

Driving with Data: Modeling and Forecasting Vehicle Fleet Maintenance in Detroit

Josh Gardner
University of Michigan
jpgard@umich.edu

Danai Koutra
University of Michigan
dkoutra@umich.edu

Jawad Mroueh
University of Michigan
jmroueh@umich.edu

Victor Pang
University of Michigan
vicpang@umich.edu

Arya Farahi
University of Michigan
aryaf@umich.edu

Sam Krassenstein
City of Detroit
Operations and Infrastructure Group
samkrass@umich.edu

Jared Webb
University of Michigan
jaredaw@umich.edu

ABSTRACT

The City of Detroit maintains an active fleet of over 2500 vehicles, spending an annual average of over \$5 million on new vehicle purchases and over \$7.7 million on maintaining this fleet. Understanding the existence of patterns and trends in this data could be useful to a variety of stakeholders, particularly as Detroit emerges from Chapter 9 bankruptcy, but the patterns in such data are often complex and multivariate and the city lacks dedicated resources for detailed analysis of this data. This work, a data collaboration between the Michigan Data Science Team¹ and the City of Detroit's Operations and Infrastructure Group, seeks to address this unmet need by analyzing data from the City of Detroit's entire vehicle fleet from 2010-2017. We utilize tensor decomposition techniques to discover and visualize unique temporal patterns in vehicle maintenance; apply differential sequence mining to demonstrate the existence of common and statistically unique maintenance sequences by vehicle make and model; and, after showing these time-dependencies in the dataset, demonstrate an application of a predictive Long Short Term Memory (LSTM) neural network model to predict maintenance sequences. Our analysis shows both the complexities of municipal vehicle fleet data and useful techniques for mining and modeling such data.

1. INTRODUCTION

The City of Detroit, like many city governments, manages and maintains a large vehicle fleet consisting of over 2500 active vehicles. These vehicles support a diverse array of government functions, including service delivery, law enforcement, and grounds maintenance. In addition to being critical to Detroit's ability to effectively serve its citizens, the maintenance required to sustain this fleet is both complex and expensive: The city spent an annual average of \$7.7 mil-

¹<http://midas.umich.edu/mdst>



Figure 1: Vehicle fleet in the City of Detroit.

lion on maintenance and over \$5 million on new vehicle purchases between 2010 and 2017. As of 2015, the city had four shops, six fuel sites, and 70 technicians to maintain the fleet; this represented a significant reduction of the city's fleet after undergoing Chapter 9 bankruptcy filing in July 2013 [28]. Most cities lack the resources and expertise to dedicate to understanding and optimizing fleet operations, and even if those resources existed, analyzing the complex patterns of vehicle use and maintenance is a challenging task. Having a more nuanced understanding of the patterns in their fleet maintenance data would allow the City of Detroit to make intelligent decisions for efficiency and cost-reduction at a critical time in the city's history. In particular, the city needs to balance concerns of cost and resource efficiency with maximizing vehicle uptime and lifetime, ensuring consistent service delivery, and reducing the city's carbon footprint.

This project, a data collaboration between the Michigan Data Science Team (MDST)—a student organization at the University of Michigan—and the City of Detroit's Operations and Infrastructure Group, a municipal entity, is an initial foray into understanding and modeling municipal vehicle maintenance data. The analysis we present constitutes an initial step toward meeting the complex needs of the city (and citizens) of Detroit.

The goals of this paper are two-fold: First, we aim to show that Detroit's fleet maintenance data contains discoverable structure, and to demonstrate methods for revealing this structure. Second, we seek to apply methods for modeling this structure to make predictions relevant to municipal decision-making and resource allocation, namely, forecasting vehicle maintenance. These predictions could reduce costs,

Table 1: Description of the *vehicles* table.

Field	Description	Example
Unit#	Unique Vehicle Identifier	026603
Dept#	Code of department vehicle is assigned to	37
Dept Desc	Description of department	POLICE
Make	Vehicle make	CHEVROLET
Model	Vehicle model	2500
Year	Model year of vehicle	2002
Last Meter	Odometer reading at last check (mi)	52738
Last Fuel Date	Most recent refuel at city refueling station	2009-11-05 15:37:25
Purchase Cost	Purchase cost, in US dollars	\$20,456
Status Code	A = Active; S = Disposed	A
Status Desc	Description of status	Active Unit
LTD Maintenance Cost	Total maintenance cost to date, in US dollars	\$5,951.04
LTD Fuel Cost	Total fuel cost to date, in US dollars	\$9,295.01
LTD Fuel Gallons	Total fuel consumption to date	\$3,646.6

fraud, and erroneous data; lead to better scheduling; and form the basis for future internal tools in the City of Detroit. In the analysis that follows, we pursue those aims by (a) exploring multidimensional patterns in vehicle maintenance using the parallel factors (PARAFAC) decomposition to reveal patterns in maintenance of automotive systems in different vehicle types over time using data *tensors*, providing a visual approach to representing patterns in fleet maintenance both over time; (b) applying a sequence mining technique to statistically identify frequent maintenance sequences by make/model, and (c) leveraging a modern neural network approach to predict vehicle maintenance in the City of Detroit’s fleet.

