Beyond Pavement: Scaling AI for Infrastructure Management

A Google Cloud Pilot Project for Automated Road Sign Inventory

MPH Data Day - May 14, 2025

Indiana Department of Transportation (INDOT)
Bryan Chisholm - Data Science & Visualization Manager
Warren Wang - Senior Data Scientist
Matthew Paradise - AI Product Owner

The AI Approval Process

Including the AI Readiness Assessment

Matthew Paradise – AI Product Owner

The Approval Process

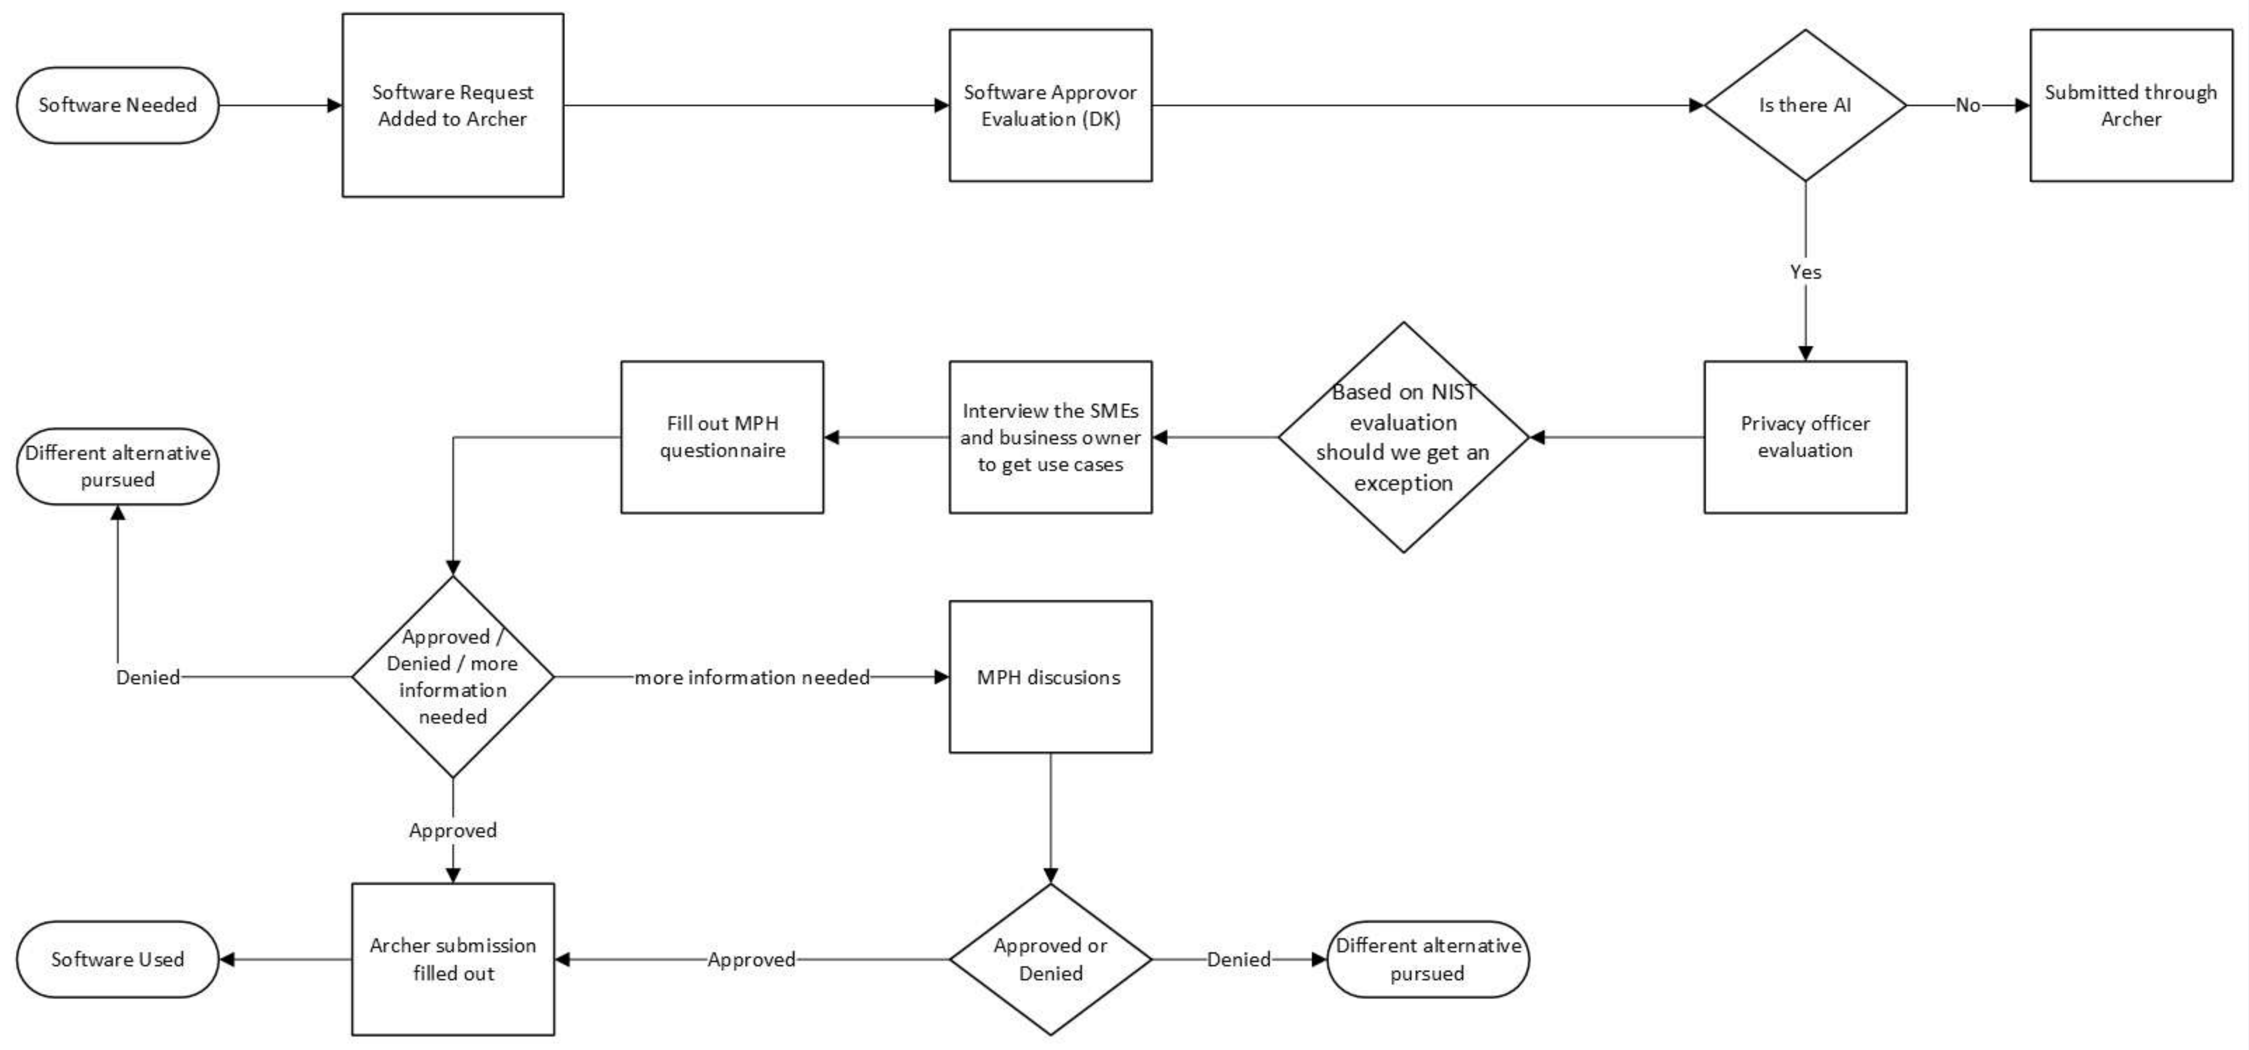

The Approval Process - Overview

- **State AI Policy Adoption**
  - https://www.in.gov/mph/AI/
- **Balancing AI Benefits and Risks**
  - https://www.nist.gov/itl/ai-risk-management-framework
- **Guidelines**
  - State Agency Standard for AI
    - https://www.in.gov/mph/cdo/files/State-of-Indiana-State-Agency-AI-Systems-Standard.pdf

Approval Process – Readiness Assessment

- **AI Readiness Assessment Requirement**
  - The OCDO issued the State Agency Artificial Intelligence Systems Standard.
  - Agencies must complete an AI Readiness Assessment before implementing or using AI tools.
  - A Readiness Assessment Questionnaire is required before AI deployment.
  - The MPH AI Review Team, led by the CPO, reviews assessments to ensure ethical and privacy standards.

Approval Process – Readiness Questions

- **Main AI Readiness Assessment Questions**
  - Will you be using the AI
  - Description and use case
  - Is there a substitute
  - Third party?
  - What domain
  - What type of AI
  - Intended audience
  - What data and where
  - What environment
  - Residence
  - Maintenance

