

H2O.ai



DOCUMENTATION

---

For questions, please contact [support@h2o.ai](mailto:support@h2o.ai)

# Contents

<b>Documentation structure and philosophy</b>	<b>8</b>
Structure . . . . .	8
Philosophy . . . . .	8
<b>What is H2O Label Genie?</b>	<b>9</b>
<b>Access H2O Label Genie</b>	<b>10</b>
Step 1: Access HAIC . . . . .	10
Step 2: Search H2O Label Genie . . . . .	10
Step 3: Run H2O Label Genie . . . . .	10
Step 4: H2O Label Genie instance . . . . .	10
Pause or terminate instance . . . . .	11
<b>Labeling flow</b>	<b>13</b>
Overview . . . . .	13
Step 1: (Optional) Explore dataset . . . . .	13
Step 2: Create an annotation task . . . . .	14
Step 3: Specify an annotation task rubric . . . . .	14
Step 4: Annotate dataset (with AI assistance) . . . . .	14
<b>Tutorials</b>	<b>17</b>
Learning path . . . . .	17
Text annotation tasks . . . . .	18
Image annotation tasks . . . . .	18
Audio annotation tasks . . . . .	18
<b>Tutorial 1A: Text classification annotation task</b>	<b>19</b>
Overview . . . . .	19
Step 1: Explore dataset . . . . .	19
Step 2: Create an annotation task . . . . .	19
Step 3: Specify an annotation task rubric . . . . .	19
Step 4: Annotate dataset . . . . .	20
Summary . . . . .	21
Next . . . . .	21
<b>Tutorial 2A: Text regression annotation task</b>	<b>22</b>
Overview . . . . .	22
Step 1: Explore dataset . . . . .	22
Step 2: Create an annotation task . . . . .	22
Step 3: Specify an annotation task rubric . . . . .	22
Step 4: Annotate dataset . . . . .	23
Summary . . . . .	24
Next . . . . .	24
<b>Tutorial 3A: Text-entity recognition annotation task</b>	<b>25</b>
Overview . . . . .	25
Step 1: Explore dataset . . . . .	25
Step 2: Create an annotation task . . . . .	25
Step 3: Specify an annotation task rubric . . . . .	25
Step 4: Annotate dataset . . . . .	25
Summary . . . . .	26
Next . . . . .	26
<b>Tutorial 4A: Text summarization annotation task</b>	<b>27</b>
Overview . . . . .	27
Step 1: Explore dataset . . . . .	27
Step 2: Create an annotation task . . . . .	27
Step 3: Specify an annotation task rubric . . . . .	27

Step 4: Annotate dataset . . . . .	28
Summary . . . . .	28
Next . . . . .	28
<b>Tutorial 5A: Text-generative AI annotation task</b>	<b>29</b>
Overview . . . . .	29
Step 1: Explore dataset . . . . .	29
Step 2: Create an annotation task . . . . .	29
Step 3: Specify an annotation task rubric . . . . .	29
Step 4: Annotate dataset . . . . .	30
Summary . . . . .	31
Next . . . . .	31
<b>Tutorial 1B: Image classification annotation task</b>	<b>32</b>
Overview . . . . .	32
Step 1: Explore dataset . . . . .	32
Step 2: Create an annotation task . . . . .	32
Step 3: Specify an annotation task rubric . . . . .	32
Step 4: Annotate dataset . . . . .	33
Summary . . . . .	34
Next . . . . .	34
<b>Tutorial 2B: Image regression annotation task</b>	<b>35</b>
Overview . . . . .	35
Step 1: Explore dataset . . . . .	35
Step 2: Create an annotation task . . . . .	35
Step 3: Specify an annotation task rubric . . . . .	35
Step 4: Annotate dataset . . . . .	36
Summary . . . . .	37
Next . . . . .	37
<b>Tutorial 3B: Object detection annotation task</b>	<b>38</b>
Overview . . . . .	38
Step 1: Explore dataset . . . . .	38
Step 2: Create an annotation task . . . . .	38
Step 3: Specify an annotation task rubric . . . . .	38
Step 4: Annotate dataset . . . . .	39
Summary . . . . .	39
Next . . . . .	40
<b>Tutorial 4B: Image instance segmentation annotation task</b>	<b>41</b>
Overview . . . . .	41
Step 1: Explore dataset . . . . .	41
Step 2: Create an annotation task . . . . .	41
Step 3: Specify an annotation task rubric . . . . .	41
Step 4: Annotate dataset . . . . .	42
Summary . . . . .	43
Next . . . . .	43
<b>Tutorial 1C: Audio classification annotation task</b>	<b>44</b>
Overview . . . . .	44
Step 1: Explore dataset . . . . .	44
Step 2: Create an annotation task . . . . .	44
Step 3: Specify an annotation task rubric . . . . .	44
Step 4: Annotate dataset . . . . .	45
Summary . . . . .	46
Next . . . . .	46
<b>Tutorial 2C: Audio regression annotation task</b>	<b>47</b>
Overview . . . . .	47

Step 1: Explore dataset . . . . .	47
Step 2: Create an annotation task . . . . .	47
Step 3: Specify an annotation task rubric . . . . .	47
Step 4: Annotate dataset . . . . .	48
Summary . . . . .	49
Next . . . . .	49
<b>Import a dataset</b>	<b>50</b>
Overview . . . . .	50
Instructions . . . . .	50
Dataset connectors . . . . .	51
Overview . . . . .	51
Supported connectors . . . . .	51
Supported <i>image</i> extensions . . . . .	51
Supported <i>audio</i> extensions . . . . .	52
<b>Import an annotated dataset</b>	<b>53</b>
Overview . . . . .	53
Instructions . . . . .	53
Required annotated dataset format . . . . .	53
Text classification . . . . .	53
Text regression . . . . .	53
Text-entity recognition . . . . .	54
Text summarization . . . . .	54
Text-generative AI . . . . .	54
Image classification . . . . .	54
Image regression . . . . .	54
Object detection . . . . .	55
Image instance segmentation . . . . .	55
Audio classification . . . . .	57
Audio regression . . . . .	57
<b>Access a demo dataset</b>	<b>58</b>
Overview . . . . .	58
Instructions . . . . .	58
Demo datasets in H2O Label Genie . . . . .	58
Amazon reviews demo . . . . .	58
Car or coffee demo . . . . .	58
Twitter demo . . . . .	58
Text readability demo . . . . .	59
CNN Daily Mail sample . . . . .	59
Plant pathology demo . . . . .	59
ESC10 audio demo . . . . .	59
Amnist demo . . . . .	59
<b>Edit a dataset</b>	<b>60</b>
Overview . . . . .	60
<b>View a dataset</b>	<b>61</b>
Overview . . . . .	61
Instructions . . . . .	61
<b>Delete a dataset</b>	<b>62</b>
Overview . . . . .	62
Instructions . . . . .	62
<b>Download an annotated dataset</b>	<b>63</b>
Overview . . . . .	63
Instructions . . . . .	63
Downloaded dataset formats . . . . .	63

Text classification . . . . .	63
Text-entity recognition . . . . .	64
Text-generative AI . . . . .	64
Image regression . . . . .	65
Image instance segmentation . . . . .	66
Audio regression . . . . .	67
<b>Download a dataset's zero-shot predictions</b>	<b>68</b>
Overview . . . . .	68
Instructions . . . . .	68
<b>Export an annotated dataset to H2O Drive</b>	<b>69</b>
Overview . . . . .	69
Instructions . . . . .	69
<b>Export a dataset's zero-shot predictions to H2O Drive</b>	<b>70</b>
Overview . . . . .	70
Instructions . . . . .	70
<b>Supported clustering tasks</b>	<b>71</b>
Overview . . . . .	71
Image clustering . . . . .	71
Text clustering . . . . .	71
<b>Create a clustering task</b>	<b>73</b>
Overview . . . . .	73
Instructions . . . . .	73
<b>View a clustering task</b>	<b>74</b>
Overview . . . . .	74
Instructions . . . . .	74
<b>Clustering task tabs</b>	<b>75</b>
Overview . . . . .	75
Gallery . . . . .	75
Map . . . . .	76
<b>Download a clustering task's cluster labels</b>	<b>78</b>
Overview . . . . .	78
Instructions . . . . .	78
<b>Delete a clustering task</b>	<b>79</b>
Overview . . . . .	79
Instructions . . . . .	79
<b>Supported annotation tasks</b>	<b>80</b>
Overview . . . . .	80
Text annotation tasks . . . . .	80
Text classification . . . . .	80
Text regression . . . . .	80
Text-entity recognition . . . . .	80
Text summarization . . . . .	80
Text-generative AI . . . . .	80
Image annotation tasks . . . . .	80
Image classification . . . . .	80
Image regression . . . . .	80
Object detection . . . . .	81
Image instance segmentation . . . . .	81
Audio annotation tasks . . . . .	81
Audio classification . . . . .	81

Audio regression . . . . .	81
<b>Hotkeys</b>	<b>82</b>
Overview . . . . .	82
Instructions . . . . .	82
<b>Zero-shot learning models</b>	<b>83</b>
Overview . . . . .	83
What are zero-shot learning models? . . . . .	83
What are zero-shot learning model predictions? . . . . .	83
Annotation tasks + zero-shot learning models . . . . .	83
Image classification . . . . .	83
Object detection . . . . .	83
Image instance segmentation . . . . .	83
Text classification . . . . .	83
Text summarization . . . . .	83
Text-generative AI . . . . .	84
Select a zero-shot learning model . . . . .	84
<b>Create an annotation task</b>	<b>85</b>
Overview . . . . .	85
Instructions . . . . .	85
<b>Access an annotation task's tabs</b>	<b>86</b>
Overview . . . . .	86
Instructions . . . . .	86
Tabs . . . . .	86
Dashboard . . . . .	86
Rubric . . . . .	86
Statistics . . . . .	87
Annotate . . . . .	87
Review . . . . .	88
Import . . . . .	88
Export . . . . .	88
<b>Specify an annotation task rubric</b>	<b>90</b>
Overview . . . . .	90
Instructions . . . . .	90
Text annotation tasks . . . . .	90
Image annotation tasks . . . . .	93
Audio annotation tasks . . . . .	95
Large language model (LLM) parameters . . . . .	96
Select API endpoint type . . . . .	96
LLM model name . . . . .	96
Max response tokens . . . . .	96
Temperature . . . . .	96
Repetition penalty . . . . .	97
Prompt template . . . . .	97
Select example prompt . . . . .	97
<b>Edit an annotation task rubric</b>	<b>98</b>
Overview . . . . .	98
Instructions . . . . .	98
<b>View an annotation task</b>	<b>99</b>
Overview . . . . .	99
Instructions . . . . .	99
<b>Delete an annotation task</b>	<b>100</b>
Overview . . . . .	100

Instructions . . . . .	100
<b>Segment objects</b>	<b>101</b>
Overview . . . . .	101
Instructions . . . . .	101
<b>Appearance settings</b>	<b>103</b>
Overview . . . . .	103
Dark mode . . . . .	103
Overview . . . . .	103
Instructions . . . . .	103
Upload annotation instructions . . . . .	103
Overview . . . . .	103
Instructions . . . . .	104
Colors for annotations . . . . .	104
Overview . . . . .	104
Instructions . . . . .	104
<b>Smart annotation settings</b>	<b>105</b>
Overview . . . . .	105
Enable zero-shot predictions . . . . .	105
Overview . . . . .	105
Instructions . . . . .	105
Enable segment anything . . . . .	105
Overview . . . . .	105
Instructions . . . . .	106
Enable keyboard support . . . . .	106
Overview . . . . .	106
Instructions . . . . .	106
<b>Image settings</b>	<b>107</b>
Overview . . . . .	107
Max image width [px] . . . . .	107
Overview . . . . .	107
Instructions . . . . .	107
Max image height [px] . . . . .	107
Overview . . . . .	107
Instructions . . . . .	107
<b>Audio settings</b>	<b>108</b>
Overview . . . . .	108
Max audio duration [s] . . . . .	108
Overview . . . . .	108
Instructions . . . . .	108
<b>OpenAI API settings</b>	<b>109</b>
Overview . . . . .	109
OpenAI API key . . . . .	109
Overview . . . . .	109
Instructions . . . . .	109
Use OpenAI API on Azure . . . . .	109
Overview . . . . .	109
Instructions . . . . .	109
OpenAI API Endpoint . . . . .	109
Overview . . . . .	109
Instructions . . . . .	109
OpenAI API deployment ID . . . . .	109
Instructions . . . . .	110
OpenAI API version . . . . .	110
Overview . . . . .	110

Instructions . . . . .	110
<b>Collaboration settings</b>	<b>111</b>
Overview . . . . .	111
Multi-user annotation . . . . .	111
Overview . . . . .	111
Instructions . . . . .	111
<b>Key terms</b>	<b>112</b>
Annotation task . . . . .	112
Annotation task rubric . . . . .	112
Clustering task . . . . .	112
<b>AI governance</b>	<b>113</b>
What is AI governance? . . . . .	113
H2O.ai's AI governance with H2O Label Genie . . . . .	113
<b>Release notes</b>	<b>114</b>
v1.2.0   March 20, 2025 . . . . .	114
Fixes . . . . .	114
<b>FAQs</b>	<b>115</b>
General . . . . .	115
What is the difference between a CPU and GPU bundle of H2O Label Genie? . . . . .	115

# Documentation structure and philosophy

## Structure

The H2O Label Genie documentation is organized into the following main sections:

- **Get started:** This section introduces H2O Label Genie.
- **Tutorials:** This section offers tutorials on H2O Label Genie.
- **Guide:** This section encompasses several subsections:
  - **Datasets:** This subsection covers datasets in H2O Label Genie.
  - **Data exploration:** This subsection covers clustering tasks in H2O Label Genie.
  - **Annotation tasks:** This subsection covers annotation tasks in H2O Label Genie.
  - **App settings:** This subsection covers the application settings of H2O Label Genie.
- **Key terms:** This section defines key terms relevant to H2O Label Genie.
- **AI governance:** This section discusses how H2O Label Genie can help with AI governance.
- **Release notes:** This section discusses the release notes of a particular H2O Label Genie version.
- **Third-party licenses:** This section includes all third-party licenses H2O Label Genie utilizes.
- **FAQs:** This section answers frequently asked questions about H2O Label Genie.

## Philosophy

- **Philosophy of alignment:** The documentation for H2O Label Genie is designed to match the product's structure. This means that information in the documentation is organized to align with the features and workflows of H2O Label Genie. Our goal is to make learning the product easy and intuitive for users.
- **Simplicity and modularity principles:** Our documentation follows a simple and modular approach. Most pages have two main sections: "Overview" and "Instructions." A page can also have other sections that heavily relate to the page's "Overview" and "Instructions." This makes it easier for users to understand and use H2O Label Genie effectively.

## What is H2O Label Genie?

H2O Label Genie is an application that enables you with zero-shot learning models to rapidly label your datasets for [annotation tasks](#) in computer vision (CV), natural language processing (NLP), and audio.

## Access H2O Label Genie

You can access H2O Label Genie through an instance that you can create on the H2O AI Cloud (**HAIC**). To access H2O Label Genie:

- [Step 1: Access HAIC](#)
- [Step 2: Search H2O Label Genie](#)
- [Step 3: Run H2O Label Genie](#)
- [Step 4: H2O Label Genie instance](#)

### Step 1: Access HAIC

Access your H2O AI Cloud (**HAIC**) account.

### Step 2: Search H2O Label Genie

1. In HAIC, click **APP STORE**.
2. In the **HAIC** search bar, search **H2O Label Genie**.
3. Select the H2O Label Genie tile. Details about H2O Label Genie appear.

The screenshot shows the H2O AI Cloud App Store interface. At the top, there's a navigation bar with 'H2O.ai' and 'APP STORE' highlighted. Below the navigation bar, there's a search bar containing 'H2O Label Genie'. The search results show a single tile for 'H2O Label Genie' by 'pascal.pfeiffer@h2o.ai'. The tile features a black background with the text 'Label Genie' in yellow. To the left of the main content, there's a 'Categories' sidebar with various industry categories like Data Science, Document, Financial Services, etc.

### Step 3: Run H2O Label Genie

1. To start an H2O Label Genie instance, click **Run**.

The screenshot shows the app details page for 'H2O Label Genie'. The app title is 'H2O Label Genie' with the subtitle 'Supercharge your labeling workflow'. There's a 'Run' button in the top right corner, along with a version dropdown set to '0.2.0' and a 'Send feedback' button. Below the app details, there's a section for 'App Details' with tabs for 'My Instances 0' and 'All Instances 2'. At the bottom, there are three preview images showing the app's interface.

### Step 4: H2O Label Genie instance

Now, the H2O AI Cloud is starting an instance of H2O Label Genie for you. While you have a starting/running instance, the **Run** button will change its name to **Visit**.

1. To open H2O Label Genie in a new tab, click **Visit**.

**Note:**

- The latest version of H2O Label Genie is preselected.
- In the H2O Label Genie instance, several items will be installed. Right after, you will be able to use H2O Label Genie. All items are automatically installed when you start an instance.

## Pause or terminate instance

You can pause or terminate an instance of H2O Label Genie.

- **Pause:** Pausing an instance reduces computational resources (and is less expensive). In other words, the cost of having an instance decreases.
- **Terminate:** Terminating an instance deletes the instance permanently.

**Note:** Customers pay for H2O AI Cloud via AI Units so that as you consume more resources, you pay more.

### Pause

You can **Pause** an instance in the **app details** page:

- In the **app details** page, click the **All instances** tab.
- Locate your instance row.
- Click **Pause**.

H2O Label Genie v0.2.0	Owner	Created At	Updated At
ID: 077740b9-695c-43ad-a7f4-c841c03a3c0b	oshini.nugapiti...	28/09/2022 12:14:09	05/10/2022 09:56:51

You can also **pause** an instance in the **My instances** page:

- In the **H2O AI Cloud** menu, click **My instances**.
- Locate your instance row.
- Click **Pause**.

**My Instances**  
You have 1 instances

Filter by app

Label Genie	H2O Label Genie v0.2.0	Created At	Updated At	Instance log	Visit	Pause
	ID: 077740b9-695c-43ad-a7f4-c841c03a3c0b Visibility: All Users   Status: Deployed   Instance will be automatically paused in one hour	28/09/2022 12:14:09	28/09/2022 12:14:09			

**Terminate** You can **terminate** an instance in the **app details** page:

- In the **app details** page, click the **All instances** tab.
- Locate your instance row.
- Click **Expand**.
- Click **Terminate**.

**H2O Label Genie**  
Supercharge your labeling workflow

OTHER

pascal.pfeiffer@h2o.ai

App Details | My Instances 1 | **All Instances 3**

Label Genie	H2O Label Genie v0.2.0	Owner	Created At	Updated At	Visit	Pause
	ID: 077740b9-695c-43ad-a7f4-c841c03a3c0b Visibility: All Users   Status: Deployed   Instance will be automatically paused in one hour	oshini.nugap...	28/09/2022 12:14:09	29/09/2022 14:21:25		
	H2O Label Genie v0.2.0	oshini.nugap...	28/09/2022	29/09/2022		

Dropdown menu options: Terminate, Postpone Auto-Pause

You can also **terminate** an instance in the **My instance** page:

- In the **H2O AI Cloud** menu, click **My instances**.
- Locate your instance row.
- Click **Expand**.
- Click **Terminate**.

**My Instances**  
You have 1 instances

Filter by app

Label Genie	H2O Label Genie v0.2.0	Created At	Updated At	Instance log	Visit	Pause
	ID: 077740b9-695c-43ad-a7f4-c841c03a3c0b Visibility: All Users   Status: Deployed   Instance will be automatically paused in one hour	28/09/2022 12:14:09	29/09/2022 14:21:25			

Dropdown menu options: Terminate, Postpone Auto-Pause

# Labeling flow

## Overview

The flow of labeling a dataset in H2O Label Genie can be summarized in the following sequential steps:

- [Step 1: \(Optional\) Explore dataset](#)
- [Step 2: Create an annotation task](#)
- [Step 3: Specify an annotation task rubric](#)
- [Step 4: Annotate dataset \(with AI assistance\)](#)

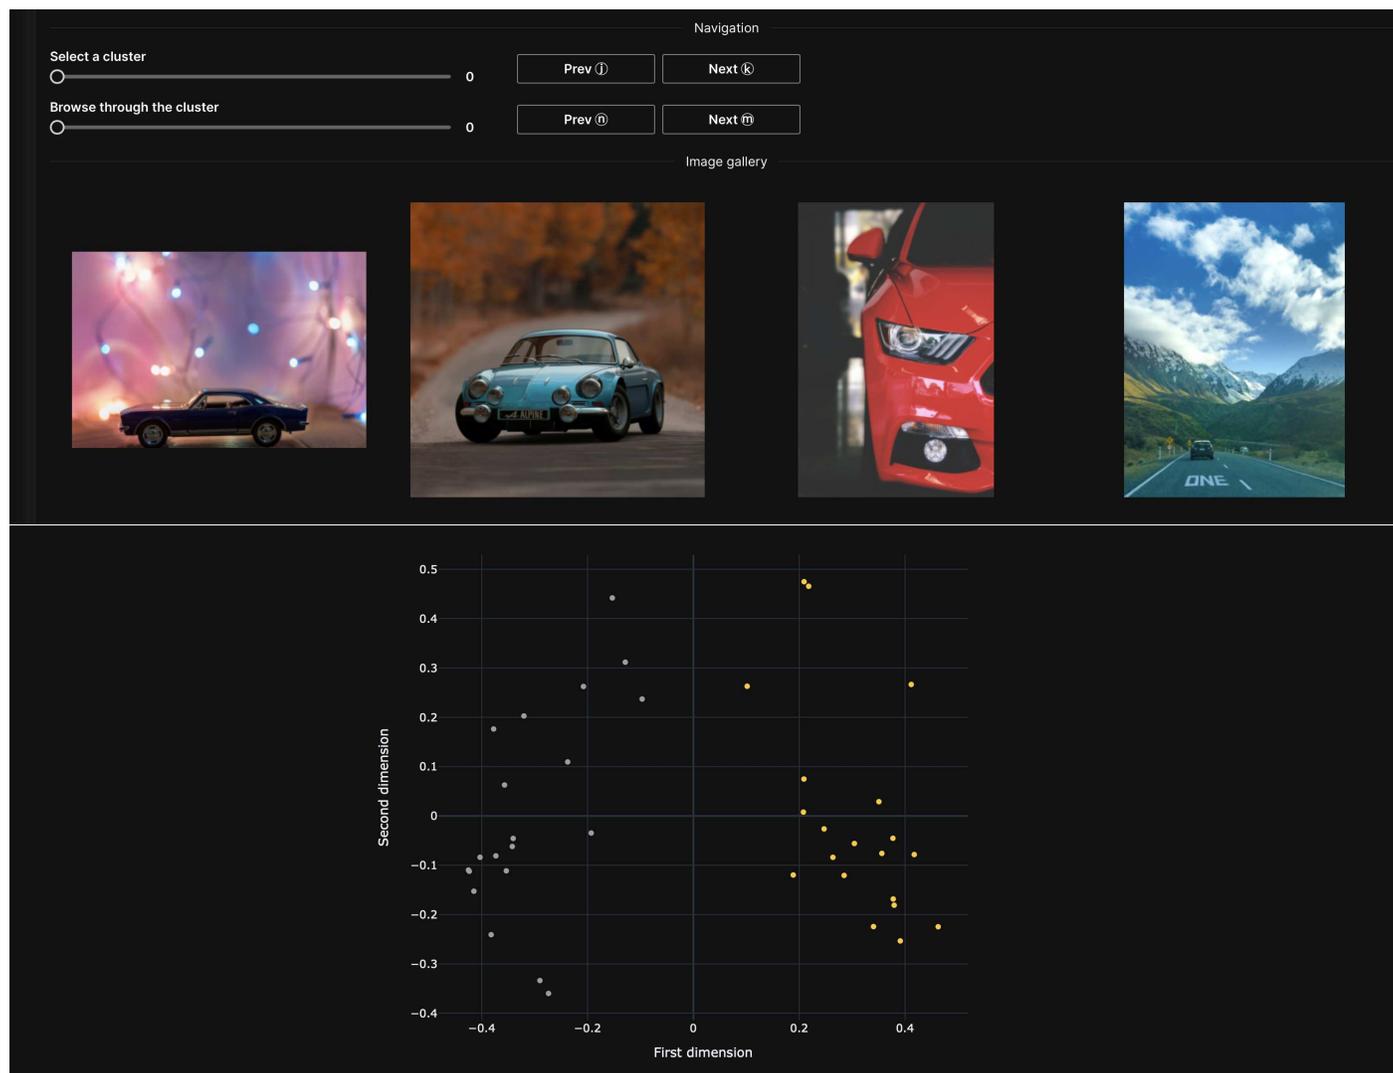
In the below sections, each step above, in turn, is summarized.

## Step 1: (Optional) Explore dataset

As the first step in the labeling flow (an optional step), before you start labeling your dataset, you can explore patterns and groups in your dataset using unsupervised methods, resulting in improved quality of your data. In particular, to provide in-depth knowledge of image and text datasets, H2O Label Genie supports the analysis of such datasets through clustering tasks that, for example, generate embeddings in 2D and 3D that can help you understand the data structure of the dataset.

**Clustering task:** A clustering task refers to finding and exploring groups in a dataset.

- To learn more about clustering tasks, see [Supported clustering tasks](#)



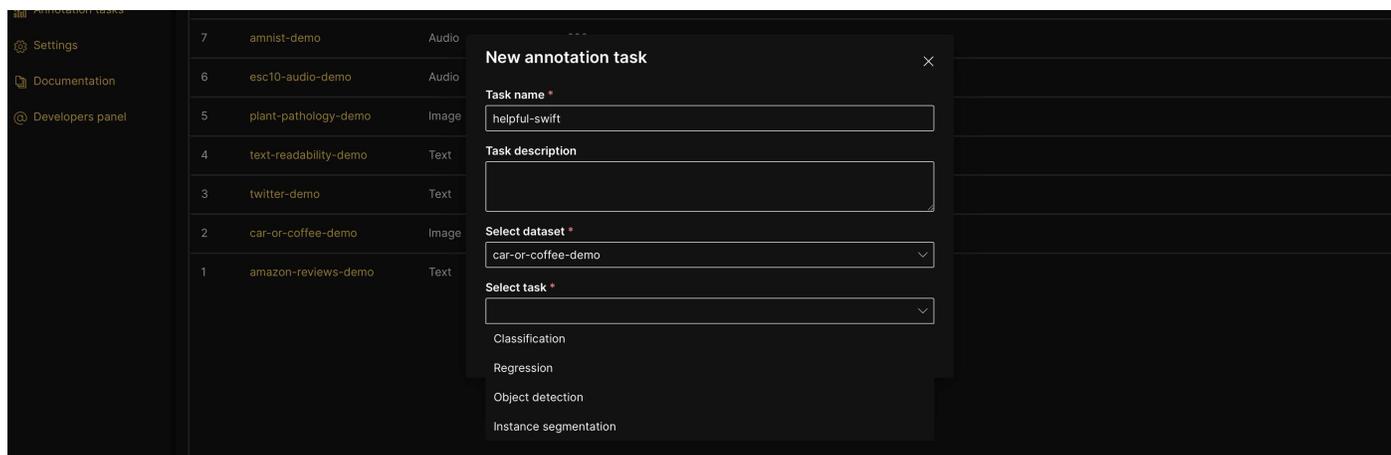
## Step 2: Create an annotation task

As the second step in the labeling flow (or the first one if you skipped the optional step 1), you need to create an annotation task.

**Annotation task:** An annotation task refers to labeling data in a manner that makes them suitable to support an array of deep learning problem types. This process can, for example, involve adding bounding boxes to images where you can attribute labels to each box.

H2O Label Genie supports various annotation tasks in computer vision (CV) (image), natural language processing (NLP) (text), and audio.

- To learn about supported annotation tasks, see [Supported annotation tasks](#)
- To learn how to create an annotation task, see [Create an annotation task](#)

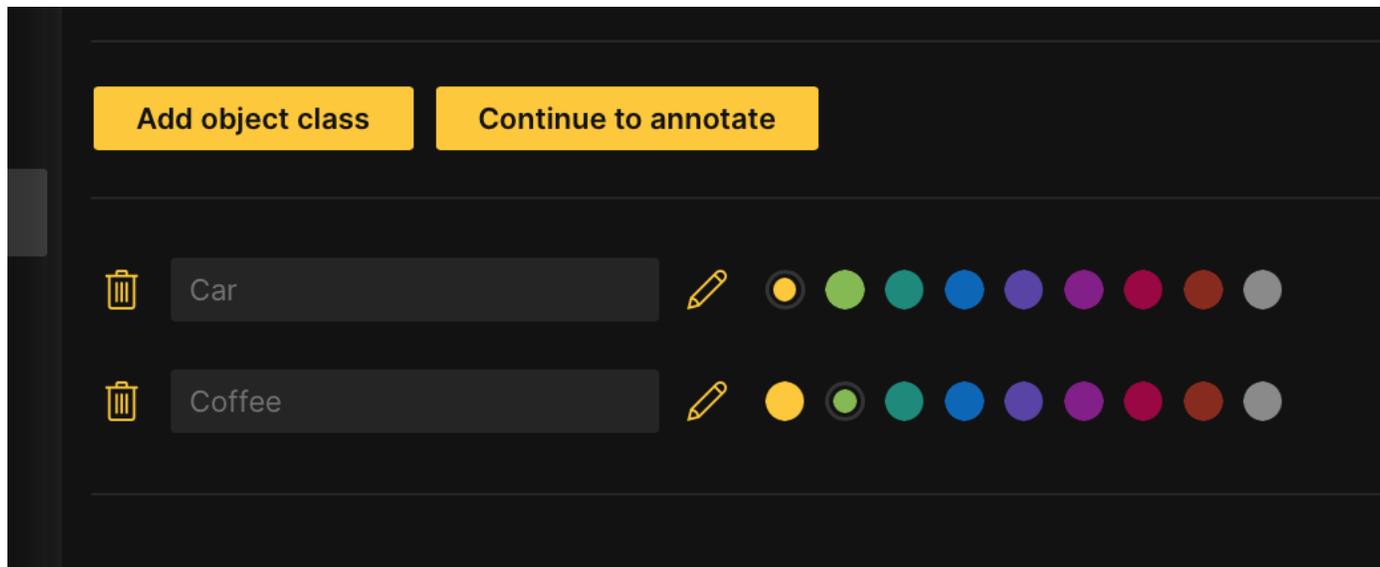


## Step 3: Specify an annotation task rubric

As the third step in the labeling flow, you need to specify an annotation task rubric.

**Annotation task rubric:** An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset.

- To learn how to create an annotation task rubric, see [Specify an annotation task rubric](#)
- To learn how to edit an annotation task rubric, see [Edit an annotation task rubric](#)



## Step 4: Annotate dataset (with AI assistance)

As the fourth step in the labeling flow, annotate your dataset. Each annotation task for a dataset requires a different process for annotating the dataset. To learn more, see [Tutorials](#).

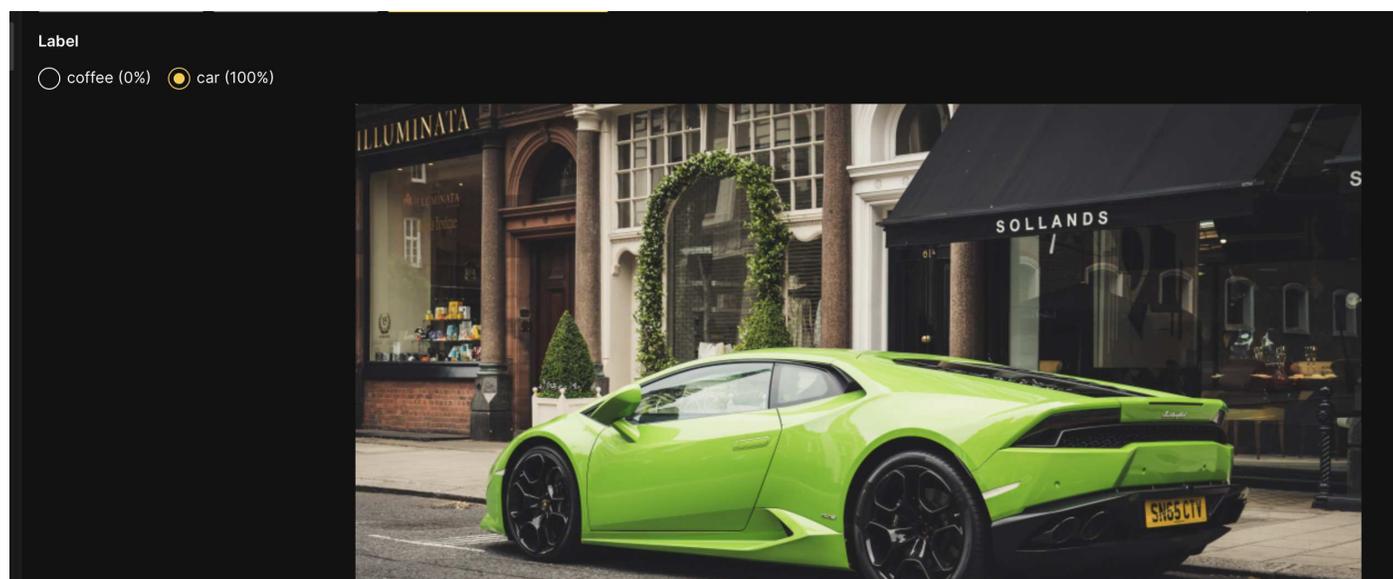
**Note:** An array of datasets labeled in H2O Label Genie are supported in [H2O Hydrogen Torch](#) and [H2O LLM Studio](#). To learn more, see [Download an annotated dataset](#).

