
PeekingDuck

Release developer

CVHub AI Singapore

Feb 01, 2023

CONTENTS

1	Introduction	1
1.1	What is PeekingDuck?	1
1.2	Features	1
1.3	How PeekingDuck Works	2
1.4	Acknowledgements	2
1.5	License	2
1.6	Communities	2
2	Getting Started	3
2.1	Documentation Convention	3
2.2	Standard Install	3
2.3	Custom Install	4
3	Tutorials	7
3.1	“Hello Computer Vision”	7
3.2	Duck Confit	13
3.3	Custom Nodes	17
3.4	Peeking Duck	32
3.5	Calling PeekingDuck in Python	53
3.6	Using Your Own Models	61
3.7	Using Model Hub Nodes	73
4	PeekingDuck Ecosystem	81
4.1	PeekingDuck Viewer	81
5	Model Resources & Information	85
5.1	Object Detection Models	85
5.2	Pose Estimation Models	89
5.3	Object Tracking Models	91
5.4	Crowd Counting Models	92
5.5	Instance Segmentation Models	93
5.6	Bibliography	97
6	Edge AI	99
6.1	Installing TensorRT	99
6.2	Using TensorRT Models	99
6.3	Performance Speedup	100
6.4	References	103
7	Use Cases	105
7.1	Privacy Protection	105

7.2	Smart Monitoring	111
7.3	COVID-19 Prevention and Control	123
8	FAQ and Troubleshooting	131
8.1	How can I post-process and visualize model outputs?	131
8.2	How can I dynamically use all prior outputs as the input at run time?	131
8.3	How do I debug custom nodes?	131
8.4	Why does <code>input.visual</code> progress stop before 100%?	131
8.5	Why does the output screen flash briefly and disappear on my second run?	132
9	Glossary	133
10	API Documentation	135
10.1	<code>input</code>	135
10.2	<code>augment</code>	135
10.3	<code>model</code>	136
10.4	<code>dabble</code>	136
10.5	<code>draw</code>	136
10.6	<code>output</code>	137
	Python Module Index	139
	Index	141

INTRODUCTION

1.1 What is PeekingDuck?

PeekingDuck is an open-source, modular framework in Python, built for Computer Vision (CV) inference. The name “PeekingDuck” is a play on: “Peeking” in a nod to CV; and “Duck” in [duck typing](#) used in Python.

1.2 Features

1.2.1 Build realtime CV pipelines

PeekingDuck enables you to build powerful CV pipelines with minimal lines of code.

1.2.2 Leverage on SOTA models

PeekingDuck comes with various state of the art (SOTA) *object detection*, *pose estimation*, *object tracking*, *crowd counting*, and *instance segmentation* models. Mix and match different nodes to construct solutions for various *use cases*.

1.2.3 Create custom nodes

You can create *custom nodes* to meet your own project’s requirements. PeekingDuck can also be *imported as a library* to fit into your existing workflows.

1.3 How PeekingDuck Works

Nodes are the building blocks of PeekingDuck. Each node is a wrapper for a pipeline function, and contains information on how other PeekingDuck nodes may interact with it.

PeekingDuck has 6 types of nodes:

A **pipeline** governs the behavior of a chain of nodes. The diagram below shows a sample pipeline. Nodes in a pipeline are called in sequential order, and the output of one node will be the input to another. For example, `input.visual` produces *img*, which is taken in by `model.yolo`, and `model.yolo` produces *bboxes*, which is taken in by `draw.bbox`. For ease of visualization, not all the inputs and outputs of these nodes are included in this diagram.

1.4 Acknowledgements

This project is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG-RP-2019-050). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the National Research Foundation, Singapore.

1.5 License

PeekingDuck is under the open source [Apache License 2.0](#) (:

Even so, your organization may require legal proof of its right to use PeekingDuck, due to circumstances such as the following:

- Your organization is using PeekingDuck in a jurisdiction that does not recognize this license
- Your legal department requires a license to be purchased
- Your organization wants to hold a tangible legal document as evidence of the legal right to use and distribute PeekingDuck

[Contact us](#) if any of these circumstances apply to you.

1.6 Communities

- [AI Singapore community forum](#)
- [Discuss on GitHub](#)

GETTING STARTED

2.1 Documentation Convention

Parts of this documentation and the tutorials are run from the command-line interface (CLI) environment, e.g., via *Terminal* in Linux/macOS, or via *Anaconda* in Windows. There will be examples of commands you need to type as inputs and text that PeekingDuck will display as outputs. The input commands can be dependent on the current folder where they are typed.

The following text color scheme is used to illustrate these different contexts:

Color	Context	Example
Blue	Current folder	[~user/src]
Green	User input: what you type in	> peekingduck --version
Black	PeekingDuck's output	peekingduck, version v1.2.0

The command prompt is assumed to be the symbol `>`, your home directory is assumed to be `~user`, and the symbol means to press the `<Enter>` key.

Putting it altogether, a sample terminal session looks like this:

Terminal Session

```
[~user/src] > peekingduck --version  
peekingduck, version v1.2.0
```

2.2 Standard Install

2.2.1 Install PeekingDuck

Then run:

PeekingDuck supports Python 3.6 to 3.9.

It is recommended to install PeekingDuck in a Python virtual environment (such as `pkd` in the above commands), as it creates an isolated environment for a Python project to install its own dependencies and avoid package version conflicts with other projects.

Note: For Apple Silicon Mac users, please see *Custom Install - Apple Silicon Mac*.

2.2.2 Verify PeekingDuck Installation

To check that PeekingDuck is installed successfully, run the following command:

Terminal Session

```
[~user] > peekingduck verify-install
```

Changed in version 1.3.0: The verify installation command has been changed from `--verify_install` to `verify-install`.

You should see a video of a person waving his hand ([taken from here](#)) with bounding boxes overlaid as shown below:

The video will auto-close when it is run to the end (about 20 seconds, depending on system speed).
To exit earlier, click to select the video window and press q.

2.3 Custom Install

This section covers advanced PeekingDuck installation steps for users with ARM64 devices or Apple Silicon Macs.

2.3.1 Arm64

To install PeekingDuck on an ARM-based device, such as a Raspberry Pi, include the `--no-dependencies` flag, and separately install the other dependencies listed in PeekingDuck's [\[requirements.txt\]](#):

Terminal Session

```
[~user] > pip install peekingduck --no-dependencies  
[~user] > [ install additional dependencies as specified within requirements.txt ]
```

Verify the installation using:

Terminal Session

```
[~user] > peekingduck verify-install
```

See [here](#) for changes to the verify installation command in version 1.3.0.

You should see a video of [a person waving his hand with bounding boxes overlaid](#).

2.3.2 Apple Silicon Mac

Apple released their advanced ARM-based [Apple Silicon M1](#) chip in late 2020, a significant change from the previous Intel processors. We've successfully installed PeekingDuck on Apple Silicon Macs running macOS Big Sur and macOS Monterey.

1. Prerequisites:

- Install [homebrew](#)
- Install miniforge using homebrew:

Terminal Session

```
[~user] > brew install miniforge
```

2. Create conda virtual environment and install base packages:

Terminal Session

```
[~user] > conda create -n pkd python=3.8  
[~user] > conda activate pkd  
[~user] > conda install click colorama opencv openblas pyyaml requests scipy shapely tqdm typeguard
```

3. Install Apple's Tensorflow build that is optimized for Apple Silicon Macs:

- For macOS Monterey:

Terminal Session

```
[~user] > conda install -c apple tensorflow-deps  
[~user] > pip install tensorflow-macos tensorflow-metal
```

- For macOS Big Sur:

Terminal Session

```
[~user] > conda install -c apple tensorflow-deps=2.6.0
```

```
[~user] > pip install tensorflow-estimator==2.6.0 tensorflow-macos==2.6.0  
[~user] > pip install tensorflow-metal==0.2.0
```

4. Install PyTorch (currently CPU-only):

Terminal Session

```
[~user] > pip install torch torchvision
```

5. Install PeekingDuck and verify installation:

Terminal Session

```
[~user] > pip install peekingduck --no-dependencies  
[~user] > peekingduck verify-install
```

See [here](#) for changes to the verify installation command in version 1.3.0.

You should see a video of a person waving his hand ([taken from here](#)) with bounding boxes overlaid as shown below:

The video will auto-close when it is run to the end (about 20 seconds, depending on system speed).
To exit earlier, click to select the video window and press q.

TUTORIALS

The tutorials are presented in order of increasing difficulty, from the basic *Hello Computer Vision* to the advanced *Peeking Duck*. It is recommended to go through these tutorials in order, especially if you are new to PeekingDuck.

3.1 “Hello Computer Vision”

Computer Vision (or CV) is a field in AI that develops techniques to help computers to “see” and “understand” the contents of digital images like photographs and videos, and to derive meaningful information. Common CV applications include object detection to detect what objects are present in the image and pose estimation to detect the position of human limbs relative to the body.

PeekingDuck allows you to build a CV pipeline to analyze and process images and/or videos. This pipeline is made up of nodes: each node can perform certain CV-related tasks.

This section presents two basic “hello world” examples to demonstrate how to use PeekingDuck for pose estimation and object detection.

3.1.1 Pose Estimation

To perform pose estimation with PeekingDuck, initialize a new PeekingDuck project using the following commands:

Terminal Session

```
[~user] > mkdir pose_estimation
[~user] > cd pose_estimation
[~user/pose_estimation] > peekingduck init
```

peekingduck init will prepare the pose_estimation folder for use with PeekingDuck. It creates a default pipeline file called pipeline_config.yml and a src folder that will be covered in the later tutorials. The pipeline_config.yml file looks like this:

```
1 nodes:
2 - input.visual:
3   source: https://storage.googleapis.com/peekingduck/videos/wave.mp4
4 - model.posenet
```

(continues on next page)

(continued from previous page)

```
5 - draw.pose
6 - output.screen
```

The above forms a **pose estimation pipeline** and it comprises four nodes that do the following:

1. `input.visual`: reads the file `wave.mp4` from PeekingDuck's cloud storage
2. `model.posenet`: runs the PoseNet pose estimation model on it
3. `draw.pose`: draws a human pose skeleton over the person tracking his hand movement
4. `output.screen`: outputs everything onto the screen for display

Now, run the pipeline using

Terminal Session

```
[~user/pose_estimation] > peekingduck run
```

You should see the following video of a person waving his hand. Skeletal poses are drawn on him which track the hand movement.



Fig. 1: PeekingDuck's Pose Estimation Screenshot

You have successfully run a PeekingDuck pose estimation pipeline!

The video will auto-close when it is completed.
To exit earlier, click to select the video window and press q.

3.1.2 Object Detection

To perform object detection, initialize a new PeekingDuck project using the following commands:

Terminal Session

```
[~user] > mkdir object_detection  
[~user] > cd object_detection  
[~user/object_detection] > peekingduck init
```

Then modify `pipeline_config.yml` as follows:

```
1 nodes:  
2 - input.visual:  
3   source: https://storage.googleapis.com/peekingduck/videos/wave.mp4  
4 - model.yolo  
5 - draw.bbox  
6 - output.screen
```

The key differences between this and the earlier pipeline are:

Line 4: `model.yolo` runs the YOLO object detection model

Line 5: `draw.bbox` draws the bounding box to show the detected person

Run the new **object detection pipeline** with `peekingduck run`.

You will see the same video with a bounding box surrounding the person.

That's it: you have created a new object detection pipeline by changing only two lines!

Note:

Try replacing `wave.mp4` with your own video file and run both models.

For best effect, your video file should contain people performing some activities.



Fig. 2: PeekingDuck's Object Detection Screenshot

3.1.3 Using a WebCam

If your computer has a webcam attached, you can use it by changing the first `input` node (line 2) as follows:

```
1 nodes:
2 - input.visual:
3   source: 0      # use webcam for live video
4 - model.posenet  # use pose estimation model
5 - draw.pose      # draw skeletal poses
6 - output.screen
```

Now do a peekingduck run and you will see yourself onscreen. Move your hands around and see PeekingDuck tracking your poses.

To exit, click to select the video window and press `q`.

Note: PeekingDuck assumes the webcam is defaulted to input source 0. If your system is configured differently, you would have to specify the input source by changing the `input.visual` configuration. See [changing node configuration](#).

3.1.4 Pipelines, Nodes and Configs

PeekingDuck comes with a rich collection of nodes that you can use to create your own CV pipelines. Each node can be customized by changing its configurations or settings.

To get a quick overview of PeekingDuck's nodes, run the following command:

Terminal Session

[~user] > peekingduck nodes

```
> peekingduck nodes

PeekingDuck has the following input nodes:
1:visual Info: https://peekingduck.readthedocs.io/en/stable/nodes/input.visual.html#module-input.visual

PeekingDuck has the following augment nodes:
1:brightness Info: https://peekingduck.readthedocs.io/en/stable/nodes/augment.brightness.html#module-augment.brightness
2:contrast Info: https://peekingduck.readthedocs.io/en/stable/nodes/augment.contrast.html#module-augment.contrast

PeekingDuck has the following model nodes:
1:yolox Info: https://peekingduck.readthedocs.io/en/stable/nodes/model.yolox.html#module-model.yolox
2:mtcnn Info: https://peekingduck.readthedocs.io/en/stable/nodes/model.mtcnn.html#module-model.mtcnn
3:csrnet Info: https://peekingduck.readthedocs.io/en/stable/nodes/model.csrnet.html#module-model.csrnet
4:yolo_license_plate Info: https://peekingduck.readthedocs.io/en/stable/nodes/model.yolo_license_plate.html#module-model
```

You will see a comprehensive list of all PeekingDuck's nodes with links to their readthedocs pages for more information.

PeekingDuck supports 6 types of nodes:

A PeekingDuck pipeline is created by stringing together a series of nodes that perform a logical sequence of operations. Each node has its own set of configurable settings that can be modified to change its behavior. An example pipeline is shown below:

3.1.5 Bounding Box vs Image Coordinates

PeekingDuck has two (x, y) coordinate systems, with top-left corner as origin $(0, 0)$:

- **Absolute image coordinates**

For an image of width W and height H , the absolute image coordinates are integers from $(0, 0)$ to $(W - 1, H - 1)$. E.g., for a 720 x 480 image, the absolute coordinates range from $(0, 0)$ to $(719, 479)$.

- **Relative bounding box coordinates**

For an image of width W and height H , the relative image coordinates are real numbers from $(0.0, 0.0)$ to $(1.0, 1.0)$. E.g., for a 720 x 480 image, the relative coordinates range from $(0.0, 0.0)$ to $(1.0, 1.0)$.

This means that in order to draw a bounding box onto an image, the bounding box relative coordinates would have to be converted to the image absolute coordinates.

Using the above figure as an illustration, the bounding box coordinates are given as $(0.18, 0.10)$ top-left and $(0.52, 0.88)$ bottom-right. To convert them to image coordinates, multiply the x-coordinates by the image width and the y-

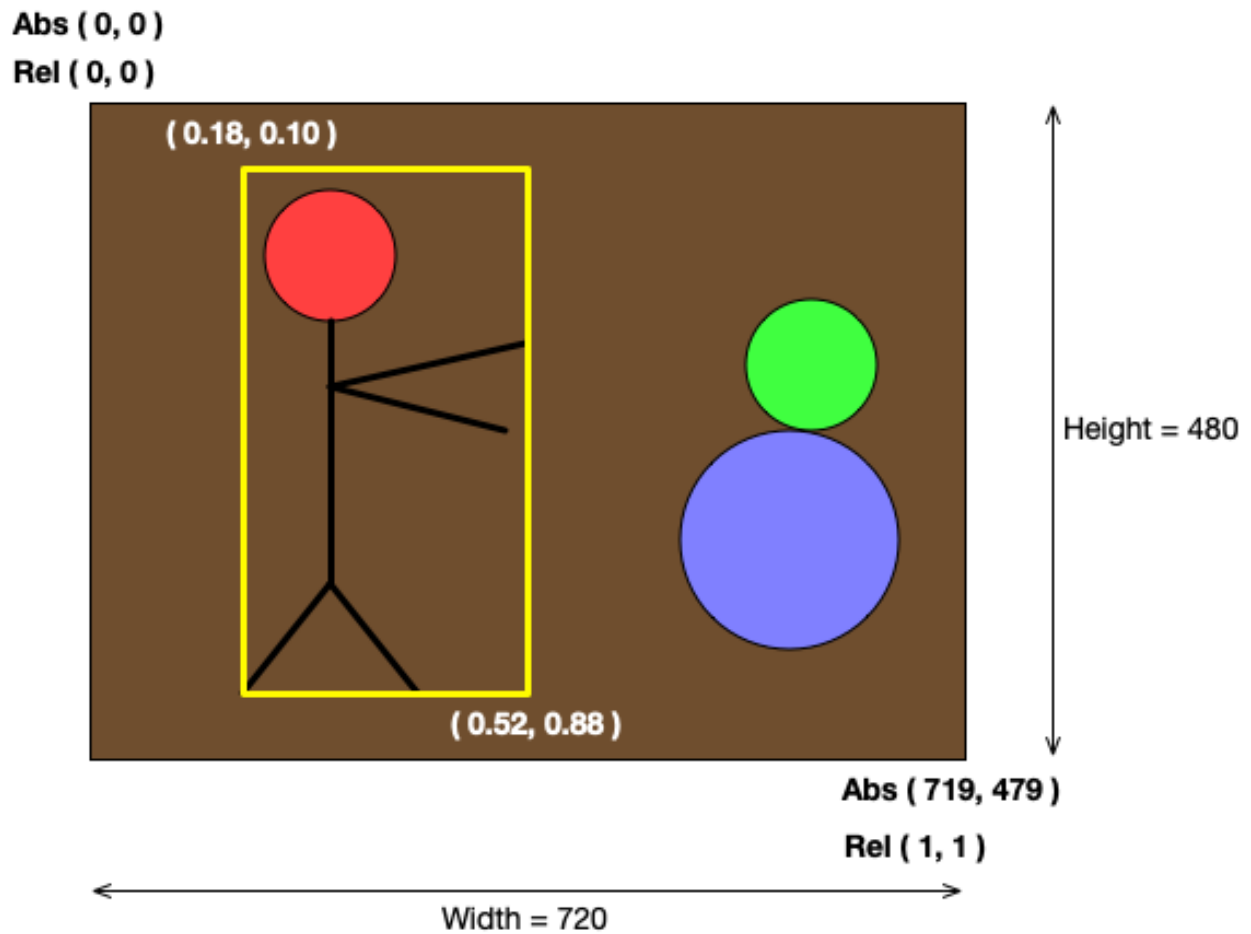


Fig. 3: PeekingDuck's Image vs Bounding Box Coordinates

coordinates by the image height, and round the results into integers.

$$\begin{aligned} 0.18 &\rightarrow 0.18 \times 720 = 129.6 = 130 & (int) \\ 0.10 &\rightarrow 0.10 \times 480 = 48.0 = 48 & (int) \end{aligned}$$

$$\begin{aligned} 0.52 &\rightarrow 0.52 \times 720 = 374.4 = 374 & (int) \\ 0.88 &\rightarrow 0.88 \times 480 = 422.4 = 422 & (int) \end{aligned}$$

Thus, the image coordinates are (130, 48) top-left and (374, 422) bottom-right.

Note: The *model* nodes return results in relative coordinates.

3.2 Duck Confit

This tutorial presents intermediate recipes for cooking up new PeekingDuck pipelines by modifying the nodes and their configs.

3.2.1 More Object Detection

This section will demonstrate how to change the settings of PeekingDuck's nodes to vary their functionalities.

If you had completed the earlier *object detection tutorial*, you will have the necessary folder and can skip to the next step. Otherwise, create a new PeekingDuck project as shown below:

Terminal Session

```
[~user] > mkdir object_detection
[~user] > cd object_detection
[~user/object_detection] > peekingduck init
```

Next, download this demo video [cat_and_computer.mp4](#) and save it into the `object_detection` folder.

The folder should contain the following:

```
object_detection/
├── cat_and_computer.mp4
├── pipeline_config.yml
└── src/
```

To perform object detection on the `cat_and_computer.mp4` file, edit the `pipeline_config.yml` file as follows:

```
1 nodes:
2   - input.visual:
3     source: cat_and_computer.mp4
4   - model.yolo:
5     detect: ["cup", "cat", "laptop", "keyboard", "mouse"]
6   - draw.bbox:
7     show_labels: True    # configure draw.bbox to display object labels
8   - output.screen
```

Here is a step-by-step explanation of what has been done:

Line 2 `input.visual`: tells PeekingDuck to load the `cat_and_computer.mp4`.

Line 4 `model.yolo`: by default, the YOLO model detects `person` only.

The `cat_and_computer.mp4` contains other classes of objects like cup, cat, laptop, etc.

So we have to change the model settings to detect the other object classes.

Line 6 `draw.bbox`: reconfigure this node to display the detected object class label.

Run the above with the command `peekingduck run`.

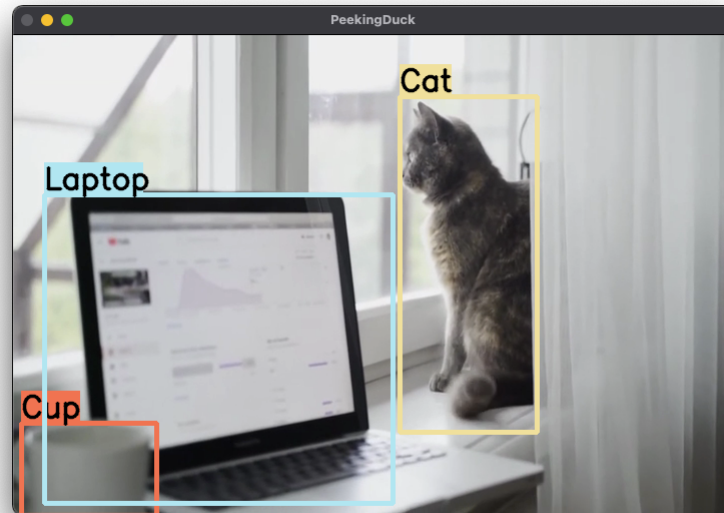


Fig. 4: Cat and Computer Screenshot

You should see a display of the `cat_and_computer.mp4` with the various objects being highlighted by PeekingDuck in bounding boxes. The 30-second video will auto-close at the end, or you can press `q` to end early.

Note: The YOLO model can detect 80 different *object classes*. By default, it only detects the `"person"` class. Use `detect: ["*"]` in the `pipeline_config.yml` to configure the model to detect all 80 classes.

3.2.2 Record Video File with FPS

This section demonstrates how to record PeekingDuck's output into a video file. In addition, we will modify the pipeline by adding new nodes to calculate the frames per second (FPS) and to show the FPS.

Edit `pipeline_config.yml` as shown below:

```
1 nodes:
2   - input.visual:
3     source: cat_and_computer.mp4
4   - model.yolo:
```

(continues on next page)

(continued from previous page)

```

5     detect: ["cup", "cat", "laptop", "keyboard", "mouse"]
6 -   draw.bbox:
7       show_labels: True
8 -   dabble.fps                                # add new dabble node
9 -   draw.legend:                             # show fps
10      show: ["fps"]
11 -   output.screen
12 -   output.media_writer:                    # add new output node
13      output_dir: /folder/to/save/video    # this is a folder name

```

The additions are:

Line 8 `dabble.fps`: adds new *dabble* node to the pipeline. This node calculates the FPS.

Line 9 `draw.legend`: adds new *draw* node to display the FPS.

Line 12 `output.media_writer`: adds new *output* node to save PeekingDuck's output to a local video file. It requires a local folder path. If the folder is not available, PeekingDuck will create the folder automatically. The filename is auto-generated by PeekingDuck based on the input source.

Run the above with the command `peekingduck run`. You will see the same video being played, but now it has the FPS counter. When the video ends, an `mp4` video file will be created and saved in the specified folder.

Note: You can view all the available nodes and their respective configurable settings in PeekingDuck's *API documentation*.

3.2.3 Configuration - Behind the Scenes

Here is an explanation on what goes on behind the scenes when you configure a node. Every node has a set of default configuration. For instance, `draw.bbox` default configuration is:

```

1 input: ["bboxes", "img", "bbox_labels"]
2 output: ["none"]
3
4 show_labels: False

```

The keys `input` and `output` are compulsory and common across every node.

`input` specifies the data types the node would consume, to be read from the pipeline.

`output` specifies the data types the node would produce, to be put into the pipeline.

By default, `show_labels` is disabled. When you enable it with `show_labels: True`, what PeekingDuck does is to override the default `show_labels: False` configuration with your specified `True` value. You will see another instance of this at work in the advanced *Peaking Duck* tutorial on *Tracking People Within a Zone*.

3.2.4 Augmenting Images

PeekingDuck has a class of *augment* nodes that can be used to perform preprocessing or postprocessing of images/videos. Augment currently lets you modify the brightness and contrast, and remove distortion from a wide-angle camera image. For more details on image undistortion, refer to the documentation on `augment.undistort` and `dabble.camera_calibration`.

The `pipeline_config.yml` below shows how to use the `augment.brightness` node within the pipeline:

```

1 nodes:
2   - input.visual:
3     source: https://storage.googleapis.com/peekingduck/videos/wave.mp4
4   - model.yolo
5   - augment.brightness:
6     beta: 50      # ranges from -100 (darken) to +100 (brighten)
7   - draw.bbox
8   - output.screen

```

The following figure shows the difference between the original vs the brightened image:



Fig. 5: Augment Brightness: Original vs Brightened Image

Note:

Royalty free video of cat and computer from: <https://www.youtube.com/watch?v=-C1TEGZavko>

Royalty free video of man waving hand from: https://www.youtube.com/watch?v=IKj_z2hgYUM

3.3 Custom Nodes

This tutorial will show you how to create your own custom nodes to run with PeekingDuck. Perhaps you'd like to take a snapshot of a video frame, and post it to your API endpoint; or perhaps you have a model trained on a custom dataset, and would like to use PeekingDuck's *input*, *draw*, and *output* nodes. PeekingDuck is designed to be very flexible — you can create your own nodes and use them with ours.

