Cloud AutoML Vision Object Detection

Cloud model creation via the UI quickstart

This quickstart walks you through the process of:

- Creating a dataset of images with labeled bounding boxes.
- Train a custom model using your dataset.

In this quickstart you use the custom user interface (UI) (https://console.cloud.google.com/vision) to interact with the Cloud AutoML API. You can also complete all the steps in this quickstart using the AutoML API (https://cloud.google.com/vision/automl/object-detection/docs/quickstart). For further directions on using the UI or API, see the how-to guides (https://cloud.google.com/vision/automl/object-detection/docs/how-to).

Set up your project

Before you can use Cloud AutoML Vision Object Detection, you must create a Google Cloud project (one never used with any other AutoML product) and enable Cloud AutoML Vision Object Detection for that project.

1. Sign in (https://accounts.google.com/Login) to your Google Account.

   If you don't already have one, sign up for a new account (https://accounts.google.com/SignUp).

2. In the Cloud Console, on the project selector page, select or create a Google Cloud project.

   ★ **Note**: If you don't plan to keep the resources that you create in this procedure, create a project instead of selecting an existing project. After you finish these steps, you can delete the project, removing all resources associated with the project.

   GO TO THE PROJECT SELECTOR PAGE (HTTPS://CONSOLE.CLOUD.GOOGLE.COM/PROJECTSELECT

3. Make sure that billing is enabled for your Google Cloud project. Learn how to confirm billing is enabled for your project (https://cloud.google.com/billing/docs/how-to/modify-project).

4. Enable the Cloud AutoML and Storage APIs.

   ENABLE THE APIS (HTTPS://CONSOLE.CLOUD.GOOGLE.COM/FLOWS/ENABLEAPI?APIID=STORAGE-

Beta requirements

1. This Beta requires you to use `us-central1` as the specified region.

Preparing a dataset

In this quickstart you use a dataset created from Open Images Dataset V4 (https://storage.googleapis.com/openimages/web/index.html). This publicly available Salads dataset is located at `gs://cloud-ml-data/img/openimage/csv/salads_ml_use.csv`.

The CSV format is as follows:

```
TRAIN,gs://cloud-ml-data/img/openimage/3/2520/3916261642_0a504acd60_o.jpg,Salad,0.0,
VALIDATE,gs://cloud-ml-data/img/openimage/3/2520/3916261642_0a504acd60_o.jpg,Seafood
TEST,gs://cloud-ml-data/img/openimage/3/2520/3916261642_0a504acd60_o.jpg,Tomato,0.0,
```

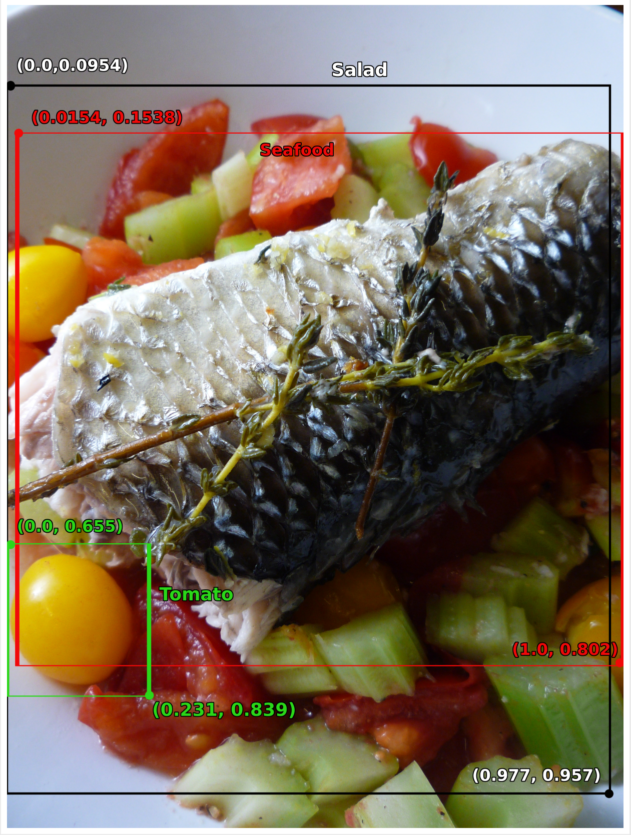

`3916261642_0a504acd60_o.jpg`

Each row corresponds to an object localized inside a larger image, with each object specifically designated as test, train, or validation data. The three lines included here indicate **three distinct**

objects located inside the same image available at `gs://cloud-ml-data/img/openimage/3/2520/3916261642_0a504acd60_o.jpg`. Each row has a different label: `Salad`, `Seafood`, `Tomato`, in addition to other rows with `Baked goods` or `Cheese` labels.

Bounding boxes are specified for each image using the top left and bottom right vertices:

- (0,0) corresponds to the top left-most vertex.
- (1,1) corresponds to the bottom right-most vertex.

For the first row shown above, the (x, y) coordinates for the top left vertex of the `Salad` labeled object are (0.0,0.0954), and the coordinates for the bottom right vertex of the object are (0.977,0.957).

For more detailed information on how to format your CSV file and the minimum requirements for creating a valid dataset, see Preparing your training data (https://cloud.google.com/vision/automl/object-detection/docs/prepare).

Create a dataset and import training images

1. Open the Cloud AutoML Vision Object Detection UI (https://console.cloud.google.com/vision) and select your project from the drop-down list in the title bar.

   The first time you open the Cloud AutoML Vision Object Detection UI you will need to **"Enable AutoML API"** if it isn't already enabled.

2. Select **"Get started"** in the pop-up window when prompted.

   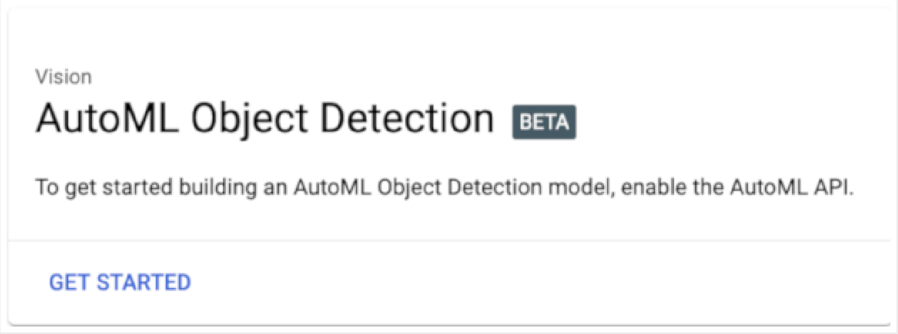

   Vision

   AutoML Object Detection BETA

   To get started building an AutoML Object Detection model, enable the AutoML API.

3. You will be taken to the datasets list page. Create a new dataset by selecting **New Dataset**.

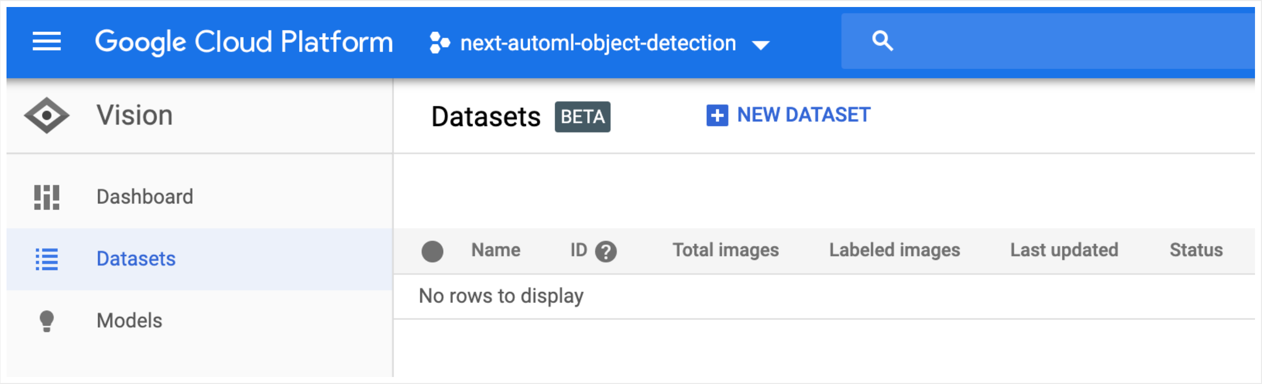

	Name	ID	Total images	Labeled images	Last updated	Status
No rows to display						

Enter a unique name for your dataset.