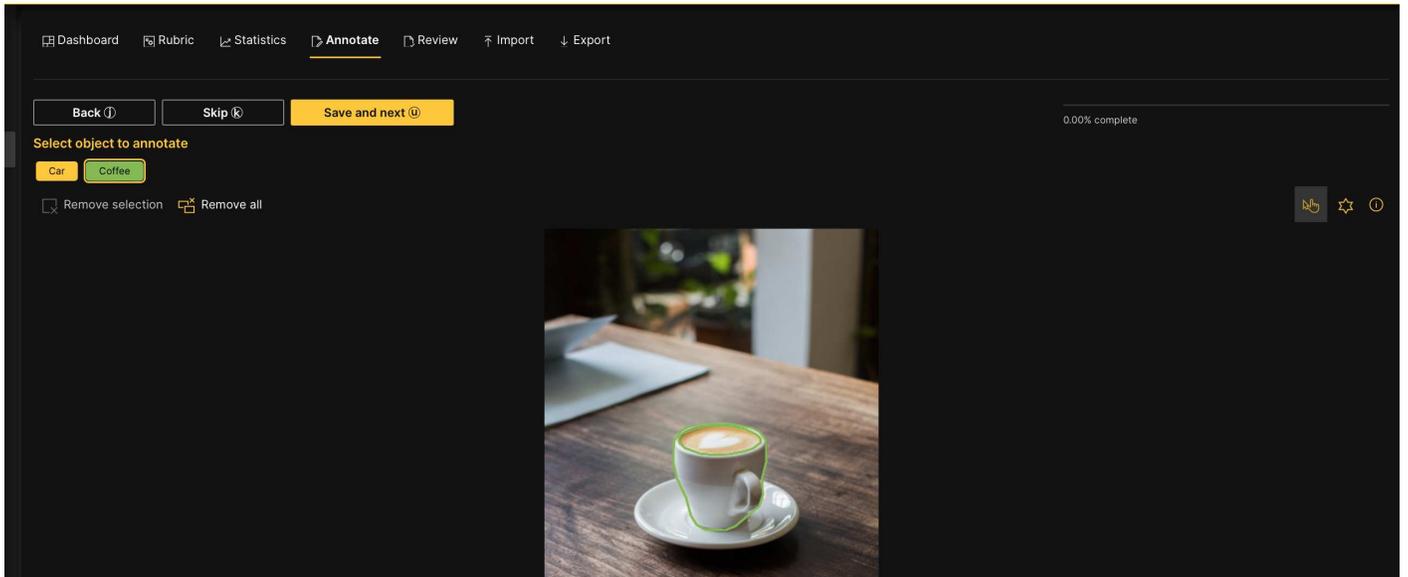
H2O Label Genie offers the following major features to speed up the labeling process:

- **Zero-shot learning models:** By default, H2O Label Genie utilizes certain zero-shot learning models to accelerate the labeling process. In particular, H2O Label Genie lets you use a zero-shot learning model for several supported annotation tasks
  - To learn more, see [Zero-shot learning models](#)
- **Hotkeys:** H2O Label Genie supports several hotkeys (keyboard shortcuts) designed to speed up a dataset's annotation (labeling) process
  - To learn more, see [Hotkeys](#)
- **Real-time multi-user support and collaboration:** By default, while annotating your dataset, H2O Label Genie enables you to work with others to complete an annotation task, resulting in real-time multi-user support and collaboration
  - To learn more, see [Collaboration settings](#)

**Zero-shot learning models:** Labeled data is crucial for supervised learning problem types in computer vision (CV), natural language processing (NLP), and audio. High-quality labeled data usually requires a lot of manual labeling that can lead to high costs and delay production or execution.

One way to accelerate the labeling process is to utilize zero-shot learning models. These models let data scientists label unlabeled data with high accuracy and speed. Zero-shot learning models are pre-trained models that have been trained on vast and distinct classes. As a result, zero-shot learning models with prior knowledge can label unlabeled data.





## Tutorials

Learn how to generate labeled datasets in H2O Label Genie

### Learning path

graph LR;

```
User[User] --> A1[Text annotation tasks];
```

```
User --> A2[Image annotation tasks];
```

```
User --> A3[Audio annotation tasks];
```

```
A1 --> B1[Tutorial 1A: Text classification annotation task];
```

```
A1 --> B2[Tutorial 2A: Text regression annotation task];
```

```
A1 --> B3[Tutorial 3A: Text-entity recognition annotation task];
```

```
A1 --> B4[Tutorial 4A: Text summarization annotation task];
```

```
A1 --> B5[Tutorial 5A: Text-generative AI annotation task];
```

```
A2 --> C1[Tutorial 1B: Image classification annotation task];
```

```
A2 --> C2[Tutorial 2B: Image regression annotation task];
```

```
A2 --> C3[Tutorial 3B: Object detection annotation task];
```

```
A2 --> C4[Tutorial 4B: Image instance segmentation annotation task];
```

```
A3 --> D1[Tutorial 1C: Audio classification annotation task];
```

```
A3 --> D2[Tutorial 2C: Audio regression annotation task];
```

```
%% Apply custom color to nodes
```

```
style User fill:#FEC925;
```

```
style A1 fill:#FEC925;
```

```
style A2 fill:#FEC925;
```

```
style A3 fill:#FEC925;
```

```
style B1 fill:#FEC925;
```

```
style B2 fill:#FEC925;
```

```
style B3 fill:#FEC925;
```

```
style B4 fill:#FEC925;
```

```
style B5 fill:#FEC925;
```

```
style C1 fill:#FEC925;
```

```
style C2 fill:#FEC925;
```

```
style C3 fill:#FEC925;
```

```
style C4 fill:#FEC925;
```

```
style D1 fill:#FEC925;
```

```
style D2 fill:#FEC925;
```

```
%% Add links to each node
```

```
click B1 "/wave-apps/h2o-label-genie/tutorials/text/tutorial-1-a" "Go to Tutorial 1A"
```

```
click B2 "/wave-apps/h2o-label-genie/tutorials/text/tutorial-2-a" "Go to Tutorial 2A"
```

```
click B3 "/wave-apps/h2o-label-genie/tutorials/text/tutorial-3-a" "Go to Tutorial 3A"
```

```
click B4 "/wave-apps/h2o-label-genie/tutorials/text/tutorial-4-a" "Go to Tutorial 4A"
```

```
click B5 "/wave-apps/h2o-label-genie/tutorials/text/tutorial-5-a" "Go to Tutorial 5A"
```

```
click C1 "/wave-apps/h2o-label-genie/tutorials/image/tutorial-1-b" "Go to Tutorial 1B"
```

```
click C2 "/wave-apps/h2o-label-genie/tutorials/image/tutorial-2-b" "Go to Tutorial 2B"
```

```
click C3 "/wave-apps/h2o-label-genie/tutorials/image/tutorial-3-b" "Go to Tutorial 3B"
```

```
click C4 "/wave-apps/h2o-label-genie/tutorials/image/tutorial-4-b" "Go to Tutorial 4B"
```

```
click D1 "/wave-apps/h2o-label-genie/tutorials/audio/tutorial-1-c" "Go to Tutorial 1C"
```

```
click D2 "/wave-apps/h2o-label-genie/tutorials/audio/tutorial-2-c" "Go to Tutorial 2C"
```

**Note:** To learn how H2O Label Genie can help annotate data to build and deploy a model with [H2O Hydrogen Torch](#) and [H2O MLOps](#), refer to the following blog: [In H2O AI Cloud, build, deploy, and score a state-of-the-art image classification model, starting with unlabeled data.](#)

## Text annotation tasks

- [Tutorial 1A: Text classification annotation task](#) > This tutorial describes the process of creating a text classification annotation task, including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing user reviews (in text format) and ratings (from 0 to 5) of Amazon products.
- [Tutorial 2A: Text regression annotation task](#) > This tutorial describes the process of creating a text regression annotation task, including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing user reviews (in text format) and ratings (from 0 to 5) of Amazon products.
- [Tutorial 3A: Text-entity recognition annotation task](#) > This tutorial describes the process of creating a text-entity recognition annotation task, including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing user reviews (in text format) and ratings (from 0 to 5) of Amazon products.
- [Tutorial 4A: Text summarization annotation task](#) > This tutorial describes the process of creating a text summarization annotation task. To highlight the process, we are going to annotate a dataset that contains human-generated abstract summaries from news stories published on the Cable News Network (CNN) and Daily Mail websites.
- [Tutorial 5A: Text-generative AI annotation task](#) > This tutorial describes the process of creating a text-generative AI annotation task, including specifying an annotation task rubric for it. To highlight the process, we will utilize the Amazon reviews dataset that contains user reviews (in text format) and ratings (from 0 to 5) of Amazon products. In particular, we will utilize an H2O.ai zero-shot learning model (large language model (LLM)) to summarize the product reviews.

## Image annotation tasks

- [Tutorial 1B: Image classification annotation task](#) > This tutorial describes the process of creating an image classification annotation task, including specifying an annotation task rubric for it. To highlight the process, we are going to annotate a dataset that contains images of cars and coffee.
- [Tutorial 2B: Image regression annotation task](#) > This tutorial describes the process of creating an image regression annotation task, including specifying an annotation task rubric for it. To highlight the process, we are going to annotate a dataset that contains images of healthy and diseased apple leaves for plant pathology recognition.
- [Tutorial 3B: Object detection annotation task](#) > This tutorial describes the process of creating an object detection annotation task, including specifying an annotation task rubric for it. To highlight the process, we are going to annotate a dataset that contains images of cars and coffee.
- [Tutorial 4B: Image instance segmentation annotation task](#) > This tutorial describes the process of creating an image instance segmentation annotation task, including specifying an annotation task rubric for it. To highlight the process, we are going to annotate a dataset that contains images of cars and coffee.

## Audio annotation tasks

- [Tutorial 1C: Audio classification annotation task](#) > This tutorial describes the process of creating an audio classification annotation task, including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing 5-second-long recordings of environmental sounds organized into ten classes (with 40 examples per class).
- [Tutorial 2C: Audio regression annotation task](#) > This tutorial describes the process of creating an audio regression annotation task, including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing 600 audio samples of spoken digits (0-9) of 60 different speakers.

## Tutorial 1A: Text classification annotation task

### Overview

This tutorial describes the process of creating a [text classification annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing user reviews (in text format) and ratings (from 0 to 5) of Amazon products.

### Step 1: Explore dataset

We are going to use the preloaded [Amazon reviews demo](#) dataset for this tutorial. The dataset contains **180** samples (text), each containing a review of an Amazon product. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **amazon-reviews-demo**.

### Step 2: Create an annotation task

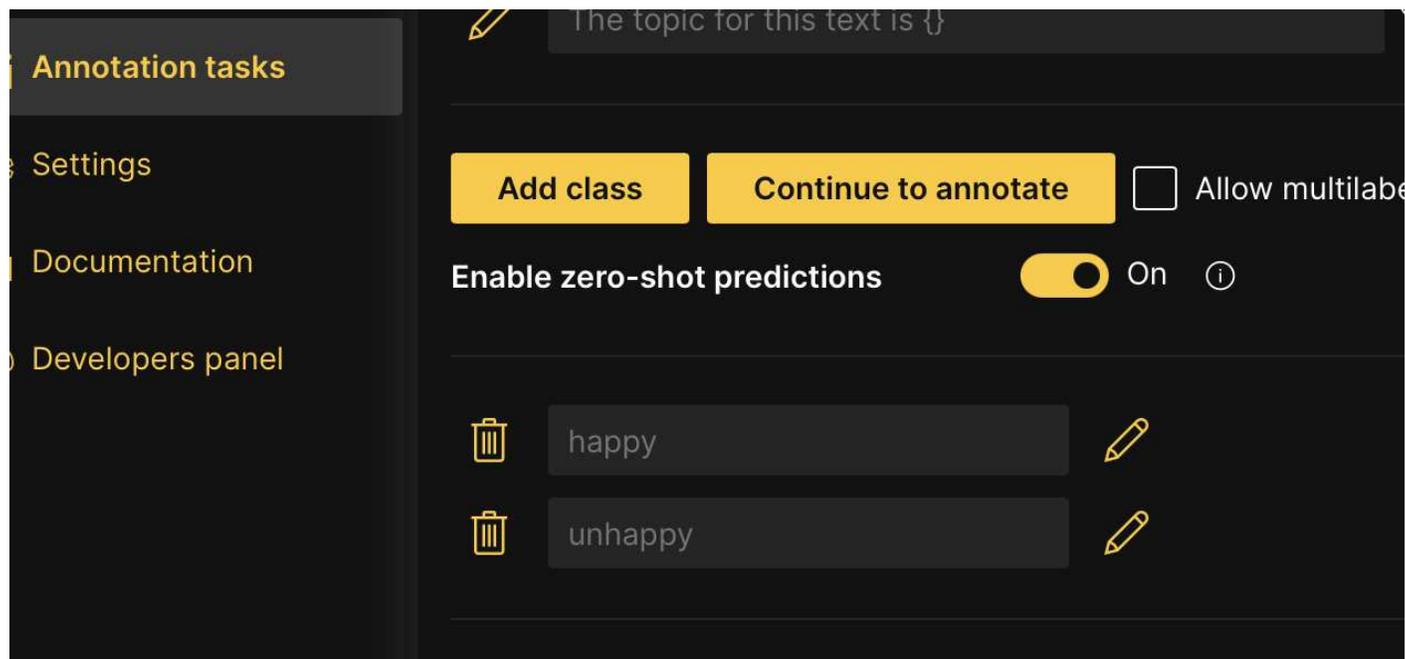
Now that we have seen the dataset let's create an annotation task that enables you to annotate the dataset. For this tutorial, the annotation task refers to a text classification annotation task assigning one or more categorical target labels to an input text.

1. Click **New annotation task**.
2. In the **Task name** box, enter **tutorial-1a**.
3. In the **Task description** box, enter **Annotate a dataset containing reviews from Amazon products**.
4. In the **Select task** list, select **Classification**.
5. In the **Select text column** box, select **comment**.
6. Click **Create task**.

### Step 3: Specify an annotation task rubric

Before annotating the dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. For the dataset, let's utilize the **comment** column to label if a product review follows under a **happy** or **unhappy** customer.

1. In the **New class name** box, enter **happy**.
2. Click **Add**.
3. Click **Add class**.
4. In the **New class name** box, enter **unhappy**.
5. Click **Add**.
6. Click **Continue to annotate**.



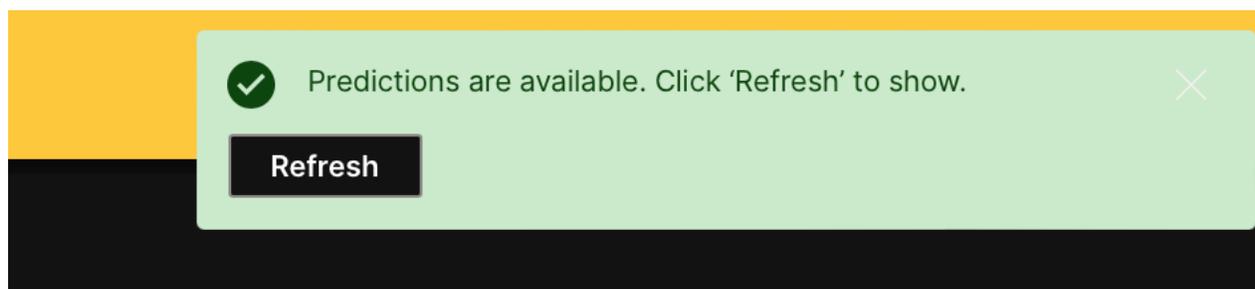
**Note:** H2O Label Genie supports multi-label text classification annotation tasks.

#### Step 4: Annotate dataset

In the **Annotate** tab, you can individually annotate each review (text) in the dataset. Let's annotate the first review.

1. Click **unhappy** (a random label for purposes of this tutorial). **Note:**

- A zero-shot learning model is *on* by default when you annotate a text classification annotation task. The model accelerates the annotation (labeling) process by providing the percentage probability of a text (in this case, a review) belonging to a certain label (one of the labels created in the **Rubric** tab). You can immediately start annotating in the **Annotate** tab or wait until the zero-shot model is ready to provide annotation suggestions. H2O Label Genie notifies you to **Refresh** the instance when zero-shot predictions (suggestions) are available.



- During the annotation process of a text classification dataset, you can download generated zero-shot predictions (probabilities) in the **Export** tab. To learn more, see [Download a dataset's zero-shot predictions](#).
- For example, after refreshing the instance, the model provides probabilities for each label.  
! [] (annotated-review-1-a.png)
- To learn about the utilized model for a text classification annotation task, see [Zero-shot learning models: Text classification](#).

2. Click **Save and next**.

**Note:**

- **\*\*Save and next\*\*** saves the annotated review (sample)
- To skip a review to annotate later: Click **\*\*Skip\*\***.
  - Skipped reviews (samples) reappear after all non-skipped reviews are annotated

3. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already approved annotated samples. You do not need to fully annotate a dataset to download already annotated samples (approved samples). To learn more, see [Download an annotated dataset](#)

## Summary

In this tutorial, we learned how to annotate and specify an annotation task rubric for a text classification task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

## Tutorial 2A: Text regression annotation task

### Overview

This tutorial describes the process of creating a [text regression annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing user reviews (in text format) and ratings (from 0 to 5) of Amazon products.

### Step 1: Explore dataset

We are going to use the preloaded [Amazon reviews demo](#) dataset for this tutorial. The dataset contains **180** samples (text), each containing a review of an Amazon product. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **amazon-reviews-demo**.

### Step 2: Create an annotation task

Now that we have seen the dataset let's create an annotation task that enables you to annotate the dataset. For this tutorial, the annotation task refers to a text regression annotation task assigning one continuous target label to each input text.

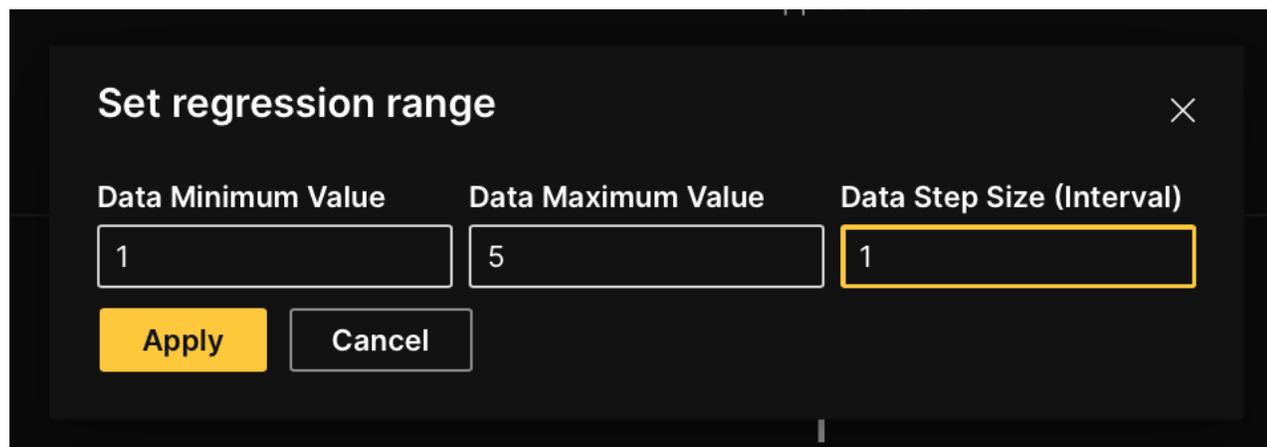
1. Click **New annotation task**.
2. In the **Task name** box, enter **tutorial-2a**.
3. In the **Task description** box, enter **Annotate a dataset containing reviews from Amazon products**.
4. In the **Select task** list, select **Regression**.
5. In the **Select text column** box, select **comment**.
6. Click **Create task**.

### Step 3: Specify an annotation task rubric

Before we can start annotating our dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. For our dataset, we are going to label each review with a value from **1** to **5**, where *1* refers to *1* star, *2* refers to *2* Stars, etc.

**Note:** The dataset has a **start** column rating a review on a start basis with the following structure: **3.0 out of 5 stars**. The new column (a label column) we will create for purposes of this tutorial will only contain a number value from **1-5**.

1. In the **Data minimum value** box, enter **1**.
  - The **Data minimum value** value refers to the minimum value in your continuous values (star ratings from 1 to 5)
2. In the **Data maximum value** box, enter **5**.
  - The **Data maximum value** value refers to the maximum value in your continuous values (star ratings from 1 to 5)
3. In the **Data step size (interval)** box, enter **1**.
  - The **Data step size (interval)** value refers to the value the label range slider interval takes (the slider is used in the next step to label a review)
4. Click **Apply**.



**Set regression range** ✕

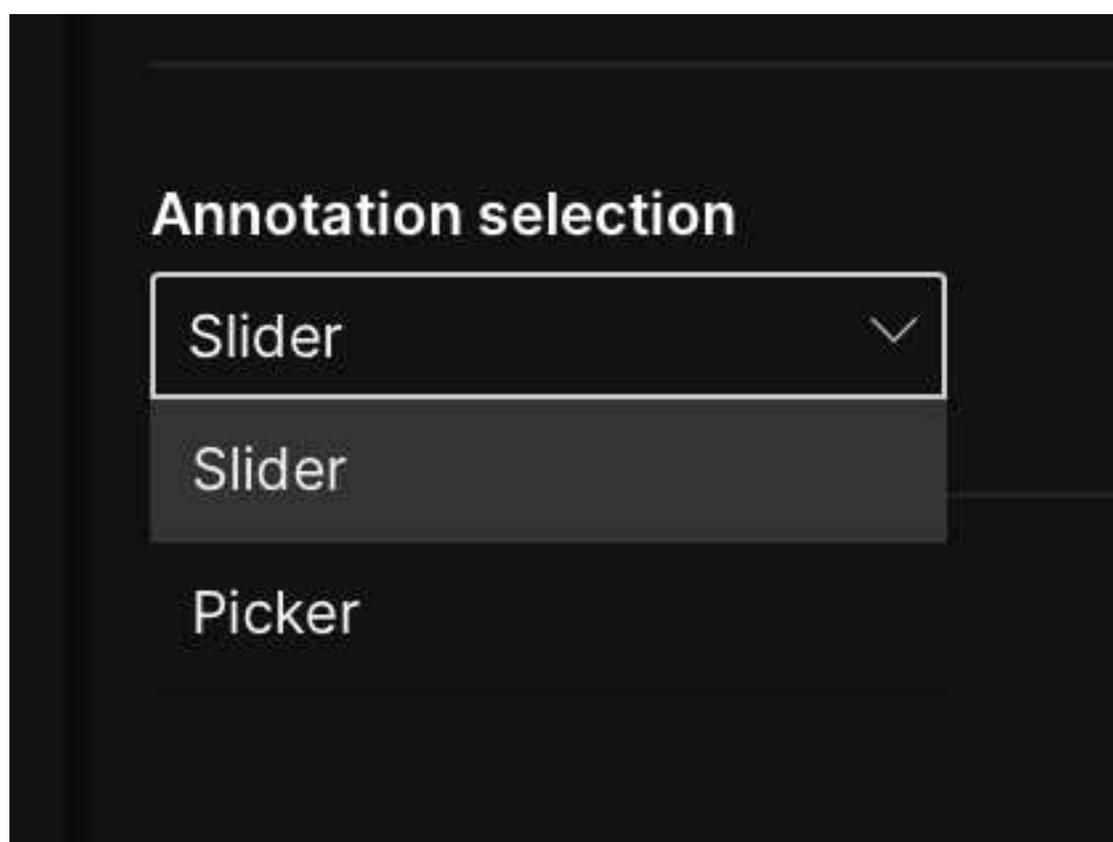
Data Minimum Value      Data Maximum Value      Data Step Size (Interval)

1      5      1

**Apply**      Cancel

Let's utilize the slider, not the picker, to annotate the samples. To enable the slider, consider the following instructions:

1. In the **Annotation selection** list, select **Slider**.



**Annotation selection**

Slider ▾

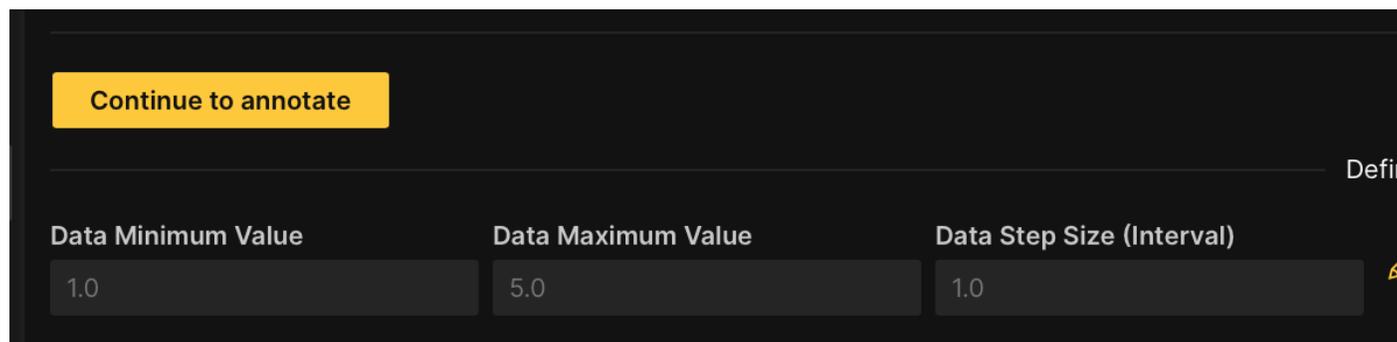
Slider

Picker

#### Step 4: Annotate dataset

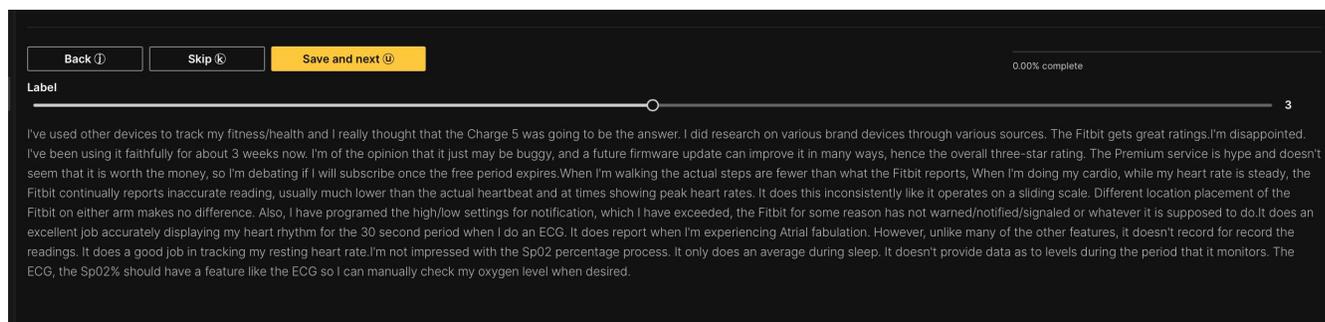
Now that we have specified the annotation task rubric, let's annotate the dataset.

1. Click **Continue to annotate**.



In the **Annotate** tab, you can individually annotate each review (text) in the dataset. Let's annotate the first review.

1. In the **Label** slider, slide to **3** (a random value for purposes of this tutorial).



2. Click **Save and next**. **Note:**
  - **Save and next** saves the annotated review
  - To skip a review to annotate later: Click **Skip**.
    - Skipped reviews (samples) reappear after all non-skipped reviews are annotated
3. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate an imported dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#).

## Summary

In this tutorial, we learned how to annotate and specify an annotation task rubric for a text regression task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

## Tutorial 3A: Text-entity recognition annotation task

### Overview

This tutorial describes the process of creating a [text-entity recognition annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing user reviews (in text format) and ratings (from 0 to 5) of Amazon products.

### Step 1: Explore dataset

We are going to use the preloaded [Amazon reviews demo](#) dataset for this tutorial. The dataset contains **180** samples (text), each containing a review of an Amazon product. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **amazon-reviews-demo**.

### Step 2: Create an annotation task

Now that we understand the dataset, let's create an annotation task that enables you to annotate the dataset. For this tutorial, the annotation task refers to a text-entity recognition annotation task locating and classifying named entities in unstructured text into pre-defined categories.

1. Click **New annotation task**.
2. In the **Task name** box, enter **tutorial-3a**.
3. In the **Task description** box, enter **Annotate a dataset containing reviews from Amazon products**.
4. In the **Select task** list, select **Entity recognition**.
5. In the **Select text column** box, select **comment**.
6. Click **Create task**.

### Step 3: Specify an annotation task rubric

Before we can start annotating our dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. For our dataset, let's define the following two entities: **Product** and **Emotion**.

**Product** refers to the Amazon product reviewed, while **Emotion** refers to one or several expressed (written) feelings during the review.

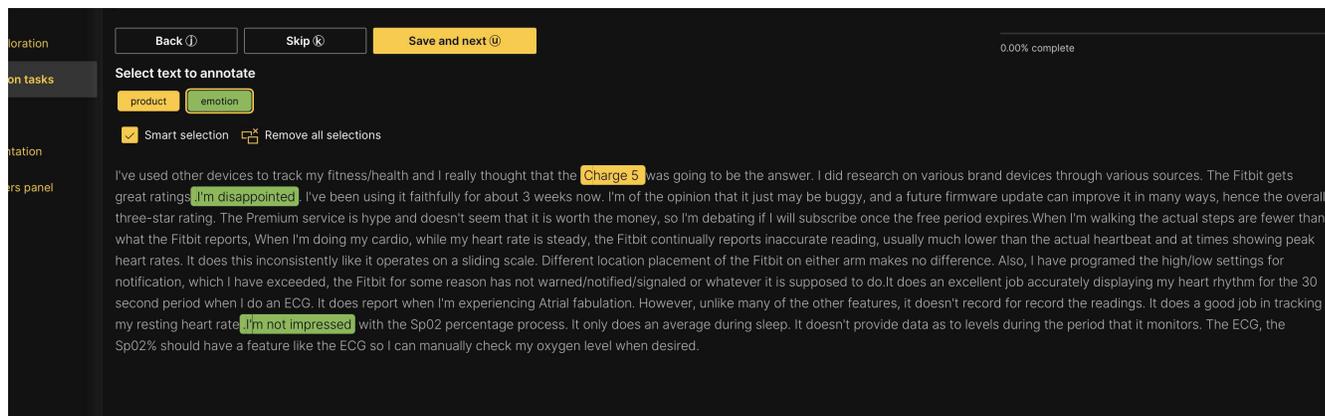
1. In the **New object name** box, enter **product**.
2. Click **Add**.
3. Click **Add entity**.
4. In the **New object name** box, enter **emotion**.
5. Click **Continue to annotate**.



### Step 4: Annotate dataset

In the **Annotate** tab, you can individually annotate each review (text) in the dataset. Let's annotate the first review.

1. Let's start by annotating the **product** entities in the review. Highlight **Charge 5**.
2. Now, let's annotate the review's **emotional (emotion)** entities. Click **Emotion**.
3. Highlight **I'm disappointed**. **Note:** You can attribute a particular entity (**Product** or **Emotion**) to a word by clicking it.
4. Highlight **I'm not impressed**.



5. Click **Save and next**. **Note:**
  - **Save and next** saves the annotated review
  - To skip a review to annotate later: Click **Skip**.
    - Skipped reviews (samples) reappear after all non-skipped reviews are annotated
6. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate a dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#).

## Summary

In this tutorial, we learned how to annotate and specify an annotation task rubric for a text-entity recognition annotation task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

## Tutorial 4A: Text summarization annotation task

### Overview

This tutorial describes the process of creating a [text summarization annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset that contains human-generated abstract summaries from news stories published on the Cable News Network (CNN) and Daily Mail websites.

### Step 1: Explore dataset

We are going to use the preloaded [CNN Daily Mail sample](#) demo dataset for this tutorial. The dataset contains **100** samples (text), each containing a summary of a CNN or Daily Mail article. Let's quickly explore the dataset.

**Note:** The dataset already contains a **summary** column. For purposes of this tutorial, we will ignore that column and create our own column to see how one can create a summarization annotation task.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **cnn-dailymail-sample**.

### Step 2: Create an annotation task

Now that we have seen the dataset let's create an annotation task that enables you to annotate the dataset. For this tutorial, the text summarization annotation task refers to writing a summary for each text input.

1. Click **New annotation task**.
2. In the **Task name** box, enter **tutorial-4a**.
3. In the **Task description** box, enter **Annotate a dataset containing summaries from news stories from CNN and the Daily Mail websites**.
4. In the **Select task** list, select **Summarization**.
5. In the **Select text column** box, select **text**.
6. Click **Create task**.