Let's start by creating a new PeekingDuck project:

Terminal Session

```
[~user] > mkdir custom_project
[~user] > cd custom_project
[~user/custom_project] > peekingduck init
```

This creates the following `custom_project` folder structure:

```
custom_project/
├── pipeline_config.yml
├── src/
│   ├── custom_nodes/
│   │   └── configs/
```

The sub-folders `src`, `custom_nodes`, and `configs` are empty: they serve as placeholders for contents to be added.

3.3.1 Recipe 1: Object Detection Score

When the YOLO object detection model detects an object in the image, it assigns a bounding box and a score to it. This score is the “confidence score” which reflects how likely the box contains an object and how accurate is the bounding box. It is a decimal number that ranges from 0.0 to 1.0 (or 100%). This number is internal and not readily viewable.

We will create a custom node to retrieve this score and display it on screen. This tutorial will use the [cat_and_computer.mp4](#) video from the earlier *object detection tutorial*. Copy it into the `custom_project` folder.

Use the following command to create a custom node: `peekingduck create-node`. It will prompt you to answer several questions. Press <Enter> to accept the default `custom_nodes` folder name, then key in `draw` for node type and score for node name. Finally, press <Enter> to answer `Y` when asked to proceed.

The entire interaction is shown here, the answers you type in are shown in green text:

Terminal Session

```
[~user/custom_project] > peekingduck create-node
Creating new custom node...
Enter node directory relative to ~user/custom_project [src/custom_nodes]:
Select node type (input, augment, model, draw, dabble, output): draw
Enter node name [my_custom_node]: score
```

```
Node directory: ~user/custom_project/src/custom_nodes
```

Node type: draw

Node name: score

Creating the following files:

Config file: `~user/custom_project/src/custom_nodes/configs/draw/score.yml`

Script file: `~user/custom_project/src/custom_nodes/draw/score.py`

Proceed? [Y/n]:

Created node!

The `custom_project` folder structure should look like this:

```
custom_project/
├── cat_and_computer.mp4
├── pipeline_config.yml
├── src/
│   ├── custom_nodes/
│   │   ├── configs/
│   │   │   └── draw/
│   │   │       └── score.yml
│   │   └── draw/
│   │       └── score.py
```

`custom_project` now contains **three files** that we need to modify to implement our custom node function.

1. **`src/custom_nodes/configs/draw/score.yml`:**

`score.yml` initial content:

```
1  # Mandatory configs
2  # Receive bounding boxes and their respective labels as input. Replace with
3  # other data types as required. List of built-in data types for PeekingDuck can
4  # be found at https://peekingduck.readthedocs.io/en/stable/glossary.html.
5  input: ["bboxes", "bbox_labels"]
6  # Example:
7  # Output `obj_attrs` for visualization with `draw.tag` node and `custom_key` for
8  # use with other custom nodes. Replace as required.
9  output: ["obj_attrs", "custom_key"]
10
11 # Optional configs depending on node
12 threshold: 0.5 # example
```

The first file `score.yml` defines the properties of the custom node. Lines 5 and 9 show the mandatory configs `input` and `output`.

`input` specifies the data types the node would consume, to be read from the pipeline. `output` specifies the data types the node would produce, to be put into the pipeline.

To display the bounding box confidence score, our node requires three pieces of input data: the bounding box, the score to display, and the image to draw on. These are defined as the data types `bboxes`, `bbox_scores`, and `img` respectively in the [API docs](#).

Our custom node only displays the score on screen and does not produce any outputs for the pipeline, so the output is `"none"`.

There are also no optional configs, so lines 11 - 12 can be removed.

score.yml updated content:

```

1 # Mandatory configs
2 input: ["img", "bboxes", "bbox_scores"]
3 output: ["none"]
4
5 # No optional configs

```

Note: Comments in yaml files start with # It is possible for a node to have input: ["none"]

2. src/custom_nodes/draw/score.py:

The second file score.py contains the boilerplate code for creating a custom node. Update the code to implement the desired behavior for the node.

Show/Hide Code for score.py

```

1 """
2 Custom node to show object detection scores
3 """
4
5 from typing import Any, Dict, List, Tuple
6 import cv2
7 from peekingduck.nodes.abstract_node import AbstractNode
8
9 YELLOW = (0, 255, 255)          # in BGR format, per opencv's convention
10
11
12 def map_bbox_to_image_coords(
13     bbox: List[float], image_size: Tuple[int, int]
14 ) -> List[int]:
15     """This is a helper function to map bounding box coords (relative) to
16     image coords (absolute).
17     Bounding box coords ranges from 0 to 1
18     where (0, 0) = image top-left, (1, 1) = image bottom-right.
19
20     Args:
21         bbox (List[float]): List of 4 floats x1, y1, x2, y2
22         image_size (Tuple[int, int]): Width, Height of image
23
24     Returns:
25         List[int]: x1, y1, x2, y2 in integer image coords
26     """
27     width, height = image_size[0], image_size[1]
28     x1, y1, x2, y2 = bbox
29     x1 *= width
30     x2 *= width
31     y1 *= height
32     y2 *= height
33     return int(x1), int(y1), int(x2), int(y2)
34
35
36 class Node(AbstractNode):

```

(continues on next page)

(continued from previous page)

```

37 """This is a template class of how to write a node for PeekingDuck,
38 using AbstractNode as the parent class.
39 This node draws scores on objects detected.
40
41 Args:
42     config (:obj:`Dict[str, Any]` | :obj:`None`): Node configuration.
43 """
44
45 def __init__(self, config: Dict[str, Any] = None, **kwargs: Any) -> None:
46     """Node initializer
47
48     Since we do not require any special setup, it only calls the __init__
49     method of its parent class.
50     """
51     super().__init__(config, node_path=__name__, **kwargs)
52
53 def run(self, inputs: Dict[str, Any]) -> Dict[str, Any]: # type: ignore
54     """This method implements the display score function.
55     As PeekingDuck iterates through the CV pipeline, this 'run' method
56     is called at each iteration.
57
58     Args:
59         inputs (dict): Dictionary with keys "img", "bboxes", "bbox_scores"
60
61     Returns:
62         outputs (dict): Empty dictionary
63     """
64
65     # extract pipeline inputs and compute image size in (width, height)
66     img = inputs["img"]
67     bboxes = inputs["bboxes"]
68     scores = inputs["bbox_scores"]
69     img_size = (img.shape[1], img.shape[0]) # width, height
70
71     for i, bbox in enumerate(bboxes):
72         # for each bounding box:
73         # - compute (x1, y1) top-left, (x2, y2) bottom-right coordinates
74         # - convert score into a two decimal place numeric string
75         # - draw score string onto image using opencv's putText()
76         # (see opencv's API docs for more info)
77         x1, y1, x2, y2 = map_bbox_to_image_coords(bbox, img_size)
78         score = scores[i]
79         score_str = f"{score:0.2f}"
80         cv2.putText(
81             img=img,
82             text=score_str,
83             org=(x1, y2),
84             fontFace=cv2.FONT_HERSHEY_SIMPLEX,
85             fontScale=1.0,
86             color=YELLOW,
87             thickness=3,
88         )

```

(continues on next page)

(continued from previous page)

```

89
90     return {}                                # node has no outputs

```

The updated node code defines a helper function `map_bbox_to_image_coords` to map the bounding box coordinates to the image coordinates, as explained in [this section](#).

The `run` method implements the main logic which processes every bounding box to compute its on-screen coordinates and to draw the bounding box confidence score at its left-bottom position.

3. pipeline_config.yml:

pipeline_config.yml initial content:

```

1 nodes:
2   - input.visual:
3     source: https://storage.googleapis.com/peekingduck/videos/wave.mp4
4   - model.posenet
5   - draw.pose
6   - output.screen

```

This file implements the pipeline. Modify the default pipeline to the one shown below:

pipeline_config.yml updated content:

```

1 nodes:
2   - input.visual:
3     source: cat_and_computer.mp4
4   - model.yolo:
5     detect: ["cup", "cat", "laptop", "keyboard", "mouse"]
6   - draw.bbox:
7     show_labels: True
8   - custom_nodes.draw.score
9   - output.screen

```

Line 8 adds our custom node into the pipeline where it will be run by PeekingDuck during each pipeline iteration.

Execute `peekingduck run` to see your custom node in action.

Note: Royalty free video of cat and computer from: <https://www.youtube.com/watch?v=-C1TEGZavko>

3.3.2 Recipe 2: Keypoints, Count Hand Waves

This tutorial will create a custom node to analyze the skeletal keypoints of the person from the `wave.mp4` video in the [pose estimation tutorial](#) and to count the number of times the person waves his hand.

The PoseNet pose estimation model outputs seventeen keypoints for the person corresponding to the different body parts as documented [here](#). Each keypoint is a pair of (x, y) coordinates, where x and y are real numbers ranging from 0.0 to 1.0 (using *relative coordinates*).

Starting with a newly initialized PeekingDuck folder, call `peekingduck create-node` to create a new `dabble.wave` custom node as shown below:

Terminal Session

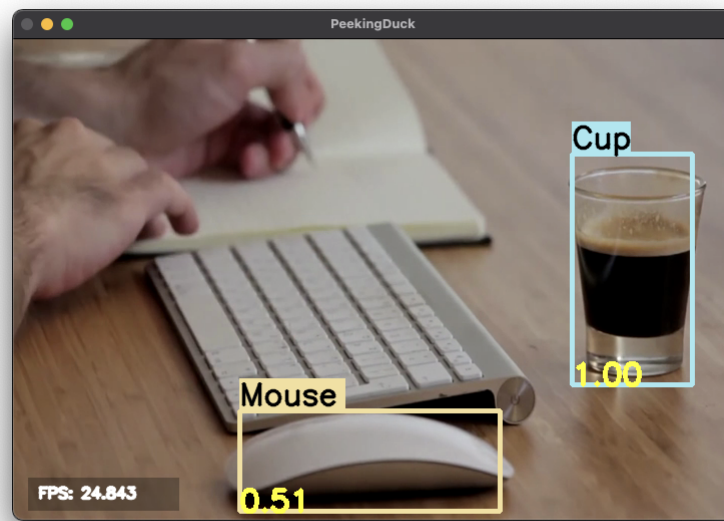


Fig. 6: Custom Node Showing Object Detection Scores

```
[~user] > mkdir wave_project
[~user] > cd wave_project
[~user/wave_project] > peekingduck init
Welcome to PeekingDuck!
2022-02-11 18:17:31 peekingduck.cli INFO: Creating custom nodes folder in ~user/wave_project/src/custom_nodes
[~user/wave_project] > peekingduck create-node
Creating new custom node...
Enter node directory relative to ~user/wave_project [src/custom_nodes]:
Select node type (input, augment, model, draw, dabble, output): dabble
Enter node name [my_custom_node]: wave

Node directory: ~user/wave_project/src/custom_nodes
Node type: dabble
Node name: wave

Creating the following files:
  Config file: ~user/wave_project/src/custom_nodes/configs/dabble/wave.yml
  Script file: ~user/wave_project/src/custom_nodes/dabble/wave.py
Proceed? [Y/n]:
Created node!
```

Also, copy `wave.mp4` into the above folder. You should end up with the following folder structure:

```
wave_project/
├── pipeline_config.yml
└── src/
```

```

├── custom_nodes/
│   ├── configs/
│   │   ├── dabble/
│   │   │   └── wave.yml
│   │   └── dabble/
│   │       └── wave.py
└── wave.mp4

```

To implement this tutorial, the **three files** `wave.yml`, `wave.py` and `pipeline_config.yml` are to be edited as follows:

1. `src/custom_nodes/configs/dabble/wave.yml`:

```

1  # Dabble node has both input and output
2  input: ["img", "bboxes", "bbox_scores", "keypoints", "keypoint_scores"]
3  output: ["none"]
4
5  # No optional configs

```

We will implement this tutorial using a custom `dabble` node, which will take the inputs `img`, `bboxes`, `bbox_scores`, `keypoints`, and `keypoint_scores` from the pipeline. The node has no output.

2. `src/custom_nodes/dabble/wave.py`:

The `dabble.wave` code structure is similar to the `draw.score` code structure in the other custom node tutorial.

Show/Hide Code for `wave.py`

```

1  """
2  Custom node to show keypoints and count the number of times the person's hand is
   ↪ waved
3  """
4
5  from typing import Any, Dict, List, Tuple
6  import cv2
7  from peekingduck.nodes.abstract_node import AbstractNode
8
9  # setup global constants
10 FONT = cv2.FONT_HERSHEY_SIMPLEX
11 WHITE = (255, 255, 255)      # opencv loads file in BGR format
12 YELLOW = (0, 255, 255)
13 THRESHOLD = 0.6              # ignore keypoints below this threshold
14 KP_RIGHT_SHOULDER = 6        # PoseNet's skeletal keypoints
15 KP_RIGHT_WRIST = 10
16
17
18 def map_bbox_to_image_coords(
19     bbox: List[float], image_size: Tuple[int, int]
20 ) -> List[int]:
21     """First helper function to convert relative bounding box coordinates to
22     absolute image coordinates.
23     Bounding box coords ranges from 0 to 1
24     where (0, 0) = image top-left, (1, 1) = image bottom-right.
25
26     Args:
27         bbox (List[float]): List of 4 floats x1, y1, x2, y2
28         image_size (Tuple[int, int]): Width, Height of image

```

(continues on next page)

(continued from previous page)

```

29
30     Returns:
31         List[int]: x1, y1, x2, y2 in integer image coords
32         """
33     width, height = image_size[0], image_size[1]
34     x1, y1, x2, y2 = bbox
35     x1 *= width
36     x2 *= width
37     y1 *= height
38     y2 *= height
39     return int(x1), int(y1), int(x2), int(y2)
40
41
42 def map_keypoint_to_image_coords(
43     keypoint: List[float], image_size: Tuple[int, int]
44 ) -> List[int]:
45     """Second helper function to convert relative keypoint coordinates to
46     absolute image coordinates.
47     Keypoint coords ranges from 0 to 1
48     where (0, 0) = image top-left, (1, 1) = image bottom-right.
49
50     Args:
51         bbox (List[float]): List of 2 floats x, y (relative)
52         image_size (Tuple[int, int]): Width, Height of image
53
54     Returns:
55         List[int]: x, y in integer image coords
56         """
57     width, height = image_size[0], image_size[1]
58     x, y = keypoint
59     x *= width
60     y *= height
61     return int(x), int(y)
62
63
64 def draw_text(img, x, y, text_str: str, color_code):
65     """Helper function to call opencv's drawing function,
66     to improve code readability in node's run() method.
67     """
68     cv2.putText(
69         img=img,
70         text=text_str,
71         org=(x, y),
72         fontFace=cv2.FONT_HERSHEY_SIMPLEX,
73         fontScale=0.4,
74         color=color_code,
75         thickness=2,
76     )
77
78
79 class Node(AbstractNode):
80     """Custom node to display keypoints and count number of hand waves

```

(continues on next page)

(continued from previous page)

```

81
82 Args:
83     config (:obj:`Dict[str, Any]` | :obj:`None`): Node configuration.
84     """
85
86 def __init__(self, config: Dict[str, Any] = None, **kwargs: Any) -> None:
87     super().__init__(config, node_path=__name__, **kwargs)
88     # setup object working variables
89     self.right_wrist = None
90     self.direction = None
91     self.num_direction_changes = 0
92     self.num_waves = 0
93
94 def run(self, inputs: Dict[str, Any]) -> Dict[str, Any]: # type: ignore
95     """This node draws keypoints and count hand waves.
96
97     Args:
98         inputs (dict): Dictionary with keys
99         "img", "bboxes", "bbox_scores", "keypoints", "keypoint_scores".
100
101     Returns:
102         outputs (dict): Empty dictionary.
103     """
104
105     # get required inputs from pipeline
106     img = inputs["img"]
107     bboxes = inputs["bboxes"]
108     bbox_scores = inputs["bbox_scores"]
109     keypoints = inputs["keypoints"]
110     keypoint_scores = inputs["keypoint_scores"]
111
112     img_size = (img.shape[1], img.shape[0]) # image width, height
113
114     # get bounding box confidence score and draw it at the
115     # left-bottom (x1, y2) corner of the bounding box (offset by 30 pixels)
116     the_bbox = bboxes[0] # image only has one person
117     the_bbox_score = bbox_scores[0] # only one set of scores
118
119     x1, y1, x2, y2 = map_bbox_to_image_coords(the_bbox, img_size)
120     score_str = f"BBox {the_bbox_score:0.2f}"
121     cv2.putText(
122         img=img,
123         text=score_str,
124         org=(x1, y2 - 30), # offset by 30 pixels
125         fontFace=cv2.FONT_HERSHEY_SIMPLEX,
126         fontScale=1.0,
127         color=WHITE,
128         thickness=3,
129     )
130
131     # hand wave detection using a simple heuristic of tracking the
132     # right wrist movement

```

(continues on next page)

(continued from previous page)

```

133 the_keypoints = keypoints[0]           # image only has one person
134 the_keypoint_scores = keypoint_scores[0] # only one set of scores
135 right_wrist = None
136 right_shoulder = None
137
138 for i, keypoints in enumerate(the_keypoints):
139     keypoint_score = the_keypoint_scores[i]
140
141     if keypoint_score >= THRESHOLD:
142         x, y = map_keypoint_to_image_coords(keypoints.tolist(), img_size)
143         x_y_str = f"({x}, {y})"
144
145         if i == KP_RIGHT_SHOULDER:
146             right_shoulder = keypoints
147             the_color = YELLOW
148         elif i == KP_RIGHT_WRIST:
149             right_wrist = keypoints
150             the_color = YELLOW
151         else:
152             # generic keypoint
153             the_color = WHITE
154
155         draw_text(img, x, y, x_y_str, the_color)
156
157 if right_wrist is not None and right_shoulder is not None:
158     # only count number of hand waves after we have gotten the
159     # skeletal poses for the right wrist and right shoulder
160     if self.right_wrist is None:
161         self.right_wrist = right_wrist           # first wrist data point
162     else:
163         # wait for wrist to be above shoulder to count hand wave
164         if right_wrist[1] > right_shoulder[1]:
165             pass
166         else:
167             if right_wrist[0] < self.right_wrist[0]:
168                 direction = "left"
169             else:
170                 direction = "right"
171
172         if self.direction is None:
173             self.direction = direction           # first direction data point
174         else:
175             # check if hand changes direction
176             if direction != self.direction:
177                 self.num_direction_changes += 1
178             # every two hand direction changes == one wave
179             if self.num_direction_changes >= 2:
180                 self.num_waves += 1
181                 self.num_direction_changes = 0   # reset direction count
182
183         self.right_wrist = right_wrist           # save last position
184         self.direction = direction

```

(continues on next page)

(continued from previous page)

```

185     wave_str = f"#waves = {self.num_waves}"
186     draw_text(img, 20, 30, wave_str, YELLOW)
187
188     return {}

```

This (long) piece of code implements our custom *dabble* node. It defines three helper functions to convert relative to absolute coordinates and to draw text on-screen. The number of hand waves is displayed at the top-left corner of the screen.

A simple heuristic is used to count the number of times the person waves his hand. It tracks the direction in which the right wrist is moving and notes when the wrist changes direction. Upon encountering two direction changes, e.g., left -> right -> left, one wave is counted.

The heuristic also waits until the right wrist has been lifted above the right shoulder before it starts tracking hand direction and counting waves.

3. pipeline_config.yml:

```

1  nodes:
2  - input.visual:
3      source: wave.mp4
4  - model.yolo
5  - model.posenet
6  - dabble.fps
7  - custom_nodes.dabble.wave
8  - draw.pose
9  - draw.legend:
10     show: ["fps"]
11 - output.screen

```

We modify `pipeline_config.yml` to run both the object detection and pose estimation models to obtain the required inputs for our custom *dabble* node.

Execute `peekingduck run` to see your custom node in action.

Note: Royalty free video of man waving from: https://www.youtube.com/watch?v=IKj_z2hgYUM

3.3.3 Recipe 3: Debugging

When working with PeekingDuck's pipeline, you may sometimes wonder what is available in the *data pool*, or whether a particular data object has been correctly computed. This tutorial will show you how to use a custom node to help with troubleshooting and debugging PeekingDuck's pipeline.

Continuing from the above tutorial, create a new `dabble.debug` custom node:

Terminal Session

```

[~user/wave_project] > peekingduck create-node
Creating new custom node...
Enter node directory relative to ~user/wave_project [src/custom_nodes]:
Select node type (input, augment, model, draw, dabble, output): dabble

```

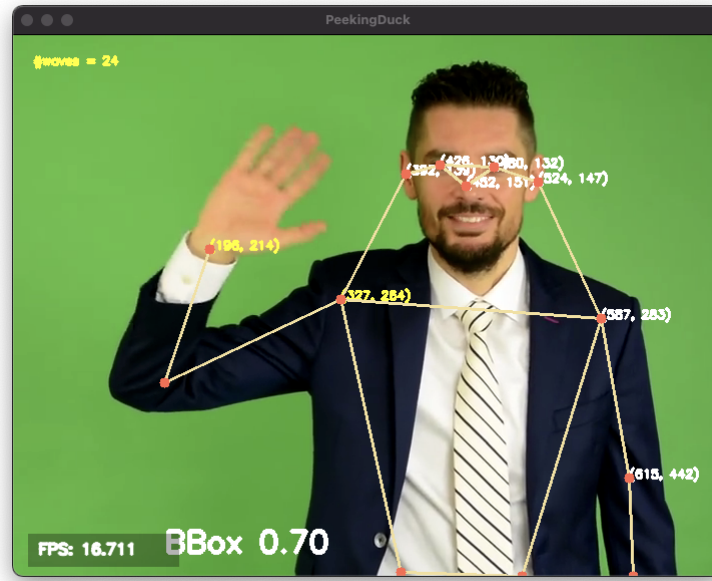


Fig. 7: Custom Node Counting Hand Waves

Enter node name [my_custom_node]: debug

Node directory: ~user/wave_project/src/custom_nodes

Node type: dabble

Node name: debug

Creating the following files:

Config file: ~user/wave_project/src/custom_nodes/configs/dabble/debug.yml

Script file: ~user/wave_project/src/custom_nodes/dabble/debug.py

Proceed? [Y/n]:

Created node!

The updated folder structure is:

```

wave_project/
├── pipeline_config.yml
├── src
│   ├── custom_nodes
│   │   ├── configs
│   │   │   └── dabble
│   │   │       ├── debug.yml
│   │   │       └── wave.yml
│   │   └── dabble
│   │       ├── debug.py
│   │       └── wave.py
└── wave.mp4
  
```

Make the following **three** changes:

1. Define `debug.yml` to receive “all” inputs from the pipeline, as follows:

```

1 # Mandatory configs
2 input: ["all"]
3 output: ["none"]
4
5 # No optional configs

```

2. Update `debug.py` as shown below:

Show/Hide Code for debug.py

```

1 """
2 A custom node for debugging
3 """
4
5 from typing import Any, Dict
6
7 from peekingduck.nodes.abstract_node import AbstractNode
8
9
10 class Node(AbstractNode):
11     """This is a simple example of creating a custom node to help with debugging.
12
13     Args:
14         config (:obj:`Dict[str, Any]` | :obj:`None`): Node configuration.
15     """
16
17     def __init__(self, config: Dict[str, Any] = None, **kwargs: Any) -> None:
18         super().__init__(config, node_path=__name__, **kwargs)
19
20     def run(self, inputs: Dict[str, Any]) -> Dict[str, Any]: # type: ignore
21         """A simple debugging custom node
22
23         Args:
24             inputs (dict): "all", to view everything in data pool
25
26         Returns:
27             outputs (dict): "none"
28         """
29
30         self.logger.info("-- debug --")
31         # show what is available in PeekingDuck's data pool
32         self.logger.info(f"input.keys={list(inputs.keys())}")
33         # debug specific data: bboxes
34         bboxes = inputs["bboxes"]
35         bbox_labels = inputs["bbox_labels"]
36         bbox_scores = inputs["bbox_scores"]
37         self.logger.info(f"num bboxes={len(bboxes)}")
38         for i, bbox in enumerate(bboxes):
39             label, score = bbox_labels[i], bbox_scores[i]
40             self.logger.info(f"bbox {i}:")
41             self.logger.info(f"  label={label}, score={score:0.2f}")
42             self.logger.info(f"  coords={bbox}")

```

(continues on next page)

(continued from previous page)

```
43  
44     return {} # no outputs
```

The custom node code shows how to see what is available in PeekingDuck's pipeline data pool by printing the input dictionary keys. It also demonstrates how to debug a specific data object, such as *bboxes*, by printing relevant information for each item within the data.

3. Update pipeline_config.yml:

```
1  nodes:  
2  - input.visual:  
3    source: wave.mp4  
4  - model.yolo  
5  - model.posenet  
6  - dabble.fps  
7  - custom_nodes.dabble.wave  
8  - custom_nodes.dabble.debug  
9  - draw.pose  
10 - draw.legend:  
11   show: ["fps"]  
12 - output.screen
```

Now, do a peekingduck run and you should see a sample debug output like the one below:

Terminal Session

```
[~user/wave_project] > peekingduck run  
2022-03-02 18:42:51 peekingduck.declarative_loader INFO: Successfully loaded pipeline_config file.  
2022-03-02 18:42:51 peekingduck.declarative_loader INFO: Initializing input.visual node...  
2022-03-02 18:42:51 peekingduck.declarative_loader INFO: Config for node input.visual is updated to: 'source':  
wave.mp4  
2022-03-02 18:42:51 peekingduck.nodes.input.visual INFO: Video/Image size: 710 by 540  
2022-03-02 18:42:51 peekingduck.nodes.input.visual INFO: Filepath used: wave.mp4  
2022-03-02 18:42:51 peekingduck.declarative_loader INFO: Initializing model.yolo node...  
[ ... many lines of output deleted here ... ]  
2022-03-02 18:42:53 peekingduck.declarative_loader INFO: Initializing custom_nodes.dabble.debug node...  
2022-03-02 18:42:53 peekingduck.declarative_loader INFO: Initializing draw.pose node...  
2022-03-02 18:42:53 peekingduck.declarative_loader INFO: Initializing draw.legend node...  
2022-03-02 18:42:53 peekingduck.declarative_loader INFO: Initializing output.screen node...  
2022-03-02 18:42:55 custom_nodes.dabble.debug INFO: - debug -  
2022-03-02 18:42:55 custom_nodes.dabble.debug INFO: input.keys=['img', 'pipeline_end', 'filename',  
'saved_video_fps', 'bboxes', 'bbox_labels', 'bbox_scores', 'keypoints', 'keypoint_scores', 'keypoint_conns',  
'hand_direction', 'num_waves', 'fps']  
2022-03-02 18:42:55 custom_nodes.dabble.debug INFO: num bboxes=1  
2022-03-02 18:42:55 custom_nodes.dabble.debug INFO: bbox 0:  
2022-03-02 18:42:55 custom_nodes.dabble.debug INFO: label=Person, score=0.91  
2022-03-02 18:42:55 custom_nodes.dabble.debug INFO: coords=[0.40047657 0.21553655 0.85199741  
1.02150181]
```

3.3.4 Other Recipes to Create Custom Nodes

This section describes two faster ways to create custom nodes for users who are familiar with PeekingDuck.

