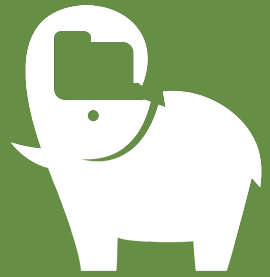

THIRDEYE

Solving Real-World Industry Problems with AI Computer Vision

Agenda

- What is AI ML DL
- Why AI ML
- Traditional Approach vs AI ML
- One Real World Problem
- Google AI Solution Platform
- Technology Stack
- Architecture Diagram
- Data Workflow Diagram
- AI Computation Layers
- The Demo
- AI Model Development methodology
- Real life Deployment
- Q & A

Artificial intelligence (AI)

AI generally refers to processes and algorithms that are able to simulate human intelligence, including mimicking cognitive functions such as perception, learning and problem solving. ML and deep learning (DL) are subsets of AI.

Some of the practical applications of AI are as follows:

- Recommendation systems
- Chatbots
- Self driving vehicles

Machine Learning (ML)

Machine learning is an application or subset of AI that allows machines to learn from data without being programmed explicitly.

The three types of machine learning algorithms are as follows:

- Supervised learning
- Unsupervised learning
- Reinforcement learning

Deep Learning (DL)

DL is a subset of machine learning that attempts to emulate human neural networks. DL algorithms are able to ingest, process and analyze vast quantities of unstructured data to learn without any human intervention.

Practical applications of deep learning includes:

- Computer vision
- Natural language processing

Traditional Approach vs AI ML

Traditional programming - One has to manually formulate or code rules.
ML - One needs to point the algorithm at the data so that it can learn powerful rules and can be used to predict future outcomes.

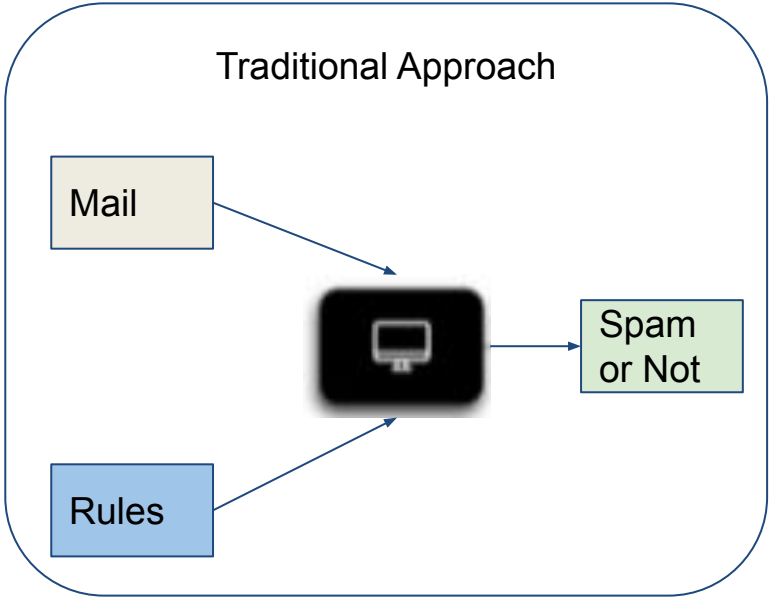

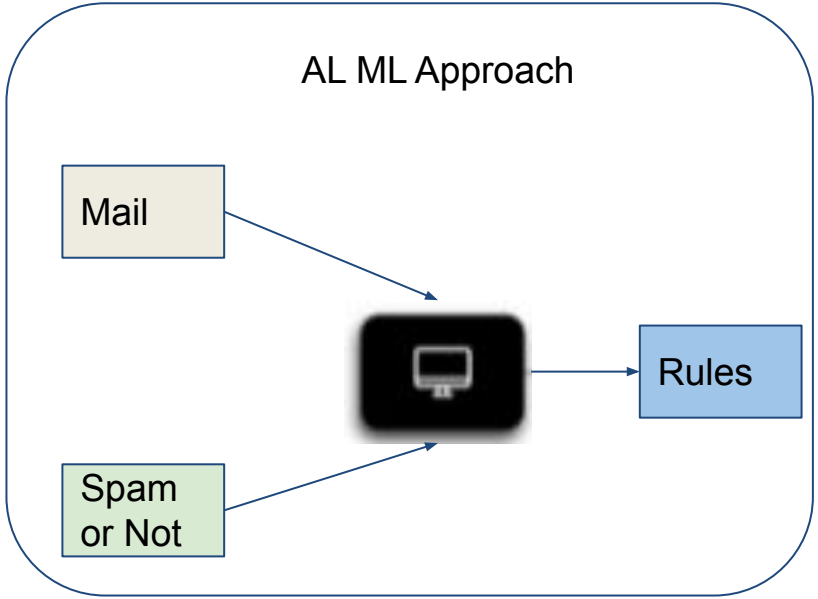

Why AI ML

Intelligent Agent: Need not to develop a decision making program with the set of rules. Using AL and ML we would build an agent that could look into a bunch of data and train itself to figure out the output.

Data Size: Over the time, world is more reliant on computers than ever before.

Cloud Services: Computers have become fast enough to be able to address these very large data sets with computational methods.

