Peak Oil: Implications Of Resource Scarcity On (National) Security, jul 2010. Acc. May 1 2012. English summary: <http://www.theoildrum.com/node/6912>.
- [Car] Tommy Carpenter. A Python Electric Vehicle Sentiment Analysis System. <http://www.tommyjcarpenter.com/code/evsentiment.html>.
- [Car10] Steve Carroll. The Smart Move Trial: Description and Initial Results. *Centre of Excellence for Low Carbon and Fuel Cell Technologies Report*, March 2010.
- [Car14] Tommy Carpenter. Webike user portal, 2014. Production: <https://bitbucket.org/tommyjcarpenter/webikewebsiteprod>, Development: <https://bitbucket.org/tommyjcarpenter/webikewebsitedev>.
- [CBS11] CBS San Francisco. San Francisco Taxi Fares Go Up, September 2011. Acc. October 20th 2012. <http://sanfrancisco.cbslocal.com/2011/09/02/111118/>.
- [CCK13] Tommy Carpenter, Andrew R. Curtis, and S. Keshav. The return on investment for taxi companies transitioning to electric vehicles. *Transportation*, 40, Jul 2013.
- [CFL10] Alexander Clark, Chris Fox, and Shalom Lappin, editors. *The Handbook of Computational Linguistics and Natural Language Processing*. Blackwell Handbooks in Linguistics. John Wiley & Sons, 2010.
- [Cha13] Chargepoint America. ChargePoint America Vehicle Charging Infrastructure Summary Report, April 2013.

- [Cha14a] Max Chafkin. A Broken Place: The Spectacular Failure Of The Startup That Was Going To Change The World, April 2014. Acc. Aug 1 2014. <http://www.fastcompany.com/3028159/a-broken-place-better-place>.
- [Cha14b] ChargePoint. Chargepoint America, 2014. <http://chargepointamerica.com/>.
- [Che14a] Chevrolet. 2014 Volt Models & Specs, September 2014. Acc. Sept 4th, 2014. <http://www.chevrolet.com/2014-volt-electric-car/specs/trims.html>.
- [Che14b] Chevrolet. Chevrolet volt, 2014. Acc. Apr 4th 2014. <http://www.chevrolet.com/volt-electric-car.html>.
- [Cit10] City Of New York. Exploring Electric Vehicle Adoption in NYC . *City Of New York Report*, January 2010.
- [CJ89] Dah-Ming Chiu and Raj Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Comput. Netw. ISDN Syst.*, 17(1), June 1989.
- [CJA14] Gonçalo Homem De Almeida Correia, Diana Ramos Jorge, and David Marques Antunes. The added value of accounting for users's flexibility and information on the potential of a station-based one-way car-sharing system: An application in lisbon, portugal. *Journal of Intelligent Transportation Systems*, 18(3), 2014.
- [CMR12] Claudio Contardo, Catherine Morency, and Louis-Martin Rousseau. Balancing a Dynamic Public Bikeshearing System. *CIRRELT Report*, Mar 2012.
- [Com14] Community Carshare. Community Carshare, 2014. <http://communitycarshare.ca/>.
- [Cra09] Crawdad. GPS Mobility Data Set, February 2009. Acc. Oct 2010. <http://crawdad.cs.dartmouth.edu/meta.php?name=epfl/mobility>.
- [CWJ11] Christopher Cherry, Stacy Worley, and David Jordan. Electric Bikeshearing System Requirements and Operational Concepts. *Annual Meeting of the Transportation Research Board*, September 2011.
- [Das13] M.S. Daskin. *Network and Discrete Location: Models, Algorithms, and Applications, 2nd Edition*. Wiley, 2013.
- [Del10] Deloitte. Gaining traction: A customer view of electric vehicle mass adoption in the U.S. automotive market. *Deloitte Report*, June 2010.
- [Del11] Deloitte. Unplugged: Electric vehicle realities versus consumer expectations. *Deloitte Report*, 2011.

- [DF14a] M. Dozza and A Fernandez. Understanding bicycle dynamics and cyclist behavior from naturalistic field data (november 2012). *Intelligent Transportation Systems, IEEE Transactions on*, 15(1), Feb 2014.
- [DF14b] Marco Dozza and Andre Fernandez. Understanding Bicycle Dynamics and Cyclist Behavior From Naturalistic Field Data (November 2012). *IEEE Transactions on Intelligent Transportation Systems*, 15(1), February 2014.
- [DLY09] Xiaowen Ding, Bing Liu, and Philip S Yu. A Holistic Lexicon-Based Approach to Opinion Mining. *WSDM*, 2009.
- [DM09] V Reddy Dondeti and Bidhu B Mohanty. Motor Pool Fleet Sizing Determination: A Simulation Case Study. *Review Of Business Research*, 9(2), 2009.
- [DRvO13] Alexandros Dimitropoulos, Piet Rietveld, and Jos N van Ommeren. Consumer valuation of changes in driving range: A meta-analysis. *Transportation Research Part A*, 55(C), September 2013.
- [DWM13] M Dozza, J Werneke, and M Mackenzie. e-BikeSAFE: A Naturalistic Cycling Study to Understand how Electrical Bicycles Change Cycling Behaviour and Influence Safety. *International Cycling Safety Conference*, November 2013.
- [E-W14] E-WALD. E-carsharing sylt, 2014. Acc. Sept 30, 2014. <http://e-wald.eu/>.
- [EATM12] Dimitrios Efthymiou, Constantinos Antoniou, Yannis Tyrinopoylos, and Evangelos Mitsakis. Spatial Exploration of Effective Electric Vehicle Infrastructure Location. *Procedia - Social and Behavioral Sciences*, 48, January 2012.
- [EBH<sup>+</sup>11] Andrew Everett, Mark Burgess, Margaret Harris, Sarah Mansbridge, Elisa Lewis, Chris Walsh, and Steve Carrol. Initial Findings from the Ultra Low Carbon Vehicle Demonstrator Programme. *Technology Strategy Board Report*, September 2011.
- [ECO13] ECOtality North America. The EV Project Q4 2012 Report, February 2013.
- [Efr79] B Efron. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 7, 1979.
- [EL12] Ona Egbue and Suzanna Long. Barriers to widespread adoption of electric vehicles: An analysis of consumer attitudes and perceptions. *Energy Policy*, 48(C), September 2012.
- [Ele09] Electrification Coalition. Electrification Roadmap: Revolutionizing Transportation And Achieving Energy Security. *Electrification Coalition Report*, November 2009.
- [Ele14] Electric Drive Transportation Association. Electric drive sales dashboard, 2014. Acc. March 29th 2014. <http://electricdrive.org/index.php?ht=d/sp/i/20952/pid/20952>.

- [EML<sup>+</sup>09] Shane B Eisenman, Emiliano Miluzzo, Nicholas D Lane, Ronald A Peterson, Gahng-Seop Ahn, and Andrew T Campbell. BikeNet: A Mobile Sensing System for Cyclist Experience Mapping. *ACM Transactions on Sensor Networks*, 6(1), December 2009.
- [Eng12] Wiebe Engelmoer. The E-bike: opportunities for Commuter Traffic . *Masters Thesis, University of Groningen*, February 2012.
- [EPR10] EPRI. Characterizing Consumers' Interest in and Infrastructure Expectations for Electric Vehicles: Research Design and Survey Results. *EPRI Report*, May 2010.
- [EPR13] EPRI. Total Cost of Ownership Model for Current Plug-in Electric Vehicles. *EPRI Report*, June 2013.
- [ePr14] eProdigy. eProdigy Whistler, July 2014. Acc. Aug15, 2014. <http://www.eprodigybikes.com/whistler.aspx>.
- [Ern10] Ernst & Young. Gauging interest for plug-in hybrid and electric vehicles in select markets. *Ernst & Young Report*, June 2010.
- [ES06a] Andrea Esuli and Fabrizio Sebastiani. Determining Term Subjectivity and Term Orientation for Opinion Mining. *EACL*, 2006.
- [ES06b] Andrea Esuli and Fabrizio Sebastiani. SENTIWORDNET: A Publicly Available Lexical Resource for Opinion Mining. *LREC*, 2006.
- [FK13] Thomas Franke and Josef F Krems. What drives range preferences in electric vehicle users? *Transport Policy*, 30(C), November 2013.
- [Flo01] J.E. Flood. *Telecommunications Switching, Traffic and Networks*. Pearson Education, 2001.
- [FR00] R Farrington and J Rugh, editors. *Impact of Vehicle Air Conditioning on Fuel Economy, Tailpipe Emissions, and Electric Vehicle Range*. *Earth Technologies Forum*. National Renewable Energy Laboratory, 2000.
- [Fri02] Jeffrey E. F. Friedl. *Mastering Regular Expressions*. O'Reilly & Associates, Inc., 2 edition, 2002.
- [FSS13] Angela Di Febbraro, Nicola Sacco, and Mahnam Saeednia. One-Way Car-sharing: Solving the Relocation Problem. *Transportation Research Record: Journal of the Transportation Research Board*, 2319, Feb 2013.
- [GA05] Michael Gamon and Anthony Aue. Automatic identification of sentiment vocabulary: exploiting low association with known sentiment terms. *The ACL workshop on feature engineering for machine learning in NLP*, 2005.