CLI Recipe

You can skip the step-by-step prompts from `peekingduck create-node` by specifying all the options on the command line, for instance:

Terminal Session

```
[~user/wave_project] > peekingduck create-node --node_subdir src/custom_nodes --node_type dabble --node_name wave
```

The above is the equivalent of the tutorial *Recipe 1: Object Detection Score* [custom node creation](#). For more information, see `peekingduck create-node --help`.

Pipeline Recipe

PeekingDuck can also create custom nodes by parsing your pipeline configuration file. Starting with the basic folder structure from `peekingduck init`:

```
wave_project/
├── pipeline_config.yml
├── src
│   ├── custom_nodes
│   │   └── configs
└── wave.mp4
```

and the following modified `pipeline_config.yml` file:

```
1 nodes:
2   - input.visual:
3     source: wave.mp4
4   - model.yolo
5   - model.posenet
6   - dabble.fps
7   - custom_nodes.dabble.wave
8   - custom_nodes.dabble.debug
9   - draw.pose
10  - draw.legend:
11    show: ["fps"]
12  - output.screen
```

You can tell PeekingDuck to parse your pipeline file with `peekingduck create-node --config_path pipeline_config.yml`:

Terminal Session

```
[~user/wave_project] > peekingduck create-node --config_path pipeline_config.yml
2022-03-14 11:21:21 peekingduck.cli INFO: Creating custom nodes declared in
~user/wave_project/pipeline_config.yml.
2022-03-14 11:21:21 peekingduck.declarative_loader INFO: Successfully loaded pipeline file.
2022-03-14 11:21:21 peekingduck.cli INFO: Creating files for custom_nodes.dabble.wave:
    Config file: ~user/wave_project/src/custom_nodes/configs/dabble/wave.yml
    Script file: ~user/wave_project/src/custom_nodes/dabble/wave.py
2022-03-14 11:21:21 peekingduck.cli INFO: Creating files for custom_nodes.dabble.debug:
    Config file: ~user/wave_project/src/custom_nodes/configs/dabble/debug.yml
    Script file: ~user/wave_project/src/custom_nodes/dabble/debug.py
```

PeekingDuck will read `pipeline_config.yml` and create the two specified custom nodes `custom_nodes.dabble.wave` and `custom_nodes.dabble.debug`. Your folder structure will now look like this:

```
wave_project/
├── pipeline_config.yml
├── src
│   └── custom_nodes
│       ├── configs
│       │   └── dabble
│       │       ├── debug.yml
│       │       └── wave.yml
│       └── dabble
│           ├── debug.py
│           └── wave.py
└── wave.mp4
```

From here, you can proceed to edit the custom node configs and source files.

3.4 Peaking Duck

PeekingDuck includes some “power” nodes that are capable of processing the contents or outputs of the other nodes and to accumulate information over time. An example is the `dabble.statistics` node which can accumulate statistical information, such as calculating the cumulative average and maximum of particular objects (like people or cars). This tutorial presents advanced recipes to showcase the power features of PeekingDuck, such as using `dabble.statistics` for object counting and tracking.

3.4.1 Interfacing with SQL

This tutorial demonstrates how to save data to an SQLite database. We will extend the tutorial for *counting hand waves* with a new custom *output* node that writes information into a local SQLite database.

Note: The above tutorial assumes `sqlite3` has been installed in your system. If your system does not have `sqlite3`, please see the [SQLite Home Page](#) for installation instructions.

First, create a new custom `output.sqlite` node in the `custom_project` folder:

Terminal Session

```
[~user/wave_project] > peekingduck create-node
Creating new custom node...
Enter node directory relative to ~user/wave_project [src/custom_nodes]:
Select node type (input, augment, model, draw, dabble, output): output
Enter node name [my_custom_node]: sqlite

Node directory: ~user/wave_project/src/custom_nodes
Node type: output
Node name: sqlite

Creating the following files:
    Config file: ~user/wave_project/src/custom_nodes/configs/output/sqlite.yml
    Script file: ~user/wave_project/src/custom_nodes/output/sqlite.py
Proceed? [Y/n]:
Created node!
```

The updated folder structure would be:

```

wave_project/
├── pipeline_config.yml
├── src/
│   ├── custom_nodes/
│   │   ├── configs/
│   │   │   ├── dabble/
│   │   │   │   └── wave.yml
│   │   │   └── output/
│   │   │       └── sqlite.yml
│   │   ├── dabble/
│   │   │   └── wave.py
│   │   └── output/
│   │       └── sqlite.py
└── wave.mp4

```

Edit the following **five** files as described below:

1. **src/custom_nodes/configs/output/sqlite.yml:**

```

1  # Mandatory configs
2  input: ["hand_direction", "num_waves"]
3  output: ["none"]
4
5  # No optional configs

```

The new `output.sqlite` custom node will take in the hand direction and the current number of hand waves to save to the external database.

2. **src/custom_nodes/output/sqlite.py:**

Show/Hide Code for `sqlite.py`

```

1  """
2  Custom node to save data to external database.

```

(continues on next page)

(continued from previous page)

```

3  """
4
5  from typing import Any, Dict
6  from datetime import datetime
7  from peekingduck.nodes.abstract_node import AbstractNode
8  import sqlite3
9
10 DB_FILE = "wave.db"           # name of database file
11
12
13 class Node(AbstractNode):
14     """Custom node to save hand direction and current wave count to database.
15
16     Args:
17         config (:obj:`Dict[str, Any]` | :obj:`None`): Node configuration.
18     """
19
20     def __init__(self, config: Dict[str, Any] = None, **kwargs: Any) -> None:
21         super().__init__(config, node_path=__name__, **kwargs)
22
23         self.conn = None
24         try:
25             # try to establish connection to database,
26             # will create DB_FILE if it does not exist
27             self.conn = sqlite3.connect(DB_FILE)
28             self.logger.info(f"Connected to {DB_FILE}")
29             sql = """ CREATE TABLE IF NOT EXISTS wavetable (
30                     datetime text,
31                     hand_direction text,
32                     wave_count integer
33                 ); """
34             cur = self.conn.cursor()
35             cur.execute(sql)
36         except sqlite3.Error as e:
37             self.logger.info(f"SQL Error: {e}")
38
39     def update_db(self, hand_direction: str, num_waves: int) -> None:
40         """Helper function to save current time stamp, hand direction and
41         wave count into DB wavetable.
42         """
43         now = datetime.now()
44         dt_str = f"{now:%Y-%m-%d %H:%M:%S}"
45         sql = """ INSERT INTO wavetable(datetime,hand_direction,wave_count)
46                 values (?, ?, ?) """
47         cur = self.conn.cursor()
48         cur.execute(sql, (dt_str, hand_direction, num_waves))
49         self.conn.commit()
50
51     def run(self, inputs: Dict[str, Any]) -> Dict[str, Any]: # type: ignore
52         """Node to output hand wave data into sqlite database.
53
54         Args:

```

(continues on next page)

(continued from previous page)

```

55         inputs (dict): Dictionary with keys "hand_direction", "num_waves"
56
57     Returns:
58         outputs (dict): Empty dictionary
59     """
60
61     hand_direction = inputs["hand_direction"]
62     num_waves = inputs["num_waves"]
63     self.update_db(hand_direction, num_waves)
64
65     return {}

```

This tutorial uses the `sqlite3` package to interface with the database.

On first run, the node initializer will create the `wave.db` database file. It will establish a connection to the database and create a table called `wavetable` if it does not exist. This table is used to store the hand direction and wave count data.

A helper function `update_db` is called to update the database. It saves the current date time stamp, hand direction and wave count into the `wavetable`.

The node's `run` method retrieves the required inputs from the pipeline's data pool and calls `self.update_db` to save the data.

3. `src/custom_nodes/configs/dabble/wave.yml`:

```

1  # Dabble node has both input and output
2  input: ["img", "bboxes", "bbox_scores", "keypoints", "keypoint_scores"]
3  output: ["hand_direction", "num_waves"]
4
5  # No optional configs

```

To support the `output.sqlite` custom node's input requirements, we need to modify the `dabble.wave` custom node to return the current hand direction `hand_direction` and the current wave count `num_waves`.

4. `src/custom_nodes/dabble/wave.py`:

```

173     ... same as previous ...
174     return {
175         "hand_direction": self.direction if self.direction is not None else "None",
176         "num_waves": self.num_waves,
177     }

```

This file is the same as the `wave.py` in the [counting hand waves](#) tutorial, except for the changes in the last few lines as shown above. These changes outputs the `hand_direction` and `num_waves` to the pipeline's data pool for subsequent consumption by the `output.sqlite` custom node.

5. `pipeline_config.yml`:

```

11     ... same as previous ...
12     - custom_nodes.output.sqlite

```

Likewise, the pipeline is the same as in the previous tutorial, except for line 12 that has been added to call the new custom node.

Run this project with `peekingduck run` and when completed, a new `wave.db` sqlite database file would be created in the current folder. Examine the created database as follows:

Terminal Session

```
[~user/wave_project] > sqlite3
SQLite version 3.37.0 2021-11-27 14:13:22
Enter ".help" for usage hints.
Connected to a transient in-memory database.
Use ".open FILENAME" to reopen on a persistent database.
sqlite> .open wave.db
sqlite> .schema wavetable
CREATE TABLE wavetable (
    datetime text,
    hand_direction text,
    wave_count integer
);
sqlite> select * from wavetable where wave_count > 0 limit 5;
2022-02-15 19:26:16|left|1
2022-02-15 19:26:16|right|1
2022-02-15 19:26:16|left|2
2022-02-15 19:26:16|right|2
2022-02-15 19:26:16|right|2
sqlite> select * from wavetable order by datetime desc limit 5;
2022-02-15 19:26:44|right|72
2022-02-15 19:26:44|right|72
2022-02-15 19:26:44|right|72
2022-02-15 19:26:44|right|72
2022-02-15 19:26:43|right|70
```

Press CTRL-D to exit from `sqlite3`.

3.4.2 Counting Cars

This tutorial demonstrates using the `dabble.statistics` node to count the number of cars traveling across a highway over time and the `draw.legend` node to display the relevant statistics.

Create a new PeekingDuck project, download the [highway cars video](#) and save it into the project folder.

Terminal Session

```
[~user] > mkdir car_project
[~user] > cd car_project
[~user/car_project] > peekingduck init
```

The car_project folder structure:

```
car_project/
├── highway_cars.mp4
├── pipeline_config.yml
├── src
│   ├── custom_nodes
│   └── configs
```

Edit pipeline_config.yml as follows:

```
1 nodes:
2 - input.visual:
3     source: highway_cars.mp4
4 - model.yolo:
5     detect: ["car"]
6     score_threshold: 0.3
7 - dabble.bbox_count
8 - dabble.fps
9 - dabble.statistics:
10     identity: count
11 - draw.bbox
12 - draw.legend:
13     show: ["fps", "count", "cum_max", "cum_min"]
14 - output.screen
```

Run it with `peekingduck run` and you should see a video of cars travelling across a highway with a legend box on the bottom left showing the realtime count of the number of cars on-screen, the cumulative maximum and minimum number of cars detected since the video started. The sample screenshot below shows:

- the count that there are currently 3 cars on-screen
- the cumulative maximum number of cars “seen” previously was 5
- the cumulative minimum number of cars was 1



Fig. 8: Counting Cars on a Highway

Note: Royalty free video of cars on highway from: <https://www.youtube.com/watch?v=8yP1gJg4b2w>

3.4.3 Object Tracking

Object tracking is the application of CV models to automatically detect objects in a video and to assign a unique identity to each of them. These objects can be either living (e.g. person) or non-living (e.g. car). As they move around in the video, these objects are identified based on their assigned identities and tracked according to their movements.

This tutorial demonstrates using `dabble.statistics` with a custom node to track the number of people walking down a path.

Create a new PeekingDuck project, download the [people walking video](#) and save it into the project folder.

Terminal Session

```
[~user] > mkdir people_walking
[~user] > cd people_walking
[~user/people_walking] > peekingduck init
```

Create the following `pipeline_config.yml`:

```
1 nodes:
2   - input.visual:
3     source: people_walking.mp4
4   - model.yolo:
5     detect: ["person"]
6   - dabble.tracking
7   - dabble.statistics:
8     maximum: obj_attrs["ids"]
9   - dabble.fps
10  - draw.bbox
11  - draw.tag:
12    show: ["ids"]
13  - draw.legend:
14    show: ["fps", "cum_max", "cum_min", "cum_avg"]
15  - output.screen
```

The above pipeline uses the YOLO model to detect people in the video and uses the `dabble.tracking` node to track the people as they walk. Each person is assigned a tracking ID and `dabble.tracking` returns a list of tracking IDs. `dabble.statistics` is used to process these tracking IDs: since each person is assigned a monotonically increasing integer ID, the maximum ID within the list tells us the number of persons tracked so far. `draw.tag` shows the ID above the tracked person. `draw.legend` is used to display the various statistics: the FPS, and the cumulative maximum, minimum and average relating to the number of persons tracked.

Do a peekingduck run and you will see the following display:

Note: Royalty free video of people walking from: <https://www.youtube.com/watch?v=du74nvmRUzo>



Fig. 9: People Walking

Tracking People within a Zone

Suppose we are only interested in people walking down the center of the video (imagine a carpet running down the middle). We can create a custom node to tell PeekingDuck to focus on the middle zone, by filtering away the detected bounding boxes outside the zone.

Start by creating a custom node `dabble.filter_bbox`:

Terminal Session

```
[~user/people_walking] > peekingduck create-node
Creating new custom node...
Enter node directory relative to ~user/people_walking [src/custom_nodes]:
Select node type (input, augment, model, draw, dabble, output): dabble
Enter node name [my_custom_node]: filter_bbox

Node directory: ~user/people_walking/src/custom_nodes
Node type: dabble
Node name: filter_bbox

Creating the following files:
  Config file: ~user/people_walking/src/custom_nodes/configs/dabble/filter_bbox.yml
  Script file: ~user/people_walking/src/custom_nodes/dabble/filter_bbox.py
Proceed? [Y/n]:
Created node!
```

The folder structure looks like this:

```
people_walking/
├── people_walking.mp4
├── pipeline_config.yml
└── src
    ├── custom_nodes
    │   ├── configs
    │   │   └── dabble
    │   │       └── filter_bbox.yml
    │   └── dabble
    │       └── filter_bbox.py
```

Change pipeline_config.yml to the following:

```
1 nodes:
2 - input.visual:
3     source: people_walking.mp4
4 - model.yolo:
5     detect: ["person"]
6 - dabble.bbox_to_btm_midpoint
7 - dabble.zone_count:
8     resolution: [720, 480]
9     zones: [
10         [[0.35,0], [0.65,0], [0.65,1], [0.35,1]],
11     ]
12 - custom_nodes.dabble.filter_bbox:
13     zones: [
14         [[0.35,0], [0.65,0], [0.65,1], [0.35,1]],
15     ]
16 - dabble.tracking
17 - dabble.statistics:
18     maximum: obj_attrs["ids"]
19 - dabble.fps
20 - draw.bbox
21 - draw.zone
22 - draw.tag:
23     show: ["ids"]
24 - draw.legend:
25     show: ["fps", "cum_max", "cum_min", "cum_avg", "zone_count"]
26 - output.screen
```

We make use of `dabble.zone_count` and `dabble.bbox_to_btm_midpoint` nodes to create a zone in the middle. The zone is defined by a rectangle with the four corners (0.35, 0.0) - (0.65, 0.0) - (0.65, 1.0) - (0.35, 1.0). (For more info, see [Zone Counting](#)) This zone is also defined in our custom node `dabble.filter_bbox` for bounding box filtering. What `dabble.filter_bbox` will do is to take the list of bboxes as input and output a list of bboxes within the zone, dropping all bboxes outside it. Then, `dabble.tracking` is used to track the people walking and `dabble.statistics` is used to determine the number of people walking in the zone, by getting the maximum of the tracked IDs. `draw.legend` has a new item `zone_count` which displays the number of people walking in the zone currently.

The `filter_bbox.yml` and `filter_bbox.py` files are shown below:

1. `src/custom_nodes/configs/dabble/filter_bbox.yml`:

```
1 # Mandatory configs
2 input: ["bboxes"]
3 output: ["bboxes"]
```

(continues on next page)

(continued from previous page)

```

4
5 zones: [
6     [[0,0], [0,1], [1,1], [1,0]],
7 ]

```

Note: The zones default value of `[[0,0], [0,1], [1,1], [1,0]]` will be overridden by those specified in `pipeline_config.yml` above. See [Configuration - Behind The Scenes](#) for more details.

2. `src/custom_nodes/dabble/filter_bbox.py`:

Show/Hide Code for `filter_bbox.py`

```

1  """
2  Custom node to filter bboxes outside a zone
3  """
4
5  from typing import Any, Dict
6  import numpy as np
7  from peekingduck.nodes.abstract_node import AbstractNode
8
9
10 class Node(AbstractNode):
11     """Custom node to filter bboxes outside a zone
12
13     Args:
14         config (:obj:`Dict[str, Any]` | :obj:`None`): Node configuration.
15     """
16
17     def __init__(self, config: Dict[str, Any] = None, **kwargs: Any) -> None:
18         super().__init__(config, node_path=__name__, **kwargs)
19
20     def run(self, inputs: Dict[str, Any]) -> Dict[str, Any]: # type: ignore
21         """Checks bounding box x-coordinates against the zone left and right borders.
22         Retain bounding box if within, otherwise discard it.
23
24         Args:
25             inputs (dict): Dictionary with keys "bboxes"
26
27         Returns:
28             outputs (dict): Dictionary with keys "bboxes".
29         """
30         bboxes = inputs["bboxes"]
31         zones = self.config["zones"]
32         zone = zones[0] # only work with one zone currently
33         # convert zone with 4 points to a zone bbox with (x1, y1), (x2, y2)
34         x1, y1 = zone[0]
35         x2, y2 = zone[2]
36         zone_bbox = np.asarray([x1, y1, x2, y2])
37
38         retained_bboxes = []
39         for bbox in bboxes:

```

(continues on next page)

(continued from previous page)

```

40     # filter by left and right borders (ignore top and bottom)
41     if bbox[0] > zone_bbox[0] and bbox[2] < zone_bbox[2]:
42         retained_bboxes.append(bbox)
43
44     return {"bboxes": np.asarray(retained_bboxes)}

```

Do a peekingduck run and you will see the following display:



Fig. 10: Count People Walking in a Zone

3.4.4 Using Callbacks

While PeekingDuck allows you to extend its capabilities through custom nodes, it may not be the most optimal/efficient approach in some cases. For example, if you are trying to bring in PeekingDuck into an existing codebase/project, it may not be desirable to refactor the existing codebase such that it can be constructed/initialized in a custom node.

Introduction to Using Callbacks in PeekingDuck

Callbacks are externally defined functions/methods which are invoked as specific points of a PeekingDuck pipeline and can interact with code/systems outside of PeekingDuck. They can be registered to nodes in a pipeline using the `callbacks` config option.

Pipeline Events

The following pipeline events and their respective event keys are implemented:

1. `run_begin`: Invoked before a Node's `run()` method.
2. `run_end`: Invoked after a Nodes's `run()` method.

For example, in the following pipeline config:

```

1 nodes:
2   - input.visual
3   - model.posenet:
4     callbacks:
5       run_begin: [<callback 1>]
6       run_end: [<callback 2>, <callback 3>, <callback 4>]
7   - draw.pose
8   - output.screen

```

After a image frame is read by `input.visual`, `<callback 1>` is invoked, followed by `model.posenet`'s `run()` method, then `<callback 2>`, `<callback 3>`, and `<callback 4>` are invoked after that.

Callback Definition

Callback definition, i.e. `<callback 1>` and `<callback 2>` in the previous example, follows a strict format. The following definition has been written using the [EBNF notation](#):

```

callback definition = module part, ":", callback part
module part         = { subdir name, "." }, script name
callback part       = function
                      | class name, ":", method
                      | instance name, ":", instance method

script name         = ? name of callback Python script ?
subdir name         = ? name of individual subdirectory folders relative to
↳ "callbacks" folder ?
class name          = ? name of class which contains the callback methods ?
instance name       = ? name of instantiated object which contains the callback
↳ method ?
function            = ? name of callback function ?
method              = ? name of class/static methods ?
instance method     = ? name of instance method ?

```

When using `<instance name>::<instance method>` for callback part, the instance object has to be in the script's global scope.

The callbacks directory, which houses the scripts containing the callback functions, is expected to be at the same location of the pipeline config file. For example, given the following directory structure and file contents:

```

callbacks_project/
├── pipeline_config.yml
├── callbacks/
│   ├── my_callback.py
│   └── cb_dir/
│       └── my_cb.py

```

1. callbacks/my_callback.py

Show/Hide Code for my_callback.py

```

1  def general_function(data_pool):
2      print("Called general_function")
3      filename = data_pool["filename"]
4      print(f"  filename={filename}")
5
6
7  class MyCallbackClass:
8      @staticmethod
9      def the_static_method(data_pool):
10         print("Called the_static_method")
11
12     @classmethod
13     def the_class_method(cls, data_pool):
14         print("Called the_class_method")
15
16     def the_instance_method(self, data_pool):
17         print("Called the_instance_method")
18         kp_scores = data_pool["keypoint_scores"][0]
19         print(f"  num_keypoint_scores={len(kp_scores)}")
20
21
22     # my_callback_obj has to be visible from my_callback.py's global scope
23     # to enable access to the_instance_method()
24     my_callback_obj = MyCallbackClass()

```

2. callbacks/cb_dir/my_cb.py

Show/Hide Code for my_cb.py

```

1  def my_cb_function(data_pool):
2      print("Called my_cb_function")

```

The various callback functions and methods can be registered in pipeline_config.yml as follows:

```

1  nodes:
2  - input.visual:
3      source: https://storage.googleapis.com/peekingduck/videos/wave.mp4
4  - model.posenet:
5      callbacks:
6          run_begin:
7              [
8                  "my_callback::MyCallbackClass::the_static_method",
9                  "my_callback::MyCallbackClass::the_class_method",
10             ]

```

(continues on next page)

(continued from previous page)

```

11     run_end:
12     [
13         "my_callback::general_function",
14         "my_callback::my_callback_obj::the_static_method",
15         "my_callback::my_callback_obj::the_class_method",
16         "my_callback::my_callback_obj::the_instance_method",
17         "cb_dir.my_cb::my_cb_function",
18     ]
19 - draw.pose
20 - output.screen

```

The following table illustrates how the callback definitions map to their respective callback functions:

Definition	Function	Script location
<code>my_callback::general_function</code>	<code>general_function()</code>	callbacks/ my_callback.py
<code>my_callback::MyCallbackClass::the_static_method</code>	<code>MyCallbackClass.the_static_method()</code>	
<code>my_callback::my_callback_obj::the_static_method</code>	<code>my_callback_obj.the_static_method()</code>	
<code>my_callback::MyCallbackClass::the_class_method</code>	<code>MyCallbackClass.the_class_method()</code>	
<code>my_callback::my_callback_obj::the_class_method</code>	<code>my_callback_obj.the_class_method()</code>	
<code>my_callback::my_callback_obj::the_instance_method</code>	<code>my_callback_obj.the_instance_method()</code>	callbacks/cb_dir/ my_cb.py
<code>cb_dir.my_cb::my_cb_function</code>	<code>my_cb_function()</code>	

Running the pipeline with peekingduck run should give the following output:

Terminal Session

[Truncated]

```

2022-12-01 17:33:23 peekingduck.declarative_loader INFO: Config for node model.posenet is updated
to: 'callbacks': {'run_begin': ['my_callback::MyCallbackClass::the_static_method',
'my_callback::MyCallbackClass::the_class_method'], 'run_end': ['my_callback::general_function',
'my_callback::my_callback_obj::the_static_method',
'my_callback::my_callback_obj::the_class_method',
'my_callback::my_callback_obj::the_instance_method', 'cb_dir.my_cb::my_cb_function']}
2022-12-01 17:33:23 peekingduck.nodes.model.posenetv1.posenet_files.predictor INFO: PoseNet model
loaded with following configs:
    Model type: resnet,
    Input resolution: (225, 225),
    Max pose detection: 10,
    Score threshold: 0.4
2022-12-01 17:33:26 peekingduck.declarative_loader INFO: Initializing draw.pose node...
2022-12-01 17:33:26 peekingduck.declarative_loader INFO: Initializing output.screen node...
Called the_static_method
Called the_class_method
Called general_function
    filename=video.mp4
Called the_static_method

```

```
Called the _class_method
Called the _instance_method
    num_keypoint_scores=17
Called my_cb_function

[Truncated]
```

Interfacing with SQL Using Callbacks

This tutorial replicates the earlier *Interfacing with SQL* tutorial but uses callbacks instead of custom nodes.

Note: The above tutorial assumes `sqlite3` has been installed in your system. If your system does not have `sqlite3`, please see the [SQLite Home Page](#) for installation instructions.