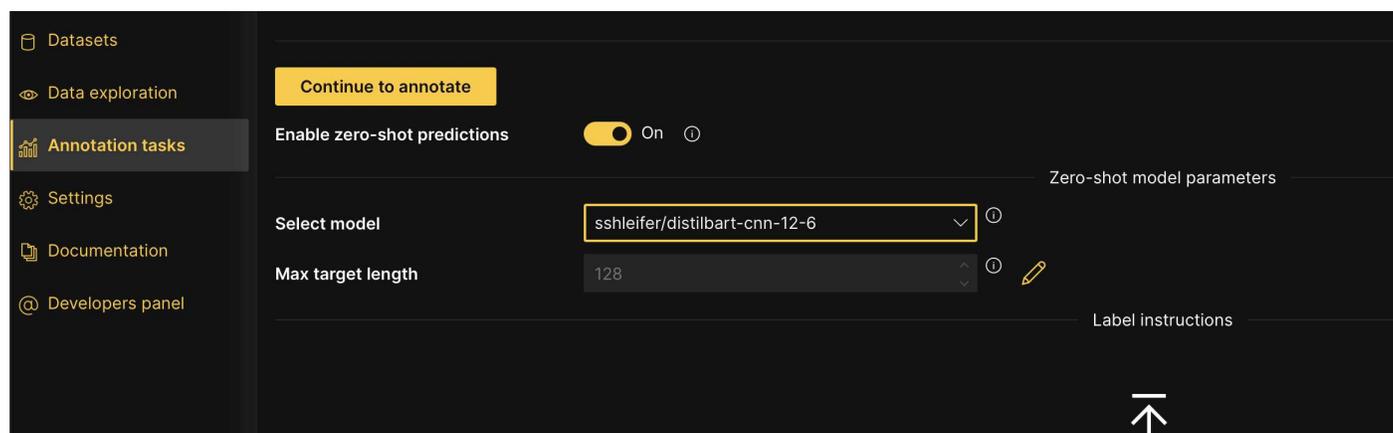
### Step 3: Specify an annotation task rubric

Before we can start annotating our dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. Generally, an annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset.

In the case of a summarization annotation task rubric, you need to specify the following two settings in the **Rubric** tab:

- **Select model**
  - The **Select model** value refers to the zero-shot learning model to utilize in your annotation task. To learn more, see [Annotation tasks + zero-shot learning models: Text summarization](#)
- **Max target length**
  - The **Max target length** value refers to the minimum character length of your summaries

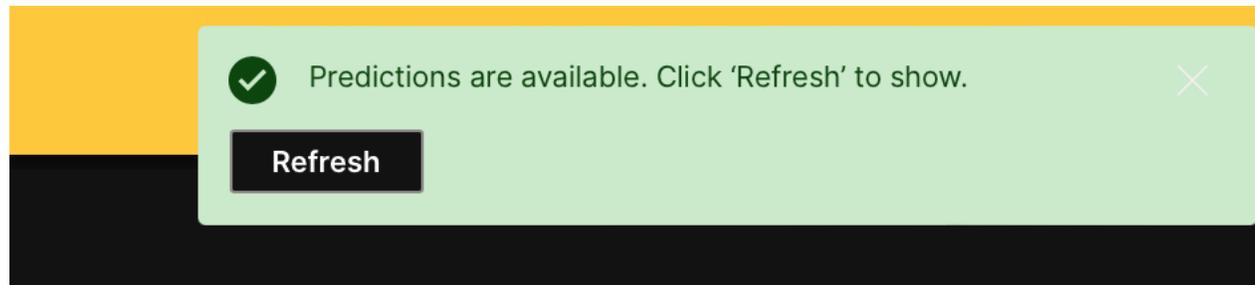
1. For purposes of this tutorial, let's utilize the default model.
2. For purposes of this tutorial, let's utilize the default maximum target length.
3. Click **Continue to annotate**.



## Step 4: Annotate dataset

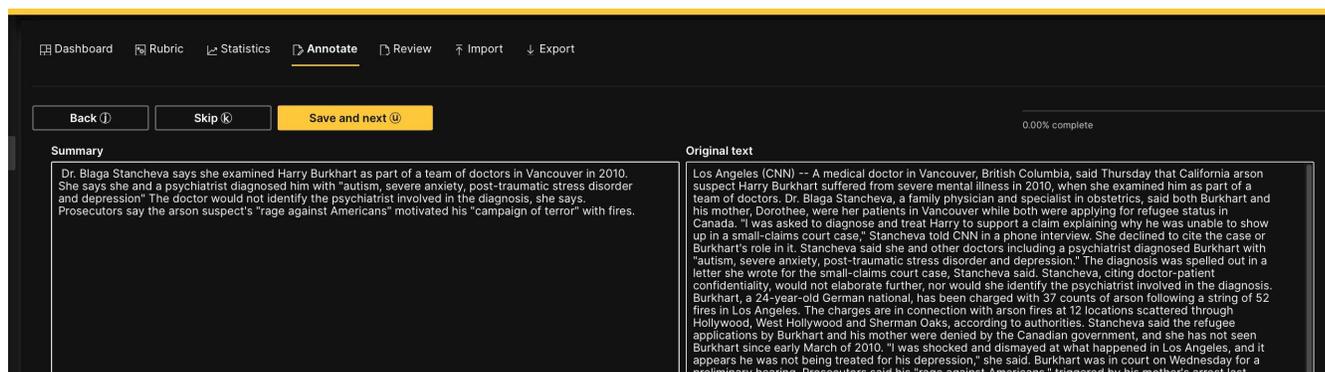
In the **Annotate** tab, you can individually annotate each summary in the dataset. Let's annotate (summarize) the first text.

You can immediately start annotating in the **Annotate** tab or wait until the zero-shot model is ready to provide annotation suggestions. H2O Label Genie notifies you to **Refresh** the instance when zero-shot predictions (suggestions) are available. A zero-shot learning model is *utilized* by default when you annotate a text summarization annotation task. The model accelerates the annotation process by summarizing a given *original text*.



**Note:** To learn about the utilized model for a text summarization annotation task, see [Zero-shot learning models: Text summarization](#).

1. Click **Refresh**.



2. Click **Save and next**. **Note:**
  - **Save and next** saves the annotated text
  - To skip a text to annotate later: Click **Skip**.
    - Skipped text samples reappear after all non-skipped summaries are annotated
3. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate an imported dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#)

## Summary

In this tutorial, we learned how to annotate and specify an annotation task rubric for a text summarization task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

## Tutorial 5A: Text-generative AI annotation task

### Overview

This tutorial describes the process of creating a text-generative AI annotation task, including specifying an annotation task rubric for it. To highlight the process, we will utilize the Amazon reviews dataset that contains user reviews (in text format) and ratings (from 0 to 5) of Amazon products. In particular, we will utilize an H2O.ai zero-shot learning model (large language model (LLM)) to summarize the product reviews.

### Step 1: Explore dataset

We are going to use the preloaded [Amazon reviews demo](#) dataset for this tutorial. The dataset contains **180** samples (text), each containing a review of an Amazon product. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **amazon-reviews-demo**.

### Step 2: Create an annotation task

Now that we understand the dataset, let's create an annotation task that enables you to annotate the dataset. For this tutorial, the annotation task refers to a text-generative AI annotation task capable of generating text in response to prompts (text inputs).

1. Click **New annotation task**.
2. In the **Task name** box, enter **tutorial-5a**.
3. In the **Task description** box, enter **Annotate a dataset containing reviews from Amazon products**.
4. In the **Select task** list, select **Generative AI**.
5. In the **Select text column** box, select **comment**.
6. Click **Create task**.

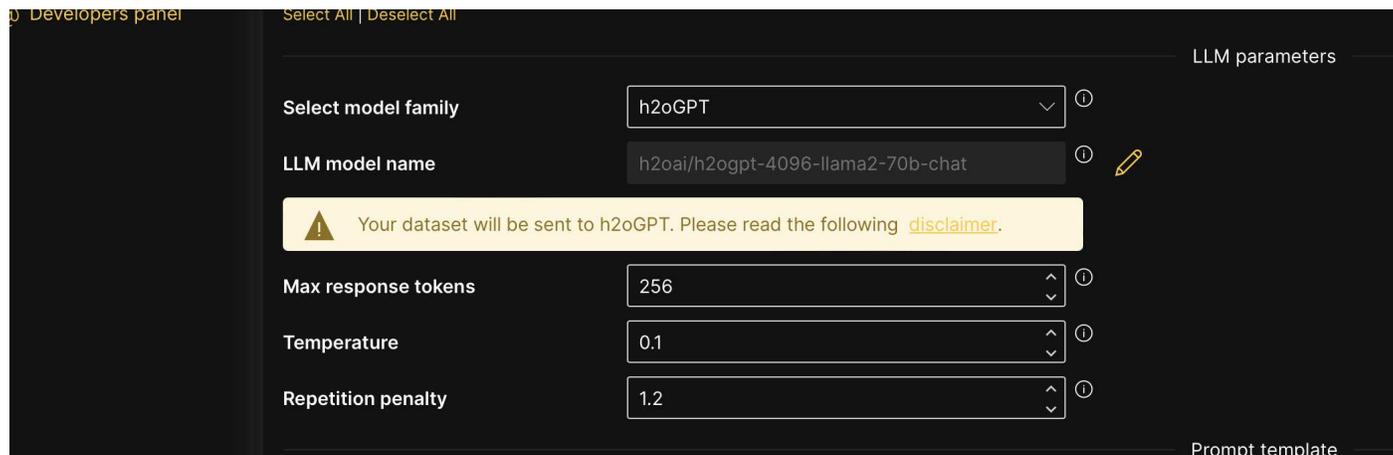
### Step 3: Specify an annotation task rubric

Before annotating our dataset, we need to specify an annotation task rubric. Generally, an annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. In the case of a text-generative AI annotation task, the rubric specifies the following:

- **(A)** A zero-shot learning model (large language model (LLM)) and its parameters
- **(B)** A prompt template (that is the input for the LLM)

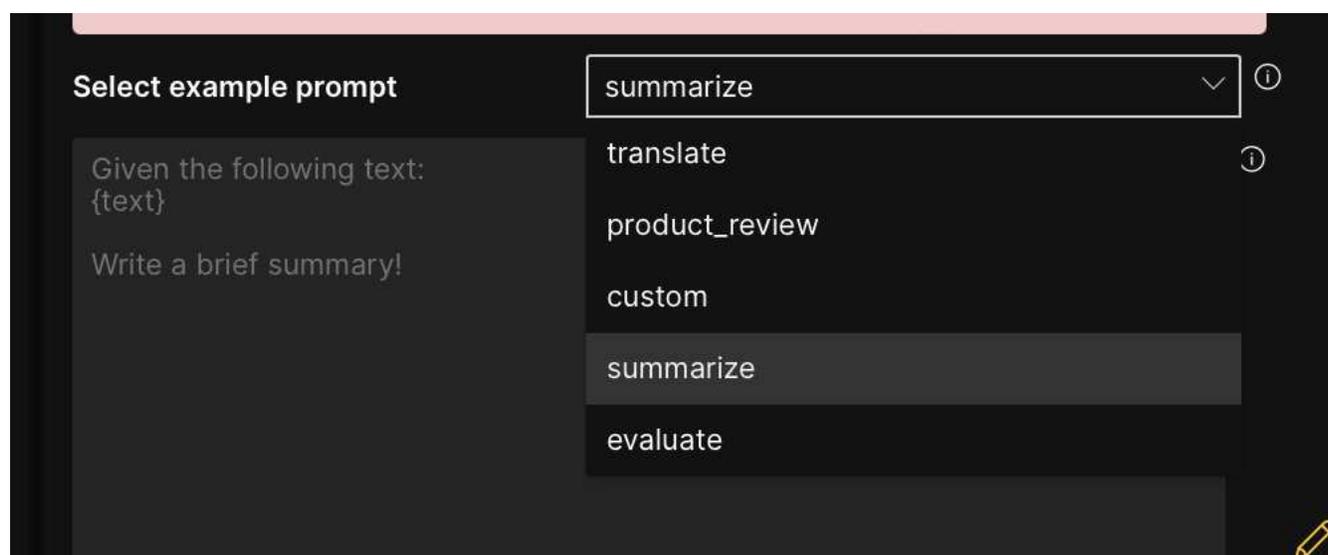
**A:** Let's begin by defining a zero-shot learning model (large language model (LLM)) and its parameters.

1. In the **Select API endpoint type** list, select **h2oGPT**.
  - The **Select API endpoint type** setting defines the zero-shot learning model family to utilize in the annotation task
2. After selecting a model family, H2O Label Genie, by default, selects an h2oGPT zero-shot model for the task. For purposes of this tutorial, let's utilize the default selected model (in the **LLM model name** setting). To learn more, see [Zero-shot learning models: Text-generative AI](#).
3. For purposes of this tutorial, let's utilize the default values for the other LLM parameters. To learn more about each parameter, see [Large language model \(LLM\) parameters](#).



**B:** Now, let's define the prompt the zero-shot learning model will receive as inputs.

1. In the **Select example prompt** list, select **summarize**.



**Note:** For purposes of this tutorial, we selected the **summarize** prompt (to submit as inputs to the LLM), but note that there are several other options, including the option to create your own custom prompt. To learn more, see [Select example prompt](#).

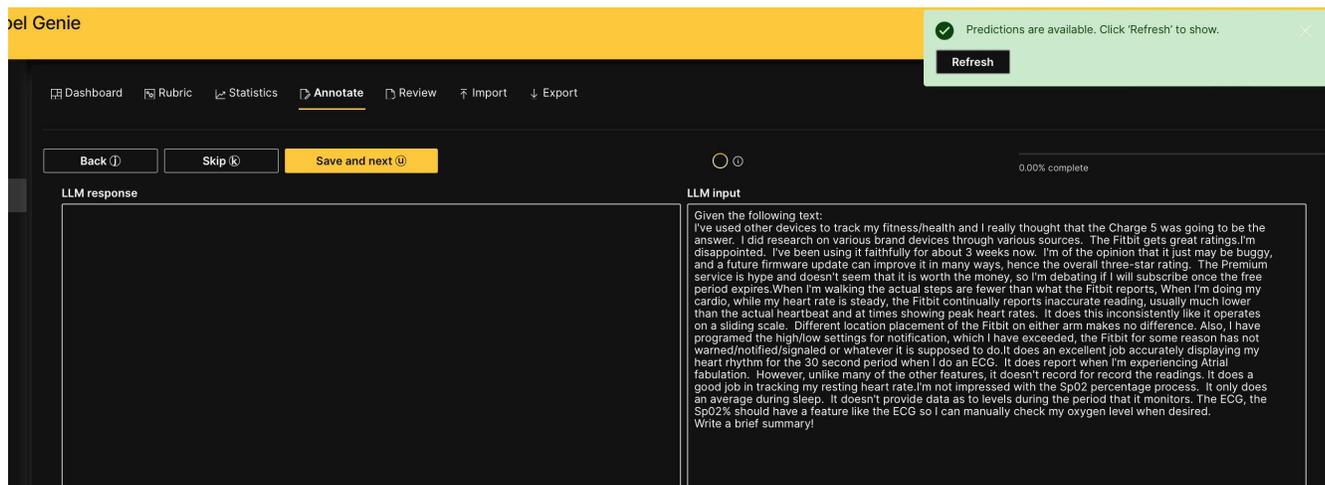
2. Click edit **Edit**.
3. In the **Prompt for large language model** write box, modified the text as follows: `Given the following text: {comment} Write a brief summary! Note:`
  - The `{comment}` placeholder refers to one of the reviews in the dataset. To summarize a review, H2O Label Genie places a review in the placeholder, which the LLM utilizes to summarize
  - The `{comment}` placeholder refers to the column name containing the Amazon product reviews
4. Click **Save**.
5. Click **Continue to annotate**.

## Step 4: Annotate dataset

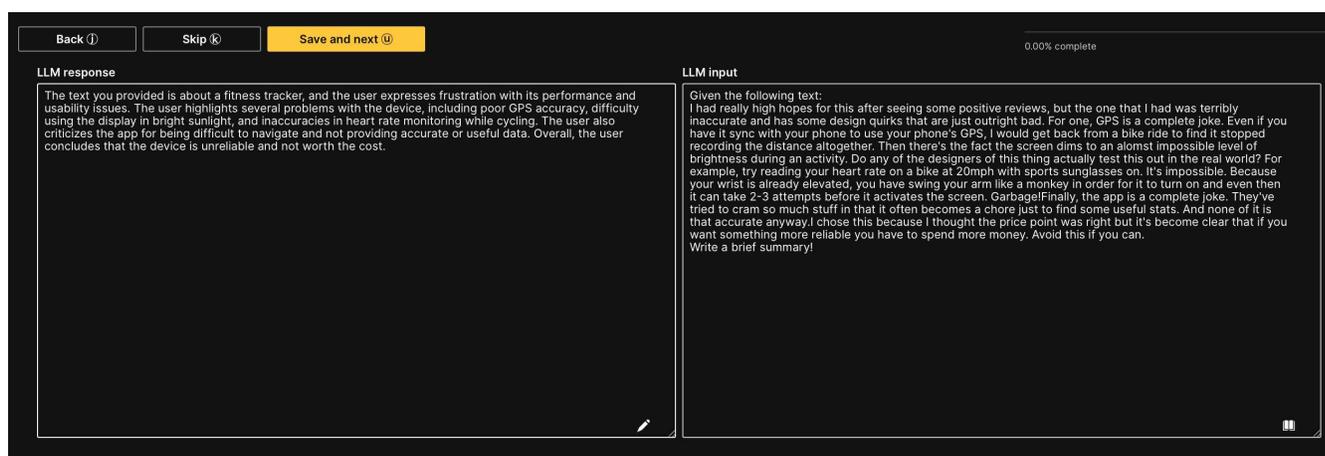
In the **Annotate** tab, you can individually annotate each review (text) in the dataset. After proceeding with the annotation task:

- H2O Label Genie sends each review (prompt) to the zero-shot learning model to summarize, right after
- H2O Label Genie notifies you to **Refresh** the instance when the responses (summaries) become available

1. Click **Refresh**.



2. (Optional) If you want to edit a review's response, in the **LLM response** write box, click edit.



3. Click **Save and next**. **Note:**

- **Save and next** saves the annotated (summarized) text (approves it)
- To skip a text to annotate later: Click **Skip**
  - Skipped summaries reappear after all non-skipped summaries are annotated

4. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate an imported dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#).

## Summary

In this tutorial, we learned how to annotate and specify an annotation task rubric for a text-generative AI annotation task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

## Tutorial 1B: Image classification annotation task

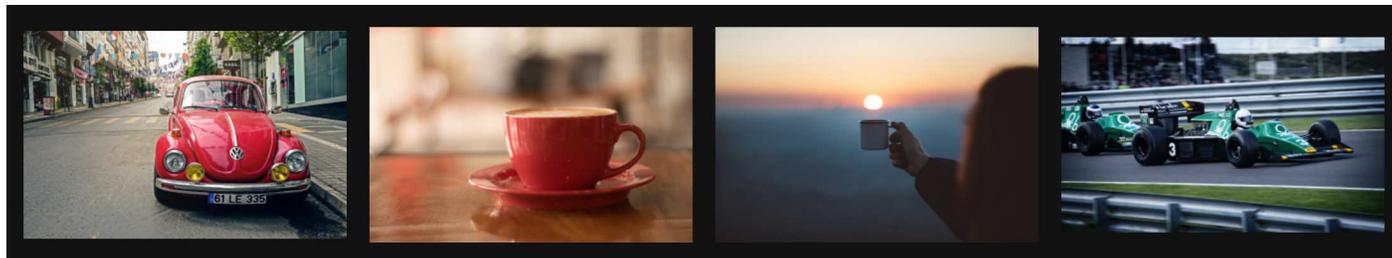
### Overview

This tutorial describes the process of creating an [image classification annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we are going to annotate a dataset that contains images of cars and coffee.

### Step 1: Explore dataset

We are going to use the preloaded [Car or coffee demo](#) dataset for this tutorial. The dataset contains **40** images, each depicting a car or coffee. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **car-or-coffee-demo**.



### Step 2: Create an annotation task

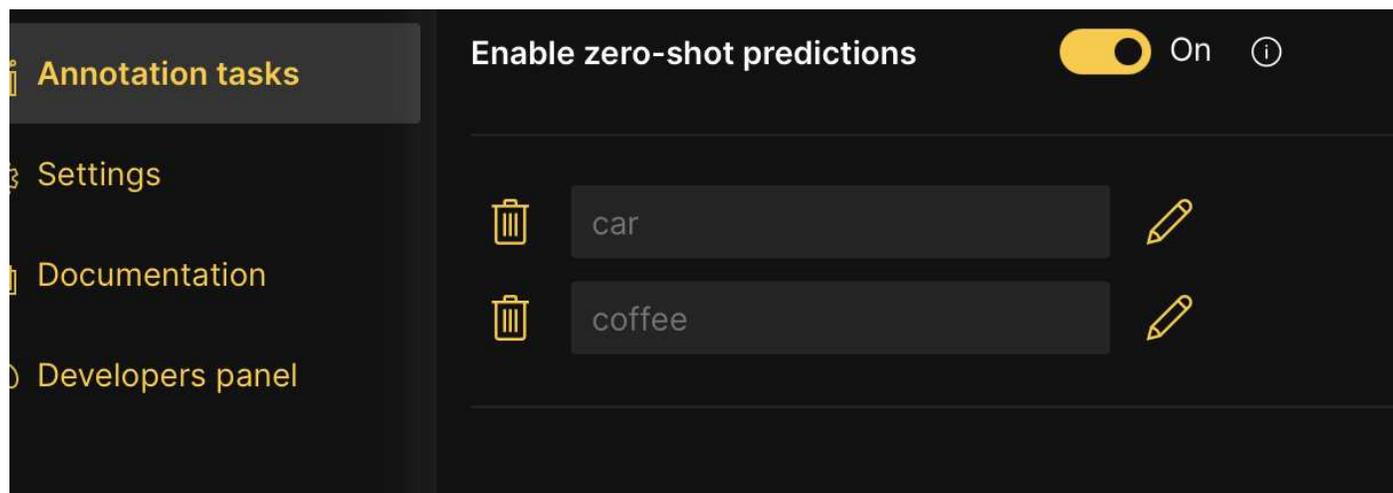
Now that we have seen the dataset let's create an annotation task that enables you to annotate the dataset. For this tutorial, an image classification annotation task refers to assigning one or more categorical target labels to an input image. Let's create an annotation task.

1. Click **New annotation task**.
2. In the **Task name** box, enter **tutorial-1b**.
3. In the **Task description** box, enter **Annotate a dataset containing images of cars and coffee**.
4. In the **Select task** list, select **Classification**.
5. Click **Create task**.

### Step 3: Specify an annotation task rubric

Before we can start annotating our dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. For our dataset, there are two categorical target labels we want to specify, **car** and **coffee**. Let's define the annotation task rubric.

1. In the **New class name** box, enter **car**.
2. Click **Add**.
3. Click **Add class**.
4. In the **New class name** box, enter **coffee**.
5. Click **Add**.
6. Click **Continue to annotate**.



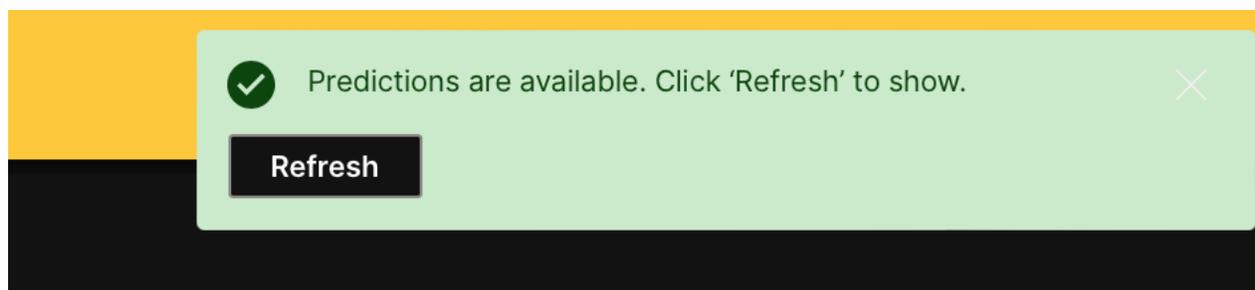
**Note:** H2O Label Genie supports multi-label image classification annotation tasks.

### Step 4: Annotate dataset

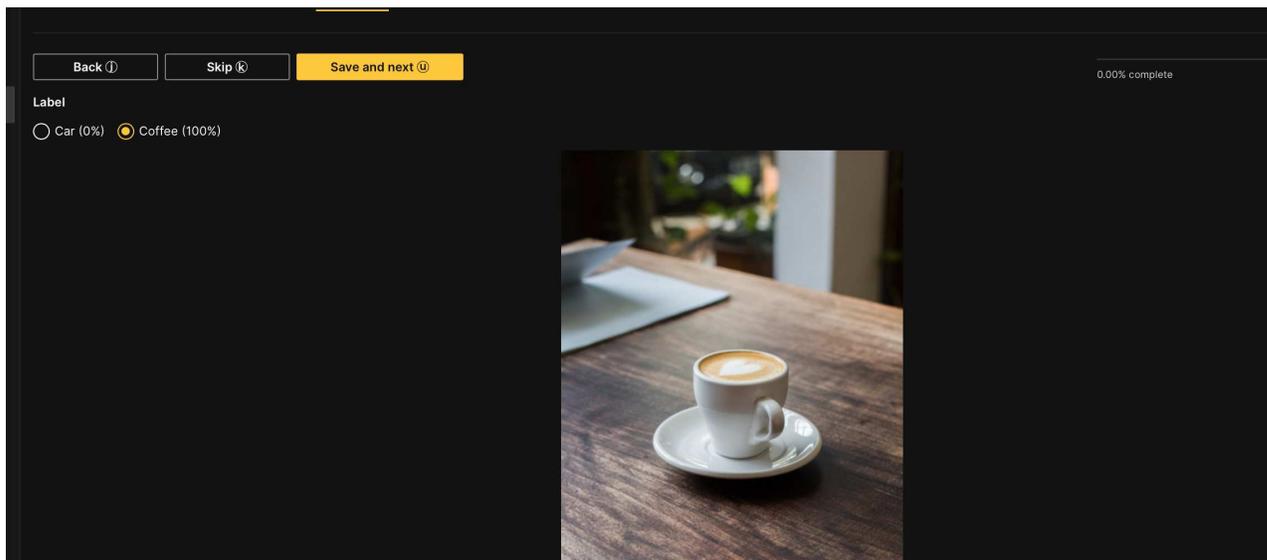
In the **Annotate** tab, you can individually annotate each image in the dataset. Let's annotate the first image.

1. Select **coffee** or **car**. **Note:**

- A zero-shot learning model is *on* by default when you annotate an image classification annotation task. The model accelerates the annotation (labeling) process by providing the percentage probability of an image (in this case, a car or coffee image) belonging to a certain label (one of the labels created in the **Rubric** tab). You can immediately start annotating in the **Annotate** tab or wait until the zero-shot model is ready to provide annotation suggestions. H2O Label Genie notifies you to **Refresh** the instance when zero-shot predictions (suggestions) are available.



- During the annotation process of an image classification dataset, you can download generated zero-shot predictions (probabilities) in the **Export** tab. To learn more, see [Download a dataset's zero-shot predictions](#).
- For example, after refreshing the instance in this tutorial, the model provided probabilities for each label.



- To learn about the utilized model for an image classification annotation task, see [Zero-shot learning models: Image classification](#).
2. Click **Save and next**. **Note:**
    - **Save and next** saves the annotated image
    - To skip an image to annotate later: Click **Skip**.
      - Skipped images (samples) reappear after all non-skipped images are annotated
  3. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate an imported dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#).

## Summary

In this tutorial, we learned the process of annotating and specifying an annotation task rubric for an image classification annotation task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

## Tutorial 2B: Image regression annotation task

### Overview

This tutorial describes the process of creating an [image regression annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we are going to annotate a dataset that contains images of healthy and diseased apple leaves for plant pathology recognition.

### Step 1: Explore dataset

We are going to use the preloaded [Plant pathology demo](#) dataset for this tutorial. The dataset contains **32** images depicting healthy and diseased apple leaves. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **plant-pathology-demo**.



### Step 2: Create an annotation task

Now that we understand the dataset let's create an annotation task that enables you to annotate the dataset. For example, an image regression annotation task refers to assigning one continuous target label to each input image.

1. Click **New annotation task**.
2. In the **Task name** box, enter `tutorial-2b`.
3. In the **Task description** box, enter `Annotate a dataset containing images of healthy and diseased apple leaves`.
4. In the **Select task** list, select **Regression**.
5. Click **Create task**.

### Step 3: Specify an annotation task rubric

Before we can start annotating our dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating a dataset. For our dataset, we are going to label each image on a scale between **0** to **9** to rank the severity of the plant disease. A **0** refers to the plant having the least disease severity, and **9** refers to the plant having the most disease severity.

1. In the **Data minimum value** box, enter `0`.
  - The **Data minimum value** value refers to the minimum value in your continuous values (digits ranging between 0 to 9)
2. In the **Data maximum value** box, enter `9`.
  - The **Data maximum value** value refers to the maximum value in your continuous values (digits ranging between 0 to 9)
3. In the **Data step size (interval)** box, enter `1`.
  - The **Data step size (interval)** value refers to the value the label range slider interval takes (the slider is utilized in [step 4](#))
4. Click **Apply**

**Set regression range** ✕

Data Minimum Value: 0      Data Maximum Value: 9      Data Step Size (Interval): 1

**Apply**      **Cancel**

Let's utilize the slider, not the picker, to annotate the samples. To enable the slider, consider the following instructions:

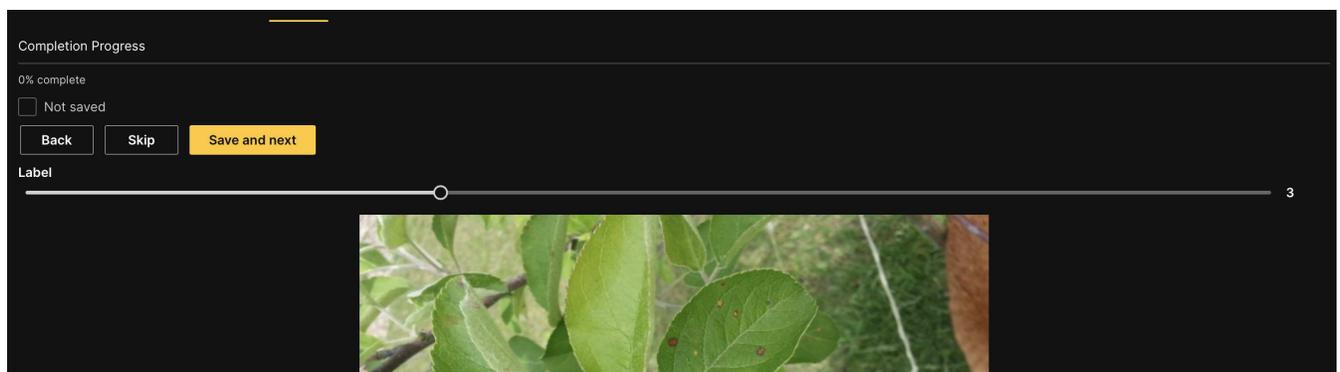
1. In the **Annotation selection** list, select **Slider**.



#### Step 4: Annotate dataset

In the **Annotate** tab, you can individually annotate each image in the dataset. Let's annotate the first image.

1. In the **Label** slider, slide to **3** (a random value for purposes of this tutorial).



## 2. Click **Save and next**. **Note:**

- **Save and next** saves the annotated image
- To skip an image to annotate later: Click **Skip**.
  - Skipped images reappear after all non-skipped images are annotated

## 3. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate an imported dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#).

## Summary

In this tutorial, we learned the process of annotating and specifying an annotation task rubric for an image regression annotation task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

## Tutorial 3B: Object detection annotation task

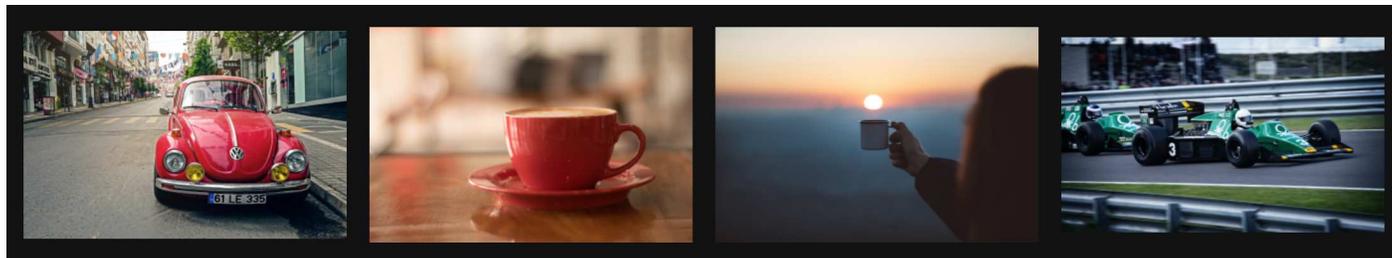
### Overview

This tutorial describes the process of creating an [object detection annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we are going to annotate a dataset that contains images of cars and coffee.