THIRDEYE

Real World Problem

Manual Electric Pole Inspection Process is an effort oriented and time-consuming task. It can take several days for humans to inspect more than 150K images of electric poles. The AI can do it in a couple of hours. The AI analyzes the image quality issues of the electric poles and makes predictions.

Google AI Solution Platform

- Vertex AI, Google Cloud's new unified ML platform provides a platform to address the problem.
- Vertex AI is the platform to build, deploy, and scale ML models faster, with pre-trained and custom tooling within a unified artificial intelligence platform.
- In Vertex AI, one can now easily train and compare models using AutoML or custom code training and all your models are stored in one central model repository.
- AutoML enables developers to train high-quality models specific to their business needs with minimal ML expertise or effort. With a centrally managed registry for all datasets across data types (vision, natural language, and tabular)
- It provide end-to-end integration for data and AI

Google AI Solution Platform

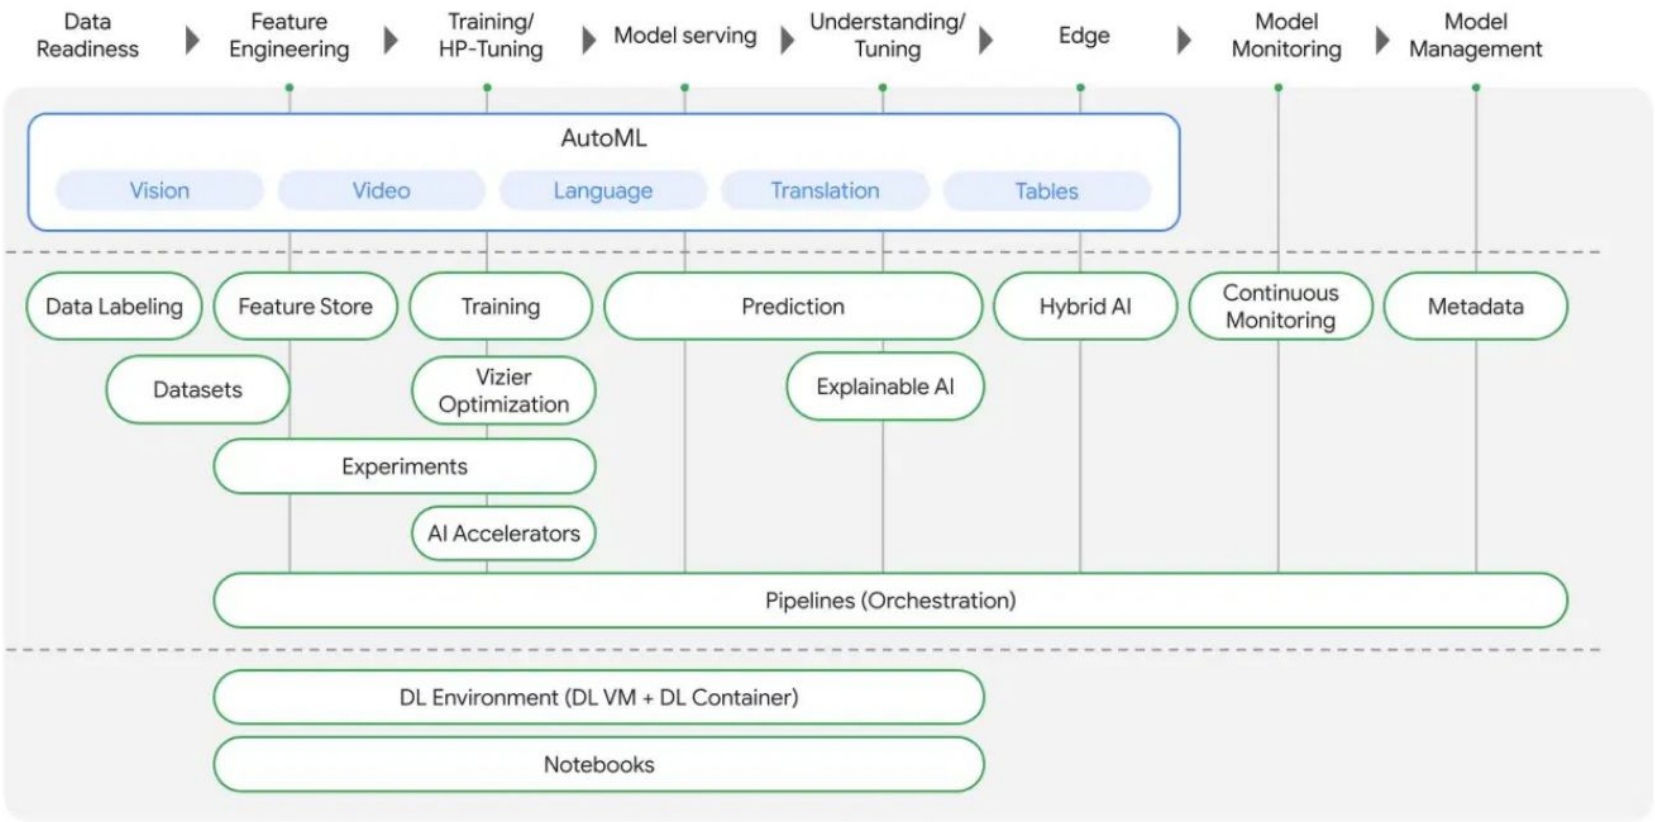

THIRDEYE

Technology Stack

Cloud Platform	Google Cloud
Storage Platform	Cloud Storage
AI Platform	Vertex AI
Programming Language	Python
Model Training	AutoML

