Data Acquisition and Processing Using Artificial Intelligence and Machine Learning

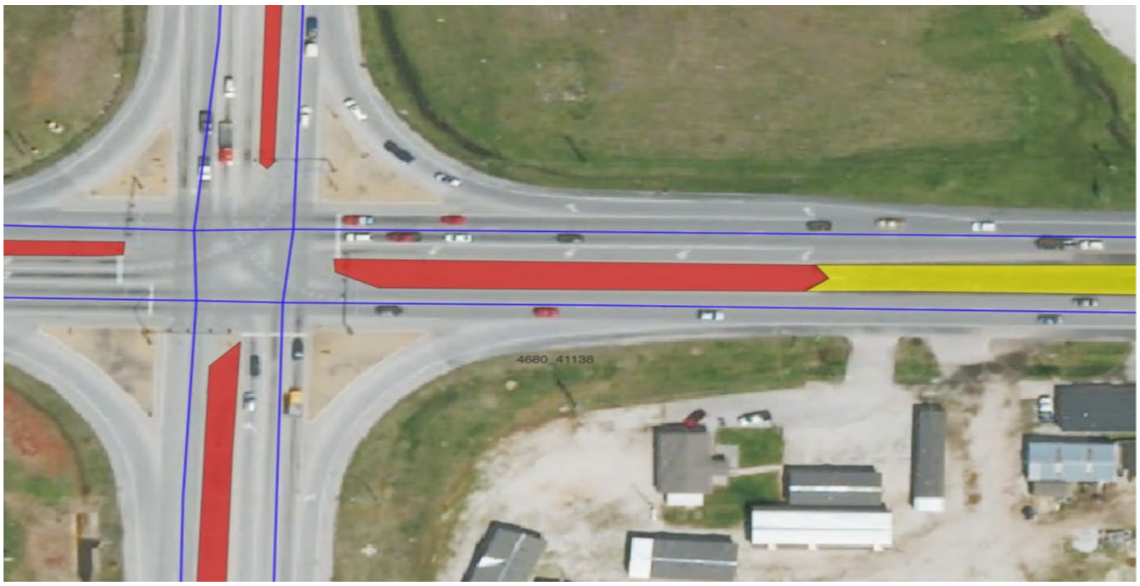

July 2024
Final Report

Project number TR202215
MoDOT Research Report number cmr 24-009

PREPARED BY:

Ryan Loos

Mark Egge

Anna Batista

High Street Consulting Group

PREPARED FOR:

Missouri Department of Transportation

Construction and Materials Division, Research Section

TECHNICAL REPORT DOCUMENTATION PAGE

<table>
<tr><td colspan="2">1. Report No.
cmr 24-009</td><td colspan="2">2. Government Accession No.</td><td colspan="2">3. Recipient's Catalog No.</td></tr>
<tr><td colspan="4" rowspan="2">4. Title and Subtitle
Data Acquisition and Processing Using Artificial Intelligence and Machine Learning</td><td colspan="2">5. Report Date
June 2024
Published: July 2024</td></tr>
<tr><td colspan="2">6. Performing Organization Code</td></tr>
<tr><td colspan="4">7. Author(s)
Ryan Loos, https://orcid.org/0009-0004-7362-3011
Mark Egge, https://orcid.org/0009-0007-9128-2099
Anna Batista, https://orcid.org/0009-0007-4503-4860</td><td colspan="2">8. Performing Organization Report No.</td></tr>
<tr><td colspan="4" rowspan="2">9. Performing Organization Name and Address
High Street Consulting Group
6937 Blenheim Ct.
Pittsburgh, PA 15208</td><td colspan="2">10. Work Unit No.</td></tr>
<tr><td colspan="2">11. Contract or Grant No.
MoDOT project # TR202215</td></tr>
<tr><td colspan="4" rowspan="2">12. Sponsoring Agency Name and Address
Missouri Department of Transportation (SPR-B)
Construction and Materials Division
P.O. Box 270
Jefferson City, MO 65102</td><td colspan="2">13. Type of Report and Period Covered
Final Report (June 2022-June 2024)</td></tr>
<tr><td colspan="2">14. Sponsoring Agency Code</td></tr>
<tr><td colspan="6">15. Supplementary Notes
Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration. MoDOT research reports are available in the Innovation Library at https://www.modot.org/research-publications.</td></tr>
<tr><td colspan="6">16. Abstract
Artificial intelligence (AI) and machine learning (ML) offer state transportation agencies promising tools to help them advance their missions by informing decision making; improving information accuracy, completeness, and timeliness; and automating tedious tasks to free up valuable DOT resources. The objective of this research was to provide MoDOT with tools, information, and examples to implement and leverage these technologies. The project reviewed existing work efforts and explored opportunities where AI and ML could replace existing work activities or augment and improve them. This was completed by developing a universe of potential DOT applications, screening them based on key criteria, scoping five potential projects, and executing two pilot projects. Findings from the idea generation process and execution of the pilots can be applied to MoDOT's future AI and ML endeavors.</td></tr>
<tr><td colspan="3">17. Key Words
Artificial intelligence; Machine learning; Deep learning; Computer vision; Factor grouping; Median inventory; Advance analytics; Automated data collection</td><td colspan="3">18. Distribution Statement
No restrictions. This document is available through the National Technical Information Service, Springfield, VA 22161.</td></tr>
<tr><td>19. Security Classif. (of this report)
Unclassified.</td><td colspan="2">20. Security Classif. (of this page)
Unclassified.</td><td colspan="2">21. No. of Pages
57</td><td>22. Price</td></tr>
</table>

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Data Acquisition and Processing Using Artificial Intelligence and Machine Learning

High Street Consulting Group

6937 Blenheim Ct., Pittsburgh, PA 15208

Authors

Ryan Loos, PE

Mark Egge

Anna Batista

Prepared for

Missouri Department of Transportation

1617 Missouri Blvd., Jefferson City, MO 65102

June 2024

Final Report

TR202215

Copyright

Disclaimer

The opinions, findings, and conclusions expressed in this document are those of the investigators. They are not necessarily those of the Missouri Department of Transportation, U.S. Department of Transportation, or Federal Highway Administration. This information does not constitute a standard or specification.

Acknowledgments

The authors would like to thank the Missouri Department of Transportation (MoDOT) for sponsoring this research. The authors would also like to acknowledge the assistance provided by MoDOT personnel, including the Technical Advisory Committee (TAC). The TAC members include Jenni Hosey, Research; Jen Harper, Research; Thomas Blair, St. Louis District; Thomas Honich, Traffic Engineering; Michael Middleton, Maintenance; Alexander Wassman, Highway Safety and Traffic; James Whaley, Transportation Planning; Jason Volkart, Information Systems.

The authors also appreciate the input from additional personnel at MoDOT whose guidance was paramount to identifying, selecting, and completing the pilot project efforts: Ploisongsaeng Intaratip, Spencer Robinson, Alexander Schroeder, Brian Reagan, Eddie Watkins, Tommy Caudle, Cheri Middendorf, Kevin Stegeman, Randy Pittman, Natalie Roark, and Nicole Hood.

Finally, we would like to thank our consultant partners at Strong Analytics, Andy Wong, Brock Ferguson, and Jacob Zweig for their integral contributions to the idea generation workshop and performance on the Median Inventory pilot project.

Table of Contents

List of Tables

List of Figures

List of Abbreviations and Acronyms

AADT ....Annual Average Daily Traffic

AI ....Artificial Intelligence

ARAN ....Automatic Road Analyzer

AVL ....Automatic Vehicle Locator

COCO ....Common Objects in Context

DOT ....Department of Transportation

KPI ....Key Performance Indicator

LiDAR ....Light Detection and Ranging

MoDOT ....Missouri Department of Transportation

ML ....Machine Learning

MMS ....Maintenance Management System

NDOT ....Nebraska Department of Transportation

SME ....Subject Matter Expert

STIP ....Statewide Transportation Improvement Program

TAC ....Technical Advisory Committee

TMC ....Transportation Management Center

Executive Summary

The convergence of trends in big data—a dramatic increase in the volume, variety, and velocity of data—and repeated breakthroughs in Artificial Intelligence (AI) and Machine Learning (ML) technologies enables decision making with unprecedent amounts of information and at a hitherto inconceivable scale. This potent combination promises to positively transform transportation planning and systems operations.

AI and ML models are good at answering only a narrow range of questions (how many, how much, when, to which group does something belong, what is shown in a photo, etc.) but, once trained, can produce millions of decisions and predictions at a very low cost.

Given the tremendous promise of AI and ML, should Missouri DOT be doing more to actively incorporate these technologies into the agency? This research provides MoDOT with an opportunity to “kick the tires” with AI and ML and better understand how these powerful technologies can help the agency meet its present and future challenges.

This report—and the research activities preceding it—endeavors to equip MoDOT with the necessary knowledge and evaluative framework to effectively assess opportunities to incorporate AI and ML technologies into the agency and make the most of these promising new technologies.

The project was executed through capacity building activities and applied evaluation methods. To build AI/ML capacity, a workshop with leaders from across the department was held to introduce AI and ML concepts and to generate ideas for their use at the department; and a peer exchange was hosted with AI/ML leaders from other agencies to share the experiences and knowledge.

Of the 50+ work ideas for AI/ML applications at MoDOT, two were ultimately implemented to provide hands on experience and a benefit cost analysis: a median inventory derived from satellite imagery using computer vision, and a new method for Annual Average Daily Traffic (AADT) factor grouping using clustering and classification machine learning techniques.

Neither pilot project produced a favorable benefit/cost ratio, but both point in the direction of circumstances where a favorable ratio could be achieved.

Key takeaways from the research include:

1. Scale matters. 10,000 repetitions of a model output should be considered a minimum viable scale for ML to be cost effective. Millions of repetitions are better. In general, this implies that ML will rarely be a good replacement for activities currently being done by MoDOT employees, as few DOT processes are currently being repeated tens of thousands of times.
2. Structure is important. Few decisions are made tens of thousands of times because human decision making is expensive. Unlocking the benefits of AI and ML will require fundamental rethinking and restructuring processes or even whole divisions around the possibility of dramatically lower decision costs.

As a result of this research Missouri DOT staff have gained exposure to AI and ML and provide an opportunity for hands on experience with these tools. This hands-on experience will aid in educating MoDOT staff to identify future opportunities where AI and ML can provide in meeting the unique challenges of current era with its unique new tools and technologies.

Chapter 1. Introduction

Background

This report is written amidst the confluence of two powerful trends that promise to radically transform the business and operations of the Missouri Department of Transportation (Missouri DOT, MoDOT): a dramatic increase in the volume, variety, and velocity of data; and repeated technological breakthroughs in Artificial Intelligence (AI) and Machine Learning (ML). The combination of big data with AI/ML promises to transform the methods of many currently-human-based tasks, including asset inventories, inspection activities, and design surveys. These activities could be effectively augmented or replaced by leveraging AI and ML capabilities.

Missouri DOT's St. Louis Transportation Management Center (TMC) perfectly embodies this ongoing transformation. The center receives a tremendous volume of data in real time from a multitude of sources including embedded roadway sensors, video feeds, third-party probe data, incident reports from other agencies, live weather data, and more. Video feeds are processed through advanced deep learning algorithms that can count vehicles, detect roadway debris, and identify crashes. Powerful deep learning methods are being deployed to draw from all of these sources combined to not just improve traffic incident response times and provide real-time information via message boards but even to predict traffic incidents before they occur.

AI encompasses the computational processes through which computers derive insights and conclusions from data inputs. While human intelligence relies on the brain to receive, store, and analyze information, AI utilizes diverse technologies to perform similar functions. Within AI, ML represents a branch where computers autonomously create processes to analyze extensive datasets, developing algorithms and procedures to discern patterns, trends, and insights from the data itself (Google, 2024).au

These advancements hold the potential for tangible benefits, including savings in personnel time, enhanced safety for both highway workers and the public, and improved access to data crucial for informed decision-making. As more state departments of transportation recognize these advantages, it becomes imperative for entities like MoDOT to embrace and integrate these technologies into their operations, ensuring they remain at the forefront of industry developments and continue delivering on their mission efficiently in an environment where requirements and compliance is increasing, and capital is static or decreasing.

Project Goals and Objectives

This project's goal is to assist Missouri DOT in better understanding and evaluate the opportunities offered by AI and ML to support attainment of the agency's goals.

The objectives to be completed as part of this project are outlined below:

Literature Review: Provide an assessment of AI and ML practices, particularly highlighting examples of where state departments of transportation are currently implementing AI/ML technologies.

Data Acquisition, AI Training, and Analysis: Reviewed MoDOT data collection and analysis efforts and recommended five work activities that could be replaced or augmented by AI and ML. Our team completed this effort by holding an in-person workshop with leaders from across the agency, meeting with the maintenance division (who maintains large data stores), and the St. Louis traffic management center staff. A set of screening criteria was used to whittle down potential projects and select pilots.

Development of AI Processes: Develop two AI/ML pilot products. The development of these pilots gave DOT subject matter experts exposure on what steps are required, and what considerations need to be made when applying AI/ML technology.

Evaluate Benefit-Cost of AI and ML Methodologies: After the development and delivery of the pilot products, the team assessed the benefits and costs associated with the pilot as compared to existing practice. While AI/ML methods can be a great tool for the right problem, it is not a magical solution for all problems. This evaluation provided additional insight as to when AI/ML can and should be applied.

Reporting: The project team provided a demonstration of each of the pilot projects and shared findings and lessons learned from each endeavor that can be leveraged for future instances in which AI/ML technologies are being considered as a potential solution.

Through these objectives, the goal of this effort is to introduce MoDOT to practical applications of AI/ML. The introduction aids the agency by providing hands on opportunity to learn about and apply AI/ML technology and will build staff intuition on what AI/ML is capable of, when it can be applied, potential shortcomings, and when it can be most cost-effective. While the scope of this research effort was narrow, the potential for these applications exists across the department and should continue to be explored.

Project Team

The research project is led by High Street Consulting Group with additional technical consulting from Strong Analytics. Key MoDOT staff formulate the Technical Advisory Committee (TAC) and included representatives from various disciplines and departments: Jenni Hosey, Research; Jen Harper, Research; Thomas Blair, St. Louis District; Thomas Honich, Traffic Engineering; Michael Middleton, Maintenance; Alexander Wassman, Highway Safety and Traffic; James Whaley, Transportation Planning; and Jason Volkart, Information Systems. The purpose of the TAC is to guide the pilot project selection process and provide technical input and documentation reviews during the pilot project delivery process. In addition to the TAC, subject matter experts (SMEs) are to assist with the development and review of pilot projects and coordination of the peer exchange. SMEs who will assist with pilot project development include: Ploisongsaeng Intaratip, Spencer Robinson, Alexander Schroeder, Brian Reagan, Eddie Watkins, Tommy Caudle, Cheri Middendorf, Kevin Stegeman, Randy Pittman, Natalie Roark, and Nicole Hood.

Chapter 2. Literature Review

The integration of AI and ML technologies has reshaped the landscape of transportation management, offering innovative solutions to traditional challenges. There exists extensive literature in the academic arena on in-depth application and assessment of ML algorithms to transportation use cases. This research provides insight into novel algorithms and performance of emerging algorithms versus traditional approaches. While this literature review includes a few examples drawn from this literature, additional examples and case studies were highlighted based on applied uses from public agencies, examples from High Street's previous work efforts, and example products from third-party vendors also working in the transportation space.

