AI SUMMIT
TOP TECH

This program is sponsored by the APWA Technology Committee

AI Summit

Wednesday, January 29, 2025

- 11:00 a.m. - 2:00 p.m. ET
- 10:00 a.m. - 1:00 p.m. CT
- 9:00 a.m. - 12:00 p.m. MT
- 8:00 a.m. - 11:00 a.m. PT

Thursday, January 30, 2025

- 11:00 a.m. - 1:00 p.m. ET
- 10:00 a.m. - 12:00 p.m. CT
- 9:00 a.m. - 11:00 a.m. MT
- 8:00 a.m. - 10:00 a.m. PT

APWA Proprietary Interest Disclosure Statement

Speakers and instructors for APWA education offerings are practitioners and experts in public works and infrastructure issues and solutions. Some may have proprietary interest in services and products used by public works agencies and professionals.

Speakers for APWA educational offerings agree to refrain from specific product or services endorsements or promotions during APWA educational offerings.

The AI Summit is eligible for .4 Continuing Education Unit (CEU) credits.

- A link to the program evaluation will be emailed to you after tomorrows program.

- To receive CEU’s, you must correctly answer 80% of the questions.

ACCREDITED IACET PROVIDER

APWA has been accredited as an Authorized Provider by the International Association for Continuing Education and Training (IACET), 1760 Old Meadow Road, Suite 500, McLean, VA 22102; (703) 506-3275.

APWA

RESOURCE CENTER

ONE PLACE where a member can access:

- Past Click, Listen & Learn (CLL) presentations
- Content from past conferences (PWX and Snow)

All searchable by topic. All without any extra fees beyond membership dues! No limits to how often you access or open the items in the Resource Center.

No travel. No scheduling hassles. No delays while waiting for someone else’s approval. No waiting for registration payment to be processed.

For more information go online to www.apwa.org/resources/resource-center.

Today's Moderator

Tracy Warner, PE

Deputy Director of Development Services

Pima County

AI SUMMIT TOP TECH

APWA

Polls

AI SUMMIT TOP TECH

APWA

How AI Systems are Developed and Trained

Day 1, Session 2

AI SUMMIT
TOP TECH

APWA

Session Two Speakers

Chris Lareau
Co-Founder, COO
Vue Robotics

Steve Cacioppo
Senior GIS Analyst
Douglas County, NE

AI SUMMIT
TOP TECH

APWA

Learning Objectives

- Describe the engines within AI systems.
- Explain what large language models (LLM) are and how transformer models are part of LLMs.
- Identify the role of data in training AI systems- how much is needed, what quality control must be done, and so forth.

Image Source: Vue Robotics

APWA AI SUMMIT TOP TECH

AI Engines

APWA AI SUMMIT TOP TECH

What is an AI Engine?

An AI engine is a subset of machine learning that processes and analyzes vast amounts of data, recognizes patterns and makes decisions or predictions. It allows the system to learn and adapt to varying conditions and data sets and provide an output.

An AI engine is like a vehicle in a fleet or a specialized tool. Each vehicle shares some similarities but serves a specialized purpose. They can work independently to solve a specific task, or multiple vehicles can work together to solve more complicated problems.

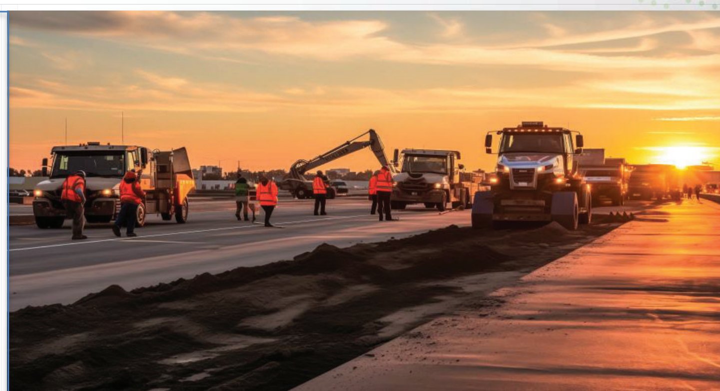

Image Source: Adobe Stock (licensed by Vue Robotics)

APWA AI SUMMIT TOP TECH

Types AI Engines

Natural Language Processing (NLP)
Understand and process human language, both text and speech.
Example: ChatGPT and customer service chatbots.