The structure of this paper is as follows: We first survey prior research on applied tensor analysis and on municipal vehicle fleets. In Section 3, we describe the Detroit dataset in detail. Section 4 presents the results of the PARAFAC evaluation and in Section 5 we conduct differential pattern mining to demonstrate the presence of statistically unique maintenance patterns by make/model and construct a predictive model utilizing a long short-term memory (LSTM) neural network to model these sequences. In Section 6, we give conclusions, challenges of data collaboration and analysis in real-world public-sector contexts. We conclude with suggestions for future work in Section 7.

2. RELATED WORK

Our analysis is based on tensor decompositions and related to other studies on municipal vehicle fleets.

2.1 Tensor Analysis and Applications

Tensor representations and various tensor decompositions have found wide applications in a variety of domains, including psychometrics [8] and brain imaging [20] (where many core techniques, such as the PARAFAC decomposition used here, were developed), the evolution of chatroom [2] and email [5] conversations over time, modeling web search [24], epidemiology [1], and anomaly detection [13]. Tensor representation is useful in a variety of problem domains because it allows for multi-way analysis of data containing multidimensional patterns.

For a more detailed overview of tensor decompositions, their mechanics, and their applications, we refer the interested reader to [12]. We describe the decompositions that are relevant to our data analysis in Section 3.2.

2.2 Municipal Vehicle Fleets and Predictive Maintenance Models

While predictive analytics, data science, and the application of such techniques to urban planning (sometimes called *urban informatics*) have dramatically expanded in recent years, these techniques have seen only limited applications to one of the largest and most substantial assets managed by many governments—their vehicles—and published research on the topic is limited. Some state and local governments conduct, but rarely publish, fleet lifecycle reports and maintenance analyses [9]. [15] reports on fleet replacement management by state Departments of Transportation across all states. [21] presents a case study of municipal fleet management in a mid-sized American city mostly focused on cost-reduction analysis.

Recent research on predictive maintenance has utilized on-board vehicle data to predict maintenance [22] the use of vehicle speed data to evaluate winter maintenance operations [16]. Other vehicle-related issues in urban areas have received significant research attention, including accident prediction [17, 18] and traffic flow prediction and optimization [25, 27, 19].

3. DATASET

In this section, we first describe the raw dataset obtained from the City of Detroit and used in the following analysis. We also describe the transformation of the raw vehicle and maintenance data into the data tensor to which we apply the tensor modeling techniques described in Section 4.

3.1 Detroit Vehicles Dataset

MDST partnered with the City of Detroit’s Operations and Infrastructure Group to obtain a comprehensive dataset from the City of Detroit. This dataset consists of two tables.

The **vehicles table** consists of 6,725 records, one per vehicle, representing every known vehicle currently or previously owned by the City of Detroit. 2,566 of these vehicles are currently active in the fleet, but the oldest vehicle purchases date to 1944. The table includes information about each vehicle’s manufacture, purchase, and use. The table includes police cars, garbage trucks, freight trucks, ambulances, boats, motorcycles, mowers, and other vehicles. Table 1 gives a description of the fields, with a sample entry.

The **maintenance table** consists of job-level records for

all maintenance performed on any vehicles owned by the City of Detroit. This table includes 229,540 records representing individual jobs, which include everything from routine inspections, tire changes, and preventive maintenance to major collision repairs, glass work, upgrades, and engine replacements. The maintenance data is described in Table 2.

Together, these tables form a dataset representing detailed job-level information about maintenance on Detroit’s entire vehicle fleet across 87 different departments, including police, airport, fire, solid waste, and grounds maintenance. There is no missing data, but there are potential concerns about the accuracy of some data due to data-entry errors or human coding of job types and descriptions. Because the City of Detroit’s fleet, and its data collection practices, have changed substantially over time, we limit our analysis to maintenance data from vehicles purchased in 2010 or later in order to utilize only the most reliable and relevant vehicle and maintenance patterns. This represents 1,087 vehicles and over 25,000 individual maintenance records.

3.2 Data Representation as a Tensor

In this section, we describe the process of representing the Detroit vehicle maintenance dataset as a series of $vehicle \times system \times time$ data tensors, providing a brief introduction to tensors and describing both the process and motivation for this approach.

A tensor is a multidimensional or N-way array [12]. Tensors provide a way of representing, analyzing, and modeling complex, multidimensional data. While tensors of arbitrary numbers of dimensions, or modes, can be evaluated using the techniques described here, the current analysis uses only 3-mode tensors which are, fortunately, straightforward to visualize and discuss.

In this analysis, we were interested in understanding how vehicle maintenance unfolded over time, and whether there were patterns and structure in how different types of vehicles were maintained. This task is of interest to our partners in Detroit in order to understand fleet maintenance, but also has the potential to inform future work on predicting vehicle maintenance, breakdowns, availability, and direct and indirect maintenance costs.

In order to represent the raw data – which consisted of a *vehicles* table and a *maintenance* table – as a data tensor, we needed to aggregate the data by vehicle, job type, and time. We assembled counts of maintenance jobs, by vehicle, system², and month/year. This produced a 3-way tensor similar to the one shown in Figure 2, where the vertical axis (the *first mode*) represents each different vehicle, sorted by year and unit number; the horizontal axis (the *second mode*) represents each distinct vehicle system that occurred for at least one job in the dataset; and the depth (*third mode*) represents time in months or years. The value at any given [vehicle, system, time] point in the tensor is the count of jobs for that particular vehicle, system, and month. An example of the data tensor we construct is shown in Figure 2.

