oduct is in a pre-release state and might change or have limited support. For more information, see the product laun
(/products/#product-launch-stages).

After you have created (trained) a model, you can request a prediction for one or more videos using the `batchPredict` (/video-intelligence/automl/docs/reference/rest/v1beta1/projects.locations.models/batchPredict) method. You supply a CSV file with a list of videos to the `batchPredict` method. The `batchPredict` method applies labels to your videos based on the primary objects of the videos that your model predicts.

The maximum lifespan for a custom model is two years. You must create and train a new model to continue classifying content after that amount of time.

For this release, only the **`us-central1`** location is supported.

To make it more convenient to run the `curl` samples in this topic, set the following environment variable. Replace ***project-id*** with the name of your GCP project.

To request a batch of predictions from AutoML Video Classification, create a CSV file that lists the Google Cloud Storage paths to the videos that you want to annotate. You can also specify a start and end time to tell AutoML Video Classification to only annotate a segment of the video. The start time must be zero or greater and must be before the end time. The end time must be greater than the start time and less than or equal to the duration of the video. You can also use `inf` to indicate the end of a video.

You must also specify an output file path where AutoML Video Classification will write the results of the predictions from your model. This path must be a Google Cloud Storage bucket and object where you have write permissions.

Each video can be up to 3 hours in duration with a maximum file size of 50GB. AutoML Video Classification can generate predictions for approximately 100 hours of video in 12 hours of processing time.

You can also set the following options in the `params` section when you request a prediction for your videos. If you do not specify any of these options, then the default score threshold is applied, and **segment-classification** is used.

- **score_threshold** - A value from 0.0 to 1.0. When the model makes predictions for a video, it will only produce results that have at least this confidence score. The default for the API is 0.5.

★ **Note:** The default for the UI is **0** and not **0.5**, therefore the user will get more results with the default UI value.

- **segment-classification** - Set to true to enable segment-level classification. AutoML Video Classification returns labels and their confidence scores for the entire segment of video that you specified in your request configuration. The default is true.
- **shot-classification** - Set to true to enable shot-level classification. AutoML Video Classification determines the boundaries for each camera shot in the entire segment of video that you specified in your request configuration. AutoML Video Intelligence then returns labels and their confidence scores for each shot detected, along with the start and end time of the shot. The default is false.
- **1s_interval** - Set to true to enable classification for a video at one-second intervals. AutoML Video Classification returns labels and their confidence scores for each second of the entire segment of video that you specified in your request configuration. The default is false.

tant: The error message “Model is warming up. Please try again later.” indicates you should wait a few minutes befor
to query the model again. This warm-up delay can occur when you initially query your model, or when you make a rec
model that has not been queried recently.