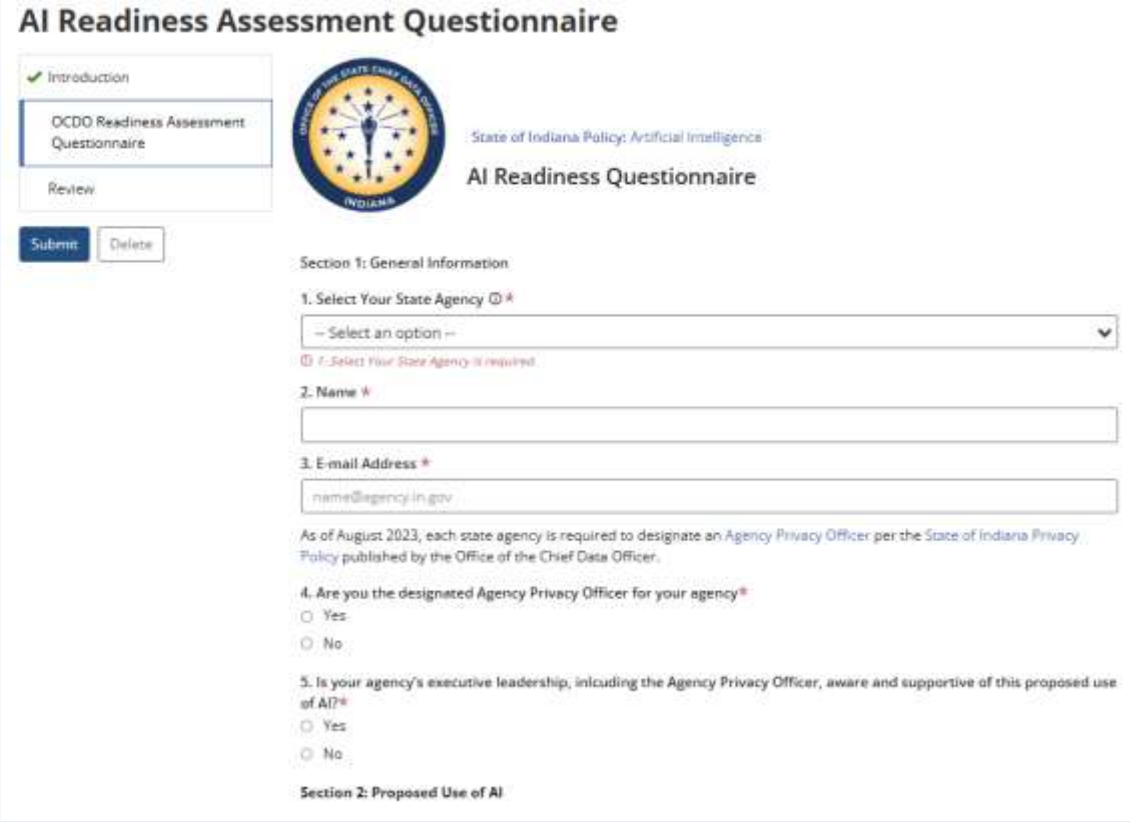
AI Readiness Assessment Questionnaire

Introduction

OCDO Readiness Assessment Questionnaire

Review

Submit Delete

State of Indiana Policy: Artificial Intelligence

AI Readiness Questionnaire

Section 1: General Information

1. Select Your State Agency *

-- Select an option --

1. Select Your State Agency is required

2. Name *

3. E-mail Address *

name@agency.in.gov

As of August 2023, each state agency is required to designate an Agency Privacy Officer per the State of Indiana Privacy Policy published by the Office of the Chief Data Officer.

4. Are you the designated Agency Privacy Officer for your agency*

Yes

No

5. Is your agency's executive leadership, inlcuding the Agency Privacy Officer, aware and supportive of this proposed use of AI?*

Yes

No

Section 2: Proposed Use of AI

Approval Process - INDOT

- **SharePoint**
  - Clone of MPH questionnaire
- **Business user meeting**
  - Go over questions
- **Submission to MPH**
  - Dialog begins
- **Optional**
  - MPH meeting

Approval Process - Results

- **Denial**
  - Terms
  - EULA
  - Failure of standards

- **Out of Scope Exception**
  - Not using the AI

- **Full Approval**
  - Approved for use following guidelines

Beyond Pavement: Scaling AI for Infrastructure Management

A Google Cloud Pilot Project for Automated Road Sign Inventory

MPH Data Day - May 14, 2025

Indiana Department of Transportation (INDOT)
Bryan Chisholm - Data Science & Visualization Manager
Warren Wang - Senior Data Scientist

The Manual Inventory Bottleneck

- MUTCD Asset (traffic sign) Identification on roadways has historically been a manual effort
  - Labor intensive
  - Costly
  - Hard to maintain up to date
- With current technologies, we can alleviate much of the manual efforts required with AI/ML assistance
  - Save Cost/Time
  - Scalable
  - Adaptable to engineer's need

An AI-Powered Solution: Automated Road Sign Inventory

- MPH-Approved AI Use Case Pilot Project
- Goal: Automate the detection, classification, and inventory of road signs
- Leveraging Google Cloud Platform (GCP) for scalability and cost-effectiveness
- Focus area: I-65 & I-70

Building on Google Cloud Platform

- Scalable data storage and processing
  - JTRP (Purdue) test environment
- Pre-trained and custom AI models.
- Serverless deployment for efficiency

From Images to Insights: Data Collection

- Image Source: Vendor (Pathways) – Vehicle-mounted cameras
- Available Metadata/Annotation:
  - Location on road network (measure) or GPS
  - Camera focal length
  - Equipment configuration setup