The purpose of these examples was to highlight the variety of transportation topic areas that AI and ML can be applied to, as well as the types of problems these technologies can be used to solve. Whether predicting values, classifying groups, clustering data, forecasting, or employing computer vision techniques for inventory management, AI/ML offers solutions across the transportation spectrum. These examples were used to introduce MoDOT staff to AI/ML with real world examples. These same staff would then draw upon these examples, their own day-to-day business efforts, and an AI/ML concepts presentation to help identify opportunities for MoDOT to apply AI/ML internally.

Wisconsin DOT Preservation Cost Model

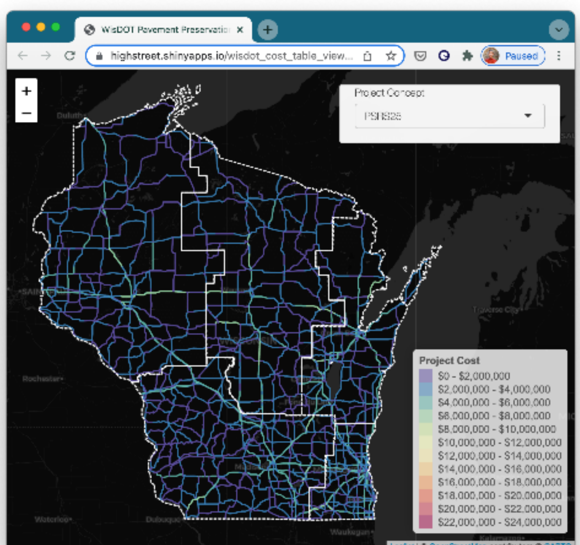

Figure 1. Statewide maintenance cost estimates developed with machine learning.

The Wisconsin Department of Transportation partnered with High Street to predict future maintenance costs for varying levels of maintenance intensity based on historical project costs, geometric information, and traffic information. The team developed a random forest regression model to predict project costs. This model was 30% more accurate than the traditional tabular look up methods used by the department. The model was delivered via simple scripts that were inserted into existing DOT data pipelines.

Roadbotics

Roadbotics (now a part of Michelin DDi) is a private company that uses roadway imagery and proprietary AI algorithms (computer vision) to develop pavement rating maps for public agencies. The company also uses phone gyroscope and AI methods to assess pavement smoothness. These maps provide consistent unbiased ratings at costs lower than that of traditional pavement inspection.

TrafficVision, Acusensus, and MioVision

TrafficVision is one of many proprietary vendors who currently offer incident detection and automated response services. These vendors leverage streaming video and AI algorithms to detect predetermined event types and/or monitor and control ramp meters to help improve response times, decrease crashes, and improve mobility. Like TrafficVision, Acusensus uses similar technology, but has focused more on the enforcement domain and has been used to capture red light runners and drivers using cell phones. Like the above, MioVision uses AI and streaming video to perform traffic counting activities.

Surtrac

Surtrac is an intelligent traffic signal control system that adapts in real time to changing traffic to optimize traffic flows. The technology leverages video detection, reinforcement learning, and machine scheduling to do this. The system is installed on a corridor/grid with signals communicating with their neighbors. The system has been working effectively to reduce travel times and crashes compared to previous optimization and coordination strategies.

Cyclomedia

Cyclopedia is a vendor that provides Light Detection and Ranging (LiDAR) based asset inventory. The company provides this data product by shooting the entire state roadway system to generate a point cloud which is processed with proprietary AI to detect and classify a predetermined set of roadway assets. Assets can include items like striping, signing, traffic control, clearance, among others. This large-scale asset inventory can be used to provide more efficient maintenance operations as well as provide much needed data for safety analyses and planning activities.

Machine Learning Solutions for Bridge Scour Forecast Based on Monitoring Data

This paper (Yousefpour et al., 2021) introduces a novel approach for scour risk management by integrating emerging AI/ML algorithms with real-time monitoring systems for early scour forecast. The research used historical data on riverbed elevation and river stage elevation time-series data. The data was used in combination with neural networks to forecast scour one hour and seven days into the future. The results from these forecasts could be used as an early warning system for critical events. The models could also be used to predict a max scour value, and if those values were over a given threshold they could be flagged for further investigation.

Nebraska Department of Transportation – Crosswalk Inventory

The Nebraska Department of Transportation (NDOT) contracted High Street to find, locate, and inventory painted-pedestrian-crosswalks (crosswalks) on its state highway system. Images taken from NDOT's pavement profiler van were used with out-of-the-box object-detection machine-learning models (e.g. a convolution neural network, or CNN, which are commonly used for computer vision), and Esri's Python API to create an inventory of crosswalks that was stored within a feature layer. This machine-learning, data-driven approach to inventory collection reduced the time and cost associated with collecting inventory information and automated complex GIS based data collection steps. This new asset is planned to be levered in future planning activities and traffic safety analyses. This project is just one example of what is possible using static imagery and ML. The use case can be extended to other roadway assets and/or modified to leverage aerial imagery.

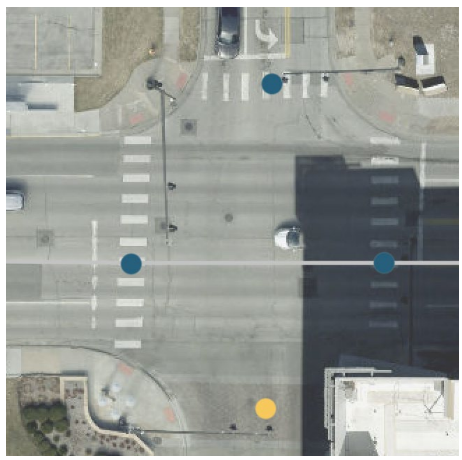

Figure 2. Example identification of crosswalks using computer vision.

Existing Uses of AI and ML at Missouri DOT

Numerous applications of AI and ML are currently in pilot or production status within the department. Two notable uses include:

Traffic Vision. Traffic vision monitors video feeds coming in from the Closed-Circuit Television (CCTV) cameras installed across the St. Louis region and the urban areas of the Kansas City District. The technology uses deep learning to identify load speeds, pedestrians on restricted access facilities, and debris. These identified events are sent as alerts to the TMC. The true positive rate on incident identification since deployment in January of 2022 is greater than 90 percent.

Rekor pre-incident intervention. Funded through an FHWA innovation grant, Missouri DOT (as of 2023) is in the midst of an ambitious AI deployment to help intervene to prevent traffic incidents before they occur. The system uses combined inputs from the many data sources available within the St. Louis traffic management center, including HERE probe data, roadway radar data, radar speed and volume data, Advanced Road Weather Information System (ARWIS) weather data, and more to deploy tools such as ramp metering and messages on dynamic message signs to warn motorists of hazardous driving conditions.

Chapter 3. Project Identification and Scoping

The objective of this task was to surface a wide range of opportunities to apply AI or ML at MoDOT and recommend five potential pilot projects. The intention of the pilot projects was to provide MoDOT staff with hands on experience with AI/ML technologies to develop intuition for where and when they may be applied and what considerations should be made when identifying processes to which AI/ML should be applied, and how to navigate implementation challenges and assess accuracy of results. To achieve this outcome, the research team completed the following objectives:

1. Education and Brainstorming
2. Idea Generation and Refinement
3. Screening and Selection

To complete the above objectives, the research team held a full-day workshop to establish baseline understanding of AI and ML and provide real world industry examples where peer agencies are leveraging these technologies. The workshop also provided a forum for attendees to brainstorm multiple ideas and opportunities that were directly applicable to their work efforts and rank their priority based on a series of quantitative and qualitative criteria. The workshop was followed by in-person site visits and interviews were held with St. Louis TMC and Southwest District Maintenance staff to take a deeper dive into data collection and storage, operations, and key performance indicators (KPIs). After the workshop and site visits a series of informational interviews with key stakeholders from different department divisions to develop understanding of data assets, repetitive and resource intensive processes, and key department outputs and initiatives. These interviews were also used as an opportunity to dive further into ideas that were developed during the workshop and site visits to gain an understanding of the need for the project, potential data sources, and staff who may be involved should an AI/ML process be developed.

Concepts developed through the activities above were qualitatively and quantitatively screened using standardized criteria to narrow the ideas to a list of promising projects. These potential projects had scopes developed and were presented to the TAC for consideration and selection to be moved forward for development.

Details and outcomes from these engagements are summarized in the sections below.

Workshop

Purpose

The goal of the workshop was to introduce and educate MoDOT professionals on what AI and ML are and help develop an intuition of when these tools may be able to be applied to existing or future business practices. The workshop would serve to solicit ideas from attendees, group them together into similar project efforts, and rank them based on attendee preferences.

Participant Identification and Coordination

Participants were identified by the MoDOT research team and were identified based on a large cross-section of DOT functional groups and divisions. Attendees covered a range of employees from subject matter experts, data and analytics managers, and decision makers.

Workshop Agenda and Content

Discussions with the research team became the basis for developing the agenda. Desired discussion and presentation topics included:

- An introduction to the project objectives, research team, and attendees. This included descriptions of the tasks and planned methodologies to achieving research objectives.
- A primer on AI and ML, including definitions, key ingredients for AI and ML, and relatable examples.
- Literature review findings with an emphasis on practical examples being used by public agencies or being offered by third-party vendors. This section was concluded with a summary of promising application areas.
- An interactive session where attendees submitted ideas for AI/ML at DOT.
- A section demonstrating the use of AI, ML, and forecasting techniques.
- A second interactive session where attendees ranked groups of similar AI/ML ideas that were submitted during the first interactive session.
- Wrap up and thank you.

Interactive Session 1 – "Idea Fire Hose"

The goal of this interactive session was to name as many potential DOT ideas as possible. Attendees were asked to answer the key question, "Can you identify opportunities where knowing a future value (or impact), classifying or grouping objects, or optimizing outcomes would serve to automate or augment existing or desired DOT processes?" Attendees were challenged to answer this question in a format that was written as a machine learning problem. That is, they have some input (e.g. a specific piece of data) and a desired output (e.g. a predicted number, item, or group).

To help facilitate the idea generation process, a series of prompts were provided. These included:

- Any immediately obvious opportunities based on any of the content covered so far, or have seen in your own work?
- What data or information do you wish you had that you don't?
- Are there cumbersome repetitive processes that you would automate?
- Are there available data sets you are aware of that aren't being used, or could be used more? What would you do with them?
- Data sets that are suspect in terms of their quality (recency, accuracy, completeness)?
- Do you work with an existing process where you wish you could predict, or group things?
- Do you have issues with data timeliness?
- Do you need to make decisions while balancing competing priorities? Are the impacts of those decisions understood?
- Are you overwhelmed with the data being used to make decisions?
- Do you have issues accessing or aggregating data?
- Can you tie these issues or opportunities back to a department goal or vision?

Interactive Session 2 – Results Ranking

Results from Interactive Session 1 were grouped into similar opportunity groups such that they could be subjectively ranked within common categories. The objective of this exercise was to group like ideas together, refine them based on group discussion, and prioritize them such that the project team could use that frame of reference in sorting down project ideas for scoping. To facilitate the collaborative ranking, the project team used interactive Mentimeter software. Results from the ranking process are illustrated below (scores are based on a vote count by rank weighting):

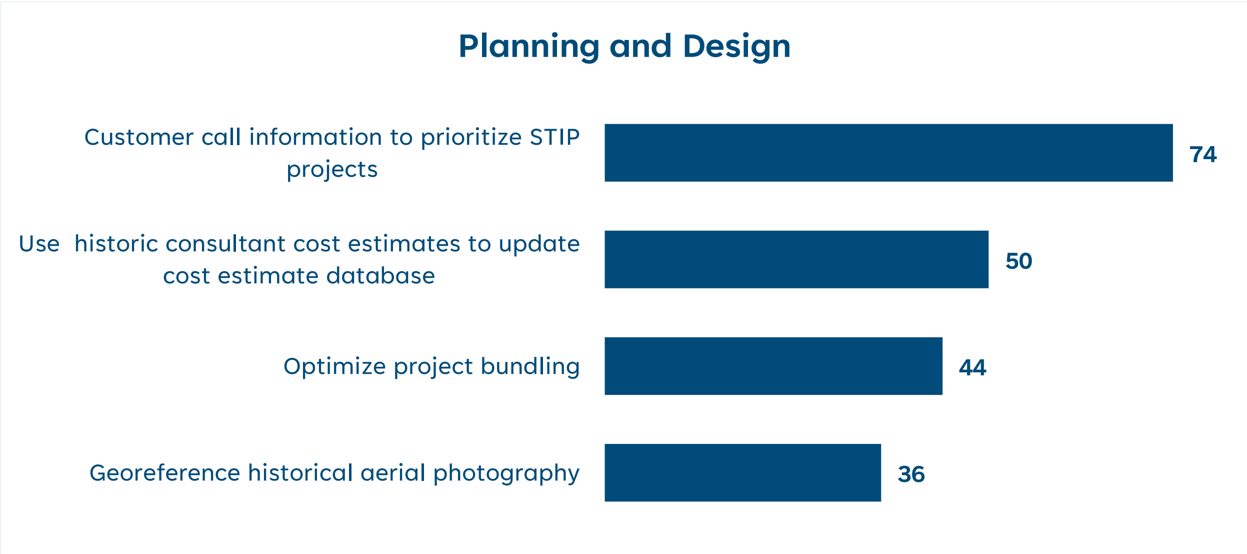

Figure 3. Workshop ranking results for planning and design ideas.

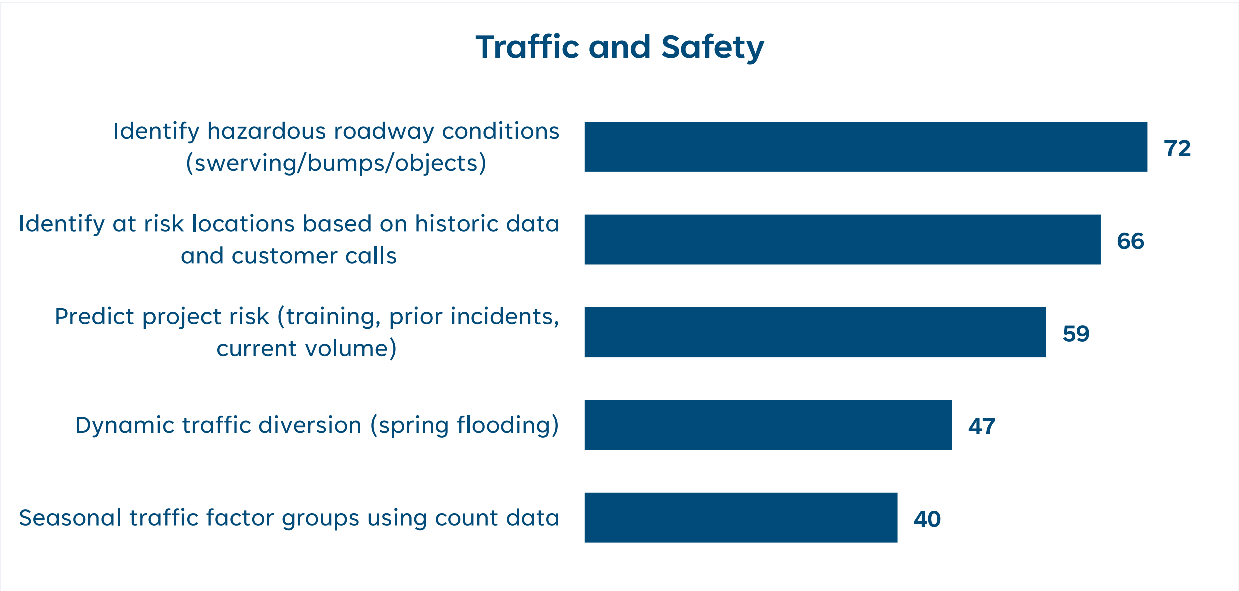

Figure 4. Workshop ranking results for traffic and safety ideas.