### Step 1: Explore dataset

We are going to use the preloaded [Car or coffee demo](#) dataset for this tutorial. The dataset contains **40** images, each depicting a **car** or **coffee**. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **car-or-coffee-demo**.



### Step 2: Create an annotation task

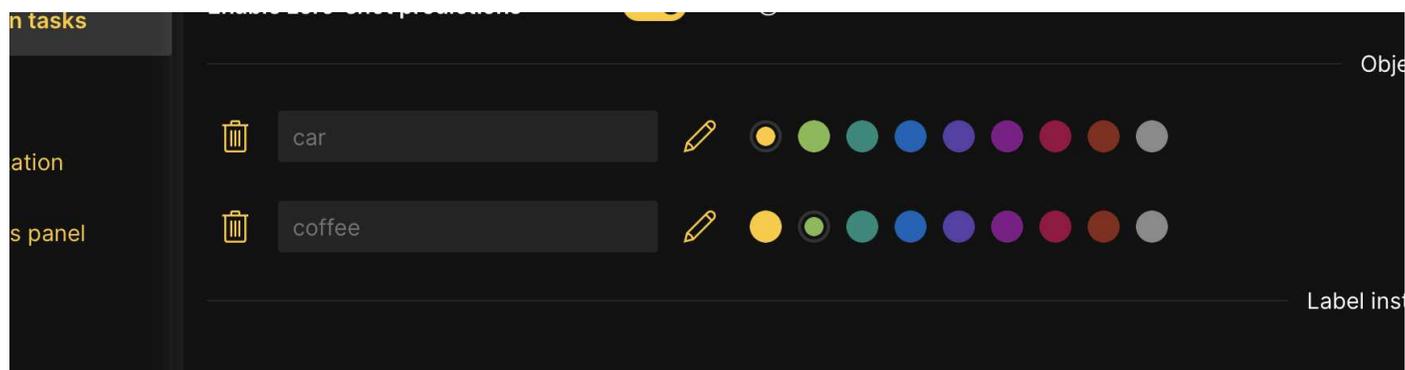
Now that we have seen the dataset let's create an annotation task that enables you to annotate the dataset. For this tutorial, an object detection annotation task refers to specifying one or more object classes (labels) to each input image. Let's create an annotation task.

1. Click **New annotation task**.
2. In the **Task name** box, enter **tutorial-3b**.
3. In the **Task description** box, enter **Annotate a dataset containing images of cars and coffee**.
4. In the **Select task** list, select **Object detection**.
5. Click **Create task**.

### Step 3: Specify an annotation task rubric

Before we can start annotating our dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. For our dataset, there are two labels (object classes) we want to specify, **car** and **coffee**. Let's define the annotation task rubric.

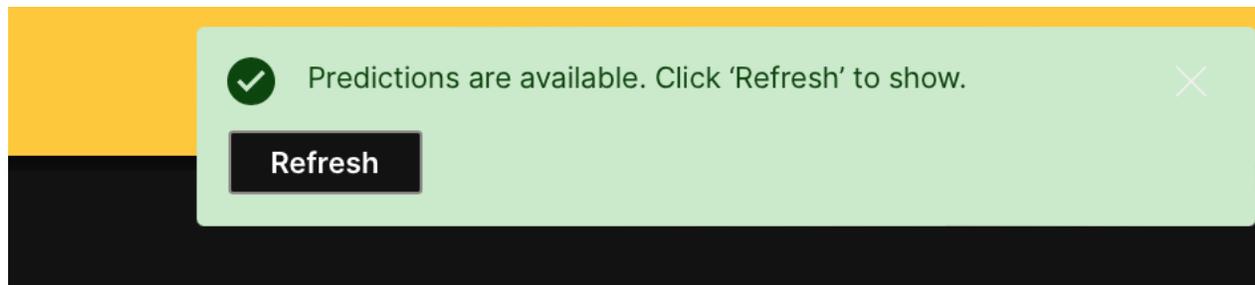
1. In the **Add new object class** box, enter **car**.
2. Click **Add object class**.
3. In the **Add new object class** box, enter **coffee**.
4. Click **Add**.
5. Click **Continue to annotate**.



## Step 4: Annotate dataset

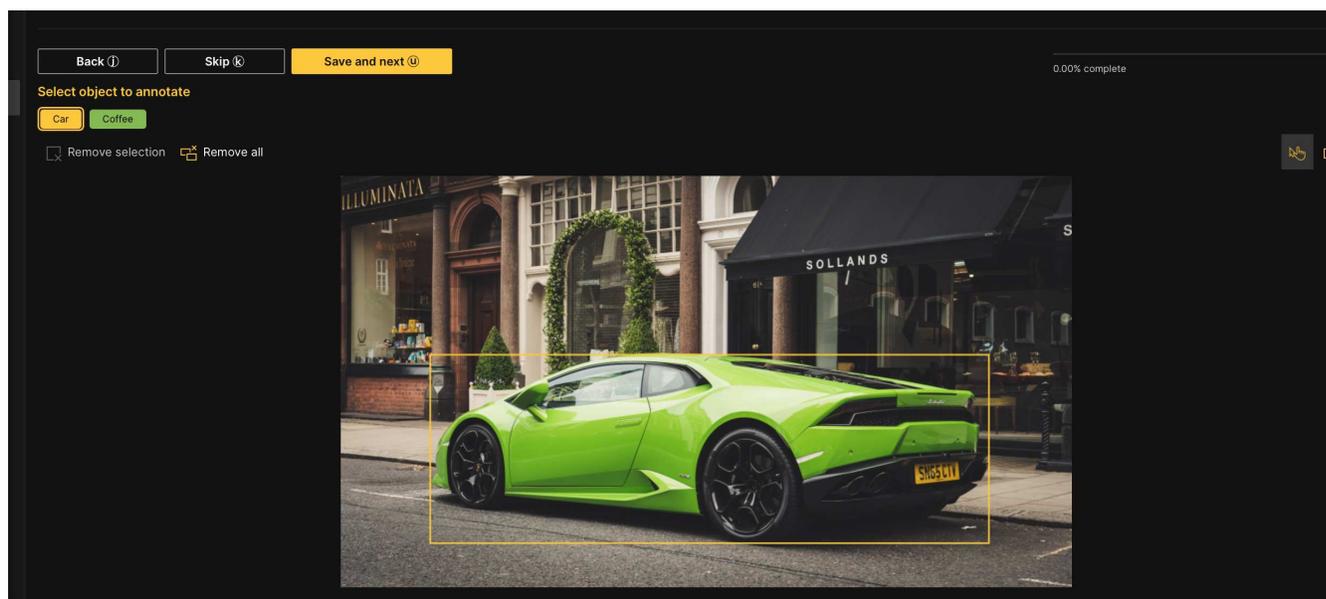
Now that we have specified the annotation task rubric, let's annotate the dataset (the first image).

1. A zero-shot learning model is *on* by default when you annotate an object detection annotation task. The model accelerates the annotation (labeling) process by capturing the object observed in the image with a bounding box. You can immediately start annotating in the **Annotate** tab or wait until the zero-shot model is ready to provide annotation suggestions. H2O Label Genie notifies you to **Refresh** the instance when zero-shot predictions (suggestions) are available.



**Note:** To learn about the utilized model for an object detection annotation task, see [Zero-shot learning models: Object detection](#).

2. Click **Refresh**.



**Note:** If the zero-shot learning model did not accurately capture the object on the image, inside the bounding box, double-click inside the bounding box. Double-clicking inside the bounding box enables you to move and resize the bounding box to accurately cover the area the car (object) takes on the image.

3. Click **Save and next**. **Note:**
  - **Save and next** saves the annotated image
  - To skip an image to annotate later: Click **Skip**.
    - Skipped images (samples) reappear after all non-skipped images are annotated
4. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate an imported dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#).

## Summary

In this tutorial, we learned the process of annotating and specifying an annotation task rubric for an object detection annotation task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

## Tutorial 4B: Image instance segmentation annotation task

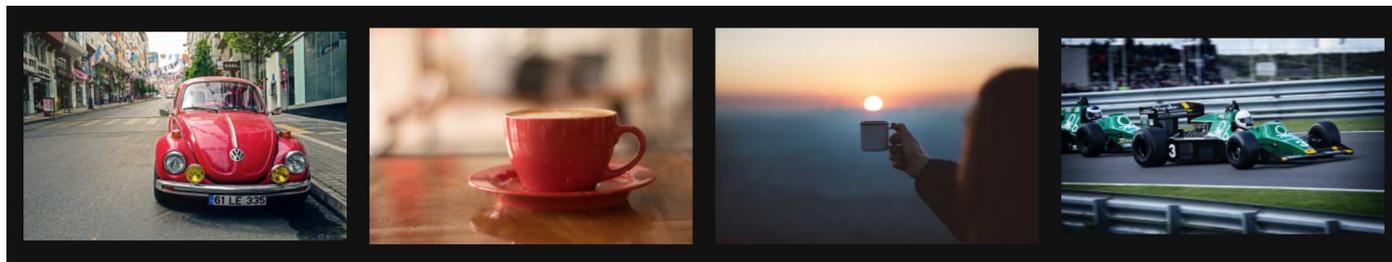
### Overview

This tutorial describes the process of creating an [image instance segmentation annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we are going to annotate a dataset that contains images of cars and coffee.

### Step 1: Explore dataset

We are going to use the preloaded [Car or coffee demo](#) dataset for this tutorial. The dataset contains **40** images, each depicting a **car** or **coffee**. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **car-or-coffee-demo**.



### Step 2: Create an annotation task

Now that we have seen the dataset, let's create an annotation task that enables you to annotate the dataset. For this tutorial, an image instance segmentation annotation task refers to:

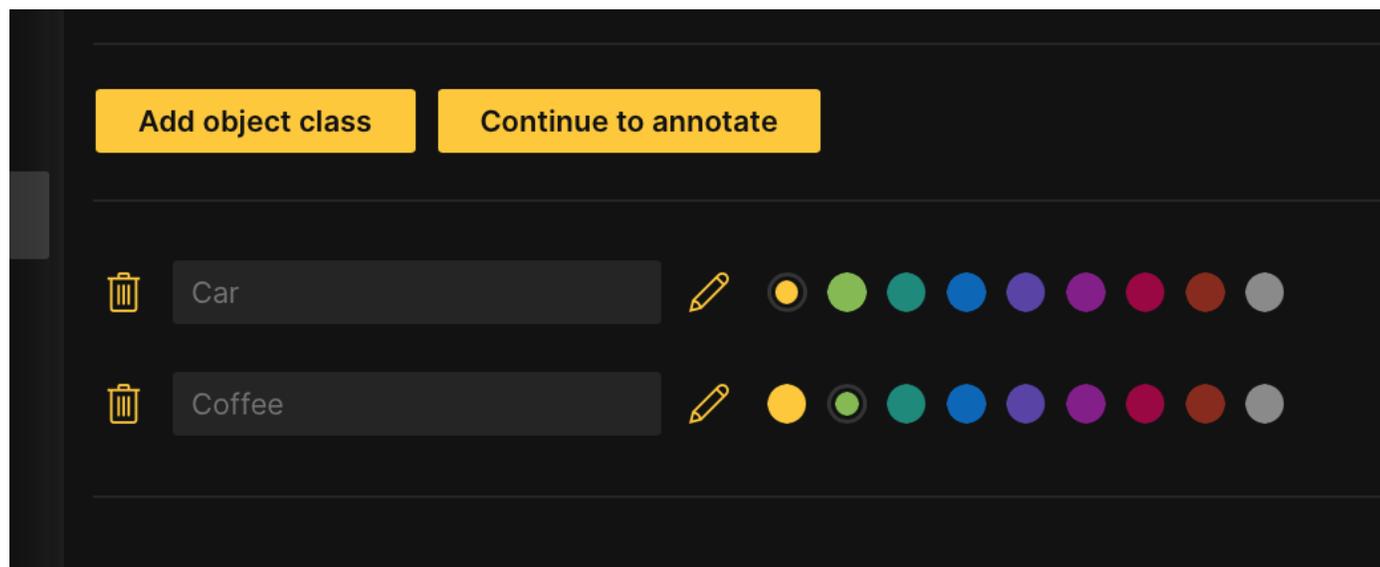
- Specify one or more object classes (labels) for each input image
- Draw a polygon mask for each instance of the object classes on all images

1. Click **New annotation task**.
2. In the **Task name** box, enter **tutorial-4b**.
3. In the **Task description** box, enter **Annotate a dataset containing images of cars and coffee**.
4. In the **Select task list**, select **Instance segmentation**.
5. Click **Create task**.

### Step 3: Specify an annotation task rubric

Before we can start annotating our dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. For our dataset, there are two labels (object classes) we want to specify, **car** and **coffee**. Let's define the annotation task rubric.

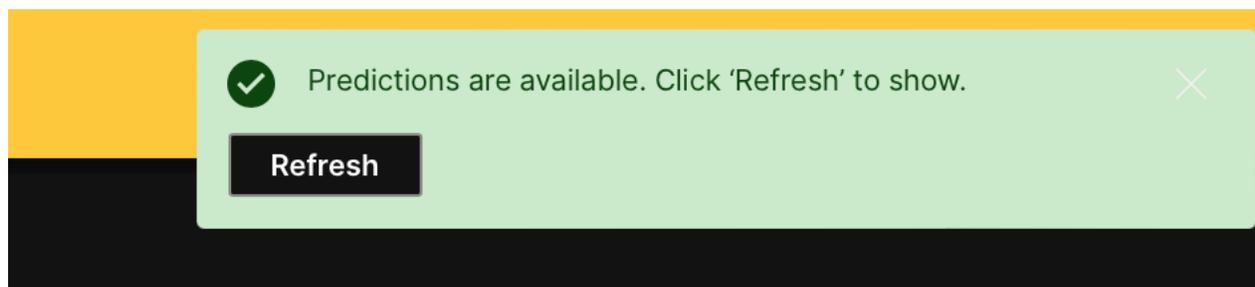
1. In the **Add a new object class** box, enter **car**.
2. Click **Add**.
3. Click **Add object class**.
4. In the **Add a new object class** box, enter **coffee**.
5. Click **Add**.
6. Click **Continue to annotate**.



## Step 4: Annotate dataset

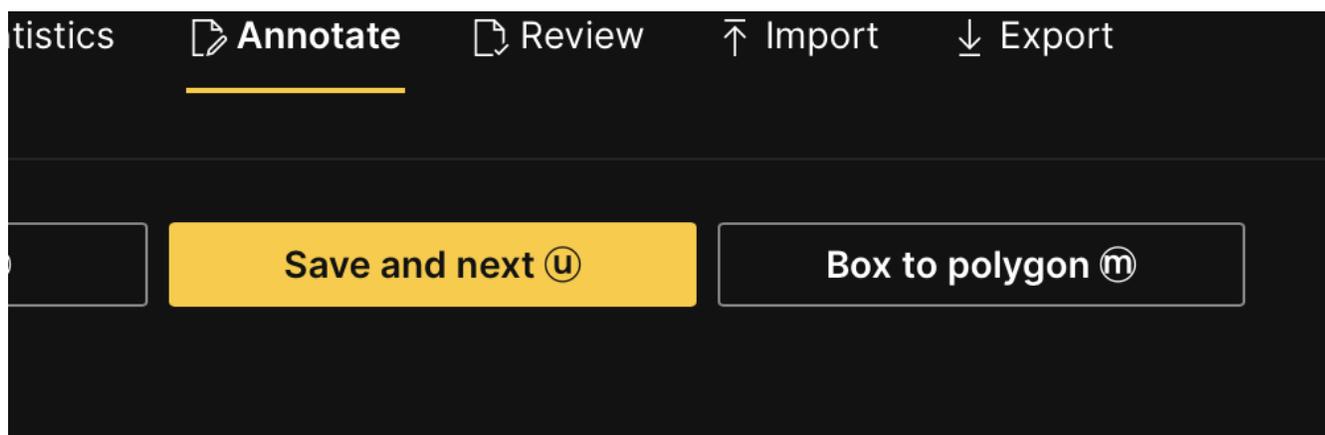
Now that we have specified the annotation task rubric, let's annotate the dataset.

1. You can immediately start annotating in the **annotate** tab or wait until the zero-shot model is ready to provide annotation suggestions. H2O Label Genie notifies you to **Refresh** the instance when zero-shot predictions (suggestions) are available. A zero-shot learning model is *on* by default for an image instance segmentation annotation task. The model accelerates the annotation (labeling) process by capturing with a polygon where in an image one of the specified object classes is located.



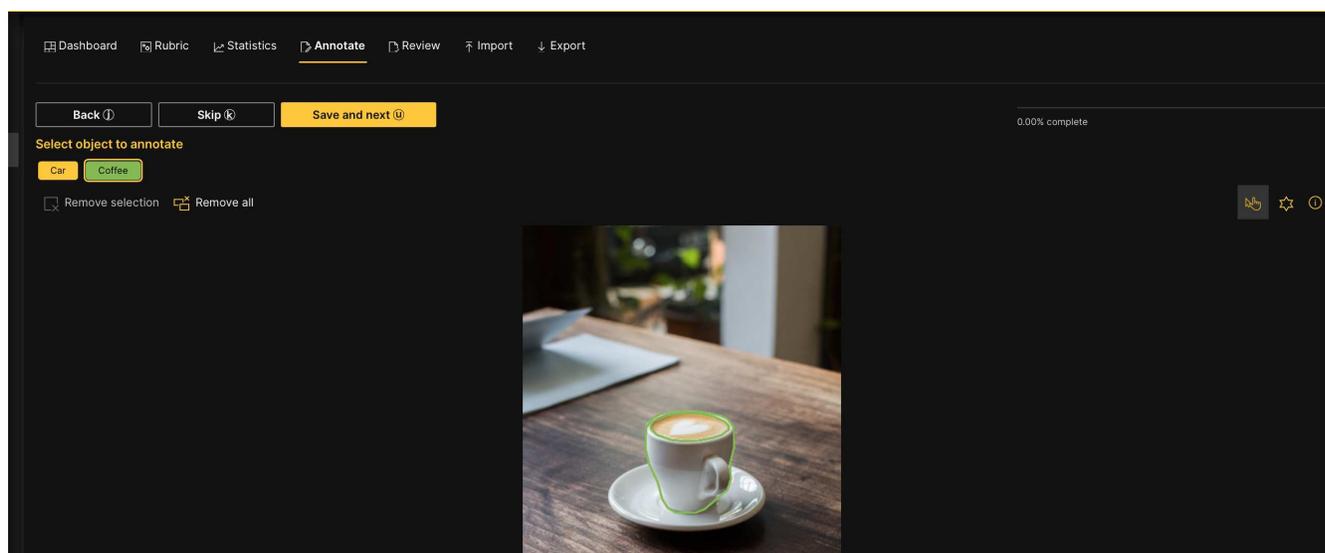
**Note:**

- **Box to polygon:** By default, for an image instance segmentation annotation task, H2O Label Genie offers a **Box to polygon** feature that enables you to improve the efficiency of manual labeling. A **Box to polygon** feature enables you to encapsulate an object class with a rectangular that is immediately replaced with a more complex polygon that captures all the borders of the object class inside the rectangular. To learn more, see [Segment objects](#)



- To learn about the utilized model for an image instance segmentation annotation task, see [Zero-shot learning models: Image instance segmentation](#).

## 2. Click **Refresh**.



**Note: Adjusting misplaced polygons:** If the zero-shot learning model did not accurately capture an object on the image with a polygon shape, left-click inside the polygon. Left-clicking inside the polygon shape enables you to move and reshape the polygon to accurately cover the area the object takes in the image.

## 3. Click **Save and next**. **Note:**

- **Save and next** saves the annotated image
- To skip an image to annotate later: Click **Skip**.
  - Skipped images (samples) reappear after all non-skipped images are annotated

## 4. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate an imported dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#)

## Summary

In this tutorial, we learned the process of annotating and specifying an annotation task rubric for an image instance segmentation annotation task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

# Tutorial 1C: Audio classification annotation task

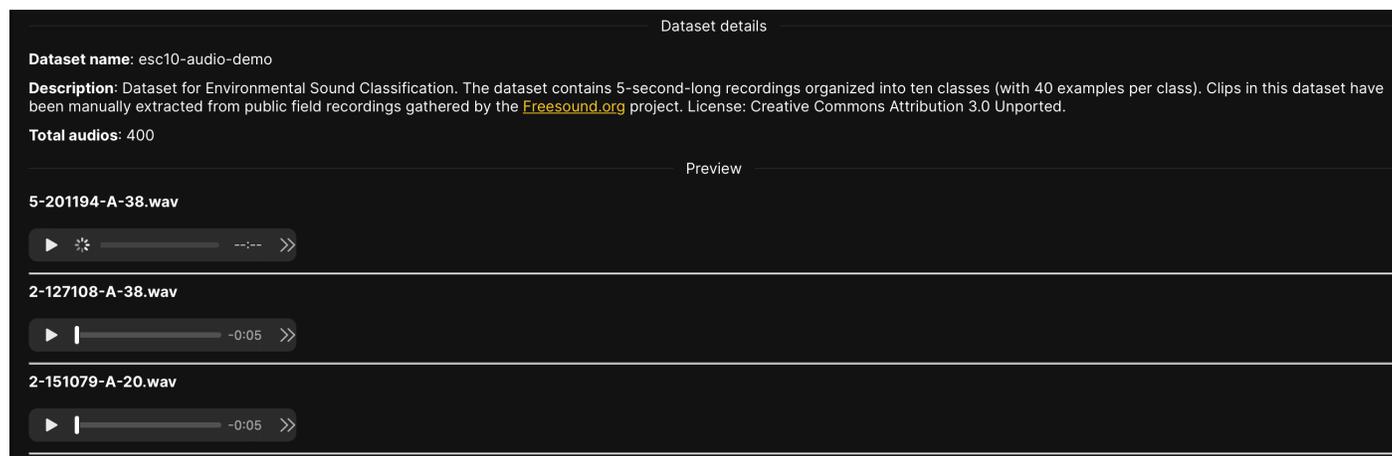
## Overview

This tutorial describes the process of creating an [audio classification annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing 5-second-long recordings of environmental sounds organized into ten classes (with 40 examples per class).

## Step 1: Explore dataset

This tutorial uses the preloaded [ESC10 audio demo](#) dataset, which contains **400** audios, each depicting the sound of a *chainsaw*, *dog*, *helicopter*, *rain*, *rooster*, etc. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **esc10-audio-demo**.



## Step 2: Create an annotation task

Now that we have seen the dataset let's create an annotation task that enables you to annotate the dataset. For this tutorial, an audio classification annotation task refers to assigning a categorical target label to an audio clip.

1. Click **New annotation task**.
2. In the **Task name** box, enter `tutorial-1c`.
3. In the **Task description** box, enter `Annotate a dataset containing samples of environmental sounds.`
4. In the **Select task** list, select **Classification**.
5. Click **Create task**.

## Step 3: Specify an annotation task rubric

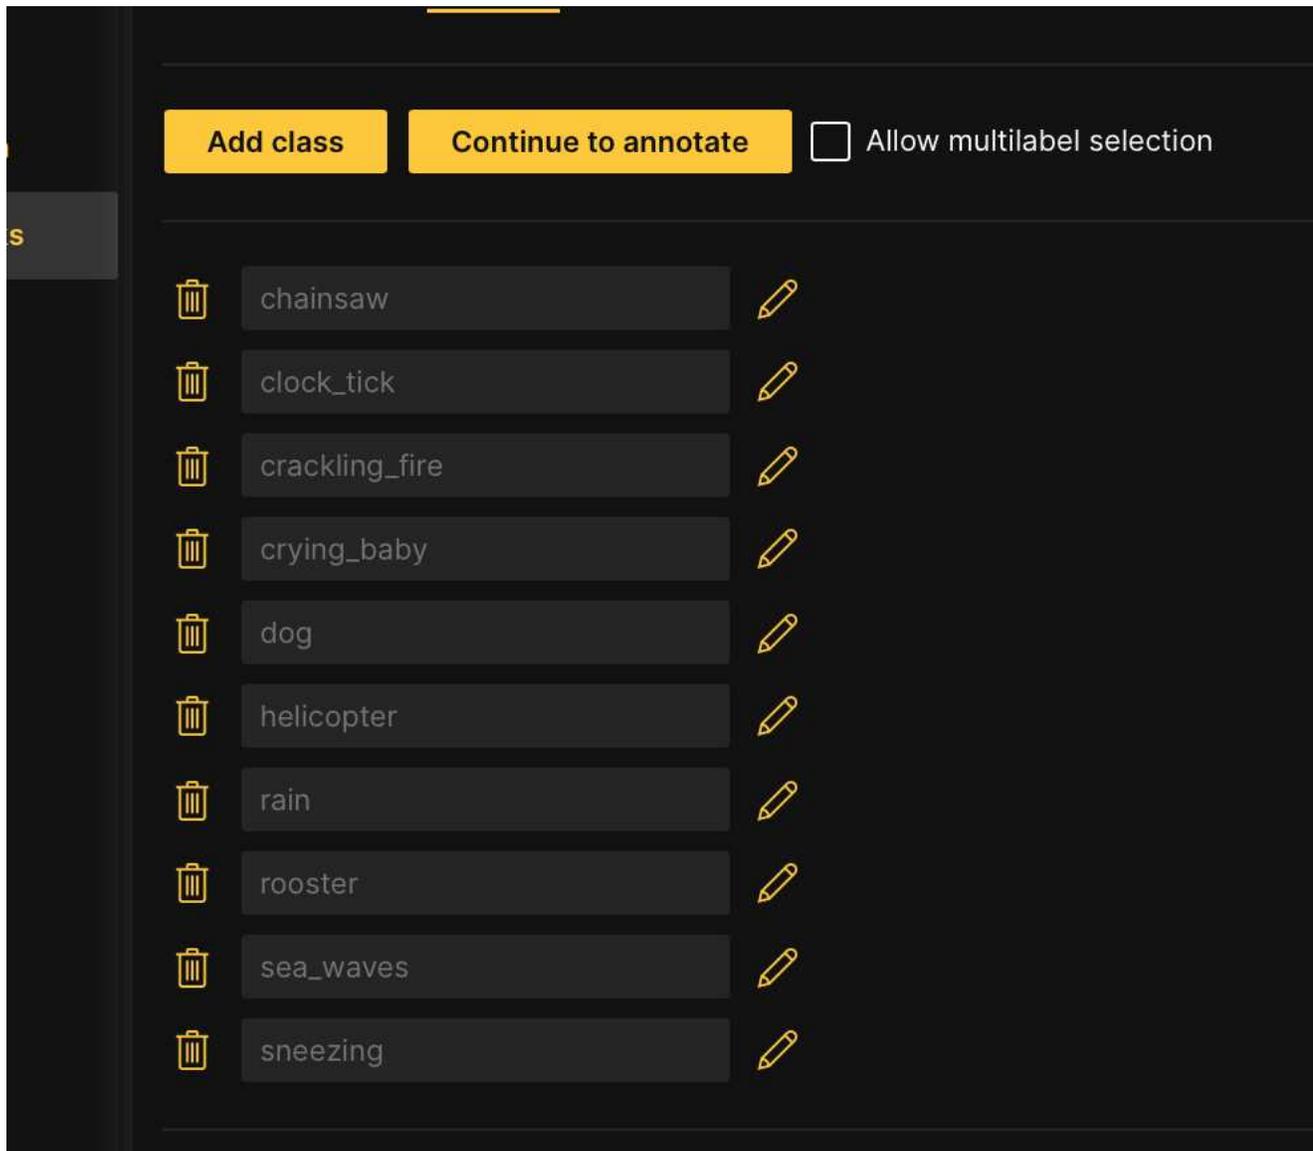
Before annotating our dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. For our dataset, the following are the multiple categorical target labels we want to specify:

- **chainsaw**
- **clock\_tick**
- **crackling\_fire**
- **crying\_baby**
- **dog**
- **helicopter**
- **rain**
- **rooster**
- **sea\_waves**
- **sneezing**

Let's define the annotation task rubric.

1. In the **New class name** box, enter `chainsaw`.

2. Click **Add**.
3. Click **Add class**.
4. In the **New class name** box, enter `clock_tick`.
5. Click **Add**.
6. Repeat the above steps until you create all labels.
7. Click **Continue to annotate**.

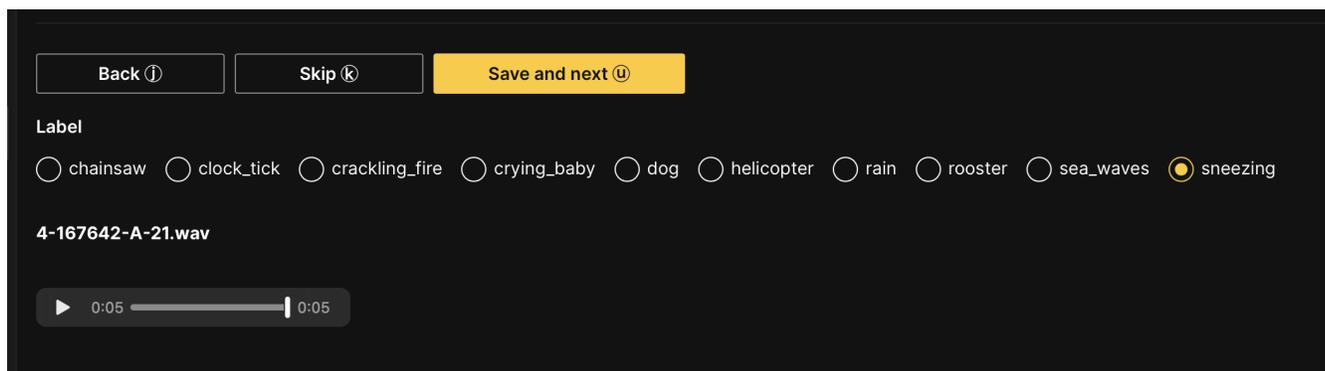


**Note:** H2O Label Genie supports multi-label audio classification annotation tasks.

#### Step 4: Annotate dataset

In the **Annotate** tab, you can individually annotate each audio clip in the dataset. Let's annotate the first audio.

1. Select the label that you associate with the sound audio you hear (for example, **sneezing**).



2. Click **Save and next**. **Note:**

- **Save and next** saves the annotated audio
- To skip an audio clip to annotate later: Click **Skip**.
  - Skipped audio clips (samples) reappear after all non-skipped audios are annotated

3. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate an imported dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#)

## Summary

In this tutorial, we learned the process of annotating and specifying an annotation task rubric for an audio classification annotation task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

## Tutorial 2C: Audio regression annotation task

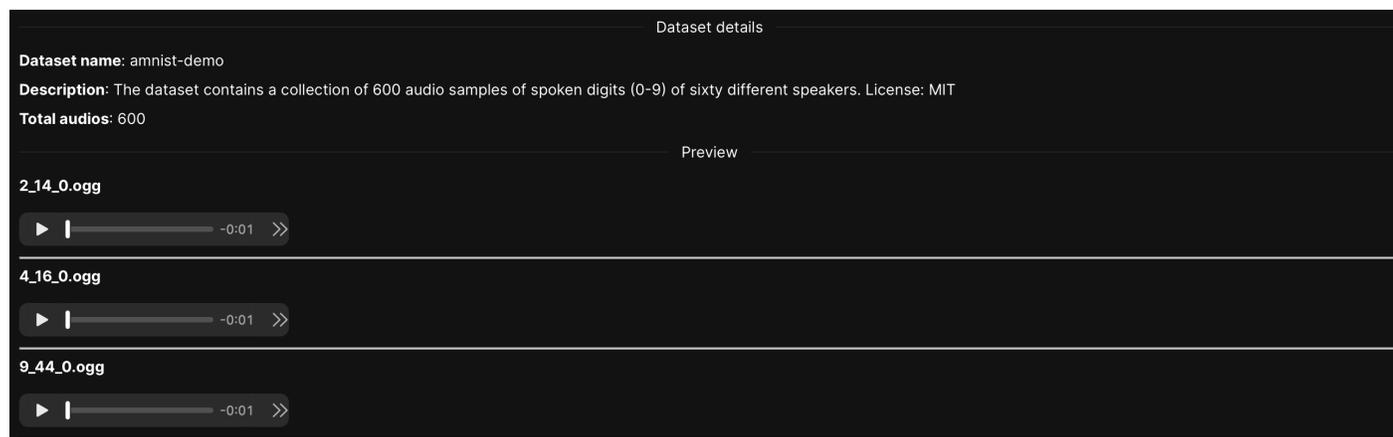
### Overview

This tutorial describes the process of creating an [audio regression annotation task](#), including specifying an annotation task rubric for it. To highlight the process, we will annotate a dataset containing 600 audio samples of spoken digits (0-9) of 60 different speakers.

### Step 1: Explore dataset

We are going to use the preloaded [Amnist demo](#) dataset for this tutorial. The dataset contains **600** samples (audio) of spoken digits (0-9) of sixty different speakers. Let's quickly explore the dataset.

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the **Datasets** table, click **amnist-demo**.



### Step 2: Create an annotation task

Now that we have seen the dataset let's create an annotation task that enables you to annotate the dataset. For this tutorial, an audio regression annotation task refers to assigning one continuous target label to each input audio.