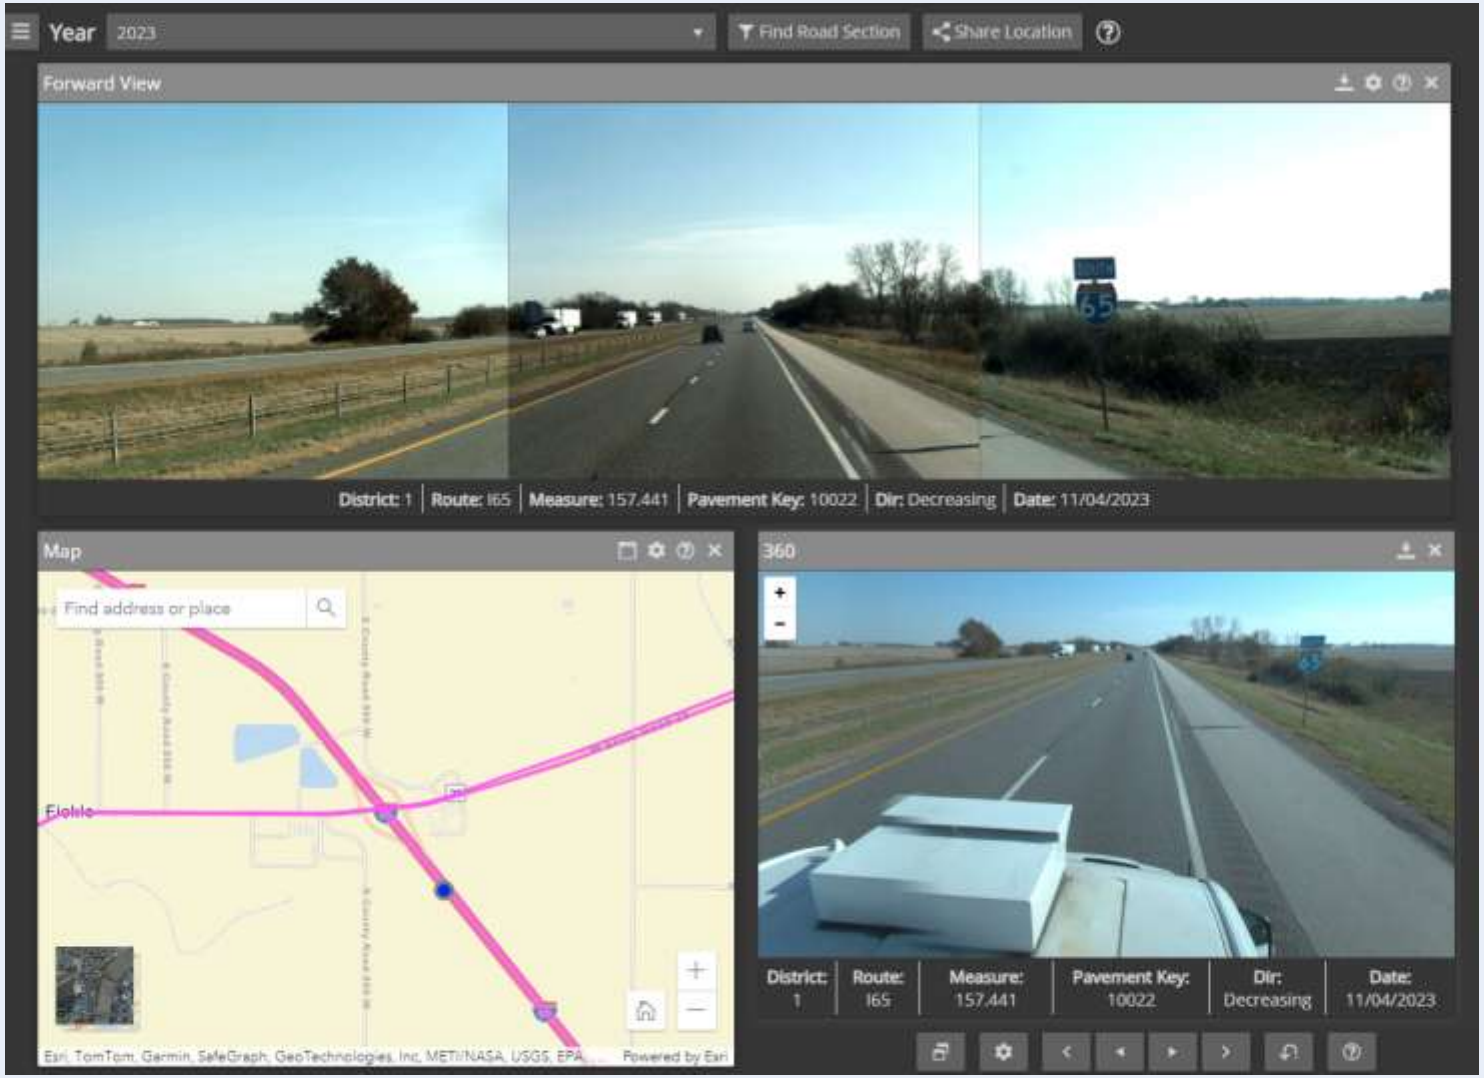

Can you find all of the signs in this video?

From Images to Insights: General Approach

- Split MUTCD into 2 main categories:
  - Text-based (identifiable by combination of unique attributes)
  - Symbolic (not identifiable by combination of unique attributes)
- Feed prompt into Gemini
- Output attributes for each identified signs, with their bounding boxes (BBOX)
  - Sign Text
  - Sign Color
  - Sign Shape
  - Sign Bounding Box
  - Sign Description

From Images to Insights: Data Preparation

1. Split images into 3 Parts - better BBOX coordinates
2. Align file directories into real world measurements
3. Assign MUTCD or save results for data labeling for training

From Images to Insights: Split the task in two

MUTCD Assignable via Text

MUTCD Not Assignable via Text

MUTCD Assignable via Text

Gemini

Sign Text : WATCH FOR ICE ON BRIDGES
Sign Color: YELLOW
Sign Shape: DIAMOND
Sign Bounding Box: [256, 256, 512, 512]

Rule based function

W8-Y8

- Performant
  - 332,853 panoramic images batch processed in 61 minutes 26 seconds
- "Cheap"
  - ~ $10K for processing 5.5 million panoramic images.