Generative AI
Create new content, such as text, images, and search results.
Example: AI writing tools generating stories or articles.

Computer Vision Engines
Analyze and interpret visual data like images or videos.
Example: Detecting objects in self-driving car camera feeds.

Deep Learning
Multiple layers of neural network "nodes" are powerful for tasks involving complex, high-dimensional, or unorganized data.
Example: Processing and translating text from images in real-time.

Speech Recognition and Generation
Convert speech to text and text back to speech.
Example: Voice assistants like Alexa or Siri understanding voice commands.
destination.

Machine Learning (ML)
Learn patterns from data to make predictions or decisions.
Example: Recommending movies on a streaming platform based on viewing history.

Autonomous Systems
Enable systems to operate independently in real-time environments.
Example: Drones navigating obstacles to deliver packages.

Decision-Making and Optimization "Experts"
Find the best solution among many possibilities.
Example: Waze or Google Maps determining the fastest route to a

Big Data and Analytics
Process and analyze large-scale datasets to extract insights.
Example: Amazon identifying buying trends from user activity.

Edge AI
Perform AI tasks on devices with limited resources, without relying on the cloud.
Example: Apple Watch analyzing heart rate and step count in real-time.

APWA AI SUMMIT TOP TECH

AI Engine Architecture Diagram

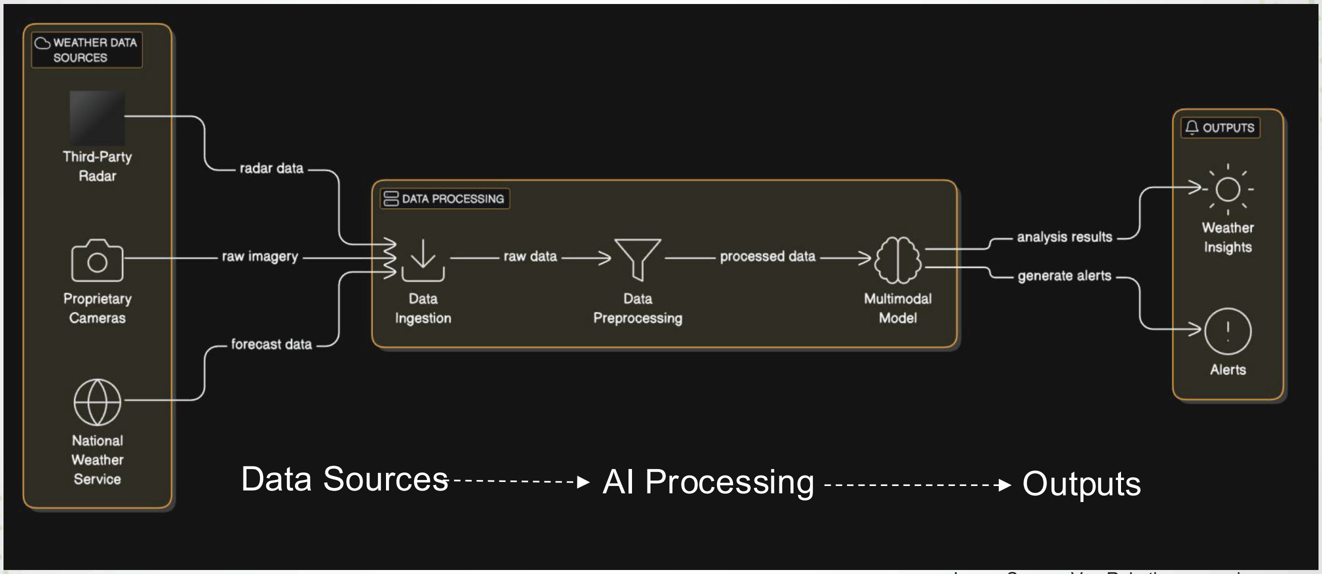

Image Source: Vue Robotics, eraser.io

AI SUMMIT TOP TECH

APWA

AI Engine Example

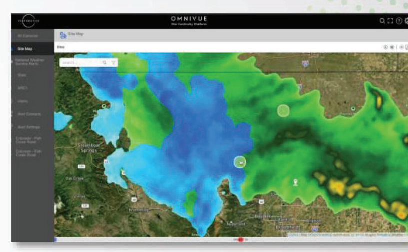

Task-Specific
A computer vision engine for weather prediction focuses on analyzing visual data, such as satellite images, radar scans, or camera feeds, to identify patterns and forecast weather.

Example: Analyzing local camera feeds to determine road weather conditions.