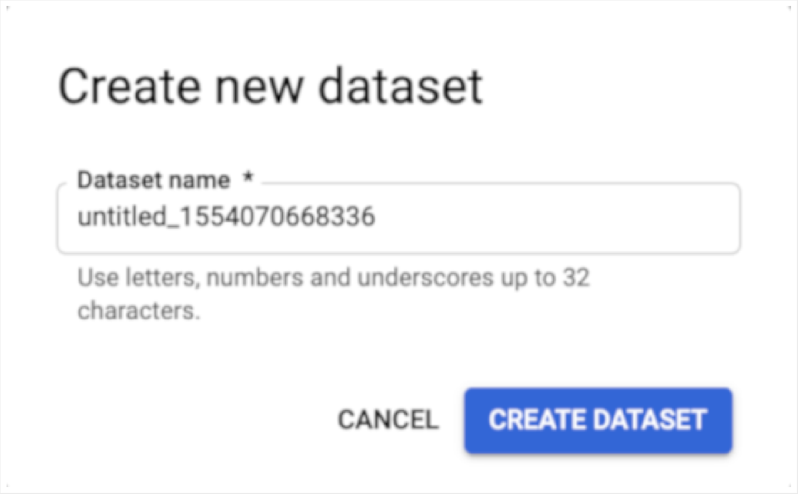

Create new dataset

Dataset name *: untitled_1554070668336

Use letters, numbers and underscores up to 32 characters.

4. Enter the location of the training data to import in the window that follows.

   In the **Select a CSV file on Cloud Storage** text box, enter the path for the sample CSV file (the `gs://` prefix is added automatically):

   ```
   cloud-ml-data/img/openimage/csv/salads_ml_use.csv
   ```

   Alternatively, you can select **Browse** and navigate to the CSV file in one of your Google Cloud Storage buckets.

   This quickstart uses sample data staged in a public Google Cloud Storage bucket. The training data are sample JPG images annotated with the bounding box and label of the objects you want the model to learn to identify. To import the training data into the dataset, use a CSV file that points to the image (JPEG, PNG, GIF, BMP, or ICO) files; see Preparing your training data (https://cloud.google.com/vision/automl/object-detection/docs/prepare) for information about the format and image specifications.

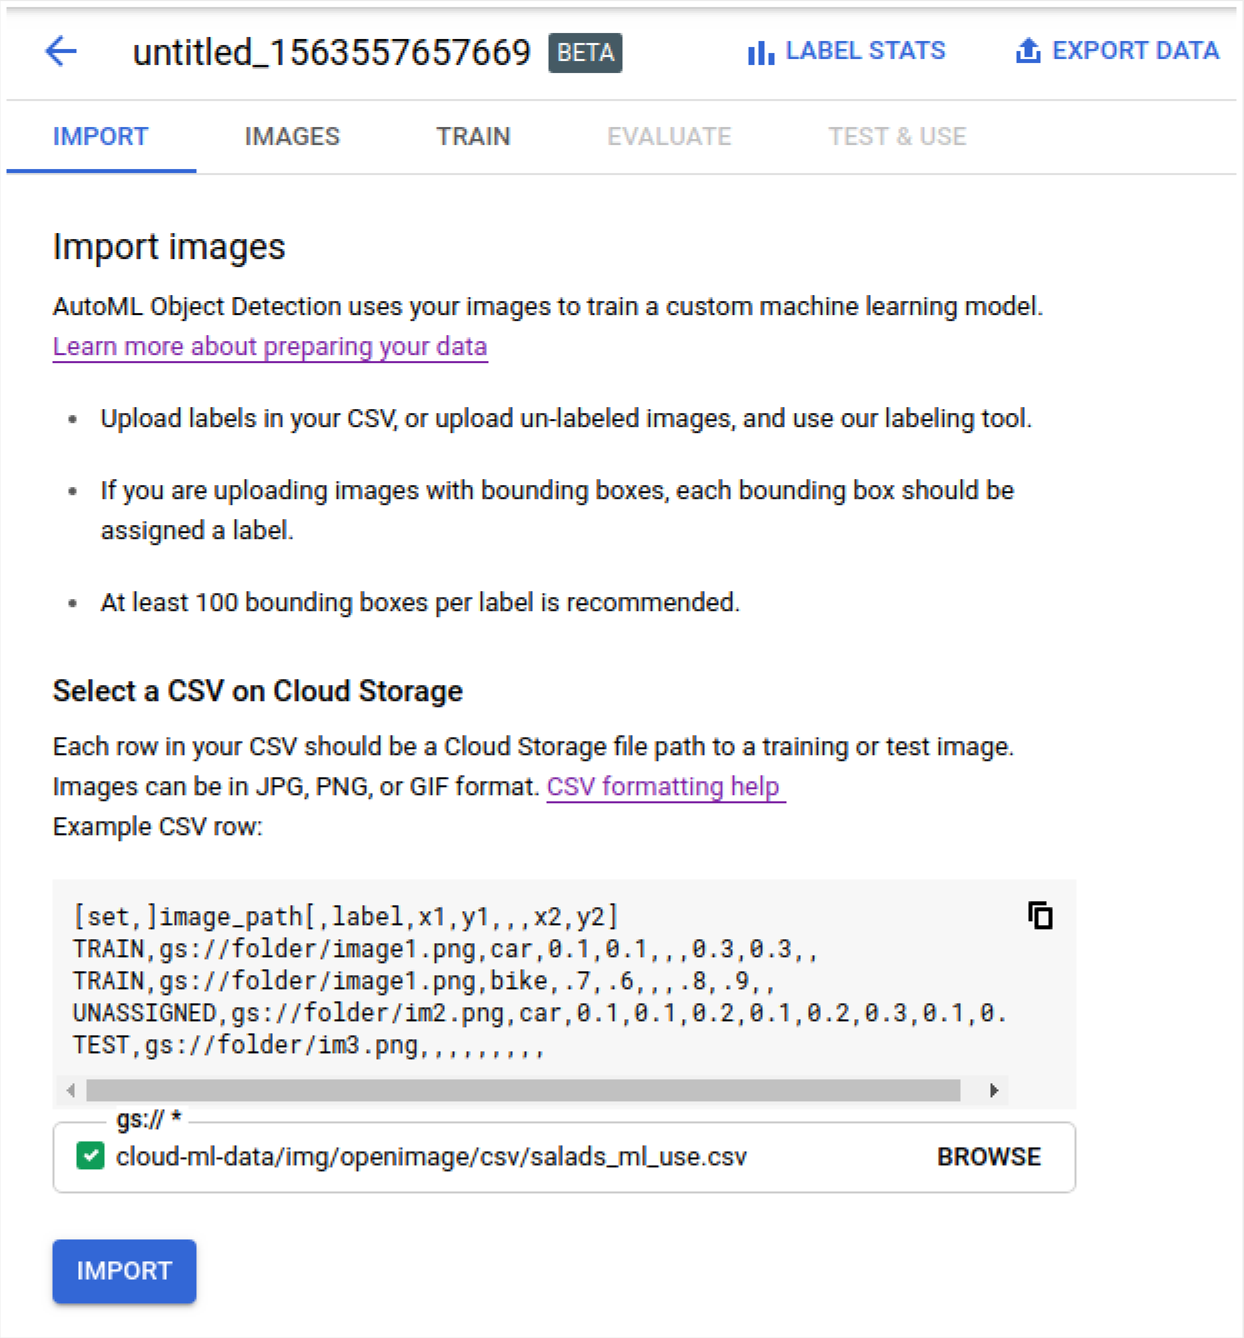

Import images

AutoML Object Detection uses your images to train a custom machine learning model.
Learn more about preparing your data

- Upload labels in your CSV, or upload un-labeled images, and use our labeling tool.
- If you are uploading images with bounding boxes, each bounding box should be assigned a label.
- At least 100 bounding boxes per label is recommended.