²*System* describes the vehicle component repaired in a job, such as brakes, lighting system, or suspension; see “System Description” in Table 2.




Figure 2: Depiction of (a) 3-mode data tensor; (b) the same tensor as a stacked series of frontal slices, or arrays; (c) an example single frontal slice of a vehicle data tensor used in this analysis (each entry corresponds to the count of a specific job type for a vehicle at a fixed time).

This 3-dimensional tensor representation allows us to model the relationships across these three dimensions in the data – and in particular, to see how patterns evolve over time. This representation is critical to answering our initial question as to whether patterns exist in maintenance over time, which would lead directly to insights about maintenance trends in Detroit’s vehicle fleet, inform approaches to modeling and prediction of vehicle maintenance, and potentially lead to changes in the city’s fleet maintenance operations.

4. UNDERSTANDING COMPLEX MAINTENANCE PATTERNS OVER TIME WITH PARAFAC DECOMPOSITION

In this section, we describe a technique for extracting insights about structure and patterns in a data tensor known as the PARAFAC (PARAllel FACTors) decomposition, and describe insights gained from applying this technique to tensors of the vehicle fleet over absolute time, as well as over vehicles’ lifetimes.

4.1 The PARAFAC Decomposition

Tensors can be thought of as higher-dimensional versions of the “flat” two-dimensional arrays common in data analysis tasks. As such, many techniques have been developed for manipulating and understanding tensors that are analogous to methods for two-dimensional data. The PARAFAC decomposition is an example of such a technique. PARAFAC decomposes a tensor into a sum of component rank-one tensors which best reconstruct the original tensor. Given a third-order tensor $\mathcal{X} \in \mathbb{R}^{I \times J \times K}$, PARAFAC decomposes the tensor as:

$$\mathcal{X} \approx \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r \quad (1)$$

where R is a positive integer, $\mathbf{a}_r \in \mathbb{R}^I$, $\mathbf{b}_r \in \mathbb{R}^J$, $\mathbf{c}_r \in \mathbb{R}^K$ for $r = 1, \dots, R$ and “ \circ ” represents the vector outer product. Thus, PARAFAC represents each element of \mathcal{X} as the vector outer product of the respective fibers of \mathcal{X} :

$$x_{ijk} = \sum_{r=1}^R a_{ir} b_{jr} c_{kr} \text{ for } i = 1, \dots, I, j = 1, \dots, J, k = 1, \dots, K. \quad (2)$$

The PARAFAC decomposition can be written compactly as the combination of three loading matrices \mathbf{A} , \mathbf{B} , \mathbf{C} :

$$\mathcal{X} = [\mathbf{A}^{I \times R}, \mathbf{B}^{J \times R}, \mathbf{C}^{K \times R}], \quad (3)$$

Table 2: Description of the *maintenance* table.

Field	Description	Example
Job ID	Unique identifier for job	847956
Year WO Completed	Year of completion	2017
Unit No	Vehicle identifier	067602
Work Order No	Unique identifier for work order	635864
WO Open Date	Work Order Open	2017-01-17
WO Completed Date	Work Order Completion	2017-01-17
Work Order Location	Location of work order	CODRF
Job Open Date	Job Open	2017-01-17
Job Reason	Job reason code	B
Job Reason Desc	Job reason description	BREAKDOWN / REPAIR
Job Open Date2	Job Open 2	2017-01-17
Job Completed Date	Job Completed	2017-01-17
Job Code	Job ID	24-13-000
Job Description	Detailed description of job	REPAIR Brakes
Labor Hours	Hours of labor completed on job	6.35
Actual Labor Cost	Cost of labor for job	\$348.16
Commercial Cost	Cost of commercial (non-city) labor	\$0
Part Cost	Cost of parts for job	\$57.55
Primary Meter	Odometer reading at time of repair (mi)	48250
Job Status	Status code; DON = Done	DON
Job WAC	Job type code	24
WACDescription	Job type description	REPAIR
Job System	Code for vehicle system repaired by job	13
System Description	Description for vehicle system repaired by job	Brakes
Job Location	Location where job was completed	CODRF




Figure 3: A visual depiction of the PARAFAC decomposition: It decomposes the $vehicle \times system \times time$ data tensor into vector products of the respective vehicle, system, and time dimensions. From this decomposition, we can obtain factor matrices, which reveal factors that are visualized to display relationships between these three variables in the Detroit vehicle-maintenance dataset. Figure based on [12].

in which the r^{th} columns correspond to the vectors \mathbf{a}_r , \mathbf{b}_r , and \mathbf{c}_r , respectively. For more information about the PARAFAC decomposition, we refer the reader to [12]. The key aspect of the PARAFAC decomposition that makes it useful for understanding the Detroit vehicle-maintenance dataset is that it yields R sets (components or factors) of \mathbf{a} , \mathbf{b} , and \mathbf{c} vectors which best reconstruct the original data tensor. These R factors can be thought of as containing the most “important” relationships between different fibers of the data tensor across all three dimensions (or modes).