MUTCD Assignable via Text

Bounding Box "magically" drawn around sign (i.e. it’s the computer)

MUTCD auto assigned

MUTCD Not Assignable via Text

1. Label Data

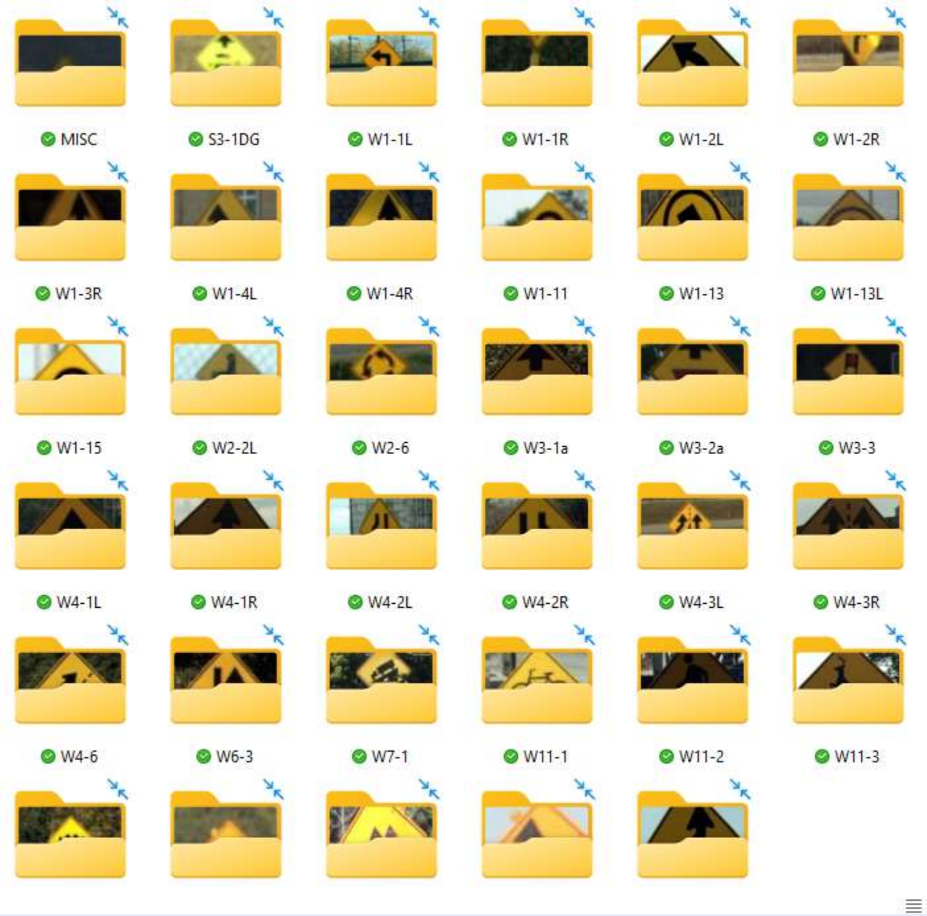

2. Create Dataset

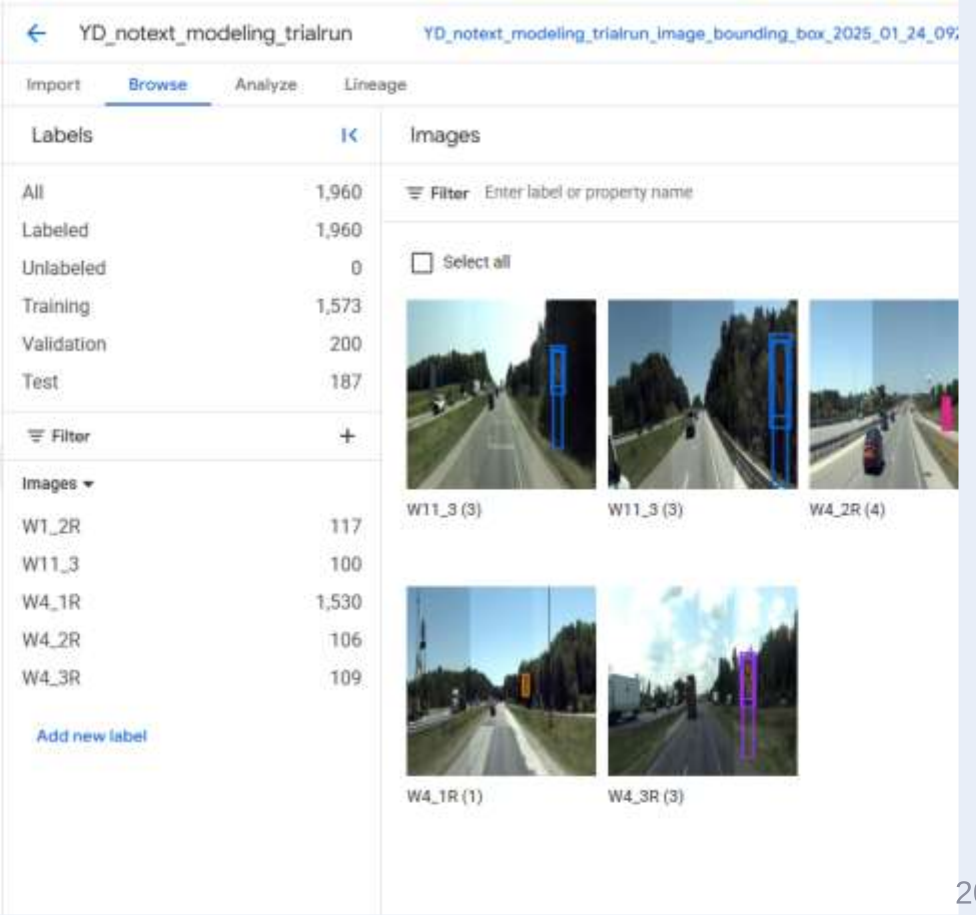

GCP AutoML: Model Training and Validation

- Reserved for symbolic signs
- Utilized Google AutoML for custom object detection.
- Iterative training process to optimize performance
- Key Metrics: Precision, Recall
  - 100% Precision sits at 0.749 confidence threshold with a 29.091% Recall rate
  - Steep Drop in Precision happened at 0.243 confidence with 91.845% Precision and 77.818% Recall

Precision vs recall explanation: You use a wide net, and catch 80 of 100 total fish in a lake. That's 80% recall. But you also get 80 rocks in your net. That means 50% precision, half of the net's contents is junk. You could use a smaller net and target one pocket of the lake where there are lots of fish and no rocks, but you might only get 20 of the fish in order to get 0 rocks. That is 20% recall and 100% precision.

INDOT SSD: Model Training and Validation

- Reserved for symbolic signs
- From scratch implementation of Single Shot Detector (SSD) neural network architecture
- Written in Python
  - Portable to any platform
  - Can be run locally