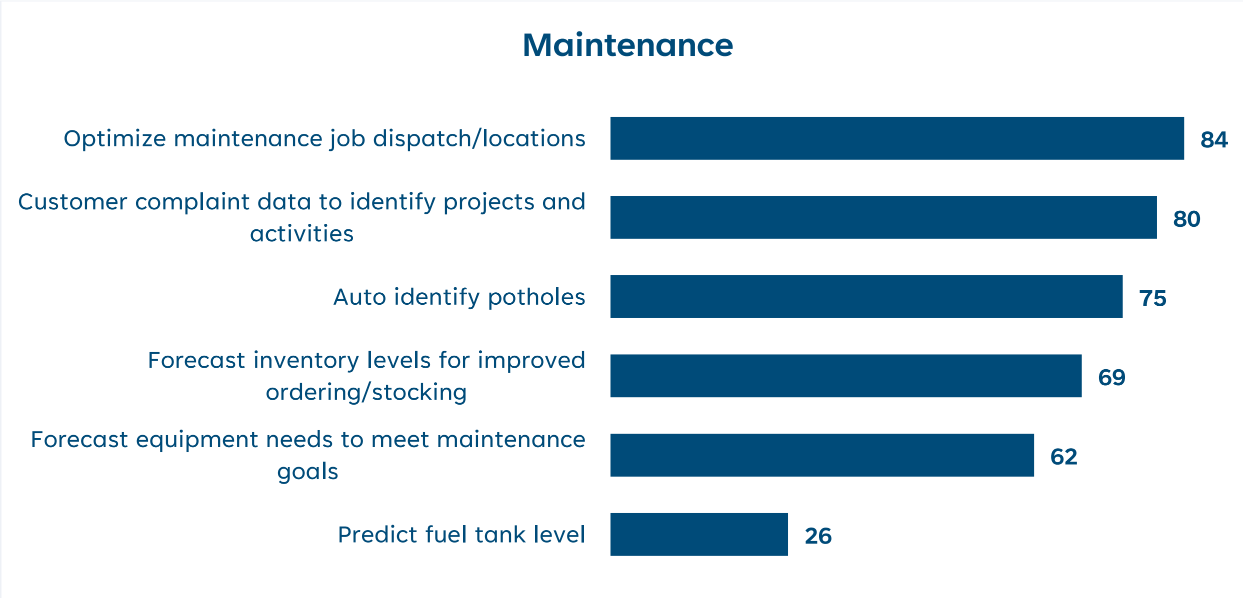

Figure 5. Workshop ranking results for maintenance ideas.

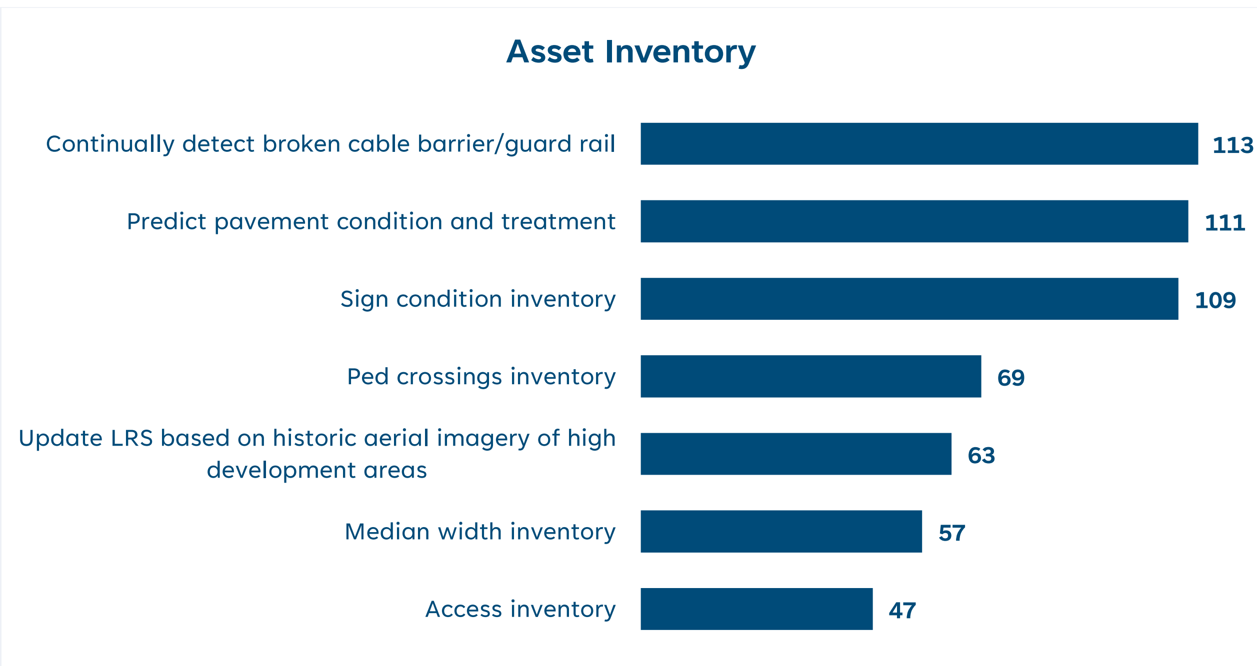

Figure 6. Workshop ranking results for asset inventory ideas.

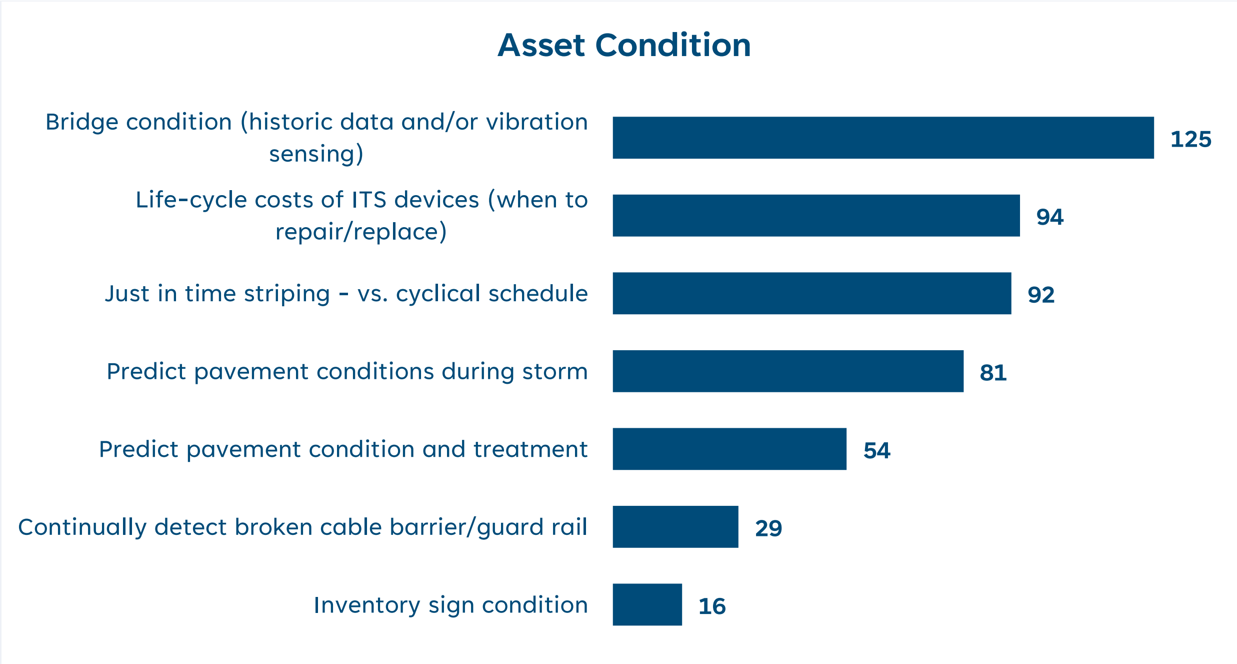

Figure 7. Workshop ranking results for asset condition ideas.

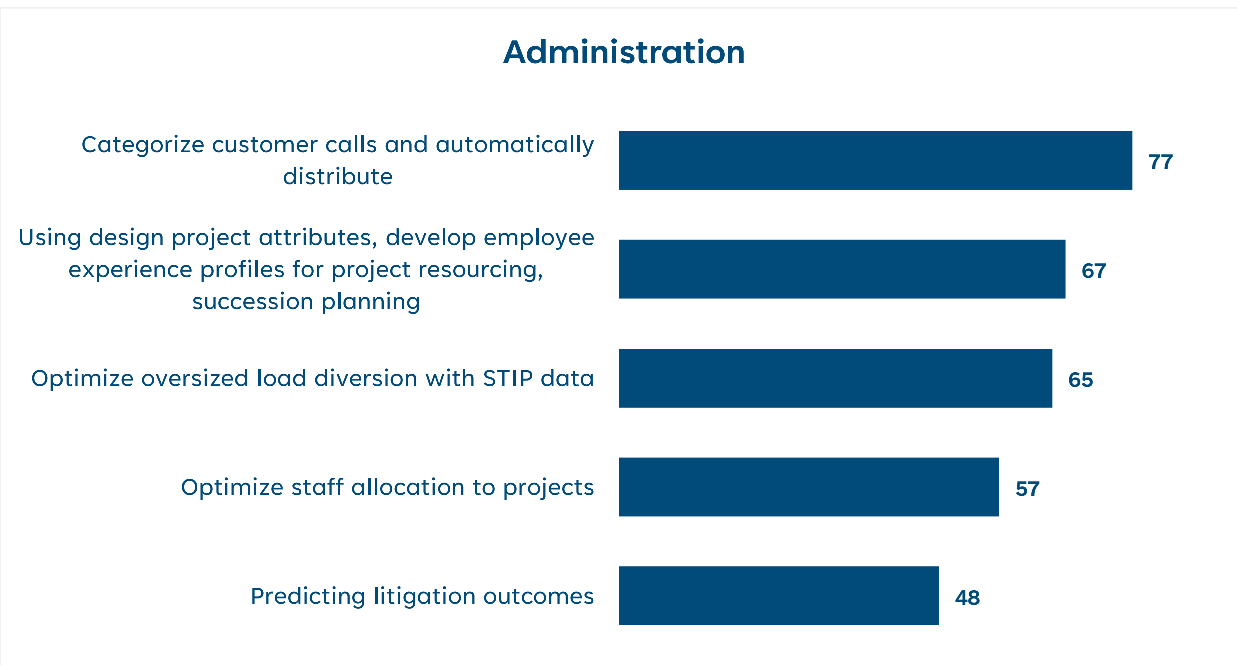

Figure 8. Workshop ranking results for administration ideas.

Results from this exercise informed the project evaluation process.

Site Visits and Informational Interviews

In addition to the in-person interactive workshop, the research team conducted numerous interviews with agency leaders and decision makers, including representatives from the St. Louis TMC and Southwest District Maintenance Team. The following interviews were conducted:

Date	Attendees	Topics
10/31/2022	Jason Volkart (Information Systems Director)	MoDOT IS staffing, infrastructure, and existing data sources
11/04/2022	Natalie Roark (State Maintenance Director) Michael Middleton (Project Director)	MoDOT Management System (MMS) asset management and maintenance activities. Tracks labor, equipment, materials
11/07/2022	Tom Blair (St. Louis District Engineer)	Road to Tomorrow effort, staffing and vacancies, existing TMC AI/ML efforts
11/07/2022	Nicole Hood (State Highway Safety and Traffic Engineer)	Vision Zero, extent of system, 700k signs and inspection, data collection
11/21/2022	James Whaley (Transportation Management System Administrator)	Automatic Road Analyzer(ARAN) vehicles, existing datasets and dashboards, MMS, pavement data, bridge data
11/30/2024	Alex Wassman, Jeff Baird, Jamie Rana, Brian Umfleet, Eddie Watkins, Ploisongsaeng Intaratip	St. Louis TMC site visit. Sensors and ITS devices, traffic incident management, Traffic Vision
12/02/2022	Michael Middleton (Project Director) Tommy Caudle (MMS Manager)	Overview of MMS, maintenance activities, customer calls

These groups and individuals were identified as having access to vast amounts of data and the opportunity to readily apply AI/ML methods and technologies to improve or stand-up new business processes. Some common themes from these interviews were:

- MoDOT has lots of data—but may not be fully leveraging that data. Lots of data in spreadsheets. Data isn't accessible and isn't appropriately targeted to the audience (e.g. may be too granular for executive staff). Not getting the right data to the right people at the right time to effectively inform decision making.
- Concern that the projects getting funded are getting picked because they got a better sales pitch than other projects—not because they move the needle more than other projects. Need to connect data to project impacts.
- Strong desire for more analytics. There are familiar refrains about having too much data and lacking the skills to parse out what the data means.
- General desire for excellence, to better and more meaningfully incorporate data into decision making, to automate repetitive processes and free up staff time (or fill vacancies)
- Self-identified limited familiarity with AI and ML
- Some instances where AI/ML already being used (but not described explicitly in those terms–implemented by a third-party vendor).

Project Identification

Through these activities the project team documented 65 unique ideas where AI and ML may be applied to aid in department business operations. With this list, a preliminary screening was completed to make an assessment if the project idea had a potentially available data set and could be completed within the research project budget in combination with another project. Using these high-level screening factors, a list of 22 materialized projects was generated. The project descriptions for these are listed below.

1. Automate AADT factor grouping.
2. Automatically suggest a maintenance call report priority based on call documentation. Classify by division and who needs to respond and automatically notify the responsible party.
3. Forecast ITS equipment life cycle and associated costs based on historical maintenance cycles.
4. Create a recommendation engine to automatically identify when a newly reported maintenance issue has already been recorded to prevent staff time being spent on resolved issues.
5. Predict traffic volumes on low volume roads.
6. Flag when an area is generating a high number of maintenance issues and should be considered for a future construction or maintenance project.
7. Predict the probability or outcomes of litigation settlements and forecast the amount of settlement money needed for each fiscal year.
8. Perform statistical evaluation of the benefits of bridge flushing and conduct a benefit-cost analysis.
9. Automatically estimate the quantity of asphalt required (in tons) per mile for various maintenance projects based on current conditions and roadway characteristics.
10. Using the 3,300 Automatic Vehicle Locator (AVL)-equipped units, estimate how many times per year each of MoDOT's 72,000 directional centerline miles get driven.
11. Forecast when a bridge will reach poor condition based on historical values.
12. Use roadway imagery and computer vision to estimate striping condition and automatically flag when roads need restriping.
13. Fit deterioration models for agency assets (roads, bridges, etc.) and perform an assessment to determine investment levels (for maintenance, capital projects) needed to maintain roadway condition, improve condition, achieve department goals.
14. Automatically set camera /zoom to focus on likely location of congestion when HERE data indicates an unexpected slow down and flash this camera on the wall.
15. Predict priority one maintenance issues to allow DOT to address proactively rather than reactively.
16. Build a statistical model to predict retroreflectivity (maintained over time with a sampling procedure). Use computer vision to identify issues with sign legibility (encroaching branches, graffiti, etc.).