Select a CSV on Cloud Storage

Each row in your CSV should be a Cloud Storage file path to a training or test image. Images can be in JPG, PNG, or GIF format. CSV formatting help
Example CSV row:

```
[set,]image_path[,label,x1,y1,,,x2,y2]
TRAIN,gs://folder/image1.png,car,0.1,0.1,,,0.3,0.3,,
TRAIN,gs://folder/image1.png,bike,.7,.6,,,.8,.9,,
UNASSIGNED,gs://folder/im2.png,car,0.1,0.1,0.2,0.1,0.2,0.3,0.1,0.
TEST,gs://folder/im3.png,,,,,,,,,
```

gs:// *
☑ cloud-ml-data/img/openimage/csv/salads_ml_use.csv

5. Select **Import**.

   Your dataset will show a **Status** of `Running:importing images` while your images are being imported. This process takes just a few minutes.

When your training data has been successfully imported, the **Status** column says `Success:Creating dataset` and the UI shows the generated **ID** for the dataset (used when making AutoML API calls) as well as the number of items imported.

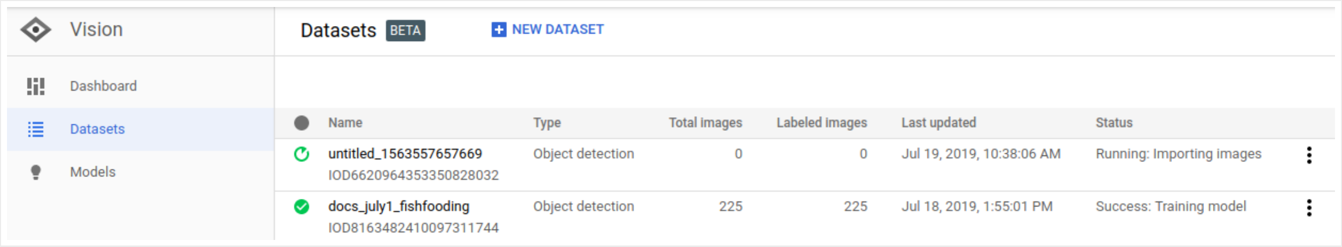

Name	Type	Total images	Labeled images	Last updated	Status
untitled_1563557657669 IOD6620964353350828032	Object detection	0	0	Jul 19, 2019, 10:38:06 AM	Running: Importing images
docs_july1_fishfooding IOD8163482410097311744	Object detection	225	225	Jul 18, 2019, 1:55:01 PM	Success: Training model

If there are any problems importing images you will see `Warning: Importing images` as the status. Select the dataset name and **Details** to see the errors with specific image import.

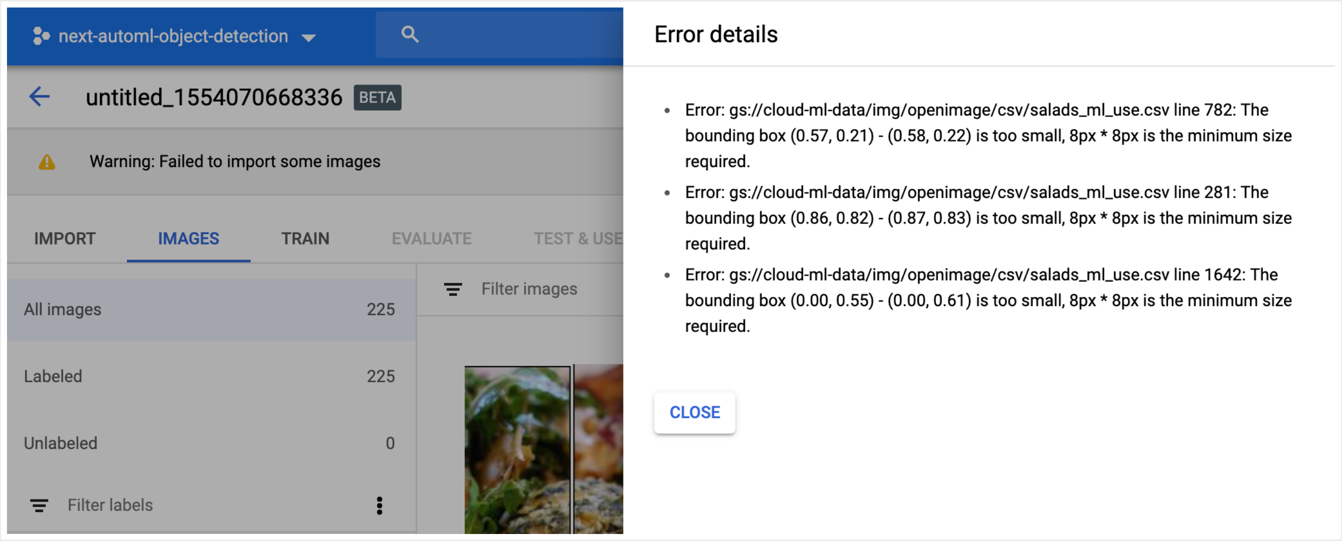

Warning: Failed to import some images

Error details

- Error: gs://cloud-ml-data/img/openimage/csv/salads_ml_use.csv line 782: The bounding box (0.57, 0.21) - (0.58, 0.22) is too small, 8px * 8px is the minimum size required.
- Error: gs://cloud-ml-data/img/openimage/csv/salads_ml_use.csv line 281: The bounding box (0.86, 0.82) - (0.87, 0.83) is too small, 8px * 8px is the minimum size required.
- Error: gs://cloud-ml-data/img/openimage/csv/salads_ml_use.csv line 1642: The bounding box (0.00, 0.55) - (0.00, 0.61) is too small, 8px * 8px is the minimum size required.

Train your model

After you have created your dataset and imported your training data into your dataset, you can train your custom model.