- [Gal09] Kate Galbraith. Better Place Unveils Battery Swap Station, May 2009. Acc. Apr 25th, 2011. <http://tinyurl.com/kbn7q96>.
- [Gas] Gasbuddy. USA National Gas Price Heat Map. [http://gasbuddy.com/gb\\_gastemperaturemap.aspx](http://gasbuddy.com/gb_gastemperaturemap.aspx).
- [Geo14] Geotab. Geotab: Measurement by management, 2014. Acc. Sept 30, 2014. [www.geotab.com/](http://www.geotab.com/).
- [GHOC11] Megan C Guilford, Charles A S Hall, Peter O'Connor, and Cutler J Cleveland. Long Term Assessment of Energy Return on Investment (EROI) for U.S. Oil and Gas Discovery and Production. *Sustainability*, 3(12), December 2011.
- [GJ79] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.
- [GKS<sup>+</sup>12] T Gehlert, M Kuhn, K Schleinitz, T Petzoldt, S Schwanitz, and R Gerike. The German Pedelec Naturalistic Cycling Study – Study Design and First Experiences. *International Cycling Safety Conference*, November 2012.
- [GL97] Fred Glover and Manuel Laguna. *Tabu Search*. Kluwer Academic Publishers, Norwell, MA, USA, 1997.
- [Glo13] Global Research. Fracked Oil Spill Disaster in North Dakota, 2013. Acc. March 25th, 2014. <http://tinyurl.com/pkr45yk>.
- [GM12] Israel García and Luis Javier Miguel. Is the Electric Vehicle an Attractive Option for Customers? *Energies*, 5(1), January 2012.
- [GMST07] Jeffrey Gonder, Tony Markel, Andrew Simpson, and M Thornton. Using GPS Travel Data to Assess the Real World Driving Energy Use of Plug-In Hybrid Electric Vehicles (PHEVs). *Transportation Research Board (TRB) 86th Annual Meeting*, January 2007.
- [Goo10] David Goodman. An Electric Boost for Bicyclists, Jan 2010. Acc. July 20th 2014. [http://www.nytimes.com/2010/02/01/business/global/01ebike.html?\\_r=0](http://www.nytimes.com/2010/02/01/business/global/01ebike.html?_r=0).
- [GRGA<sup>+</sup>12] Ella Graham-Rowe, Benjamin Gardner, Charles Abraham, Stephen Skippon, Helga Dittmar, Rebecca Hutchins, and Jenny Stannard. Mainstream consumers driving plug-in battery-electric and plug-in hybrid electric cars. *Transportation Research Part A*, 46(1), January 2012.
- [Gur14] Gurobi Optimization. Gurobi, 2014. <http://www.gurobi.com/>.
- [GX11] David K George and Cathy H Xia. Fleet-sizing and service availability for a vehicle rental system via closed queueing networks. *European Journal of Operational Research*, 211(1), May 2011.



- [Ham11] James D Hamilton. Nonlinearities And The Macroeconomic Effects Of Oil Prices. *Macroeconomic Dynamics*, 2011.
- [HBW05] Robert L Hirsch, Roger Bezdek, and Robert Wendling. Peaking of world oil production: Impacts, mitigation, & risk management. *U.S. Department of Energy Report*, March 2005.
- [Hei09] Richard Heinberg. *Searching For A Mircale: 'Net Energy' Limits & the Fate of Industrial Society*. ChromaGraphic, 2009.
- [HG10] Robert C Hampshire and Craig Gaites. Peer-to-Peer Carsharing: Market Analysis and Potential Growth. *TRB 2011 Annual Meeting*, 2010.
- [HK88] M.G. Hluchyj and M.J. Karol. Queueing in high-performance packet switching. *Selected Areas in Communications, IEEE Journal on*, 6(9), Dec 1988.
- [HL04a] Mingqing Hu and Bing Liu. Mining and Summarizing Customer Reviews. *KDD*, 2004.
- [HL04b] Mingqing Hu and Bing Liu. Mining Opinion Features in Customer Reviews. *AAAI*, 2004.
- [HM12] Vasileios Hatzivassiloglou and Kathleen R McKeown. Predicting the Semantic Orientation of Adjectives. *ACL*, 2012.
- [Hon14] Honda. Honda civic, 2014. Acc. Apr 4th 2014. <http://automobiles.honda.com/civic-sedan/>.
- [IBM12] IBM. The Battery 500 Project, January 2012. Acc. Jul 15, 2014. [http://www.ibm.com/smarterplanet/us/en/smart\\_grid/article/battery500.html](http://www.ibm.com/smarterplanet/us/en/smart_grid/article/battery500.html).
- [IBM14] IBM. Cplex optimization studio, 2014. <http://www.ibm.com/developerworks/downloads/ws/ilogcplex/>.
- [IEA13] IEA. Global ev outlook. *IEA Whitepaper*, April 2013. Acc. July 31, 2014. [http://www.iea.org/publications/globalevoutlook\\_2013.pdf](http://www.iea.org/publications/globalevoutlook_2013.pdf).
- [Ing13] Anthony Ingram. BMW Electric Car Owners Get Gas-Powered Loaner For Longer Trips, March 2013. Acc. Apr 20th, 2013, [http://www.greencarreports.com/news/1082784\\_bmw-electric-car-owners-get-gas-powered-loaner-for-longer-trips?utm\\_source=GreenCarReports&utm\\_medium=twitter](http://www.greencarreports.com/news/1082784_bmw-electric-car-owners-get-gas-powered-loaner-for-longer-trips?utm_source=GreenCarReports&utm_medium=twitter).
- [Ing14] Antony Ingram. Electric Car Price Guide: Every 2014-2015 Plug-In Car, With Specs, June 2014. Acc. Jul 15, 2014. <http://tinyurl.com/kfrvv39>.



- [Inm13] Mason Inman. The true cost of fossil fuels. *Scientific American*, 308, 2013. Acc. March 29th 2014. Summary: [http://www.nature.com/scientificamerican/journal/v308/n4/box/scientificamerican0413-58\\_BX2.html](http://www.nature.com/scientificamerican/journal/v308/n4/box/scientificamerican0413-58_BX2.html).
- [Int12] International Energy Agency. Faqs: Oil, 2012. Acc. March 25th, 2014. <http://www.iea.org/aboutus/faqs/oil/>.
- [Inv] Investor Campus. How to analyze a car rental company. Available, <http://www.investorcampus.com/Courses/TheHowtoAnalyseSeries/Commerceandindustry/HowtoAnalyseaCarRentalCompanyFree/502.pdf>.
- [ISS14a] ISS4E. WeBike Hardware, July 2014. Acc. Aug15, 2014. [http://blizzard.cs.uwaterloo.ca/iss4e/?page\\_id=3763](http://blizzard.cs.uwaterloo.ca/iss4e/?page_id=3763).
- [ISS14b] ISS4E. WeBike Software, July 2014. Acc. Aug15, 2014. [http://blizzard.cs.uwaterloo.ca/iss4e/?page\\_id=3765](http://blizzard.cs.uwaterloo.ca/iss4e/?page_id=3765).
- [Ive09] Villy B Iversen. Teletraffic engineering and network planning (2009 version). *Course of Technical University of Denmark*, 2009.
- [Jan10] Jan Owen Jansson. Road pricing and parking policy. *Research in Transportation Economics*, 29(1), 2010.
- [JHS09] Wei Jin, Hung Hay Ho, and Rohini K Srihari. OpinionMiner: A Novel Machine Learning System for Web Opinion Mining and Extraction. *KDD*, 2009.
- [Kar13] Sten Karlsson. The Swedish car movement data project. *Chalmers University of Technology PRT Report*, November 2013.
- [KCLG13] Rachel M Krause, Sanya R Carley, Bradley W Lane, and John D Graham. Perception and reality: Public knowledge of plug-in electric vehicles in 21 U.S. cities. *Energy Policy*, 63(C), December 2013.
- [Kes12] Srinivasan Keshav. *Mathematical Foundations of Computer Networking*. Addison-Wesley, 2012.
- [KF09] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.
- [KGV83] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by Simulated Annealing. *Science*, 220, 4598, 1983.
- [KGWS13] Christopher King, Wynita Griggs, Fabian Wirth, and Robert Shorten. Using A Car Sharing Model To Alleviate Electric Vehicle Range Anxiety. *Yale Workshop on Adaptive and Learning Systems*, June 2013.
- [KH04] Soo-Min Kim and Eduard Hovy. Determining the Sentiment of Opinions. *COLING*, 2004.