The following steps are necessary to prepare the files and folder structure required for this tutorial:

1. Create and initialize a new PeekingDuck project in a folder named `callbacks_wave_project`.
2. Create a `dabble.wave` custom node.
3. Create a `sql_callback.py` in the `callbacks/` folder.
4. Copy `wave.mp4` into the folder.

You should end up with the following folder structure:

```
callbacks_wave_project/
├── pipeline_config.yml
├── callbacks/
│   └── sql_callback.py
├── src/
│   ├── custom_nodes/
│   │   ├── configs/
│   │   │   └── dabble/
│   │   │       └── wave.yml
│   │   └── dabble/
│   │       └── wave.py
└── wave.mp4
```

Edit the following **files** as described below:

1. `src/custom_nodes/configs/dabble/wave.yml`:

```
1 # Dabble node has both input and output
2 input: ["img", "bboxes", "bbox_scores", "keypoints", "keypoint_scores"]
3 output: ["hand_direction", "num_waves"]
4
5 callbacks: {}
```

The additional `callbacks: {}` config allows us to override it later in `pipeline_config.yml` to define the callback functions to use.

2. `src/custom_nodes/dabble/wave.py`:

Show/Hide Code for wave.py

```

1  """
2  Custom node to show keypoints and count the number of times the person's hand is_
   ↪waved
3  """
4
5  from typing import Any, Dict, List, Tuple
6
7  import cv2
8  from peekingduck.pipeline.nodes.abstract_node import AbstractNode
9
10 # setup global constants
11 FONT = cv2.FONT_HERSHEY_SIMPLEX
12 WHITE = (255, 255, 255) # opencv loads file in BGR format
13 YELLOW = (0, 255, 255)
14 THRESHOLD = 0.6 # ignore keypoints below this threshold
15 KP_RIGHT_SHOULDER = 6 # PoseNet's skeletal keypoints
16 KP_RIGHT_WRIST = 10
17
18
19 def map_bbox_to_image_coords(
20     bbox: List[float], image_size: Tuple[int, int]
21 ) -> Tuple[int, ...]:
22     """First helper function to convert relative bounding box coordinates to
23     absolute image coordinates.
24     Bounding box coords ranges from 0 to 1
25     where (0, 0) = image top-left, (1, 1) = image bottom-right.
26
27     Args:
28         bbox (List[float]): List of 4 floats x1, y1, x2, y2
29         image_size (Tuple[int, int]): Width, Height of image
30
31     Returns:
32         Tuple[int, ...]: x1, y1, x2, y2 in integer image coords
33     """
34     width, height = image_size[0], image_size[1]
35     x1, y1, x2, y2 = bbox
36     x1 *= width
37     x2 *= width
38     y1 *= height
39     y2 *= height
40     return int(x1), int(y1), int(x2), int(y2)
41
42 def map_keypoint_to_image_coords(
43     keypoint: List[float], image_size: Tuple[int, int]
44 ) -> Tuple[int, ...]:
45     """Second helper function to convert relative keypoint coordinates to
46     absolute image coordinates.
47     Keypoint coords ranges from 0 to 1
48     where (0, 0) = image top-left, (1, 1) = image bottom-right.
49
50     Args:

```

(continues on next page)

(continued from previous page)

```

51         bbox (List[float]): List of 2 floats x, y (relative)
52         image_size (Tuple[int, int]): Width, Height of image
53
54     Returns:
55         Tuple[int, ...]: x, y in integer image coords
56     """
57     width, height = image_size[0], image_size[1]
58     x, y = keypoint
59     x *= width
60     y *= height
61     return int(x), int(y)
62
63 def draw_text(img, x, y, text_str: str, color_code):
64     """Helper function to call opencv's drawing function,
65     to improve code readability in node's run() method.
66     """
67     cv2.putText(
68         img=img,
69         text=text_str,
70         org=(x, y),
71         fontFace=cv2.FONT_HERSHEY_SIMPLEX,
72         fontScale=0.4,
73         color=color_code,
74         thickness=2,
75     )
76
77 class Node(AbstractNode):
78     """Custom node to display keypoints and count number of hand waves
79
80     Args:
81         config (:obj:`Dict[str, Any]` | :obj:`None`): Node configuration.
82     """
83
84     def __init__(self, config: Dict[str, Any] = None, **kwargs: Any) -> None:
85         super().__init__(config, node_path=__name__, **kwargs)
86         # setup object working variables
87         self.right_wrist = None
88         self.direction = None
89         self.num_direction_changes = 0
90         self.num_waves = 0
91
92     def run(self, inputs: Dict[str, Any]) -> Dict[str, Any]: # type: ignore
93         """This node draws keypoints and count hand waves.
94
95         Args:
96             inputs (dict): Dictionary with keys
97                 "img", "bboxes", "bbox_scores", "keypoints", "keypoint_scores".
98
99         Returns:
100             outputs (dict): Empty dictionary.
101         """
102         # get required inputs from pipeline

```

(continues on next page)

(continued from previous page)

```

103     img = inputs["img"]
104     bboxes = inputs["bboxes"]
105     bbox_scores = inputs["bbox_scores"]
106     keypoints = inputs["keypoints"]
107     keypoint_scores = inputs["keypoint_scores"]
108
109     img_size = (img.shape[1], img.shape[0]) # image width, height
110
111     # get bounding box confidence score and draw it at the
112     # left-bottom (x1, y2) corner of the bounding box (offset by 30 pixels)
113     the_bbox = bboxes[0] # image only has one person
114     the_bbox_score = bbox_scores[0] # only one set of scores
115
116     x1, y1, x2, y2 = map_bbox_to_image_coords(the_bbox, img_size)
117     score_str = f"BBox {the_bbox_score:0.2f}"
118     cv2.putText(
119         img=img,
120         text=score_str,
121         org=(x1, y2 - 30), # offset by 30 pixels
122         fontFace=cv2.FONT_HERSHEY_SIMPLEX,
123         fontScale=1.0,
124         color=WHITE,
125         thickness=3,
126     )
127
128     # hand wave detection using a simple heuristic of tracking the
129     # right wrist movement
130     the_keypoints = keypoints[0] # image only has one person
131     the_keypoint_scores = keypoint_scores[0] # only one set of scores
132     right_wrist = None
133     right_shoulder = None
134
135     for i, keypoints in enumerate(the_keypoints):
136         keypoint_score = the_keypoint_scores[i]
137
138         if keypoint_score >= THRESHOLD:
139             x, y = map_keypoint_to_image_coords(keypoints.tolist(), img_size)
140             x_y_str = f"({x}, {y})"
141
142             if i == KP_RIGHT_SHOULDER:
143                 right_shoulder = keypoints
144                 the_color = YELLOW
145             elif i == KP_RIGHT_WRIST:
146                 right_wrist = keypoints
147                 the_color = YELLOW
148             else: # generic keypoint
149                 the_color = WHITE
150
151             draw_text(img, x, y, x_y_str, the_color)
152
153     if right_wrist is not None and right_shoulder is not None:
154         # only count number of hand waves after we have gotten the

```

(continues on next page)

(continued from previous page)

```

155     # skeletal poses for the right wrist and right shoulder
156     if self.right_wrist is None:
157         self.right_wrist = right_wrist # first wrist data point
158     else:
159         # wait for wrist to be above shoulder to count hand wave
160         if right_wrist[1] > right_shoulder[1]:
161             pass
162         else:
163             if right_wrist[0] < self.right_wrist[0]:
164                 direction = "left"
165             else:
166                 direction = "right"
167
168         if self.direction is None:
169             self.direction = direction # first direction data point
170         else:
171             # check if hand changes direction
172             if direction != self.direction:
173                 self.num_direction_changes += 1
174             # every two hand direction changes == one wave
175             if self.num_direction_changes >= 2:
176                 self.num_waves += 1
177                 self.num_direction_changes = 0 # reset direction count
178
179         self.right_wrist = right_wrist # save last position
180         self.direction = direction
181
182     wave_str = f"#waves = {self.num_waves}"
183     draw_text(img, 20, 30, wave_str, YELLOW)
184
185     return {
186         "hand_direction": self.direction if self.direction is not None else
↪ "None",
187         "num_waves": self.num_waves,
188     }

```

Note that we still use a custom node for the “wave” logic as we need to produce custom outputs such as `hand_direction` and `num_waves`.

3. callbacks/sql_callback.py:

Show/Hide Code for `sql_callback.py`

```

1  """
2  Callbacks for interacting with an external database.
3  """
4
5  import logging
6  import sqlite3
7  from datetime import datetime
8  from typing import Any, Dict
9
10 DB_FILE = "wave.db" # name of database file

```

(continues on next page)

(continued from previous page)

```

11
12
13 class SQLCallback:
14     """Interacts with a SQLite database and provides a callback method to
15     trigger updates.
16     """
17
18     def __init__(self) -> None:
19         self.logger = logging.getLogger(__name__)
20         self.conn = None
21         try:
22             self.conn = sqlite3.connect(DB_FILE)
23             self.logger.info(f"Connected to {DB_FILE}")
24             sql = """ CREATE TABLE IF NOT EXISTS wavetable (
25                     datetime text,
26                     hand_direction text,
27                     wave_count integer
28                 ); """
29             cur = self.conn.cursor()
30             cur.execute(sql)
31         except sqlite3.Error as e:
32             self.logger.info(f"SQL Error: {e}")
33
34     def track_wave_callback(self, data_pool: Dict[str, Any]) -> None:
35         hand_direction = data_pool["hand_direction"]
36         num_waves = data_pool["num_waves"]
37         self.update_db(hand_direction, num_waves)
38
39     def update_db(self, hand_direction: str, num_waves: int) -> None:
40         """Helper function to save current time stamp, hand direction and
41         wave count into DB wavetable.
42         """
43         now = datetime.now()
44         dt_str = f"{now:%Y-%m-%d %H:%M:%S}"
45         sql = """ INSERT INTO wavetable(datetime,hand_direction,wave_count)
46                 values (?,?,?) """
47         cur = self.conn.cursor()
48         cur.execute(sql, (dt_str, hand_direction, num_waves))
49         self.conn.commit()
50
51
52 sql_callback_obj = SQLCallback()

```

We have recreated the `output.sqlite` node into a standalone class. As it no longer inherits from `AbstractNode`, having a `run()` method is not necessary anymore. Instead, we implement a `track_wave_callback()` method which will be used update the database with the require information. As `track_wave_callback()` is an instance method, we need to instantiate `SQLCallback` as `sql_callback_obj` in order to use it.

4. `pipeline_config.yml`:

```

1 nodes:
2 - input.visual:

```

(continues on next page)

(continued from previous page)

```
3     source: wave.mp4
4 - model.yolo
5 - model.posenet
6 - custom_nodes.dabble.wave:
7     callbacks:
8         run_end: ["sql_callback::sql_callback_obj::track_wave_callback"]
9 - draw.pose
10 - output.screen
```

To use the callback function, we override the `callbacks` config option in `custom_nodes.dabble.wave`. The key `run_end` specifies that the callback function will be invoked after the node's `run()` method. To specify the callback function to be invoked we format the string as `<script name>::<instance name>::<instance method name>`.

Run this project with `peekingduck run`, and when completed, a `wave.db` database file will be generated. We can examine the created database using similar commands as the *Interfacing with SQL* tutorial:

Terminal Session

```
[~user/wave_project] > sqlite3
SQLite version 3.39.2 2022-07-21 15:24:47
Enter ".help" for usage hints.
Connected to a transient in-memory database.
Use ".open FILENAME" to reopen on a persistent database.
@sqlite> .open wave.db
@sqlite> .schema wavetable
CREATE TABLE wavetable (
    datetime text,
    hand_direction text,
    wave_count integer
);
@sqlite> select * from wavetable where wave_count > 0 limit 5;
2022-11-08 13:30:06|left|1
2022-11-08 13:30:06|right|1
2022-11-08 13:30:06|left|2
2022-11-08 13:30:06|left|2
2022-11-08 13:30:06|left|2
@sqlite> select * from wavetable order by datetime desc limit 5;
2022-11-08 13:30:36|right|70
2022-11-08 13:30:36|left|71
2022-11-08 13:30:35|left|69
2022-11-08 13:30:35|left|69
2022-11-08 13:30:35|left|69
```

Press CTRL-D to exit from `sqlite3`.

3.5 Calling PeekingDuck in Python

3.5.1 Using PeekingDuck's Pipeline

As an alternative to running PeekingDuck using the command-line interface (CLI), users can also import PeekingDuck as a Python module and run it in a Python script. This demo corresponds to the *Record Video File with FPS* Section of the *Duck Confit* tutorial.

In addition, we will demonstrate basic debugging techniques which users can employ when troubleshooting PeekingDuck projects.

Setting Up

Create a PeekingDuck project using:

Terminal Session

```
[~user] > mkdir pkd_project
[~user] > cd pkd_project
[~user/pkd_project] > peekingduck init
```

Then, download the [cat and computer video](#) to the `pkd_project` folder and create a Python script `demo_debug.py` in the same folder.

You should have the following directory structure at this point:

```
pkd_project/
├── cat_and_computer.mp4
├── demo_debug.py
├── pipeline_config.yml
└── src/
```

Creating a Custom Node for Debugging

Run the following to create a [dabble](#) node for debugging:

Terminal Session

```
[~user/pkd_project] > peekingduck create-node --node_subdir src/custom_nodes --node_type dabble --node_name debug
```

The command will create the `debug.py` and `debug.yml` files in your project directory as shown:

```
pkd_project/
├── cat_and_computer.mp4
├── demo_debug.py
└──
```

```
├── pipeline_config.yml
├── src/
│   └── custom_nodes/
│       ├── configs/
│       │   └── dabble/
│       │       └── debug.yml
│       └── dabble/
│           └── debug.py
```

Change the content of debug.yml to:

```
1 input: ["all"]
2 output: ["none"]
```

Line 1: The data type `all` allows the node to receive all outputs from the previous nodes as its input. Please see the [Glossary](#) for a list of available data types.

Change the content of debug.py to:

Show/Hide Code

```
1 from typing import Any, Dict
2
3 import numpy as np
4
5 from peekingduck.nodes.abstract_node import AbstractNode
6
7
8 class Node(AbstractNode):
9     def __init__(self, config: Dict[str, Any] = None, **kwargs: Any) -> None:
10         super().__init__(config, node_path=__name__, **kwargs)
11         self.frame = 0
12
13     def run(self, inputs: Dict[str, Any]) -> Dict[str, Any]: # type: ignore
14         if "cat" in inputs["bbox_labels"]:
15             print(
16                 f"{self.frame} {inputs['bbox_scores'][np.where(inputs['bbox_labels'] ==
17                 ↪ 'cat')]}]"
18             )
19             self.frame += 1
20         return {}
```

Lines 14 - 17: Print out the frame number and the confidence scores of bounding boxes which are detected as “cat”.

Line 18: Increment the frame number each time `run()` is called.

Creating the Python Script

Copy over the following code to `demo_debug.py`:

Show/Hide Code

```

1 from pathlib import Path
2
3 from peekingduck.nodes.dabble import fps
4 from peekingduck.nodes.draw import bbox, legend
5 from peekingduck.nodes.input import visual
6 from peekingduck.nodes.model import yolo
7 from peekingduck.nodes.output import media_writer, screen
8 from peekingduck.runner import Runner
9 from src.custom_nodes.dabble import debug
10
11
12 def main():
13     debug_node = debug.Node(pkd_base_dir=Path.cwd() / "src" / "custom_nodes")
14
15     visual_node = visual.Node(source=str(Path.cwd() / "cat_and_computer.mp4"))
16     yolo_node = yolo.Node(detect=["cup", "cat", "laptop", "keyboard", "mouse"])
17     bbox_node = bbox.Node(show_labels=True)
18
19     fps_node = fps.Node()
20     legend_node = legend.Node(show=["fps"])
21     screen_node = screen.Node()
22
23     media_writer_node = media_writer.Node(output_dir=str(Path.cwd() / "results"))
24
25     runner = Runner(
26         nodes=[
27             visual_node,
28             yolo_node,
29             debug_node,
30             bbox_node,
31             fps_node,
32             legend_node,
33             screen_node,
34             media_writer_node,
35         ]
36     )
37     runner.run()
38
39
40 if __name__ == "__main__":
41     main()

```

Lines 9, 13: Import and initialize the debug custom node. Pass in the path/to/project_dir/src/custom_nodes via `pkd_base_dir` for the configuration YAML file of the custom node to be loaded properly.

Lines 15 - 23: Create the PeekingDuck nodes necessary to replicate the demo shown in the [Record Video File with FPS](#) tutorial. To change the node configuration, you can pass the new values to the `Node()` constructor as keyword arguments.

Lines 25 - 37: Initialize the PeekingDuck Runner from `runner.py` with the list of nodes passed in via the `nodes`

argument.

Note: A PeekingDuck node can be created in Python code by passing a dictionary of config keyword - config value pairs to the `Node()` constructor.

Running the Python Script

Run the `demo_debug.py` script using:

Terminal Session

[~user/pkd_project] > python demo_debug.py

You should see the following output in your terminal:

```
1 2022-02-24 16:33:06 peekingduck.nodes.input.visual INFO: Config for node input.visual_
   ↳ is updated to: 'source': ~user/pkd_project/cat_and_computer.mp4
2 2022-02-24 16:33:06 peekingduck.nodes.input.visual INFO: Video/Image size: 720 by 480
3 2022-02-24 16:33:06 peekingduck.nodes.input.visual INFO: Filepath used: ~user/pkd_
   ↳ project/cat_and_computer.mp4
4 2022-02-24 16:33:06 peekingduck.nodes.model.yolo INFO: Config for node model.yolo is_
   ↳ updated to: 'detect': [41, 15, 63, 66, 64]
5 2022-02-24 16:33:06 peekingduck.nodes.model.yolov4.yolo_files.detector INFO: Yolo_
   ↳ model loaded with following configs:
6     Model type: v4tiny,
7     Input resolution: 416,
8     IDs being detected: [41, 15, 63, 66, 64]
9     Max Detections per class: 50,
10    Max Total Detections: 50,
11    IOU threshold: 0.5,
12    Score threshold: 0.2
13 2022-02-24 16:33:07 peekingduck.nodes.draw.bbox INFO: Config for node draw.bbox is_
   ↳ updated to: 'show_labels': True
14 2022-02-24 16:33:07 peekingduck.nodes.dabble.fps INFO: Moving average of FPS will be_
   ↳ logged every: 100 frames
15 2022-02-24 16:33:07 peekingduck.nodes.output.media_writer INFO: Config for node output.
   ↳ media_writer is updated to: 'output_dir': ~user/pkd_project/results
16 2022-02-24 16:33:07 peekingduck.nodes.output.media_writer INFO: Output directory used_
   ↳ is: ~user/pkd_project/results
17 0 [0.90861976]
18 1 [0.9082737]
19 2 [0.90818006]
20 3 [0.8888804]
21 4 [0.8877487]
22 5 [0.9071386]
23 6 [0.870267]
24
25 [Truncated]
```

Lines 17 - 23: The debugging output showing the frame number and the confidence score of bounding boxes predicted as “cat”.

3.5.2 Integrating with Your Workflow

The modular design of PeekingDuck allows users to pick and choose the nodes they want to use. Users are also able to use PeekingDuck nodes with external packages when designing their pipeline.

In this demo, we will show how users can construct a custom PeekingDuck pipeline using:

- Data loaders such as `tf.keras.preprocessing.image_dataset_from_directory` (available in `tensorflow>=2.3.0`),
- External packages (not implemented as PeekingDuck nodes) such as `easyocr`, and
- Visualization packages such as `matplotlib`.

The notebook corresponding to this tutorial, `calling_peekingduck_in_python.ipynb`, can be found in the `notebooks` folder of the PeekingDuck repository and is also available as a [Colab notebook](#).

Show/Hide Instructions for Linux/Mac (Intel)/Windows

Note: The uninstallation step is necessary to ensure that the proper version of OpenCV is installed.

You may receive an error message about the incompatibility between `awscli` and `colorama==0.4.4`. `awscli` is conservative about pinning versions to maintain backward compatibility. The code presented in this tutorial has been tested to work and we have chosen to prioritize PeekingDuck’s dependency requirements.

Show/Hide Instructions for Mac (Apple Silicon)

Note: We install the problematic packages `easyocr` and `oidv6` first and then uninstall the pip-related OpenCV packages which were installed as dependencies. Mac (Apple silicon) needs conda’s OpenCV.

There will be a warning that `easyocr` needs some version of Pillow which can be ignored.

We are using `Open Images Dataset V6` as the dataset for this demo. We recommend using the third-party `oidv6` PyPI package to download the images necessary for this demo.

Run the following command after installing the prerequisites:

Terminal Session

```
[~user] > mkdir pkd_project
[~user] > cd pkd_project
[~user/pkd_project] > oidv6 downloader en --dataset data/oidv6 --type_data train --classes car --limit 10 --yes
```

Copy `calling_peekingduck_in_python.ipynb` to the `pkd_project` folder and you should have the following directory structure at this point:

```
pkd_project/
├── calling_peekingduck_in_python.ipynb
├── data/
│   └── oidv6/
```

```
├ boxes/
├ metadata/
├ train/
└ car/
```

Import the Modules

Show/Hide Code

```
1 import os
2 from pathlib import Path
3
4 import cv2
5 import easyocr
6 import matplotlib.pyplot as plt
7 import numpy as np
8 import tensorflow as tf
9 from peekingduck.nodes.draw import bbox
10 from peekingduck.nodes.model import yolo_license_plate
11
12 %matplotlib inline
```

Lines 9 - 10: You can also do:

```
from peekingduck.nodes.draw import bbox as pkd_bbox
from peekingduck.nodes.model import yolo_license_plate as pkd_yolo_license_plate

bbox_node = pkd_bbox.Node()
yolo_license_plate_node = pkd_yolo_license_plate.Node()
```

to avoid potential name conflicts.

Initialize PeekingDuck Nodes

Show/Hide Code

```
1 yolo_lp_node = yolo_license_plate.Node()
2
3 bbox_node = bbox.Node(show_labels=True)
```

Lines 3: To change the node configuration, you can pass the new values to the `Node()` constructor as keyword arguments.

Refer to the [API Documentation](#) for the configurable settings for each node.

Create a Dataset Loader

Show/Hide Code

```

1 data_dir = Path.cwd().resolve() / "data" / "oidv6" / "train"
2 dataset = tf.keras.preprocessing.image_dataset_from_directory(
3     data_dir, batch_size=1, shuffle=False
4 )

```

Lines 2 - 4: We create the data loader using `tf.keras.preprocessing.image_dataset_from_directory()`; you can also create your own data loader class.

Create a License Plate Parser Class

Show/Hide Code

```

1 class LPReader:
2     def __init__(self, use_gpu):
3         self.reader = easyocr.Reader(["en"], gpu=use_gpu)
4
5     def read(self, image):
6         """Reads text from the image and joins multiple strings to a
7         single string.
8         """
9         return " ".join(self.reader.readtext(image, detail=0))
10
11 reader = LPReader(False)

```

We create the license plate parser class in a Python class using `easyocr` to demonstrate how users can integrate the PeekingDuck pipeline with external packages.

Alternatively, users can create a custom node for parsing license plates and run the pipeline through the CLI instead. Refer to the *custom nodes* tutorial for more information.

The Inference Loop

Show/Hide Code

```

1 def get_best_license_plate(frame, bboxes, bbox_scores, width, height):
2     """Returns the image region enclosed by the bounding box with the highest
3     confidence score.
4     """
5     best_idx = np.argmax(bbox_scores)
6     best_bbox = bboxes[best_idx].astype(np.float32).reshape((-1, 2))
7     best_bbox[:, 0] *= width
8     best_bbox[:, 1] *= height
9     best_bbox = np.round(best_bbox).astype(int)
10
11     return frame[slice(*best_bbox[:, 1]), slice(*best_bbox[:, 0])]
12
13 num_col = 3
14 # For visualization, we plot 3 columns, 1) the original image, 2) image with
15 # bounding box, and 3) the detected license plate region with license plate

```

(continues on next page)

(continued from previous page)

```

16 # number prediction shown as the plot title
17 fig, ax = plt.subplots(
18     len(dataset), num_col, figsize=(num_col * 3, len(dataset) * 3)
19 )
20 for i, (element, path) in enumerate(zip(dataset, dataset.file_paths)):
21     image_orig = cv2.imread(path)
22     image_orig = cv2.cvtColor(image_orig, cv2.COLOR_BGR2RGB)
23     height, width = image_orig.shape[:2]
24
25     image = element[0].numpy().astype("uint8")[0].copy()
26
27     yolo_lp_input = {"img": image}
28     yolo_lp_output = yolo_lp_node.run(yolo_lp_input)
29
30     bbox_input = {
31         "img": image,
32         "bboxes": yolo_lp_output["bboxes"],
33         "bbox_labels": yolo_lp_output["bbox_labels"],
34     }
35     _ = bbox_node.run(bbox_input)
36
37     ax[i][0].imshow(image_orig)
38     ax[i][1].imshow(image)
39     # If there are any license plates detected, try to predict the license
40     # plate number
41     if len(yolo_lp_output["bboxes"]) > 0:
42         lp_image = get_best_license_plate(
43             image_orig, yolo_lp_output["bboxes"],
44             yolo_lp_output["bbox_scores"],
45             width,
46             height,
47         )
48         lp_pred = reader.read(lp_image)
49         ax[i][2].imshow(lp_image)
50         ax[i][2].title.set_text(f"Pred: {lp_pred}")

```

Lines 1 - 11: We define a utility function for retrieving the image region of the license plate with the highest confidence score to improve code clarity. For more information on how to convert between bounding box and image coordinates, please refer to the [Bounding Box vs Image Coordinates](#) tutorial.

Lines 27 - 35: By carefully constructing the input for each of the nodes, we can perform the inference loop within a custom workflow.

Lines 37 - 38: We plot the data using matplotlib for debugging and visualization purposes.

Lines 41 - 48: We integrate the inference loop with external packages such as the license plate parser we have created earlier using easyocr.