4.2 Insights From PARAFAC Application to Detroit Vehicle-Maintenance Dataset

This analysis was intended to answer the question of whether vehicle-system-time relationships exist in the Detroit dataset. To this end, we generate so-called “three-way” plots of the three factor matrices from the PARAFAC decomposition [13], using the MATLAB tensor toolkit provided by [4, 6] for PARAFAC. Each plot visualizes the vectors \mathbf{a}_r , \mathbf{b}_r and \mathbf{c}_r , which correspond to the r^{th} factor and represent the dif-

ferent modes participating in the factor (i.e., vehicle, system description and time, respectively). We explored two different representations of time for the third mode of the data tensors: one which used *absolute time* (month and year, January 2010 - present, measured by the start date of a maintenance job) and another using *vehicle lifetime* (year, starting with year 0 as the vehicle’s purchase year). The absolute time analysis allows us to model seasonality and other real-time trends in fleet maintenance, and could be more useful in forecasting future maintenance, while the vehicle lifetime analysis allows us to measure trends and changes in vehicles’ maintenance over the course of their lifetime in the Detroit fleet, and could be useful for vehicle make/model reliability analyses. We examine $R = 25$ components for each analysis.

Examples of the results from the absolute time analysis are shown in Figures 4 - 6. These results demonstrate clear patterns across vehicles, systems under repair, and time, underscoring the importance of taking a multivariate approach. For example, fire trucks and ambulances (the Terrastar Horton in Figure 4 and Smeal SST Pumper in Figure 5, respectively) both show strong evidence of patterns in their maintenance, but with very different groups of systems and across different time bands. The riding mower owned by the GSD - Grounds Maintenance Department shown in Figure 6, however, displays an entirely different maintenance pattern, with a focus on only two systems (mowing blades and tires/tubes/liners) and strong seasonality which reflects the seasonal use of mowers in a northern city such as Detroit.

Examples of the results from the PARAFAC vehicle lifetime analysis are shown in Figures 7 - 9. This analysis demonstrates a different set of patterns – this time, across the lifetime of vehicles, beginning when they are purchased.

The PARAFAC analysis via three-way factor plots demonstrates the variety of insights that can be gained from using tensor decomposition to understand complex multidimensional data.

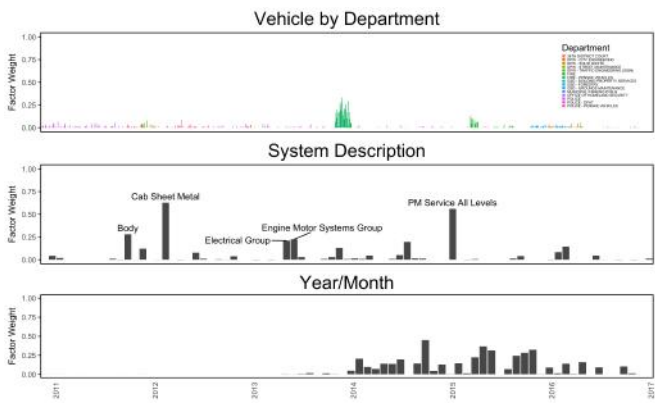


Figure 4: PARAFAC 3-way plot of absolute-time analysis. High factor weights in the top panel are for 2014 Terrastar Horton vehicles, an ambulance. The bottom two panels show systems (Body, Cab/Sheet Metal, Engine and Motor, and Preventive Maintenance Service) and time frames where this maintenance most often occurs.

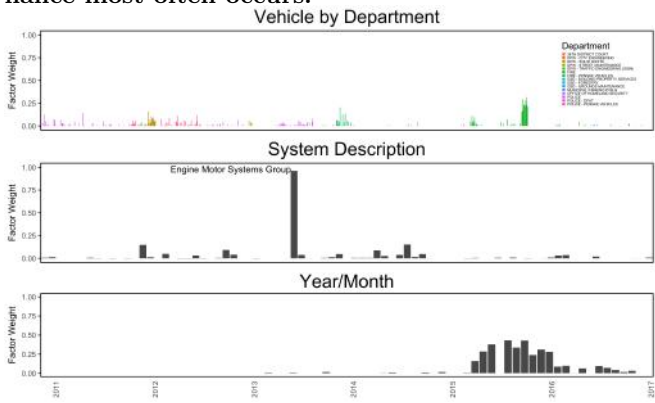


Figure 5: PARAFAC 3-way plot of absolute-time analysis. This plot demonstrates strong and specific maintenance patterns for the 2015 Smeal SST Pumper fire truck. It shows extensive and specific repair to the engine systems with little other maintenance, from late 2015 through 2016.




Figure 6: PARAFAC 3-way plot of absolute-time analysis. This plot shows factors related to the 2014 Hustler X-One, a Grounds Maintenance Department riding mower. There are strong system and time patterns to X-One maintenance, with repairs almost exclusively to mower blades, with repairs to tires/tubes/liners/valves during seasons of high usage.