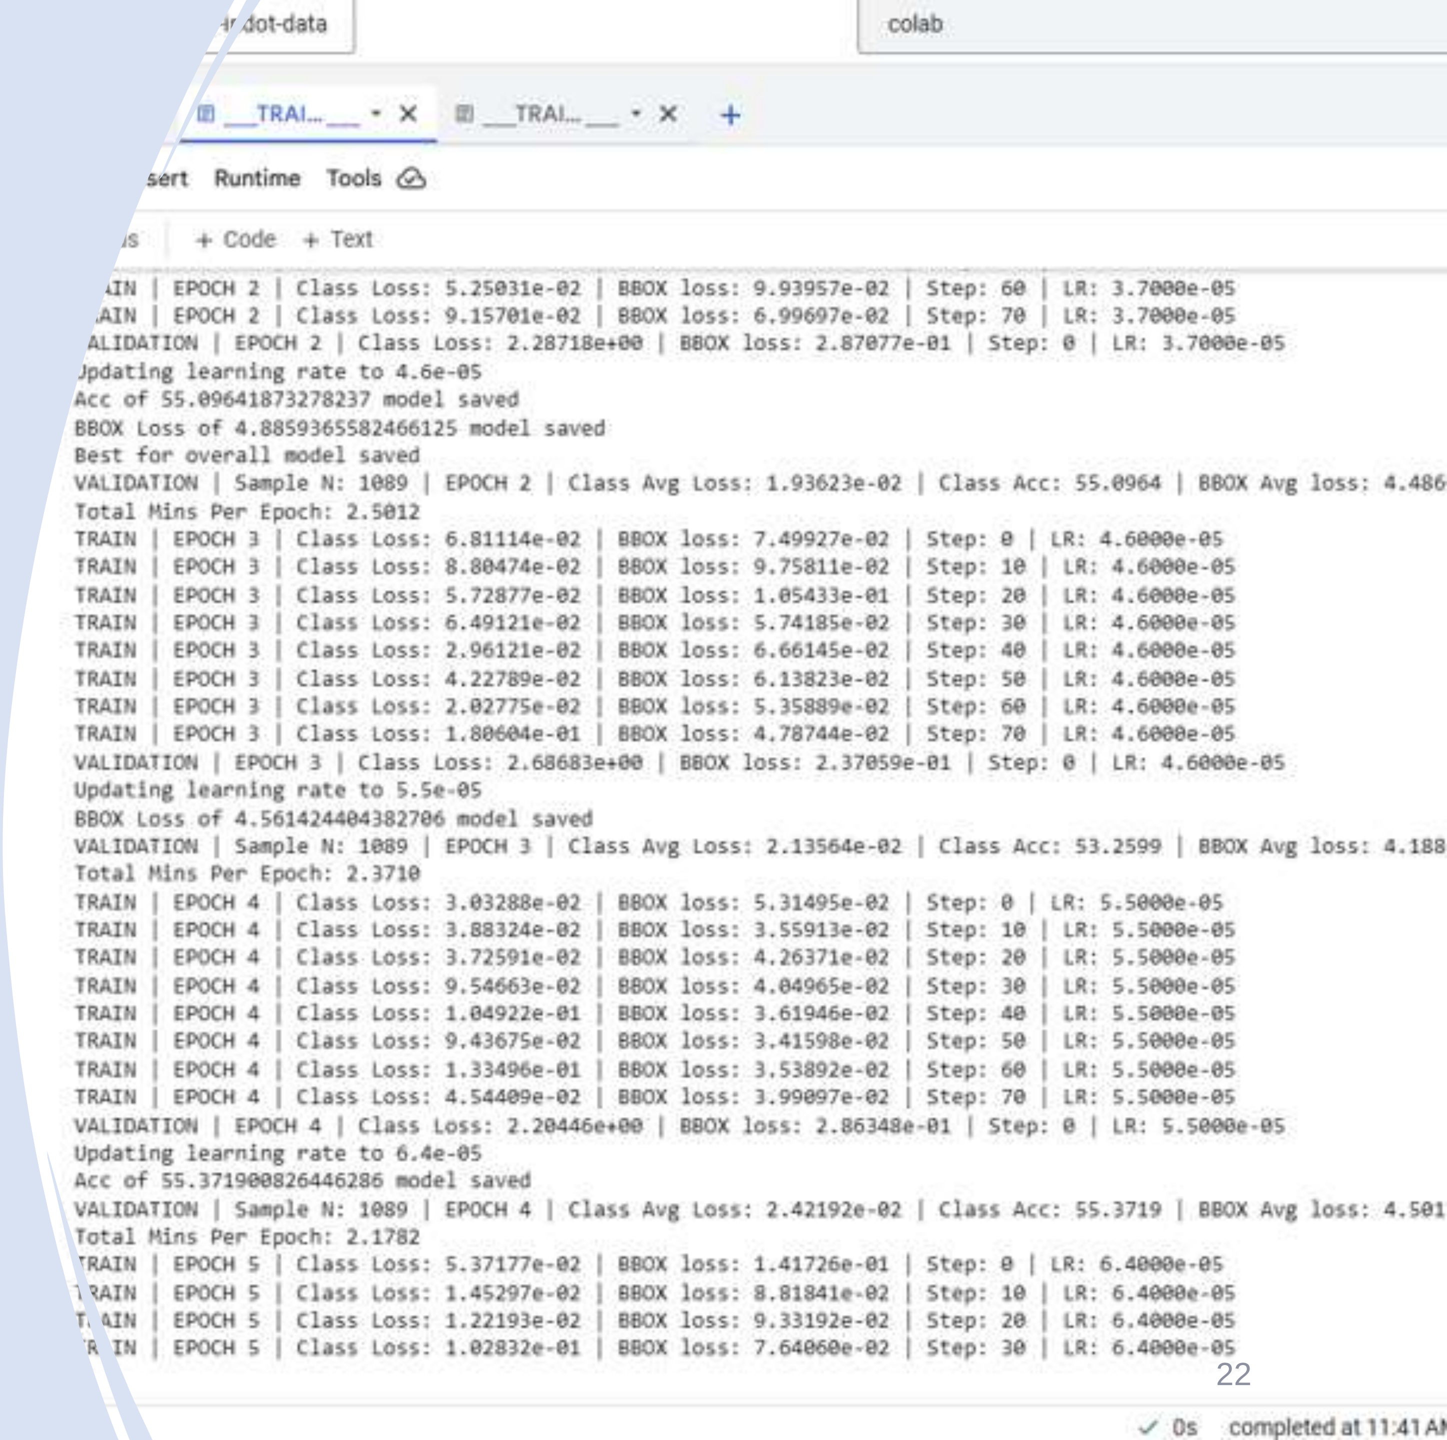

Pilot Project Success

- "Text" Signs:
  - ~100% precision in identifying the correct MUTCD assets (spot checking over 1,000 + signs identified thus far)
    - 108 + MUTCD codes programmed in thus far

- Symbolic Signs:
  - Manual Validation on Hold Out set (model never saw this in training):
    - 565/570 labeled with correct MUTCD code

Pilot Project Success: Accurate Detection and Classification

Next Challenge: Locating the sign “IRL”

Requirements:

- **Travel Bearing** (Current and next image GPS or stored metric)
- MUTCD sign **real world dimensions**
- **Tight Bounding Box** (pixel dimensions)
- Lots of math to reverse calc sign location based on;
    - Image GPS
    - Sign bearing
    - Sign distance