3.6 Using Your Own Models

PeekingDuck offers pre-trained *model* nodes that can be used to tackle a wide variety of problems, but you may need to train your own model on a custom dataset sometimes. This tutorial will show you how to package your model into a custom *model* node, and use it with PeekingDuck. We will be tackling a manufacturing use case here - classifying images of steel castings into “defective” or “normal” classes.

Casting is a manufacturing process in which a material such as metal in liquid form is poured into a mold and allowed to solidify. The solidified result is also called a casting. Sometimes, defective castings are produced, and quality assurance departments are responsible for preventing defective pieces from being used downstream. As inspections are usually done manually, this adds a significant amount of time and cost, and thus there is an incentive to automate this process.

The images of castings used in this tutorial are the front faces of steel *pump impellers*. From the comparison below, it can be seen that the defective casting has a rough, uneven edges while the normal casting has smooth edges.

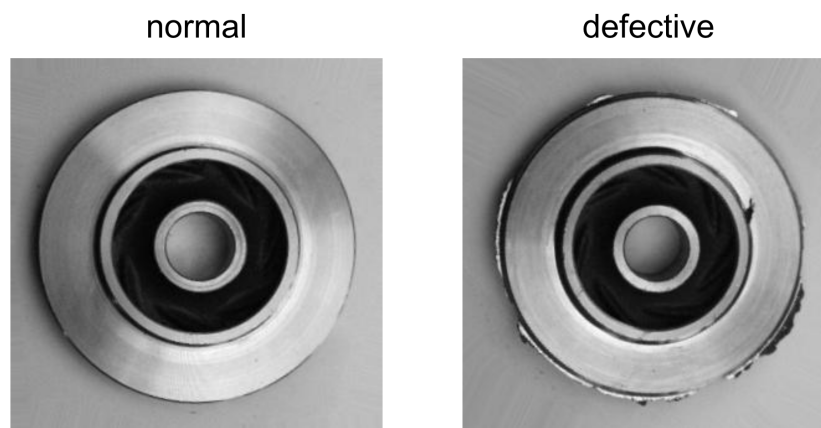


Fig. 11: Normal Casting Compared to Defective Casting

3.6.1 Model Training

PeekingDuck is designed for model *inference* rather than model *training*. This optional section shows how a simple Convolutional Neural Network (CNN) model can be trained separately from the PeekingDuck framework. If you have already trained your own model, the *following section* describes how you can convert it to a custom model node, and use it within PeekingDuck for inference.

Setting Up

Install the following prerequisite for visualization.

```
> conda install matplotlib
```

Create the following project folder:

Terminal Session

```
[~user] > mkdir castings_project
```

```
[~user] > cd castings_project
```

Download the [castings dataset](#) and unzip the file to the `castings_project` folder.

Note: The castings dataset used in this example is modified from the original dataset from [Kaggle](#).

You should have the following directory structure at this point:

```
castings_project/
├── castings_data/
│   ├── inspection/
│   ├── train/
│   └── validation/
```

Update Training Script

Create an empty `train_classifier.py` file within the `castings_project` folder, and update it with the following code:

train_classifier.py:

Show/Hide Code for train_classifier.py

```
1  """
2  Script to train a classification model on images, save the model, and plot the
   ↪ training results
3
4  Adapted from: https://www.tensorflow.org/tutorials/images/classification
5  """
6
7  import pathlib
8  from typing import List, Tuple
9
10 import matplotlib.pyplot as plt
11 import tensorflow as tf
12 from tensorflow.keras import layers
13 from tensorflow.keras.models import Sequential
14 from tensorflow.keras.layers.experimental.preprocessing import Rescaling
15
16 # setup global constants
17 DATA_DIR = "./castings_data"
18 WEIGHTS_DIR = "./weights"
19 RESULTS = "training_results.png"
20 EPOCHS = 10
21 BATCH_SIZE = 32
22 IMG_HEIGHT = 180
23 IMG_WIDTH = 180
24
25
26 def prepare_data() -> Tuple[tf.data.Dataset, tf.data.Dataset, List[str]]:
27     """
```

(continues on next page)

(continued from previous page)

```

28     Generate training and validation datasets from a folder of images.
29
30     Returns:
31         train_ds (tf.data.Dataset): Training dataset.
32         val_ds (tf.data.Dataset): Validation dataset.
33         class_names (List[str]): Names of all classes to be classified.
34     """
35
36     train_dir = pathlib.Path(DATA_DIR, "train")
37     validation_dir = pathlib.Path(DATA_DIR, "validation")
38
39     train_ds = tf.keras.preprocessing.image_dataset_from_directory(
40         train_dir,
41         image_size=(IMG_HEIGHT, IMG_WIDTH),
42         batch_size=BATCH_SIZE,
43     )
44
45     val_ds = tf.keras.preprocessing.image_dataset_from_directory(
46         validation_dir,
47         image_size=(IMG_HEIGHT, IMG_WIDTH),
48         batch_size=BATCH_SIZE,
49     )
50
51     class_names = train_ds.class_names
52
53     return train_ds, val_ds, class_names
54
55
56 def train_and_save_model(
57     train_ds: tf.data.Dataset, val_ds: tf.data.Dataset, class_names: List[str]
58 ) -> tf.keras.callbacks.History:
59     """
60     Train and save a classification model on the provided data.
61
62     Args:
63         train_ds (tf.data.Dataset): Training dataset.
64         val_ds (tf.data.Dataset): Validation dataset.
65         class_names (List[str]): Names of all classes to be classified.
66
67     Returns:
68         history (tf.keras.callbacks.History): A History object containing
69         ↪ recorded events from
70         ↪ model training.
71     """
72
73     num_classes = len(class_names)
74
75     model = Sequential(
76         [
77             Rescaling(1.0 / 255, input_shape=(IMG_HEIGHT, IMG_WIDTH, 3)),
78             layers.Conv2D(16, 3, padding="same", activation="relu"),
79             layers.MaxPooling2D(),

```

(continues on next page)

(continued from previous page)

```

79         layers.Conv2D(32, 3, padding="same", activation="relu"),
80         layers.MaxPooling2D(),
81         layers.Conv2D(64, 3, padding="same", activation="relu"),
82         layers.MaxPooling2D(),
83         layers.Dropout(0.2),
84         layers.Flatten(),
85         layers.Dense(128, activation="relu"),
86         layers.Dense(num_classes),
87     ]
88 )
89
90 model.compile(
91     optimizer="adam",
92     loss=tf.keras.losses.SparseCategoricalCrossentropy(from_logits=True),
93     metrics=["accuracy"],
94 )
95
96 print(model.summary())
97 history = model.fit(train_ds, validation_data=val_ds, epochs=EPOCHS)
98 model.save(WEIGHTS_DIR)
99
100 return history
101
102
103 def plot_training_results(history: tf.keras.callbacks.History) -> None:
104     """
105     Plot training and validation accuracy and loss curves, and save the plot.
106
107     Args:
108         history (tf.keras.callbacks.History): A History object containing
109         ↪ recorded events from
110             model training.
111     """
112     acc = history.history["accuracy"]
113     val_acc = history.history["val_accuracy"]
114     loss = history.history["loss"]
115     val_loss = history.history["val_loss"]
116     epochs_range = range(EPOCHS)
117
118     plt.figure(figsize=(16, 8))
119     plt.subplot(1, 2, 1)
120     plt.plot(epochs_range, acc, label="Training Accuracy")
121     plt.plot(epochs_range, val_acc, label="Validation Accuracy")
122     plt.legend(loc="lower right")
123     plt.title("Training and Validation Accuracy")
124
125     plt.subplot(1, 2, 2)
126     plt.plot(epochs_range, loss, label="Training Loss")
127     plt.plot(epochs_range, val_loss, label="Validation Loss")
128     plt.legend(loc="upper right")
129     plt.title("Training and Validation Loss")
130     plt.savefig(RESULTS)

```

(continues on next page)

(continued from previous page)

```

130
131
132 if __name__ == "__main__":
133     train_ds, val_ds, class_names = prepare_data()
134     history = train_and_save_model(train_ds, val_ds, class_names)
135     plot_training_results(history)

```

Training the Model

Train the model by running the following command.

Terminal Session

```
[~user/castings_project] > python train_classifier.py
```

Note: For macOS Apple Silicon, the above code only works on macOS 12.x Monterey with the latest tensorflow-macos and tensorflow-metal versions. It will crash on macOS 11.x Big Sur due to bugs in the outdated tensorflow versions.

The model will be trained for 10 epochs, and when training is completed, a new `weights` folder and `training_results.png` will be created:

```

castings_project/
├── train_classifier.py
├── training_results.png
├── castings_data/
│   ├── inspection/
│   ├── train/
│   └── validation/
├── weights/
│   ├── keras_metadata.pb
│   ├── saved_model.pb
│   ├── assets/
│   └── variables/

```

The plots from `training_results.png` shown below indicate that the model has performed well on the validation dataset, and we are ready to create a custom `model` node from it.

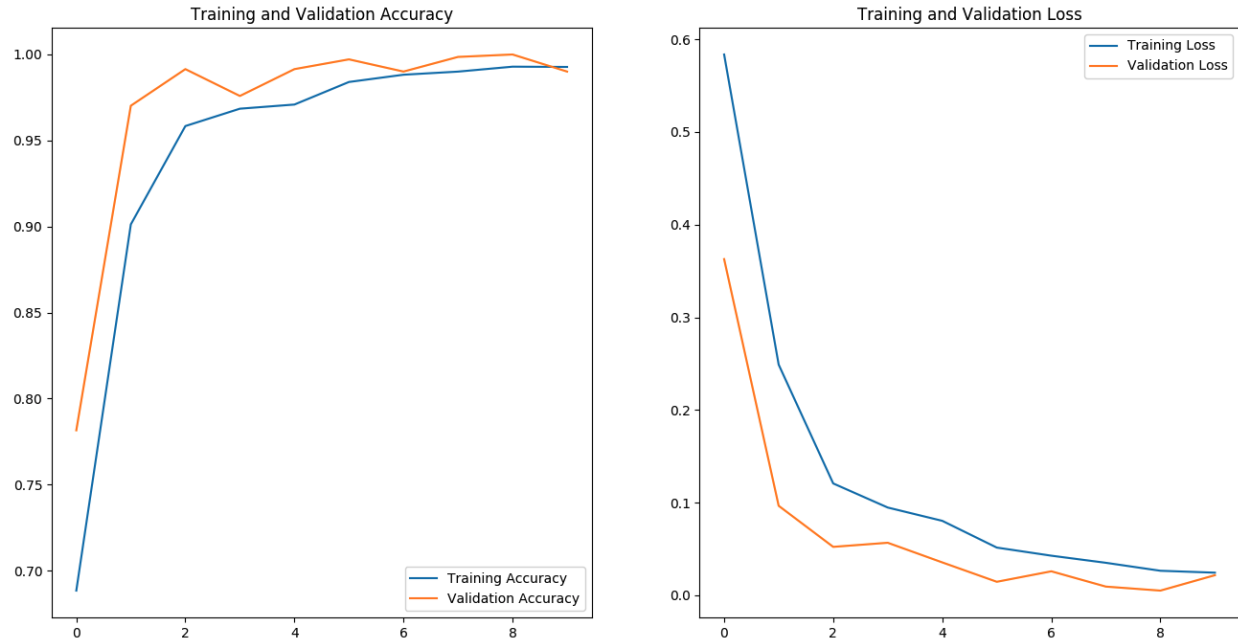


Fig. 12: Model Training Results

3.6.2 Using Your Trained Model with PeekingDuck

This section will show you how to convert your trained model into a custom PeekingDuck node, and give an example of how you can integrate this node in a PeekingDuck pipeline. It assumes that you are already familiar with the process of creating custom nodes, covered in the earlier *custom node* tutorial.

Converting to a Custom Model Node

First, let's create a new PeekingDuck project within the existing `castings_project` folder.

Terminal Session

```
[~user/castings_project] > peekingduck init
```

Next, we'll use the `peekingduck create-node` command to create a custom node:

Terminal Session

```
[~user/castings_project] > peekingduck create-node
Creating new custom node...
Enter node directory relative to ~user/castings_project [src/custom_nodes]:
Select node type (input, augment, model, draw, dabble, output): model
Enter node name [my_custom_node]: casting_classifier
```

Node directory: ~user/castings_project/src/custom_nodes

Node type: model

Node name: casting_classifier

Creating the following files:

Config file: ~user/castings_project/src/custom_nodes/configs/model/casting_classifier.yml

Script file: ~user/castings_project/src/custom_nodes/model/casting_classifier.py

Proceed? [Y/n]:

Created node!

The castings_project folder structure should now look like this:

```
castings_project/
├── pipeline_config.yml
├── train_classifier.py
├── training_results.png
├── castings_data/
│   ├── inspection/
│   ├── train/
│   └── validation/
├── src/
│   └── custom_nodes/
│       ├── configs/
│       │   └── model/
│       │       └── casting_classifier.yml
│       └── model/
│           └── casting_classifier.py
└── weights/
    ├── keras_metadata.pb
    ├── saved_model.pb
    ├── assets/
    └── variables/
```

castings_project now contains **two files** that we need to modify to implement our custom node.

1. src/custom_nodes/configs/model/casting_classifier.yml:

casting_classifier.yml updated content:

```
1 input: ["img"]
2 output: ["pred_label", "pred_score"]
3
4 weights_parent_dir: weights
5 class_label_map: {0: "defective", 1: "normal"}
```

2. src/custom_nodes/model/casting_classifier.py:

casting_classifier.py updated content:

Show/Hide Code for casting_classifier.py

```
1  """
2  Casting classification model.
3  """
4
5  from typing import Any, Dict
6
7  import cv2
8  import numpy as np
9  import tensorflow as tf
10
11  from peekingduck.nodes.node import AbstractNode
12
13  IMG_HEIGHT = 180
14  IMG_WIDTH = 180
15
16
17  class Node(AbstractNode):
18      """Initializes and uses a CNN to predict if an image frame shows a normal
19      or defective casting.
20      """
21
22      def __init__(self, config: Dict[str, Any] = None, **kwargs: Any) -> None:
23          super().__init__(config, node_path=__name__, **kwargs)
24          self.model = tf.keras.models.load_model(self.weights_parent_dir)
25
26      def run(self, inputs: Dict[str, Any]) -> Dict[str, Any]:
27          """Reads the image input and returns the predicted class label and
28          confidence score.
29
30          Args:
31              inputs (dict): Dictionary with key "img".
32
33          Returns:
34              outputs (dict): Dictionary with keys "pred_label" and "pred_score".
35          """
36          img = cv2.cvtColor(inputs["img"], cv2.COLOR_BGR2RGB)
37          img = cv2.resize(img, (IMG_WIDTH, IMG_HEIGHT))
38          img = np.expand_dims(img, axis=0)
39          predictions = self.model.predict(img)
40          score = tf.nn.softmax(predictions[0])
41
42          return {
43              "pred_label": self.class_label_map[np.argmax(score)],
44              "pred_score": 100.0 * np.max(score),
45          }
```

The custom node takes in the built-in PeekingDuck *img* data type, makes a prediction based on the image, and produces two custom data types: `pred_label`, the predicted label (“defective” or “normal”); and `pred_score`, which is the confidence score of the prediction.

Using the Classifier in a PeekingDuck Pipeline

We'll now pair this custom node with other PeekingDuck nodes to build a complete solution. Imagine an automated inspection system like the one shown below, where the castings are placed on a conveyor belt and a camera takes a picture of each casting and sends it to the PeekingDuck pipeline for prediction. A report showing the predicted result for each casting is produced, and the quality inspector can use it for further analysis.

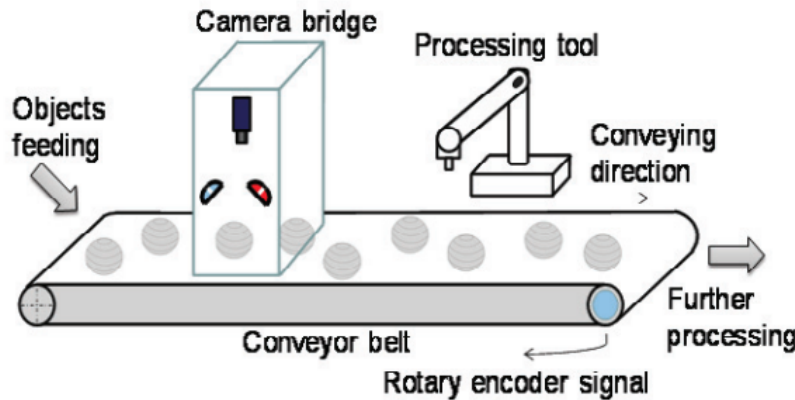


Fig. 13: Vision Based Inspection of Conveyed Objects (Source: [ScienceDirect](#))

Edit the `pipeline_config.yml` file to use the `input.visual` node to read in the images, and the `output.csv_writer` node to produce the report. We will test our solution on the 10 casting images in `castings_data/inspection`, where each image's filename is a unique casting ID such as `28_4137.jpeg`.

pipeline_config.yml:

pipeline_config.yml updated content:

```

1 nodes:
2   - input.visual:
3     source: castings_data/inspection
4   - custom_nodes.model.casting_classifier
5   - output.csv_writer:
6     stats_to_track: ["filename", "pred_label", "pred_score"]
7     file_path: casting_predictions.csv
8     logging_interval: 0

```

Line 2 `input.visual`: tells PeekingDuck to load the images from `castings_data/inspection`.

Line 4 Calls the custom model node that you have just created.

Line 5 `output.csv_writer`: produces the report for the quality inspector in a CSV file `castings_predictions_DDMMYY-hh-mm-ss.csv` (time stamp appended to `file_path`). This node receives the `filename` data type from `input.visual`, the custom data types `pred_label` and `pred_score` from the custom model node, and writes them to the columns of the CSV file.

Run the above with the command `peekingduck run`.

Open the created CSV file and you would see the following results. Half of the castings have been predicted as defective with high confidence scores. As the file name of each image is its unique casting ID, the quality inspector would be

able to check the results with the actual castings if needed.

Time	filename	pred_label	pred_score
11:50:31	28_3123.jpeg	defective	100
11:50:31	28_3416.jpeg	defective	100
11:50:31	28_4137.jpeg	defective	99.9213099
11:50:31	28_5297.jpeg	defective	100
11:50:31	28_7653.jpeg	defective	100
11:50:31	28_9918.jpeg	normal	99.6197104
11:50:31	28_9920.jpeg	normal	99.99367
11:50:31	28_9925.jpeg	normal	99.9608338
11:50:31	28_9926.jpeg	normal	97.5584447
11:50:31	28_9928.jpeg	normal	97.8244841

Fig. 14: Casting Prediction Results

To visualize the predictions alongside the casting images, create an empty Python script named `visualize_results.py`, and update it with the following code:

visualize_results.py:

Show/Hide Code for `visualize_results.py`

```

1  """
2  Script to visualize the prediction results alongside the casting images
3  """
4
5  import csv
6
7  import cv2
8  import matplotlib.pyplot as plt
9
10 CSV_FILE = "casting_predictions_280422-11-50-30.csv" # change file name_
    ↳ accordingly
11 INSPECTION_IMGS_DIR = "castings_data/inspection/"
12 RESULTS_FILE = "inspection_results.png"
13
14 fig, axs = plt.subplots(2, 5, figsize=(50, 20))
15
16 with open(CSV_FILE) as csv_file:
17     csv_reader = csv.reader(csv_file, delimiter=",")
18     next(csv_reader, None)
19     for i, row in enumerate(csv_reader):
20         # csv columns follow this order: 'Time', 'filename', 'pred_label', 'pred_score'
21         image_path = INSPECTION_IMGS_DIR + row[1]
22         image_orig = cv2.imread(image_path)
23         image_orig = cv2.cvtColor(image_orig, cv2.COLOR_BGR2RGB)
24
25         row_idx = 0 if i < 5 else 1

```

(continues on next page)

(continued from previous page)

```

26     axs[row_idx][i % 5].imshow(image_orig)
27     axs[row_idx][i % 5].set_title(row[1] + " - " + row[2], fontsize=35)
28     axs[row_idx][i % 5].axis("off")
29
30 fig.savefig(RESULTS_FILE)

```

In Line 10, replace the name of `CSV_FILE` with the name of the CSV file produced on your system, as a timestamp would have been appended to the file name.

Run the following command to visualize the results.

Terminal Session

```
[~user/castings_project] > python visualize_results.py
```

An `inspection_results.png` would be created, as shown below. The top row of castings are clearly defective, as they have rough, uneven edges, while the bottom row of castings look normal. Therefore, the prediction results are accurate for this batch of inspected castings. The quality inspector can provide feedback to the manufacturing team to further investigate the defective castings based on the casting IDs.

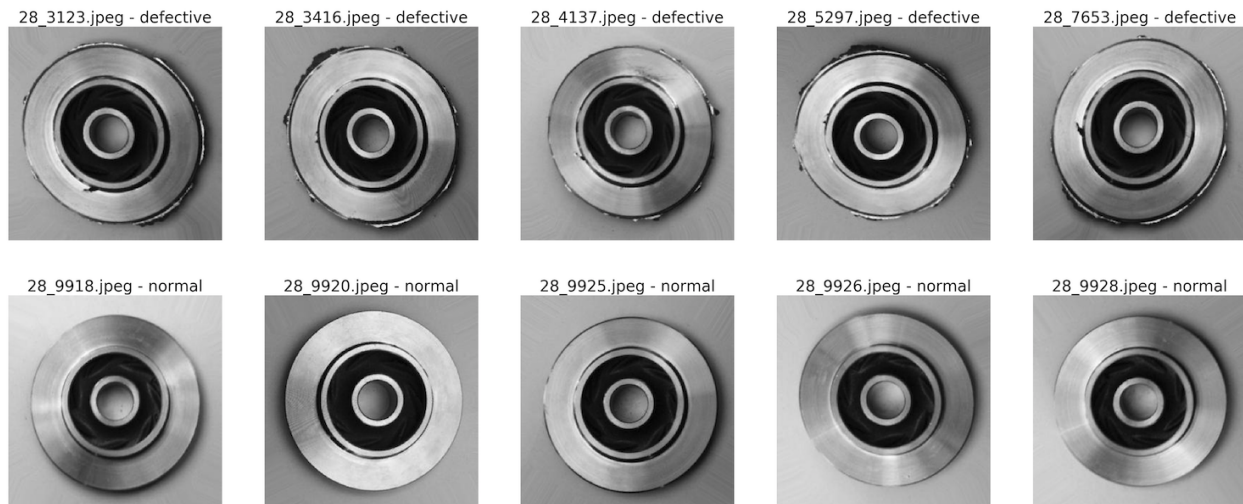


Fig. 15: Casting Prediction Visualization

This concludes the guided example on using your own custom models.

3.6.3 Custom Object Detection Models

The previous example was centered on the task of image classification. *Object detection* is another common task in Computer Vision. PeekingDuck offers several pre-trained *object detection* model nodes which can detect up to 80 different types of objects, such as persons, cars, and dogs, just to name a few. For the complete list of detectable objects, refer to the *Object Detection IDs* page. Quite often, you may need to train a custom object detection model on your own dataset, such as defects on a printed circuit board (PCB) as shown below. This section discusses some important considerations for the object detection task, supplementing the guided example above.

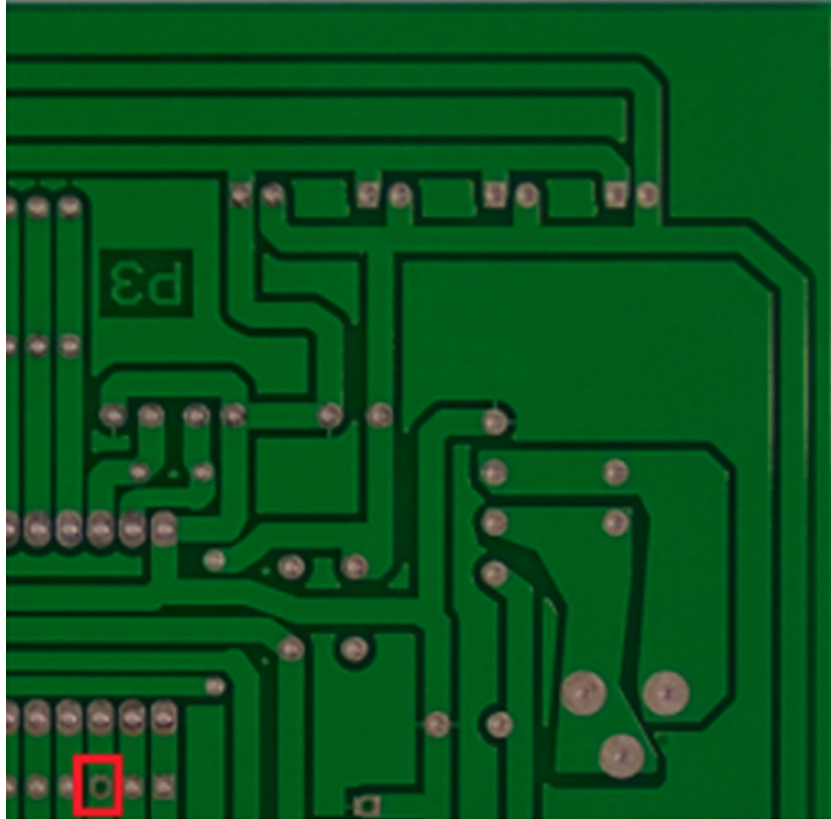


Fig. 16: Object Detection of Defects on PCB (Source: The Institution of Engineering and Technology)

PeekingDuck’s object detection model nodes conventionally receive the *img* data type, and produce the *bboxes*, *bbox_labels*, and *bbox_scores* data types. An example of this can be seen in the API documentation for a node such as `model.efficientdet`. We strongly recommend keeping to these data type conventions for your custom object detection node, ensuring that they adhere to the described format, e.g. *img* is in BGR format, and *bboxes* is a NumPy array of a certain shape.

This allows you to leverage on PeekingDuck’s ecosystem of existing nodes. For example, by ensuring that your custom model node receives *img* in the correct format, you are able to use PeekingDuck’s `input.visual` node, which can read from multiple visual sources such as a folder of images or videos, an online cloud source, or a CCTV/webcam live feed. By ensuring that your custom model node produces *bboxes* and *bbox_labels* in the correct format, you are able to use PeekingDuck’s `draw.bbox` node to draw bounding boxes and associated labels around the detected objects.