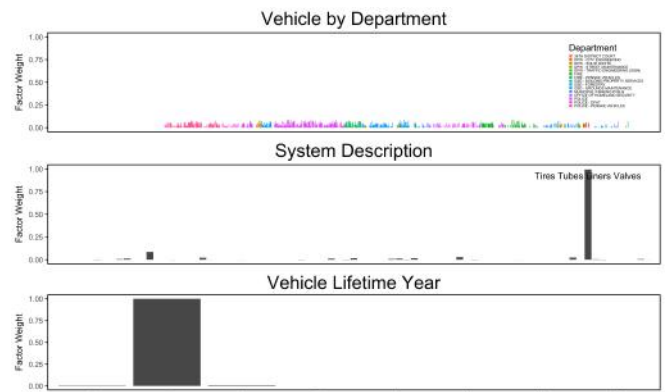


Figure 7: PARAFAC 3-way plot of vehicle lifetime analysis revealing a simple pattern common to almost all vehicles, as demonstrated by the consistent loading across the vehicle factor (top panel): having tires/tubes/valves/liners replaced during the second year of their lifetime, with few repairs to this system either before or after.

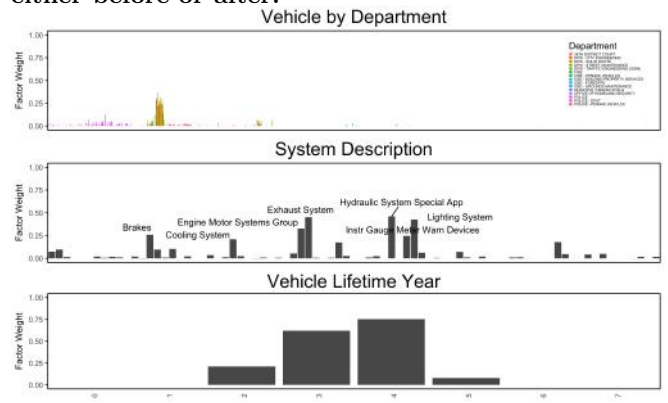


Figure 8: PARAFAC 3-way plot of vehicle lifetime analysis showing the 2012 Freightliner M2112V, a Department of Solid Waste garbage truck. This plot reveals a strong pattern of increased maintenance in years 2-4 after purchase, focusing on a variety of technical systems: hydraulics, lighting, gauges and warning devices, and cooling systems.




Figure 9: PARAFAC 3-way plot of vehicle lifetime analysis. This plot demonstrates maintenance patterns for the 2013 Hustler Z 60 2013 (which is a riding mower), which have their moving blades serviced frequently in the second and third years of their lifetime.



Figure 10: 2015 Smeal SST Pumper in Detroit.

mensional data. The analysis shown above reveals common trends across the entire Detroit vehicle fleet, as well as unique trends specific to certain vehicles, systems, and times. Additionally, the use of two different measures of time—month/year, and vehicle lifetime—allows us to demonstrate two different modes of time-bound pattern in the data. Such an approach demonstrates that there are unique patterns in the Detroit vehicle-maintenance dataset by vehicle, system, and time, suggesting that analysis and modeling approaches which can capture these patterns are likely to be effective.

5. MINING FREQUENT MAINTENANCE PATTERNS

In Section 4, our results demonstrate discoverable structure in the Detroit vehicle-maintenance data, particularly by vehicle make/model. In this section, we expand on these results, statistically verifying the existence of unique patterns in the *sequences* of systems repaired by vehicle make and model and applying a sequence modeling approach to build a predictive maintenance model.

5.1 Sequence Mining By Vehicle Make/Model

Sequential pattern mining is a constellation of techniques used to identify and evaluate sequences of events [3]. *Differential* sequence mining compares differences in sequences between two groups, statistically identifying different pathways unique to each group. We apply a methodology adapted from [11] to a subset of vehicles identified as having potentially unique maintenance patterns in the tensor decomposition analysis above, both in order to statistically verify these unique patterns, and to determine whether we can ignore time (and simply focus on order) in modeling maintenance sequences. Specifically, the method consists of three steps:

- **Step 1:** Find the top n most frequent sequences (restricting these to sequences of length 3 or longer) for a given make/model and normalize by the total number of maintenance sequences of the same length for that make/model (this generates the *left support* and the *left normalized support*), using a general algorithm to find the most frequent sequences [26].
- **Step 2:** Calculate the same ratio for all other make/models as a separate group (the *right support* and *right normalized support*).

- **Step 3:** Compare these two normalized frequencies by (a) calculating the left:right ratio, the *i-ratio* [11], and (b) by conducting a test of the difference between two population proportions to test the null hypothesis $H_0 : p_1 = p_2$ against $H_a : p_1 \neq p_2$ where p_1 and p_2 are the left and right normalized supports, respectively.³

The result of this analysis is shown in Table 3. These results demonstrate strong and statistically significant distinctions in maintenance sequences by vehicle type, suggesting that common maintenance sequences for these vehicles are unique to their make/model and statistically uncommon across the rest of the fleet. All sequences evaluated for the Dodge Charger, Ford Crown Victoria, and Smeal SST Pumper exceed any reasonable significance threshold, with only the Hustler X-One demonstrating less significant results for three of seven sequences tested (due to similarity with other models of Hustler mowers)⁴. Furthermore, because the sequential pattern mining approach simply looks at the *order* of the maintenance jobs, but not their actual timing, these results demonstrate that there is strong correlation between maintenance sequence and make/model, even when actual timing is ignored, and that sequential models—even those which ignore the time between maintenance events—may be effective in modeling vehicle maintenance. This analysis informs the approach adopted in the following section.