End to End Architecture Diagram

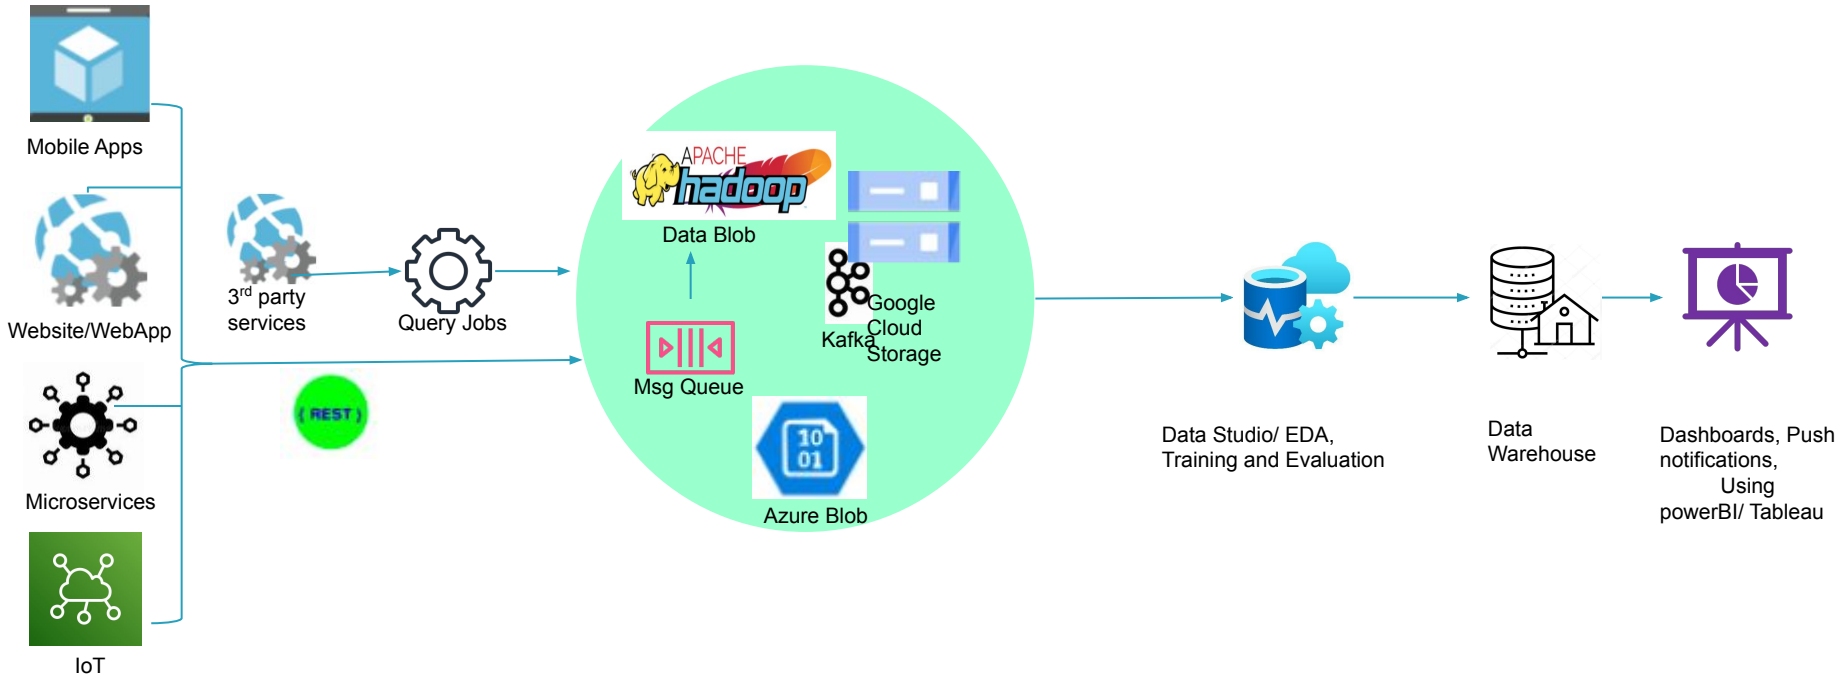

Data Workflow Pipeline

- **Data Collection**: Various Data sources (mobile apps, websites, web apps, microservices, IoT devices, etc.) are instrumented to collect relevant data.

- **Data Ingestion**: All this data gets collected into a Data Lake.

- **Data Preparation**: It is the extract, transform, load (ETL) operation to cleanse, conform, shape, transform, and catalog the data blobs and streams in the data lake; making the data ready-to-consume for ML and store it in a Data Warehouse.

Data Workflow Pipeline

- **Computation**: This is where analytics, data science, and machine learning happen. Computation can be a combination of batch and stream processing. Models and insights (both structured data and streams) are stored back in the Data Warehouse.

- **Presentation**: The insights are delivered through dashboards, emails, SMSs, push notifications, and microservices. The ML model inferences are exposed as microservices.