Powered by Models
AI engines use one or more models to process data and produce outputs.

Example: A model detects snow or flooding being present on the serviceable area of a road and classifies the conditions.

Integrated into Larger Systems
Computer vision engine works alongside other weather forecasting engines and engines that live on device (physical hardware) to prove a comprehensive condition analysis and forecast.

Example: The engine identifies flooding patterns which are combined with tidal, rainfall and flood plain data for detailed, predictive forecasting

Images Source: Vue Robotics

AI SUMMIT TOP TECH

APWA

Large Language Models
LLMs

AI SUMMIT TOP TECH

Large Language Models (LLMs)

- A **large language model (LLM)** is a type of artificial intelligence (AI) designed to process and generate human-like text. The "large" in LLM refers to the billions of parameters in the model and the data it has been trained on, which enables it to learn and represent complex language patterns, structures, and contexts.
- Early LLMs were used to simply predict the next word based on the text before it.
- **Transformers** are the underlying architecture of what makes LLM today special. Helps these models more effectively understand the order of words in a sentence, which is crucial for understanding the context and meaning of a sentence.

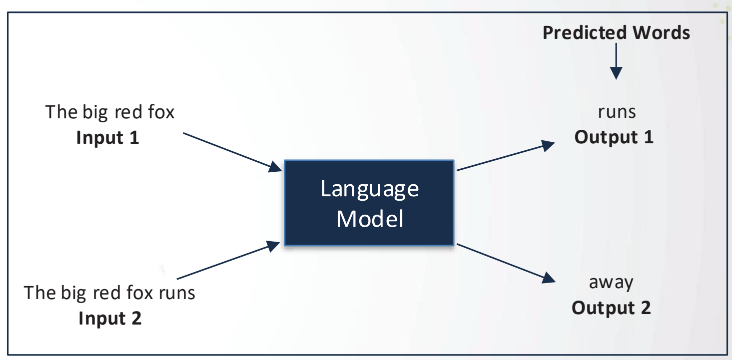

In LLMs, text is converted into numbers, which enables generative and natural language processing models to process them.

Data Source: OpenAI

AI SUMMIT TOP TECH

The Role of Data in Training

AI SUMMIT TOP TECH

Role of Data in Training

- Proprietary data is your greatest differentiator.
- Good data = good models = meaningful outputs.
- Depending on the type of data, sourcing, scrubbing, segmenting, and labeling can still be very manual processes when developing a new mode or tuning an existing one.