5.2 Predictive Model for Maintenance Sequences

In this section, we build off of the findings of our prior analysis to construct a predictive model to predict the next maintenance job, given a vehicles' previous jobs, by learning from the maintenance histories of similar make/models.

Having demonstrated strong sequential patterns in maintenance for vehicle make and models, we developed an exploratory model to predict vehicle maintenance – one of the potential applications of our data collaboration identified by Detroit's Operations and Infrastructure Group. From the raw data, we assemble a dataset consisting of the complete sequence of system repair jobs for each vehicle. Each vehicle's sequence is considered a separate observation. We train a probabilistic model which assigns probabilities to various repair sequences using a Long Short-Term Memory (LSTM) neural network [10]. We specifically implement the architecture used in [29] for predicting words in sentences because of its ability to model complex sequences while avoiding overfitting.

An LSTM model reads over a sequence, one item at a time, and computes probabilities of the possible values for the next item. In theory, an LSTM is capable of learning arbitrarily long-distance dependencies across a sequence; in our implementation, the LSTM considers a step size of 20 items. This

³Note that in the original implementation, a *t*-test was used; here, we use a difference-in-proportions *z*-test because our analysis tests whether the *normalized* support differ, not the raw counts [11].

⁴Note that we report *p*-values unadjusted for multiple comparisons, as this test is a heuristic to search for differences in patterns as in [11] and not a strict statistical test; however, even conservative *p*-value adjustments would still yield highly significant results in most cases.

Table 3: Results of differential sequence mining for selected make/models from Section 4. The table shows the top 5 most frequent maintenance sequences of length ≥ 3 for each make/model (‘Left’), which are compared to the frequency of that maintenance sequence across the rest of the fleet (‘Right’). Large i-ratios and z -statistics indicate sequences unique to a given make/model. Abbreviations: TTLV = Tires, Tubes, Liners & Valves; PM = PM Service All Levels; EX = Exhaust; MOW = MOW; ENG/MS = Engine / Motor Systems Group; PUMP = Pump - Product Transfer; CSM = Cab & Sheet Metal.

Vehicle	Sequence	Left Support	Left Norm Support	Right Support	Right Norm Support	i-Ratio	z	P(z)
Dodge Charger	(PM, TTLV, PM)	187	0.0377	126	0.0067	5.6	-10.4	< 0.0001
	(PM, PM, TTLV)	186	0.0375	81	0.0043	8.67	-9.9	< 0.0001
	(PM, PM, PM)	185	0.0373	97	0.0052	7.2	-10.3	< 0.0001
	(TTLV, PM, PM)	185	0.0373	82	0.0044	8.51	-10.1	< 0.0001
	(PM, TTLV, TTLV)	183	0.0369	158	0.0085	4.37	-11.3	< 0.0001
	(TTLV, TTLV, PM)	183	0.0369	168	0.009	4.11	-11.4	< 0.0001
	(TTLV, PM, TTLV)	182	0.0367	180	0.0096	3.82	-11.7	< 0.0001
	(PM, TTLV, PM, TTLV)	180	0.0378	40	0.0022	17.03	-9.0	< 0.0001
Ford Crown Victoria	(TTLV, PM, TTLV)	101	0.0247	365	0.0187	1.32	-18.4	< 0.0001
	(PM, TTLV, TTLV)	99	0.0242	333	0.017	1.42	-18.6	< 0.0001
Victoria	(PM, PM, TTLV)	99	0.0242	130	0.0066	3.64	-19.3	< 0.0001
	(TTLV, TTLV, TTLV)	99	0.0242	285	0.0146	1.66	-18.6	< 0.0001
	(PM, TTLV, PM)	97	0.0237	248	0.0127	1.87	-18.9	< 0.0001
	(TTLV, TTLV, PM)	97	0.0237	295	0.0151	1.57	-18.8	< 0.0001
Hustler X-One	(MOW, MOW, TTLV)	49	0.0486	37	0.0016	29.72	-1.6	0.1128
	(MOW, MOW, MOW)	48	0.0476	70	0.0031	15.39	-2.4	0.0149
	(MOW, TTLV, MOW)	48	0.0476	39	0.0017	27.62	-2.1	0.0331
	(MOW, TTLV, TTLV)	48	0.0476	28	0.0012	38.47	-2.0	0.0442
	(TTLV, MOW, MOW)	47	0.0466	34	0.0015	31.02	-2.6	0.0088
	(MOW, MOW, MOW, MOW)	47	0.049	36	0.0017	29.7	-1.3	0.1841
	(MOW, MOW, TTLV, TTLV)	47	0.049	9	0.0004	118.81	-0.9	0.3905
Smeal SST	(EX, EX, EX)	12	0.0198	11	0.0005	41.41	-24.0	< 0.0001
	(EX, ‘PUMP’, EX)	11	0.0181	0	0.0000	10000.0	-36.0	< 0.0001
Pumper	(EX, EX, EX, EX)	11	0.0185	3	0.0001	136.93	-30.7	< 0.0001
	(CSM, EX, EX)	11	0.0181	2	0.0001	208.78	-33.2	< 0.0001
	(CSM, EX, EX, EX)	11	0.0185	0	0.0000	10000.0	-34.5	< 0.0001
	(ENG/MS, EX, EX)	11	0.0181	7	0.0003	59.65	-28.4	< 0.0001
	(ENG/MS, EX, EX, EX)	11	0.0185	3	0.0001	136.93	-30.7	< 0.0001

means that the model considers up to 20 previous items in the sequence (if they exist) when predicting the next job. This model uses a dense representation of the input features, which allows it to learn about relationships between repairs to different systems. A feature that makes this model particularly well-suited to our problem is that it utilizes a technique called *dropout* [29] to regularize the model and avoid overfitting, which allows the network to model the vehicle data accurately without learning spurious or irrelevant patterns in the relatively small training dataset.

