

Lab 1: Computer Vision Applications


CSE 185 Introduction to Computer Vision, Fall 2023

Description: The goal of this lab is to try and get familiar with various computer vision applications. Run computer vision demons through <https://huggingface.co> which is an online community for open-sourced models. See the results for fun!

1 Image classification

Image classification is the task of assigning a label or class to an entire image. For simplicity in this demo, each images is expected to have only one class. Image classification models take an image as input and return a prediction about which class the image belongs to.

URL: <https://huggingface.co/tasks/image-classification>




Upload images of your choice and check which class has the highest probability.

2 Text to image

Text-to-image is the task of generating an image based on a text description (prompt). Recent advances in deep learning has enabled image generation that is highly realistic.

URL: <https://huggingface.co/tasks/text-to-image>


Try the online demo and generate desired images based on input text prompts. Change prompts and see how the generated images are changed accordingly.



3 Visual question answering

Visual Question Answering is the task of answering open-ended questions based on an image. They output natural language responses to natural language questions.

URL: <https://huggingface.co/tasks/visual-question-answering>




Upload images of your choice and ask questions about uploaded images, e.g., "what's in this image and how many people are there in this image".

4 Image to text

Image to text models output a text from a given image. Image captioning or optical character recognition can be considered as the most common applications of image to text.

URL: <https://huggingface.co/tasks/image-to-text>




Upload images of your choice and see what descriptions are generated.

5 Object detection

Object detection models allow users to identify objects of certain defined classes. Object detection models receive an image as input and output the images with bounding boxes and labels on detected objects.

URL: <https://huggingface.co/tasks/object-detection>



Upload images of your choice for the demo. Each bounding box is associated with a probability. Hover over each bounding box to see the output probability.