1. From the **Datasets** listing page, click the dataset name.

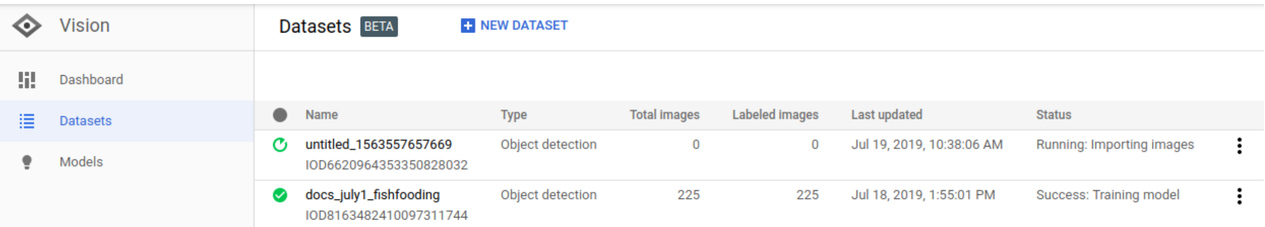

Name	Type	Total images	Labeled images	Last updated	Status
untitled_1563557657669 IOD6620964353350828032	Object detection	0	0	Jul 19, 2019, 10:38:06 AM	Running: Importing images
docs_july1_fishfooding IOD8163482410097311744	Object detection	225	225	Jul 18, 2019, 1:55:01 PM	Success: Training model

2. Select the **Train** tab. This will show you all labels and their division into train, test, and validation sets.

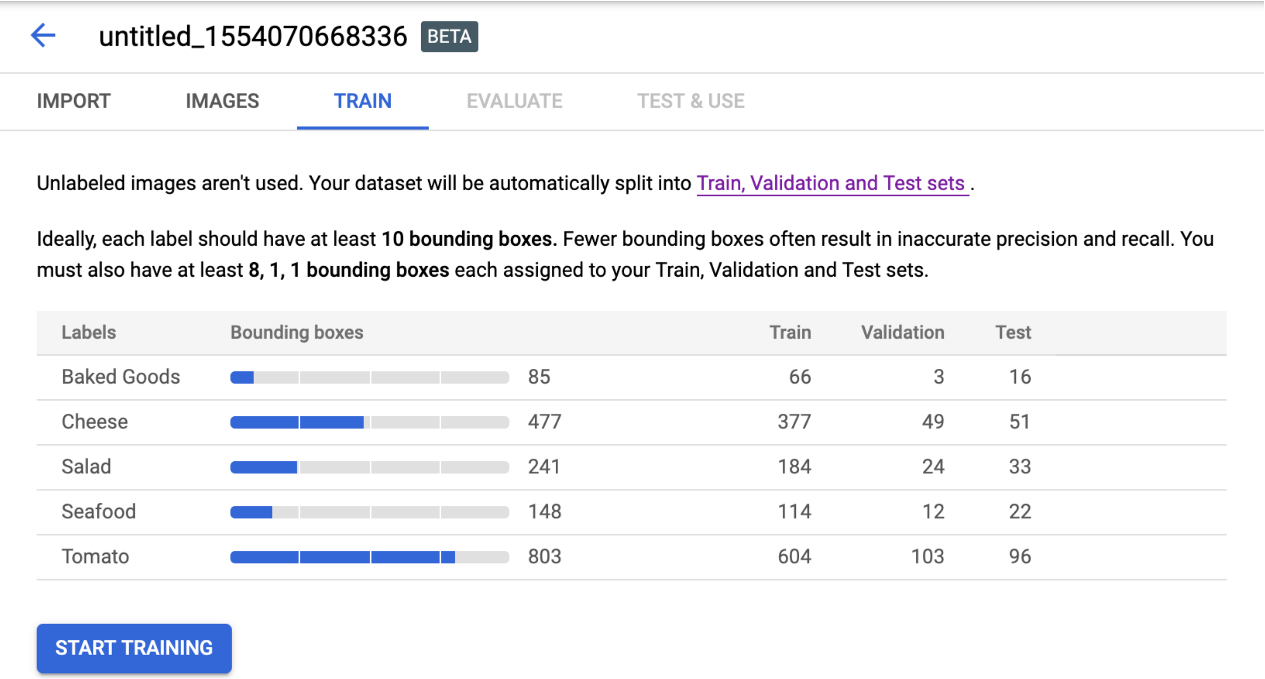

untitled_1554070668336 BETA

Unlabeled images aren't used. Your dataset will be automatically split into Train, Validation and Test sets .

Ideally, each label should have at least **10 bounding boxes.** Fewer bounding boxes often result in inaccurate precision and recall. You must also have at least **8, 1, 1 bounding boxes** each assigned to your Train, Validation and Test sets.

Labels	Bounding boxes	Train	Validation	Test
Baked Goods	85	66	3	16
Cheese	477	377	49	51
Salad	241	184	24	33
Seafood	148	114	12	22
Tomato	803	604	103	96

3. Select **Start training** to open a side panel with training options.

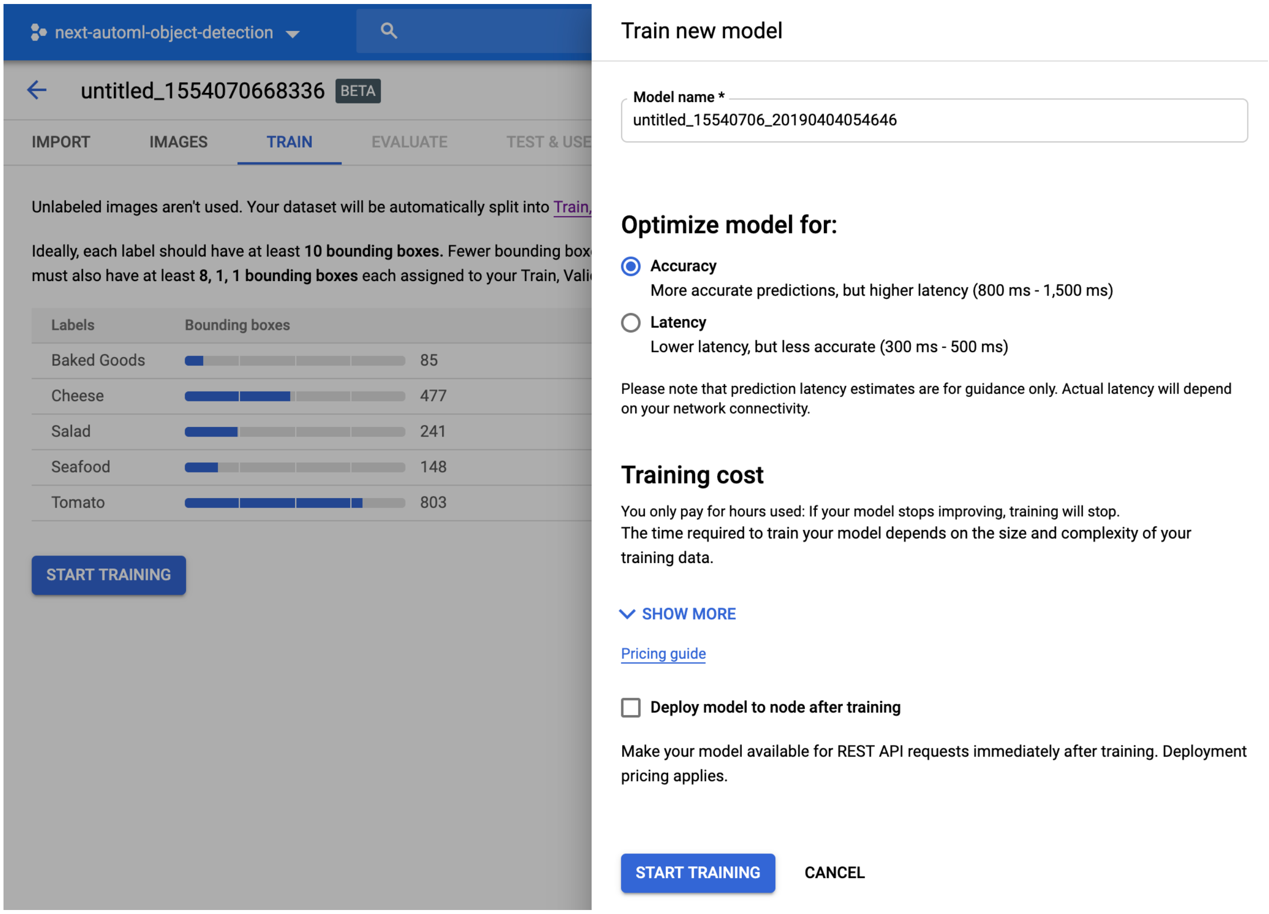

next-automl-object-detection

untitled_1554070668336 BETA

IMPORT IMAGES TRAIN EVALUATE TEST & USE

Unlabeled images aren't used. Your dataset will be automatically split into Train,