1. Click **New annotation task**.
2. In the **Task name** box, enter `tutorial-2c`.
3. In the **Task description** box, enter `Annotate a dataset containing samples of spoken digits (0-9) of sixty different speakers.`
4. In the **Select task** list, select **Regression**.
5. Click **Create task**.

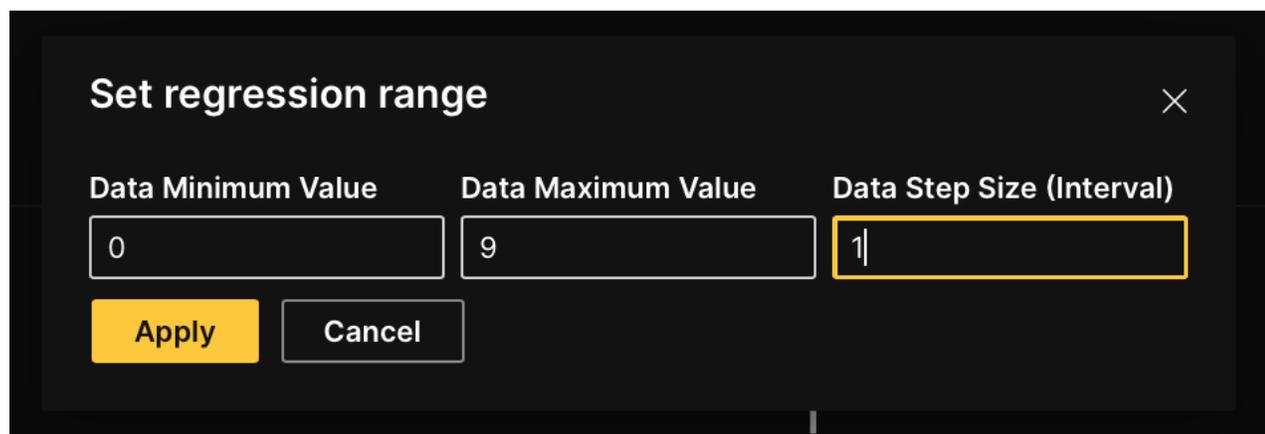
### Step 3: Specify an annotation task rubric

Before we can start annotating our dataset, we need to specify an annotation task rubric. An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. For our dataset, let's label each audio clip with a value from **0** to **9**, where *0* refers to number *0*, *1* refers to number *1*, etc.

1. In the **Data minimum value** box, enter `0`.
  - The **Data minimum value** value refers to the minimum value in your continuous values (in this case, digits ranging from 0 to 9)
2. In the **Data maximum value** box, enter `9`.
  - The **Data maximum value** value refers to the maximum value in your continuous values (in this case, digits ranging from 0 to 9)
3. In the **Data step size (interval)** box, enter `1`.

- The **Data step size (interval)** value refers to the value the label range slider interval takes (the slider its utilize in [step 4](#))

#### 4. Click **Apply**



**Set regression range** ×

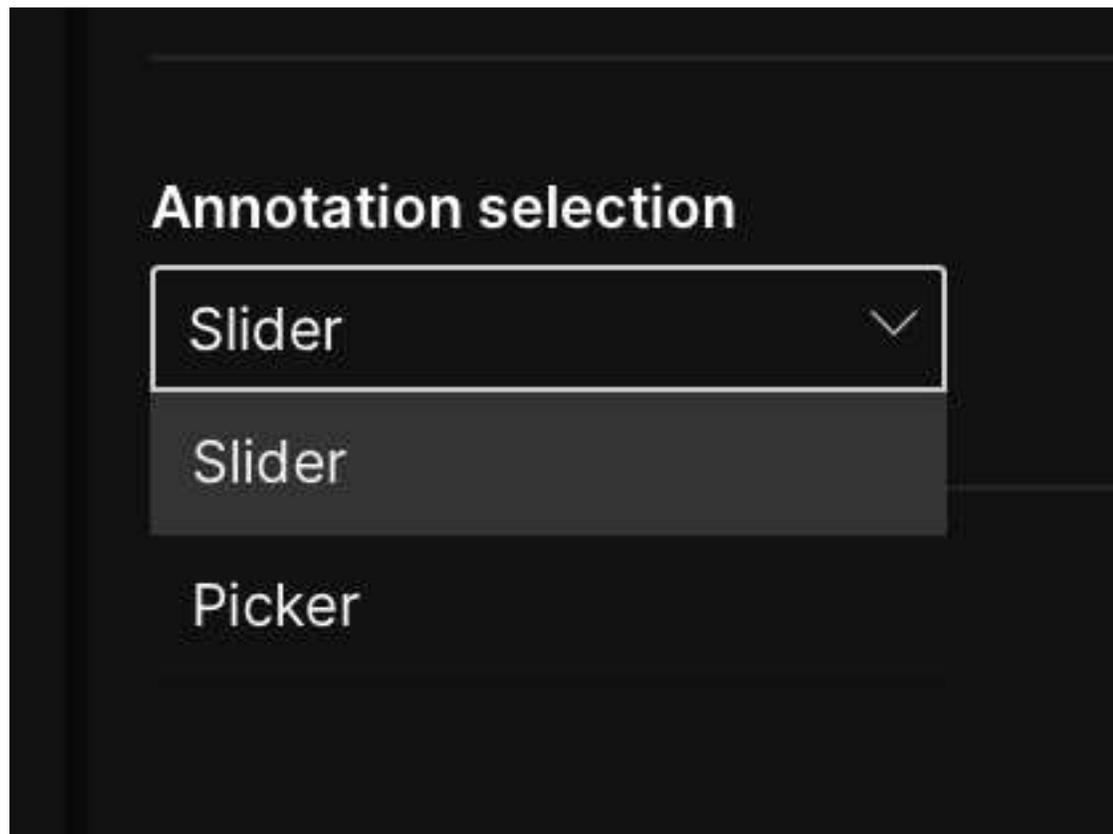
Data Minimum Value    Data Maximum Value    Data Step Size (Interval)

0    9    1

**Apply**    Cancel

Let's utilize the slider, not the picker, to annotate the samples. To enable the slider, consider the following instructions:

1. In the **Annotation selection** list, select **Slider**.



**Annotation selection**

Slider ✓

Slider

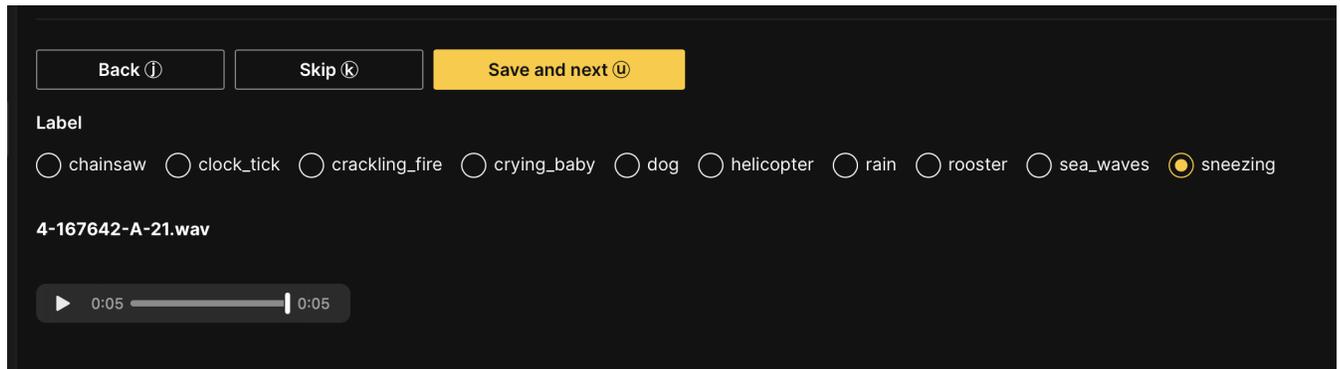
Picker

### Step 4: Annotate dataset

Now that we have specified the annotation task rubric, let's annotate the dataset. 1. Click **Continue to annotate**.

In the **Annotate** tab, you can individually annotate each audio clip in the dataset. Let's annotate the first audio.

1. In the **Label** slider, slide to the label you associate with the audio sound (for example, if you hear the number 6, slide to **6**).



2. Click **Save and next**. **Note:**

- **Save and next** saves the annotated audio
- To skip an audio clip to annotate later: Click **Skip**.
  - Skipped audio clips (samples) reappear after all non-skipped audio clips are annotated

3. Annotate all dataset samples. **Note:** At any point in an annotation task, you can download the already annotated (approved) samples. You do not need to fully annotate an imported dataset to download already annotated samples. To learn more, see [Download an annotated dataset](#).

## Summary

In this tutorial, we learned the process of annotating and specifying an annotation task rubric for an audio regression task.

## Next

To learn the process of annotating and specifying an annotation task rubric for other various annotation tasks in computer vision (CV), natural language processing (NLP), and audio, see [Tutorials](#).

# Import a dataset

## Overview

Before you can import your dataset to H2O Label Genie, the dataset needs to meet the following requirements:

1. The dataset data type needs to be **text**, **image**, or **audio**.
2. Dataset format:
  - The dataset (data) for an image or audio annotation task must be in a **zip** file containing the images or audios.  
**Note:**
    - You can have any nested folder structure inside the **zip** file.
    - All images need to have an image extension. Images can contain a mix of supported image extensions. To learn about supported image extensions, see [Supported image extensions](#).
    - All audios need to have an audio extension. audios can contain a mix of supported audio extensions. To learn about supported audio extensions, see [Supported audio extensions](#).
  - The dataset (data) for a text annotation task must be in a **CSV** file or a **zip** containing a single **CSV** file.
    - One column needs to hold the text data.

### Note:

- To learn how to annotate your dataset, see [Create an annotation task](#).
- To learn how to download an annotated dataset, see [Download an annotated dataset](#).
- To learn how to import an already annotated dataset to H2O Label Genie to review further or edit, see [Import an annotated dataset](#).

## Instructions

To import your dataset (data) to H2O Label Genie, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. Click **Import data**.
3. In the **Name** box, enter a name for the dataset.
4. (Optional) In the **Description** box, enter a description for the dataset.
5. For **Data type**, choose an option.
  - If the data type of the dataset you are importing is **text**: Select **Text**.
  - If the data type of the dataset you are importing is **image**: Select **Image**.
  - If the data type of the dataset you are importing is **audio**: Select **Audio**.
6. In the **Source** list, select the source (data connector) that you want to use (for example, S3).

### Upload

- Click **Browse...**

Or drag and drop the file (dataset)

### S3

- In the **S3 bucket name** box, enter the name of the S3 bucket name.
- In the **AWS access key** box, enter the AWS access key.
- In the **AWS access key** box, enter the AWS access key.

You don't need to enter the AWS access key if the S3 bucket is public.

- In the **AWS secret key** box, enter the AWS secret key.

You don't need to enter the AWS secret key if the S3 bucket is public.

- In the **File name** list, select the file you want to use.

### H2O Drive

- In the **File name** list, select the file you want to use.

7. Click **Import**.

## Dataset connectors

### Overview

H2O Label Genie provides **2+** dataset connectors to access external data sources.

### Supported connectors

**Upload (Standard upload feature)** The following parameter is required:

- **File**

**S3 (Amazon AWS S3)** The following parameters are required:

- **S3 bucket name**
- **AWS access key**
- **AWS secret key**
- **File name**

**H2O Drive (H2O.ai's data storage)** The following parameter is required:

- **File name**

**Note:** To learn more about H2O Drive, see [H2O Drive | Docs](#).

### Supported *image* extensions

The following is a list of supported image extensions for image annotation tasks in H2O Label Genie:

- **Windows bitmaps**
  - **.bmp** (Bitmap Image)
- **JPEG files:**
  - **.jpeg** (JPEG Image)
  - **.jpg** (JPEG Image)
  - **.jpe** (JPEG Image)
- **JPEG 2000 files**
  - **.jp2** (JPEG 2000 Core Image)
- **Portable Network Graphics**
  - **.png** (Portable Network Graphic)
- **WebP:**
  - **.webp** (WebP Image)

- **Portable image format:**
  - .pbm (Portable Bitmap Image)
  - .pgm (Portable Gray Map Image)
  - .ppm (Portable Pixmap Image)
  - .pnm (Portable Any Map Image)
- **TIFF files**
  - .tiff (Tagged Image File Format)
  - .tif (Tagged Image File)
- **Radiance HDR**
  - .hdr (High Dynamic Range Image File)

## Supported *audio* extensions

The following is a list of supported audio extensions for audio annotation tasks in H2O Label Genie:

- **Uncompressed**
  - .wav (WAVE Audio)
  - .aiff (Audio Interchange File Format)
- **Lossless compressed**
  - .flac (Free Lossless Audio Codec File)
- **Lossy compressed**
  - .mp3 (MP3 Audio)
  - .ogg (Ogg Vorbis Audio)

## Import an annotated dataset

### Overview

After creating an annotation task, you can import to the annotation task already annotated samples (for example, annotated images) from the same dataset you used to create the annotation task.

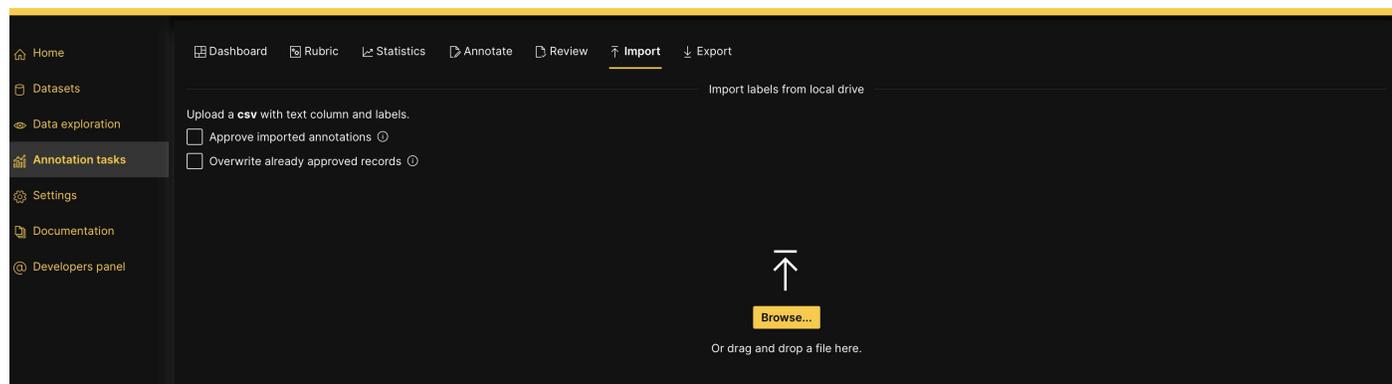
### Instructions

To import an annotated dataset to an annotation task, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the *annotation tasks* table, double-click the row where the annotation task you want to use is located.
3. Click **Import** tab.
4. Upload your dataset with existing annotations.
5. Click **Import annotated samples**.

### Note:

- H2O Label Genie allows you to review and approve the imported annotations later. - To import them pre-approved: Click the **Approve imported labels** checkbox (so that you don't have to approve them in the **Annotate** tab). - H2O Label Genie does not overwrite already approved records by default. - To enable overwrites: Click the **Overwrite already approved records** checkbox. - Before importing an annotated dataset to an existing annotation task, the annotated dataset needs to be formatted in a certain way depending on the problem type it aims to solve. To learn more, see [Required annotated dataset format](#).



### Required annotated dataset format

Before importing an annotated dataset to an existing annotation task, the annotated dataset needs to be formatted in a certain way depending on the problem type it aims to solve.

#### Text classification

The annotated dataset to import to a text classification annotation task needs to be formatted as follows:

A **CSV** file containing a *text* and *label* column. Columns:

- **text:** The *text* column needs to contain the original text input. **Note:** The **text** column is used to merge the labels. Please make sure to use the original text column without any duplications. The original text column refers to the text column present in the dataset you used to create the annotation task.
- **label:** The *label* column needs to contain the labels attribute to the texts specified in the *text* column.

#### Text regression

The annotated dataset to import to a text regression annotation task needs to be formatted as follows:

A **CSV** file containing a *text* and *label* column. Columns:

- **text:** The *text* column needs to contain the original text input. **Note:** The **text** column is used to merge the labels. Please make sure to use the original text column without any duplications. The original text column refers to the text column present in the dataset you used to create the annotation task.
- **label:** The label column needs to contain the labels attributed to the texts specified in the *text* column.

### Text-entity recognition

The annotated dataset to import to a text-entity recognition annotation task needs to be formatted as follows:

A **Parquet** file containing a *raw\_text*, *text*, and *label* column. Columns:

- **raw\_text:** The *text* column needs to contain the original text input. **Note:** The **raw\_text** column is used to merge the labels. Please make sure to use the original text column without any duplications. The original text column refers to the text column present in the dataset you used to create the annotation task.
- **text:** The *text* column needs to contain the lists of entities.
- **label:** The *label* column needs to contain the lists of labels for the entities.

### Text summarization

The annotated dataset to import to a text summarization annotation task needs to be formatted as follows:

A **CSV** file containing a *text* and *label* column. Columns:

- **text:** The *text* column needs to contain the original text input. **Note:** The **text** column is used to merge the labels. Please make sure to use the original text column without any duplications. The original text column refers to the text column present in the dataset you used to create the annotation task.
- **label:** The *label* column needs to contain the summaries attributed to the texts specified in the *text* column.

### Text-generative AI

The annotated dataset to import to a text-generative AI annotation task needs to be formatted as follows:

A **CSV** file containing the following columns:

- **text:** The *text* column needs to contain the original text input. **Note:** The **text** column is used to merge the labels. Please make sure to use the original text column without any duplications. The original text column refers to the text column present in the dataset you used to create the annotation task.
- **label:** The *label* column needs to contain the answers attributed to the texts specified in the *text* column.

### Image classification

The annotated dataset to import to an image classification annotation task needs to be formatted as follows:

A **CSV** file containing an *image\_path* and *label* column. Columns:

- **image\_path:** The *image\_path* column needs to specify the locations of the images. **Note:** The **image\_path** column is used to merge the labels. Please make sure to use the correct relative paths without any duplications.
- **label:** The *label* column needs to contain the labels attributed to the images specified in the *image\_path* column.

### Image regression

The annotated dataset to import to an image regression annotation task needs to be formatted as follows:

A **CSV** file containing an *image\_path* and *label* column. Columns:

- **image\_path:** The *image\_path* column needs to specify the locations of the images. **Note:** The **image\_path** column is used to merge the labels. Please make sure to use the correct relative paths without any duplications.
- **label:** The label column needs to contain the labels attributed to the images specified in the *image\_path* column.

## Object detection

The annotated dataset to import to an object detection annotation task needs to be formatted as follows:

A **Parquet** file containing an *image\_path* and *class\_id* column; the file should also contain an *x\_min*, *x\_max*, *y\_min*, and *y\_max* column corresponding to the bounding box locations. Columns:

- **image\_path**: The *image\_path* column needs to specify the locations of the images. **Note**: The **image\_path** column is used to merge the labels. Please make sure to use the correct relative paths without any duplications.
- **class\_id**: The *class\_id* column needs to contain the class IDs of the bounding box(es).
- **x\_min, x\_max, y\_min, and y\_max**: The *x\_min*, *x\_max*, *y\_min*, and *y\_max* specify the spatial location of the bounding box(es).

## Image instance segmentation

### Format

The annotated dataset to import to an image instance segmentation annotation task needs to be formatted as follows:

A **JSON** file in a COCO format containing the following keys:

- **images**: The image ID and *file\_name* (location). **Note**: The **file\_name** is used to merge the labels. Please make sure to use the correct relative paths without any duplications.
- **annotations**: The *annotations* in a COCO polygon format.
- **categories**: The *categories* ID and name.

### Example

```
{
  "images": [
    {
      "id": 0,
      "file_name": "car_or_coffee_sample/pexels-anna-tukhfatullina-food-photographerstylist-2648988.jpg",
      "width": 3423,
      "height": 4279
    }
  ],
  "annotations": [
    {
      "id": "1f138007-30e6-4fdc-b855-bad58ea5cce0",
      "image_id": 0,
      "category_id": 1,
      "iscrowd": 0,
      "segmentation": [
        [
          1377.548780487805,
          2156.21484375,
          1385.8975609756098,
          2198.001953125,
          1444.3390243902438,
          2256.50390625,
          1594.6170731707318,
          2306.6484375,
          1711.5,
          2323.36328125,
          1895.1731707317074,
          2315.005859375,
          2045.451219512195,
          2289.93359375,
          2145.6365853658535,
          2256.50390625,

```

```
2212.4268292682927,
2198.001953125,
2212.4268292682927,
2156.21484375,
2170.682926829268,
2114.427734375,
2053.7999999999997,
2055.92578125,
1861.778048780488,
2022.49609375,
1719.8487804878048,
2022.49609375,
1561.2219512195122,
2047.568359375,
1419.2926829268292,
2106.0703125
]
],
},
{
  "id": "bbe67236-ced4-475a-b964-9c36d7511fd6",
  "image_id": 0,
  "category_id": 1,
  "iscrowd": 0,
  "segmentation": [
    [
      1477.7341463414634,
      2039.2109375,
      1402.5951219512194,
      2072.640625,
      1327.4560975609756,
      2181.287109375,
      1327.4560975609756,
      2415.294921875,
      1360.8512195121953,
      2565.728515625,
      1452.6878048780488,
      2766.306640625,
      1502.780487804878,
      2950.169921875,
      1561.2219512195122,
      3042.1015625,
      1644.709756097561,
      3092.24609375,
      1853.429268292683,
      3092.24609375,
      1953.6146341463414,
      3058.81640625,
      2204.0780487804877,
      2849.880859375,
      2237.4731707317073,
      2732.876953125,
      2229.1243902439023,
      2565.728515625,
      2270.868292682927,
      2390.22265625,
      2270.868292682927,
      2172.9296875,
      2204.0780487804877,
```

```
2080.998046875,  
2128.939024390244,  
2047.568359375,  
1920.2195121951222,  
2005.78125,  
1678.1048780487804,  
2005.78125  
]  
]  
}  
],  
"categories": [  
{  
  "id":0,  
  "name":"Car"  
},  
{  
  "id":1,  
  "name":"Coffee"  
}  
]  
}
```

### Audio classification

The annotated dataset to import to an audio classification annotation task needs to be formatted as follows:

A **CSV** file containing an *audio\_path* and *label* column. Columns:

- **audio\_path**: The *audio\_path* column needs to specify the locations of the audios. **Note:** The **audio\_path** column is used to merge the labels. Please make sure to use the correct relative paths without any duplications.
- **label**: The *label* column needs to contain the labels attributed to the audios specified in the *audio\_path* column.

### Audio regression

The annotated dataset to import to an audio regression annotation task needs to be formatted as follows:

A **CSV** file containing an *audio\_path* and *label* column. Columns:

- **audio\_path**: The *audio\_path* column needs to specify the locations of the audios. **Note:** The **audio\_path** column is used to merge the labels. Please make sure to use the correct relative paths without any duplications.
- **label**: The label column needs to contain the labels attributed to the audios specified in the *audio\_path* column.

## Access a demo dataset

### Overview

In H2O Label Genie, you can use demo datasets to explore [supported annotation tasks](#).

### Instructions

To access a demo dataset in H2O Label Genie, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the *datasets* table, select one of the [demo datasets in H2O Label Genie](#).

### Note:

- After selecting a demo dataset, click **New annotation task** to annotate the dataset.
  - To learn how to annotate your dataset, see [Tutorials](#)

id	name	type	# samples
8	cnn-dailymail-sample	Text	100
7	amnist-demo	Audio	600
6	esc10-audio-demo	Audio	400
5	plant-pathology-demo	Image	32
4	text-readability-demo	Text	100
3	twitter-demo	Text	90
2	car-or-coffee-demo	Image	40
1	amazon-reviews-demo	Text	180

## Demo datasets in H2O Label Genie

### Amazon reviews demo

- **Dataset name:** amazon-reviews-demo
- **Description:** The dataset contains user reviews (in text format) and ratings (from 0 to 5) of Amazon products.
- **Dataset columns:** stars, comment
- **Problem type:** Text-generative AI, text classification, text regression, text-entity recognition
- **License:** [CC0 1.0 Universal \(CC0 1.0\)](#)

### Car or coffee demo

- **Dataset name:** car-or-coffee-demo
- **Description:** The dataset contains images of cars and coffee.
- **Problem type:** Image classification, object detection
- **License:** [Pexels license](#)

### Twitter demo

- **Dataset name:** twitter-demo
- **Description:** The dataset contains tweets that can be used to analyze tweet sentiments and recognize the emotion in text tweets.
- **Dataset columns:** text, sentiment
- **Problem type:** Text-generative AI, text classification, text-entity recognition
- **License:** [Attribution 4.0 International \(CC BY 4.0\)](#)

### Text readability demo

- **Dataset name:** text-readability-demo
- **Description:** This dataset contains excerpts, and it is part of the CLEAR Corpus.
- **Dataset columns:** id, excerpt
- **Problem type:** Text-generative AI, text regression, text-entity recognition
- **License:** [MIT license](#)

### CNN Daily Mail sample

- **Dataset name:** cnn-dailymail-sample
- **Description:** The dataset contains human-generated abstract summaries from news stories published on the CNN and Daily Mail websites.
- **Dataset columns:** id, text, summary
- **Problem type:** Text-generative AI, text summarization, text classification
- **License:** [MIT license](#)

### Plant pathology demo

- **Dataset name:** plant-pathology-demo
- **Description:** This dataset contains images of healthy and diseased apple leaves for plant pathology recognition.
- **Problem type:** Image classification, image regression, object detection
- **License:** [Attribution 4.0 International \(CC BY 4.0\)](#)

### ESC10 audio demo

- **Dataset name:** esc10-audio-demo
- **Description:** This dataset contains 5-second-long recordings of environmental sounds organized into ten classes (with 40 examples per class). Clips in this dataset have been manually extracted from public field recordings gathered by the [Freesound.org](#) project.
- **Problem type:** Audio classification
- **License:** [Attribution 3.0 Unported \(CC BY 3.0\)](#)

### Amnist demo

- **Dataset name:** amnist-demo
- **Description:** The dataset contains a collection of 600 audio samples of spoken digits (0-9) of sixty different speakers.
- **Problem type:** Audio regression
- **License:** [MIT license](#)

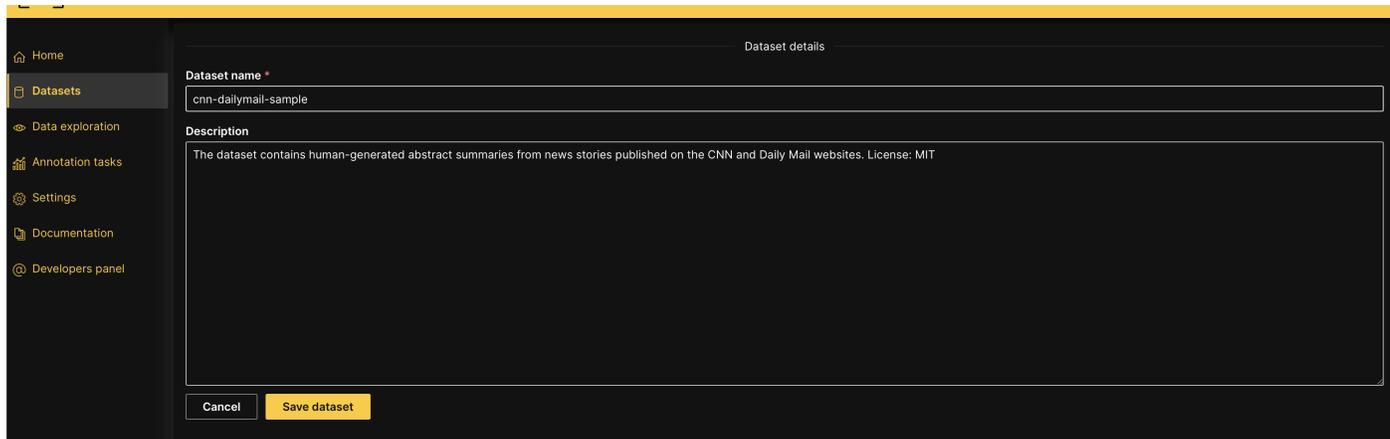
## Edit a dataset

### Overview

H2O Label Genie enables you to edit the name and description of an imported dataset. ## Instructions

To edit a dataset's name and description, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the *datasets* table, double-click the row where the dataset you want to edit is located.
3. Click **Edit dataset**.
4. In the **Dataset name** box, edit the dataset name.
5. (Optional) In the **Description** box, edit the dataset description.
6. Click **Save dataset**.

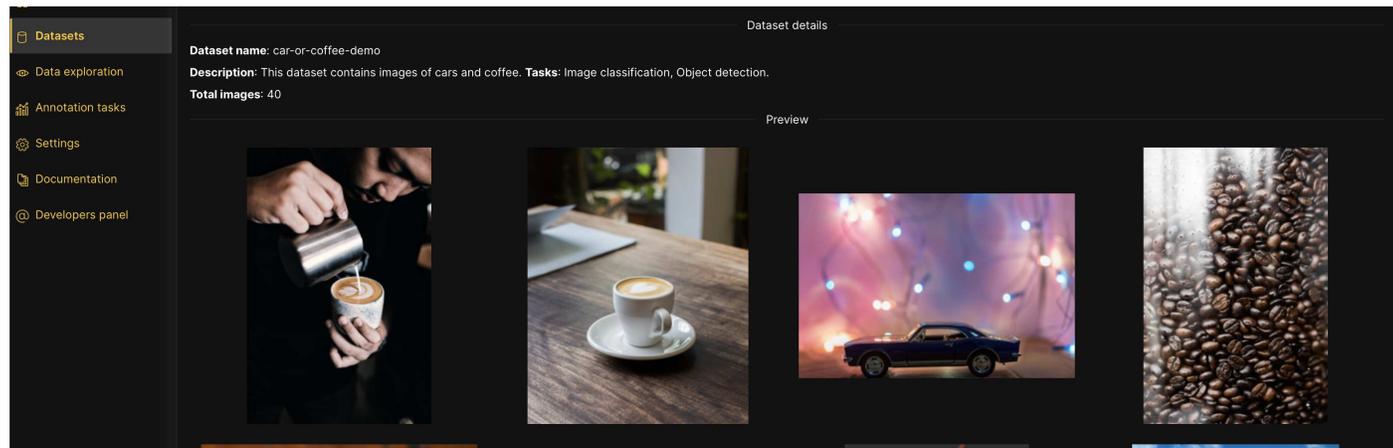


The screenshot shows the 'Dataset details' form in the H2O Label Genie interface. On the left is a navigation sidebar with options: Home, Datasets (selected), Data exploration, Annotation tasks, Settings, Documentation, and Developers panel. The main content area has a title 'Dataset details' and a 'Dataset name' field containing 'cnn-dailymail-sample'. Below it is a 'Description' text area containing the text: 'The dataset contains human-generated abstract summaries from news stories published on the CNN and Daily Mail websites. License: MIT'. At the bottom of the form are two buttons: 'Cancel' and 'Save dataset'.

# View a dataset

## Overview

Viewing a dataset or all your datasets lets you review, edit, or start a new annotation task.



## Instructions

To view an imported dataset, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. In the *datasets* table, double-click the row where the dataset you want to view is located.

# Delete a dataset

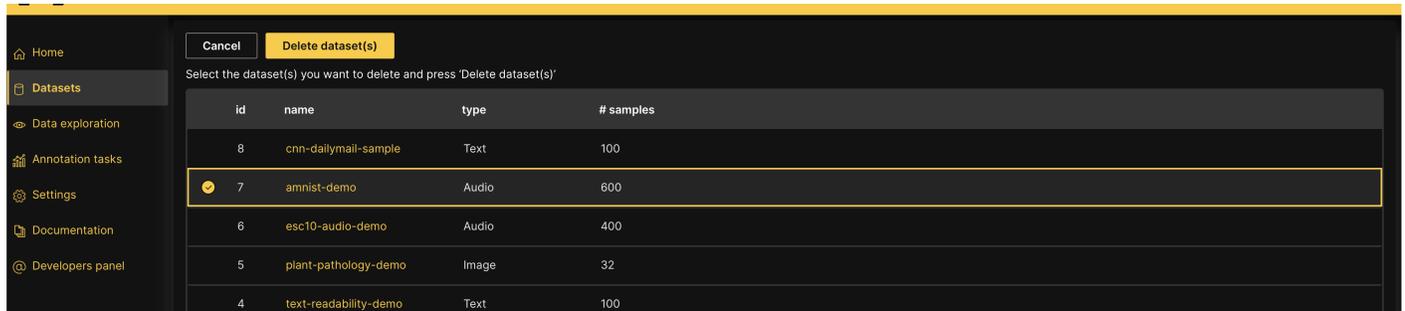
## Overview

When a dataset is no longer needed, you can delete it. Deleted datasets are permanently removed from the H2O Label Genie instance.

## Instructions

To delete one or multiple datasets, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Datasets**.
2. Click **Delete dataset(s)**.
3. Select the dataset(s) you want to delete.
4. Click **Delete dataset(s)**.