By doing so, you would have saved a significant amount of development time, and can focus more on developing and finetuning your custom object detection model. This was just a simple example, and you can find out more about PeekingDuck’s nodes from our *API Documentation*, and PeekingDuck’s built-in data types from our *Glossary*.

3.7 Using Model Hub Nodes

PeekingDuck has support for models from external model hubs. You can leverage the model hub nodes to use these external models with the PeekingDuck pipeline. The model hub nodes differ in design from traditional model nodes as the external models perform a variety of computer vision tasks. This tutorial demonstrates how to use these nodes through some sample pipelines.

3.7.1 List of Model Hub Nodes

The table below shows the model hub nodes available.

Documentation	Model Hubs
<code>model.huggingface_hub</code>	Hugging Face Hub
<code>model.mediapipe_hub</code>	MediaPipe Solutions

3.7.2 Hugging Face Hub

The `model.huggingface_hub` node supports transformer models which perform the following computer vision tasks:

1. Object detection
2. Instance segmentation

You can use the following command to get a list of supported computer vision tasks:

Terminal Session

```
[~user/project_dir] > peekingduck model-hub huggingface tasks
Supported computer vision tasks: ['instance segmentation', 'object detection']
```

Object Detection

This example shows how Hugging Face Hub's object detection models can be used to blur computer related objects from `cat_and_computer.mp4`.

Supported Models

Use the following command to get a list of supported object detection models:

Terminal Session

```
[~user/project_dir] > peekingduck model-hub huggingface models --task 'object detection'
Supported Hugging Face `object_detection` models:
facebook/detr-resnet-50
facebook/detr-resnet-50-dc5
```

hustvl/yolos-base
hustvl/yolos-small

<long list truncated>

Class Labels

The models are trained on a variety of datasets which may not share the same set of class labels, e.g., the [PubTables-1M](#) dataset contains the label “table column” which is not found in the [COCO](#) dataset. As such, it may be necessary to look through the class labels supported by the model. You can do so with the following command:

Terminal Session

```
[~user/project_dir] > peekingduck model-hub huggingface detect-ids --model_type 'hustvl/yolos-small'
```

The detect ID-to-label mapping for `facebook/hustvl/yolos-small` can be found at <https://huggingface.co/hustvl/yolos-small/blob/main/config.json> under the `id2label` key.

Pipeline

The following `pipeline_config.yml` can be used to blur computer related objects:

```
1 nodes:
2   - input.visual:
3     source: /path/to/cat_and_computer.mp4
4   - model.huggingface_hub:
5     task: object_detection
6     model_type: hustvl/yolos-small
7     detect: ["laptop", "keyboard", "mouse"]
8   - draw.blur_bbox
9   - output.screen
```

Here is a step-by-step explanation of what has been done:

Line 2 `input.visual` tells PeekingDuck to load `cat_and_computer.mp4`.

Line 4 `model.huggingface_hub` is set to perform the `object_detection` task and load the `hustvl/yolos-small` model. We also set the model to detect only laptops, keyboards, and mice.

Line 8 `draw.blur_bbox` is used to apply a blur effect on the bounding boxes detected by the model.

Run the above with the command `peekingduck run`.

You should see a display of the `cat_and_computer.mp4` with the computer related objects blurred. The 30-second video will auto-close at the end, or you can press `q` to end early.



Fig. 17: Cat and Computer Screenshot with Computer Related Objects Blurred

Instance Segmentation

This example also demonstrates how to blur computer related objects but with Hugging Face Hub's instance segmentation models instead.

Supported Models

To get a list of supported instance segmentation models, use the following CLI command:

Terminal Session

```
[~user/project_dir] > peekingduck model-hub huggingface models --task 'instance segmentation'
Supported hugging face `instance_segmentation` models:
facebook/detr-resnet-101-panoptic
facebook/detr-resnet-50-dc5-panoptic
facebook/detr-resnet-50-panoptic
facebook/maskformer-swin-tiny-coco

<long list truncated>
```

Class Labels

Similar to the object detection models, the instance segmentation models are also trained on a variety of datasets, the same `detect-ids` command can be used to retrieve the file containing the model's class labels:

Terminal Session

```
[~user/project_dir] > peekingduck model-hub huggingface detect-ids --model_type  
'facebook/maskformer-swin-tiny-coco'  
The detect ID-to-label mapping for `facebook/maskformer-swin-tiny-coco` can be found at  
https://huggingface.co/facebook/maskformer-swin-tiny-coco/blob/main/config.json under the `id2label`  
key.
```

Pipeline

The task of blurring computer related objects can be achieved with the following `pipeline_config.yml`:

```
1 nodes:  
2 - input.visual:  
3     source: /path/to/cat_and_computer.mp4  
4 - model.huggingface_hub:  
5     task: instance_segmentation  
6     model_type: facebook/maskformer-swin-tiny-coco  
7     detect: ["laptop", "keyboard", "mouse"]  
8 - draw.instance_mask:  
9     effect:  
10     blur: 50  
11 - output.screen
```

Here is a step-by-step explanation of what has been done:

Line 2 `input.visual` tells PeekingDuck to load `cat_and_computer.mp4`.

Line 4 `model.huggingface_hub` is set to perform the `instance_segmentation` task and load the `facebook/maskformer-swin-tiny-coco` model. We also set the model to detect only laptops, keyboards, and mice.

Line 8 `draw.instance_mask` is used to visualize the output of the Hugging Face Hub model. A blur effect is applied on the instance mask outputs.

Run the above with the command `peekingduck run`.

You should see a display of the `cat_and_computer.mp4` with the computer related objects blurred. The 30-second video will auto-close at the end, or you can press `q` to end early.



Fig. 18: Cat and Computer Screenshot with Computer Related Objects Blurred

3.7.3 MediaPipe Solutions

The `model.mediapipe_hub` node supports MediaPipe solutions which perform the following computer vision tasks:

1. Object detection
2. Pose estimation

Note: The object detection model only detects faces as the generic object detection model is not available in Python at the time of writing.

You can use the following command to get a list of supported computer vision tasks and their respective subtasks:

Terminal Session

```
[~user/project_dir] > peekingduck model-hub mediapipe tasks
Supported computer vision tasks and respective subtasks:
pose estimation
  body
object detection
  face
```

Object Detection (Face)

This example demonstrates the usage of MediaPipe object detection (face) solution.

Supported Model Types

To get a list of supported model types for the pose estimation (body) task, use the following command:

Terminal Session

```
[~user/project_dir] > peekingduck model-hub mediapipe model-types --task 'object detection' --subtask 'face'
```

Supported model types for 'object detection/face'

0: A short-range model that works best for faces within 2 meters from the camera.

1: A full-range model best for faces within 5 meters.

Pipeline

```
1 nodes:
2   - input.visual:
3     source: https://storage.googleapis.com/peekingduck/videos/wave.mp4
4   - model.mediapipe_hub:
5     task: object_detection
6     subtask: face
7     model_type: 0
8   - draw.bbox
9   - output.screen
```

Here is a step-by-step explanation of what has been done:

Line 2 `input.visual` tells PeekingDuck to load `wave.mp4`.

Line 4 `model.mediapipe_hub` is set to perform the `object_detection` task and `face` subtask. Model type 0 is selected.

Line 8 `draw.bbox` is used to visualize the output of the MediaPipe model.

Run the above with the command `peekingduck run`.

You should see a display of the `wave.mp4` with a bounding box drawn around the person's face. The video will auto-close at the end, or you can press `q` to end early.



Fig. 19: Object Detection Screenshot

Pose Estimation (Body)

This example demonstrates the usage of MediaPipe pose estimation (body) solution.

Supported Model Types

To get a list of supported model types for the pose estimation (body) task, use the following command:

Terminal Session

```
[~user/project_dir] > peekingduck model-hub mediapipe model-types --task 'pose estimation' --subtask 'body'
```

Supported model types for 'pose estimation/body'

0: BlazePose GHUM Lite, lower inference latency at the cost of landmark accuracy.

1: BlazePose GHUM Full.

2: BlazePose GHUM Heavy, higher landmark accuracy at the cost of inference latency.

Pipeline

```
1 nodes:
2   - input.visual:
3     source: https://storage.googleapis.com/peekingduck/videos/wave.mp4
4   - model.mediapipe_hub:
5     task: pose_estimation
6     subtask: body
7     model_type: 1
8   - draw.pose
9   - output.screen
```

Here is a step-by-step explanation of what has been done:

Line 2 `input.visual` tells PeekingDuck to load `wave.mp4`.

Line 4 `model.mediapipe_hub` is set to perform the `pose_estimation` task and `body` subtask. Model type 1 is selected.

Line 8 `draw.pose` is used to visualize the output of the MediaPipe model.

Run the above with the command `peekingduck run`.



Fig. 20: Pose Estimation Screenshot

You should see a display of the `wave.mp4` with the skeletal poses drawn on the person. The video will auto-close at the end, or you can press `q` to end early.

PEEKINGDUCK ECOSYSTEM

This section covers the extensions to the PeekingDuck ecosystem.

4.1 PeekingDuck Viewer

The PeekingDuck Viewer offers you an interactive GUI application to manage and run PeekingDuck pipelines, and to view and analyze the output video.

4.1.1 Running the Viewer

The PeekingDuck Viewer can be activated using the CLI `--viewer` option:

Terminal Session

```
[~user] > peekingduck run --viewer
```

A screenshot of the Viewer and its GUI components is shown below:

Once the Viewer screen appears, PeekingDuck will begin executing the current pipeline. The pipeline output is displayed as a video in the center of the screen, with a progress bar below it.

If pipeline input length is deterministic (e.g. using a video file as the source), the progress bar functions like a normal progress bar moving from start to end. Upon completion, the progress bar will be replaced with a slider that you can use to navigate the output video.

If the length is non-deterministic (e.g. capturing a webcam video), then the progress bar will function in a non-deterministic manner: animating itself to indicate progress but without an end point (as PeekingDuck has no idea how long the webcam video will be). In this case, click the **Play/Stop** button to end the webcam video capture, and the progress bar will become a slider.

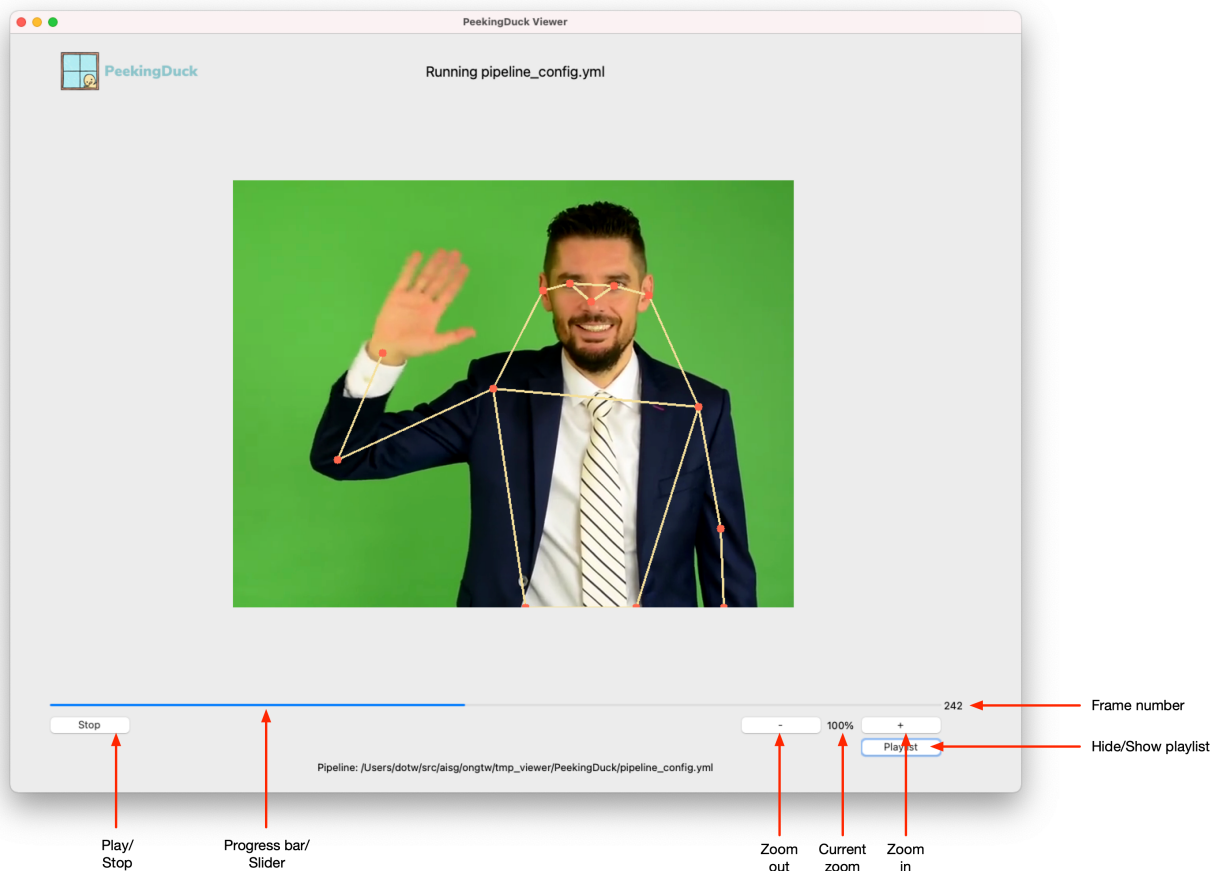


Fig. 1: The PeekingDuck Viewer screen, with explanations for the main controls.

4.1.2 Navigating the Output Video

You can examine the output video of the executed pipeline by using the Play/Stop button to replay the entire video.

You may also scrub through the video using the slider to go directly to the frames of interest. The current video frame number is shown to the right of the slider, serving as a position indicator. To “jump” to a particular position on the slider, click the *right* mouse button on that position. To move frame-by-frame forward/backward, click the *left* mouse button anywhere to the right/left of the current slider position.

The + (zoom in) and - (zoom out) buttons allow you to adjust the video size. You may also use keyboard shortcuts to adjust the zoom: CTRL - - zoom out, CTRL - + zoom in, CTRL - = reset zoom

4.1.3 Using the Pipeline Playlist

Clicking the Playlist button will show/hide the playlist.

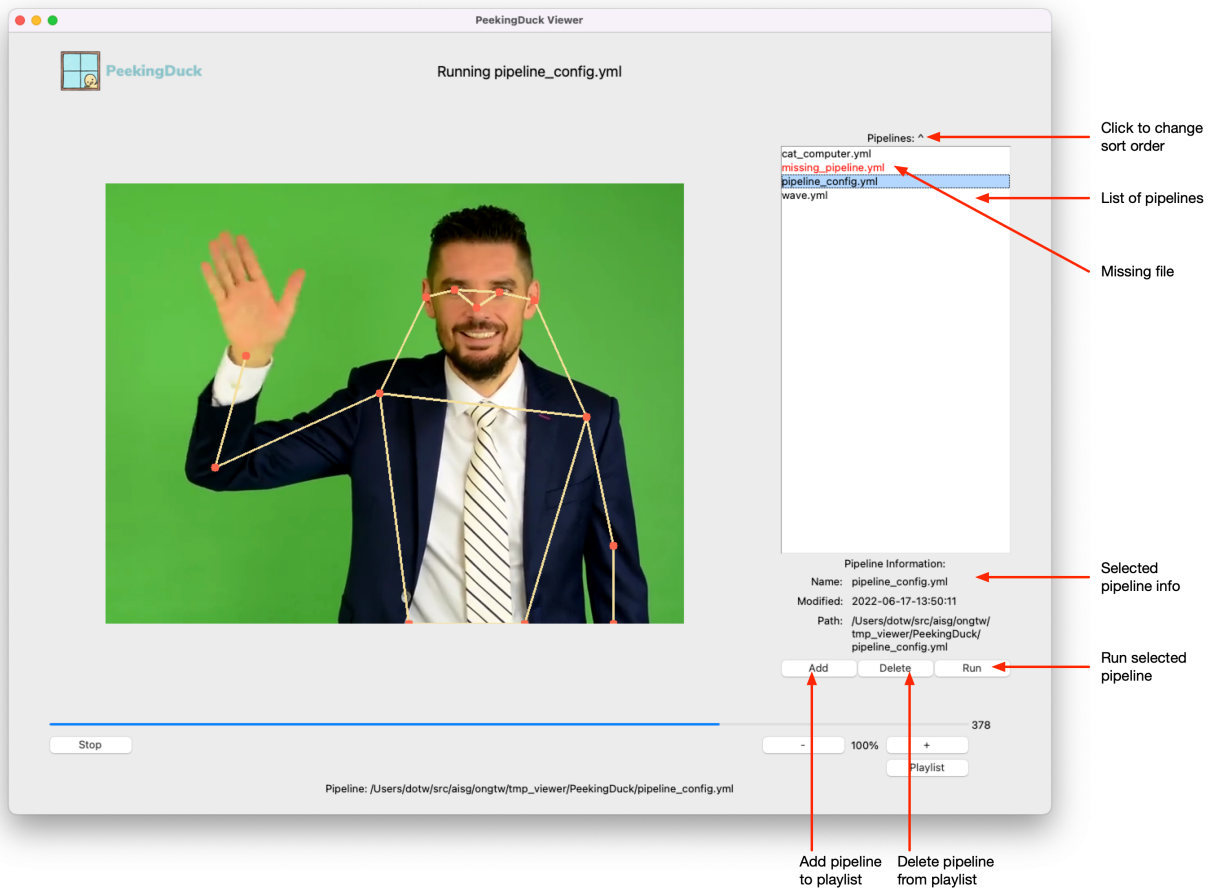


Fig. 2: PeekingDuck Viewer with playlist shown.

The above screenshot shows the playlist on the right. The playlist is a collection of pipeline files that can be run with PeekingDuck. The current pipeline is automatically added to the playlist. This playlist is specific to you and is saved across different PeekingDuck Viewer runs.

Click to select a pipeline in the playlist. The pipeline’s information will be displayed in the Pipeline Information panel below. It shows the pipeline’s name, last modified date/time, and full file path.

To run the currently selected pipeline, click the Run button.

The Add button lets you manually add a pipeline file to the playlist. It will display a File Explorer dialog. Use it to select a PeekingDuck pipeline YAML file and it will be added to your playlist.

The Delete button will remove the currently selected pipeline from the playlist, after you have confirmed the deletion.

If the pipeline in the playlist is red, it means the pipeline YAML file is missing. This could mean the pipeline had been added earlier, but its YAML file had since been deleted or moved to another folder. Delete the missing pipeline entry to remove it from the playlist.

The list of pipelines can be sorted in reverse order by clicking the playlist header.

Note: The playlist is saved in `~/.peekingduck/playlist.yml`, where `~` is the user's home folder.

4.1.4 Exiting the Viewer

To exit the Viewer, close the Viewer window.

MODEL RESOURCES & INFORMATION

5.1 Object Detection Models

5.1.1 List of Object Detection Models

The table below shows the object detection models available for each task category.

Category	Model	Documentation
General	EfficientDet	<code>model.efficientdet</code>
	YOLOv4	<code>model.yolo</code>
	YOLOX	<code>model.yolox</code>
Face	MTCNN	<code>model.mtcnn</code>
	YOLOv4 (Face)	<code>model.yolo_face</code>
License plate	YOLOv4 (License Plate)	<code>model.yolo_license_plate</code>

5.1.2 Benchmarks

Inference Speed

The table below shows the frames per second (FPS) of each model type.

Model	Type	Size	CPU		GPU	
			single	multiple	single	multiple
YOLO	v4tiny	416	22.42	21.71	65.24	57.50
	v4	416	2.62	2.59	30.40	28.71
EfficientDet	0	512	5.24	5.25	29.51	29.39
	1	640	2.53	2.49	23.79	24.44
	2	768	1.54	1.50	19.86	20.51
	3	896	0.78	0.75	14.69	14.84
	4	1024	0.43	0.42	11.74	11.88
MTCNN	—	—	32.42	18.53	56.35	51.45
YOLOX	yolox-tiny	416	19.43	19.29	55.36	55.38
	yolox-s	640	15.10	15.44	53.81	53.74
	yolox-m	640	8.29	8.04	42.79	43.83
	yolox-l	640	4.59	4.75	35.30	36.08

Hardware

The following hardware were used to conduct the FPS benchmarks:

- CPU: 2.8 GHz 4-Core Intel Xeon (2020, Cascade Lake) CPU and 16GB RAM
- GPU: NVIDIA A100, paired with 2.2 GHz 6-Core Intel Xeon CPU and 85GB RAM

Test Conditions

The following test conditions were followed:

- `input.visual`, the model of interest, and `dabble.fps` nodes were used to perform inference on videos
- 2 videos were used to benchmark each model, one with only 1 human (`single`), and the other with multiple humans (`multiple`)
- Both videos are about 1 minute each, recorded at ~30 FPS, which translates to about 1,800 frames to process per video
- 1280×720 (HD ready) resolution was used, as a bridge between 640×480 (VGA) of poorer quality webcams, and 1920×1080 (Full HD) of CCTVs

Model Accuracy

The table below shows the performance of our object detection models using the detection evaluation metrics from COCO. Description of these metrics can be found [here](#).

Model	Type	Size	AP	AP IoU=.50	AP IoU=.75	AP small	AP medium	AP large	AR max=1	AR max=10	AR max=100	AR small	AR medium	AR large
YOLO	v4tiny	416	17.4	32.7	16.6	6.4	20.1	25.6	16.7	22.8	21.1	6.1	23.7	32.1
	v4	416	43.7	64.0	48.1	23.1	49.6	60.9	33.3	49.1	50.0	26.2	56.1	70.3
Efficient-Det	0	512	29.7	44.3	32.4	7.4	34.4	49.2	25.3	34.5	34.8	7.8	39.7	58.4
	1	640	35.2	50.8	38.8	14.3	40.1	53.9	28.8	40.5	40.9	15.6	46.3	62.8
	2	768	38.5	54.4	42.1	18.9	42.7	57.1	30.9	43.9	44.4	20.8	48.9	65.5
	3	896	41.1	57.0	45.2	22.2	45.1	58.7	32.6	46.7	47.3	24.8	51.5	66.9
	4	1024	43.4	59.2	47.8	24.2	47.6	60.4	33.8	49.1	49.7	27.3	53.9	68.7
YOLOX	yolox-tiny	416	32.4	50.5	33.9	13.4	35.4	49.5	28.2	43.5	45.7	20.7	51.7	65.9
	yolox-s	416	35.6	53.4	37.8	14.0	39.3	55.7	30.3	46.0	48.1	20.9	54.7	70.8
	yolox-m	416	41.6	59.7	44.4	18.8	46.9	62.8	33.9	51.6	53.7	26.9	60.9	76.8
	yolox-l	416	44.5	62.5	47.6	21.9	50.6	65.5	35.5	54.2	56.3	31.0	64.0	78.1

Dataset

The [MS COCO](#) (val 2017) dataset is used. We integrated the COCO API into the PeekingDuck pipeline for loading the annotations and evaluating the outputs from the models. All values are reported in percentages.

All images from the 80 object categories in the MS COCO (val 2017) dataset were processed.

Test Conditions

The following test conditions were followed:

- The tests were performed using [pycocotools](#) on the MS COCO dataset
- The evaluation metrics have been compared with the original repository of the respective object detection models for consistency

5.1.3 Object Detection IDs

General Object Detection

The tables below provide the associated indices for each class in object detectors.

To detect all classes, specify `detect: ["*"]` under the object detection node configuration in `pipeline_config.yml`.

Class name	ID		Class name	ID	
	YOLO / YOLOX	EfficientDet		YOLO / YOLOX	EfficientDet
person	0	0	elephant	20	21
bicycle	1	1	bear	21	22
car	2	2	zebra	22	23
motorcycle	3	3	giraffe	23	24
aeroplane	4	4	backpack	24	26
bus	5	5	umbrella	25	27
train	6	6	handbag	26	30
truck	7	7	tie	27	31
boat	8	8	suitcase	28	32
traffic light	9	9	frisbee	29	33
fire hydrant	10	10	skis	30	34
stop sign	11	12	snowboard	31	35
parking meter	12	13	sports ball	32	36
bench	13	14	kite	33	37
bird	14	15	baseball bat	34	38
cat	15	16	baseball glove	35	39
dog	16	17	skateboard	36	40
horse	17	18	surfboard	37	41
sheep	18	19	tennis racket	38	42
cow	19	20	bottle	39	43

Class name	ID		Class name	ID	
	YOLO / YOLOX	EfficientDet		YOLO / YOLOX	EfficientDet
wine glass	40	45	dining table	60	66
cup	41	46	toilet	61	69
fork	42	47	tv	62	71
knife	43	48	laptop	63	72
spoon	44	49	mouse	64	73
bowl	45	50	remote	65	74
banana	46	51	keyboard	66	75
apple	47	52	cell phone	67	76
sandwich	48	53	microwave	68	77
orange	49	54	oven	69	78
broccoli	50	55	toaster	70	79
carrot	51	56	sink	71	80
hot dog	52	57	refrigerator	72	81
pizza	53	58	book	73	83
donut	54	59	clock	74	84
cake	55	60	vase	75	85
chair	56	61	scissors	76	86
couch	57	62	teddy bear	77	87
potted plant	58	63	hair drier	78	88
bed	59	64	toothbrush	79	89

Face Detection

This table provides the associated indices for the `model.yolo_face` node.

Class name	ID
no mask	0
mask	1

Model Hub Models

Use the following CLI command to access the class name to ID mapping of model hub models:

Terminal Session

```
[~user] > peekingduck model-hub <repository name> detect-ids \  
--model_type '<model identifier>'
```

where `repository name` is the model hub name, e.g., `huggingface` and `model identifier` is one of the models from the specified model hub, e.g., `facebook/detr-resnet-50`. Please note that `model identifier` is case sensitive.

5.2 Pose Estimation Models

5.2.1 List of Pose Estimation Models

The table below shows the pose estimation models available for each task category.