To assemble training, validation, and testing datasets for the model, we use all data from three make/models all often used as police cars, with similar maintenance patterns: Dodge Charger, Ford Crown Victoria, and Chevrolet Impala (see Table 3 for an illustration of the similarities between frequent repair sequences for the Charger and Crown Victoria). Ideally, a model would be fit on only a single make/model; however, due to the relatively small size of the dataset (in total, this consists of only 329 vehicles), it was necessary to combine multiple make/models. We assemble the repair sequences for each vehicle and train on a random subset of 50% of these vehicles, using 25% for model validation and 25% for testing.

We assess the performance of our model using average per-item perplexity, which evaluates the probability that the model assigns to the test data (an effective model would

assign a high probability to unseen data):

$$e^{-\frac{1}{N} \sum_{i=1}^N \ln(p_{\text{target}_i})} = e^{\text{loss}} \quad (4)$$

The performance of our model, which achieves an average test perplexity score of 15.7, demonstrate that even this relatively simple, computationally lightweight model with a small dataset is able to achieve a reasonable performance on testing data. While perplexity benchmarks vary considerably by task, we can compare this with the perplexity of a ‘random’ model which assigns a probability proportional to the frequency of item in the list to any given sequence of the 50 different system types observed in the training data. According to Equation (4) such a model would achieve a test perplexity of 260 ± 40 , substantially larger than the LSTM model. We can also compare these results to the original application of our model, which achieved perplexity of 23.7 on the Penn Treebank dataset, and the state-of-the-art performance benchmarks on the Google Billion Words dataset [7, 23, 14] (perplexity of 43.8, 28.0, and 24.29, respectively). We note that our model’s low perplexity score cannot be directly compared to model performance on other corpora, however it reflects the relatively high degree of predictability (and the relatively low number of candidate items in the sequence – 81 unique systems in the entire vehicles dataset compared to many thousands in text corpora).

6. CONCLUSIONS AND CHALLENGES

In this analysis, we describe the initial results of a data collaboration between MDST and the City of Detroit’s Operations and Infrastructure Group. This work demonstrates that there is significant, but highly complex, structure in the City of Detroit’s vehicle-maintenance data. The complexity arises from inter-relationships between vehicle type, system repair type, and time (both absolute time and vehicle life time). We employ PARAFAC tensor decomposition to uncover and visualize these relationships. Sequential pattern-mining adapted in this work verifies the time-vehicle relationship, and we find a statistically significant sequential patterns in maintenance. We note that a predictive model can accurately capture this sequence structure to make effective predictions using the available (modest-in-size) data.

This collaboration demonstrates a small sample of the insights that can be gained from detailed multivariate analysis of municipal data, and illustrates several of the challenges of working with such data. Many aspects of the data—its observational nature; overlapping or difficult-to-decipher descriptions; error and incompleteness which are likely systematic and nonrandom⁵—underscore the challenges of working with real-world municipal data often generated as “data exhaust” and not with the express aim of providing insights or accurate measurements. Additionally, the disconnect between our analytical team and the users generating the data (vehicle drivers, technicians, and clerical staff) highlights how challenging it can be to understand data context.

The tools used and generated in this analysis are open-source, including the MATLAB code used to generate the PARAFAC decompositions [4, 6] and the Python and R code used to clean, analyze, and model the data⁶. We hope that this will lead to further similar data explorations in other domains, extensions of our methodology.

7. FUTURE WORK

There are several promising avenues for future research. While we apply the PARAFAC decomposition to *vehicle* × *system* × *time* data tensors, nothing about this approach limits it to evaluating these three specific variables. This analysis should be applied to several other dimensions, including WAC and Job Descriptions, other measures of vehicle lifetime (i.e., mileage), garage (location) and technician. Future work can utilize this and other information to build more robust predictive models to predict demand, maintenance costs, and vehicle downtime (repair duration), and to assess maintenance effectiveness. Furthermore, as is demonstrated by the wide variety of applications of tensor decompositions discussed in Section 2, this approach could be extended to any civic data where complex multivariate relationships exist by using the open-source code provided.

Acknowledgements

This work is partially supported by National Science Foundation, grant IIS-1453304. It would not have happened

⁵For example, technicians subjectively choose between several job codes: i.e., “Adjust brakes” vs. “Repair brakes” vs. “Overhaul brakes”; many older vehicles and jobs are believed to be missing from this data.

⁶<https://gitlab.eecs.umich.edu/mdst/D4GX-2017-Detroit-Vehicles>

without the support of the broader Michigan Data Science Team. The authors recognize the support of Michigan Institute for Data Science (MIDAS) and computational support from NVIDIA. We would like to thank the General Services Department of the City of Detroit for bringing this project to our attention and making the data available for use.

