DeepRob

Discussion 3
PROPS Dataset and Retraining
University of Michigan | Department of Robotics

PROPS Datasets

- Downsampled data from the ProgressLabeller annotation tool (Chen et al., 2022)
- Focuses on table-top scenes

PROPS Classification

- Image classification
- Based on CIFAR-10 dataset
- 10 object categories with 50K training images and 10K testing images.
- 32x32 RGB color image

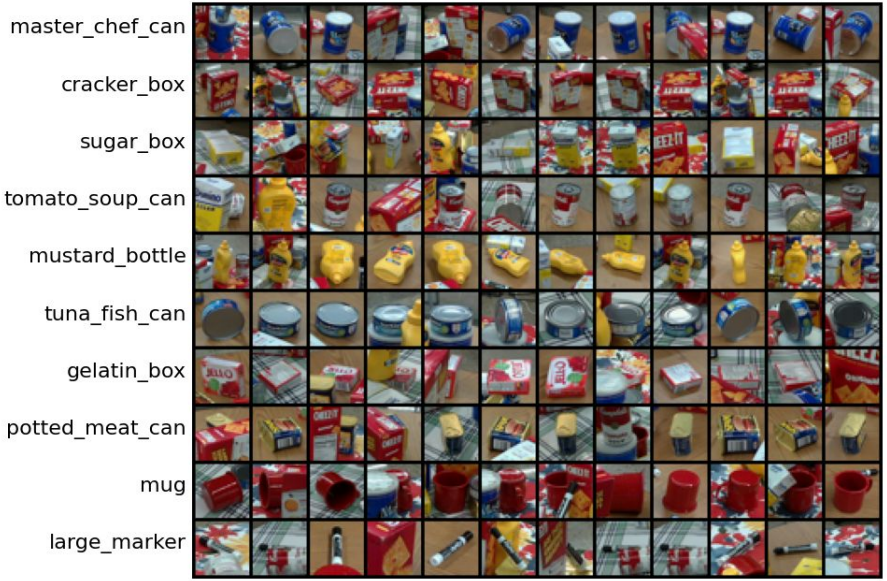

PROPS Detection

- Object detection tasks
- 10 object categories with 2.5K training images and 2.5K validation images
- 640x480 RGB color image

PROPS Pose Estimation

- 6 degrees-of-freedom rigid body object pose estimation.
- 10 object categories with 500 training images and 500 validation images
- 640x480 RGB color image with depth images and segmentation masks

Simple Classification

- Use Resnet to classify objects in the PROPS dataset
- What was the highest accuracy for the training and test that you achieved in the homework?
- We well tune a pretrained Resnet to achieve at least a 70% training and 60% test accuracy

ResNet

- From the Paper “Deep Residual Learning for Image Recognition”
- We will talk in more detail in later classes
- ResNet comes is several variants e.g. ResNet-18, ResNet-34
- The higher the number the deeper the network

Fine Tuning

- Technique that is commonly use if there exits pre-trained models
- Theory is that the pretrained models some useful knowledge/feature recognition, then we only need to retrain a little usually the last layer
- Example Use case: Retraining on a new images or different domain

Live Demo

- Finetune ResNet-18 to classify the PROPS dataset

DeepRob

Discussion 3
PROPS Dataset and Retraining
University of Michigan | Department of Robotics