17. Calculate excess expected recordable worker incidents by county or maintenance office. Identify maintenance shops that have more incidents than their exposure justifies. Estimate worker risk based on experience, type of tasks, prior incidents, time of day, AADT, functional class, etc.
18. Use computer vision to develop an asset inventory for a single asset (profiler van imagery).
19. Use aerial imagery to delineate and measure median widths and mowable acres.
20. For all the consumable materials tracked in MMS (over 2000), create seasonal forecasts for materials consumption to help optimize inventory levels.
21. How do we resource materials and maintenance staff to minimize downside and cost? Ensure adequate inventory of critical materials and labor are maintained.
22. Use travel speeds, incident information, and traffic cameras to detect/classify roadway condition and populate Traveler Information Map.

To narrow this down to the desired five projects for detailed scoping, the project team screened the projects and solicited ranking input from the project TAC. To narrow the world of potential pilot projects, qualitative screening criteria were applied to each idea. These criteria included:

- Data is currently available.
- Feasible within the scope of the research contract.
- Large in scale and is a continuously repeated process.
- Will serve as a good example of how to apply generalized AI/ML practices to MoDOT workflows.
- Has a high magnitude of benefits as perceived relative to other pilot project ideas.
- Has balanced technical and implementation difficulty.
- Has an actionable outcome.
- Has a motivated department champion.

Using these criteria, the project team reduced the field of potential projects to seven strong ideas and a handful of potential runners-up projects that remained strong contenders. These projects had scopes developed and were presented to the TAC for discussion and refinement. After suggestions from the TAC were incorporated into the project's scopes, the TAC was asked to anonymously rank the potential projects on a scale of one to five .

Results from the TAC ranking exercise are shown in Table 1 below:

Pilot Project Concept	Ranking
Sign Inventory	4.38
Project Bundling *(1)	4.00
AADT Factor Grouping	3.77
Bridge Maintenance Optimization	3.75
Access Inventory	3.63
Maintenance Hotspot Identification *(4)	3.50
Median Inventory	3.00
Materials Forecasting	2.99

Table 1. Results from TAC ranking of potential projects.

While the TAC ranking provides a clear prioritization of the final set of projects, there are also some additional considerations for the final project selection. These included aligning the selected suite of projects with the research teams' individual budget and skillsets, ensuring that the projects provide unique examples of AI/ML applications for the department, and considering TAC concerns raised during the ranking process. It should be emphasized that with the right underlying data and adequate budgets, each of these projects are viable AI/ML opportunities that the department could pursue as individual standalone projects. A brief discussion of each project and selection considerations for the nominated project concepts is provided below:

Sign Inventory: Originally this project was devised to forecast sign retroreflectivity degradation over time. This evolved into developing a statewide sign inventory using computer vision and comparing this year over year to detect missing and damaged signs. Issues were identified with the quality of the existing sign inventory and that the Automatic Road Analyzer (ARAN) van imagery that would be used to detect roadway signage may only be collected in one roadway direction. Combined with the technical complexity of training a ML program to detect any sign this project concept began to exceed that which could be accomplished as part of this study.

Project Bundling: The project concept would be to use historical projects letting information, time of year, and project characteristics to group projects in a manner that minimized bid costs to the department. This project was prioritized by a single TAC member and the ranking should not carry the same weight obtained for projects rated by all TAC members.

AADT Factor Grouping: This pilot project would use ML to predict AADT values from short term counts, which could be used to augment or replace the existing standardized process. It is believed that this project would improve the accuracy of the estimates, produce them more rapidly, could be used repeatedly, and would serve as an example to other state agencies who are also federally required to produce this information.

Bridge Maintenance Optimization: Originally this project was devised to develop bridge deterioration models. This research was recently completed by MoDOT, so the scope was adjusted to leverage those models to optimize bridge conditions based on funding or determine required funding to maintain a specified bridge condition. While it was determined that this would be useful for the department, this also has some overlap with the existing AASHTOWare Bridge Management Software that MoDOT currently uses.

Access Inventory: The pilot would utilize ARAN, satellite imagery, and computer vision to identify and inventory local access driveways. The risk identified for this pilot by the project team is due to the potential quality issues of the imagery and the relative size of driveways. This could result in being unable to train an ML model to detect driveways. While it is unknown if this would come to fruition, the project team advised selecting a project with higher likelihood of success.

Maintenance Hotspot Identification: The purpose of the project would be to use customer call logs, travel time data, pavement condition information, and maintenance activity data to identify maintenance hotspots and determine if more intense maintenance activity should be pursued. This project was prioritized by four TAC members and the ranking should not carry the same weight obtained for projects rated by all TAC members.

Median Inventory: The project team would use satellite imager to train an ML model to detect and delineate median areas on the state highway system. The product from this project could be leveraged to estimate costs associated with maintaining these facilities and conducting safety analyses. The project team has determined that there is high quality imagery to produce this result from the southern half of the state. Models would also be available to finish this product once high-resolution imagery becomes available for the rest of the state. While this project is more of a 'one-and-done' application of AI and ML, the approach would serve as an example for doing more of these inventories in the future.

Materials Forecasting: The project team would use historical material inventory data to forecast future inventory allowing maintenance staff to order materials when they are needed. This would reduce storage needs and waste from obsolete and spoiled items. There were no issues raised with this project concept but given the low ranking it was likely to deliver less value or less of a priority to the department than other project concepts.

Final Pilot Project Recommendations

Taking into consideration all the information gathered during the project discovery process, the project team recommends the following projects:

AADT Estimating: Automate and improve estimation of annualized average daily traffic counts from short-term traffic counts. This project will showcase the scalable nature of machine learning (thousands of decisions made annually). The resulting process can be used repeatedly by the department, replacing a more manual and likely less accurate process. This work is required for all state DOTs and this project fits the bill in demonstrating how AI and ML can be leveraged to improve business operations across the industry. The project team believes there is a high likelihood of success and lower risk for this project than some of the other concepts, while also fitting into the study budget.

Median Inventory: Create an inventory of roadway medians (and mowable areas). This project will produce a data asset that does not currently exist at MoDOT and will be useful to multiple departments. This process demonstrates the capacity of machine learning to perform a repetitive process at scale (scanning and delineating thousands of miles of roadway), and the methodologies learned through this effort could be applied to produce additional data assets. This project scored lower than the Access Inventory project, but the project team believes that the data available to complete them is more likely to result in better outcomes for this effort.

Detailed scoping information and results from the pilot projects are discussed in Chapter 4.

Chapter 4. Development of AI and ML Processes

To demonstrate the use of AI and ML the research team executed two pilot projects. Details and findings from these projects are presented below.

Median Inventory

Background

MoDOT currently conducts maintenance activities within roadway medians and ditches (e.g. mowing, garbage clean up, etc.) and levels of effort for these activities are measured in area. MoDOT also does not have an inventory of median facilities to accurately make these estimations. Further, having a median inventory available would aid in conducting follow up analysis that align with MoDOT's vision (e.g. safety). A combination of aerial photography and/or satellite images will be used to apply machine learning techniques to delineate and measure median widths.

Data

This project utilized imagery sourced from the Missouri Spatial Data Information Service (MSDIS). The imagery is 30 cm resolution aerial orthoimagery (that is, an object one meter wide corresponds to roughly three pixels). At the time of the project, imagery was only available for the southern half of Missouri.

Methodology

For this pilot, instance segmentation was used. This method is a computer vision technique that assigns class labels to each pixel in an image based on the corresponding object category. To accomplish this, 5,000 x 5,000 pixel tiles were extracted from the complete satellite view of the state, and instance segmentation was applied to automatically label areas containing unpaved and paved medians. Instance segmentation was selected as it is a more advanced computer vision technique compared to traditional semantic segmentation approaches. While both techniques aimed to categorize objects in an image at the pixel level, instance segmentation goes a step further by differentiating between multiple instances of the same object class. In contrast, semantic segmentation focused solely on assigning class labels to each pixel without distinguishing between individual instances of objects belonging to the same category.

Data Export

Data were extracted from the MSDIS archives. A map of divided highways was provided by MoDOT. Using this, 2,369 square image tiles were identified as containing portions of divided highway. Each tile has an edge length of 5,000 pixels, or 1.5 kilometers. For data annotation by humans, the divided highways were identified and overlaid on these images. For model training and inference, highways were not overlaid on the images.

Data Annotation

To develop a deep learning model capable of predicting the presence of both unpaved and paved medians within a given area, it is essential to first convert satellite imagery tiles into a format amenable to data annotation, as described above. Subsequently, a total of 1,100 image samples were annotated to facilitate the labeling process for segmentation purposes. For the purposes of this project, the dataset is annotated with two distinct classes:

1. Median: Represents unpaved medians.
2. PMedian: Represents paved medians.

An example of this is provided in Figure 9 below, where the red (or dark) shading is the PMedian (paved) and the yellow (or light) shading is the Median (unpaved).

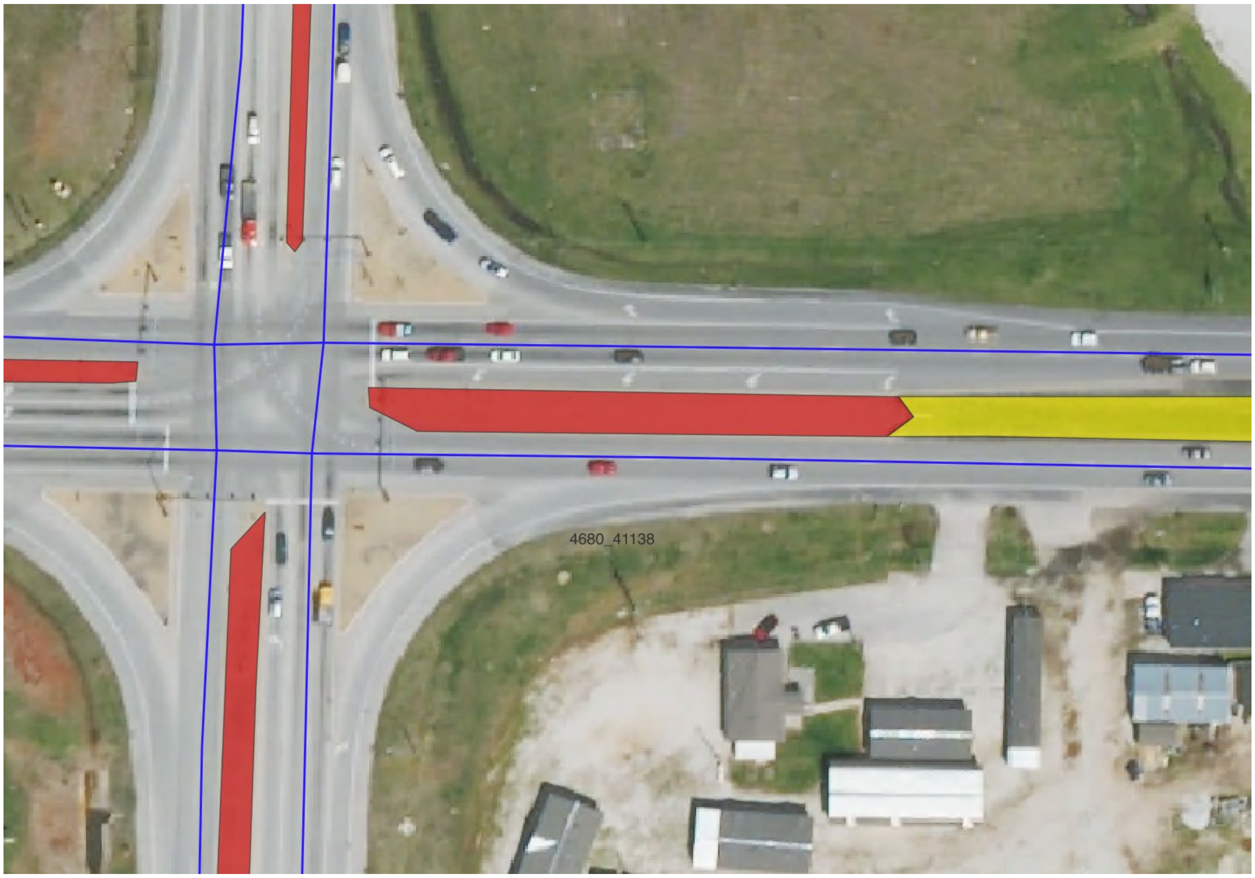

Figure 9. Sample masking of medians used to train computer vision model.

Additional requirements for data labeling followed the instructions below:

1. Median area was designated between roadway shoulders.
2. Abutments, area under bridges, waterways and other facilities were not labeled.
3. Shadows from bridges should not form label boundaries.