8. REFERENCES

- [1] [PDF]Mining and forecasting of big time-series data.
- [2] E. Acar, S. A. Çamtepe, and B. Yener. Collective sampling and analysis of high order tensors for chatroom communications. In *Intelligence and Security Informatics*, pages 213–224. Springer, Berlin, Heidelberg, 2006.
- [3] R. Agrawal and R. Srikant. Mining sequential patterns. In *Proceedings of the Eleventh International Conference on Data Engineering*, pages 3–14, 1995.
- [4] B. Bader and T. Kolda. Efficient MATLAB computations with sparse and factored tensors. *SIAM J. Sci. Comput.*, 30(1):205–231, 2007.
- [5] B. W. Bader, M. W. Berry, and M. Browne. Discussion tracking in enron email using PARAFAC. In M. W. Berry and M. Castellanos, editors, *Survey of Text Mining II*, pages 147–163. Springer London, 2008.
- [6] B. W. Bader and T. G. Kolda. MATLAB tensor toolbox version 2.0, 2006.
- [7] C. Chelba, T. Mikolov, M. Schuster, Q. Ge, T. Brants, P. Koehn, and T. Robinson. One billion word benchmark for measuring progress in statistical language modeling. 2013.
- [8] J. Douglas Carroll and J.-J. Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of “Eckart-Young” decomposition. *Psychometrika*, 35(3):283–319, 1 Sept. 1970.
- [9] D. D. Gransberg and P. Investigator. Major equipment life-cycle cost analysis.
- [10] S. Hochreiter and J. Schmidhuber. LSTM can solve hard long time lag problems. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems 9*, pages 473–479. MIT Press, 1997.
- [11] J. S. Kinnebrew, K. M. Loretz, and G. Biswas. A contextualized, differential sequence mining method to derive students’ learning behavior patterns. *JEDM*, 5(1):190–219, 2013.
- [12] T. Kolda and B. Bader. Tensor decompositions and applications. *SIAM Rev.*, 51(3):455–500, 2009.
- [13] D. Koutra, E. E. Papalexakis, and C. Faloutsos. TensorSplat: Spotting latent anomalies in time. In *2012 16th Panhellenic Conference on Informatics*, pages 144–149. ieeexplore.ieee.org, 2012.
- [14] O. Kuchaiev and B. Ginsburg. Factorization tricks for LSTM networks. 2017.
- [15] P. T. Lauria and D. T. Lauria. *State Department of Transportation Fleet Replacement Management Practices*. Transportation Research Board, 2014.
- [16] C. Lee, W.-Y. Loh, X. Qin, and M. Sproul. Development of new performance measure for winter maintenance by using vehicle speed data. *Transportation Research Record: J. of the Transportation Research Board*, 2055:89–98, 2008.
- [17] N. Levine, K. E. Kim, and L. H. Nitz. Spatial analysis of honolulu motor vehicle crashes: I. spatial patterns. *Accid. Anal. Prev.*, 27(5):663–674, Oct. 1995.
- [18] N. Levine, K. E. Kim, and L. H. Nitz. Spatial analysis of honolulu motor vehicle crashes: II. zonal generators. *Accid. Anal. Prev.*, 27(5):675–685, 1995.
- [19] Y. Lv, Y. Duan, W. Kang, Z. Li, and F. Y. Wang. Traffic flow prediction with big data: A deep learning approach. *IEEE Trans. Intell. Transp. Syst.*, 16(2):865–873, 2015.
- [20] J. Mocks. Topographic components model for event-related potentials and some biophysical considerations. *IEEE Transactions on Biomedical Engineering*, 35(6):482–484,

1988.

- [21] E. B. Osborne. From measurement to management: A Performance-Based approach to improving municipal fleet operations in burlington, north carolina. Master's thesis, The University of North Carolina at Chapel Hill, 2012.
- [22] R. Prytz. *Machine learning methods for vehicle predictive maintenance using off-board and on-board data*. PhD thesis, Halmstad University, 2014.
- [23] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean. Outrageously large neural networks: The Sparsely-Gated Mixture-of-Experts layer. 2017.
- [24] J.-T. Sun, H.-J. Zeng, H. Liu, Y. Lu, and Z. Chen. CubeSVD: A novel approach to personalized web search. In *Proceedings of the 14th International Conference on World Wide Web*, WWW '05, pages 382–390, New York, NY, USA, 2005. ACM.
- [25] E. I. Vlahogianni, M. G. Karlaftis, and J. C. Golias. Optimized and meta-optimized neural networks for short-term traffic flow prediction: A genetic approach. *Transp. Res. Part C: Emerg. Technol.*, 13(3):211–234, 2005.
- [26] J. Wang, J. Han, and C. Li. Frequent closed sequence mining without candidate maintenance. *IEEE Trans. Knowl. Data Eng.*, 19(8):1042–1056, 2007.
- [27] Weizhong Zheng, Der-Hong Lee, and Qixin Shi. Short-Term freeway traffic flow prediction: Bayesian combined neural network approach. *J. Transp. Eng.*, 132(2):114–121, 1 Feb. 2006.
- [28] J. White. How detroit's fleet survived city bankruptcy. *Government Fleet*, Nov. 2015.
- [29] W. Zaremba, I. Sutskever, and O. Vinyals. Recurrent neural network regularization. 2014.