Category	Model	Documentation
Whole body	HRNet	<code>model.hrnet</code>
	PoseNet	<code>model.posenet</code>
	MoveNet	<code>model.movenet</code>

5.2.2 Benchmarks

Inference Speed

The table below shows the frames per second (FPS) of each model type.

Model	Type	Size	CPU		GPU	
			single	multiple	single	multiple
PoseNet	50	225	64.46	51.95	136.31	89.37
	75	225	57.62	47.01	132.84	83.73
	100	225	44.70	37.60	132.73	81.24
	resnet	225	18.77	17.21	73.15	51.65
HRNet (YOLO)	(v4tiny)	256 × 192 (416)	5.86	1.09	21.91	13.86
MoveNet	SinglePose Lightning	192	40.78	40.54	99.47	–
	SinglePose Thunder	256	25.13	24.87	92.05	–
	MultiPose Lightning	256 or multiple of 32	25.33	24.90	80.64	79.32

Hardware

The following hardware were used to conduct the FPS benchmarks:

- CPU: 2.8 GHz 4-Core Intel Xeon (2020, Cascade Lake) CPU and 16GB RAM
- GPU: NVIDIA A100, paired with 2.2 GHz 6-Core Intel Xeon CPU and 85GB RAM

Test Conditions

The following test conditions were followed:

- `input.visual`, the model of interest, and `dabble.fps` nodes were used to perform inference on videos
- 2 videos were used to benchmark each model, one with only 1 human (`single`), and the other with multiple humans (`multiple`)
- Both videos are about 1 minute each, recorded at ~30 FPS, which translates to about 1,800 frames to process per video
- 1280×720 (HD ready) resolution was used, as a bridge between 640×480 (VGA) of poorer quality webcams, and 1920×1080 (Full HD) of CCTVs

Model Accuracy

The table below shows the performance of our pose estimation models using the keypoint evaluation metrics from COCO. Description of these metrics can be found [here](#).

Model	Type	Size	AP	AP OKS=.50	AP OKS=.75	AP medium	AP large	AR	AR OKS=.50	AR OKS=.75	AR medium	AR large
PoseNet	50	225	5.2	15.5	2.7	0.8	11.8	9.6	22.7	7.1	1.4	20.7
	75	225	7.2	19.7	3.6	1.3	15.9	12.1	26.5	9.3	2.2	25.5
	100	225	7.7	20.8	4.4	1.5	17.1	12.6	27.7	10.1	2.3	26.5
	resnet	225	11.9	27.4	8.3	2.2	25.3	17.3	32.5	15.9	2.9	36.8
HRNet (YOLO)	(v4tiny)	256 × 192 (416)	35.8	61.5	37.5	30.1	44.0	40.2	64.4	42.7	33.0	50.2
MoveNet	single- pose_lightning	256 x 256	7.3	15.7	5.7	1.3	15.4	8.8	17.6	7.7	1.1	19.2
	single- pose_thunder	256 x 256	11.6	21.3	10.7	3.0	23.1	13.1	22.5	12.8	2.8	27.1
	multi- pose_lightning	256 x 256	18.7	36.8	16.3	9.0	31.8	21.0	38.5	19.2	9.3	37.0

Dataset

The [MS COCO](#) (val 2017) dataset is used. We integrated the COCO API into the PeekingDuck pipeline for loading the annotations and evaluating the outputs from the models. All values are reported in percentages.

All images from the “person” category in the MS COCO (val 2017) dataset were processed.

Test Conditions

The following test conditions were followed:

- The tests were performed using [pycocotools](#) on the MS COCO dataset
- The evaluation metrics have been compared with the original repository of the respective pose estimation models for consistency

5.2.3 Keypoint IDs

Whole Body

Keypoint	ID	Keypoint	ID
nose	0	left wrist	9
left eye	1	right wrist	10
right eye	2	left hip	11
left ear	3	right hip	12
right ear	4	left knee	13
left shoulder	5	right knee	14
right shoulder	6	left ankle	15
left elbow	7	right ankle	16
right elbow	8		

5.3 Object Tracking Models

5.3.1 List of Object Tracking Models

The table below shows the object tracking models available for each task category.

Category	Model	Documentation
General	IoU Tracker	<code>dabble.tracking</code>
	OpenCV MOSSE Tracker	<code>dabble.tracking</code>
Human	JDE	<code>model.jde</code>
	FairMOT	<code>model.fairmot</code>

5.3.2 Benchmarks

Inference Speed

The table below shows the frames per second (FPS) of each model type.

Model	Object Detector Type	Input Size	CPU	GPU
IoU Tracker with YOLOX	yolox-m	–	7.87	36.18
OpenCV MOSSE Tracker with YOLOX	yolox-m	–	6.74	21.45
JDE	–	–	1.86	26.32
FairMOT	–	864 × 480	0.30	22.60

Hardware

The following hardware were used to conduct the FPS benchmarks:

- CPU: 2.8 GHz 4-Core Intel Xeon (Cascade Lake) CPU and 16GB RAM
- GPU: NVIDIA A100, paired with 2.2 GHz 6-Core Intel Xeon CPU and 85GB RAM

Test Conditions

The following test conditions were followed:

- `input.visual`, the model of interest, and `dabble.fps` nodes were used to perform inference on videos
- A video sequence from the MOT Challenge dataset (MOT16-04) was used
- The video sequence has 1050 frames and is encoded at 30 FPS, which translates to about 35 seconds
- 1280×720 (HD ready) resolution was used, as a bridge between 640×480 (VGA) of poorer quality webcams, and 1920×1080 (Full HD) of CCTVs

Model Accuracy

The table below shows the performance of our object tracking models using multiple object tracker (MOT) metrics from MOT Challenge. Description of these metrics can be found [here](#).

Model	Object Detector Type	MOTA	IDF1	ID Sw.	FP	FN
IoU Tracker with YOLOX	yolox-m	34.1	40.9	960	8997	62830
OpenCV MOSSE Tracker with YOLOX	yolox-m	32.8	38	2349	7695	65268
JDE	–	70.1	65.1	1321	6412	25292
FairMOT	–	81.8	80.9	536	3663	15903

Dataset

The [MOT16](#) (train) dataset is used. We integrated the MOT Challenge API into the PeekingDuck pipeline for loading the annotations and evaluating the outputs from the models. *MOTA* and *IDF1* are reported in percentages while *IDS*, *FP*, and *FN* are raw numbers.

Only the “pedestrian” category in MOT16 (train) was processed.

5.4 Crowd Counting Models

5.4.1 List of Crowd Counting Models

The table below shows the crowd counting models available.

Model	Documentation
CSRNet	<code>model.csrnet</code>

5.4.2 Benchmarks

Model Accuracy

The table below shows the performance of CSRNet obtained from the original [GitHub repo](#), using Mean Absolute Error (MAE) as the metric. The reported metrics are close to the results from the [CSRNet paper](#).

Model	Type	Dataset	MAE
CSRNet	dense	ShanghaiTech Part A	65.92
	sparse	ShanghaiTech Part B	11.01

Dataset

The [ShanghaiTech](#) dataset was used. It contains 1,198 annotated images split into 2 parts: Part A contains 482 images with highly congested scenes, while Part B contains 716 images with relatively sparse crowd scenes.

5.5 Instance Segmentation Models

5.5.1 List of Instance Segmentation Models

The table below shows the instance segmentation models available.

Model	Documentation
Mask R-CNN	<code>model.mask_rcnn</code>
YolactEdge	<code>model.yolact_edge</code>

5.5.2 Benchmarks

Inference Speed

The table below shows the frames per second (FPS) of each model type.

Model	Type	Size	CPU		GPU	
			single	multiple	single	multiple
Mask R-CNN	r50-fpn	800-1333	0.76	0.72	22.30	18.58
	r101-fpn	800-1333	0.61	0.57	17.14	14.83
YolactEdge	r50-fpn	550	2.99	2.93	40.84	33.94
	r101-fpn	550	2.32	2.27	29.55	25.89
	mobilenetv2	550	4.93	4.64	48.59	36.66

Hardware

The following hardware were used to conduct the FPS benchmarks:

- CPU: 2.8 GHz 4-Core Intel Xeon (2020, Cascade Lake) CPU and 16GB RAM
- GPU: NVIDIA A100, paired with 2.2 GHz 6-Core Intel Xeon CPU and 85GB RAM

Test Conditions

The following test conditions were followed:

- `input.visual`, the model of interest, and `dabble.fps` nodes were used to perform inference on videos
- 2 videos were used to benchmark each model, one with only 1 human (`single`), and the other with multiple humans (`multiple`)
- Both videos are about 1 minute each, recorded at ~30 FPS, which translates to about 1,800 frames to process per video
- 1280×720 (HD ready) resolution was used, as a bridge between 640×480 (VGA) of poorer quality webcams, and 1920×1080 (Full HD) of CCTVs

Model Accuracy

The table below shows the performance of our Instance Segmentation models using the detection evaluation metrics from COCO. Description of these metrics can be found [here](#).

Evaluation on masks

Model	Type	Size	AP	AP IoU=.50	AP IoU=.75	AP small	AP medium	AP large	AR max=1	AR max=10	AR max=100	AR small	AR medium	AR large
Mask R- CNN	r50- fpn	800- 1333	34.5	56.0	36.7	17.8	37.9	47.1	29.7	45.6	47.6	27.4	51.4	63.8
	r101- fpn	800- 1333	37.1	59.0	39.6	20.4	41.1	49.8	31.4	49.1	51.4	31.9	55.6	67.3
Yolact- Edge	r50- fpn	550	27.8	45.6	28.9	10.4	30.0	43.9	26.3	37.5	38.2	16.3	41.9	57.2
	r101- fpn	550	29.6	47.8	31.1	11.3	32.3	46.3	27.4	38.9	39.7	17.4	43.6	59.6
	mo- bilenetv2	550	21.9	37.2	22.6	7.0	22.9	34.7	22.5	31.7	32.3	12.0	34.8	48.3

Evaluation on bounding boxes

Model	Type	Size	AP	AP IoU=.50	AP IoU=.75	AP small	AP medium	AP large	AR max=1	AR max=10	AR max=100	AR small	AR medium	AR large
Mask R- CNN	r50- fpn	800- 1333	37.8	59.2	41.1	21.6	41.2	49.3	31.4	49.5	51.9	32.6	55.7	66.6
	r101- fpn	800- 1333	41.8	62.2	45.4	24.9	45.8	54.3	34.4	54.6	57.3	38.2	61.4	72.4
Yolact- Edge	r50- fpn	550	30.3	49.8	32.2	14.4	32.1	44.6	27.4	40.1	41.2	21.6	43.7	57.5
	r101- fpn	550	32.6	52.5	34.9	15.2	35.0	47.6	28.6	41.8	42.9	22.6	45.9	59.9
	mo- bilenetv2	550	23.2	40.8	23.8	9.3	23.4	35.1	22.9	33.5	34.5	15.8	35.2	49.1

“””

Dataset

The [MS COCO](#) (val 2017) dataset is used. We integrated the COCO API into the PeekingDuck pipeline for loading the annotations and evaluating the outputs from the models. All values are reported in percentages.

All images from the 80 object categories in the MS COCO (val 2017) dataset were processed.

Test Conditions

The following test conditions were followed:

- The tests were performed using [pycocotools](#) on the MS COCO dataset
- The evaluation metrics have been compared with the original repository of the respective instance segmentation models for consistency

5.5.3 Instance Segmentation IDs

General Instance Segmentation

The tables below provide the associated indices for each class.

To detect all classes, specify detect: `["*"]` under the instance segmentation node configuration in `pipeline_config.yml`.

Class name	ID		Class name	ID	
	Mask R-CNN	YolactEdge		Mask R-CNN	YolactEdge
person	0	0	elephant	21	20
bicycle	1	1	bear	22	21
car	2	2	zebra	23	22
motorcycle	3	3	giraffe	24	23
aeroplane	4	4	backpack	26	24
bus	5	5	umbrella	27	25
train	6	6	handbag	30	26
truck	7	7	tie	31	27
boat	8	8	suitcase	32	28
traffic light	9	9	frisbee	33	29
fire hydrant	10	10	skis	34	30
stop sign	12	11	snowboard	35	31
parking meter	13	12	sports ball	36	32
bench	14	13	kite	37	33
bird	15	14	baseball bat	38	34
cat	16	15	baseball glove	39	35
dog	17	16	skateboard	40	36
horse	18	17	surfboard	41	37
sheep	19	18	tennis racket	42	38
cow	20	19	bottle	43	39

Class name	ID		Class name	ID	
	Mask R-CNN	YolactEdge		Mask R-CNN	YolactEdge
wine glass	45	40	dining table	66	60
cup	46	41	toilet	69	61
fork	47	42	tv	71	62
knife	48	43	laptop	72	63
spoon	49	44	mouse	73	64
bowl	50	45	remote	74	65
banana	51	46	keyboard	75	66
apple	52	47	cell phone	76	67
sandwich	53	48	microwave	77	68
orange	54	49	oven	78	69
broccoli	55	50	toaster	79	70
carrot	56	51	sink	80	71
hot dog	57	52	refrigerator	81	72
pizza	58	53	book	83	73
donut	59	54	clock	84	74
cake	60	55	vase	85	75
chair	61	56	scissors	86	76
couch	62	57	teddy bear	87	77
potted plant	63	58	hair drier	88	78
bed	64	59	toothbrush	89	79

5.6 Bibliography

This document contains links, references, academic literature, and github repositories for related Computer Vision technologies and projects.

5.6.1 Legend

Symbol	Remarks
	Available in PeekingDuck
	Singapore-based research

5.6.2 Object Detection

Reference	Paper	Code
YOLOX	✓	✓
YOLOv4	✓	✓
EfficientDet	✓	✓
MTCNN	✓	✓
Recent advances in deep learning for object detection (2020)	✓	NA

5.6.3 Pose Estimation

Reference	Paper	Code
HRNet	✓	✓
PoseNet	✓	✓
MoveNet	TF Blog	TF Hub
NTU RGB+D Dataset (2016)	✓	NA

5.6.4 Crowd Counting

Reference	Paper	Code
CSRNet	✓	✓

5.6.5 Object Tracking

Reference	Paper	Code
JDE	✓	✓
FairMOT	✓	✓

5.6.6 Instance Segmentation

Reference	Paper	Code
Mask R-CNN	✓	Torchvision Models
YolactEdge	✓	✓

PeekingDuck supports running optimized TensorRT¹ models on devices with NVIDIA GPUs. Using the TensorRT model on these devices provides a speed boost over the regular TensorFlow/PyTorch version. A potential use case is running PeekingDuck on an NVIDIA Jetson device for Edge AI inference.

Currently, PeekingDuck includes TensorRT versions of the following models:

1. MoveNet model for pose estimation,
2. YOLOX model for object detection.

6.1 Installing TensorRT

The following packages are required to run PeekingDuck's TensorRT models:

1. TensorFlow
2. PyTorch
3. PyCUDA

As the actual installation steps vary greatly depending on the user's device, operating system, software environment, and pre-installed libraries/packages, we are unable to provide step-by-step installation instructions.

The user may refer to [NVIDIA's TensorRT Documentation](#) for detailed TensorRT installation information.

6.2 Using TensorRT Models

To use the TensorRT version of a model, change the `model_format` of the model configuration to `tensorrt`.

The following `pipeline_config.yml` shows how to use the MoveNet TensorRT model for pose estimation:

```
1 nodes:
2 - input.visual:
3   source: https://storage.googleapis.com/peekingduck/videos/wave.mp4
4 - model.movenet:
5   model_format: tensorrt
6   model_type: singlepose_lightning
7 - draw.pose
8 - dabble.fps
9 - draw.legend:
```

(continues on next page)

¹ NVIDIA TensorRT Reference

(continued from previous page)

```

10   show: ["fps"]
11 - output.screen

```

The following pipeline_config.yml shows how to use the YOLOX TensorRT model for object detection:

```

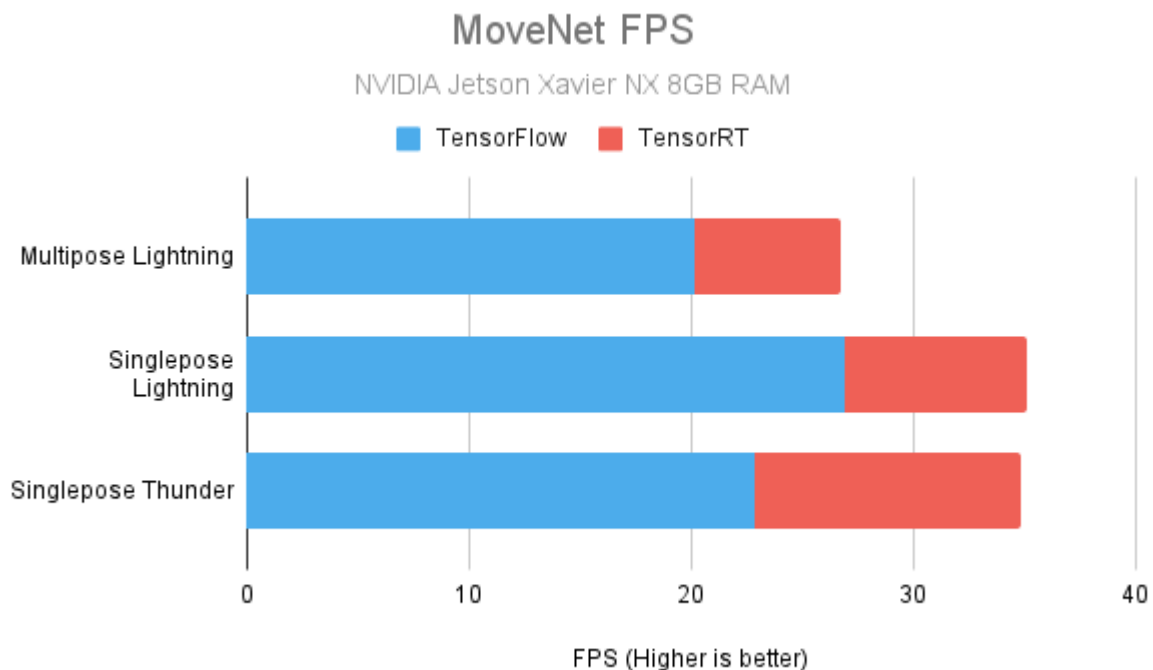
1 nodes:
2 - input.visual:
3   source: https://storage.googleapis.com/peekingduck/videos/cat_and_computer.mp4
4 - model.yolox:
5   detect: ["cup", "cat", "laptop", "keyboard", "mouse"]
6   model_format: tensorrt
7   model_type: yolox-tiny
8 - draw.bbox:
9   show_labels: True    # configure draw.bbox to display object labels
10 - dabble.fps
11 - draw.legend:
12   show: ["fps"]
13 - output.screen

```

6.3 Performance Speedup

The following charts show the speed up obtainable with the TensorRT models. The numbers were obtained from our in-house testing with the actual devices.

6.3.1 NVIDIA Jetson Xavier NX with 8GB RAM



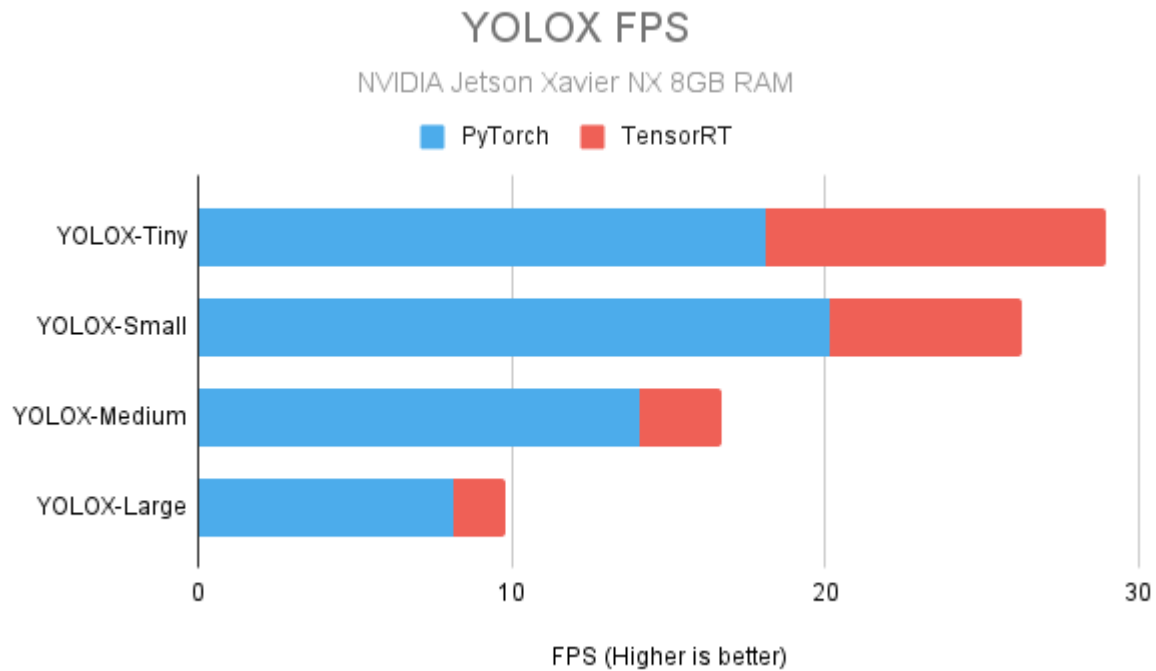


Fig. 1: Jetson Xavier NX specs used for testing:^{Page 101, 2}
CPU: 6 cores (6MB L2 + 4MB L3) GPU: 384-core Volta, 48 Tensor cores RAM: 8 GB

6.3.2 NVIDIA Jetson Xavier AGX with 16GB RAM

Test Conditions

The following test conditions were followed:

- `input.visual`, the model of interest, and `dabble.fps` nodes were used to perform inference on videos
- 2 videos were used to benchmark each model, one with only 1 human (`single`), and the other with multiple humans (`multiple`)
- Both videos are about 1 minute each, recorded at ~30 FPS, which translates to about 1,800 frames to process per video
- 1280×720 (HD ready) resolution was used, as a bridge between 640×480 (VGA) of poorer quality webcams, and 1920×1080 (Full HD) of CCTVs
- FPS numbers are averaged over 5 separate runs

² NVIDIA Jetson Xavier NX Tech Specs

³ NVIDIA Jetson Xavier AGX Tech Specs

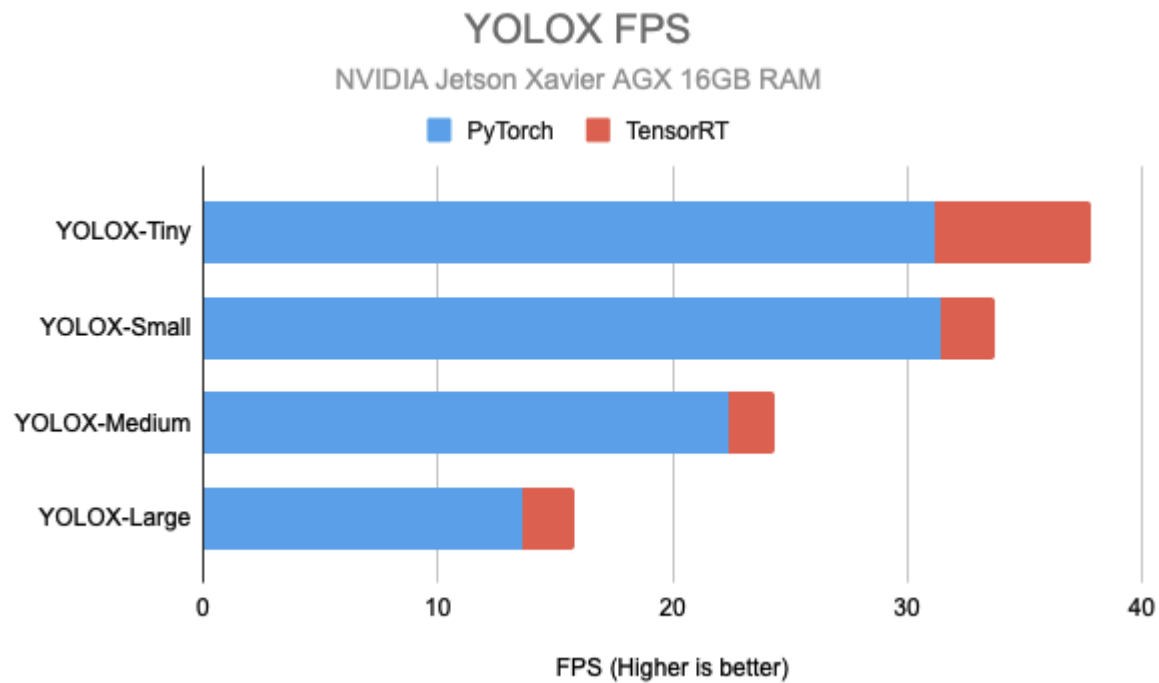
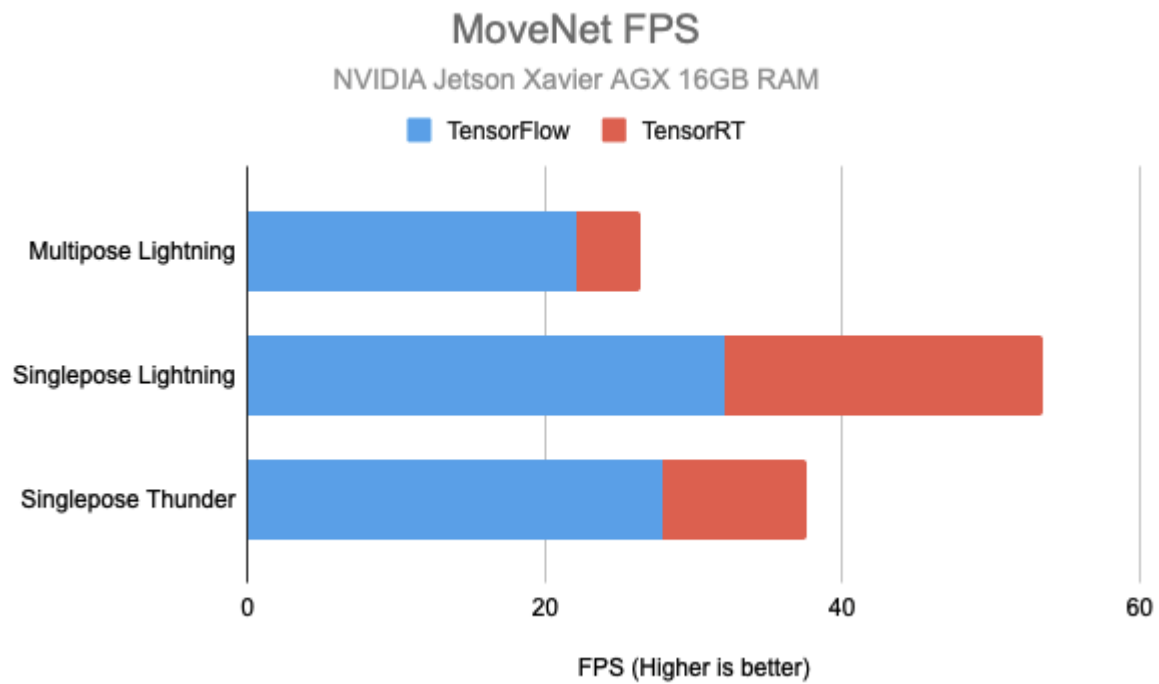


Fig. 2: Jetson Xavier AGX specs used for testing:³
CPU: 8 cores (8MB L2 + 4MB L3) GPU: 512-core Volta, 64 Tensor cores RAM: 16 GB

6.4 References

USE CASES

Computer Vision (CV) problems come in various forms, and the gallery below shows common CV use cases which can be tackled by PeekingDuck right out of the box. Areas include privacy protection, smart monitoring, and COVID-19 prevention and control. Users are encouraged to mix and match different PeekingDuck nodes and create your own *custom nodes* for your specific use case - the only limit is your imagination!