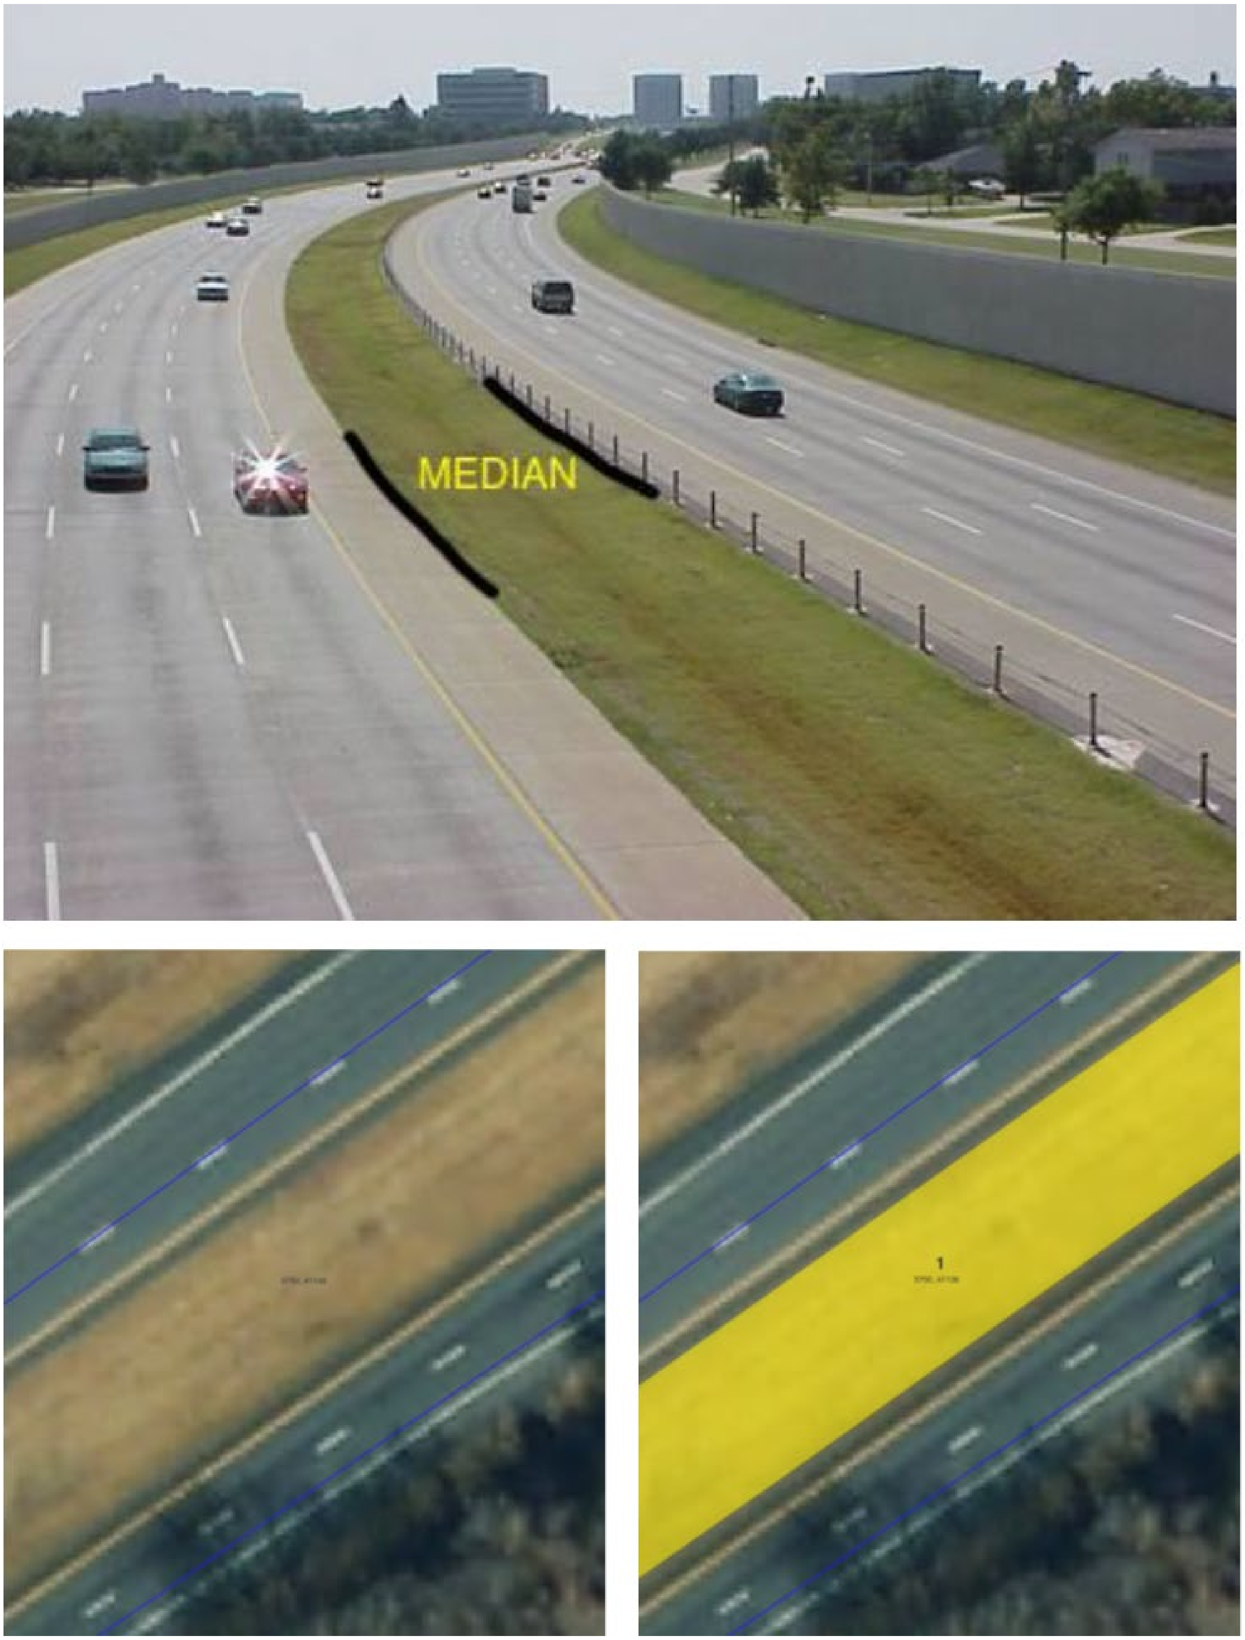

Figure 10. Representation of median labeling between shoulders.

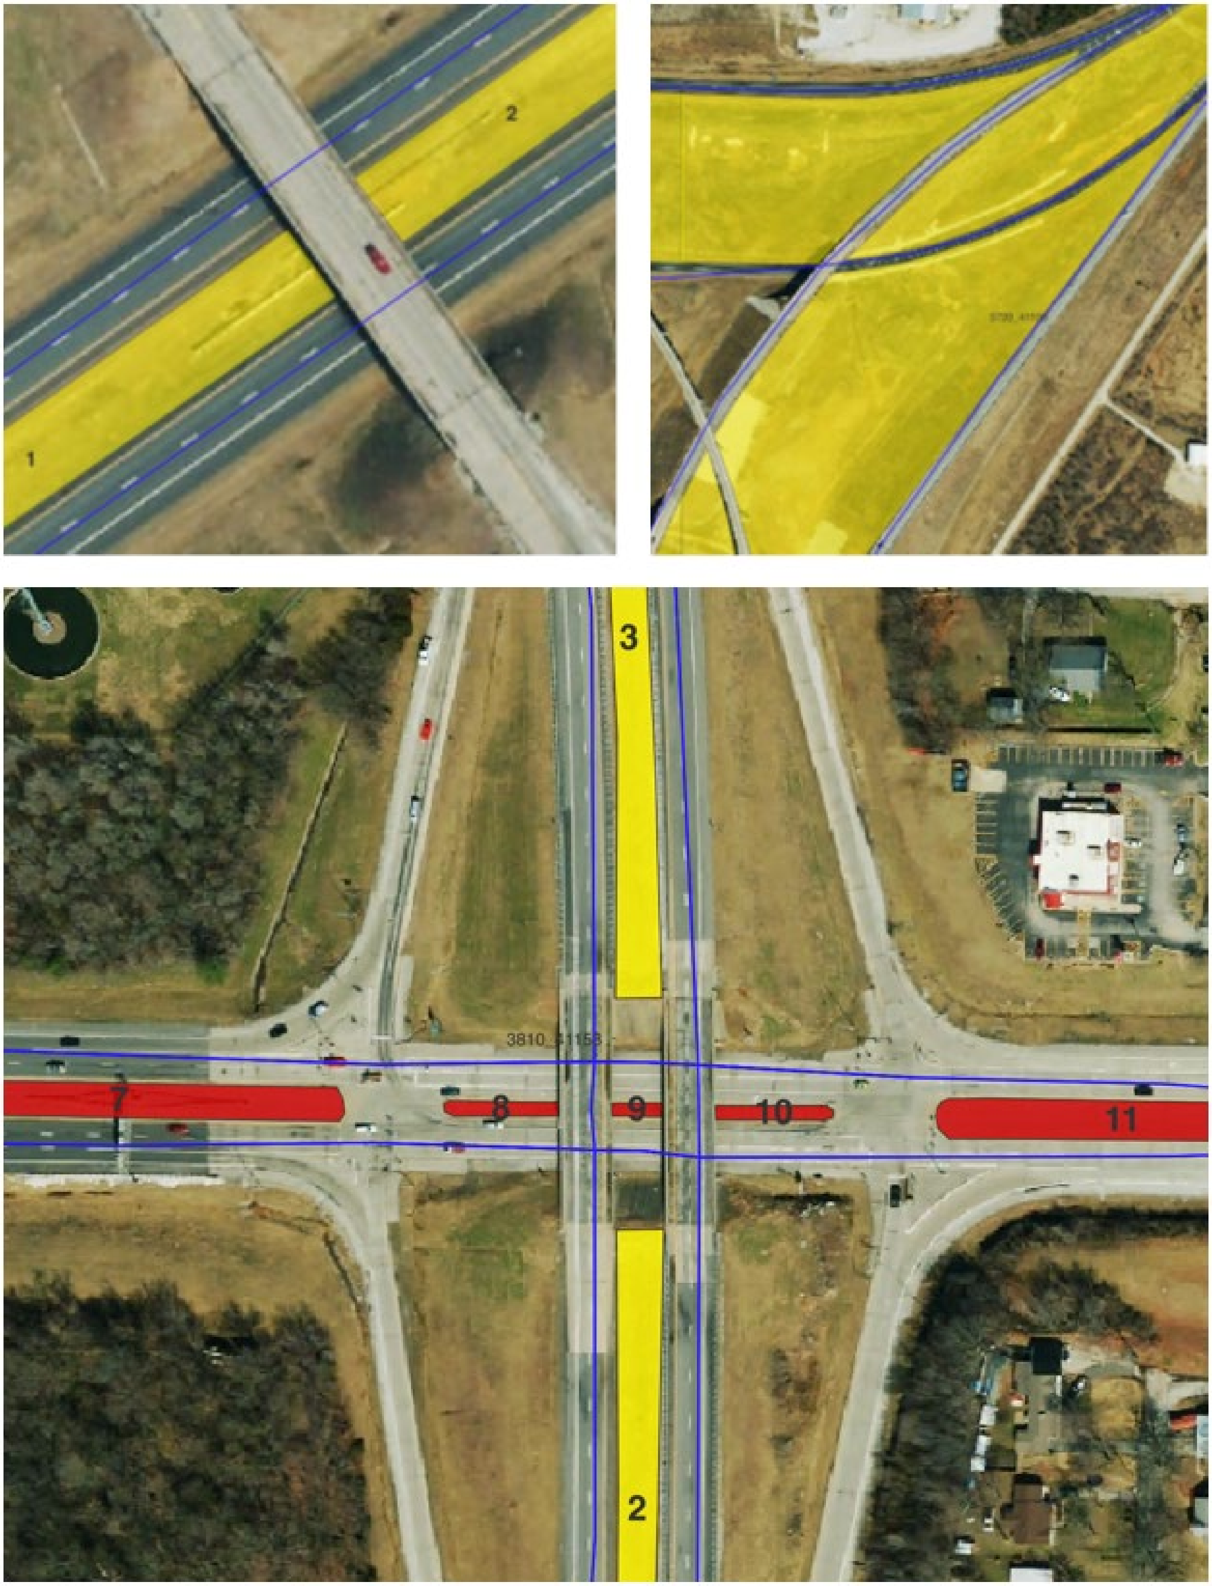

Figure 11. Representation of median labeling near abutments and in interchange gore areas.

Figure 12. Representation of median labeling around bridges and waterways.

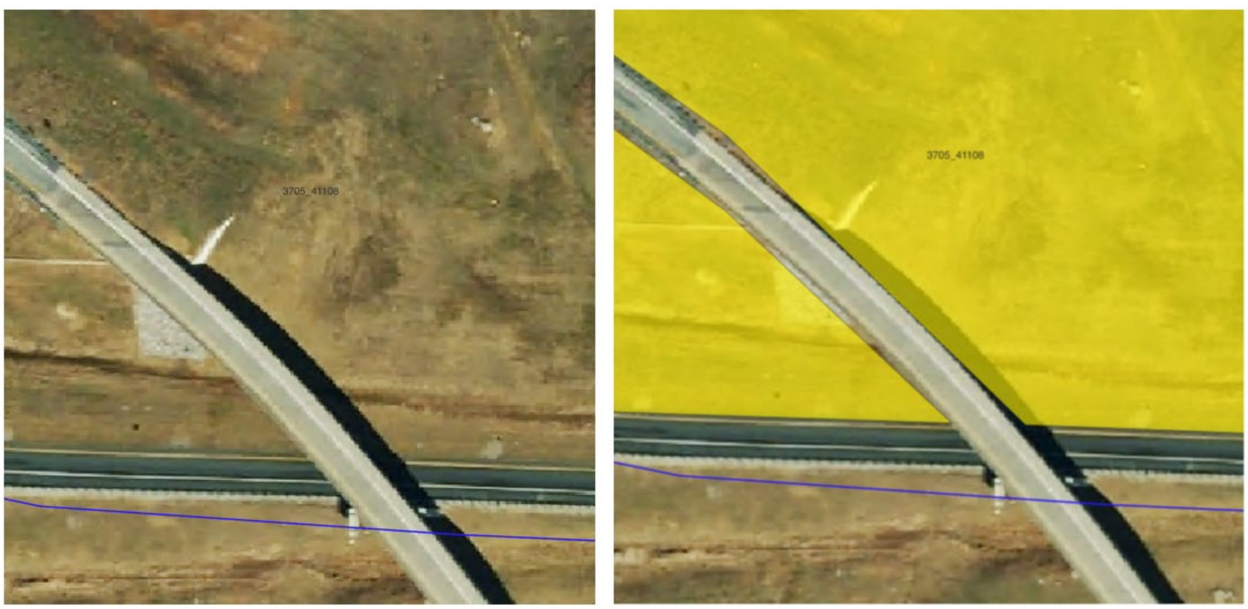

Figure 13. Representation of median labeling accounting for shadowing.

As a result of the data annotation process, a file in the Common Objects in Context (COCO) format is generated, which maps each image in the dataset to the corresponding locations of areas of interest, encompassing both paved and unpaved medians within the image. Utilizing annotations in the COCO format enables seamless training of computer vision models that typically necessitate labels in this format.

Model Training

To solve the challenge of segmenting highway medians, a state-of-the-art instance segmentation model called SparseInst (Sparse Instance Activation for Real-Time Instance Segmentation) was used. The motivation for this decision was its extremely accurate performance, the ability to handle objects of varying scales and spatial locations, and its efficient processing capabilities, making it an excellent choice for this application. The speed and performance of this model also enables rapid processing and iteration and opens the door to real-time applications. SparseInst efficiently recognizes and segments objects in each image. Its speed and simplicity make it an ideal tool for processing images rapidly, while its ability to handle objects of different scales ensures accurate segmentation of highway medians in diverse environments.

The model hyperparameters used for training are shown in the table below:

Parameter	Value
Learning Rate	0.0001
Batch Size	1
Backbone	Resnet50
Weight initialization	Imagenet weights
Training epochs	50
Image size	3,000 x 3,000 (rescaled from 5000x5000)
Augmentations	Random flip/resize
Optimizer	Adam with weight decay (AdamW)
Loss function	SparseInst loss: (Focal, cross-entropy, Dice)

Table 2. Training parameters used for SparseInst model.

Inferencing

Once trained, the model was applied to the 10,000 image tiles (spanning 9,000 square miles). This required approximately one week of processing time. Each image was passed into the model, and the model output a mask (see Figure 14 below) indicating the areas determined to be medians.

Figure 14. SpareseInst model output

Additional Geoprocessing

Output from the models were in the form of 2,369 image masks. These masks needed to be converted to vector polygons, which was completed using built in modules from ArcGIS. These vector polygons were then merged to form a single data layer. The resulting geometries appeared rough due to the natural pixelation associated with the satellite imagery. Vectors were smoothed with ArcGIS packages. Median areas that were falsely identified outside of DOT ROW were also quickly removed based on layer referencing.

Results

The performance of the model has proven to be highly encouraging, providing robust and accurate results even in challenging scenarios. The model achieved a mean pixel accuracy of 99.4 percent and average correct area prediction percentage of 93 percent. These metrics are indicative of its high precision and reliability. This highlights the efficacy of our algorithmic approach in extracting both unpaved and paved medians from satellite imagery, ultimately offering a powerful tool for various applications for MoDOT.

A few samples of the model predictions compared to human labeled examples are included below. The left-hand side are ground truth (human labeled) images and the ML model images are on the right.

Figure 15. Ground truth median identification.

Figure 16. ML model median identification.

Figure 17. Correctly identified medians.

Figure 18. Correctly identified medians with additional gore area.

Figure 19. Correctly identified median absence.