Ideally, each label should have at least **10 bounding boxes.** Fewer bounding box
must also have at least **8, 1, 1 bounding boxes** each assigned to your Train, Vali

Labels	Bounding boxes
Baked Goods	85
Cheese	477
Salad	241
Seafood	148
Tomato	803

START TRAINING

Train new model

Model name *
untitled_15540706_20190404054646

Optimize model for:

● **Accuracy**
More accurate predictions, but higher latency (800 ms - 1,500 ms)

○ **Latency**
Lower latency, but less accurate (300 ms - 500 ms)

Please note that prediction latency estimates are for guidance only. Actual latency will depend on your network connectivity.

Training cost

You only pay for hours used: If your model stops improving, training will stop.
The time required to train your model depends on the size and complexity of your training data.

□ **Deploy model to node after training**

Make your model available for REST API requests immediately after training. Deployment pricing applies.

4. Enter a name for your custom model or accept the default name. In this panel you can also select model optimization (for latency or accuracy).

Select ☑ **Deploy model to node after training** to opt-in to automatic deployment. Otherwise, you will be prompted to manually deploy your model (https://cloud.google.com/vision/automl/object-detection/docs/predict#deploy_your_model) after training before you are able to make predictions.

5. Select **Start Training**.

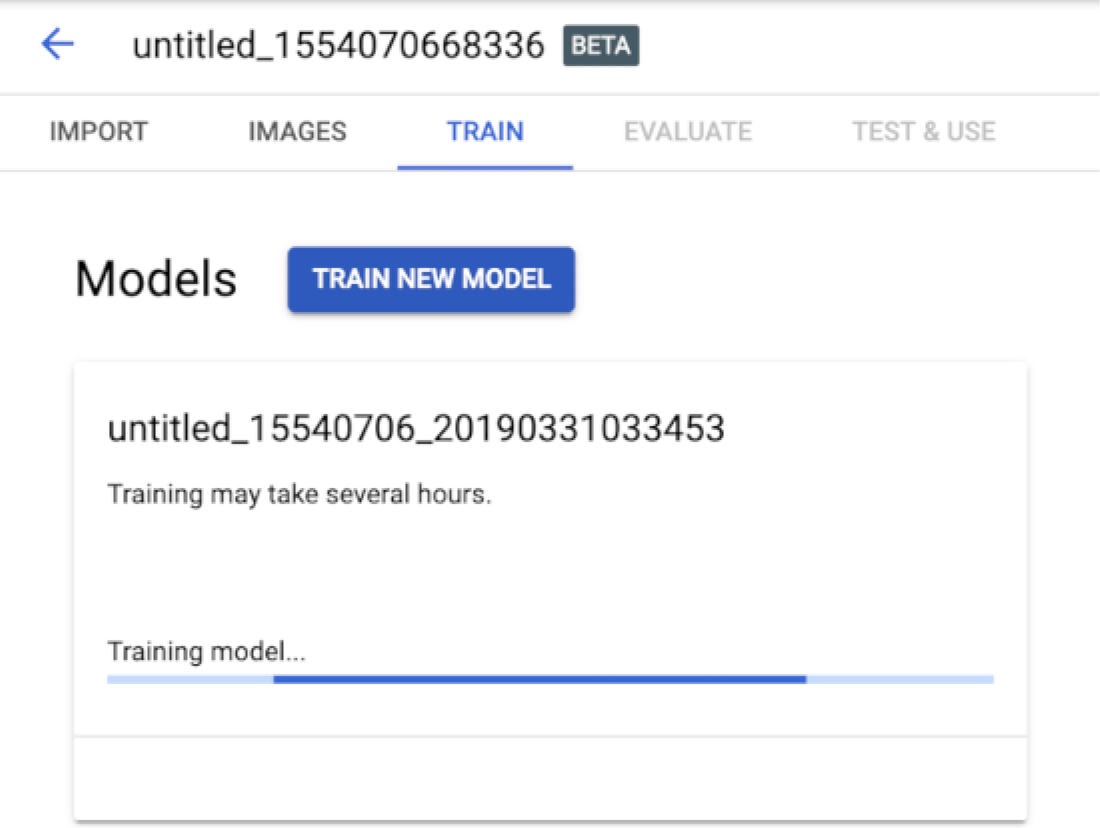

untitled_15540706_20190331033453

Training may take several hours.

Training model...

Training a model can take several hours to complete. Typical training time for the sample is about one hour.

After the model is successfully trained, you will receive a message at the email address you used to sign up for the program.

Evaluate the custom model

After training a model, Cloud AutoML Vision Object Detection evaluates the quality and accuracy of the new model. To see the evaluation metrics for your model:

1. Open the Cloud AutoML Vision Object Detection UI (https://console.cloud.google.com/vision) and click the **Models** tab (with lightbulb icon) in the left navigation bar.
2. Click the name of the model you want to evaluate.
3. If necessary, click the **Evaluate** tab just below the title bar.

   If training has been completed for the model, Cloud AutoML Vision Object Detection shows its evaluation metrics. It provides precision and recall scores for the model as a whole and for each object label at different score and intersection over union (https://cloud.google.com/vision/automl/object-detection/docs/evaluate#intersection-over-union) (IoU) thresholds. To see the metrics for a particular object label, choose the label from the "Filter labels" list.