- [KI06] ALISTAIR KENNEDY and DIANA INKPEN. Sentiment Classification of Movie Reviews Using Contextual Valence Shifters . *Computational Intelligence*, 2006.
- [KIM<sup>+</sup>04] Nozomi Kobayashi, Kentaro Inui, Yuji Matsumoto, Kenji Tateishi, and Toshikazu Fukushima. Collecting Evaluative Expressions for Opinion Extraction. *IJCNLP*, 2004.
- [KK12] Mobashwir Khan and Kara M Kockelman. Predicting the market potential of plug-in electric vehicles using multiday GPS data. *Energy Policy*, 46(C), July 2012.
- [KKN03] Peter Kochel, Sophie Kunze, and Ulf Nielander. Optimal control of a distributed service system with moving resources: Application to the fleet sizing and allocation problem. *International Journal of Production Economics*, 81-82(1), 2003.
- [KL05] Michael Kuby and Seow Lim. The flow-refueling location problem for alternative-fuel vehicles. *Socio-Economic Planning Sciences*, 39(2), June 2005.
- [Kle75a] Leonard Kleinrock. *Queueing Systems Volume 1*. Wiley-Interscience, 1975.
- [Kle75b] Leonard Kleinrock. *Theory, Volume 1, Queueing Systems*. Wiley-Interscience, 1975.
- [Koc97] Peter Kochel. An Approximate Model For Fleet Sizing And Redistribution. *Scientific Journal on Traffic and Transportation Research*, 9(5-6), 1997.
- [KR09] James F. Kurose and Keith W. Ross. *Computer Networking: A Top-Down Approach*. Addison-Wesley Publishing Company, 5th edition, 2009.
- [Kum13] Aswin Kumar. The lithium battery recycling challenge. *Waste Management World*, 2013. Acc. Apr 4th 2014. <http://tinyurl.com/llu6fad>.
- [Kur09] Kenneth Kurani. Plug-In Hybrid Electric Vehicle (PHEV) Demonstration and Consumer Education, Outread, and Market Research Program. *University of California Davis Institute of Transportation Studies Report*, UCD-ITS-RR-09-21, August 2009.
- [Lap92] Gilbert Laporte. The vehicle routing problem: An overview of exact and approximate algorithms. *European Journal of Operational Research*, 59(3), Jun 1992.
- [LB06] Annie Lubinsky and Chris Brown. Is Car Sharing A Threat To The Auto Rental Business? *ARN*, December 2006.

- [LCY<sup>+</sup>13] Brian Casey Langford, Christopher Cherry, Taekwan Yoon, Stacy Worley, and David Smith. North America's First E-bikeshare: A Year of Experience. In *Bicycles 2013: Planning, Design, Operations, and Infrastructure*. 2013.
- [LH] Bing Liu and Mingqing Hu. Opinion Mining, Sentiment Analysis, and Opinion Spam Detection. <http://www.cs.uic.edu/~liub/FBS/sentiment-analysis.html>.
- [LHC05] Bing Liu, Mingqing Hu, and Junsheng Cheng. Opinion Observer: Analyzing and Comparing Opinions on the Web. *WWW*, May 2005.
- [Lis12] Lindblad Liselotte. Deployment Methods For Electric Vehicle Infrastructure. Master's thesis, Uppsala University, July 2012.
- [Llo10] Lloyds Of London. Sustainable Energy Security; Strategic risks and opportunities for business. *LLOYDS Report*, June 2010.
- [LM94] S J Liebowitz and Stephen E Margolis. Network externality: An uncommon tragedy. *Journal of Economic Perspectives*, 8(2), Spring 1994.
- [LML<sup>+</sup>13] Kenneth Lebeau, Joeri Van Mierlo, Philippe Lebeau, Olivier Mairesse, and Cathy Macharis. Consumer attitudes towards battery electric vehicles: a large-scale survey. *International Journal of Electric and Hybrid Vehicles*, 5(1), 2013.
- [Lov11] Eric Loveday. Real-world results from CABLED electric vehicle trial, jul 2011. <http://tinyurl.com/mmuf9jh>.
- [LS10] Richard C Larson and Katsunobu Sasanuma. Congestion Pricing: A Parking Queue Model. *Industrial and Systems Engineering*, 4(1), September 2010.
- [LT09] Zhaolin Li and Feng Tao. On determining optimal fleet size and vehicle transfer policy for a car rental company . *Computers and Operation Research*, 37(2), Sep 2009.
- [Mac07] Ritch Macefield. Usability Studies and the Hawthorne Effect. *Journal Of Usability Studies*, 2(3), May 2007.
- [Mac09] D.J.C. MacKay. *Sustainable Energy: Without the Hot Air*. Uit Cambridge Limited, 2009.
- [Mat04] Matthew Barth, Michael Todd, and Lei Xue. User-Based Vehicle Relocation Techniques for Multiple-Station Shared-Use Vehicle Systems. *80th Annual Meeting of the Transportation Research Board*, 04-4161, Jan 2004.
- [Mat14] MathWorks. Matlab, 2014. <http://www.mathworks.com/products/matlab/>.

- [MbMS<sup>+</sup>05] Adam Millard-ball, Gail Murray, Jessica Ter Schure, Christine Fox, and Jon Burkhardt. Car-Sharing: Where and How It Succeeds. *Transit Cooperative Research Report*, 2005.
- [Mer14] Mercedes-Benz. E-Class Sedan, 2014. Acc. Aug 4th 2014. <https://victoriastar.mercedes-benz.ca/en-CA/Vehicles/E-Class-Sedan>.
- [MHN10] Elise Miller-Hooks and Rahul Nair. Fleet Management for Vehicle Sharing Operations. *Mid-Atlantic Universities Transportation Center Report*, UMD-2008-02, May 2010.
- [MHN<sup>+</sup>12] McKinsey&Company, Russell Hensley, John Newman, Matt Rogers, and Mark Shahinian. Battery technology charges ahead. *McKinsey Whitepaper*, July 2012.
- [Mis14a] Misc. Forum: Tesla motors club, 2014. <http://www.teslamotorsclub.com/>.
- [Mis14b] Misc. Forum: Volt ownership forum, 2014. <http://gm-volt.com/forum/forumdisplay.php?16-Volt-Ownership-Forum>.
- [Mis14c] Misc. My Nissan Leaf Discussion Forums, 2014. <http://www.mynissanleaf.com/>.
- [MK11] Sashank Musti and Kara M Kockelman. Evolution of the household vehicle fleet: Anticipating fleet composition, PHEV adoption and GHG emissions in Austin, Texas. *Transportation Research Part A*, 45(8), October 2011.
- [MSA11] Craig Morton, Geertje Schuitema, and Jillian Anable. EVs: Will Consumers get Charged Up? . *University of Aberdeen Centre of Transport Research Report*, January 2011.
- [MW05] Naoto Mukai and Toyohide Watanabe. Dynamic Location Management for On-Demand Car Sharing System. *Lecture Notes in Computer Science*, 3681, 2005.
- [Nat10] National Research Council. Transitions to Alternative Transportation Technologies—Plug-in Hybrid Electric Vehicles. *The National Academies Press*, 978-0-309-14850-4, 2010.
- [Nat13] National Bicycles Dealers Association. A Look at the Bicycle Industry's Vital Statistics, Dec 2013. Acc. July 20th 2014. <http://nbda.com/articles/industry-overview-2012-pg34.htm>.
- [NBW12] Jeremy Neubauer, Aaron Brooker, and Eric Wood. Sensitivity of battery electric vehicle economics to drive patterns, vehicle range, and charge strategies. *Journal of Power Sources*, 209, July 2012.

- [Nil11] Maria Nilsson. Electric Vehicles: The Phenomenon Of Range Anxiety. *ELVIRE Report*, June 2011.
- [Nis14a] Nissan. Compare 2014 nissan leaf prices and specs, 2014. Acc. Apr 4th 2014. <http://www.nissanusa.com/electric-cars/leaf/>.
- [Nis14b] Nissan. The Nissan Leaf, September 2014. Acc. Sept 4th, 2014. <http://www.nissanusa.com/leaf-electric-car/faq/list/charging#/leaf-electric-car/faq/list/charging>.
- [NLT13] NLTK Project. NLTK 2.0 documentation, 2013. <http://nltk.org/>.
- [NLT14] NLTK Project. Nltk 3.0 documentation—source code for nltk.tag.stanford, 2014. Acc. Sept 1 2014. [http://www.nltk.org/\\_modules/nltk/tag/stanford.html](http://www.nltk.org/_modules/nltk/tag/stanford.html).
- [NPK<sup>+</sup>11] Roni A Neff, Cindy L Parker, Frederick L Kirschenmann, Jennifer Tinch, and Robert S Lawrence. Peak Oil, Food Systems, and Public Health. *American Journal of Public Health*, 101(9), September 2011.
- [NW14] Jeremy Neubauer and Eric Wood. The impact of range anxiety and home, workplace, and public charging infrastructure on simulated battery electric vehicle lifetime utility. *Journal of Power Sources*, 257(c), July 2014.
- [OABM<sup>+</sup>10] Michael OKeefe, Caley Johnson Aaron Brooker, Mike Mendelsohn, Jeremy Neubauer, and Ahmad Pesaran. Battery ownership model: A tool for evaluating the economics of electrified vehicles and related infrastructure. *EVS-25*, November 2010.
- [OAEG11] SylvesterOlubolu Orimaye, SaadatM. Alhashmi, and Siew Eu-Gene. Using predicate-argument structures for context-dependent opinion retrieval. volume 7121. 2011.
- [Ohn14] Alan Ohnsman. Tesla rises after model s sales in 2013 exceed forecast, 2014. Acc. March 29th 2014. <http://tinyurl.com/lrr6okv>.
- [Ope] OpenDover. Opinion Miner. <http://tinyurl.com/kyh7g8y>.
- [OPE12] OPEC. Oil reservesopec share of world crude oil reserves, 2012. Acc. March 25th, 2014. [http://www.opec.org/opec\\_web/en/data\\_graphs/330.htm](http://www.opec.org/opec_web/en/data_graphs/330.htm).
- [Opi] OpinionEQ. Sentiment Analysis with Unmatched Accuracy. <http://www.opinioneq.com/>.
- [Par77] S C Parikh. On a Fleet Sizing and Allocation Problem. *Management Science*, 23(9), May 1977.
- [Par97] C. Paradis. *Degree modifiers of adjectives in spoken British English*. Lund studies in English. Lund University Press, 1997.