Data Workflow Pipeline

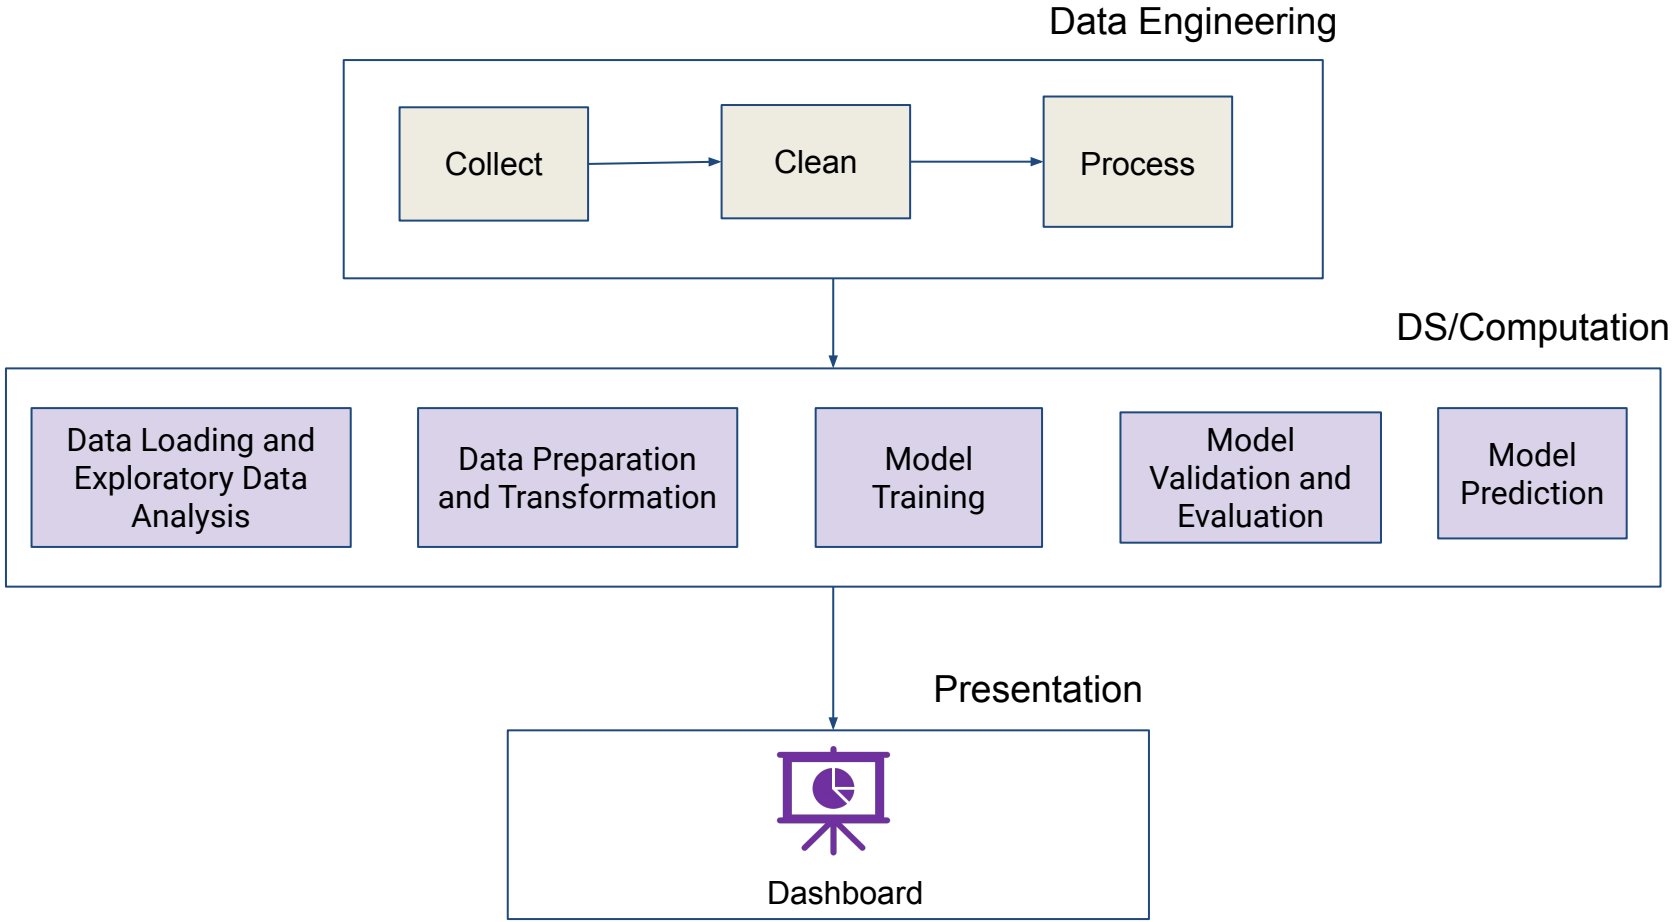

AI Computation Layers

Data Exploration and Model Development

- Define business and technical requirements
- Identify various algorithms as per the definitions
- Perform a lot of experimentation and iteration with all available data sets
- Explore data to understand the underlying correlations between different variables

AI Computation Layers

- **Model Training**
    - Information gathered from the data exploration process will help us to choose the best model
    - Develop the right model by iterating through its parameters and tuning it to get optimal outputs
    - Split the available data into three parts:
        - training,
        - validation, and
        - testing