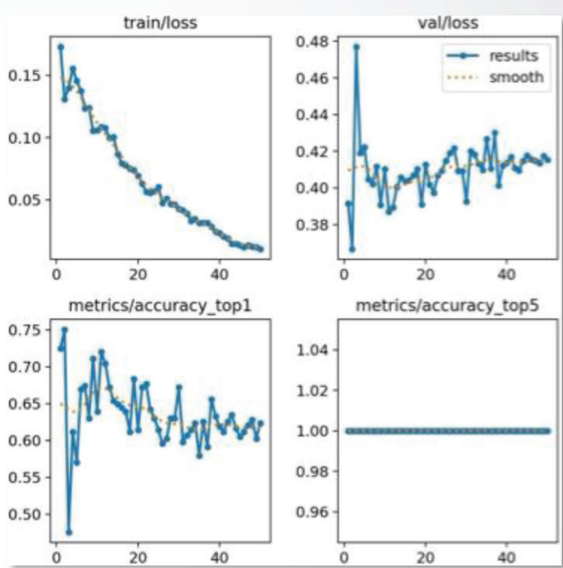

Image Source: Vue Robotics

AI SUMMIT TOP TECH

Detecting Precipitation Example

Data Set Overview

- 1747 Images Labeled
- 45 Unsure (Removed - Bad Data)
- 113 Precipitation (~6.4%)
- 669 Maybe (~38.2%)

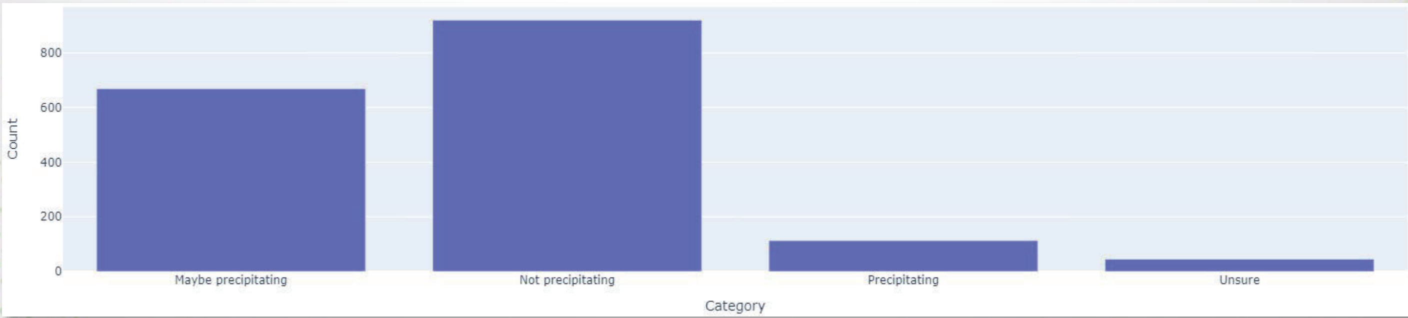

Image Source: Vue Robotics

Ground Truth (L)

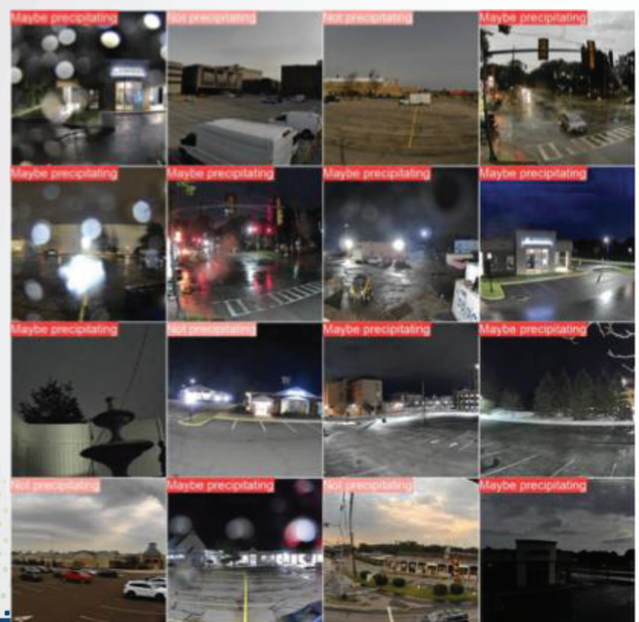

Prediction (R)