Figure 20. Narrow medians.

While the medians are delineated, there is a fair amount of error. This is due to the fact that a one-meter-wide paved median is only three pixels wide in the imagery, and closer to two pixels after images are resized for model training.

Figure 21. Large error in complex areas.

Discussion

The outputs from this computer vision process both exceeded expectations and failed to meet expectations.

Exceeded Expectations

This project was a technical success. On the whole, the model was very successful in delineating medians. In a matter of weeks, 48 million square feet of median across the southern half of the state (high resolution imagery was not available for the northern half of the state at the time of the study) was delineated. The model was frequently accurate down to the pixel in delineating medians in imagery. From a technical standpoint, the model performed very well. Additional training and refinement (specifically around labeling gore areas) would further improve model performance.

Given an objective of quickly finding and counting the mowable acres of median across southern Missouri DOT maintenance districts, the model performed admirably, and there's reason to believe that the sum of acres per district is accurate within a degree of approximation for inventoried roads.

Failed to Meet Expectations

Despite the technical success of the model training and application, the project failed to produce outputs that were of use to MoDOT. The project originated with the concept of producing estimates of mowable acres per district. To make the project more tractable and to improve the likelihood of a technically successful project, this scope was limited to delineating only medians and excluding shoulders. Thus, while the estimates of medians are accurate within a degree of approximation, the broader question of total mowable acres remains unanswered.

Moreover, there was a hope that this project would yield an inventory of medians that could be added to MoDOT's roadway management systems. Unfortunately, the data is not of sufficient quality of engineering purposes in two regards:

1) The inventory suffers from both false positives (that is, areas incorrectly labeled as median) and false negatives (areas of actual median not labeled in the data). While false positives could be somewhat expediently removed by a human analyst, identifying missing segments of medians, and adding these would require an analyst to painstakingly review the entire network.
2) Precision. The inventory was produced using 30 cm aerial imagery, meaning each pixel in the aerial imagery corresponds to 30 centimeters of actual land. While the model was generally quite accurate, misclassification of even a few pixels would result in an error of a meter or more. In general, 15 cm imagery is typically considered the minimum resolution for delineating roadway features with an acceptable degree of engineering precision. Accuracy within plus or minus a few meters is not sufficiently accurate for an engineering GIS asset.

Could these shortcomings have been overcome?

This project demonstrates the power of computer vision algorithms for inventorying roadway features, but also its current limitations. With additional resources, the model could also have been trained to delineate paved areas from grassy areas, to identify gore areas, and to identify shoulders. An additional round of model training with thousands of additional labeled images would improve model accuracy. Coupled with higher resolution aerial imagery, it is possible to imagine a scenario where the model outputs would meet or exceed expectations.

Several private sector firms have developed a specialty of identifying roadway features from aerial imagery. Their models benefit from a wide variety of high-quality training imagery, and complimentary pre- and post-processing algorithms to enrich imagery with complimentary contextual information. Due to their specialization (and the scalable nature of models, once trained) these private sector firms are able to offer inventory services at a cost-effective price point.

Deliverables

This pilot produced an output GIS layer of delineated medians and a source code repository containing both the Python code used to train the model and the trained model itself. At the time of the pilot, imagery was only available for the southern half of the state. If desired, the trained model could be used to complete the inventory for the northern half of the state when imagery becomes available.

Code for this project can be found in the source code repository: https://bitbucket.org/high-street/median_inventory/

The resulting feature layer of delineated medians is published (at the time of writing) to: https://maps.arcgis.com/apps/mapviewer/index.html?layers=2a3381e9ff234494a88e73c6267e3192

AADT Factor Grouping

Background

Annual Average Daily Traffic (AADT) values express how much traffic a road receives. These values are used for many planning purposes. For the majority of MoDOT's system, AADT values are updated by extrapolating short term count sample count data to annual averages by using daily factors and seasonal and functional classification adjustment factors. AADT values are often further manually adjusted based on underlying roadway characteristics and geography. The pilot project will create an automated process that estimates AADT creates and defines adjustment factor groups using machine learning algorithms and historical AADT and short-term count data. The traffic patterns change across the state over time, this process automates the assignment of factor groups (an otherwise tedious and time-consuming manual process), improving the accuracy of AADT estimates. The project would eliminate the need for any manual review and remove unnecessary subjectivity in developing system wide AADT values.

Data

The project used short term vehicle count data, continuous count data, and roadway inventory data. Data was obtained from MoDOT's own data warehouse for the last seven years. Count data included the location of each count location, urban and rural designation, and functional classification. MoDOT's 170 continuous count stations provided approximately 9.8 million hourly observations and 408,000 daily observations. Short-term counts were obtained for 19,828 locations and contained 1.7 million hourly observations and 110,000 daily observations. Population density was also obtained from available census data and assigned to each count location.

Methodology

To facilitate the AADT factor grouping process, a series of steps was completed:

1. Process data to develop traffic volume profiles (percent of total volume by hour of day, day of week, and month of year) for each count location.
2. Create factor groups by clustering continuous count stations using volume profile data.
3. Assign short-term count locations to factor groups by classifying each location based on functional classification, population density, urban/rural designation, sample volumes.

Clustering

The overall purpose of the process is to group together count stations that have similar traffic volume profiles. Urban commuter routes often have distinct morning and afternoon peaks and lower volumes on weekends. Rural recreation routes have higher volumes on weekends and during the summer. Clustering is an unsupervised machine learning approach in which observations are grouped together based on the similarity in their attributes. Clustering is similar to classification, except that in classification the dataset contains explicit labels, whereas with clustering the assumption is that the labels are latent in the dataset.

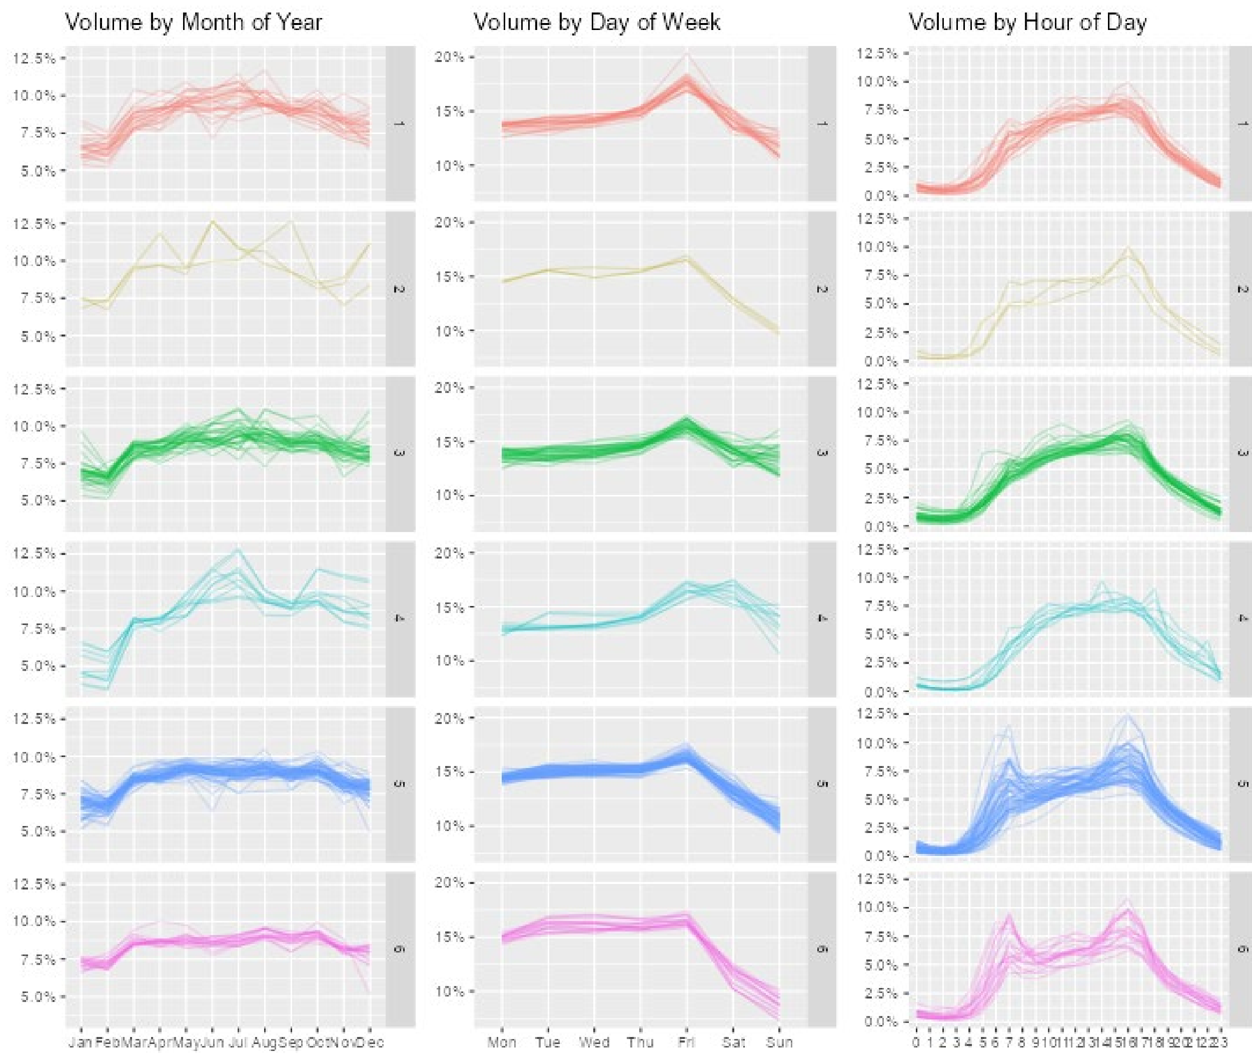

Figure 22. Sample trace plots using six distinct clusters.

The image above shows the volume profiles of the 170 continuous count stations assigned to six clusters. In the image above, we might call group 4 (the 4th row) a recreation-related route (high volumes in summer months and busier on weekends than weekdays, likely serving recreation or tourism-related areas), whereas group 6 (the 6th row) resembles a classic urban commuter route (basically no seasonality, distinct morning and afternoon peaks, distinctly lower volumes on Saturdays and Sundays).

Clustering of continuous count stations was completed using k-means clustering. K-means clustering is an unsupervised machine learning algorithm that groups similar observations (count stations) based on a distance from the centroid between a fixed number of clusters. The algorithm iteratively assigns observations to the nearest centroid based on their distance, then recalculates the centroids based on the newly formed clusters. The process repeats until the centroids do not significantly change. This clustering process was repeated using two to 10 clusters.

How Many Factor Groups (Clusters)?

Review of scree plots, profile plots, engineering judgement and input from Subject Matter Experts (SMEs) were used to determine the number of clusters that should be used for continuous count stations. Additionally, weights were assigned to the hour of day, day of week, and month of year profiles. It was recommended that a weight of zero be used for an hour of day. The reasoning for this is that the clustering algorithm would assign inbound and outbound count stations at the same location to different clusters given that one would have a distinct AM peak and the other a distinct PM peak.

The trace plot sample demonstrates the difference between continuous count station clusters. Each cluster generally has something unique to distinguish it from another. This could be a higher proportion of traffic on Fridays, or a distinct Sunday volume increase, or highly seasonal summer traffic, to highlight just a few examples.

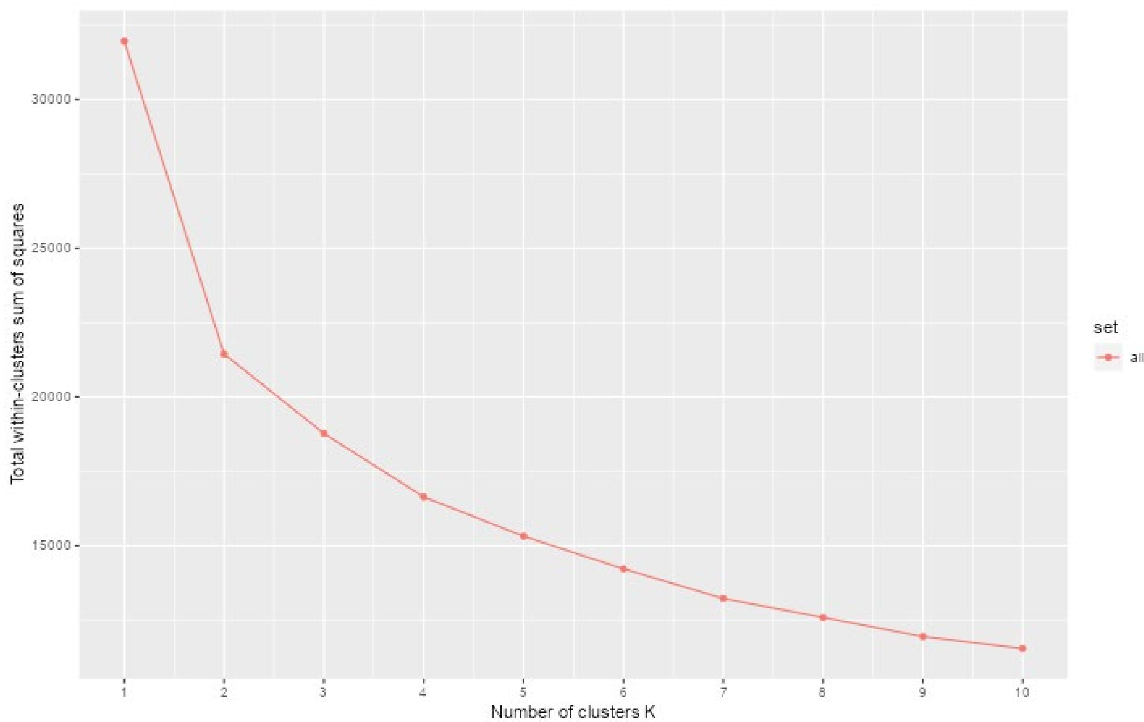

Figure 23. Sample scree plot.

Scree plots are a tool to make an assessment as to what the ‘correct’ number of clusters should be. Where the scree plot bends or knees is typically a good starting point for cluster determination. In the chart above, two, four, or seven have this effect. While this tool can be a good starting point, other factors must be weighed when determining the number of clusters used.

Based on iterative testing and input from SMEs, the clustering was completed using six clusters. A weight of 0.7 was assigned to the day of week profile, and a weight of 0.3 was assigned to month of year for the continuous count station classification.

Results from the clustering algorithm are shown in Figure 24 and Figure 25. The symbol colors indicate which cluster group the count station is assigned to. Each cluster group has a distinct traffic profile shared by the stations within the group.