AI Computation Layers

- **Model Deployment**
    - Validate the models based on both human feedback and outcomes analysis,
    - Deploy the models for downstream consumption
    - Keep checking the model outcomes for model percisicon, accuracy and drift purposes.
    - Retrain as and when needed

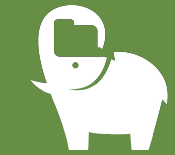

Demo

Demo Steps

1. Data Collection from open source
2. Data Annotation in Vertex AI
3. Train the model using Auto ML or custom model.
4. Import the models
5. Deploy the models to the endpoint
6. Run the object detection prediction script
7. If the below criteria meets, then send the cropped image to Occlusion and Blurriness code.
   If pole is cross arm and detection score >0.5

THIRDEYE

Flow of Process

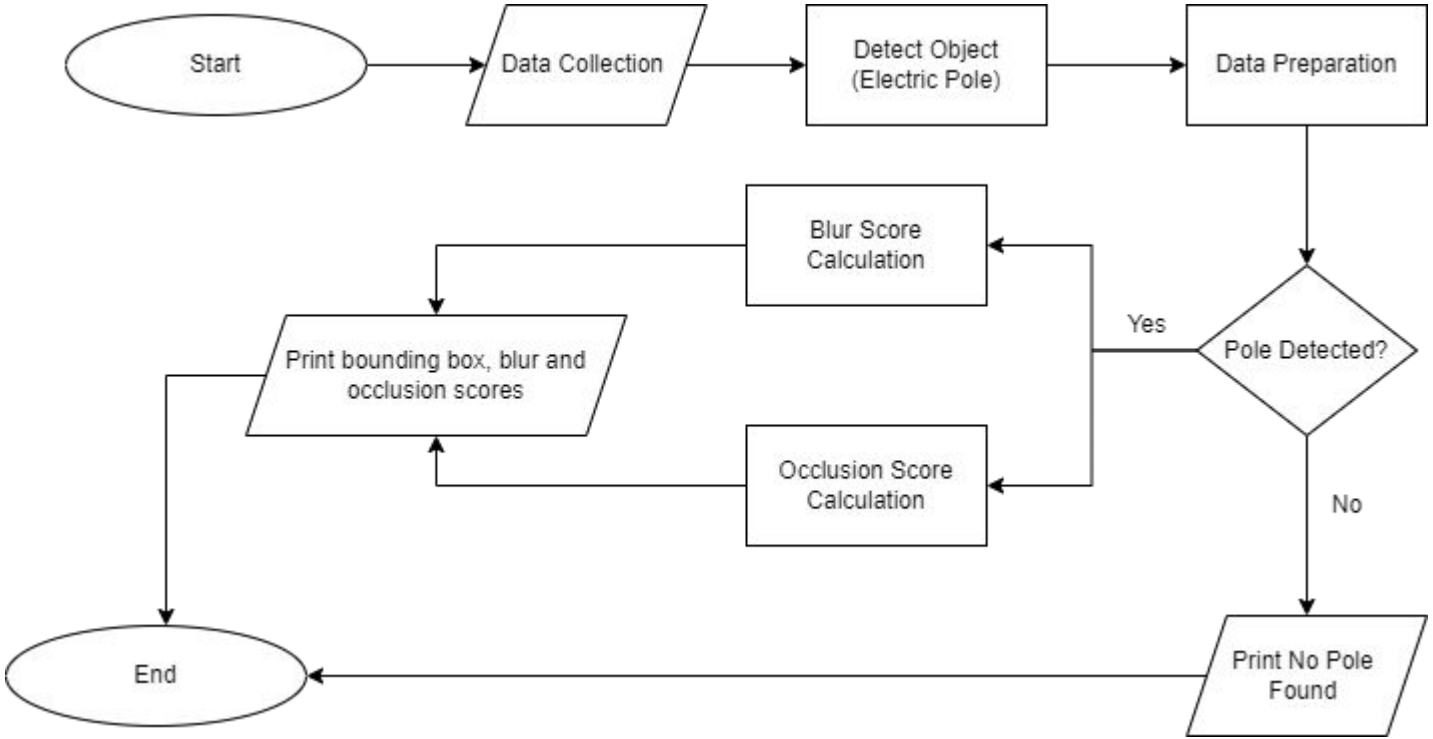

Data Collection

Data collection is done using *Image Downloader* extension

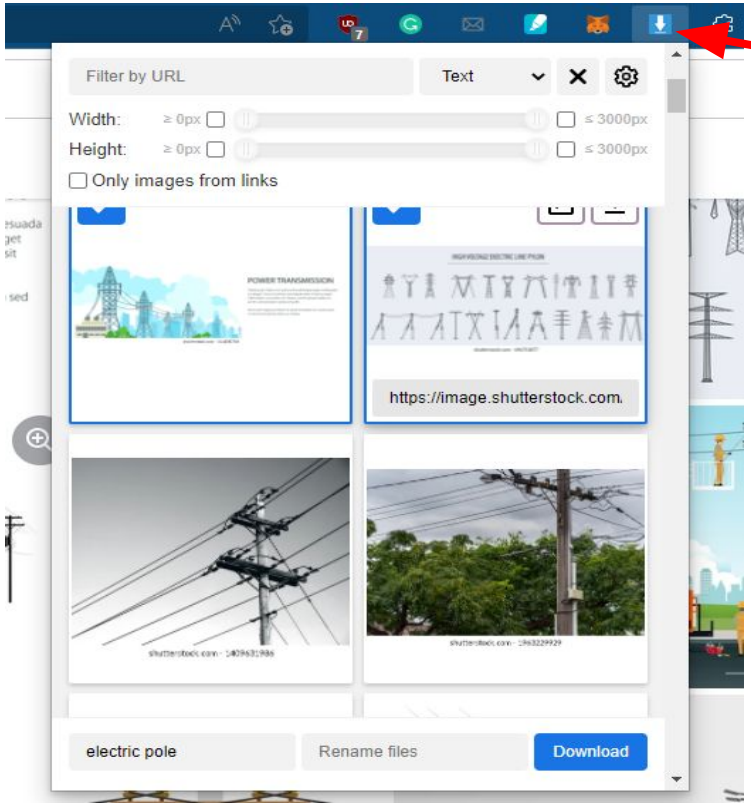

The extension allows us to easily download images from any website, without the hassle of going through 100's of lines of html

Data Annotation on Vertex AI

To annotate data on Vertex AI, we just need to upload the data, and start annotating:

Green lines denote the bounding area

Tag for the image

Evaluating the Model

After training the model, we can see the metrics too, for each label and average metrics too:

All labels

Average precision	0.285
Precision	64.4%
Recall	19.3%
Created	Mar 31, 2022, 8:28:11 PM
Total images	1,015
Training images	817
Validation images	99
Test images	99

To evaluate your model, set the **confidence threshold** to see how precision and recall are affected. The best confidence threshold depends on your use case. Read some example scenarios to learn how evaluation metrics can be used.