- [Par06] Alan A Parker. Electric Power-Assisted Bicycles Reduce Oil Dependence and Enhance the Mobility of the Elderly . *29th Australian Transport Research Forum*, September 2006.
- [PD14] Anja Peters and Elisabeth Dütschke. How do Consumers Perceive Electric Vehicles? A Comparison of German Consumer Groups. *Journal of Environmental Policy & Planning*, January 2014.
- [PE05] Ana-Maria Popescu and Oren Etzioni. Extracting Product Features and Opinions from Reviews. *HLT*, 2005.
- [PGS+14] Natalie Popovich, Elizabeth Gordon, Zhenying Shao, Yan Xing, Yunshi Wang, and Susan Handy. Experiences of electric bicycle users in the Sacramento, California area. *Travel Behaviour and Society*, 1(2), May 2014.
- [PJO95] Warren B Powell, Patrick Jaillet, and Arnedeo Odoni. Stochastic and Dynamic Networks and Routing. *Handbooks In Operations Research and Management Sciences*, 8, 1995.
- [PKGE11] Nathaniel Pearre, Willett Kempton, Randall L Guensler, and Vetri V Elango. Electric vehicles: How much range is required for a days driving? *Transportation Research Part C*, 19(6), December 2011.
- [Plo14] Patrick Plotz. How to estimate the probability of rare long-distance trips. *Working Paper*, 2014.
- [PM10] Johannes Paefgen and Florian Michahelles. Inferring Usage Characteristics of Electric Bicycles from Position Information. *LocWeb*, November 2010.
- [Pri14] Princeton University. WordNet: a lexical database for English, 2014. <http://wordnet.princeton.edu/>.
- [Pru10] Rémy Prud'homme. Electric vehicles: a tentative economic and environmental evaluation. *International Transport Forum*, November 2010.
- [PT08] Felix Papier and Ulrich W Thonemann. Queuing Models for Sizing and Structuring Rental Fleets. *Transportation Science*, 42(3), Aug 2008.
- [Pug08] Puget Sound Regional Council. Traffic choices reports, 2008. Acc. March 29th 2014. <http://www.psrc.org/transportation/traffic/reports>.
- [Ras09] Edie Rasmussen. Stoplists. In LING LIU and M.TAMER ĀŪZSU, editors, *Encyclopedia of Database Systems*. Springer US, 2009.
- [Rea14] Richard Read. New vehicle fuel economy hits 25.8 mpg in august, another all-time high, August 2014. Acc. Sept 4th, 2014. <http://tinyurl.com/lstrasg4>.

- [Ree93] Colin R. Reeves, editor. *Modern Heuristic Techniques for Combinatorial Problems*. John Wiley & Sons, Inc., 1993.
- [RFI10] RFI. Bp oil spill could be worst us environmental disaster, 2010. Acc. March 25th, 2014. <http://tinyurl.com/lhlgu9ky>.
- [Rio51] John Riordan. Telephone Traffic Time Averages. *The Bell System Technical Journal*, 30, October 1951.
- [RN03] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2003.
- [Rog03] Everett M. Rogers. *Diffusion of innovations*. Free Press, New York, NY [u.a.], 5th edition, 2003.
- [Rol11] Roland Berger Strategy Consultants. Automotive insights no. 01\_2011, Jan 2011. Available, [http://www.rolandberger.com/media/pdf/Roland\\_Berger\\_Automotive\\_Insights\\_1\\_2011\\_20110908.pdf](http://www.rolandberger.com/media/pdf/Roland_Berger_Automotive_Insights_1_2011_20110908.pdf).
- [Rul13] Clayton Rullo. Hertz: Increasing utilization feels good, Apr 2013. Acc. Apr 10th 2013, <http://seekingalpha.com/article/1327391-hertz-increasing-utilization-feels-good>.
- [SB08] Neeraj Shidore and Theodore Bohn. Evaluation of Cold Temperature Performance of the JCS-VL41M PHEV Battery Using Battery HIL, January 2008. Argonne National Laboratory, USA. [http://www.autonomie.net/docs/6%20-%20Papers/CIL/evaluation\\_of\\_cold\\_temperature.pdf](http://www.autonomie.net/docs/6%20-%20Papers/CIL/evaluation_of_cold_temperature.pdf).
- [SBC<sup>+</sup>07] Christopher Scaffidi, Kevin Bierhoff, Eric Chang, Mikhael Felker, Herman Ng, and Chun Jin. Red Opal: Product-Feature Scoring from Reviews. *EC*, 2007.
- [SCC09] Susan A Shaheen, Adam P Cohen, and Melissa S Chung. North American Carsharing: 10-Year Retrospective. *Journal of the Transportation Research Board*, 2110, 2009.
- [Sch06] Schaller Consulting. The New York City Taxi Fact Book, March 2006. Acc. July 4th, 2011. <http://www.schallerconsult.com/taxi/taxifb.pdf>.
- [Sch13] Schatz Energy Research Center. Macro-scale Siting of Plug-in Electric Vehicle Infrastructure for the North Coast of California. *Schatz Energy Research Center Report*, May 2013.
- [SCR06] Susan A Shaheen, Adam P Cohen, and J Darius Roberts. Carsharing in North America: Market growth, current developments and future potential. *Journal of the Transportation Research Board*, 2006.
- [Scr14] Scrapy Project. Scrapy, 2014. <http://scrapy.org/>.



- [SDT<sup>+</sup>09] Shane Slater, Mike Dolman, Phil Taylor, Pavlos Trichakis, and Jonathan Shine. Strategies for the uptake of electric vehicles and associated infrastructure implications. *Element Energy Report*, October 2009.
- [SG11] Stephen Skippon and Mike Garwood. Responses to battery electric vehicles: UK consumer attitudes and attributions of symbolic meaning following direct experience to reduce psychological distance. *Transportation Research Part D*, 16(7), October 2011.
- [SHvH13] Jasper Schuijbroek, Robert Hampshire, and Willem-Jan van Hove. Inventory Rebalancing and Vehicle Routing in Bike Sharing Systems. *Working Paper*, Feb 2013.
- [SM07] E Sullivan and L Magid. Bringing Car-Sharing to Your Community. *City Car-share*, 2007. Acc. June 16, 2014. [http://citycarshare.org/wp-content/uploads/2011/12/CCS\\_BCCtYC\\_Short.pdf](http://citycarshare.org/wp-content/uploads/2011/12/CCS_BCCtYC_Short.pdf).
- [SPV11] Aviral Shukla, Joseph Pekny, and Venkat Venkatasubramanian. An optimization framework for cost effective design of refueling station infrastructure for alternative fuel vehicles. *Computers and Chemical Engineering*, 35(8), August 2011.
- [SPV13] Beena M Varghese Shamna P V. State-of-charge estimation of lithium-ion batteries combining current and voltage based estimation. *International Journal of Engineering and Innovative Technology*, 2, June 2013.
- [SS98] Conrad Wagner Susan Shaheen, Darnel Sperling. Carsharing in europe and north. america: Past, present and future. *Transportation Quarterly*, 52, 1998.
- [SSW04] Susan A Shaheen, Andrew Schwartz, and Kamill Wipyewski. Policy Considerations for Carsharing and Station Cars. *Journal of the Transportation Research Board*, 1887, 2004.
- [Sta14] Stanford Natural Language Processing Group. Stanford log-linear part-of-speech tagger, 2014. Acc. Sept 1 2014. <http://nlp.stanford.edu/software/tagger.shtml>.
- [TAV06] Maite Taboada, Caroline Anthony, and Kimberly Voll. Methods for Creating Semantic Orientation Dictionaries. *LREC*, 2006.
- [TC03] Egidio Terra and C. L. A. Clarke. Frequency estimates for statistical word similarity measures. In *NAACL*, volume 1, 2003.
- [Ten10] Jim Tenant. How to improve utilization  $\text{\$}$  and still satisfy the customer, Feb 2010. Acc. Apr 10th, 2013, <http://tinyurl.com/n5znmwb>.
- [Tes14a] Tesla. Efficiency, 2014. Acc. March 29th 2014. <http://www.teslamotors.com/goelectric/efficiency>.