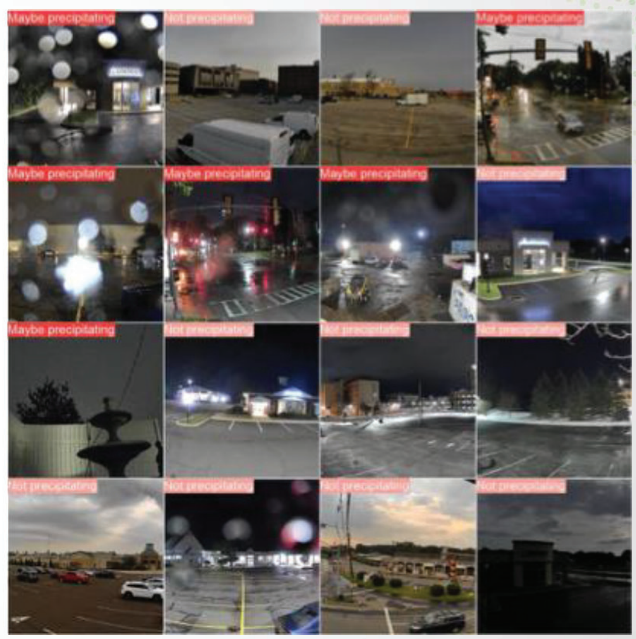

Images Source: Vue Robotics

Ground Truth (L)

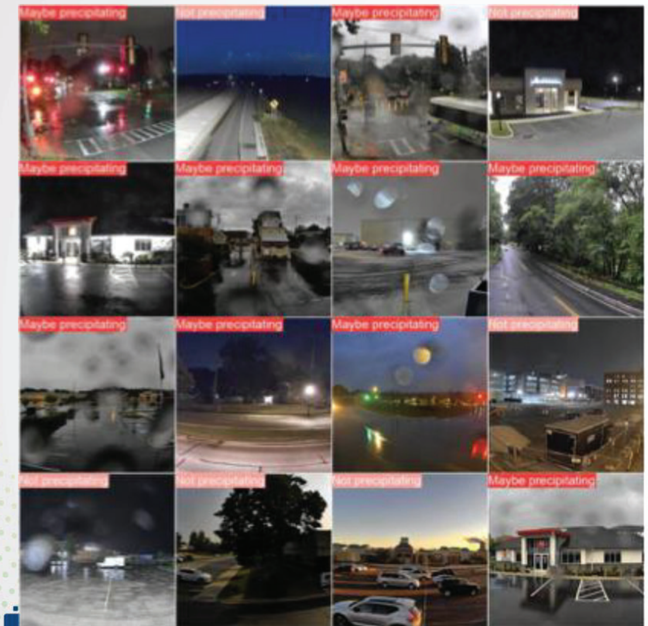

Prediction (R)

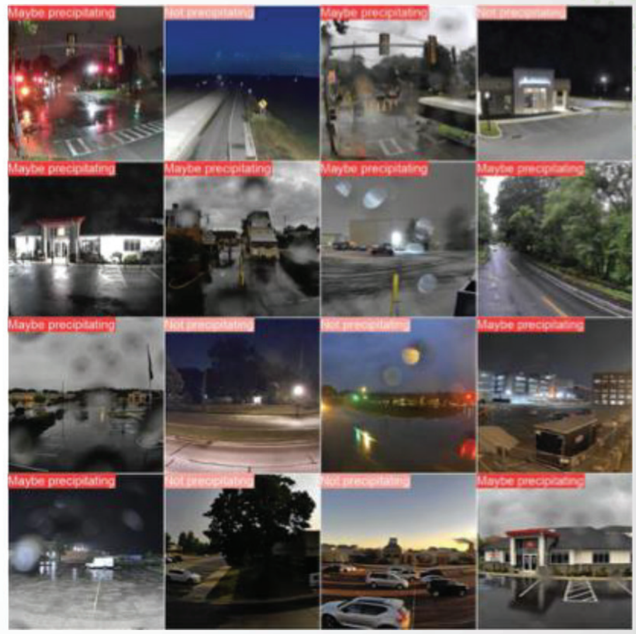

Images Source: Vue Robotics

AI SUMMIT TOP TECH

Ground Truth (L)