Precision-recall curve

Precision-recall by threshold

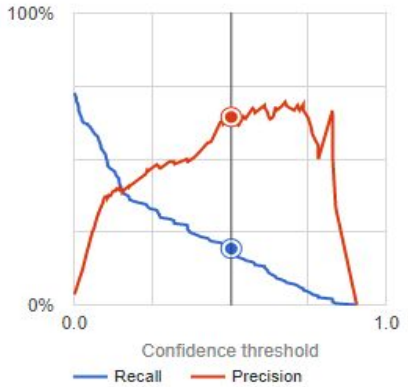

Evaluating the Model

Recall : TP/(TP+FN)
It says, from all the positive classes, how many we predicted correctly.

Precision: TP/(TP+FP)
It says, from all the classes we have Predicted as positive, how many are actually positive.

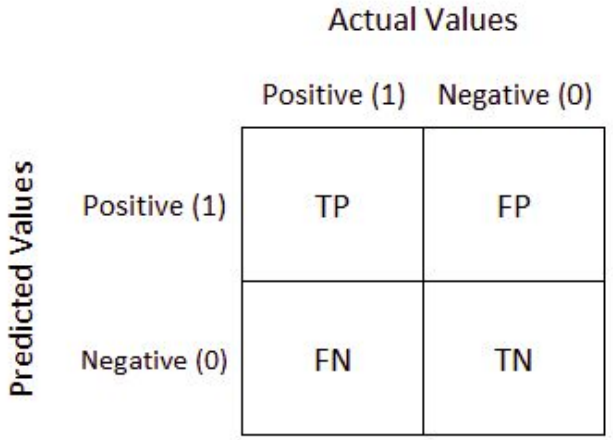

True Positive (TP): *You predicted that a woman is pregnant and she actually is.*
True Negative (TN): *You predicted that a man is not pregnant and he actually is not.*
False Positive (FP): *You predicted that a man is pregnant but he actually is not.*
False Negative (FN): *You predicted that a woman is not pregnant but she actually is.*

Evaluating the Model

We can also see instances for True Positive, False Positive and False Negative

False negatives

Your model should have predicted electricPole for these images:

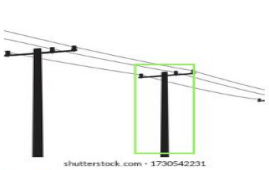

Score: 0.478


Score: 0.477

Score: 0.475

Score: 0.469


Score: 0.436


Score: 0.429

False positives

Your model incorrectly predicted electricPole on these images:

Score: 0.075


Score: 0.076

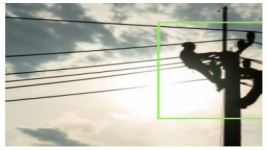
Score: 0.077

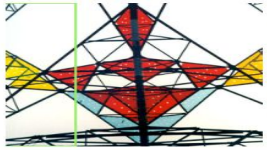
Score: 0.077


Score: 0.077

Score: 0.078

True positives

Your model correctly predicted electricPole on these images:

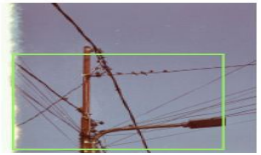
Score: 0.08

Score: 0.083

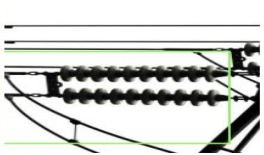
Score: 0.103


Score: 0.104

Score: 0.106

Score: 0.107

AutoML Object Detection Model

AutoML object detection model is built to identify the electric pole/structure and its location from the given image.