The screenshot shows the H2O Label Genie interface with the 'Delete dataset(s)' dialog box open. The dialog box has a 'Cancel' button and a 'Delete dataset(s)' button. Below the buttons, there is a prompt: 'Select the dataset(s) you want to delete and press 'Delete dataset(s)''.

id	name	type	# samples
8	cnn-dailymail-sample	Text	100
7	ammist-demo	Audio	600
6	esc10-audio-demo	Audio	400
5	plant-pathology-demo	Image	32
4	text-readability-demo	Text	100

## Download an annotated dataset

### Overview

At any point in an annotation task, you can download the already approved annotated samples. You do not need to *fully* annotate an imported dataset to download already annotated samples (approved samples).

**Note:** An array of datasets labeled in H2O Label Genie are supported in [H2O Hydrogen Torch](#) and [H2O LLM Studio](#).

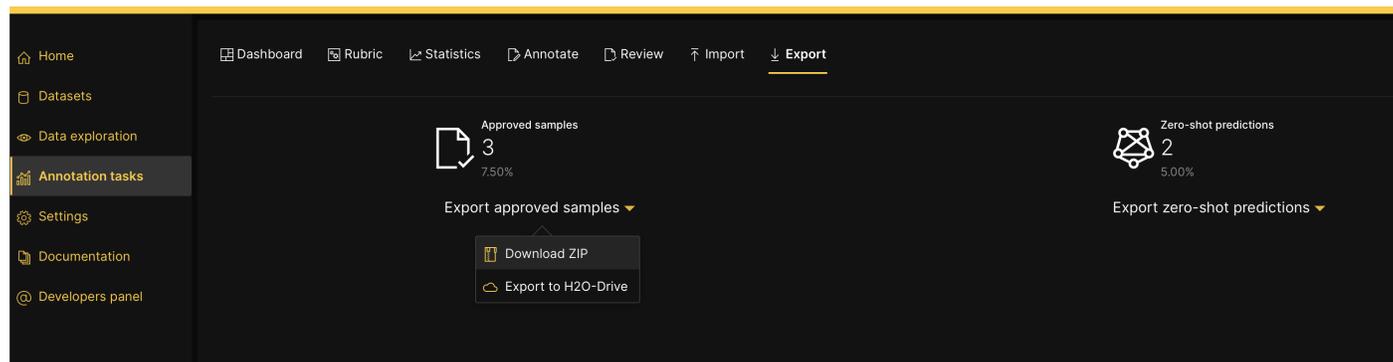
### Instructions

To download an annotated dataset (approves samples), consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the *annotation tasks* table, double-click the row where the annotation task you want to download is located.
3. Click the **Export** tab.
4. In the **Export approved samples** list, select **Download ZIP**. **Note:** H2O Label Genie downloads a **zip** file containing the annotated dataset in a particular format depending on the dataset's annotation task. To learn more, see [Downloaded dataset formats](#).

#### Note:

- **Zero-shot learning models:** H2O Label Genie generates zero-shot predictions for specific annotation tasks that can be downloaded or exported to H2O Drive.
  - To learn how to download a dataset's zero-shot predictions, see [Download a dataset's zero-shot predictions](#)
  - To learn how to export a dataset's zero-shot predictions to H2O Drive, see [Export a dataset's zero-shot predictions to H2O Drive](#)
- **Export an annotated dataset to H2O Drive:** You can export an annotated dataset (approved samples) to H2O Drive. To learn more, see [Export an annotated dataset to H2O Drive](#).



## Downloaded dataset formats

### Text classification

A downloaded text classification dataset (with approved samples) follows the following dataset format: A **zip** file (1) containing a **CSV** file (2):

- folder\_name.zip (1)
  - csv\_name.csv (2)

1. A **zip** file.
2. A **CSV** file containing a **text** and **label** column:

- **text**: The **text** column contains the text input
- **label**: The **label** column contains the labels attributed to the texts specified in the **text** column

#### **Note:**

[H2O Hydrogen Torch] (<https://docs.h2o.ai/h2o-hydrogen-torch/>) supports the downloaded dataset of a text class

### ### Text regression

A downloaded text regression dataset (with approved samples) follows the following dataset format: A **\*\*zip\*\***

- folder\_name.zip (1)
  - csv\_name.csv (2)

1. A **zip** file.
2. A **CSV** file containing a **text** and **label** column:
  - **text**: The *text* column contains the text input
  - **label**: The label column contains the labels attributed to the texts specified in the *text* column

**Note:** [H2O Hydrogen Torch](#) supports the downloaded dataset of a text regression annotation task. In particular, H2O Hydrogen Torch supports the downloaded dataset for a text regression experiment.

### Text-entity recognition

A downloaded text entity recognition dataset (with approved samples) follows the following dataset format: A **zip** file (1) containing a .pq file (2):

- folder\_name.zip (1)
  - pq\_name.pq (2)

1. A **\*\*zip\*\*** file.
2. A `.pq`` file containing a **\*\*text\*\*** and **\*\*label\*\*** column:

- **\*\*text\*\***: The *\*text\** column contains the text input
- **\*\*label\*\***: The label column contains the labels attributed to the text-entities specified in the *\*text\** column

### ### Text summarization

A downloaded text summarization dataset (with approved samples) follows the following dataset format: A **\*\*zip\*\***

- folder\_name.zip (1)
  - csv\_name.csv (2)

1. A **zip** file.
2. A **CSV** file containing a **text** and **label** column:
  - **text**: The *text* column contains the text input
  - **label**: The *label* column contains the summaries attributed to the texts specified in the *text* column

### Text-generative AI

A downloaded text-generative AI dataset (with approved samples) follows the following dataset format: A **ZIP** file (1) containing a **CSV** file (2):

- folder\_name.zip (1)
  - csv\_name.csv (2)

1. A **\*\*zip\*\*** file.
2. A **\*\*CSV\*\*** file containing the following columns:

- **\*\*llm\_input\*\***: The final prompts sent to the large language model (LLM) as an input query
- **\*\*llm\_output\*\***: The answers provided to the query (**\*\*llm\_input\*\***)
- **\*\*original columns\*\***: All the original columns

**\*\*Note:\*\***

[H2O LLM Studio](https://docs.h2o.ai/h2o-llmstudio/) supports the downloaded dataset of a text-generative AI annotation task.

**### Image classification**

A downloaded image classification dataset (with approved samples) follows the following dataset format: A **\*\*zip\*\***

- folder\_name.zip (1)
  - csv\_name.csv (2)
  - image\_folder\_name (3)
    - name\_of\_image.image\_extension
    - name\_of\_image.image\_extension
    - name\_of\_image.image\_extension
  - ...
- 1. A **zip** file.
- 2. A **CSV** file containing an **image\_path** and **label** column:
  - **image\_path**: The *image\_path* column specifies the location of the image
  - **label**: The *label* column contains the labels attributed to the images specified in the *image\_path* column
- 3. Image folder that contains all the images specified in the *image\_path* column.

**Note:** [H2O Hydrogen Torch](https://docs.h2o.ai/h2o-hydrogen-torch/) supports the downloaded dataset of an image classification annotation task. In particular, H2O Hydrogen Torch supports the downloaded dataset for an image classification experiment.

**Image regression**

A downloaded image regression dataset (with approved samples) follows the following dataset format: A **zip** file (1) containing a **CSV** file (2) and an *image* folder (3).

- folder\_name.zip (1)
  - csv\_name.csv (2)
  - 
  - image\_folder\_name (3)
    - name\_of\_image.image\_extension
    - name\_of\_image.image\_extension
    - name\_of\_image.image\_extension
  - ...
- 1. A **zip** file.
- 2. A **CSV** file containing an **image\_path** and **label** column:
  - **image\_path**: The *image\_path* column specifies the location of the image
  - **label**: The *label* column contains the labels attributed to the images specified in the *image\_path* column
- 3. Image folder that contains all the images specified in the *image\_path* column

**\*\*Note:\*\***

[H2O Hydrogen Torch](https://docs.h2o.ai/h2o-hydrogen-torch/) supports the downloaded dataset of an image regression task.

**### Object detection**

A downloaded object detection dataset (with approved samples) follows the following dataset format: A **\*\*zip\*\***

- folder\_name.zip (1)
    - pq\_name.pq (2)
  - image\_folder\_name (3)
    - name\_of\_image.image\_extension
    - name\_of\_image.image\_extension
    - name\_of\_image.image\_extension
  - ...
1. A **zip** file.
  2. A **.pq** file containing an **image\_path** and **class\_id** column; the file also contains an **x\_min**, **x\_max**, **y\_min**, and **y\_max** column corresponding to the bounding box locations
    - **image\_path**: The *image\_path* column specifies the location of the image
    - **class\_id**: The *class\_id* column contains the class IDs of the bounding box(es)
    - **x\_min,x\_max,y\_min**, and **y\_max**: The *x\_min*, *x\_max*, *y\_min*, and *y\_max* specify the spatial location of the bounding box(es)
  3. Image folder that contains all the images specified in the *image\_path* column.

**Note:** [H2O Hydrogen Torch](#) supports the downloaded dataset of an object detection task. In particular, H2O Hydrogen Torch supports the downloaded dataset for an image object detection experiment.

### Image instance segmentation

A downloaded image instance segmentation dataset (with approved samples) follows the following dataset format: A **zip** file (1) containing a **JSON** file (2) and an *image* folder (3).

- folder\_name.zip (1)
    - json\_name.json (2)
  - image\_folder\_name (3)
    - name\_of\_image.image\_extension
    - name\_of\_image.image\_extension
    - name\_of\_image.image\_extension
  - ...
1. A **\*\*zip\*\*** file.
  2. A **\*\*JSON\*\*** file containing the labels in a COCO format (+ the annotated dataset's original columns)
  3. Image folder that contains all the images specified in the **\*\*JSON\*\*** file.

**\*\*Note:\*\***

[H2O Hydrogen Torch] (<https://docs.h2o.ai/h2o-hydrogen-torch/>) supports the downloaded dataset of an image instance segmentation task.

### ### Audio classification

A downloaded audio classification dataset (with approved samples) follows the following dataset format: A **\*\*zip\*\***

- folder\_name.zip (1)
  - csv\_name.csv (2)
- audio\_folder\_name (3)
  - name\_of\_audio.audio\_extension
  - name\_of\_audio.audio\_extension
  - name\_of\_audio.audio\_extension

- ...
1. A **zip** file.
  2. A **CSV** file containing an **audio\_path** and **label** column:
    - **audio\_path**: The *audio\_path* column specifies the location of the audio
    - **label**: The *label* column contains the labels attributed to the audios specified in the *audio\_path* column
  3. Audio folder that contains all the audios specified in the *audio\_path* column.

**Note:** [H2O Hydrogen Torch](#) supports the downloaded dataset of an audio classification annotation task. In particular, H2O Hydrogen Torch supports the downloaded dataset for an audio classification experiment.

### Audio regression

A downloaded audio regression dataset (with approved samples) follows the following dataset format: A **zip** file (1) containing a **CSV** file (2) and an *audio* folder (3).

“ - folder\_name.zip (1) - csv\_name.csv (2) - - audio\_folder\_name (3) - name\_of\_audio.audio\_extension - name\_of\_audio.audio\_extension - name\_of\_audio.audio\_extension - ...

1. A **zip** file.
2. A **CSV** file containing an **audio\_path** and **label** column:
  - **audio\_path**: The *audio\_path* column specifies the location of the audio
  - **label**: The *label* column contains the labels attributed to the audios specified in the *audio\_path* column
3. Audio folder that contains all the audios specified in the *audio\_path* column.

**Note:** [H2O Hydrogen Torch](#) supports the downloaded dataset of an audio regression annotation task. In particular, H2O Hydrogen Torch supports the downloaded dataset for an audio regression experiment.

## Download a dataset's zero-shot predictions

### Overview

H2O Label Genie generates zero-shot predictions for specific annotation tasks that can be downloaded. In particular, H2O Label Genie can generate zero-shot predictions for the following supported annotation tasks:

- Image classification
- Object detection
- Image instance segmentation
- Text classification
- Text summarization
- Text-generative AI

#### What are zero-shot learning model predictions?

The labels or suggested labels for a given sample that are provided by a zero-shot learning model are called zero-shot learning model predictions. For example, for an image and text classification annotation task, H2O Label Genie, with a zero-shot learning model activated, offers a percentage probability of an image or text belonging to a certain label (class). For an object detection annotation task, it populates an image with bounding boxes where the bounding boxes capture the desired objects (for example, a car).

To learn more about zero-shot predictions, see [Zero-shot learning models](#).

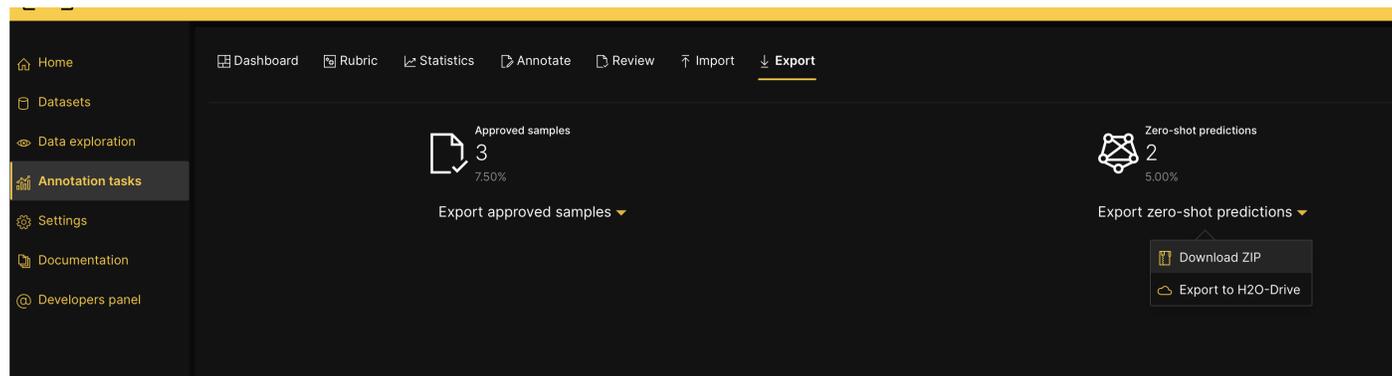
### Instructions

To download the zero-shot predictions of an annotation task, consider the following instructions:

#### Caution:

- If the **Enable zero-shot predictions** setting is turned **Off**, the zero-shot learning model utilized for the annotation is not available during the annotation process while preventing the generation of zero-shot predictions. To turn **On** the **Enable zero-shot predictions** setting, see [Smart annotation settings: Enable zero-shot predictions](#).
- If the annotation task was recently created (less than a minute ago), you might need to wait a bit before the zero-shot predictions become available for download.
  - H2O Label Genie notifies you to **Refresh** the instance when zero-shot predictions (suggestions) are available.

1. In the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the **Annotation tasks** table, double-click the row where the annotation task you want to download its zero-shot predictions is located.
3. Click the **Export** tab.
4. In the **Export zero-shot predictions** list, select **Download ZIP**.



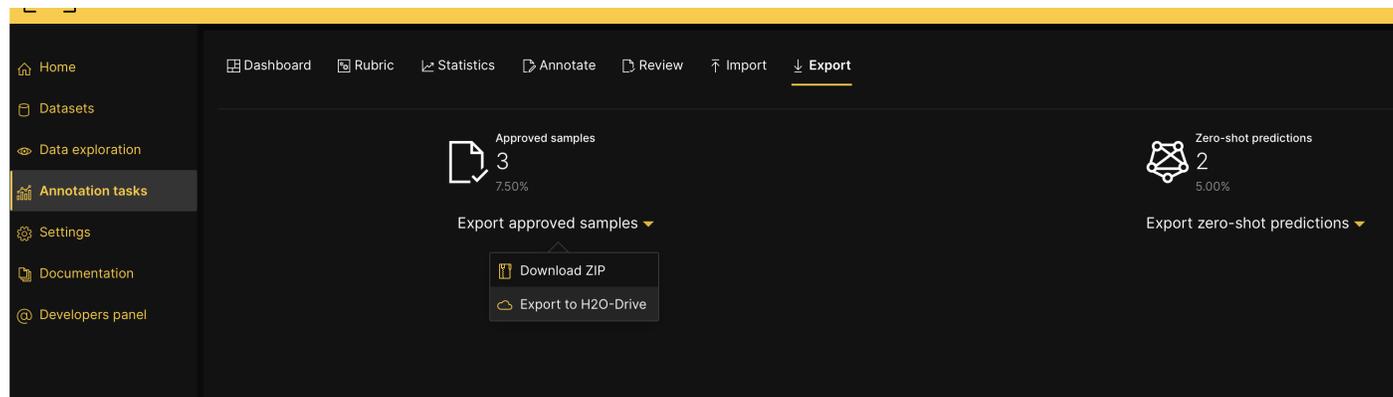
# Export an annotated dataset to H2O Drive

## Overview

H2O Label Genie allows you to export an annotated dataset to H2O Drive. H2O Label Genie only includes approved samples in the dataset to be imported to H2O Drive.

## Instructions

1. In the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the *annotation tasks* table, double-click the row where the annotation task (annotated dataset (approved samples)) you want to export to H2O Drive is located.
3. Click the **Export** tab.
4. In the **Export approved samples** list, select **Export to H2O Drive**.



# Export a dataset's zero-shot predictions to H2O Drive

## Overview

You can export an annotation task's zero-shot predictions to H2O Drive.

### What are zero-shot learning model predictions?

The labels or suggested labels for a given sample that are provided by a zero-shot learning model are called zero-shot learning model predictions. For example, for an image and text classification annotation task, H2O Label Genie, with a zero-shot learning model activated, offers a percentage probability of an image or text belonging to a certain label (class). For an object detection annotation task, it populates an image with bounding boxes where the bounding boxes capture the desired objects (for example, a car).

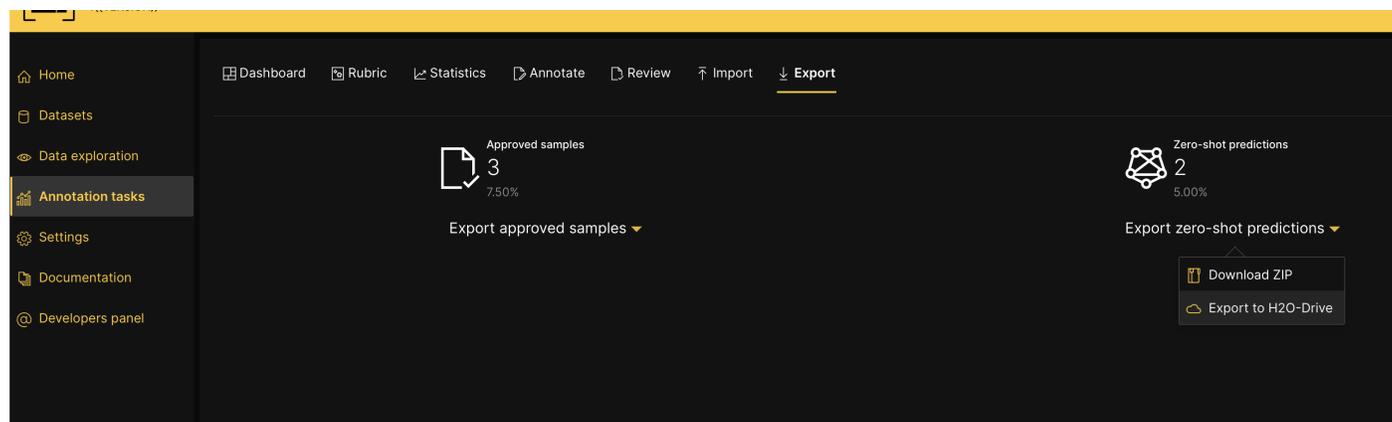
To learn more about zero-shot predictions, see [Zero-shot models](#).

## Instructions

### Caution:

- If the **Enable zero-shot predictions** setting is turned **Off**, the zero-shot learning model utilized for the annotation is not available during the annotation process while preventing the generation of zero-shot predictions. To turn **On** the **Enable zero-shot predictions** setting, see [Smart annotation](#).
- If the annotation task was recently created (less than a minute ago), you might need to wait a bit before the zero-shot predictions become available for download.
  - H2O Label Genie notifies you to **Refresh** the instance when zero-shot predictions (suggestions) are available.

1. In the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the *annotation tasks* table, double-click the row where the annotation task you want to export its zero-shot predictions is located.
3. Click the **Export** tab.
4. In the **Export zero-shot predictions** list, select **Export to H2O Drive**



## Supported clustering tasks

### Overview

H2O Label Genie supports the following clustering tasks:

- [Image clustering](#)
- [Text clustering](#)

**Note:** To learn how to create a clustering task, see [Create a clustering task](#).

### Image clustering

An image clustering task groups images into clusters where a particular cluster contains similar images.

The screenshot displays the H2O Label Genie interface for an image clustering task. At the top, there are two buttons: "Edit cluster task" and "Export cluster labels". Below these, the "Cluster details" section shows "Selected cluster: 0", "Number of samples: 20", and "Percentage of dataset: 50 %". The "Navigation" section includes two rows of controls: "Select a cluster" with a slider set to 0 and "Prev" and "Next" buttons, and "Browse through the cluster" with a slider set to 0 and "Prev" and "Next" buttons. The "Image gallery" section at the bottom displays three images: a cup of coffee on a wooden table, a hand pouring coffee from a metal pitcher into a mug on a wooden deck, and a close-up of coffee beans.

### Text clustering

A text clustering task groups text into clusters where a particular cluster contains a similar text.

Cluster details

**Selected cluster:** 0  
**Number of samples:** 104  
**Percentage of dataset:** 58 %

Navigation

Select a cluster

0

Prev ⌂

Next ⌂

Text clustering results

stars	comment	cluster_label
1.0 out of 5 stars	I had really high hopes	0
1.0 out of 5 stars	I have had my Charge 4	0
1.0 out of 5 stars	I ordered this Contact 1	0
3.0 out of 5 stars	After my Charge 2 wen	0
3.0 out of 5 stars	Most of the advertised	0
5.0 out of 5 stars	I didn't buy the Charge	0
3.0 out of 5 stars	Good:-Tracks steps an	0
1.0 out of 5 stars	The battery doesn't las	0
4.0 out of 5 stars	I had a Charge 3 that a	0

1 to 10 of 104 << < Page 1 of 11 > >>

## Create a clustering task

### Overview

H2O Label Genie enables you to explore an image or text dataset through a clustering task. A clustering task refers to finding and exploring groups in a dataset.

**Note:** To learn about supported clustering tasks, see [Supported clustering tasks](#).

### Instructions

To explore a dataset through a *clustering task*, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Data exploration**.
2. Click **New clustering task**.
3. In the **Task name** box, enter a name for the clustering task.
4. In the **Task description** box, enter a description for the clustering task.
5. In the **Select dataset** list, select an image or text dataset (the dataset you want to explore).
  - If the data type of the selected dataset is *text*, proceed with the following instructions:
    1. In the **Select text column** list, select the text column in the dataset (data).
6. In the **Number of clusters** box, enter the number of clusters to be used by the clustering algorithm.
7. In the **Type** list, select a clustering algorithm for the clustering task. **Note:** H2O Label Genie supports Gaussian mixture and K-means clustering for image and text datasets. The clustering is performed on the data embeddings generated with the OpenCLIP learning model. OpenCLIP is an adaptation of OpenAI's Contrastive Language-Image Pre-training (CLIP). To learn more about OpenCLIP, see [OpenCLIP](#).
8. Click **Start clustering**. **Note:**
  - Several tabs appear when viewing or right after creating a clustering task. To learn more, see [Clustering task tabs](#).
  - For a clustering task, you can download the cluster labels in the **Gallery** and the **Map** tab once clustering results are available. To learn more, see [Download a clustering task's cluster labels](#).

The screenshot shows the 'New clustering task' dialog box. The fields are as follows:

- Task name \*:** slick-tuna
- Task description:** (empty)
- Select dataset \*:** (dropdown menu)
- Number of clusters \*:** 2
- Clustering model:** Gaussian-Mixture
- Embedding model:** ViT-B-32-quickgelu
- Dimension reduction:** PCA
- Start clustering:** (yellow button)

## View a clustering task

### Overview

Viewing a clustering task lets you review the clustering task's gallery and map.

### Instructions

To view a clustering task, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Data exploration**.
2. In the *clustering tasks* table, double-click the row where the clustering task you want to view is located. **Note:** When viewing a particular clustering task, several tabs appear containing information about the task. To learn more, see [Clustering task tabs](#).

## Clustering task tabs

### Overview

Several tabs appear when viewing or right after creating a clustering task. Such tabs are as follows:

- [Gallery](#)
- [Map](#)

Each tab above, in turn, is explained below.

#### Note:

- You can download the cluster labels in the **Gallery** and in the **Map** tab once clustering results are available. To learn more, see [Download a clustering task's cluster labels](#).
- To learn how to create a clustering task, see [Create a clustering task](#).
- To learn how to view a clustering task, see [View a clustering task](#).

### Gallery

The **Gallery** tab contains details about the task (for example, task type). In the tab, you can explore the clustering results by browsing through different clusters or samples of an individual cluster.

The screenshot shows the 'Gallery' tab interface. At the top, there are two tabs: 'Edit cluster task' and 'Export cluster labels'. Below the tabs, the 'Cluster details' section displays: 'Selected cluster: 0', 'Number of samples: 20', and 'Percentage of dataset: 50 %'. The 'Navigation' section includes two rows of controls. The first row has a 'Select a cluster' dropdown menu with a radio button, a 'Prev' button with a left arrow, and a 'Next' button with a right arrow. The second row has a 'Browse through the cluster' dropdown menu with a radio button, a 'Prev' button with a left arrow, and a 'Next' button with a right arrow. The 'Image gallery' section at the bottom displays three images: a cup of coffee on a wooden table, a hand pouring coffee from a metal pitcher into a cup on a wooden table, and a close-up of coffee beans.

Cluster details

Selected cluster: 0  
 Number of samples: 104  
 Percentage of dataset: 58 %

Navigation

Select a cluster 0

Prev ⌂
Next ⌂

Text clustering results

stars	comment	cluster_label
1.0 out of 5 stars	I had really high hopes	0
1.0 out of 5 stars	I have had my Charge 1	0
1.0 out of 5 stars	I ordered this Contact 1	0
3.0 out of 5 stars	After my Charge 2 wen	0
3.0 out of 5 stars	Most of the advertised	0
5.0 out of 5 stars	I didn't buy the Charge	0
3.0 out of 5 stars	Good:-Tracks steps an	0
1.0 out of 5 stars	The battery doesn't las	0
4.0 out of 5 stars	I had a Charge 3 that a	0

1 to 10 of 104
⏪ ⏩
Page 1 of 11
⏪ ⏩

## Map

The **Map** tab contains details about the task (for example, task type). In the tab, you can explore the clustering task's embeddings in 2D and 3D (right after H2O Label Genie reduces the dimensionality of the data). Observing the embeddings in 2D and 3D can help you understand the data structure.



## Download a clustering task's cluster labels

### Overview

Once clustering results are available, you can download a clustering task's cluster labels. H2O Label Genie downloads a **CSV** file containing the clustering task's cluster labels. For example, the **CSV** file can be utilized to:

- Annotate a cluster in the dataset separately
- Remove a cluster completely
- Train an [H2O Hydrogen Torch](#) model

**Note:** To learn more about a clustering task's tabs, see [Clustering task tabs](#).

### Instructions

To download a clustering task's cluster labels (once clustering results are available), consider the following instructions:

1. In the **H2O Label Genie** navigation menu, click **Data exploration**.
2. In the *clustering tasks* table, double-click the row where the clustering task you want to download its clustering labels is located.
3. Click **Export cluster labels**. **Note:** You can also download the cluster labels in the **Map** tab.
  1. Click the **Map** tab.
  2. Click **Export cluster labels**.

## Delete a clustering task

### Overview

When a clustering task is no longer needed, you can delete it. A deleted clustering task removes clustering results permanently from the H2O Label Genie instance. Deleting a clustering task does not delete the data embeddings H2O Label Genie used to create the clustering task.

### Instructions

To delete one or multiple clustering tasks, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Data exploration**.
2. Click **Delete clustering task(s)**.
3. Select the clustering task(s) you want to delete.
4. Click **Delete clustering task(s)**.

## Supported annotation tasks

### Overview

H2O Label Genie supports various [annotation tasks](#) in computer vision (CV) (image), natural language processing (NLP) (text), and audio.

**Note:** To learn how to create an annotation task, see [Create an annotation task](#).

### Text annotation tasks

#### Text classification

- **Description:** A text classification annotation task specifies one or more categorical target labels for each input text.

**Note:**

- H2O Label Genie supports multi-label text classification annotation tasks.
- To accelerate the labeling process for text classification annotation tasks, H2O Label Genie offers the ability to utilize a zero-shot learning model. To learn more, see [Zero-shot learning models](#).

#### Text regression

- **Description:** A text regression annotation task specifies one continuous target label for each input text.

#### Text-entity recognition

- **Description:** A text-entity recognition annotation task specifies one or more defined entities for each unstructured input text.

#### Text summarization

- **Description:** A text summarization annotation task specifies a summary for each input text.

**Note:** To accelerate the labeling process, H2O Label Genie offers the ability to utilize pretrained text summarization models. To learn more, see [Zero-shot learning models](#).

#### Text-generative AI

- **Description:** A text-generative AI annotation task specifies a particular text response for the text column(s).

**Note: Zero-shot learning models:** To accelerate the labeling process for a text-generative AI annotation tasks, H2O Label Genie offers the ability to utilize several zero-shot learning models. To learn more, see [Zero-shot learning models: Text-generative AI](#)

## Image annotation tasks

### Image classification

- **Description:** An image classification annotation task specifies one or more categorical target labels for each input image.

**Note:**

- H2O Label Genie supports multi-label image classification annotation tasks.
- To accelerate the labeling process for image classification annotation tasks, H2O Label Genie offers the ability to utilize a zero-shot learning model. To learn more, see [Zero-shot learning models](#).

### Image regression

- **Description:** An image regression annotation task specifies one continuous target label for each input image.

## Object detection

- **Description:** An object detection annotation task specifies one or more object classes (labels) for each input image.

**Note:** To accelerate the labeling process for object detection annotation tasks, H2O Label Genie offers the ability to utilize a zero-shot learning model. To learn more, see [Zero-shot learning models](#).

## Image instance segmentation

In H2O Label Genie, for an image instance segmentation annotation task, you need to assign one or more object classes (labels) to each input image.

**Note:** To accelerate the labeling process for image instance segmentation annotation tasks, H2O Label Genie offers the ability to utilize a zero-shot learning model. To learn more, see [Zero-shot learning models](#).

## Audio annotation tasks

### Audio classification

- **Description:** An audio classification annotation task specifies one or more categorical target labels for each input audio.

**Note:** H2O Label Genie supports multi-label audio classification annotation tasks.

### Audio regression

- **Description:** An audio regression annotation task specifies one continuous target label for each input audio.

# Hotkeys

## Overview

H2O Label Genie supports several hotkeys (keyboard shortcuts) designed to speed up a dataset's annotation (labeling) process.

## Instructions

The below table describes the supported hotkeys.

Hotkey	Description
s	Activate selection tool
r	Select rectangle tool
p	Select polygon tool
b	Toggle drawing function
l	Toggle label
Shift + Click	Select/deselect multiple shapes when in selection mode
a	Select all shapes
Mouse wheel	Zoom in/out
Arrow keys (↑ ↓ → ←)	Move selected shapes by 1px (or 10px while holding the Shift key)
c	Copy selected shape(s)
v	Paste selected shape(s)
Delete	Delete selected shapes
Backspace	Delete last polygon vertex
Backspace	Delete selected shapes (if not drawing a polygon)
Esc	Cancel the creation process of a "New annotation task"
Enter	Finish drawing polyshape
j	Move to the previous sample/cluster
k	Move to the next sample/cluster
u	Approve and move to the next sample/cluster
n	Move to the previous inner cluster (a sample within a cluster)
m	Move to the next inner cluster (a sample within a cluster)

# Zero-shot learning models

## Overview

H2O Label Genie enables you to utilize zero-shot learning models to accelerate the labeling process. In particular, H2O Label Genie lets you use a zero-shot learning model for the following supported annotation tasks:

- [Image classification](#)
- [Object detection](#)
- [Image instance segmentation](#)
- [Text classification](#)
- [Text summarization](#)
- [Text-generative AI](#)

## What are zero-shot learning models?

Labeled data is crucial for supervised learning problem types in computer vision (CV), natural language processing (NLP), and audio. High-quality labeled data usually requires a lot of manual labeling that can lead to high costs and delay production or execution.

One way to accelerate the labeling process is to utilize zero-shot learning models. These models let data scientists label unlabeled data with high accuracy and speed. Zero-shot learning models are pre-trained models that have been trained on vast and distinct classes. As a result, zero-shot learning models with prior knowledge can label unlabeled data.

## What are zero-shot learning model predictions?

The labels or suggested labels for a given sample that are provided by a zero-shot learning model are called *zero-shot learning model predictions*. For example, for an image and text classification annotation task, H2O Label Genie, with a zero-shot learning model activated, offers a percentage probability of an image or text belonging to a certain label (class). For an object detection annotation task, it populates an image with bounding boxes where the bounding boxes capture the desired objects (for example, a car).