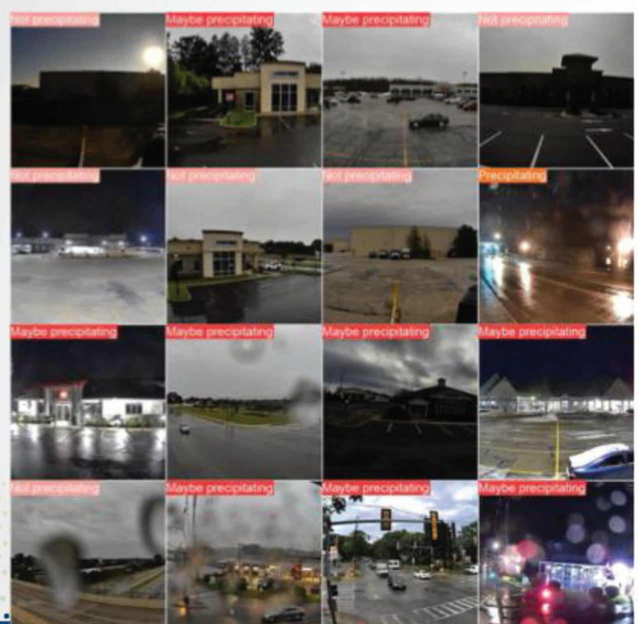

Prediction (R)

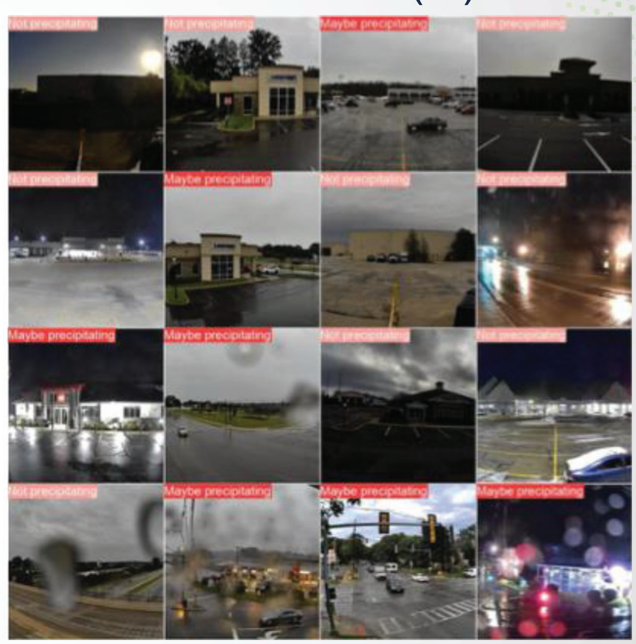

Images Source: Vue Robotics

AI SUMMIT TOP TECH

End Result

Images Source: Vue Robotics

AI SUMMIT TOP TECH

ADA Curb Ramp feature extraction using ESRI's Deep Learning tools

Steve Cacioppo - Douglas County, NE GIS

scacioppo@douglascounty-ne.gov

AI SUMMIT TOP TECH

City of Omaha Construction Division needed assistance identifying all of their ADA Curb ramps

- They had an incomplete and inaccurate inventory of ADA curb ramps
- This might be a good opportunity to try using ESRI's Deep Learning tools
- ADA curb ramps are easy for a human to pick out from an aerial photo, should be easy for the computer too

AI SUMMIT TOP TECH

APWA

Where to begin?