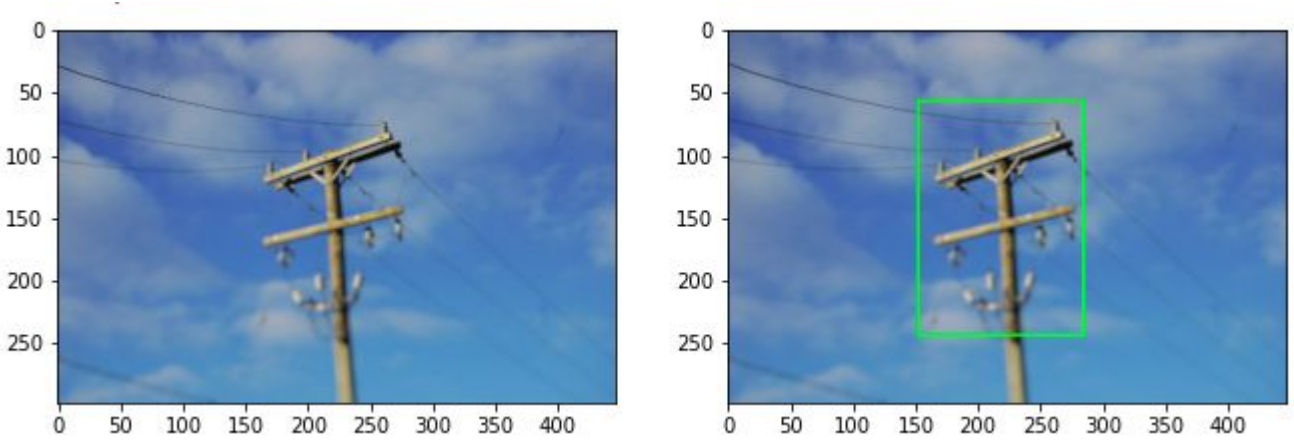

Blurriness Detection

Definition of Blur:

- An image is said to be blur when the house looks less sharp and less detailed in a way such that the edges that particular are not accurately perceived.
- In the blurred part of an image, the variance of the laplacian will be less as compared to the sharp image.

Threshold for binary output:

Non Blur : Variance score >= 260
Blur : Variance score < 260

Blurriness Detection

Blur mathematical model

- Laplacian operator - Used to measure 2nd derivative of an image
- Highlights the rapid changes in intensity of the image
    - Low variance - less change in intensity (smoother image edges)
    - High variance - More edges and less smoother

Laplacian Filter

Figure: Convolving the input image with the Laplacian operator

- We take a single channel of an image (grayscale) and convolve it with the following 3 x 3 kernel.
- Then take the variance (i.e. standard deviation squared) of the response.
- If the variance falls below a predefined threshold, then the image is considered blurry; otherwise, the image is not blurry.
- Laplacian operator itself, which is used to measure the 2nd derivative of an image. The Laplacian highlights regions of an image containing rapid intensity changes, much like the Sobel and Scharr operators.

Occlusion Detection

Definition of Occlusion:

- As per the business definition, it is required that the pole top part should be clearly and completely visible. That's why the tensorflow model is used to detect if the view is covered or blocked by any trees or bushes around.
- Additional, it tells if there is any tree present inside the property area. So that the appropriate measures can be taken.

Threshold for binary output:

Not Occluded : Confidence score >= 0.50
Occluded : Confidence score < 0.50

Occlusion Detection

Some of the sample training images are given below:

Non-Occluded

Occluded

Occlusion Detection

Model Framework

- Tensorflow based deep learning model
- A convolutional neural network architecture using residual bottleneck layers

Comparison with Transfer Learning Approaches

Method	Accuracy	Precision	Recall	F1-Score	AUC
Mobile Net V2	**98.67%**	**99.00%**	**99.00%**	**99.00%**	**0.987**
VGG16	82.78%	83.00%	83.00%	83.00%	0.827
Deep CNN	95.36%	95.00%	95.00%	95.00%	0.953

Deploying to Endpoint

Once the model is trained, we can easily deploy it to the endpoint, which can be used for online and batch prediction.