## Annotation tasks + zero-shot learning models

### Image classification

By default, H2O Label Genie utilizes the **OpenCLIP** zero-shot learning model for image classification annotation tasks. OpenCLIP is an adaptation of OpenAI's Contrastive Language-Image Pre-training (CLIP).

- To learn more, see [Open-clip](#)

### Object detection

By default, H2O Label Genie utilizes the **Detic** zero-shot learning model for object detection annotation tasks.

- To learn more, see [Detecting Twenty-thousand Classes using Image-level Supervision](#)

### Image instance segmentation

By default, H2O Label Genie utilizes the **Detic** zero-shot learning model for image instance segmentation annotation tasks.

- To learn more, see [Detecting Twenty-thousand Classes using Image-level Supervision](#)

### Text classification

By default, H2O Label Genie utilizes the **bart-large-mnli** zero-shot learning model for text classification annotation tasks.

- To learn more, see [Bart-large-mnli](#)

### Text summarization

H2O Label Genie allows you to utilize the following zero-shot learning models for text summarization annotation tasks:

- **Bart-large-cnn**
  - To learn more, see [Bart-large-cnn](#)

- **Distilbart-cnn-12-6**
  - To learn more, see [Distilbart-cnn-12-6](#)
- **Pegasus-large**
  - To learn more, see [Pegasus-large](#)

**Note:**

- **Select a particular model:** To learn how to select a particular zero-shot learning model for a text summarization annotation task, see [Select a zero-shot learning model](#)

**Text-generative AI**

H2O Label Genie allows you to utilize zero-shot learning models (LLMs) from h2oGPTe and OpenAI:

- **h2oGPTe:** The available h2oGPTe LLMs depend on the available version of h2oGPTe connected within the H2O Label Genie environment.  
**Note:** H2O Label Genie v1.0.0 can only access h2oGPTe v1.5.26 or higher.
- **OpenAI:** The available OpenAI LLMs depend on the type of your OpenAI API key. For example, free-tier or trial users might have limited access, often to GPT-3.5-turbo.

**Note: Select a particular model:** To learn how to select a particular LLM for a text-generative AI annotation task, see [Select a zero-shot learning model](#).

**Select a zero-shot learning model****Text summarization**

You can select the zero-shot learning model for a [Text summarization](#) annotation task.

**Caution:** The below instructions assume you have already created a text summarization annotation task. To learn how to create an annotation task, see [Create an annotation task](#).

1. In the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the **Annotation tasks** table, click the name of the annotation task you want to define a zero-shot learning model for.
3. Click the **Rubric** tab.
4. In the **Select model** list, select a zero-shot learning model. **Note:** To learn about available models, see [Text summarization](#)

**Text-generative AI**

You can select a zero-shot learning model (LLM) for a [Text-generative AI](#) annotation task.

**Caution:** The below instructions assume you have already created a text-generative AI annotation task. To learn how to create an annotation task, see [Create an annotation task](#).

1. In the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the **Annotation tasks** table, click the name of the annotation task you want to define an LLM for.
3. Click the **Rubric** tab.
4. In the **Select API endpoint type** list, select an LLM family (**h2oGPT** or **OpenAI**). **Note:** To utilize an OpenAI model, you need to define the OpenAI API settings linking to your OpenAI account. To learn more, see [OpenAI API settings](#).
5. Click edit **Edit**.
6. In the **Model name for large language model** list, select (or enter) a model name for the annotation task.  
**Note:** The available LLMs (zero-shot learning models) are based on the selected LLM family in the following setting: **Select API endpoint type**.

**Note:** To learn about the LLM parameter and prompt template settings, see [Large language model \(LLM\) parameters](#).

# Create an annotation task

## Overview

You need to create an [annotation task](#) to annotate (label) a dataset.

An annotation task refers to labeling data in a manner that makes them suitable to support an array of deep-learning problem types. This process can, for example, involve adding bounding boxes to images where you can attribute labels to each box.

### Note:

- To learn how to download your annotated dataset, see [Download an annotated dataset](#).
- To learn about supported annotation tasks, see [Supported annotation tasks](#).

## Instructions

To create an *annotation task*, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. Click **New annotation task**.
3. In the **Task name** box, enter a name for the annotation task.
4. In the **Task description** box, enter a description for the annotation task.
5. In the **Select dataset** list, select a dataset (the dataset you want to label).
  - If the data type of the selected dataset is *text*, proceed with the following instructions:
    1. In the **Select task** list, select the type of annotation you want to create (for example, classification, entity recognition, etc.).
    2. In the **Select text column** list, select the text column in your dataset (data).
  - If the data type of the selected dataset is *image* or *audio*, proceed with the following instructions:
    1. In the **Select task** list, select the type of annotation you want to create (for example, classification, regression, etc.).
6. Click **Create task**. **Note:** After creating a new annotation task, specify an annotation task rubric in the **Rubric** tab. An [annotation task rubric](#) refers to the labels (for example, object classes) to use when annotating a dataset. For example, after creating a new annotation task for an object detection dataset, you have to specify the object classes to use when labeling the dataset in the annotation task rubric. To learn more, see [Specify an annotation task rubric](#).

The screenshot shows a dark-themed modal window titled "New annotation task" with a close button (X) in the top right corner. The form contains the following fields:

- Task name \***: A text input field containing "boisterous-crab".
- Task description**: A large text area that is currently empty.
- Select dataset \***: A dropdown menu showing "amnist-demo" with a downward arrow.
- Select task \***: A dropdown menu that is currently empty.
- Create task**: A yellow button at the bottom of the form.

## Access an annotation task's tabs

### Overview

Several tabs appear when viewing or right after creating an annotation task. Such tabs are as follows:

- [Dashboard](#)
- [Rubric](#)
- [Statistics](#)
- [Annotate](#)
- [Review](#)
- [Import](#)
- [Export](#)

Each tab above, in turn, is explained below.

#### Note:

- To learn how to create an annotation task, see [Create an annotation task](#).
- To learn how to view an annotation task, see [View an annotation task](#).

### Instructions

To access the tabs of an annotation task, consider the following instructions:

1. In the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the **Annotation tasks** table, click the name of the annotation task you want to view its tabs.

### Tabs

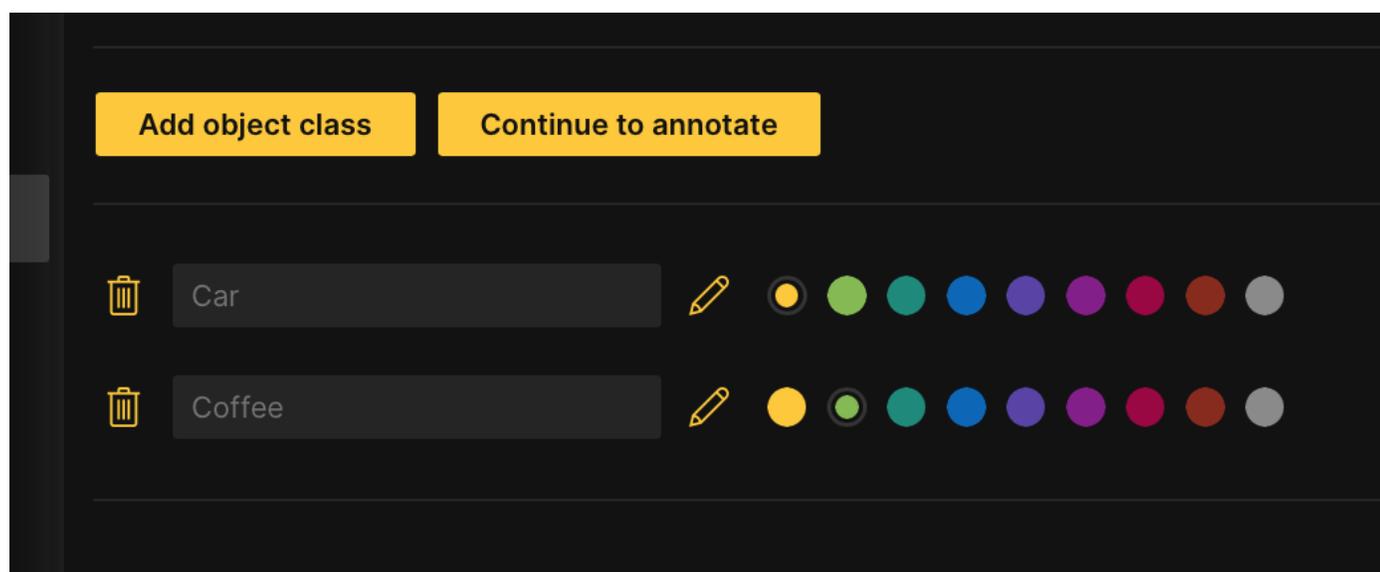
#### Dashboard

The **Dashboard** tab contains details about the task (for example, task type) while highlighting the annotation status through a table.

**Note:** In the **Dashboard** tab, you can download the annotated dataset at any point (progress). To learn more, see [Download an annotated dataset](#).

#### Rubric

The **Rubric** tab contains the specified [annotation task rubric](#) for a dataset.



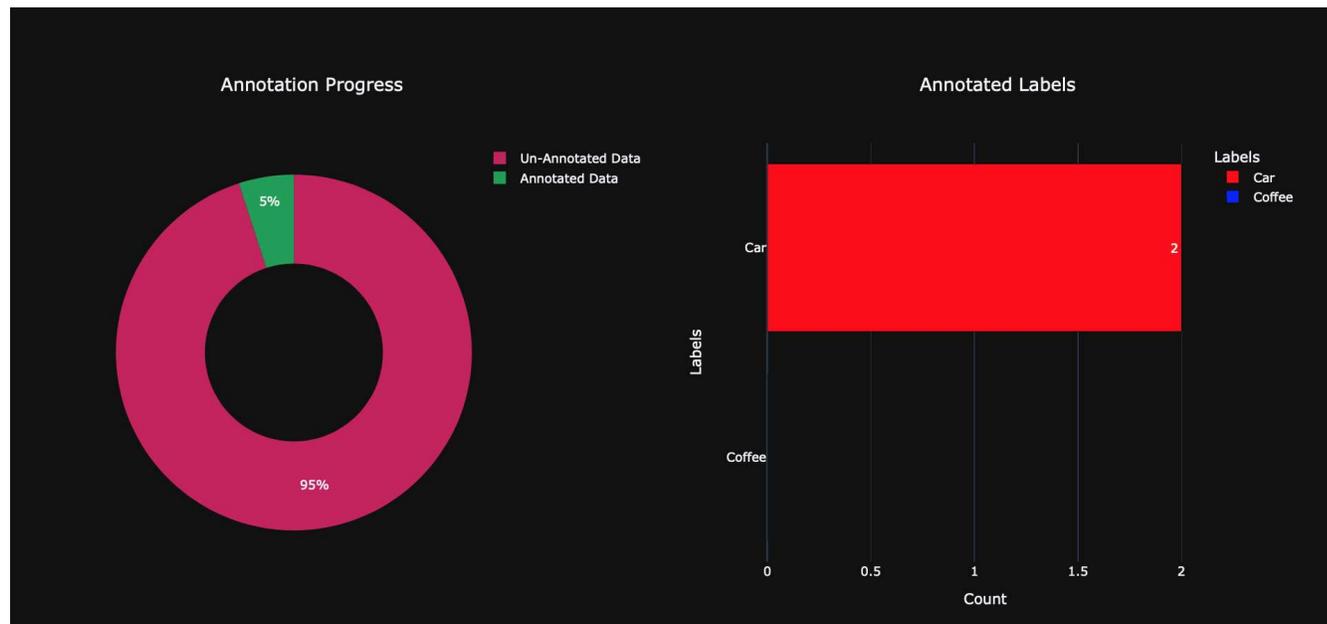
#### Note:

- At any point during the annotation process, you can edit the rubric of the annotation task (for example, add one more object class for an object detection annotation task). To learn how to edit the annotation task rubric, see [Edit an annotation task rubric](#).
- In the **Rubric** tab, you can upload a **PDF** file containing annotation instructions. To disable the option to upload a **PDF**, see [Appearance settings: Upload annotation instructions](#).

## Statistics

- **Annotation task statistics**

By default, H2O Label Genie displays graphical statistics about the annotation task in the **Statistics** tab.



- **An annotation task's zero-shot statistics (about its predictions)**

**Caution:** The time it takes H2O Label Genie to generate zero-shot predictions depends on the computational resources of the instance.

For annotation tasks that support a zero-shot learning model, in the **Statistics** tab, you can observe graphical statistics about the generated zero-shot predictions. To view the graphical statistics about generated zero-shot predictions, consider the following instructions:

1. In the **Statistics** tab, click the **Show zero-shot statistics** toggle.

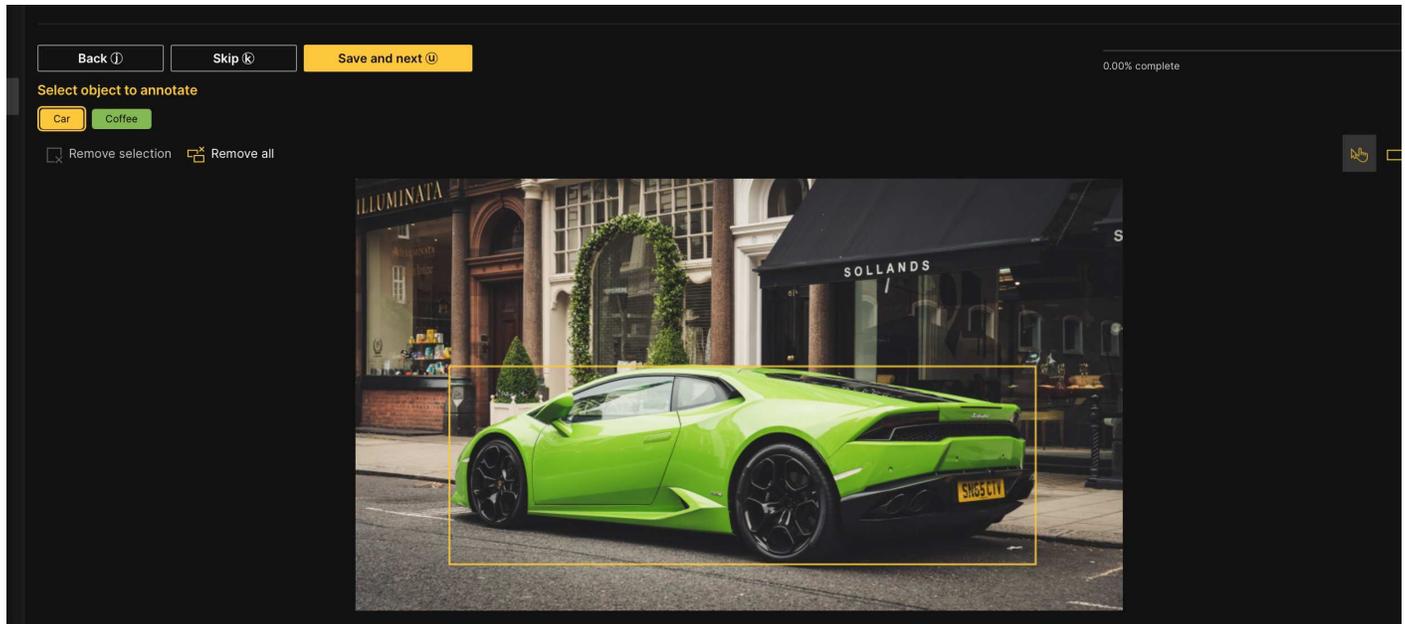


## Annotate

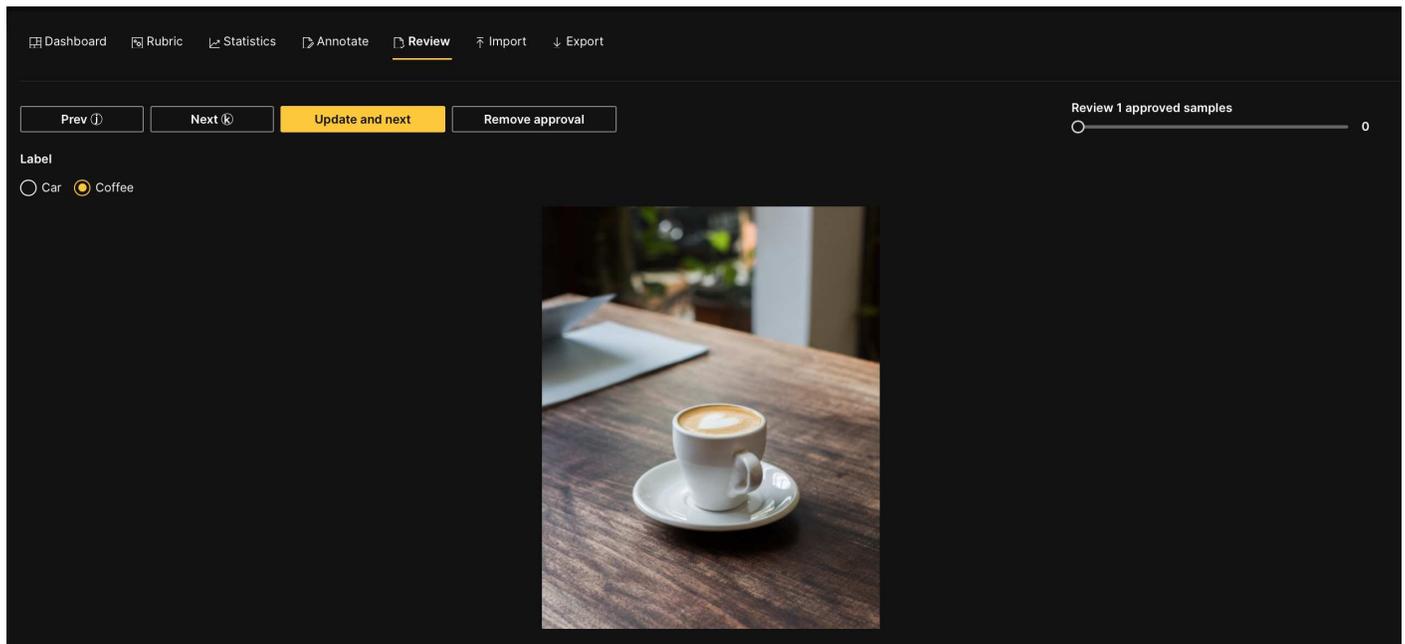
In the **Annotate** tab, you can annotate your dataset right after specifying your annotation task rubric.

**Note:** Attention

The **Annotate** tab will not be available until you specify an annotation task rubric in the **Rubric** tab.

**Review**

In the **Review** tab, you can review already approved samples. You can also remove an approved sample's approval.

**Import**

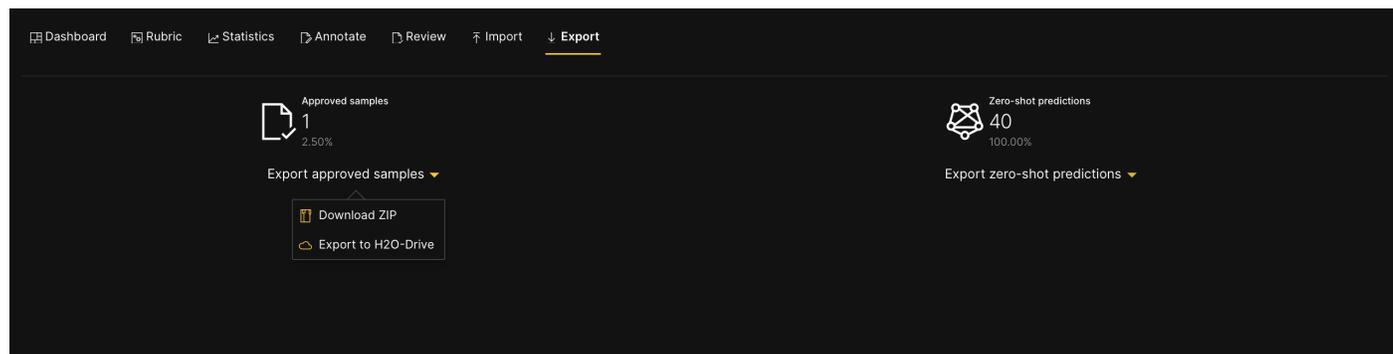
In the **Import** tab, you can import existing annotations. To learn more, see [Import an annotated dataset](#).

**Export**

In the **Export** tab, you can:

- Download or export to H2O Drive already approved samples
  - To learn how to download already approved samples, see [Download an annotated dataset \(approved samples\)](#).

- To learn how to export already approved samples to H2O Drive, see [Export an annotated dataset \(approved samples\) to H2O Drive](#).
- Download or export to H2O Drive generated zero-shot predictions (for certain annotation tasks (datasets)) **Caution:** The time it takes H2O Label Genie to generate zero-shot predictions depends on the computational resources of the instance.
  - To learn how to download an annotation task's zero-shot predictions, see [Download a dataset's zero-shot predictions](#).
  - To learn how to export an annotation task's zero-shot predictions to H2O Drive, see [Export a dataset's zero-shot predictions to H2O Drive](#).



## Specify an annotation task rubric

### Overview

After creating a new annotation task, specify an annotation task rubric in the **Rubric** tab. An **annotation task rubric** refers to the labels (for example, object classes) to use when annotating a dataset. For example, after creating a new annotation task for an object detection dataset, you have to specify the object classes to use when labeling the dataset in the annotation task rubric.

### Instructions

An annotation task rubric differs based on the specified task type of the dataset used to create the annotation task.

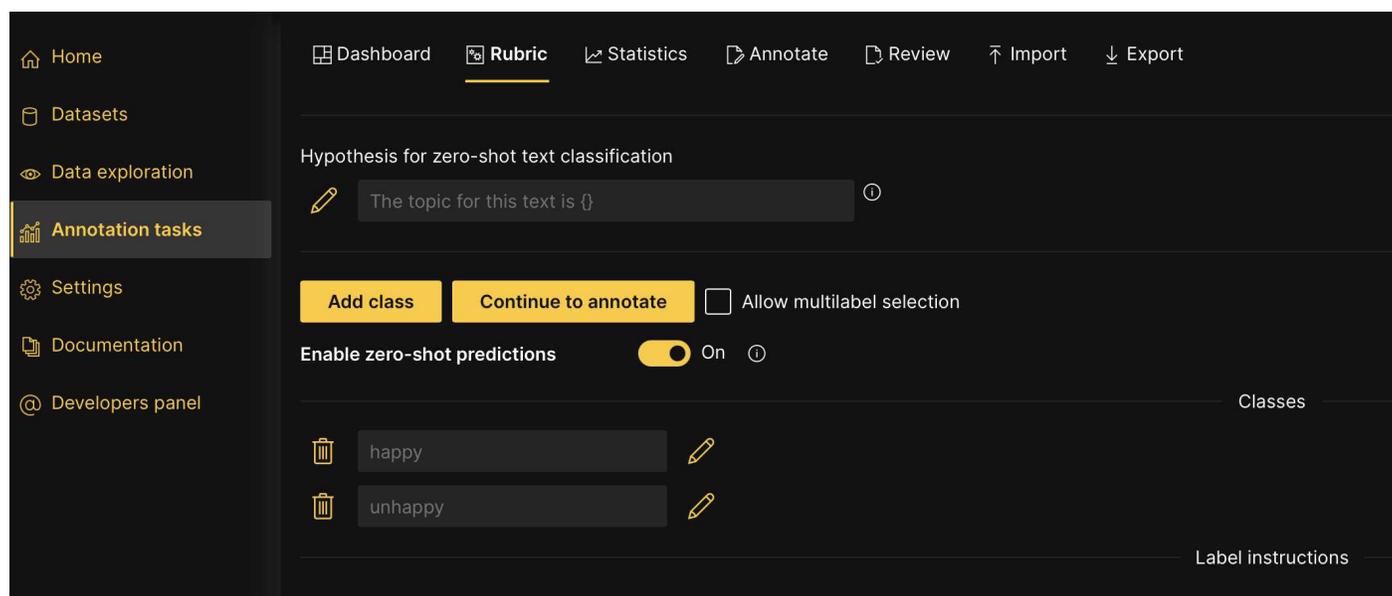
#### Text annotation tasks

##### Text classification

- **Instructions:** Specify one or more categorical target labels for a text classification task rubric.
- **Example:** To specify **happy** and **unhappy** as labels, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **New class name** box, enter **happy**.
2. Click **Add**.
3. Click **Add class**.
4. In the **New class name** box, enter **unhappy**.
5. Click **Add**.



**Note:** To learn more, see [Tutorial 1A: Text classification annotation task](#).

##### Text regression

- **Instructions:** You need to specify one continuous target label for a text regression task rubric.
- **Example:** To create one continuous target label from **1** to **5**, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **Data minimum value** box, enter **1**.
  - The Data minimum value value refers to the minimum value in your continuous values (star ratings from 1 to 5)
2. In the Data maximum value box, enter **5**.
  - The Data maximum value value refers to the maximum value in your continuous values (star ratings from 1 to 5)

3. In the **Data step size (interval)** box, enter 1.

- The Data step size (interval) value refers to the value the label range slider interval takes (the slider is used in the next step to label a review)

4. Click **Apply**.

**Note:** To learn more, see [Tutorial 2A: Text regression annotation task](#).

### Text-entity recognition

- **Instructions:** You need to specify one or more defined entities.
- **Example:** To create a `product` and `emotion` entity, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **New object name** box, enter `product`.
2. Click **Add**.
3. Click **Add entity**.
4. In the **New object name** box, enter `emotion`.

**Note:** To learn more, see [Tutorial 3A: Text-entity recognition annotation task](#).

### Text summarization

- **Instructions:** You need to specify a zero-shot learning model and a maximum target length.
- **Example:** To specify a zero-shot learning model and a maximum target length, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **Select model** box, select `sshleifer/distilbart-cnn-12-6`.
- The **Select model** value refers to the zero-shot learning model to utilize in your annotation task. To learn more, see [Annotation tasks + zero-shot learning models: Text summarization](#)

2. In the **Max target length** box, enter 128.

- The **Max target length** value refers to the minimum character length of your summaries

**Note:** To learn more, see [Tutorial 4A: Text summarization annotation task](#).

### Text-generative AI

- **Instructions:** You need to:
  - Specify a zero-shot learning model (LLM) and its parameters.
  - To learn more, see [Large language model \(LLM\) parameters](#)
  - Define a prompt template (that is the input for the LLM).
  - To learn more, see [Prompt template](#)
  - **Example:** To specify a zero-shot learning model and define a prompt template, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **Select API endpoint type** list, select **h2oGPT**.

- The **Select API endpoint type** value refers to the zero-shot learning model to utilize in the annotation task. To learn more, see [Zero-shot learning models: Text-generative AI](#)

2. Define the large language model (LLM) parameters.

- To learn more about each parameter, see [Large language model \(LLM\) parameters](#)

3. In the **Select example prompt** list, select **summarize**.

**Note:** To learn more, see [Tutorial 5A: Text-generative AI annotation task](#).

## Image annotation tasks

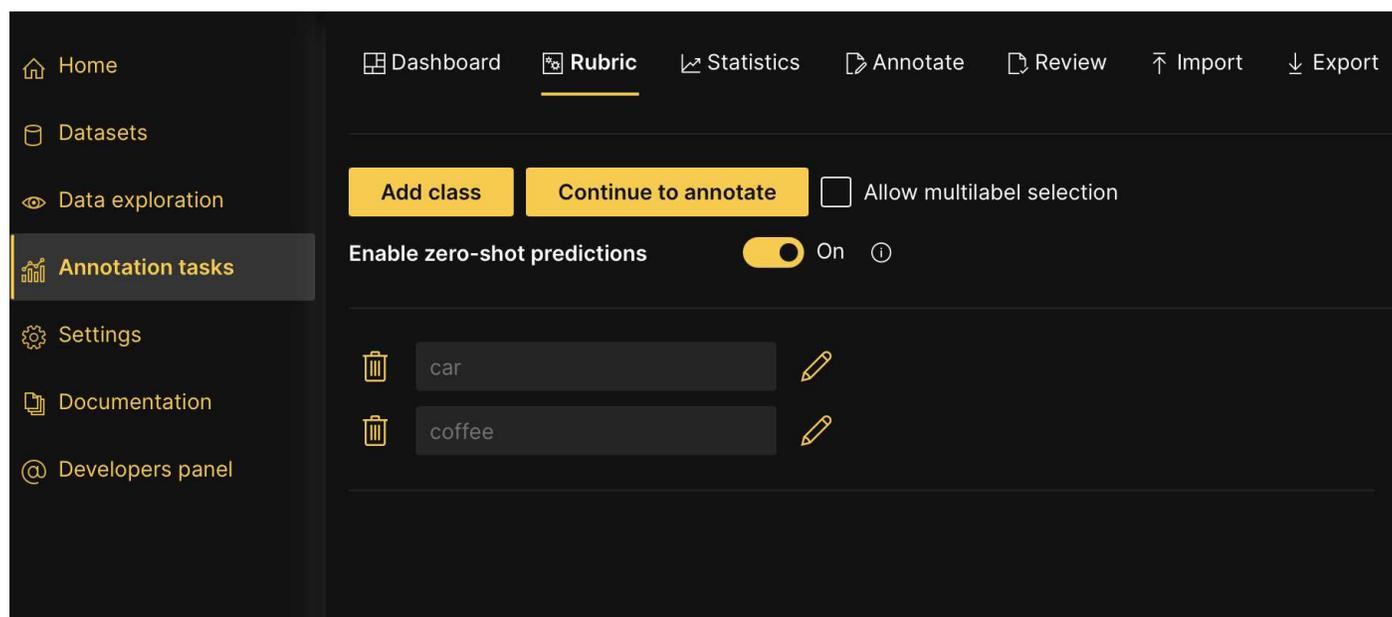
### Image classification

For an image classification task rubric, you need specify one or more categorical target labels in the annotation task rubric for an image classification annotation task. To learn more, see [Tutorial 1B: Annotation task: Image classification](#)

- **Instructions:** Specify one or more categorical target labels.
- **Example:** To create a `car` and `coffee` label, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **New class name** box, enter `car`.
2. Click **Add**.
3. Click **Add class**.
4. In the **New class name** box, enter `coffee`.
5. Click **Add**.



**Note:** To learn more, see [Tutorial 1B: Image classification annotation task](#).

### Image regression

- **Instructions:** Specify one continuous target label.
- **Example:** To create one continuous target label from `0` to `9`, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **Data minimum value** box, enter `0`.
  - The **Data minimum value** refers to the minimum value in your continuous values (digits ranging between 0 to 9)
2. In the **Data maximum value** box, enter `9`.
  - The **Data maximum value** refers to the maximum value in your continuous values (digits ranging between 0 to 9)
3. In the **Data step size (interval)** box, enter `1`.
  - The **Data step size (interval)** value refers to the value the label range slider interval takes
4. Click **Apply**.

**Set regression range** ✕

Data Minimum Value: 0      Data Maximum Value: 9      Data Step Size (Interval): 1

**Apply**      **Cancel**

**Note:** To learn more, see [Tutorial 2B: Image regression annotation task](#).

### Object detection

- **Instructions:** Specify one or more object classes (labels).
- **Example:** To specify `car` and `coffee` as a label, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **New object name** box, enter `car`.
2. Click **Add**.
3. Click **Add object class**.
4. In the **New object name** box, enter `coffee`.
5. Click **Add**.

Home | Dashboard | **Rubric** | Statistics | Annotate | Review | Import | Export

Annotations tasks

**Add object class**      **Continue to annotate**

Enable zero-shot predictions  On ⓘ

Objects

- car [color selection]
- coffee [color selection]

Label instructions

**Note:** To learn more, see [Tutorial 3B: Object detection annotation task](#).

### Image instance segmentation

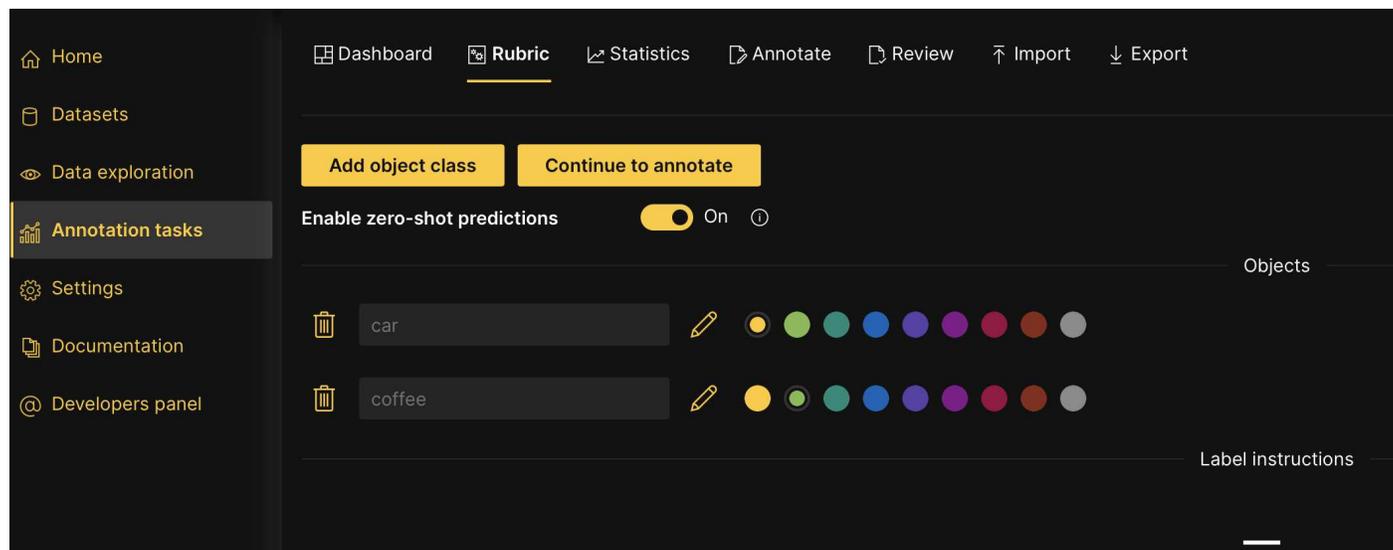
You need to specify one or more object classes (labels) in the annotation task rubric for an image instance segmentation annotation task. To learn more, see [Tutorial 4B: Annotation task: Image instance segmentation](#).