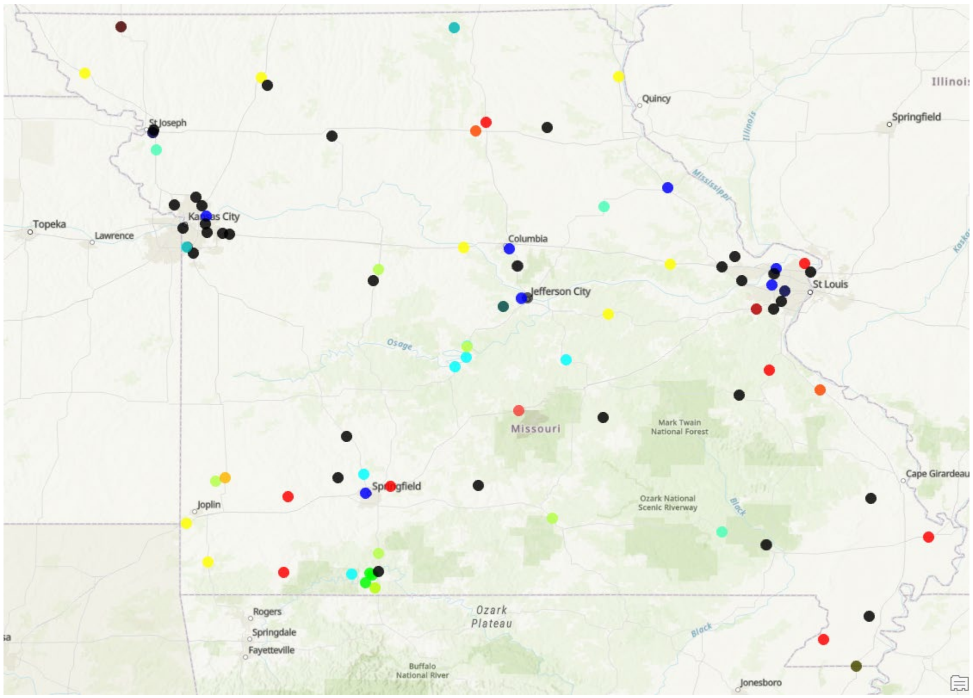

Figure 24. Geospatial continuous count clustering results.

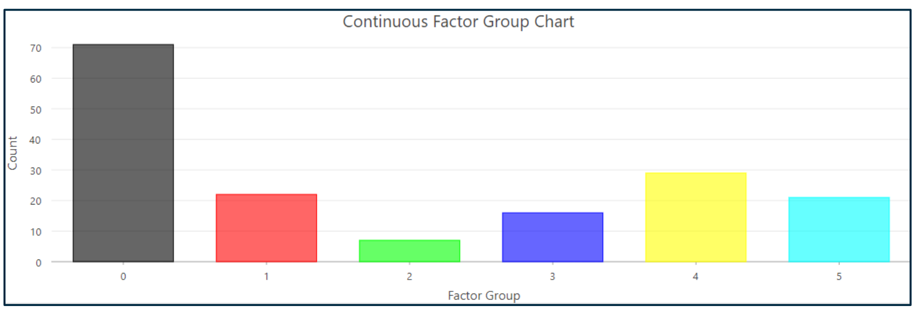

Figure 25. Continuous count station count by cluster.

Results from the clustering process were reviewed with SMEs and it was determined that this clustering result was satisfactory. The results designated a cluster that was represented outside the St. Louis and Kansas City areas with consistent Monday through Friday commuter volumes. Another cluster was characterized by unique recreational summer traffic outside of Branson. The other clusters demonstrated similar characteristics but were differentiated by their urban/rural and functional class characteristics.

Classification

The continuous count clusters were used to perform a classification of the short-term count locations based on urban/rural designation, functional classification, and population density.

Classification, also called "supervised machine learning," assigns class membership (also called labels) to individuals. Common examples of classification include email spam filters ("spam"/"not spam") and speech recognition (each spoken word is an individual class).

In this case, classification is used to assign short term count sites to factor groups. First, assumed latent classes to create groups based on continuous count sites (to which we could assign post-hoc labels like "urban commuter arterial"). Then, we use classification to assign short term count sites to the continuous count sites based on similarities in other attributes such as urban/rural location, adjacent population density, and functional classification.

Results are shown in Figure 26 and Figure 27. The color of each symbol indicates which factor group the count site is assigned to.

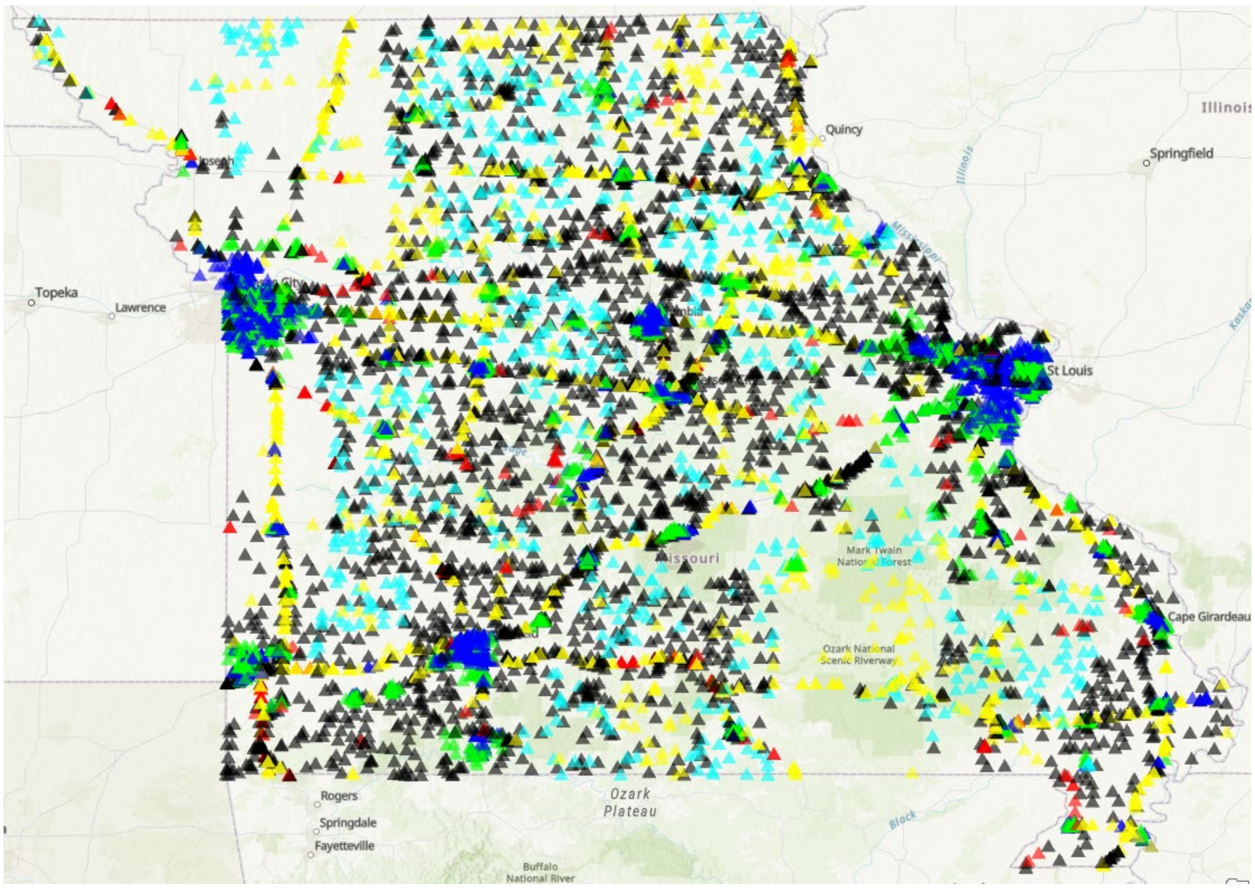

Figure 26. Geospatial short-term count classification results.

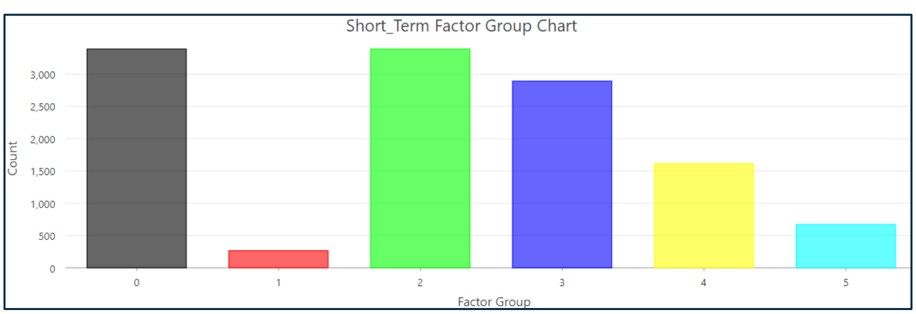

Figure 27. Short-term count station counts by cluster.

Overall, the project team was satisfied with these results. The classification included a grouping of count locations surrounding the larger metropolitan areas, a group for urban areas, another group on high functional classification roadways and clusters for lower functional classification rural areas.

Deliverables

The machine learning workflow and all associated data processing was packaged into an Esri ArcGIS Pro toolbox that could easily be transferred, modified, and updated. The tool was set up with a straightforward file structure such that new data extracts could be loaded into the tool. File structure is shown in Figure 28.

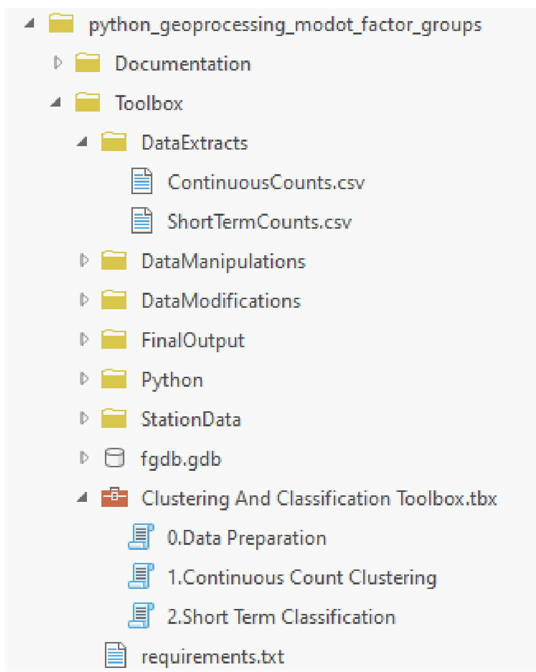

Figure 28. Esri toolbox file structure.

Within the toolbox is a series of Python scripts that can be run in succession that process the raw files for the machine learning steps, perform the clustering, and then the classification.

Project Repository

The toolbox and code for this project can be found in the source code repository:

https://bitbucket.org/high-street/python_geoprocessing_modot_factor_groups/

Ongoing Value

Traffic and travel patterns change over time. Factor groups are an important part of the traffic monitoring process to produce accurate annual average daily traffic estimates from short-term counts. Manually updating and reassigning factor groups to nearly 20,000 short term count sites is a very expensive human task, but a trivial machine task. Going forward, MoDOT can use this machine learning tool to update its factor group assignments every few years to ensure the most accurate annualized averages possible.

Chapter 5. Peer Exchange

During its deployment of AI and ML-based approaches to safety and traffic operations analysis, MoDOT staff encountered several challenges they could not find easy solutions for. Since documented solutions were not broadly available given the newness of these innovative approaches, MoDOT staff identified the need to discuss these challenges with peers at other agencies who may have encountered or worked through similar circumstances. The agency therefore convened a roundtable of transportation agency AI/ ML practitioners who would be able to share experiences of applying AI and ML in a public sector transportation agency. Key discussion topics MoDOT practitioners wanted to cover included both the technical parameters around predictive analytics and data quality, along with organizational factors like staffing, licensing, and getting stakeholder buy-in to apply AI and ML outputs. Across all topics, participants were invited to share their experiences through formal presentations or informal accounts so the group could identify common challenges, effective solutions, and agreement on the institutional changes and support needed for broader and transformative application of AI and ML to solve transportation problems.

Planning for the peer exchange spanned three months and was where the majority of effort for the task was focused. The research team took a structured approach to planning the event to ensure that all needs of the sponsoring MoDOT staff were met by the resulting event. Planning fell into five main categories:

1. Peer Exchange Purpose
2. Participant Identification and Coordination
3. Agenda and Content Development
4. Logistics
5. Findings

Peer Exchange Purpose

Without a clear vision of what needs to come out as a final product or findings, events like this can meander and easily lose focus. The team therefore began by discussing the most crucial outcomes of the event with the sponsoring MoDOT staff to clarify and confirm desired products or findings, preferred format, logistical considerations, and potential participants known to MoDOT staff already who would help to meet the vision. This was also a chance for the facilitator to understand the nature of specific challenges encountered at MoDOT that insights on would be particularly helpful for. Coming out of this meeting or meetings, the team summarized the key elements to ensure that all parties could see and agree on the goals and work toward the same end under the same assumptions.

Participant Identification and Coordination

Once the team had a clear idea of the topics of interest and what MoDOT staff wanted to achieve, research on suitable participants could begin. Using the initial contacts provided by MoDOT combined with web research and the team's existing network of practitioners, the team developed an initial list of practitioners likely to have AI- and ML-related experience related to the topics of interest to MoDOT staff. The team developed an informational flyer and sent preliminary emails to these contacts to explain what MoDOT hoped to achieve and solicit their interest in attending.

Since neither MoDOT nor the team knew what the level of interest would be, how many suitable practitioners would be found, or whether we had identified the right individuals, the group decided to use a survey with which to gather information from interested practitioners. This would allow MoDOT to review everyone expressing interest and make sure they had experience that would be relevant to the event's purpose. It would also allow the team to better control the number of final invitations should there be more interest than the budget could support. The target number of participants was 10-15 individuals from other agencies, plus interested MoDOT staff who would travel using funds from their own offices' budgets.

There were 11 responses to the survey from out-of-state practitioners, but only some of the participants were able to travel due to restrictions, and a few did not have exactly the right experience. The team reached out to all those with appropriate experience and ability to participate to formally extend an invitation and confirm their attendance. Four participants were identified from the survey.

Since there was more availability for participation, the team continued research and networking to identify additional participants interested in and suitable for the event. Several more participants were identified and confirmed this way. In addition to the team's efforts, MoDOT's traffic management staff shared the event information with a national mailing list of their peers, from which an additional two participants responded with interest. In the end, there were nine out-of-state participants in addition to interested MoDOT staff. A full table of participants and their agencies is below.

Participant	Agency
Gene Donaldson	Delaware DOT
Thien Tran	Colorado DOT
Steve Cohn	Colorado DOT
Jeremy Dilmore	Florida DOT
Alexandra Lopez	Florida DOT
Matt Haubrich	Iowa DOT
Juan Hernandez	Nevada DOT
Garrett Schreiner	Minnesota DOT
Bing Wang	Virginia DOT

Participant	Agency
Alex Wassman	Missouri DOT
Rick Zygowicz	Missouri DOT
Marc Lewis	Missouri DOT
Ploisongsaeng Intaratip	Missouri DOT
Tawanda Bryant	Missouri DOT
Kelly Alvarez	Missouri DOT
Michael Middleton	Missouri DOT
Jenni Hosey	Missouri DOT
Ryan Hale	Missouri DOT

Table 3. MoDOT AI-ML roundtable participants.

Agenda and Content Development

The most relevant topics coming out of the initial purpose discussions became the basis for developing the agenda. The survey sent to potential participants also included a request for the topics of most interest to cover, which were summarized and integrated into the initial MoDOT-specified topics. From this range of issues, the team generated four overarching categories that would become the main sessions. These included:

- Predictive Analytics
- Building Trust in AI/ ML for Agency Decisions
- Staffing, Implementation, and Organizational Challenges
- Improving Data Quality

The final agenda is available in Appendix A at the end of this report.

The team worked with MoDOT staff to develop a series of thought-provoking questions on each of these topics to facilitate the discussion in each. These discussion questions would form the bulk of the content for the event, but there was still a strong desire to share practices and lessons from other agencies. Four presenters were identified from the out-of-state participants so that each session was kicked off with a deeper dive into a specific application of AI/ ML at another DOT.