- Not a ton of resources when I began my journey
- ESRI Blogs to the rescue!
  - Sangeet Matthew from ESRI wrote a detailed blog about how he used ESRI Deep Learning tools to identify Caribou from aerial photography in Alaska. I was able to use his methodology to create my own Deep Learning model to extract ADA curb ramps

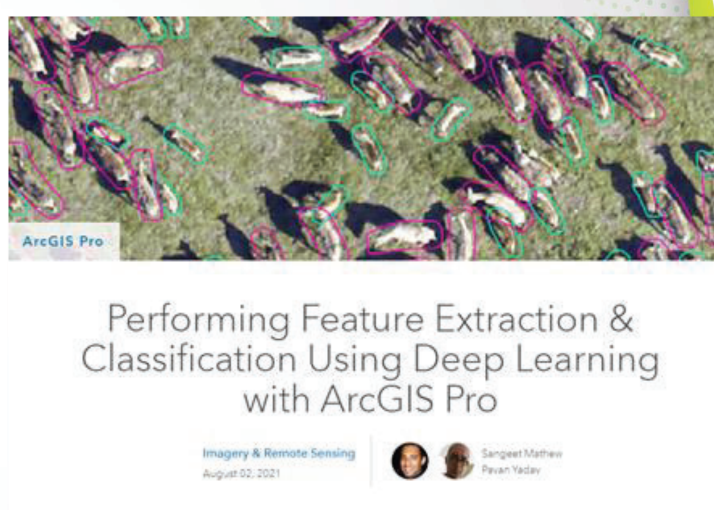

AI SUMMIT TOP TECH

APWA

Basic steps I used for creating my Deep Learning model

- Read Sangeet’s blog and follow his instructions
- Create a LOT of training samples. ESRI recommends 300-500 samples but more is better
- The more training samples you have, the better, more accurate your results will be. I created a total of 753 training samples for my model
- Run your model on a small data set first
- Once you get a model that yields decent results, think about where you should run it

AI SUMMIT TOP TECH

How to run my new Deep Learning model

- My model was trained on our commercial 1” resolution imagery.
- The area of Douglas County, NE is 339 square miles - that a lot of processing!
- Limit the processing to where you would find ADA curb ramps
  - Buffer street centerlines
  - Commercial parcels - have sidewalks & parking lots
  - Parcels that contain apartments & condos
- To speed up processing I also broke up the county into a grid and used an iterator to process the individual squares. Afterwards I appended them into one layer

AI SUMMIT TOP TECH

Results

The results from my model were better than I expected. Whereas I did have some missed ramps and some that were errant, the overall results were good

Some of the errant examples are:

- Inlet grates/manhole covers
- HVAC on commercial roofs
- Sun roofs on autos

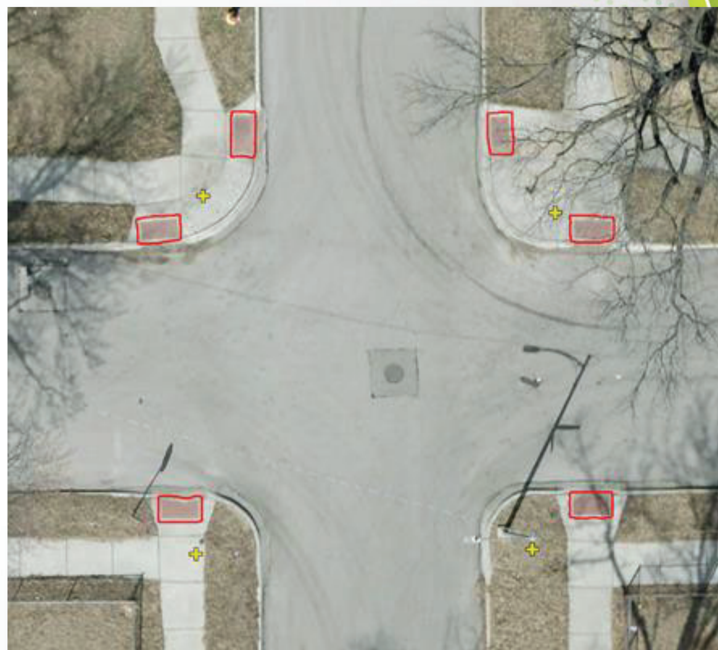

Results - continued

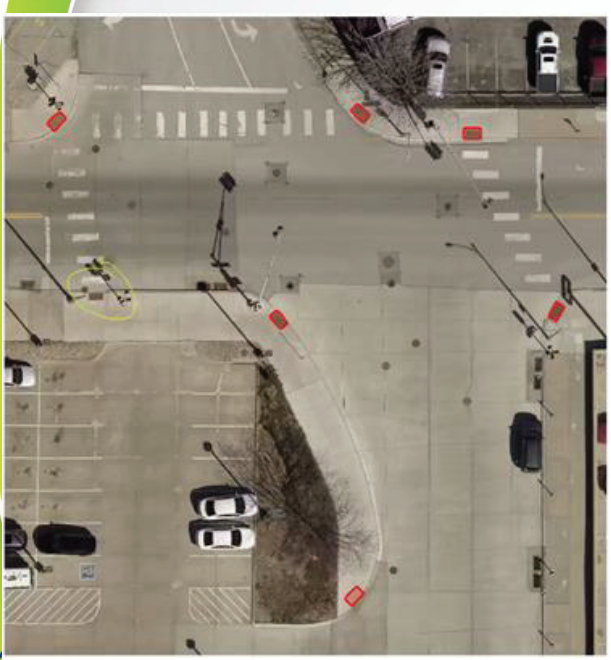

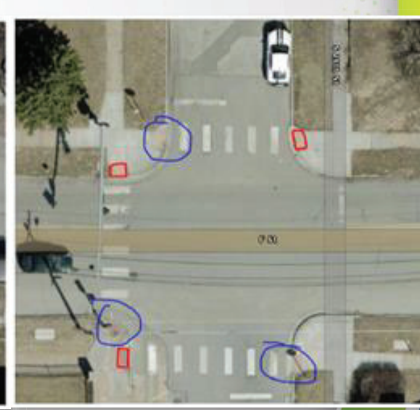

Lessons Learned

- There were lots of ‘technical’ bumps along the way - lots of trial and error with the inputs - stick with it!
- 1” imagery is overkill for ADA curb ramps
  - Next time use 3” Mr. SID imagery
- Use of image server or cloud processing
- There is no ‘one size fits all’ for these Deep Learning models

AI SUMMIT TOP TECH

AI discussions with management

- ChatGPT has given AI a negative connotation with some
- Remember that AI (Deep Learning) is a collaboration between the computer and the human
- The only time there wasn’t human interaction with my process is when the computer was actually running my model
- Computers/AI/Deep Learning are only as good as the data you give them

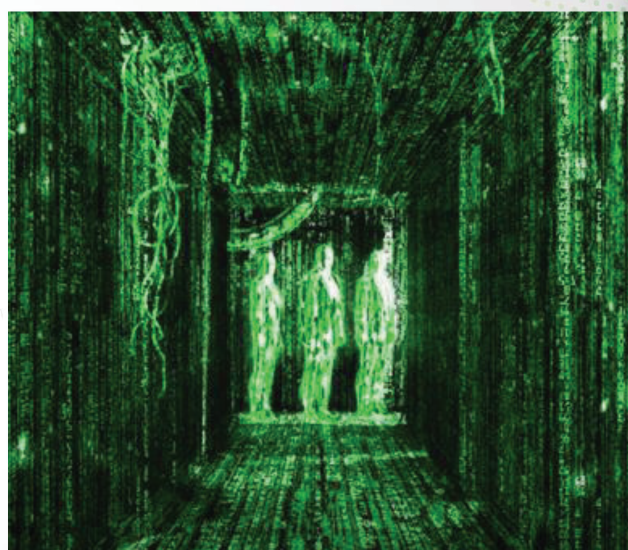

AI SUMMIT TOP TECH

Next Steps

Now that I've successfully created my first model, here's what I want to try and do next:

- Swimming pools
- Pavement crack detection
- Car detection/object tracking
- Trees/species?

In Closing…

ESRI now has 94 pre-trained deep learning models available for download from the ArcGIS Living Atlas. That's up from the 8 pre-trained models they had when I started this journey 2 years ago!

Before you create your own model, download one of these pre-trained models and try it yourself.

If you have any questions, feel free to contact me!

APWA AI SUMMIT TOP TECH