- **Instructions:** Specify one or more object classes (labels).
- **Example:** To specify `car` and `coffee` as object classes, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **New object name** box, enter `car`.
2. Click **Add**.

3. Click **Add object class**.
4. In the **New object name** box, enter `coffee`.
5. Click **Add**.



**Note:** To learn more, see [Tutorial 4B: Image instance segmentation annotation task](#).

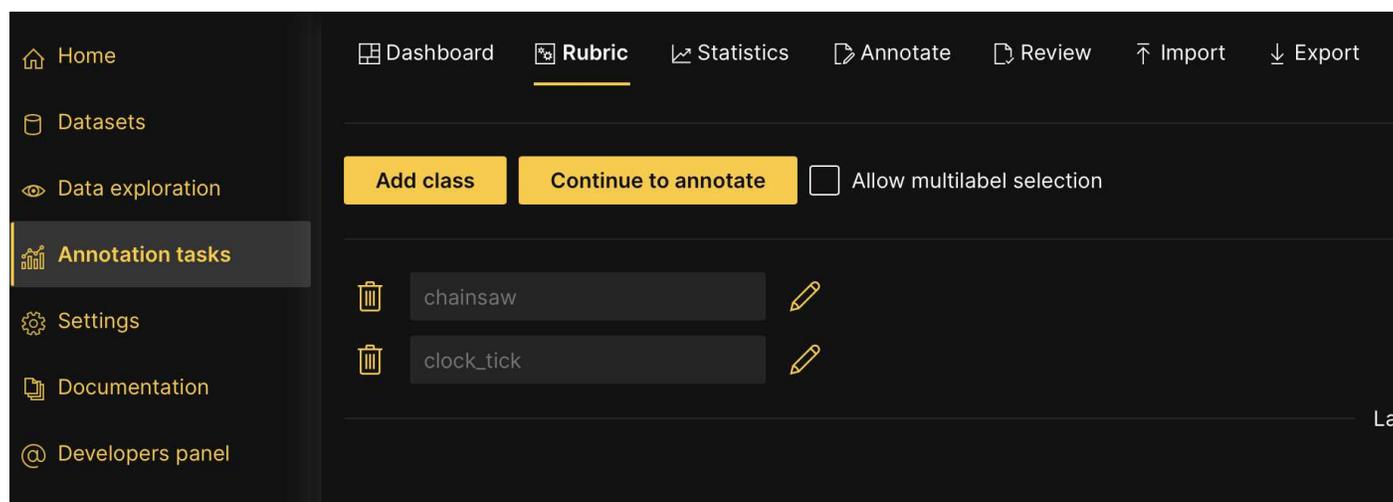
## Audio annotation tasks

### Audio classification

- **Instructions:** Specify one or more categorical target labels.
- **Example:** To specify `chainsaw` and `clock_tick` as labels, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **New class name** box, enter `chainsaw`.
2. Click **Add**.
3. Click **Add class**.
4. In the **New class name** box, enter `clock_tick`.
5. Click **Add**.



**Note:** To learn more, see [Tutorial 1C: Audio classification annotation task](#).

### Audio regression

- **Instructions:** Specify one continuous target label.

- **Example:** To create one continuous target label from **0** to **9**, one can consider the following instructions in the **Rubric** tab of the annotation task:

**Note:** To learn how to access the **Rubric** tab of an annotation task (or other tabs), see [Access an annotation task's tabs](#).

1. In the **Data minimum value** box, enter 0.
  - The **Data minimum value** refers to the minimum value in your continuous values (in this case, digits ranging from 0 to 9)
2. In the **Data maximum value** box, enter 9.
  - The **Data maximum value** refers to the maximum value in your continuous values (in this case, digits ranging from 0 to 9)
3. In the **Data step size (interval)** box, enter 1.
  - The **Data step size (interval)** value refers to the value the label range slider interval takes
4. Click **Apply**.

The screenshot shows a dark-themed dialog box titled "Set regression range" with a close button (X) in the top right corner. It features three input fields: "Data Minimum Value" containing "0", "Data Maximum Value" containing "9", and "Data Step Size (Interval)" containing "1". The "Data Step Size (Interval)" field is highlighted with a yellow border. Below the input fields are two buttons: a yellow "Apply" button and a grey "Cancel" button.

**Note:** To learn more, see [Tutorial 2C: Audio regression annotation task](#).

## Large language model (LLM) parameters

### Select API endpoint type

This LLM parameter defines the zero-shot learning model family (LLM) to utilize for the text-generative AI annotation task.

#### Options:

- **h2oGPT**
  - This option enables h2oGPTe LLMs to be available for a text-generative AI annotation task.
- **OpenAI**
  - This option enables OpenAI LLMs in your OpenAI account (API key) to be available for a text-generative AI annotation task. To connect to your OpenAI LLMs, see [OpenAI API settings](#)

### LLM model name

This LLM parameter defines the zero-shot learning model name (LLM) to utilize for a text-generative AI annotation task.

### Max response tokens

This LLM parameter defines the maximum number of tokens for a response; a low number can result in short responses, which might limit the responses.

### Temperature

This LLM parameter defines the randomness of predictions by scaling the logits. Higher temperature values increase creativity on the part of the model while producing more diverse outputs. In other words, the temperature makes the distribution more random.

### **Repetition penalty**

This LLM parameter defines the penalty value of tokens frequently reappearing in the text (response). For example, a token that has already appeared ten times can be penalized more than a token that has appeared only two times. A 1.0 value means no penalty.

**Tip:** This setting can be helpful when attempting to reduce the model's tendency to generate verbatim/identical text.

### **Prompt template**

#### **Select example prompt**

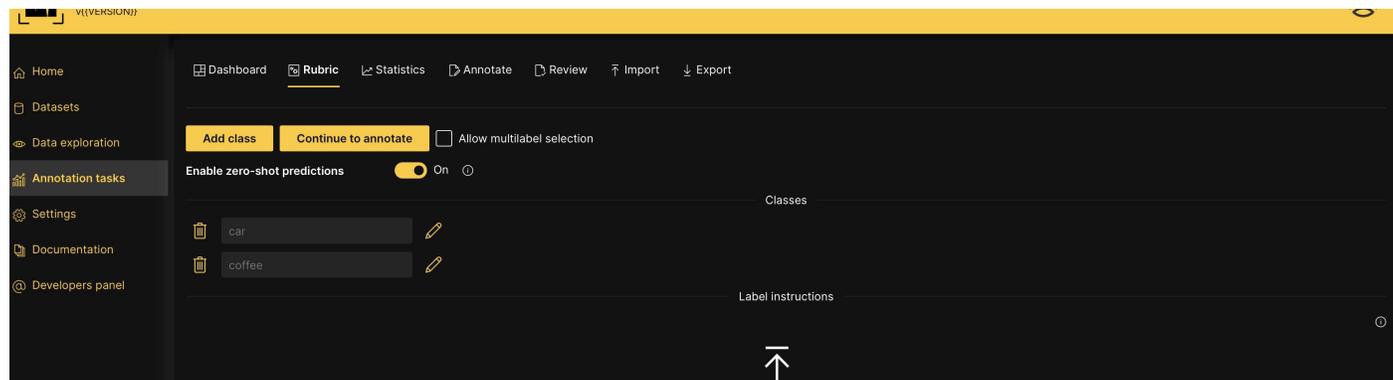
This setting defines the input format for the selected model. There are several options, including the option to create your own custom input format (**custom**).

# Edit an annotation task rubric

## Overview

An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset. You can edit the annotation task rubric at any point during the annotation process. An annotation task rubric becomes available right after creating an annotation task.

**Note:** To learn how to create an annotation task, see [Create an annotation task](#).



## Instructions

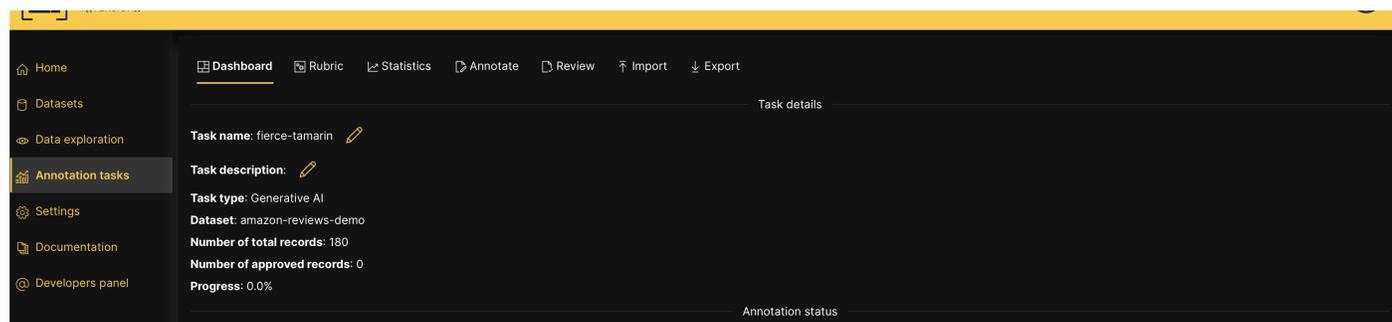
To edit an annotation task rubric, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the *annotation tasks* table, double-click the row where the annotation task you want to access its rubric is located.
3. Click the **Rubric** tab.

# View an annotation task

## Overview

Viewing an annotation task lets you review an annotation task's dashboard, rubric, and statistics.



## Instructions

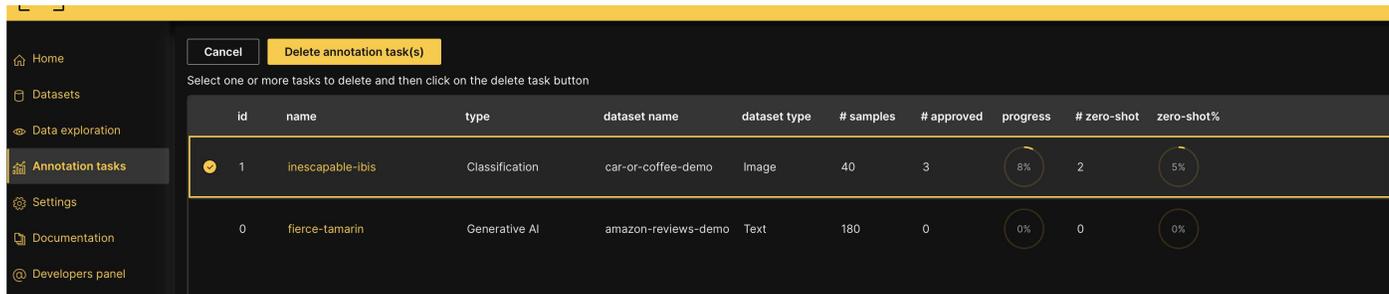
To view an annotation task, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. In the *annotation tasks* table, double-click the row where the annotation task you want to view is located. **Note:** When viewing a particular annotation task, several tabs appear containing information about the task. To learn more, see [Access an annotation task's tabs](#).

# Delete an annotation task

## Overview

When an annotation task is no longer needed, you can delete it. A deleted annotation task is permanently removed from the H2O Label Genie instance.



## Instructions

To delete one or multiple annotation tasks, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Annotation tasks**.
2. Click **Delete annotation task(s)**.
3. Select the annotation task(s) you want to delete.
4. Click **Delete annotation task(s)**.

## Segment objects

### Overview

By default, for an image instance segmentation annotation task, H2O Label Genie offers a **Box to polygon** feature that enables you to improve the efficiency of manual labeling. A **Box to polygon** feature enables you to encapsulate an object class with a rectangular that is immediately replaced with a more complex polygon that captures all the borders of the object class inside the rectangular.

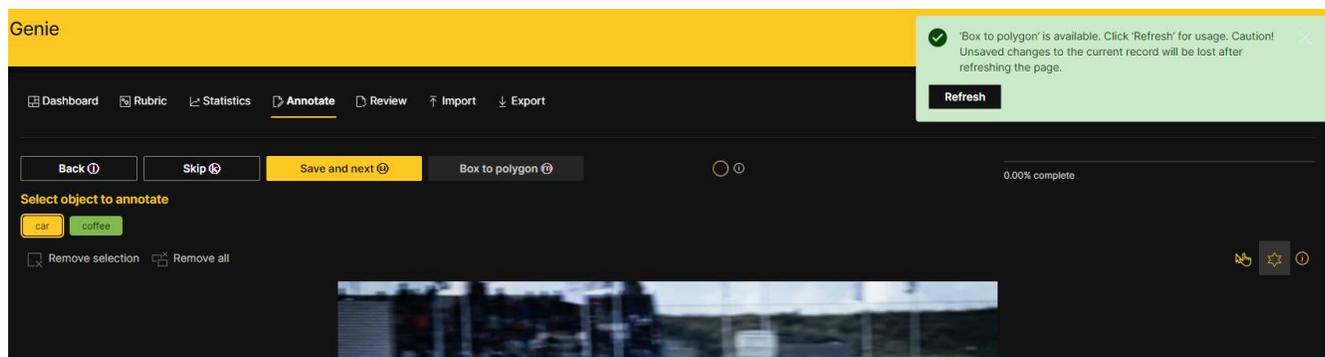


### Instructions

To segment objects during an image instance segmentation annotation task, consider the following instructions:

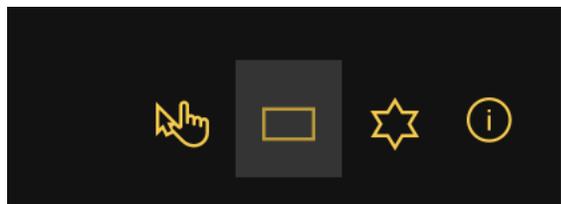
**Caution:** If the **Enable segment anything** setting is turned **Off**, the **Box to polygon** feature is not available during the annotation process. To turn **On** the **Enable segment anything** setting, see [Enable segment anything](#).

1. After moving to the **Annotate** tab, H2O Label Genie initiates a *segment anything* model that enables the **Box to polygon** feature. H2O Label Genie notifies you to **Refresh** the instance when the **Box to polygon** feature becomes available.



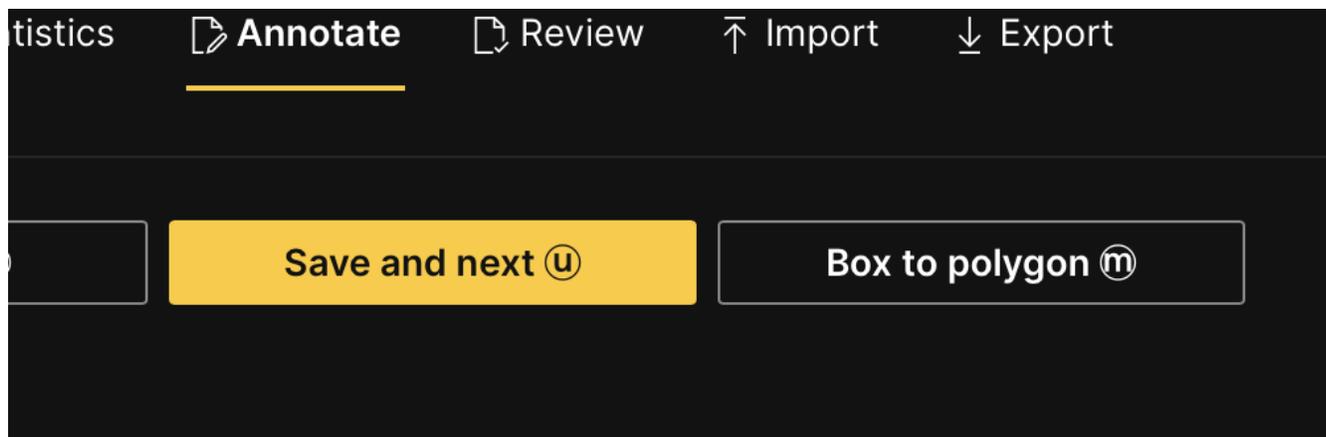
**Note:** To learn about the *segment anything* model, see [Segment anything](#).

2. In the **Annotate** tab, click the **Rectangle** tool.



**Note:** After refreshing the instance, the **Rectangle** tool appears next to the **Polygon** tool.

3. Encapsulate an object (for example, a car) with a rectangular.
4. Click **Box to polygon**.



# Appearance settings

## Overview

H2O Label Genie offers the following appearance settings:

- [Dark mode](#)
- [Upload annotation instructions](#)

In the below sections, each setting above, in turn, is discussed.

## Dark mode

### Overview

By default, H2O Label Genie uses the **Dark mode** theme (appearance). Available themes are as follows:

- **Dark mode**
- **Light mode**

The image displays two screenshots of the H2O Label Genie settings page, illustrating the 'Dark mode' toggle. The top screenshot shows the 'Dark mode' toggle set to 'On', while the bottom screenshot shows it set to 'Off'. Both screenshots show the same settings page with a sidebar on the left and a main content area on the right. The sidebar includes links for Home, Datasets, Data exploration, Annotation tasks, Settings (highlighted), and Documentation. The main content area is divided into sections: Appearance, Smart annotation, Image settings, and Audio settings. The 'Appearance' section contains the 'Dark mode' toggle, 'Upload annotation instructions' toggle, 'Enable zero-shot predictions' toggle, and 'Enable keyboard support' toggle. The 'Image settings' section contains 'Max image width [px]' (1024), 'Max image height [px]' (512), and 'Max audio duration [s]' (60).

## Instructions

To change the default theme, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Settings**.
2. Click the **Dark mode** toggle.

## Upload annotation instructions

### Overview

By default, H2O Label Genie allows you to upload a **PDF** file containing instructions for an annotation task rubric. In particular, you can upload the **PDF** file in an annotation task's **Rubric** tab.

## Instructions

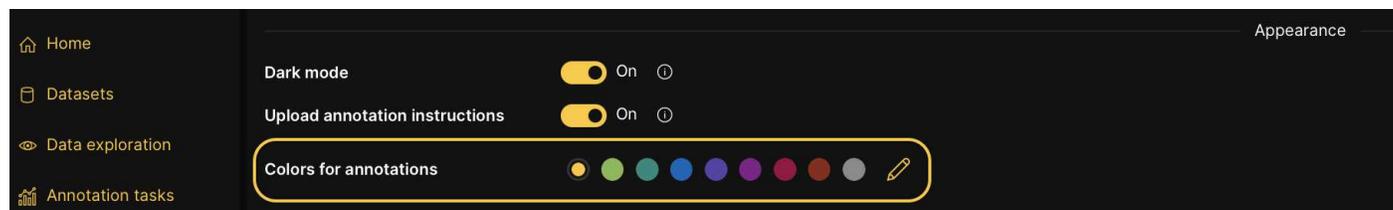
To disable or enable you to upload a **PDF** file containing annotation instructions for an annotation task, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Settings**.
2. Click the **Upload annotation instructions** toggle.

## Colors for annotations

### Overview

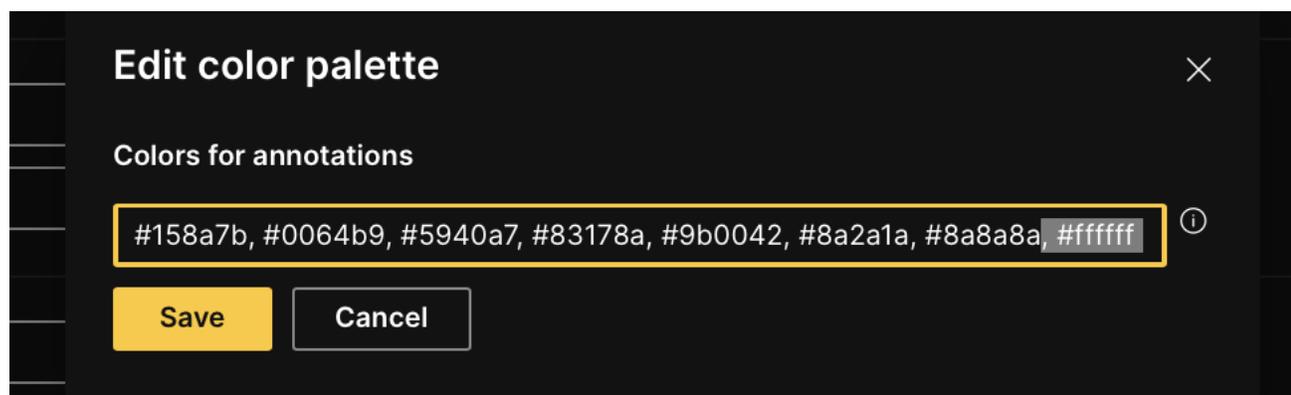
By default, H2O Label Genie provides a color palette with nine colors for annotating classes and objects. You can also add custom colors to the palette as needed.



### Instructions

To add more colors to the color palette, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Settings**.
2. Click the create Pencil icon.
3. In the provided text field, enter the hex code of the new color. When adding multiple colors, separate them with commas.



4. Click



## Smart annotation settings

### Overview

H2O Label Genie offers the following smart annotation settings:

- [Enable zero-shot predictions](#)
- [Enable segment anything](#)
- [Enable keyboard support](#)

In the below sections, each setting above, in turn, is discussed.

### Enable zero-shot predictions

#### Overview

H2O Label Genie enables you to allow H2O Label Genie to utilize zero-shot learning models to accelerate the labeling process for certain supported annotation tasks. You can turn this feature **On** or **Off**.

#### Note:

- H2O Label Genie lets you utilize a zero-shot learning model for the following supported annotation tasks:
  - [Image classification](#)
    - See [Annotation tasks + zero-shot learning models: Image classification](#)
  - [Object detection](#)
    - See, [Annotation tasks + zero-shot learning models: Object detection](#)
  - [Image instance segmentation](#)
    - See [Annotation tasks + zero-shot learning models: Image instance segmentation](#)
  - [Text classification](#)
    - See [Annotation tasks + zero-shot learning models: Text classification](#)
  - [Text summarization](#)
    - See [Annotation tasks + zero-shot learning models: Text summarization](#)
  - [Text-generative AI](#)
    - See [Annotation tasks + zero-shot learning models: Text-generative AI](#)

#### Instructions

To turn the *zero-shot learning models* on or off, consider the following instructions:

**Caution:** Turning this setting off does not allow H2O Label Genie to generate zero-shot predictions for annotation tasks supporting a zero-shot learning model(s).

1. On the **H2O Label Genie** navigation menu, click **Settings**.
2. Click the **Enable zero-shot predictions** toggle.

### Enable segment anything

#### Overview

By default, for an image instance segmentation annotation task, H2O Label Genie offers a **Box to polygon** feature that enables you to improve the efficiency of manual labeling. A **Box to polygon** feature enables you to encapsulate an object class with a rectangular that is immediately replaced with a more complex polygon that captures all the borders of the object class inside the rectangular.



### Instructions

To turn on or off the ability of H2O Label Genie to offer the **Box to polygon** feature for an image instance segmentation annotation task, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Settings**.
2. Click the **Enable segment anything** toggle.

## Enable keyboard support

### Overview

By default, H2O Label Genie enables you to utilize several keyboard shortcuts to speed up, for example, the labeling of a dataset (for example, the **U** key approves the annotated (labeled) sample to move to the next sample).

**Note:** To learn about available keyboard shortcuts, see [Hotkeys](#).

### Instructions

To disable or enable the keyboard shortcuts, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Settings**.
2. Click the **Enable keyboard support** toggle.

## Image settings

### Overview

H2O Label Genie offers the following image settings to determine how images are rendered for image annotation tasks:

- [Max image width \[px\]](#)
- [Max image height \[px\]](#)

In the below sections, each setting above, in turn, is discussed.

### Max image width [px]

#### Overview

By default, H2O Label Genie sets the width of the images you annotate (for example, for an object detection dataset).

#### Instructions

To change the default width (to adjust images to fit your screen), consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Settings**.
2. In the **Max image width [px]** box, enter an image width.

### Max image height [px]

#### Overview

By default, H2O Label Genie sets the height of the images you annotate (for example, for an object detection dataset).

#### Instructions

To change the default height (to adjust images to fit your screen), consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Settings**.
2. In the **Max image height [px]** box, enter an image width.

## Audio settings

### Overview

H2O Label Genie offers the following audio setting to determine how audios are rendered for audio annotation tasks:

- [Max audio duration \[s\]](#)

In the below section, the above setting is discussed.

### Max audio duration [s]

#### Overview

By default, H2O Label Genie sets an audio max duration of **60** seconds to the audio clips being previewed or annotated. You can change this limit at any time.

#### Instructions

To change the max audio duration limit, consider the following instructions:

1. On the **H2O Label Genie** navigation menu, click **Settings**.
2. In the **Max audio duration [s]** box, enter an audio duration limit.

# OpenAI API settings

## Overview

H2O Label Genie offers the following OpenAI API settings to connect to your OpenAI large language models (LLMs) for [Text summarization](#) and [Text-generative AI](#) annotation tasks:

- [OpenAI API key](#)
- [Use OpenAI API on Azure](#)
- [OpenAI API Endpoint](#)
- [OpenAI API Deployment ID](#)
- [OpenAI API version](#)

In the below sections, each setting above, in turn, is discussed.

## OpenAI API key

### Overview

This setting defines your OpenAI API key.

### Instructions

To define your OpenAI API key, consider the following instructions:

1. In the **H2O Label Genie** navigation menu, click **Settings**.
2. In the **OpenAI API key** box, enter your OpenAI API key.

## Use OpenAI API on Azure

### Overview

This setting allows you to connect to OpenAI through your Azure account. By default, this setting is turned **off**.

**Caution:** After turning this setting **on**, you need to define the following settings to connect to Azure successfully:

- [OpenAI API Endpoint](#)
- [OpenAI API deployment ID](#)
- [OpenAI API version](#)

### Instructions

To enable H2O Label Genie to allow you to access OpenAI through your Azure account, consider the following instructions:

1. In the **H2O Label Genie** navigation menu, click **Settings**.
2. Click the **Use OpenAI API on Azure** toggle.

## OpenAI API Endpoint

### Overview

This setting defines an OpenAI API endpoint on an Azure account (required when you turn on the following setting: **Use OpenAI API on Azure**).

### Instructions

To define an OpenAI API endpoint, consider the following instructions:

1. In the **H2O Label Genie** navigation menu, click **Settings**.
2. In the **OpenAI API endpoint** box, enter an OpenAI API endpoint on Azure.

## OpenAI API deployment ID

This setting defines an OpenAI API deployment ID on an Azure account (required when you turn on the following setting: **Use OpenAI API on Azure**).

## Instructions

To define an OpenAI API deployment ID, consider the following instructions:

1. In the **H2O Label Genie** navigation menu, click **Settings**.
2. In the **OpenAI API deployment ID** box, enter an OpenAI API deployment ID on Azure.

## OpenAI API version

### Overview

This setting defines an OpenAI API version on an Azure account (required when you turn on the following setting: **Use OpenAI API on Azure**).

### Instructions

To define an OpenAI API version, consider the following instructions:

1. In the **H2O Label Genie** navigation menu, click **Settings**.
2. In the **OpenAI API version** box, enter an OpenAI API version on Azure.

## Collaboration settings

### Overview

H2O Label Genie offers the following collaboration setting:

- [Multi-user annotation](#)

In the below section, the above setting is discussed.

### Multi-user annotation

#### Overview

H2O Label Genie enables you to work with others to improve/speed up the data labeling process. You and your team can complete an annotation task in parallel within the same app instance, resulting in real-time multi-user support and collaboration.

To collaborate, one user must create the app instance and share it with others. The dataset used for the annotation task is stored in this shared app instance.

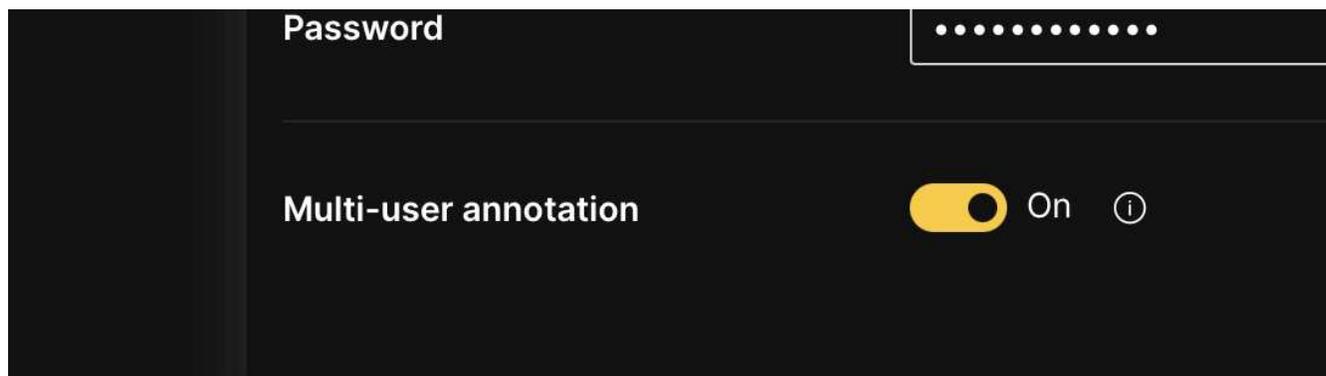
By default, H2O Label Genie enables others to annotate the same annotation task. You can change this at any time.

**Note:** When turned **On**, H2O Label Genie displays a different sample for each user.

#### Instructions

To turn on or off the ability for others to annotate the same annotation task, consider the following instructions:

1. In the **H2O Label Genie** navigation menu, click **Settings**.
2. Click the **Multi-user annotation** toggle.



## Key terms

H2O Label Genie uses several key terms across its documentation, and each, in turn, is explained in the sections below.

### Annotation task

An annotation task refers to labeling data in a manner that makes them suitable to support an array of deep learning problem types. This process can, for example, involve adding bounding boxes to images where you can attribute labels to each box.

### Annotation task rubric

An annotation task rubric refers to the labels (for example, object classes) you want to use when annotating your dataset.

### Clustering task

A clustering task refers to finding and exploring groups in a dataset.

## AI governance

### What is AI governance?

AI Governance is the set of frameworks, rules, and best practices to ensure responsible adoption and usage of artificial intelligence. This framework encourages organizations to curate and use bias free data, consider societal and end-user impact, and produce unbiased models; the framework also enforces controls on model progression through deployment stages. AI Governance is critical for organizations to realize the maximum value out of AI projects while mitigating risks. It enables organizations to not only develop AI projects in a responsible way, but also ensure that there is consistency and transparency across the entire organization.

### H2O.ai's AI governance with H2O Label Genie

H2O.ai's AI governance framework recommends four stages and a total of 11 topics. Organizations are encouraged to adopt the topics and processes most relevant to their unique needs. The H2O.ai's AI governance framework is discussed in the following guide: [Guidelines for Effective AI Governance with Applications in H2O AI Cloud](#).

For example, based on the content of the guide, H2O Label Genie can help with AI governance in one of the following ways:

1. **Data quality and annotation:** H2O Label Genie can generate high-quality datasets by annotating multiple dataset types, including image, audio, and text. This can help ensure that the data used for AI model training is accurate and well-curated, which is essential for effective AI governance.
2. **Data privacy and security:** H2O Label Genie provides data annotation capabilities that can help organizations protect sensitive data and maintain data privacy. This is particularly important in AI governance, as organizations must handle personal data responsibly and securely.

Overall, H2O Label Genie can help organizations implement effective AI governance by providing capabilities that support data quality, privacy, security, explainability, and continuous monitoring and improvement.

## Release notes

### v1.2.0 | March 20, 2025

#### Fixes

- Fixed potential XSS vulnerability.
- Upgraded transformers dependency (CVE fix).

## FAQs

H2O Label Genie is an application that enables you with zero-shot learning models to rapidly label your datasets for [annotation tasks](#) in computer vision (CV), natural language processing (NLP), and audio.

---

The below sections provide answers to frequently asked questions. If you have additional questions, please send them to [cloud-feedback@h2o.ai](mailto:cloud-feedback@h2o.ai).

### General

#### **What is the difference between a CPU and GPU bundle of H2O Label Genie?**

With H2O Label Genie **v0.3.0+**, you can obtain a CPU or GPU-based bundle of the application. H2O Label Genie offers a GPU-based bundle to users with sufficient resources to experience better performance of supported [zero-shot models](#) and [clustering tasks](#). In the case of clustering tasks, a GPU-based bundle accelerates the creation of the 2D and 3D embeddings.