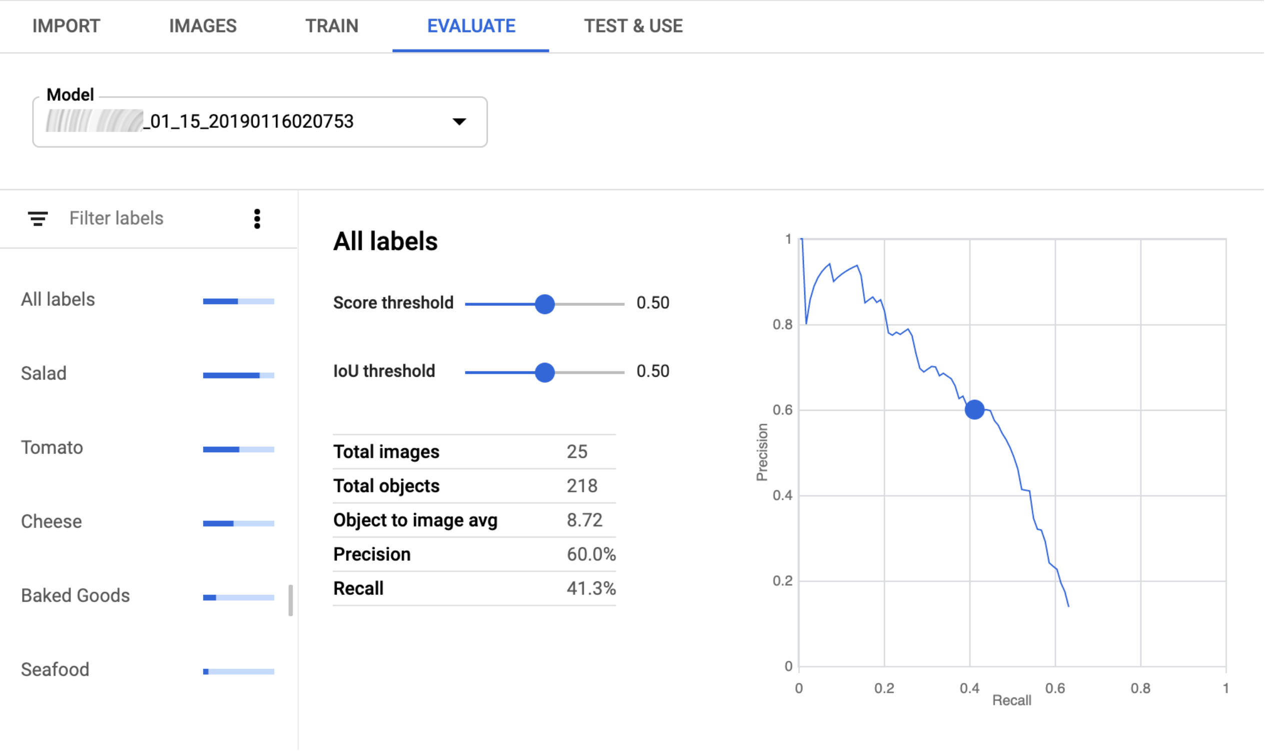

Model: [redacted]_01_15_20190116020753

All labels

Score threshold 0.50

IoU threshold 0.50

Total images	25
Total objects	218
Object to image avg	8.72
Precision	60.0%
Recall	41.3%

Precision and **recall** measure how well the model is capturing information, and how much it's leaving out. Precision indicates, from all objects annotated with a specific label, how many actually were supposed to be assigned to that label. Recall indicates, from all the objects that *should* have been identified as a particular entity, how many were *actually* assigned to that entity.

Use this data to evaluate your model's readiness:

- **Low precision or recall** scores can indicate that your model needs additional training data.

- **Perfect precision and recall** can indicate that the data is too easy and may not generalize well.

Use the custom model

After your model has been successfully trained, you can use it to identify objects in an image with bounding boxes and labels using your custom model. Select the **Test and Use** tab.

If you didn't opt-in for auto-deploy you will be prompted to deploy your model before you can make predictions.

Manually deploy model

AutoML Vision requires you to deploy your trained model before you can send prediction requests to it.

1. In the **Test and Use** tab manually deploy your model by selecting the **Deploy Model** option from the banner beneath your model name.

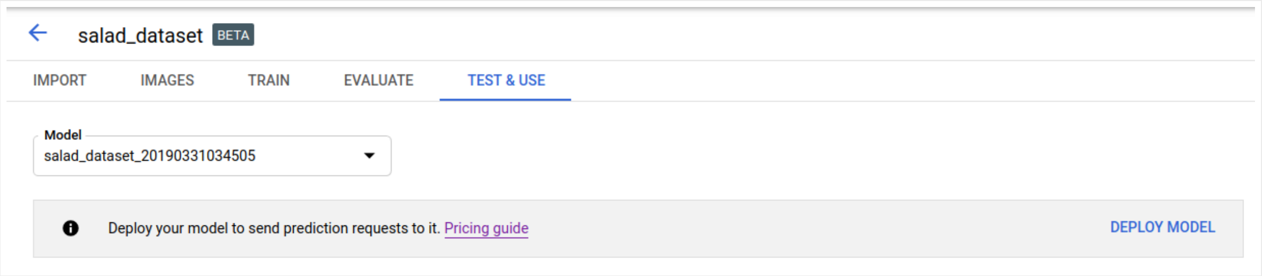
salad_dataset BETA
IMPORT
IMAGES
TRAIN
EVALUATE
TEST & USE
Model
salad_dataset_20190331034505
Deploy your model to send prediction requests to it. Pricing guide
DEPLOY MODEL

   Selecting **Deploy Model** opens the model deployment option window

2. In the model deployment option window you can select the number of nodes to deploy on and view the available prediction queries per second (QPS).

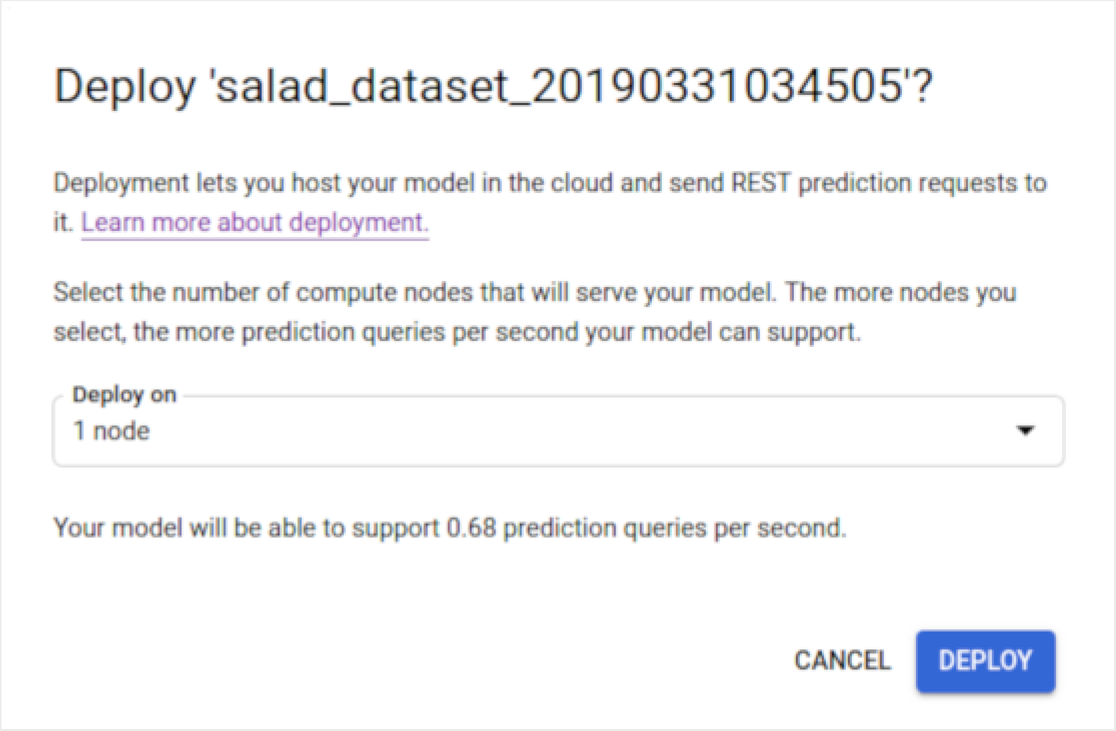

Deploy 'salad_dataset_20190331034505'?

Deployment lets you host your model in the cloud and send REST prediction requests to it. Learn more about deployment.

Select the number of compute nodes that will serve your model. The more nodes you select, the more prediction queries per second your model can support.

Deploy on: 1 node

Your model will be able to support 0.68 prediction queries per second.