- [Tes14b] Tesla. Gigafactory, February 2014. Acc. Jul 15, 2014. <http://www.teslamotors.com/blog/gigafactory>.
- [Tes14c] Tesla. The Tesla Model S, 2014. Acc. Apr 4th 2014. <http://www.teslamotors.com>.
- [Tes15] Tesla. Supercharger, 2015. Acc. Jan 10th 2015. <http://www.teslamotors.com/supercharger>.
- [TGLW11] Tom Turrentine, Dahlia Garas, Andy Lentz, and Justin Woodjack. The UC Davis MINI E Consumer Study. *University of California Davis Institute of Transportation Studies Report*, UCD-ITS-RR-11-05, May 2011.
- [Tha06] Paul Thagard. *Hot Thought: Mechanisms and Applications of Emotional Cognition*. Cambridge MA: Bradford Book/MIT Press, 2006.
- [The08] The Exploratorium. Cabspotting, August 2008. Acc. Feb 12th 2011. <http://cabspotting.org/>.
- [Tij03] H.C. Tijms. *A First Course in Stochastic Models*. Wiley, 2003.
- [Tim10] Time. 100 days of the bp spill: A timeline, 2010. Acc. March 25th 2014. <http://content.time.com/time/interactive/0,31813,2006455,00.html>.
- [Tim13] Heather Timmons. Consider the e-bike: Can 200 million chinese be wrong?, October 2013. Acc. July 20 2014. <http://tinyurl.com/q8dqjcs>.
- [TMLK11] Xidong Tang, Xiaofeng Mao, Jian Lin, and B. Koch. Li-ion battery parameter estimation for state of charge. *American Control Conference*, June 2011.
- [TN13] Gil Tal and Michael A Nicholas. Studying the PEV Market in California: Comparing the PEV, PHEV and Hybrid Markets. *EVS-27*, November 2013.
- [Tom07] Tom Gorin and Kurt Pisor. California's Residential Electricity Consumption 1980-2005. Acc. July 5th, 2011. <http://tinyurl.com/l283ff6>, 2007.
- [Toy14a] Toyota. Prius 2014, June 2014. Acc. Jul 15, 2014. <http://www.toyota.com/prius/#!/Welcome>.
- [Toy14b] Toyota. Toyota corolla, 2014. Acc. Apr 4th 2014. <http://www.toyota.com/corolla/#!/Welcome>.
- [TP10] Mikalai Tsytarau and Themis Palpanas. Disambiguating Dynamic Sentiment Ambiguous Adjectives. *Proceedings of the 23rd International Conference on Computational Linguistics*, 2010.
- [Tur02] Peter D Turney. Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. *ACL*, 2002.

- [UMW07] Kentaro Uesugi, Naoto Mukai, and Toyohide Watanabe. Optimization of Vehicle Assignment for Car Sharing System. *Lecture Notes in Computer Science*, 4693, 2007.
- [Uni10] United States Department Of Labor. Average energy prices in the San Francisco area. Acc. July 1st, 2011. [http://www.bls.gov/ro9/cpisanf\\_energy.pdf](http://www.bls.gov/ro9/cpisanf_energy.pdf), 2010.
- [US 13] US Department of Transportation. National household travel survey, 2013. Acc. March 25th 2014. <http://nhts.ornl.gov/>.
- [U.S12] U.S. Energy Information Administration. How much electricity is lost in transmission and distribution in the united states?, 2012. Acc. Apr 1st, 2014. <http://www.eia.gov/tools/faqs/faq.cfm?id=105&t=3>.
- [U.S13a] U.S. Department of Energy, Ecotality. The EV Project, 2013. <http://www.theevproject.com/>.
- [U.S13b] U.S. Department of State. Oil embargo, 1973Ð1974, 2013. Acc. March 29th 2014. <http://history.state.gov/milestones/1969-1976/oil-embargo>.
- [U.S13c] U.S. Energy Information Administration. How much oil does the united states consume per year?, 2013. Acc. Apr 1st, 2014. <http://www.eia.gov/tools/faqs/faq.cfm?id=33&t=6>.
- [U.S13d] U.S. Energy Information Administration. Transportation sector energy consumption, 2013. Acc. Apr 1st, 2014. <http://www.eia.gov/forecasts/ieo/transportation.cfm>.
- [U.S14a] U.S. Bureau of Labor Statistics. Average Energy Prices, San Francisco-Oakland-San Jose Ð July 2014, July 2014. Acc. Sept 4th 2014. [http://www.bls.gov/ro9/cpisanf\\_energy.htm](http://www.bls.gov/ro9/cpisanf_energy.htm).
- [U.S14b] U.S. Energy Information Administration. Electricity Explained: Factors Affecting Electricity Prices, 2014. [http://www.eia.gov/energyexplained/index.cfm?page=electricity\\_factors\\_affecting\\_prices](http://www.eia.gov/energyexplained/index.cfm?page=electricity_factors_affecting_prices).
- [U.S14c] U.S. Energy Information Administration. What was the highest u.s. average retail price of regular gasoline?, 2014. Acc. March 29th 2014. <http://www.eia.gov/tools/faqs/faq.cfm?id=247&t=10>.
- [VN14] Richard Van Noorden. A Better Battery. *Nature*, 507, March 2014.
- [VNR12] Erik Verhoff, Peter Nijkamp, and Piet Rietveld. The economics of regulatory parking policies. *Transportation Research Part A*, 29, 2012.

- [Voe14] John Voelcker. Plug-in electric car sales for 2013 near double previous year's, 2014. Acc. March 29th 2014. <http://tinyurl.com/lkptcqf>.
- [Was04] Larry Wasserman. *All of Statistics: A Concise Course in Statistical Inference (Springer Texts in Statistics)*. Springer, October 2004.
- [WDL14] Xing Wu, Jing Dong, and Zhenhong Lin. Cost analysis of plug-in hybrid electric vehicles using GPS-based longitudinal travel data. *Energy Policy*, 68(C), May 2014.
- [WF79] William W Williams and Oscar S Fowler. Model Formulation For Fleet Size Analysis Of A University Motor Pool. *Decision Sciences*, 10(3), 1979.
- [WHS14] Lena Winslott Hiselius and Åse Svensson. Could the increased use of e-bikes (pedelecs) in Sweden contribute to a more sustainable transport system? In *The 9th International Conference on Environmental Engineering*. Vilnius Gediminas Technical University Press, 2014.
- [Wik14a] Wikipedia. Lists of common misspellings, 2014. [http://en.wikipedia.org/wiki/Wikipedia:Lists\\_of\\_common\\_misspellings/For\\_machines](http://en.wikipedia.org/wiki/Wikipedia:Lists_of_common_misspellings/For_machines).
- [Wik14b] Wikipedia. Wikipedia:list of english contractions, 2014. [http://en.wikipedia.org/wiki/Wikipedia:List\\_of\\_English\\_contractions](http://en.wikipedia.org/wiki/Wikipedia:List_of_English_contractions).
- [Wil] Theresa Wilson. Mpqa opinion corpus. <http://www.cs.pitt.edu/mpqa/>.
- [Wil08] Theresa Ann Wilson. *Fine-grained Subjectivity and Sentiment Analysis: Recognizing the Intensity, Polarity, and Attitudes of Private States*. PhD thesis, University of Pittsburgh, 2008.
- [WLK12] Johannes Wirgesa, Susanne Linderb, and Alois Kessler. Modelling the Development of a Regional Charging Infrastructure for Electric Vehicles in Time and Space . *EJTIR*, 12(4), 2012.
- [WSNdH14] Ingo Wolf, Tobias Schroeder, Jochen Neumann, and Gerhard de Haan. Changing minds about electric cars: An empirically grounded agent-based modeling approach. *CoRR*, abs/1405.6230, 2014.
- [WU13] WXY Architecture and Urban Design. EVSE Cluster Analysis: Electric Vehicle Supply Equipment Support Study. *Transportation and Climate Initiative Report*, January 2013.
- [WW11] Miaomiao Wen and Yunfang Wu. Mining the Sentiment Expectation of Nouns Using Bootstrapping Method. *IJCNLP*, 2011.
- [WWH04] Theresa Wilson, Janyce Wiebe, and Rebecca Hwa. Just how mad are you? Finding strong and weak opinion clauses. *AAAI*, 2004.

