ARTIFICIAL INTELLIGENCE IMAGE CLASSIFICATION:

PROOF OF CONCEPT FOR UTILITY ASSET IDENTIFICATION

WRITTEN BY ANDREW WAHRER, CSPO, BUSINESS CONSULTANT

Primera

Primera's Consulting division is at the forefront of harnessing the power of artificial intelligence (AI) to improve its workflow, ultimately benefiting clients. Among our successes is AI image classification, which has proven to be a game-changer for efficiency. By employing this technique, Primera's business consultants have transformed the way we handle visual data for our utilities clients by streamlining processes, reducing manual effort, and ensuring precision. Let's delve deeper into this topic by reviewing a case study for our smart LED streetlights program for large utility clients.

First, some background on what AI image classification is. Image classification teaches a computer to recognize objects in photos, just like we would, for example, teach a friend to identify different types of trees. With supervised classification, we tell the machine what to look for by training it with labeled examples. Key components include object detection (finding where an object is) and localization (identifying the objects in an image).

Effective management and tracking of utility asset data and analytics—such as those related to streetlights—are increasingly critical due to the introduction of smart connected devices within the electrical distribution network. Utilities planning to deploy large-scale smart streetlights need an accurate assessment of their current assets. Leveraging AI image classification can be an effective and cost-efficient alternative to manual field walkdowns and asset data collection. A well-trained AI image classification model can accurately identify over 20 attributes on a single asset on a distribution pole.

Primera understands the challenges that stem from performing field walkdowns and collecting data on several thousand streetlights. For one utility client, Primera was charged with managing the surveying of over 150,000 streetlights. To do this more efficiently, Primera created an AI image classification proof of concept. We trained and created our own model to classify photos of streetlights by fixture style. The model was able to successfully identify between five different streetlight fixture styles with an average precision of 92%.

Primera's unique advantage stemmed from our extensive streetlight experience, exposure and access to streetlight data, and field data collection knowledge. Our team led smart LED streetlight programs for multiple investor-owned utilities and public municipalities. We developed the business cases, monitored the conversion projects, and resolved data discrepancies and billing issues.

Throughout the process, there were several best practices and lessons learned that were uncovered when creating and implementing the image classification model. In this article, we'll briefly review the four steps to creating the model and reflect on the lessons learned while creating the proof of concept. The four steps include:

PREPARING DATA

TRAINING THE MODEL

EVALUATING THE MODEL'S PERFORMANCE

DEPLOYING AND USING THE MODEL

PREPARING DATA

The first step to building a model is to gather and prepare the data needed for training. Make sure to optimize your model for real world photos. For instance, if the training dataset contains some photos that need to be rotated to be viewed correctly, don't rotate them! This will train the model to account for various angles you may receive back during production and helps you avoid rotating those later. If your use case involves grainy low-light images, then your training images should be equally as fuzzy. Include a variety of photos in your training datasets that will cover the variety of conditions you expect to receive during the project.

Best practices recommend about 1,000 training images per attribute, but that requires a lot of data preparation and setup time. To save time during the proof of concept, we used a little over 100 images per attribute and our model's average precision was still 92%. We also automatically resized our images to a lower resolution and file size for better training and model efficiency.

The model works best when the most common label has more images than the least common label. To improve performance, we completely removed very low frequency categories (< 1%) for our training data.

Before proceeding, keep in mind that current machine learning capabilities mean that if a human can't recognize an object in an image by looking at it for a couple seconds, the machine likely won't be able to provide predictions either.

TRAINING THE MODEL

Once the training data (images) are prepared, it's time to feed this data to the model. The model learns to map input images to their corresponding categories. This process can take a long time depending on computing resources. There are several techniques for training the model. AutoML (Automated Machine Learning) Platforms like Google AutoML automate most of the process. Python libraries such as TensorFlow and PyTorch are popular for building custom image classification models and provide fine-grained control.

With any technique, it's important to consider using cloud computing resources. Training a machine learning model with a large dataset is resource intensive and will take a long time if you don't have a powerful graphics processing unit (GPU).Cloud services allow you to pay only for what you use, reducing upfront costs. However, once you're finished with the cloud resources, remember to shut down your cloud instance or you could be surprised with a large bill.

EVALUATING THE MODEL'S PERFORMANCE

There are several metrics that measure how well a model performs. Precision simply measures the percentage of predictions that were correct. The higher the precision, the fewer false positives predicted. If the model identifies whether an image contains a tree or not, a precision score of 0.5 means the model is correct 50% of the time. Depending on your need, you might have a lower tolerance for false positives than for false negatives or the other way around.

Precision-Recall Curve

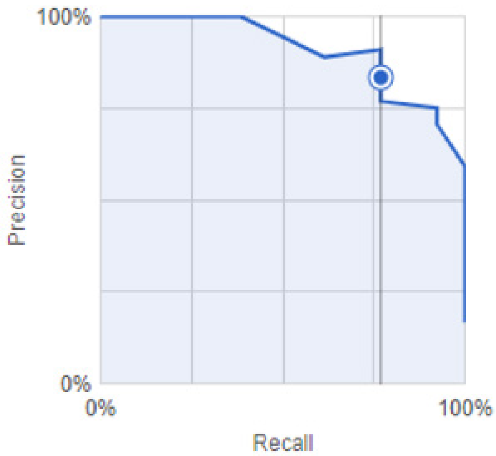

DEPLOYING AND USING THE MODEL

Once the model is trained, it's time to start using it. If you want real-time predictions, the model can be deployed to an online endpoint. Say you want to take a picture of a tree on your phone and instantly see what species of tree it is, you could create an application programming interface (API) for a mobile app that connects to the model. On the other hand, batch predictions can be made without needing to deploy the model to an endpoint. Use batch predictions when you don't need an immediate response.

SUMMARY

The primary use cases from this proof of concept include capturing assets on streetlights, other smart devices, utility poles, mast arms, and other distribution equipment. The goal was to capture attributes from these assets more accurately and more efficiently than possible with a traditional field walkdown. Simply capturing a photo of the asset and allowing AI to do the rest of the work saves on costs, improves accuracy, and increases productivity with the time saved. It also improves safety and risk without the need to leave a vehicle. Even when performing field walkdowns, AI image processing can be used as a quality check when referencing photos and checking against attributes manually captured.

If you'd like more details on using AI image classification for your project or have questions, feel free to visit ***www.primeraeng.com/consulting*** or reach out to Andrew Wahrer, Primera's data and analytics expert.

ABOUT THE AUTHOR

Andrew Wahrer is one of Primera's leading business analysis experts passionate about supporting the development and continuous improvement of the firm's business and data analytics work. As a key member of Primera's Consulting team, he enjoys learning about new technology and implementing techniques at the forefront of industry capability. His expertise is the result of over 15 years of experience with improving work practices, executive level reporting, benchmarking multi-billion-dollar investment portfolios, automating processes, and increasing productivity through analytical insights in the utility and engineering industries.