3. Select **Deploy** to begin model deployment.

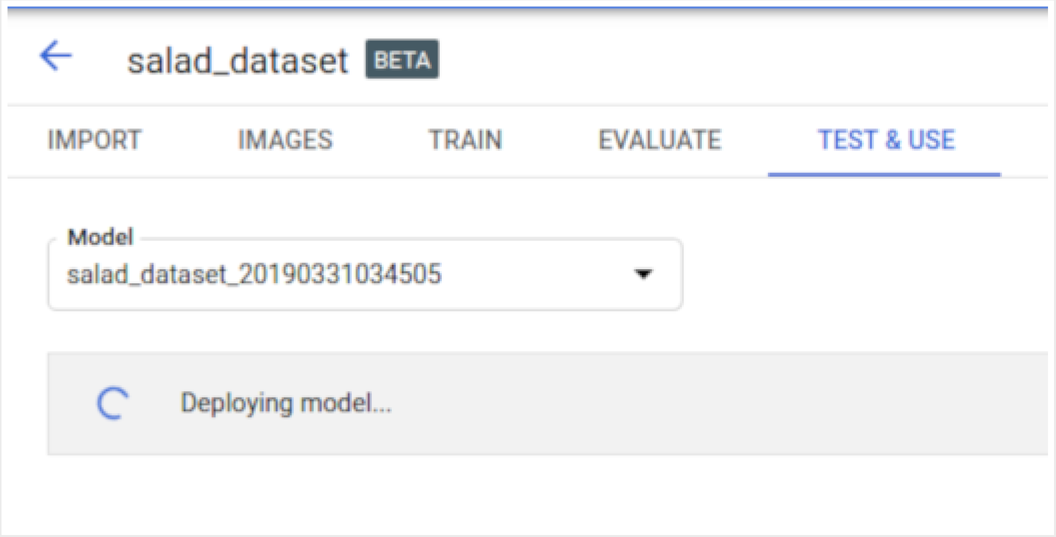

Model: salad_dataset_20190331034505

Deploying model...

4. You will receive an email when model deployment has completed.

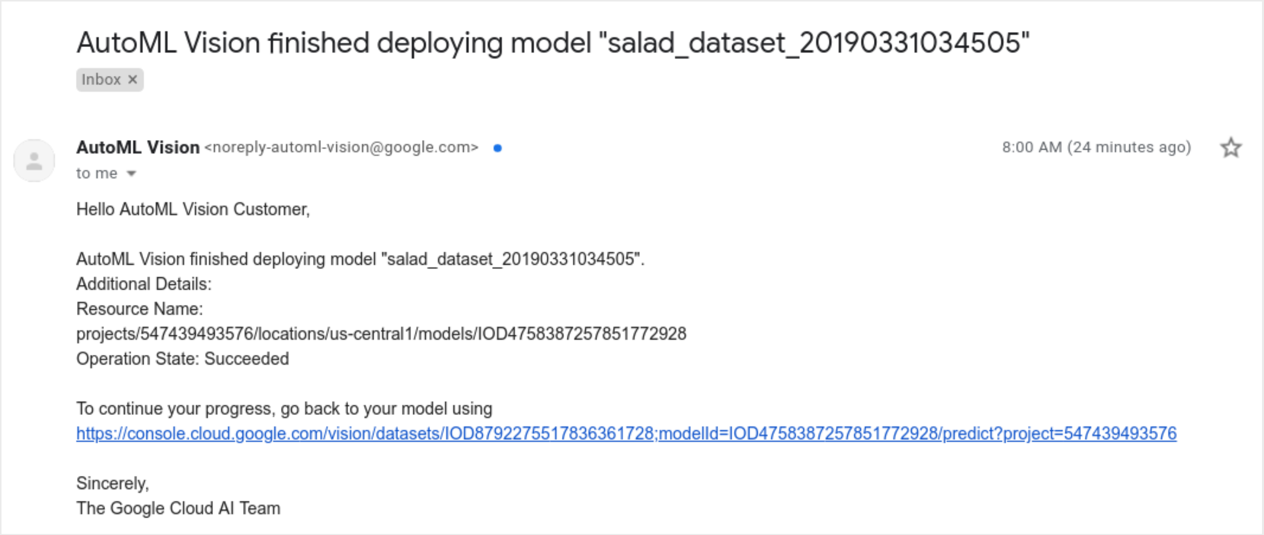

AutoML Vision finished deploying model "salad_dataset_20190331034505"

AutoML Vision <noreply-automl-vision@google.com>

to me

Hello AutoML Vision Customer,

AutoML Vision finished deploying model "salad_dataset_20190331034505".
Additional Details:
Resource Name:
projects/547439493576/locations/us-central1/models/IOD4758387257851772928
Operation State: Succeeded

To continue your progress, go back to your model using
https://console.cloud.google.com/vision/datasets/IOD8792275517836361728;modelId=IOD4758387257851772928/predict?project=547439493576

Sincerely,
The Google Cloud AI Team

Note: You can also use the UI to update the node count of a deployed model.

Make a prediction

After your model has been deployed, indicate the path to your test image on the **Test and use** page (such as this image (https://storage.cloud.google.com/cloud-ml-data/img/openimage/3357/3611049291_394ec30616_o.jpg?_ga=2.215351711.-900183802.1529534661), saved locally). Cloud AutoML Vision Object Detection analyzes the image using your model and displays labeled objects and their bounding boxes in the image.

Warning: Minimum bounding box size provided during data preparation (https://cloud.google.com/vision/automl/object-detection/docs/prepare) and training is 8px by 8px; model prediction performance can suffer when the predicted boxes are less than 8px by 8px.