- [Yar09] Jay Yarow. The Cost Of A Better Place Battery Swapping Station: \$500,000, April 2009. Acc. Jan 11th 2011. <http://tinyurl.com/crgqzx>.
- [Yel11] Yellow Cab San Francisco. Yellow cab san francisco rates, July 2011. Acc. July 4th, 2011. <http://www.yellowcabsf.com/our-service/cab-fares/>.
- [YL00] Y. R. Yang and S. S. Lam. General aimd congestion control. In *Proceedings of the 2000 International Conference on Network Protocols, ICNP '00*. IEEE Computer Society, 2000.
- [You13] Angelo Young. Avis buys zipcar to compete better against hertz and enterprise, Jan 2013. Acc. Apr 10th 2013, <http://tinyurl.com/myh77e2>.
- [YWWS12] Kwo Young, Caisheng Wang, Le Yi Wang, and Kai Strunz. Electric Vehicle Battery Technologies. In *Electric Vehicle Integration into Modern Power Networks*. Springer New York, New York, NY, October 2012.
- [Zfa10] Zfacts. Current Gas Prices and Price History, September 2010. Figures from Department of Energy. Acc. July 12th 2011. <http://zfacts.com/p/35.html>.
- [ZJZ06] Li Zhuang, Feng Jing, and Xiao-Yan Zhu. Movie Review Mining and Summarization. *CIKM*, 2006.
- [ZNC10] Kunpeng Zhang, Ramanathan Narayanan, and Alok Choudhary. Voice of the Customers: Mining Online Customer Reviews for Product Feature-based Ranking. *WOSN*, 2010.

# Appendix A

## Sentiment Analysis Chunking Grammar

Here we give our chunking grammar referenced in §4.6.

```
impt = "<IgnoreThisChunk|PRESENCE|VS|VB.*|DT|RB.*|PRP.*|CD|PDT|POS>*"
op = "<LINTM|HINTM|NONADJOP|IMPLICITFEAT|JJ.*>+"
feat = "<POSFEAT|NEGFEAT|REGFEAT|IMPLICITFEAT>"
grammar = """vs-feat:    {{<VS><NEGFEAT>}}    #e.g, no range anxiety
            op-feat-op  {{{op}{feat}{op}}}
            op-feat:    {{{impt}{op}{impt}{feat}}}
            feat-op:    {{{impt}{feat}{impt}{op}}}
            """
```

The nonstandard tags:

- "HINTM" and "LINTM" represent higher and lower intensity modifiers (see §4.7.1).
- "POSFEAT" and "NEGFEAT" refer to oriented features (see §4.7.1).
- "IgnoreThisChunk" represents the querying exception phrases (see §4.7.2).
- "Presence" represents words indicating the presence of something, e.g., "have" and "has".

For information about the other (standard NLP) tags, see the Stanford Tagger [Sta14] and the Penn Treebank tag set [Buc02]. We note that the removal of stopwords, as discussed in §4.6.1 helps write more concise grammars. For example, before filtering stopwords, we had six rules and the "impt" list included more tag types.

## **Appendix B**

### **Metric Over Time Figures For Chapter 6**

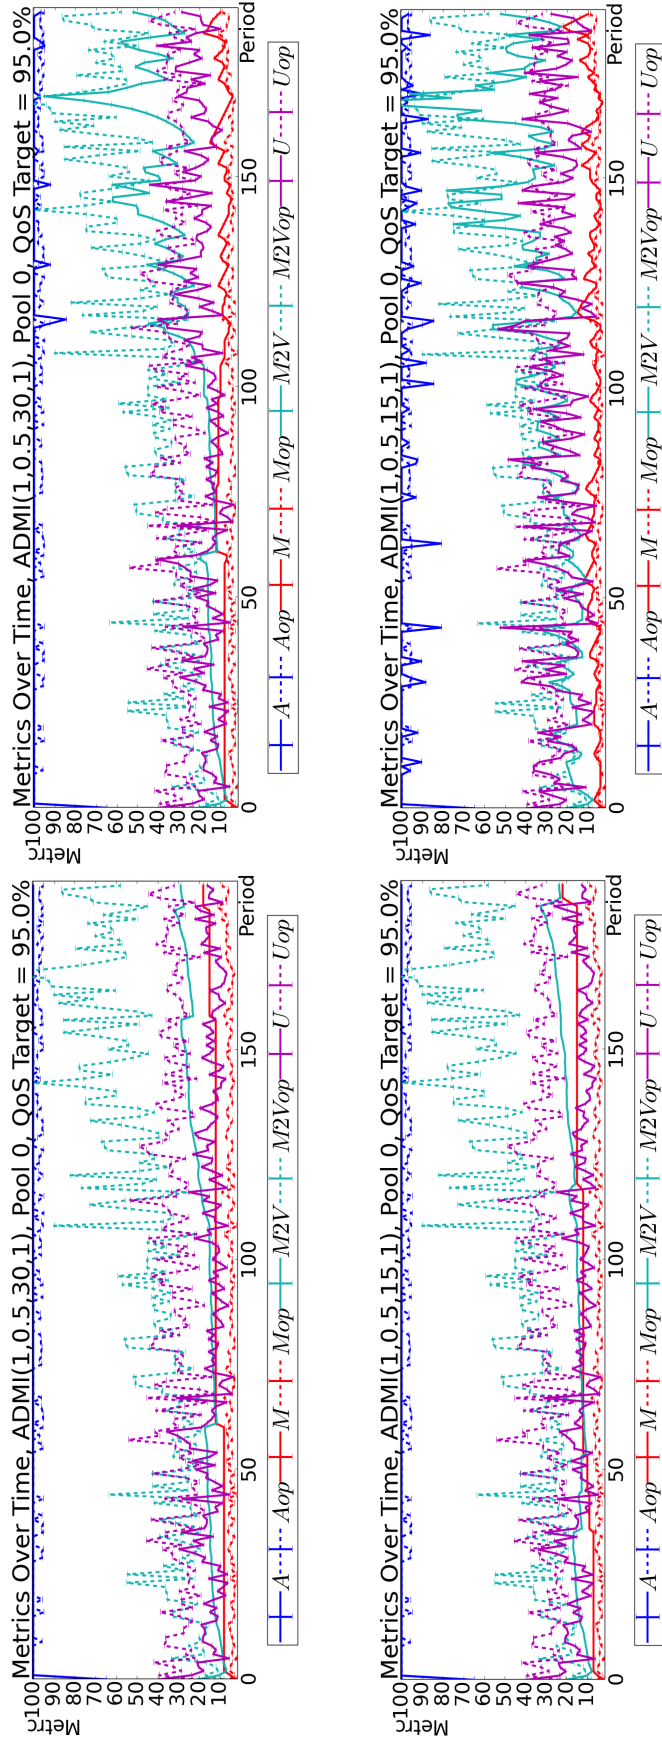










Figure B.1: Top left: ADM(1,0.5,30,1) without removals, top right: ADM(1,.5,30,1) with removals implemented, bottom left: ADM(1,0.5,15,1) without removals, bottom right: ADM(1,.5,15,1) with removals implemented. These are plotted together with the offline optimal baseline for a 95% QoS target. The averages of each metric achieved over all periods are summarized in Table 6.6.









(a) Performance metrics for two ELM and two ADM1 sizing configurations. All results shown are for Pool 0. ELM results are for a QoS of 95%, and ADM1 is plotted together with the offline optimal baseline for a 95% QoS target. The averages of each metric achieved over all periods are summarized in Table 6.3.







(b) Performance metrics for ELM(96) and ELM(168), the same two ELM configurations shown in Figure B.2a, for QoS targets of 99% (left) and 90% (right). All results shown are for Pool 0. The averages of each metric achieved over all periods are summarized in Table 6.4.





(a) Metrics for both carshare pools for one ELM and one ADMI sizing configuration (different configurations than those shown in Figures B.2a and B.2b), for a QoS target of 95%. The averages of each metric achieved over all periods are summarized in Table 6.5.



(b) In order from top left to bottom right: ELM(8) without removals, ELM(8) with removals implemented, ELM(72) without removals, ELM(72) with removals implemented. QoS=95%.

# Appendix C

## Proof Of Convergence For Algorithm 1

**Credit:** *this proof was co-authored with Parsiad Azimzadeh after Chapter 5 was published. The results given in §C.3 are not intended as contributions of the author's thesis, but are presented here for completeness in discussing Algorithm 1 terminates.*

In this section, we prove the convergence conditions of Algorithm 1 discussed in §5.2.2. Recall that algorithm 1 computes  $p(b|m)$  using the recurrence given by Eq(5.4) for numerical stability. However, Equations 5.4 and 5.2 are algebraically equivalent [lve09], so we use Eq(5.2), without loss of correctness, directly for simplicity.