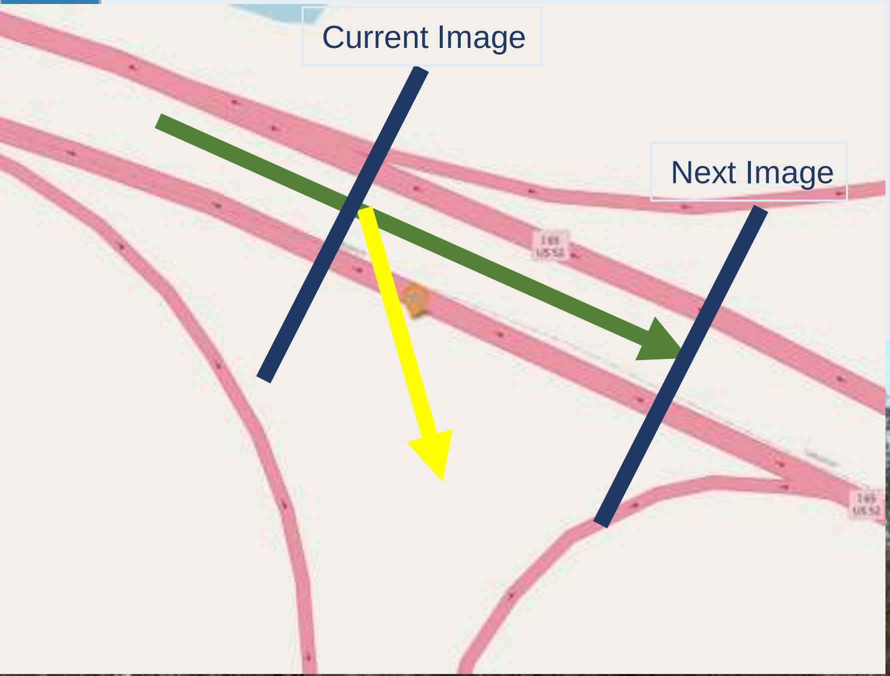

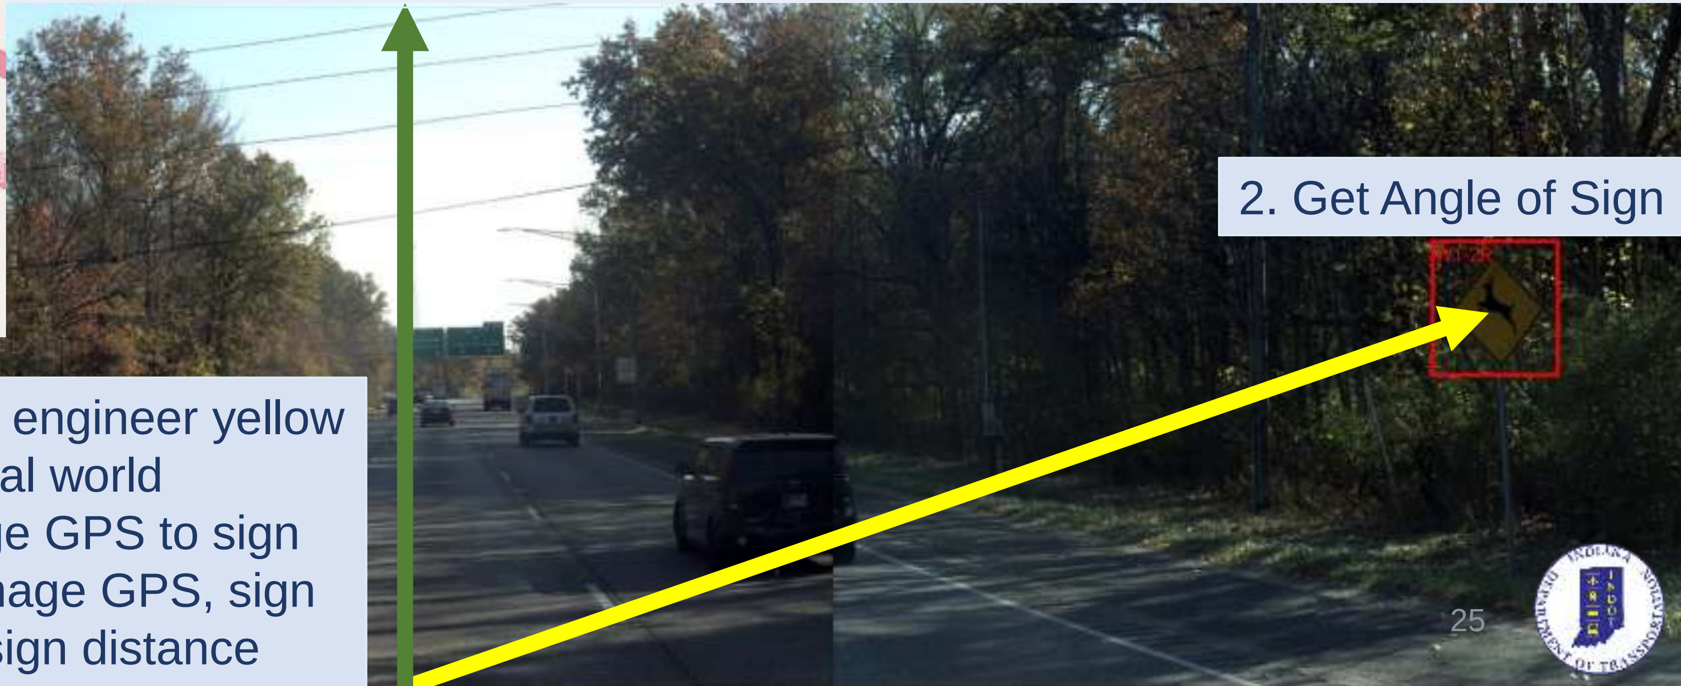

Pilot Project Success: Location from image metadata vs GIS

GIS-based sign inventory location

Estimated location from algorithm

- Sign code identified
- Location calculated from last “observed” sign

Next Steps

- More robust tracking algorithm to track MUTCD signs for the "last" sign
- More accurate sign GPS location
  - Use Pathways calibration files and not BBOX pixel dependent
- Gemini output and SSD model output exploration
  - Improve filtering criteria
  - Expected result vs actual output

Questions?

Indiana Department of Transportation (INDOT)

- Bryan Chisholm - Data Science & Visualization Manager
- Warren Wang - Senior Data Scientist
- Matthew Paradise - AI Product Owner