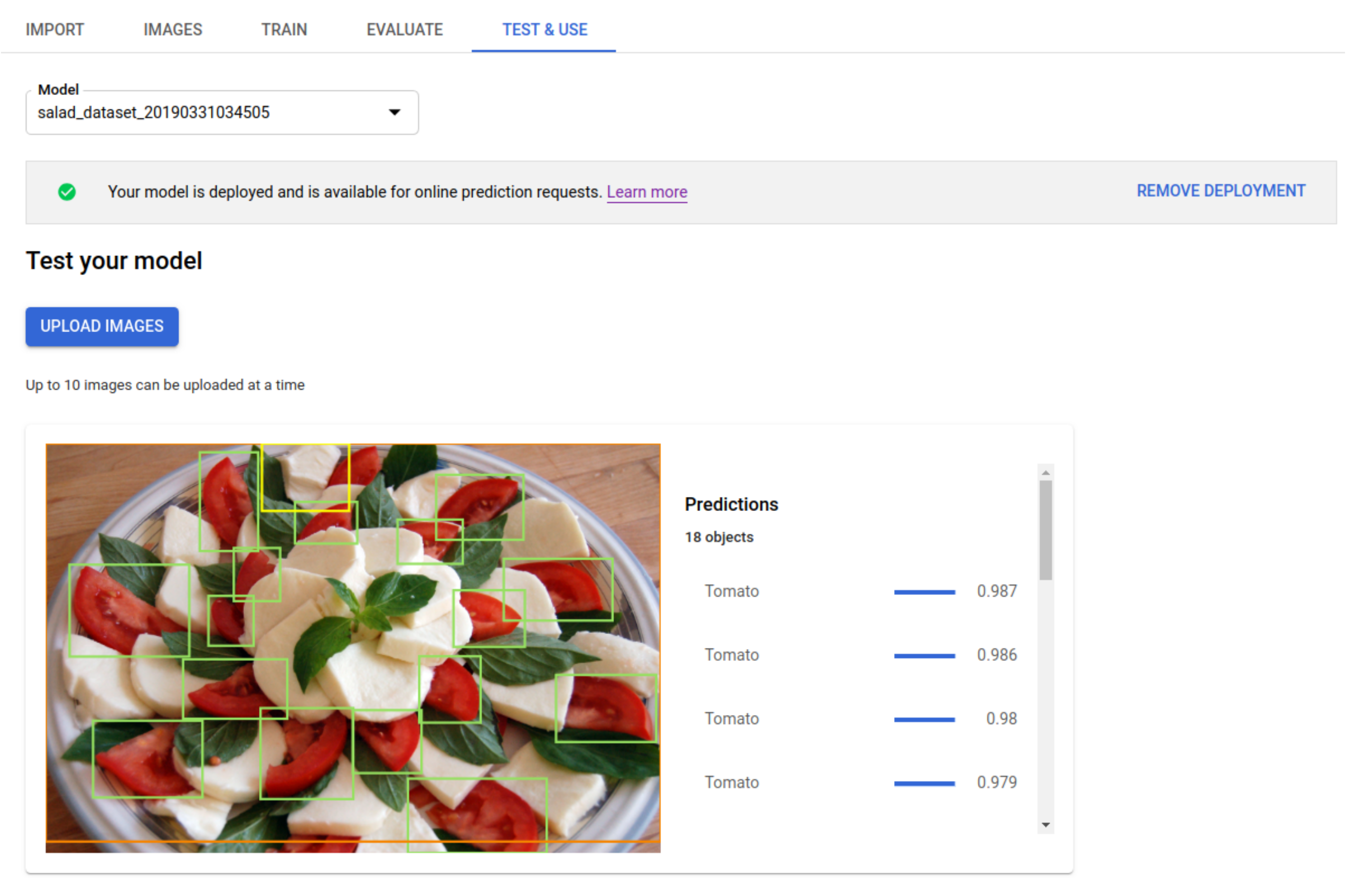

Model: salad_dataset_20190331034505

Your model is deployed and is available for online prediction requests. Learn more

Test your model

Up to 10 images can be uploaded at a time

Predictions

18 objects

Tomato	0.987
Tomato	0.986
Tomato	0.98
Tomato	0.979

Cleanup

If you no longer need your custom model or dataset, you can delete them.

To avoid unnecessary Google Cloud Platform charges, use the GCP Console (https://console.cloud.google.com/vision) to delete your project if you do not need it.

Undeploy your model

Your model incurs charges while it is deployed.

1. Select the **Test & Use** tab just below the title bar.
2. Select **Remove deployment** from the banner beneath your model name to open the undeploy option window.

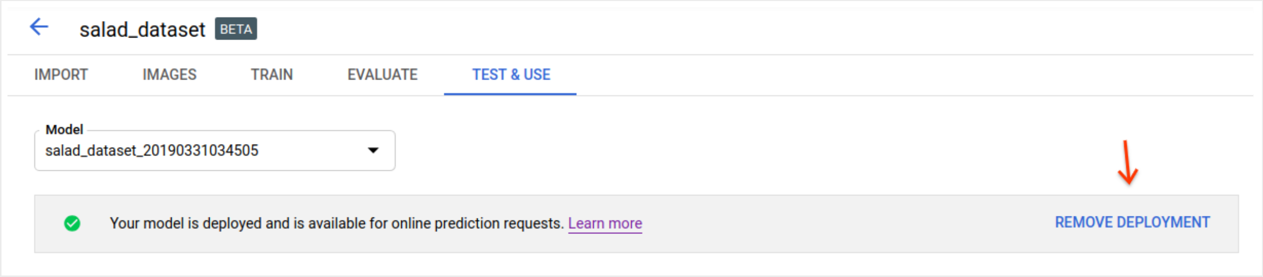

Model: salad_dataset_20190331034505

Your model is deployed and is available for online prediction requests. Learn more

3. Select **Remove deployment** to undeploy the model.

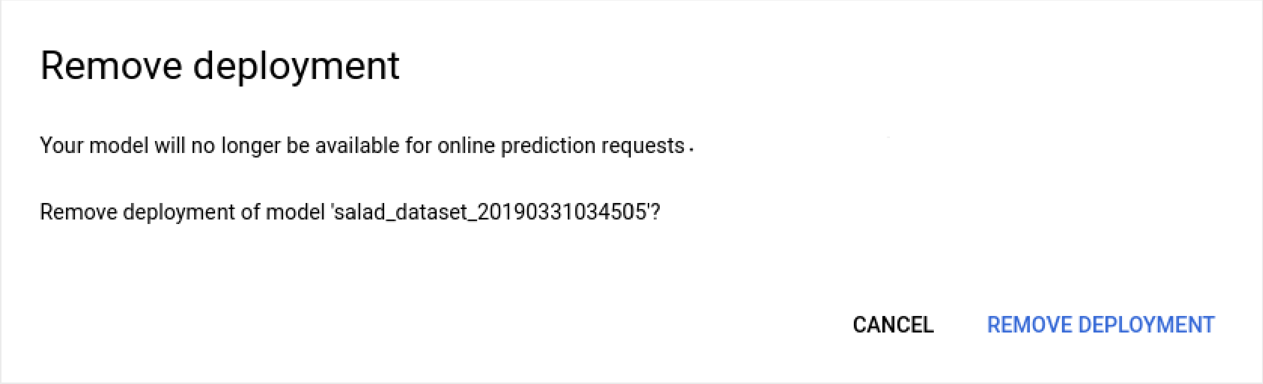

Remove deployment

Your model will no longer be available for online prediction requests.

Remove deployment of model 'salad_dataset_20190331034505'?

4. You will receive an email when model undeployment has completed.

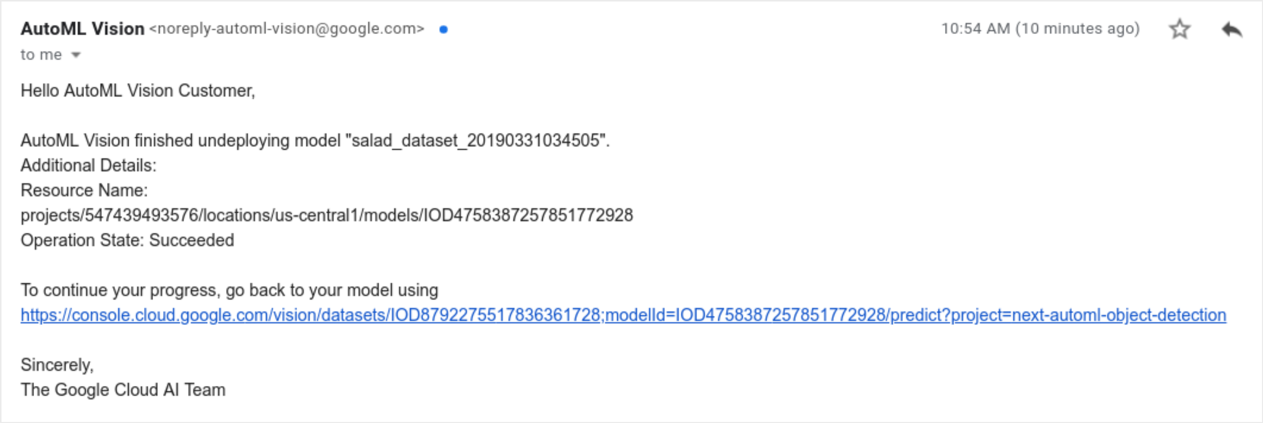
AutoML Vision <noreply-automl-vision@google.com> 10:54 AM (10 minutes ago)
to me

Hello AutoML Vision Customer,

AutoML Vision finished undeploying model "salad_dataset_20190331034505".
Additional Details:
Resource Name:
projects/547439493576/locations/us-central1/models/IOD4758387257851772928
Operation State: Succeeded

To continue your progress, go back to your model using
https://console.cloud.google.com/vision/datasets/IOD8792275517836361728;modelId=IOD4758387257851772928/predict?project=next-automl-object-detection

Sincerely,
The Google Cloud AI Team

Delete your project (optional)

To avoid unnecessary Google Cloud Platform charges, use the Cloud Console (https://console.cloud.google.com/) to delete your project if you do not need it.

Last updated November 20, 2019.