**Assumption 1.**  $m$  and  $S$  are integers satisfying  $1 \leq m \leq S$ .

### C.1 Blocking Probability As A Fixed Point

We substitute Eq(5.14) into Eq(5.15) to yield:

$$\rho = \frac{1/\mu}{\frac{SK}{n_B} - (1 - p(b|m)) \frac{1}{\mu}} = [v + p(b|m)]^{-1} \quad (\text{C.1})$$

$$\text{where } v = \frac{\mu SK}{n_B} - 1 \quad (\text{C.2})$$

Substituting the above into Eq(5.2) yields

$$p(b|m) = \frac{\binom{S}{m} (v + p(b|m))^{-m}}{\sum_{i=0}^m \binom{S}{i} (v + p(b|m))^{-i}} = \left[ \sum_{i=0}^m \frac{\binom{S}{i}}{\binom{S}{m}} (v + p(b|m))^{m-i} \right]^{-1}$$

We then write  $p(b|m)$  as a fixed-point of a function:

$$f(x) = \left[ \sum_{i=0}^m \frac{\binom{S}{i}}{\binom{S}{m}} (v + x)^{m-i} \right]^{-1}$$

## C.2 Existence and Uniqueness of a Fixed Point

Here we verify that a fixed point of  $f$  is a probability.

**Lemma 1.** *If  $v \geq 0$ ,  $f$  admits a unique fixed point on  $[0, \infty)$ . Furthermore, this fixed point lies in  $(0, 1]$ .*

*Proof.* This follows immediately from  $f$  being nonincreasing on  $[0, \infty)$ . The second part of the claim is verified by noting that  $0 < f(x) \leq 1$  for any  $x$  in  $[0, \infty)$ .  $\square$

## C.3 Fixed Point Iteration

To prove that the fixed point iteration converges, we find sufficient conditions for  $f$  to be a contraction. In particular, since  $f$  is continuously differentiable on  $[0, 1]$ , we consider  $|f'|$ .

**Definition 1.** For a nonnegative integer  $n$ , the *rising Pochhammer symbol* is

$$x^{(n)} = \begin{cases} x(x+1)\dots(x+n-1) & n > 0 \\ 1 & n = 0 \end{cases}$$

Similarly, the *falling Pochhammer symbol* is

$$x_{[n]} = \begin{cases} x(x-1)\dots(x-n+1) & n > 0 \\ 1 & n = 0 \end{cases}$$

**Assumption 2.**  $v > 0$  and  $S \geq 2m - 1$ .

**Lemma 2.** *Under Assumption 2,  $|f'| < 1$  on  $[0, \infty)$ .*

*Proof.* Let  $n = S - m + 1$  and  $y = v + x$ . For succinctness, rewrite  $f$  as

$$f(x) = \left[ \sum_{i=0}^m y^i \frac{m_{[i]}}{n^{(i)}} \right]^{-1}$$

Differentiating yields  $|f'| = A/B$  where

$$A = \sum_{i=0}^{m-1} y^i \frac{m_{[i+1]}}{n^{(i+1)}} (i+1)$$

and, expanding using the Cauchy product,

$$B = \left[ \sum_{i=0}^m y^i \frac{m_{[i]}}{n^{(i)}} \right]^2 = \sum_{i=0}^{2m} y^i \sum_{k=0}^i \frac{m_{[k]}}{n^{(k)}} \frac{m_{[i-k]}}{n^{(i-k)}}.$$

To arrive at the desired result, we need only show  $B - A > 0$ . Since  $y > 0$ , it is sufficient to show that the (finite) power series (in  $y$ ):

$$B - A = \sum_{i=0}^{m-1} y^i \left[ \sum_{k=0}^i \frac{m_{[k]}}{n^{(k)}} \frac{m_{[i-k]}}{n^{(i-k)}} - \frac{m_{[i+1]}}{n^{(i+1)}} (i+1) \right] + \sum_{i=m}^{2m} y^i \sum_{k=0}^i \frac{m_{[k]}}{n^{(k)}} \frac{m_{[i-k]}}{n^{(i-k)}}$$

has nonnegative coefficients  $c_i$  (with at least one being positive). In this form, we see that the  $y$  terms of order  $m$  to  $2m$  have positive coefficients. We direct our attention to the terms of lower order. In particular, for fixed  $i$  satisfying  $0 \leq i < m$ :

$$\begin{aligned} c_i &= \sum_{k=0}^i \left[ \frac{m_{[k]}}{n^{(k)}} \frac{m_{[i-k]}}{n^{(i-k)}} \right] - \frac{m_{[i+1]}}{n^{(i+1)}} (i+1) = \sum_{k=0}^i \left[ \frac{m_{[k]}}{n^{(k)}} \frac{m_{[i-k]}}{n^{(i-k)}} - \frac{m_{[i+1]}}{n^{(i+1)}} \right] \\ &= \sum_{k=0}^i \frac{m_{[k]}}{n^{(k)}} \left[ \frac{m_{[i-k]}}{n^{(i-k)}} - \frac{(m-k)_{[i-k+1]}}{(n+k)^{(i-k+1)}} \right]. \end{aligned}$$

It can be shown that:

$$\frac{m_{[i-k]}}{n^{(i-k)}} - \frac{(m-k)_{[i-k+1]}}{(n+k)^{(i-k+1)}} \geq 0$$

under  $S \geq 2m - 1$  (this inequality implies  $n \geq m$ ) to conclude that  $c_i \geq 0$ .  $\square$

An iteration of  $f$ , conditioned on an initial guess  $x_0$ , is defined by iterates  $x_k = f(x_{k-1})$  for  $k > 0$ . The iteration is said to converge if  $x_k \rightarrow x$  for some  $x$ . The following is an application of the above lemma and the Banach fixed point theorem:

**Theorem 1.** *Under Assumption 2, an iteration of  $f$  conditioned on an initial guess  $x_0 \geq 0$  converges to the unique fixed point of  $f$  in  $(0, 1]$ .*

*Remark 1.* The stronger inequality  $S/m \geq 2$  implies the weaker one appearing in Assumption 2. We recognize  $S/m$  as the *member to vehicle ratio* in the context of carshares. This result means the algorithm converges for all practical cases—as discussed in §5.3.2 and shown in §5.4.2, the *member to vehicle ratio*  $S/m$  is 30-50 in practice. A ratio of less than two means that more than one car is needed for every two members, defeating the purpose of a carshare.

The results in this section (in particular, Lemma 1) do not hold when  $v$  is negative. We end our analysis with the following practical result:

**Corollary 1.** *The operator can increase  $K$ , an input sizing parameter, until the conditions of Lemma 1 are satisfied, guaranteeing a fixed point.*

*Proof.*  $v$  is increasing with  $K^1$  as shown in Eq(C.2).  $\square$

---

<sup>1</sup>Note that  $n_B$  is an increasing function of  $K$  as show in Eq(5.10). As the size of the busy period  $K$  increases, the number of subscribers expected to arrive during the busy period ( $n_B$ ) also increases. Thus, increasing  $K$  may not immediately increase  $v$  as needed. However,  $n_B$  is limited by  $S$  and  $K$  is unbounded, so  $K$  can always be increased such that  $v$  increases sufficiently.

# Appendix D

## WeBike Field Trial Participation Survey

**Credit:** *this survey was co-authored with Tobias Schroeder and S. Keshav.*

# E-Bike Survey University of Waterloo

## 2. How would you characterize your knowledge of electric bicycles?

Know nothing about it	Have heard about it	Have read some articles about it	I am quite knowledgeable	I have expert knowledge
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## 3. Which of the following experiences have you had personally with e-bikes?

- None.
- I used one once.
- I have occasionally used one.
- I regularly use one.
- I own one and use it occasionally.
- I own one and use it regularly.

Other (please specify)

## 4. Please let us know to what extent the following statements apply to you:

	Does not apply to me	Does mostly not apply to me	Somewhat applies to me	Mostly applies to me	Fully applies to me
I am interested in e-bikes.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I plan to buy an e-bike within the next couple of years.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I can imagine substituting my current bicycle with an e-bike.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I can imagine substituting my car with an e-bike.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Optional comment

## Thoughts About E-Bikes

## E-Bike Survey University of Waterloo

**5. In the blanks below, please write up to twenty different statements that come to mind when you think of electric bicycles. There are no right or wrong answers. Just write down your spontaneous thoughts about e-bikes, in the order that they occur to you.**

e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>
e-bikes are	<input type="text"/>

## Transportation Behaviour

We would like to learn more about your transportation-related attitudes and behaviours. Please answer the following questions.

### 6. How often do you use the following means of transportation?

	(almost) daily	1-3 days per week	1-3 days per month	less than once per month	(almost) never
Car	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycle	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric bicycle	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public transport	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

optional comment

## E-Bike Survey University of Waterloo

**7. On average, how many kilometers do you spend per day on trips related to the following activities? Please enter the appropriate number of kilometers in each line.**

Work	<input type="text"/>
School/education	<input type="text"/>
Shopping	<input type="text"/>
Private Business (e.g. visits to doctors/other offices)	<input type="text"/>
Get or take other persons to places	<input type="text"/>
Leisure activities	<input type="text"/>
Other	<input type="text"/>

### Transportation Decisions 1

**8. When people choose a means of transportation, they want to satisfy different needs. For example, some people might choose the car to enjoy independence (ability to travel spontaneously), comfort, and safety, while others might prefer the bicycle to save costs, be healthy, and preserve the environment. How important are the following criteria for you when you choose a means of transportation?**

	very unimportant	quite unimportant	somewhat unimportant	somewhat important	quite important	very important
Independence	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Stress-free travel	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Costs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
My social status	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Fun	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Environmental friendliness	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Reliability	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Comfort	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Safety	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Healthiness	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

### Transportation Decisions 2

Next, we would like to know to what extent different means of transportation are compatible with your needs. Please tell us how much you agree with the following statements.