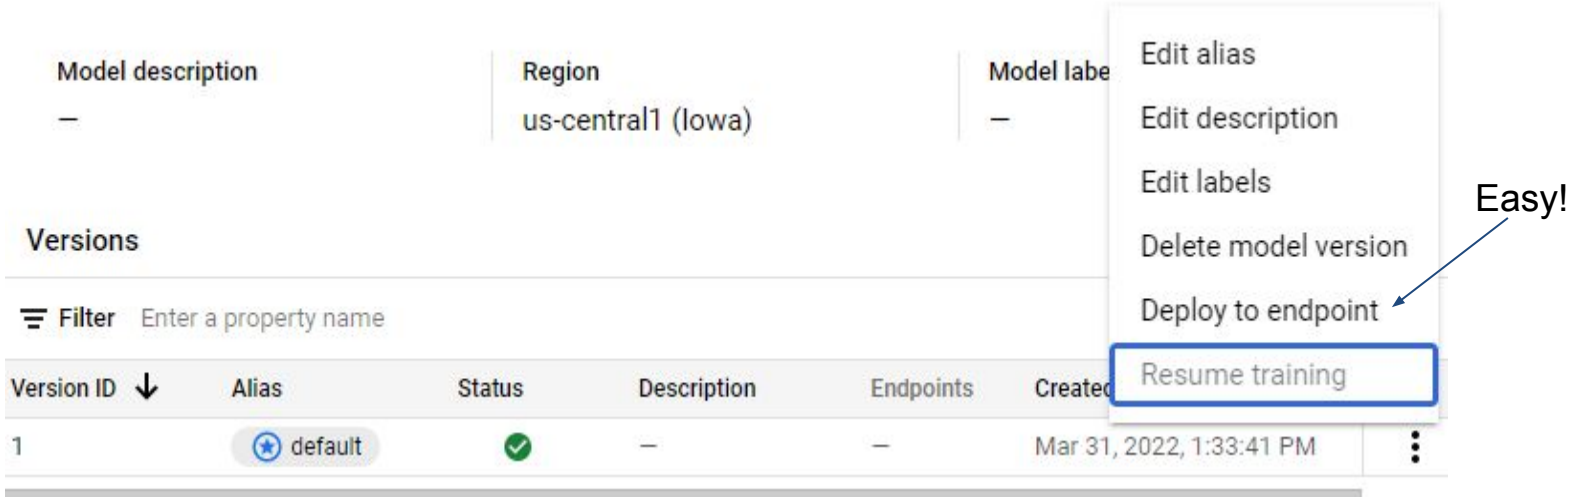

Deploying to Endpoint URL

Every endpoint has its own id, which we need if we want to access it outside of the console

Endpoint - Desired endpoint - Sample Request

Chamber of secrets

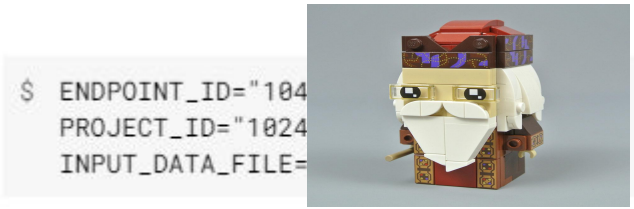

Prediction using Endpoint URL

Once we have the id, we can start using for prediction, we can either do it from console, in which it predicts one picture at a time. Or, write our own code:

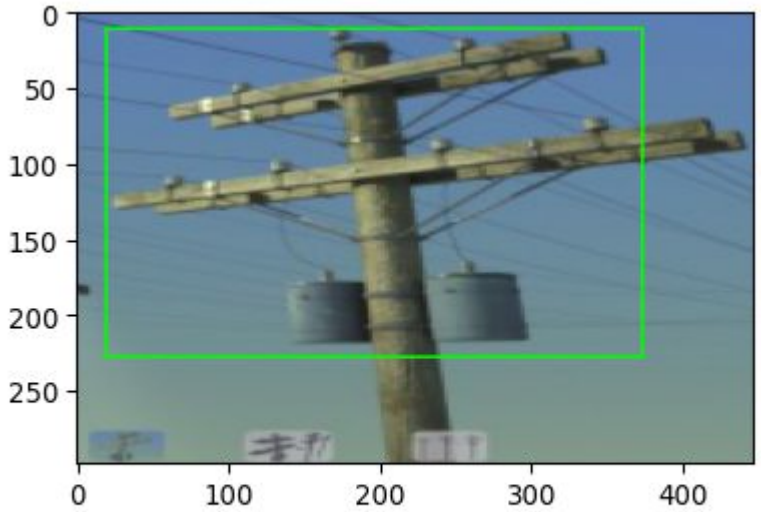

Output!

```
('Not Occluded', 0.49930325150489807)
('Not Blur', 225.17870455392827)
```

Skills required for Data Scientist

1. **Math and Statistics**
   In order to help make recommendations and decisions for the data driven businesses, one needs to understand different approaches to statistics - including maximum likelihood estimators, distributors, and statistical tests

1. **Analytics and Modeling**
   Data Scientist should be able to analyze data, run tests, and create models to gather new insights and predict possible outcomes.

Skills required for Data Scientist

3. **Machine Learning Methods**
   You should be familiar with various algorithms like Decision trees, logistic regression, random forest etc.

4. **R/Python Programming**
   To create practical applications, programming skills are must. Object-oriented programming, basic syntax, and functions, flow control statements

Skills required for Data Scientist

5. **Data Visualization**
   Understanding how to break down complex data into smaller, digestible pieces as well as using a variety of visual aids (charts, graphs, and more) is one skill any Data Scientist will need to be proficient

6. **Communication and Intellectual Curiosity**
   Data Scientist needs to have strong communication skills. Data on its own doesn't mean a whole lot, so a great Data Scientist is fueled by a desire to understand more about what the data is telling them, and how that information can be used on a broader scale.

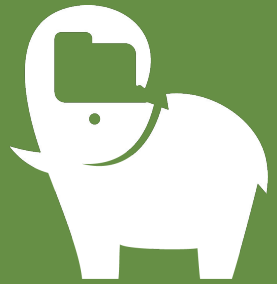

THIRD EYE

Q & A

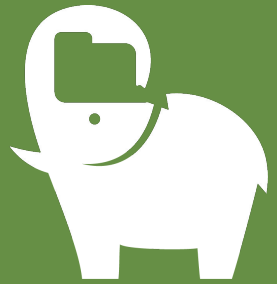

THIRD EYE

THANK YOU!

ThirdEyeData.io