To capture the whole of all participants' experience, the team invited confirmed participants to create a one-page profile of their background, relevant AI/ ML experience, and personal get-to-know you information. These profiles were compiled into a comprehensive participant book so that all attendees could learn about their peers before and during the event.

The research team developed an overarching presentation to set the stage for the discussions, including a big-picture vision of what AI and ML application could do for DOTs in the future. In addition, the team developed guidelines for the four session presenters to align their content as much as possible to the time allotted and focus areas of interest. Participants submitted their presentations in advance for review and compilation.

Logistics

An in-person gathering of like-minded professionals is a rare opportunity for individuals normally siloed in their own agencies to connect with peers and expand their networks. For this reason, special attention was given to organizing logistics to allow for socializing. The team worked with an outside logistics service to assist in identifying the best accommodations and meal vendors/ locations to allow for this while also minimizing costs and increasing overall convenience. The final hotel was selected due to its below-per-diem rate, availability of a shuttle, inclusion of breakfast, and proximity to external dining locations. Several restaurants were identified within walking distance or a short drive for the two evening meals. An affordable catering company was selected to deliver lunches to the meeting location each day to reduce unnecessary travel during the event. On the whole, participants commented on the good sociable atmosphere and quality of the food that allowed for in-depth conversations with peers and forging of new connections.

Findings

A summary of the most significant, high-level takeaways from the presentations and discussions are below. A full summary of the event is available in Appendix A: Peer Exchange.

Predictive Analytics

- Allow sufficient time to vet options and improve predictive solutions.
- Messaging for predictive analytics is important yet challenging.
- Cybersecurity remains a concern.
- Validate using both quantitative and qualitative approaches.
- Use and understand underlying metrics.

Building Trust for Decision-Making

- Education is essential for non-practitioners.
- Find open-minded allies.
- Monitor the broader discussion and watch for preconceived notions.
- Show the value.
- Success can come from the top or bottom.
- Define the problem first – don’t start with the AI “solution.”
- Understand realities on the ground.
- Consider bias and ethical responsibilities to build external support.
- You may need to manage up.

Staffing, Implementation, and Organizational Issues

- Full in-house AI/ ML staffing is rare. Address so many projects at a time.
- Most agencies rely on contractors.
- Learn to hire for the right skills.
- Build multidisciplinary teams.
- Supplement staffing through universities.
- Training agency staff for AI has limited success.
- General training for AI-adjacent roles is helpful – if you can get people’s interest.

Data and Vendor Issues

- Look behind the metrics.
- Consider the ability to customize software.
- Understand geographic, temporal, and other limitations.
- Insist on fixing existing issues before rolling out new features.
- Understand the contracting arrangement.
- Learn from peers’ experience with similar vendors.

Chapter 6. Cost Effectiveness of AI and ML Methodologies

To evaluate the effectiveness of AI and ML methods as applied to MoDOT business practices, the project team reviewed the costs of the pilot projects as it would relate to continuing with status quo approaches. It was expected that these data analytics projects would have large up-front costs and extremely low operating overhead between IT needs and minimal human involvement compared to large indefinite human involvement often seen in traditional solutions. AI and ML projects best derive their value from automating routine processes, improving prediction accuracy to aid in decision making, and removing subjectivity.

Median Inventory

The median inventory pilot had a cost of $80,000. It was determined that to manually create the same inventory by hand would take a GIS analyst an entire year (2,080 hours). The current median salary for a GIS analyst at MoDOT is $61,110. At straight face value, it would be more cost effective to not use AI/ML for this specific inventory.

Further, it was determined that given the 93 percent accuracy of the ML model, the output from the AI/ML process could not serve as a median inventory. Given this, it was also determined that to find and correct the seven percent error would take between six to nine months. This would further increase the cost of using the ML model approach.

As for updating the inventory it was determined that whether the ML approach or manual approach was used to generate the original inventory, it would be more cost effective to update the inventory by hand. Thus, the continued costs would be the same.

In this case, the training effort was high, and the scale of application was limited. Computer vision algorithms require a large training set produced through laborious hand-labeling. Once trained, the model is highly scalable, e.g. given available imagery, expanding the geographic extent from the southern half of Missouri to the include all of the adjoining states would have only marginally increased the amount of effort required. In this case, since the area of application was only a few thousand miles of roadways, and since this is a one-off inventory (not something repeated daily, weekly, or annually), the scale of the problem was actually relatively small (something a human could do in less than a year, and which only needs to be done once), yielding a poor value proposition. Once trained, the model was only applied to 10,000 images. If, instead, the model was applied to 1,000,000 images (or applied to 10,000 images every week) the value proposition would be far better. This reinforces the notion that machine learning only begins to provide a strong value proposition when the number of inferences or predictions grows into the tens of thousands to millions.

Additional Considerations

While the cost of the ML process for this use case did not prove to be cost effective, it should be noted that there was still value derived from this approach. The inventory generated is still of high enough quality to use for some department use cases. For example, the spatial file could be used for safety analysis, or to calculate costs associated with mowing or maintaining median facilities.

Further, the source code and pipeline could be used to train additional models (e.g. to delineate forests or wetlands) at a significantly lower cost. Costs would mainly be associated with image labeling and then rerunning the ML code. The ability to iteratively scale this way over even just a handful of additional spatial assets of interest would tip the favor to the AI/ML approach.

AADT Factor Grouping

The cost of developing the AADT Factor Grouping pilot was $45,000 for the ArcGIS toolbox. Currently, MoDOT essentially assigns factor groups based only on functional classification and the cost to do this from scratch would be minimal. However, if MoDOT desired to develop a factor grouping system based on performance it was determined that this would take at least two years of staff time. At an average cost of $30 per hour this effort would cost $124,800.

In its current form, the AADT Factor Grouping pilot would require some additional work to plug in to MoDOT's existing system and could cost between $10,000 and $30,000.

After that, the ongoing costs for assigning short-term count locations each year would likely be at least an order of magnitude lower for the automated ML process. The ML process takes seconds, or at max minutes, to complete versus many hours for a human to assign thousands of short-term sites to performance groups.

The machine learning process automatically assigns labels to 19,828 short-term count sites. This process, ideally, is repeatedly annually. Over a five-year period, this machine learning application, which was relatively inexpensive to create and deploy, will make 100,000 assignments. It would not be cost effective to have a human traffic analyst update Missouri's short-term count factor group assignments each year, but it's very cost effective to do so using a machine learning approach.

Additional Considerations

In addition to the straight cost comparison of these approaches, the AI/ML approach of assigning to clusters should have additional intangible benefits compared to the current approach of simply assigning count stations based on functional classification. This could include:

- More accurate AADT estimations stemming from the new factor grouping approach could have ramifications on future construction projects and ensure the correct infrastructure is put in place (number of lanes, safety treatments, etc.)
- The automated process would free up staff to work on more human intensive activities.

Chapter 7. Recommendations and Conclusions

Many technologies do not deliver their full benefit until the structures in which they are deployed are remade to take advantage of their capabilities. The electrification of factories is one such example. Eighteen years after the first introduction of electric motors into factories, fewer than five percent of factories used them. As BBC Business (Harford, 2017) explains:

> *Steam-powered factories had to be arranged on the logic of the driveshaft. Electricity meant you could organise factories on the logic of a production line. ... You couldn't get these results simply by ripping out the steam engine and replacing it with an electric motor. You needed to change everything: the architecture and the production process.*

It took decades before electricity fully transformed manufacturing. As factory machines each gained their own intendent electric motor and factories were rearranged to become long, lateral assembly lines, electricity finally transformed manufacturing.

Artificial Intelligence and Machine Learning are likewise in their infancy. Simply replacing an existing process with machine learning model will likely find few cost-effective applications. Human beings excel at reasoning but this faculty scales poorly. Though improving by the day, machine learning reasoning is still quite frail and limited compared to human cognition—but is almost unlimited in its ability to scale. For example, a major part of the recent breakthroughs in large language models like GPT 4 is using millions of books and a large portion of the content on the internet as a training source—a volume of information that not even a group of scholars could ingest in a lifetime.

If follows that great machine learning implementations will by identified not by observing existing processes arising from conditions of scarcity but rather by imaging entirely different structures that might exist where resources for decision-making and prediction were abundant.

AI and ML are gradually being integrated throughout the department. These technologies are perhaps most visible in traffic operations where advanced AI is being used to improve operations, but also in less visible ways, such as the use of dictation to send emails or text messages, use of ChatGPT or other generative AI technologies to answer questions in lieu of an internet search, and behind the scenes in enterprise decision support systems such as asset management and safety analysis systems.

While there are some promising early applications of machine learning, the true transformational potential of AI and ML will not be realized by substituting algorithms into existing processes, but rather by imaging entirely new processes and structures to take advantage of the capabilities of AI and ML.

Demonstration Projects

The pilot projects executed through research demonstrate some of the capabilities of AI and ML.

AADT factor group demonstrates the automation of a routine process. It also provided a most nuanced way to assign traffic factor groups while eliminating the large amount of human time that would be needed for this exercise otherwise.

The median inventory project demonstrated how aerial (or horizontal) imagery can be used in combination with computer vision AI/ML techniques to aid in the development and maintenance of DOT asset inventories. These processes can be used in lieu of repetitive human collection techniques that can be costly and expose agency staff to safety risks.

Staffing and Procurement

Like most DOTs, Missouri DOT has limited internal capacity to implement or develop bespoke AI and ML solutions. The skillsets required for doing so are not typically part of a civil engineering training or typical DOT roles. Most DOTs do not have defined roles for data scientists or machine learning engineers, and in many cases state DOT pay bands to not correspond to prevailing market wages for the specialists with these skills. Strategies for DOTs to access AI and ML capabilities include establishing new internal positions, providing training opportunities for existing staff, or establishing a committee than can review and evaluate the costs and value provided by third-party solutions.

Where AI and ML are being applied within DOT, it's most often in the form of a service provided by a vendor who specializes in their particular application of AI or ML.

As demonstrated through this project, it's possible for the DOT to contract for services implementing custom AI and ML solutions to specifically address a particular need.

Technology Environment

The Python and R open-source programming languages are by far the most commonly used platforms for machine learning. The Information Systems policies governing the development of custom software and algorithms prescribe the use of Microsoft's C# programming language and .NET framework. While enterprise grade machine learning applications can be implemented in C#, the ecosystem of algorithms and tools is far more limited than Python. Under the current IS environment, the most likely path for procuring bespoke machine learning in the future would be for a contractor to develop and prove out a prototype using the language or tool of their choice (most likely Python) and then implement the solution into production using C#. This approach is viable but will meaningfully reduce the number of contractors capable of bidding on such a project and would likely increase total project costs considerably versus an IS environment where Python was an approved technology.

Similarly, many distributed systems for machine learning implementation run best (or in limited cases, exclusively) on Unix-based operating systems. Early coordination with Information Systems is necessary to ensure that solutions are selected and implemented in a manner that is compliant with MoDOT's enterprise I.T. environment.

One key project lesson learned is that Information Systems should be involved early and throughout any AI/ML project to ensure compliance with department standards.

Cost Effectiveness of AI and ML Solutions

The cost effectiveness of AI and ML solutions hinges on a few key elements:

1. **The quality and accuracy required.** Generally speaking, the higher the precision required, the more the ML implementation will cost. These higher costs arise both from requiring greater training costs (e.g. needing a larger and more detailed set of training data) and likely invocation of more sophisticated (see: complicated and expensive) algorithms. When a good estimate is good enough (e.g. planning level cost estimates), machine learning can be a great tool. When absolute accuracy and precision are required, ML may be unable to provide the required level of accuracy or achieving that accuracy may be very costly.
2. **The scale of application.** The findings of this research suggest that 10,000 repetitions of a decision are a good rule of thumb for the minimum scale at which machine learning becomes cost effective. Due to the cost of human decision making, few decisions within an agency are currently exercised at this scale. The promising areas for application are those where, if the marginal cost of an additional decision or prediction was much lower, many more decisions or predictions would be made. For instance, an algorithm could easily predict costs for multiple scopes of work for thousands of segments of pavement to prioritize asset preservation spending. If each project cost estimate needed to be performed by an engineer, it wouldn’t make sense to produce thousands of cost estimates, and a different prioritization approach might be used.
3. **Goodness of fit**. Machine learning is good at predicting quantities and assigning class labels. Processes that result in a quantity and a label are much more easily implemented through machine learning than decisions that are subjective or qualitative.

In summary, the department would expect to obtain good value from a custom machine learning application where the desired decision or prediction is clear and quantitative; where robust training data is readily available; and, where the resulting algorithm would be used hundreds of thousands or millions of times.

References

Google. *Artificial Intelligence (AI) vs. Machine Learning (ML): How Do They Differ?* Accessed March 28, 2024. https://cloud.google.com/learn/artificial-intelligence-vs-machine-learning

Yousefpour, Negin, Downie, Steve, Walker, Steve, Perkins, Nathan, Dikanski, Hristo. 2021. *Machine Learning Solutions for Bridge Scour Forecast Based on Monitoring Data*. National Academy of Sciences: Transportation Research Record. Accessed September 21, 2022. https://journals.sagepub.com/doi/10.1177/03611981211012693

Weiss, Todd. *Machine Learning Helping Indiana DOT Save Millions of Dollars by Bundling Expensive Road Projects*. 2021. Accessed September 21, 2022. https://www.enterpriseai.news/2021/05/03/machine-learning-helping-indiana-dot-save-millions-of-dollars-by-bundling-expensive-road-projects/

Harford, Tim. *Why didn't electricity immediately change manufacturing?* 2017. Accessed July 13, 2023. https://www.bbc.com/news/business-40673694

Appendix A: Peer Exchange

AI & ML Roundtable

Date: June 27 and 28, 2023
MoDOT St. Louis District TMC Building, Room 209

Event Agenda

Day 1 – Full Day

Missouri AI/ ML Roundtable Agenda: Day 1	
8:30-8:45	Welcome
8:45-9:15	Big Thinking with AI: How can AI *transform* our work?
9:15-9:45	Introductions
9:45-10:15	Opening Presentation • MoDOT's AL/ ML Experience • Purpose and goals for the event
10:15-10:30	Break with coffee
10:30-11:00	**Participant Presentation:** Thien Tran, Colorado DOT
11:00-12:00	Discussion Session 1: Predictive Analytics
12:00-1:00	Lunch provided
1:00-2:00	**Activity:** Group Photo & TMC Tour!
2:00-2:30	**Participant Presentation:** Jeremy Dilmore and Alexandra Lopez, Florida DOT
2:30-3:00	**Participant Presentation:** Steve Cohn, Colorado DOT
3:00-3:15	Break with snacks
3:15-4:15	**Discussion Session 2:** Building Trust in AI/ ML for Agency Decisions
4:15-4:30	Wrap-Up
4:30-6:00	Personal Time
6:00-8:00	Group Dinner - EdgeWild

Day 2 – Morning Only

Missouri AI/ ML Roundtable Agenda: Day 2	
8:30-9:00	**Participant Presentation:** Gene Donaldson, Delaware DOT
9:00-10:00	**Discussion Session 3:** Staffing, Implementation, and Organizational Challenges
10:00-10:15	Break with coffee
10:15-10:45	**Participant Presentation:** Juan Hernandez, Nevada DOT
10:45-11:45	**Discussion Session 4:** Improving Data Quality
11:45-12:00	Wrap-Up
12:00	Departures - Boxed lunch provided