Cloud Vision API Product Search

General tips

This page describes troubleshooting steps that you might find helpful if you run into problems using Cloud Vision Product Search.

SEARCH RESULTS MORE

What is the maximum number of results returned?

The maximum number of results returned will vary from query to query due to implementation details. The maximum guaranteed amount of results returned is 500. If more are requested, the requested number may not be met.

Can I set a score threshold for results?

The search result scores are not calibrated. This means that the result scores for a given query will be comparable (this is how results are ranked). However, the result scores for query A and query B are *not necessarily* comparable. This means there is no one fixed threshold suitable for a given use case.

It is possible to set a threshold based on your own experiments. If you want to set a threshold, you should set this threshold on a per item label basis. This means that you should set a results score threshold for "Shoes" independently of "Bags", etc. for each label.

Is there a way to know why a particular result was ranked higher than another?

Currently the Vision API Product Search does not support any features that explain why a particular result was returned.

What is the difference between using `productGroupedResults` and using the Object Localizer feature of the Vision API?

The `productGroupedResults` response will return a list of detected items, their predicted label and score, their bounding boxes, and a ranked list of product search results for each item. This response is equivalent to first sending the query image to the Object Localizer feature of the Vision API and then sending the image along with each product item bounding box (filtered to the relevant labels) to Vision API Product Search.

One benefit of using `productGroupedResults` is that it is faster and more direct because it only makes a single API call. However, if there is some custom application logic that needs to be implemented between

the object detection and the search query then a two step method is a good approach. An example of such a case would be adding Vision Product Search filters based on object labels.

If I know the location of a product in an image should I crop the image before querying the API with it?

This situation might arise if, for example, a user is allowed to draw a bounding box around the item that they want to search. In general you should not crop the image before sending it to the API. The reason is that our API will use both the query bounding box to return targeted results as well as the overall image context, which is frequently helpful for improving the results.

Last updated November 12, 2019.