# E-Bike Survey University of Waterloo

## 9. Independence: In my opinion, the following modes of transport allow spontaneous travel whenever needed.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## 10. Stress-free travel: In my opinion, the following modes of transport allow stress-free travel.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## 11. Costs: In my opinion, the following modes of transport are expensive.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## 12. Social status: The following modes of transport fit with my social status.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## 13. Fun: In my opinion, the following modes of transport are fun to use.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public Transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

# E-Bike Survey University of Waterloo

## Transportation Decisions 3

We would like to know to what extent different means of transportation are compatible with your needs. Please tell us how much you agree with the following statements.

### 14. Environmental friendliness: In my opinion, the following modes of transport are eco-friendly.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public Transport	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

### 15. Reliability: In my opinion, the following modes of transport are reliable.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public Transport	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

### 16. Comfort: In my opinion, the following modes of transport are comfortable.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public Transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

### 17. Safety: In my opinion, the following modes of transportation are safe.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public Transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

# E-Bike Survey University of Waterloo

## 18. Health: In my opinion, using the following modes of transport is good for my health.

	Strongly disagree	Disagree	Neutral	Agree	Strongly Agree
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public Transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## Emotions

Now, we would like to learn more about your emotional reactions to certain transportation options. This part of the study will help us understand people's transport decisions better.

Research has shown that emotions have three different components:

- (1) How good or nice versus bad or awful are things?
- (2) How weak and powerless versus strong and powerful are things?
- (3) How calm and quiet versus arousing and active are things?

We ask you to rate the following concepts related to transport decisions. Please answer as quickly as possible. There are no right or wrong answers. We are simply interested in your intuitions.

## 19. Please rate the badness versus goodness of the following modes of transport:

	extremely bad/awful -4	-3	-2	-1	neutral	+1	+2	+3	extremely good/nice +4
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public Transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## 20. Please rate the weakness versus strength of the following modes of transport:

	extremely powerless -4	-3	-2	-1	neutral	+1	+2	+3	extremely powerful +4
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public Transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

# E-Bike Survey University of Waterloo

## 21. Please rate the passivity versus activity of the following modes of transport:

	extremely calm -4	-3	-2	-1	neutral	+1	+2	+3	extremely arousing +4
Cars	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Electric Bicycles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Public Transit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Walking	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## Emotions - Needs related to transport

Please rate the following concepts related to transport decisions. Please answer as quickly as possible. There are no right or wrong answers. We are simply interested in your intuitions.

## 22. Please rate the badness versus goodness of the following transport-related needs:

	extremely bad/awful -4	-3	-2	-1	neutral	+1	+2	+3	extremely good/nice +4
Independence	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Stress-free travel	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Costs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
My social status	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Fun	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Eco-friendliness	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Reliability	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Comfort	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Safety	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Health	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## 23. Please rate the weakness versus strength of the following transport-related needs:

	extremely powerless -4	-3	-2	-1	neutral	+1	+2	+3	extremely powerful +4
Independence	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Stress-free travel	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Costs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
My social status	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Fun	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Eco-friendliness	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Reliability	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Comfort	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Safety	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Health	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## E-Bike Survey University of Waterloo

### 24. Please rate the passivity versus activity of the following transport-related needs:

	extremely calm -4	-3	-2	-1	neutral	+1	+2	+3	extremely arousing +4
Independence	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Stress-free travel	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Costs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
My social status	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Fun	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Eco-friendliness	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Reliability	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Comfort	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Safety	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Health	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## Emotions - Self

You have almost made it through the survey. On the next page, we will explain the details about the e-bike field study.

Prior to that, we ask you to use the three emotion scales you have encountered previously to rate your feelings about yourself. Your answer will help us understand the role of intuitions in people's transport decisions. Again, answer as quickly as possible. There is no right or wrong answer.

(you may choose not to answer these questions if too personal)

### 25. Please rate your own badness versus goodness:

	extremely bad/awful - 4	-3	-2	-1	neutral	+1	+2	+3	extremely good/nice +4	prefer not to answer
Myself as I really am	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

### 26. Please rate your own weakness versus strength:

	extremely powerless - 4	-3	-2	-1	neutral	+1	+2	+3	extremely powerful +4	prefer not to answer
Myself as I really am	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

### 27. Please rate your own passivity versus activity:

	extremely calm -4	-3	-2	-1	neutral	+1	+2	+3	extremely arousing +4	prefer not to answer
Myself as I really am	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

## E-Bike Questions

## E-Bike Survey University of Waterloo

**28. Now let's imagine you had an electric bicycle. For example, you might have agreed to take part in our three-year study and we provided you with a free e-bike. How often would you use it?**

	(almost) daily	1-3 days per week	1-3 days per month	less than once per month	(almost) never	question does not apply - I own an e-bike already
during the summer months (May - October)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
during the winter months (November - April)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

**29. Please estimate the number of kilometers you would likely ride the e-bike per day, on average. Write the number of kilometers in the following lines.**

during the summer months (May - October)

during the winter months (November - April)

**30. If you had an electric bicycle, how often would you use it to replace your current favoured means of transportation for trips related to the following activities? For example, if you now use the car to get to work, how often would you use the e-bike instead?**

	would (almost) never use e-bike	would rarely use e-bike	would sometimes use e-bike	would often use e-bike	would (almost) always use e-bike	I already use an e-bike most of the time
Work	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
School/education	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Shopping	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Private Business (e.g., visits to doctors or other offices)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Get or take other persons to places	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Leisure activities	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Other	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

**31. What questions, if any, do you have about electric bicycles?**

## Participation in the Study

## E-Bike Survey University of Waterloo

As mentioned before, we plan to invite selected participants of this survey to take part in a three-year study of electric bicycle use. We will select participants to represent four distinct populations of e-bike users: (i) Car users who would consider using e-bikes instead of cars for some trips. (ii) Bike users who would consider using e-bikes instead of cars for some trips (iii) Public-transit users who would consider using e-bikes instead of cars for some trips. (iv) Walkers who would consider using e-bikes instead of cars for some trips. We are not targeting any specific gender, age-range, or special characteristics: our selection of participants aims to be a representative sample of the UW population. However, all participants must: a) be at least 18 years old b) wear a helmet when operating the e-bike c) anticipate residing in the KW region and retain a connection with the University of Waterloo for the next three years and d) undergo a training session on safe e-bike use.

If you are chosen, you will be offered a new, free e-bike, which would be yours to keep after the study is over. In return, you would agree to provide us with certain data during the time of the study. Most of it would be collected automatically with sensors (e.g., information such as speed, location, and battery performance) and transmitted to us automatically whenever your e-bike is in the range of a wireless internet connection. In addition, we would ask you every 3-6 months to fill out an online survey like the present one. The goal of this project is both to optimize e-bike technology and to understand more about the decision processes of humans who adopt new technologies.

### 32. Are you interested in participating in our three-year e-bike study?

- Yes, let me know the details about the study.
- I am not interested in operating an e-bike, but I would be willing to take part in future surveys.
- No, thank you.

### 33. If you are interested or possibly interested in the e-bike study or future surveys on this topic, please provide your email address, so we can get in touch with you about the details. (Please note: you can still decide later not to take part if you change your mind or are unhappy about the terms and conditions that we will offer you)

### 34. If you are interested in our study, please provide up to five reasons for why you would (might) like to take part in it. If you are not interested, please provide up to five reasons for why not:

1	<input type="text"/>
2	<input type="text"/>
3	<input type="text"/>
4	<input type="text"/>
5	<input type="text"/>

## Personal Information

### 35. What is your gender?

- male
- female
- other
- prefer not to say

## E-Bike Survey University of Waterloo

**\*36. What is your age? (please write your age in years in the line below)**

**\*37. What is the nature of your affiliation with the University of Waterloo?**

- Faculty
- Staff
- Undergraduate Student
- Graduate Student
- Postdoc
- Other
- I am not affiliated with the University of Waterloo

**\*38. Do you anticipate living in the K-W area until June 30, 2017?**


- Yes
- No



# Appendix E

## WeBike To EV Mapping

The following diagram or “mind map” shows some of the questions relevant to EV researchers that we hope to answer using the WeBike data, and what data/sensors we use to compute them.



If all charging stations occupied, and more vehicles are parked at location with a SOC < ROC, at that location they consider them as queued

If plugged in, full battery, and there are others who could be charging, charge range. Need driver ID and SOC because if plugged in and ROC == 0, could indicate a charging problem instead of charge range

this cannot be strictly inferred from ROC, because ROC could be 0 but person is full and plugged in or a charging prob. etc.

won't have this for standard outlets, only networked charging stations