7.1 Privacy Protection

7.1.1 Privacy Protection (Faces)

Overview

As organizations collect more data, there is a need to better protect the identities of individuals in public and private places. Our solution performs face anonymization, and can be used to comply with the General Data Protection Regulation (GDPR) or other data privacy laws.

Our solution automatically detects and mosaics (or blurs) human faces. This is explained in the *How It Works* section.

Demo

To try our solution on your own computer, *install* and run PeekingDuck with the configuration file `privacy_protection_faces.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/privacy_protection_faces.yml>
```

How It Works

There are two main components to face anonymization:

1. Face detection, and
2. Face de-identification.

1. Face Detection

We use an open source face detection model known as [MTCNN](#) to identify human faces. This allows the application to identify the locations of human faces in a video feed. Each of these locations is represented as a pair of x, y coordinates in the form $[x_1, y_1, x_2, y_2]$, where (x_1, y_1) is the top left corner of the bounding box, and (x_2, y_2) is the bottom right. These are used to form the bounding box of each human face detected. For more information on how to adjust the MTCNN node, check out the MTCNN configurable parameters.

2. Face De-Identification

To perform face de-identification, we pixelate or gaussian blur the areas bounded by the bounding boxes.

Nodes Used

These are the nodes used in the earlier demo (also in [privacy_protection_faces.yml](#)):

```
nodes:  
- input.visual:  
  source: 0  
- model.mtcnn  
- draw.mosaic_bbox  
- output.screen
```

1. Face Detection Node

As mentioned, we use the MTCNN model for face detection. It is able to detect human faces with face masks. Please take a look at the [benchmarks](#) of object detection models that are included in PeekingDuck if you would like to use a different model or model type better suited to your use case.

2. Face De-Identification Nodes

You can mosaic or blur the faces detected using the `draw.mosaic_bbox` or `draw.blur_bbox` in the run config declaration.



Fig. 1: De-identification with mosaic (left) and blur (right).

3. Adjusting Nodes

With regard to the MTCNN model, some common node behaviors that you might want to adjust are:

- **min_size**: Specifies in pixels the minimum height and width of a face to be detected. (default = 40) You may want to decrease the minimum size to increase the number of detections.
- **network_thresholds**: Specifies the threshold values for the Proposal Network (P-Net), Refine Network (R-Net), and Output Network (O-Net) in the MTCNN model. (default = [0.6, 0.7, 0.7]) Calibration is performed at each stage in which bounding boxes with confidence scores less than the specified threshold are discarded.
- **score_threshold**: Specifies the threshold value in the final output. (default = 0.7) Bounding boxes with confidence scores less than the specified threshold in the final output are discarded. You may want to lower **network_thresholds** and **score_threshold** to increase the number of detections.

In addition, some common node behaviors that you might want to adjust for the `dabble.mosaic_bbox` and `dabble.blur_bbox` nodes are:

- **mosaic_level**: Defines the resolution of a mosaic filter (*width × height*); the value corresponds to the number of rows and columns used to create a mosaic. (default = 7) For example, the default value creates a 7 × 7 mosaic filter. Increasing the number increases the intensity of pixelization over an area.
- **blur_level**: Defines the standard deviation of the Gaussian kernel used in the Gaussian filter. (default = 50) The higher the blur level, the greater the blur intensity.

7.1.2 Privacy Protection (License Plates)

Overview

Posting images or videos of our vehicles online might lead to others misusing our license plate numbers to reveal our personal information. Our solution performs license plate anonymization, and can also be used to comply with the General Data Protection Regulation (GDPR) or other data privacy laws.

Our solution automatically detects and blurs vehicles' license plates. This is explained in the [How It Works](#) section.

Demo

To try our solution on your own computer, [install](#) and run PeekingDuck with the configuration file `privacy_protection_license_plates.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/privacy_protection_license_plates.yml>
```

How It Works

There are two main components to license plate anonymization:

1. License plate detection, and
2. License plate de-identification.

1. License Plate Detection

We use open-source object detection models under the [YOLOv4](#) family to identify the locations of the license plates in an image/video feed. Specifically, we offer the YOLOv4-tiny model, which is faster, and the YOLOv4 model, which provides higher accuracy. The locations of detected license plates are returned as an array of coordinates in the form $[x_1, y_1, x_2, y_2]$, where (x_1, y_1) is the top left corner of the bounding box, and (x_2, y_2) is the bottom right. These are used to form the bounding box of each license plate detected. For more information on how to adjust the license plate detector node, check out the license plate detector configurable parameters.

2. License Plate De-Identification

To perform license plate de-identification, the areas bounded by the bounding boxes are blurred using a Gaussian blur function.

Nodes Used

These are the nodes used in the earlier demo (also in [privacy_protection_license_plates.yml](#)):

```
nodes:
- input.visual:
  source: <path/to/video with cars>
- model.yolo_license_plate
- draw.blur_bbox
- output.screen
```

1. License Plate Detection Node

By default, `model.yolo_license_plate` uses the `v4` model type to detect license plates. If faster inference speed is required, the `v4tiny` model type can be used instead.

2. License Plate De-Identification Nodes

You can choose to mosaic or blur the detected license plate using the `draw.mosaic_bbox` or `draw.blur_bbox` node in the run config declaration.



Fig. 2: De-identification with mosaic (left) and blur (right).

3. Adjusting Nodes

With regard to the YOLOv4 model, some common node configurations that you might want to adjust are:

- `score_threshold`: The bounding boxes with confidence score less than the specified score threshold are discarded. (default = 0.1)
- `iou_threshold`: The overlapping bounding boxes above the specified Intersection over Union (IoU) threshold are discarded. (default = 0.3)

In addition, some common node behaviors that you might want to adjust for the `dabble.mosaic_bbox` and `dabble.blur_bbox` nodes are:

- `mosaic_level`: Defines the resolution of a mosaic filter (*width × height*); the value corresponds to the number of rows and columns used to create a mosaic. (default = 7) For example, the default value creates a 7×7 mosaic filter. Increasing the number increases the intensity of pixelization over an area.
- `blur_level`: Defines the standard deviation of the Gaussian kernel used in the Gaussian filter. (default = 50) The higher the blur level, the greater the blur intensity.

7.1.3 Privacy Protection (People & Computer Screens)

Overview

Videos and pictures often contain people and other sensitive visual information (e.g., the display on computer screens), even though this information might not be needed at all for visual processing. Our solution performs full body anonymization and computer screen blurring to protect the identities of individuals and the sensitive information on computer screens. It can be used to comply with the General Data Protection Regulation (GDPR) or other data privacy laws.

In this example use case, we want to count the number of people in the office, but also want to avoid compromising the privacy of the office inhabitants or information displayed on computer screens.

Our solution automatically detects people, laptop and computer screens, and then blurs them. This is explained in the *How It Works* section.

Demo

To try our solution on your own computer, [install](#) and run PeekingDuck with the configuration file `privacy_protection_people_screens.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/privacy_protection_people_screens.yml>
```

How It Works

There are 2 main components to our solution:

1. Person and computer screen segmentation, and
2. Person and computer screen blurring.

1. Person and Computer Screen Segmentation

We use an open source instance segmentation model known as [Mask R-CNN](#) to obtain the masks of persons, computer screens and laptops. The masks are akin to the input frames or images, except that it only has a single channel and each pixel on the mask is a binary of either 1 or 0, which indicates whether a specific class of object is present (1) or absent (0) in a particular location of the image. For more information on how to adjust the `mask_rcnn` node, check out its configurable parameters.

2. Person and Computer Screen Blurring

To blur the people and computer screens, we pixelate or gaussian blur the image pixels where the pixel values of the relevant masks are equal to 1 (indicating presence of object).

Nodes Used

These are the nodes used in the earlier demo (also in [privacy_protection_people_screens.yml](#)):

```
nodes:
- input.visual:
  source: <path/to/video>
- model.mask_rcnn:
  detect: ["tv", "laptop"]
- draw.instance_mask:
  effect: {blur: 50}
- model.mask_rcnn:
  detect: ["person"]
- dabble.bbox_count
- draw.instance_mask:
  effect: {blur: 50}
- draw.bbox:
  show_labels: True
- draw.legend:
  show: ["count"]
- output.screen
```

This config includes the use of two `model.mask_rcnn` and `draw.instance_mask` nodes to separate the detected instances of “person” class from the “tv” and “laptop” classes, so that drawing and counting of bboxes are only performed on the “person” class. This repetition is not required if only anonymization is performed.

1. Instance Segmentation Node

In this example use case, we used the Mask R-CNN model for instance segmentation. It can detect persons as well as computer monitors. Please take a look at the [benchmarks](#) of instance segmentation models that are included in PeekingDuck if you would like to use a different model or model type better suited to your use case.

2. People and Screens De-Identification Node

The detected people and screens are blurred using the `draw.instance_mask` node.

3. Object Counting Node

`dabble.bbox_count` counts the total number of detected bounding boxes. This node has no configurable parameters.

4. Display Bounding Box Node

Then we draw bounding boxes around the detected persons using the `draw.bbox` node.

5. Person Count Display Node

The total number of detected persons are shown using the `draw.legend` node.

6. Adjusting Nodes

With regard to the Mask R-CNN model, some common node behaviors that you might want to adjust are:

- `model_type`: Defines the type of backbones to be used.
- `score_threshold`: Bounding boxes with classification score below the threshold will be discarded.
- `mask_threshold`: The confidence threshold for binarizing the masks' pixel values, whether an object is detected at a particular pixel.

In addition, some common node behaviors that you might want to adjust for the `draw.instance_mask` node are:

- `blur`: Blurs the area using this value as the “`blur_kernel_size`” parameter. Larger values gives more intense blurring.
- `mosaic`: Mosaics the area using this value as the resolution of a mosaic filter (*width* \times *height*). The value corresponds to the number of rows and columns used to create a mosaic. For example, the setting (`mosaic: 25`) creates a 25×25 mosaic filter. Increasing the number increases the intensity of pixelation over an area.

<i>Privacy Protection (Faces)</i>	<i>Privacy Protection (License Plates)</i>
<i>Privacy Protection (People and Screens)</i>	

7.2 Smart Monitoring

7.2.1 Crowd Counting

Overview

In Computer Vision (CV), crowd counting refers to the technique of counting or estimating the number of people in a crowd. This can be used to estimate the number of people attending an event, monitor crowd levels and prevent human stampedes.

Our solution utilizes CSRNet to estimate the size of a crowd. In addition, it generates a heat map that can be used to pinpoint possible bottlenecks at a venue. This is explained in the [How It Works](#) section.

Demo

To try our solution on your own computer, *install* and run PeekingDuck with the configuration file `crowd_counting.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/crowd_counting.yml>
```

You may like to try it on this [sample video](#).

How It Works

There are two main components to our solution:

1. Crowd counting, and
2. Heat map generation.

1. Crowd Counting

We use an open source crowd counting model known as [CSRNet](#) to predict the number of people in a sparse or dense crowd. The solution uses the sparse crowd model by default and can be configured to use the dense crowd model if required. The dense and sparse crowd models were trained using data from [ShanghaiTech](#) Part A and Part B respectively.

As a rule of thumb, you might want to use the dense crowd model if the people in a given image or video frame are packed shoulder to shoulder, e.g., stadiums. For more information on how to adjust the CSRNet node, check out its configurable parameters.

2. Heat Map Generation (Optional)

We generate a heat map using the density map estimated by the model. Areas that are more crowded are highlighted in red while areas that are less crowded are highlighted in blue.

Nodes Used

These are the nodes used in the earlier demo (also in `crowd_counting.yml`):

```
nodes:
- input.visual:
    source: <path/to/video with crowd>
- model.csrnet:
    model_type: dense
- draw.heat_map
- draw.legend:
    show: ["count"]
- output.screen
```

1. Crowd Counting Node

As mentioned, we use CSRNet to estimate the size of a crowd. As the models were trained to recognize congested scenes, the estimates are less accurate if the number of people is low, i.e., below ten. In such scenarios, you should consider using the *object detection models* included in our repo.

2. Heat Map Generation Node (Optional)

The heat map generation node superimposes a heat map over a given image or video frame.

3. Adjusting Nodes

Some common node behaviors that you might want to adjust are:

- `model_type`: This specifies the model to be used, i.e., `sparse` or `dense`. By default, our solution uses the sparse crowd model. As a rule of thumb, you might want to use the dense crowd model if the people in a given image or video frame are packed shoulder to shoulder, e.g., stadiums.
- `width`: This specifies the input width. By default, the width of an image will be resized to 640 for inference. The height of the image will be resized proportionally to preserve its aspect ratio. In general, decreasing the width of an image will improve inference speed. However, this might impact the accuracy of the model.

7.2.2 Object Counting (Present)

Overview

Object counting (present) is a solution within PeekingDuck's suite of *smart monitoring* use cases. It counts the number of objects detected by PeekingDuck's object detection models at the present point in time, and calculates statistics such as the cumulative average, maximum and minimum for further analytics. Up to *80 types* of objects can be counted, including humans, vehicles, animals and even household objects. Thus, this can be applied to a wide variety of scenarios, from traffic control to counting livestock.

See also:

For advanced counting tasks such as counting tracked objects over time or counting within specific zones, refer to PeekingDuck's other *smart monitoring* use cases.

In the GIF above, the count and statistics change as the number of detected persons change. This is explained in the *How It Works* section.

Demo

To try our solution on your own computer, *install* and run PeekingDuck with the configuration file `object_counting_present.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/object_counting_present.yml>
```

How It Works

There are 3 main components to this solution:

1. Object detection,
2. Count detections, and
3. Calculate statistics.

1. Object Detection

We use an open source object detection model known as [YOLOv4](#) and its smaller and faster variant known as YOLOv4-tiny to identify the bounding boxes of chosen objects we want to detect. This allows the application to identify where objects are located within the video feed. The location is returned as two x, y coordinates in the form $[x_1, y_1, x_2, y_2]$, where (x_1, y_1) is the top left corner of the bounding box, and (x_2, y_2) is the bottom right. These are used to form the bounding box of each object detected. For more information on how to adjust the yolo node, check out its configurable parameters.

2. Count Detections

To count the number of objects detected, we simply take the sum of the number of bounding boxes detected for the object category.

3. Calculate Statistics

The cumulative average, minimum and maximum over time is calculated from the count from each frame.

Nodes Used

These are the nodes used in the earlier demo (also in [object_counting_present.yml](#)):

```
nodes:
- input.visual:
  source: 0
- model.yolo:
  detect: ["person"]
- dabble.bbox_count
- dabble.statistics:
  identity: count
- draw.bbox
- draw.legend:
  show: ["count", "cum_avg", "cum_max", "cum_min"]
- output.screen
```

1. Object Detection Node

By default, the node uses the YOLOv4-tiny model for object detection, set to detect people. Please take a look at the [benchmarks](#) of object detection models that are included in PeekingDuck if you would like to use a different model or model type better suited to your use case.

2. Object Counting Node

`dabble.bbox_count` takes the detected bounding boxes and outputs the total count of bounding boxes. This node has no configurable parameters.

3. Statistics Node

The `dabble.statistics` node calculates the *cum_avg*, *cum_max* and *cum_min* from the output of the object counting node.

4. Adjusting Nodes

For the object detection model used in this demo, please see the documentation for adjustable behaviors that can influence the result of the object counting node.

For more adjustable node behaviors not listed here, check out the [API Documentation](#).

7.2.3 Object Counting (Over Time)

Overview

Object counting over time involves detecting and tracking unique objects, and incrementing the count when new objects appear. When applied to the vehicles in the GIF below, it can count the total number of vehicles passing by over a period of time, aiding transportation planning by identifying periods of peak traffic. This use case is not limited to just vehicles, as up to *80 types* of objects can be monitored (including animals), giving rise to a wide breadth of potential applications.

See also:

While it is also possible to count people over time with this use case, more accurate results can be obtained by using the [People Counting \(Over Time\)](#) use case.

See also:

If you wish to only count the number objects at an instance rather than a cumulative total over a period of time, the simpler [Object Counting \(Present\)](#) use case without requiring object tracking would be more suitable.

Object counting over time is achieved by detecting the objects using an object detection model, then tracking each unique object. As a new object appears, the number of counted objects is incremented. This is explained in the [How It Works](#) section.

Demo

To try our solution on your own computer, [install](#) and run PeekingDuck with the configuration file `object_counting_over_time.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/object_counting_over_time.yml>
```

How It Works

Object counting over time comprises three main components:

1. Object detection,
2. Tracking the outputs of object detection, and
3. Incrementing the count.

1. Object Detection

The EfficientDet model is used here to predict the bounding boxes of objects of interest. This allows the application to identify where each object is located within the video feed. The location is returned as a pair of x, y coordinates in the form $[x_1, y_1, x_2, y_2]$, where (x_1, y_1) is the top-left corner of the bounding box, and (x_2, y_2) is the bottom right.

2. Tracking the Outputs of Object Detection

An Intersection over Union (IoU) tracker adapted from [this paper](#) is used on the bounding boxes from the object detection model to produce tracked identities (IDs) for each bounding box. The IoU tracker continues a track by associating the detection with the highest IoU to the last detection in the previous frame. For example, Car 8 in frame n continues to be tracked as Car 8 in frame $n+1$ as both instances of Car 8 are within close proximity (high IoU) of each other. This assumes that the object detector correctly predicts a bounding box per frame for each object to be tracked, and also assumes that the frame rate of the video is high enough to allow unambiguous IoU overlaps between consecutive frames.

Another available option is the Minimum Output Sum of Squared Error (MOSSE) tracker which we have adapted from the OpenCV package. It is a correlation filter based tracker which uses Fast Fourier Transform (FFT) to perform operations in the frequency domain, reducing computational complexity. More details can be found from [this paper](#).

3. Incrementing the Count

Monotonically increasing integer IDs beginning from 0 are assigned to new unique objects. For example, the first tracked object is assigned an ID of 0, the second tracked object is assigned an ID of 1, and so on. Thus the total number of unique objects that have appeared in the entire duration is simply the cumulative maximum.

Nodes Used

These are the nodes used in the earlier demo (also in `object_counting_over_time.yml`):

```
nodes:
- input.visual:
  source: <path/to/video with cars>
- model.efficientdet:
  detect: ["car"]
- dabble.tracking:
  tracking_type: "iou"
- dabble.statistics:
  maximum: obj_attrs["ids"]
- draw.bbox
- draw.tag:
  show: ["ids"]
- draw.legend:
  show: ["cum_max"]
- output.screen
```

1. Object Detection Node

In the demo, the `model.efficientdet` node is used for object detection, set to detect cars. As mentioned in the earlier [How It Works](#) section, for object tracking to work well, the upstream object detector needs to produce predictions which are as accurate as possible. Please take a look at the [benchmarks](#) of object detection models that are included in PeekingDuck if you would like to use a different model or model type better suited to your use case.

2. Tracking Node

The `dabble.tracking` node used here is not an AI model but uses heuristics, hence it falls under the category of dabble nodes instead of model nodes. It needs to be paired with an upstream object detector node, but this also gives it a key advantage - it can track any of the *80 types* of detectable objects. In contrast, the [People Counting \(Over Time\)](#) use case uses a single model node purpose-built for both human detection and tracking, giving it more accuracy but limiting its usage to only humans.

3. Statistics Node

The `dabble.statistics` node retrieves the maximum detected ID for each frame. If the ID exceeds the previous maximum, the `cum_max` (cumulative maximum) is updated. As monotonically increasing integer IDs beginning from 0 are assigned to new unique objects, the maximum ID is equal to the total number of unique objects over time.

4. Adjusting Nodes

Some common node behaviors that you might need to adjust are:

For `model.efficientdet`:

- `model_type`: 0, 1, 2, 3, or 4. The larger the number, the higher the accuracy, at the cost of inference speed.
- `detect`: Object class IDs to be detected. Refer to [Object Detection IDs table](#) for the class IDs for each model.

For `dabble.tracking`:

- `tracking_type`: Choose either ["iou", "mosse"], described earlier in the [How It Works](#) section.

For more adjustable node behaviors not listed here, check out the [API Documentation](#).

Counting Objects Within Zones

It is possible to extend this use case with the [Zone Counting](#) use case. For example, if the road were a dual carriageway and we are only interested counting the vehicles on one side of the road, we could split the video into 2 different zones and only count the vehicles within the chosen zone. An example of how this can be done is given in the [Tracking People within a Zone](#) tutorial.

7.2.4 People Counting (Over Time)

Overview

People counting over time involves detecting and tracking different persons, and incrementing the count when a new person appears. This use case can reduce dependency on manual counting, and be applied to areas such as retail analytics, queue management, or occupancy monitoring.

Our solution automatically detects, tracks and counts people over time. This is explained in the [How It Works](#) section.

Demo

To try our solution on your own computer, *install* and run PeekingDuck with the configuration file `people_counting_over_time.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/people_counting_over_time.yml>
```

You may like to try it on this [sample video](#).

How It Works

People counting over time comprises three main components:

1. Human detection,
2. Appearance embedding tracking, and
3. Incrementing the count.

1. Human Detection

We use an open source detection model known as **JDE** to detect persons. JDE has been trained on pedestrian detection and person search datasets. This allows the application to identify the locations of persons in a video feed. Each of these locations is represented as a pair of x, y coordinates in the form $[x_1, y_1, x_2, y_2]$, where (x_1, y_1) is the top left corner of the bounding box, and (x_2, y_2) is the bottom right. These are used to form the bounding box of each person detected. For more information on how to adjust the JDE node, check out the JDE configurable parameters.

2. Appearance Embedding Tracking

To learn appearance embeddings for tracking, a metric learning algorithm with triplet loss is used. Observations are assigned to tracklets using the Hungarian algorithm. The Kalman filter is used to smooth the trajectories and predict the locations of previous tracklets in the current frame. The model outputs an ID for each detection based on the appearance embedding learned.

3. Incrementing the Count

Monotonically increasing integer IDs beginning from 0 are assigned to new unique persons. For example, the first tracked person is assigned an ID of 0, the second tracked person is assigned an ID of 1, and so on. Thus the total number of unique persons that have appeared in the entire duration is simply the cumulative maximum.

Nodes Used

These are the nodes used in the earlier demo (also in `people_counting_over_time.yml`):

```
nodes:
- input.visual:
    source: <path/to/video with people>
- model.jde
- dabble.statistics:
    maximum: obj_attrs["ids"]
- draw.bbox
```

(continues on next page)

(continued from previous page)

```

- draw.tag:
  show: ["ids"]
- draw.legend:
  show: ["cum_max"]
- output.screen

```

1. JDE Node

This node employs a single network to **simultaneously** output detection results and the corresponding appearance embeddings of the detected boxes. Therefore JDE stands for Joint Detection and Embedding. Please take a look at the [benchmarks](#) of object tracking models that are included in PeekingDuck if you would like to use a different model or model type better suited to your use case.

2. Statistics Node

The `dabble.statistics` node retrieves the maximum detected ID for each frame. If the ID exceeds the previous maximum, the `cum_max` (cumulative maximum) is updated. As monotonically increasing integer IDs beginning from 0 are assigned to new unique persons, the maximum ID is equal to the total number of unique persons over time.

3. Adjusting Nodes

With regard to the `model.jde` node, some common behaviors that you might want to adjust are:

- `iou_threshold`: Specifies the threshold value for Intersection over Union of detections (default = 0.5).
- `score_threshold`: Specifies the threshold values for the detection confidence (default = 0.5). You may want to lower this value to increase the number of detections.
- `nms_threshold`: Specifies the threshold value for non-maximal suppression (default = 0.4). You may want to lower this value to increase the number of detections.
- `min_box_area`: Specifies the minimum value for area of detected bounding box. Calculated by $width \times height$ (default = 200).
- `track_buffer`: Specifies the threshold to remove track if track is lost for more frames than this value (default = 30).

Counting People Within Zones

It is possible to extend this use case with the [Zone Counting](#) use case. For example, if a CCTV footage shows the entrance of a mall as well as a road, and we are only interested to apply people counting to the mall entrance, we could split the video into 2 different zones and only count the people within the chosen zone. An example of how this can be done is given in the [Tracking People within a Zone](#) tutorial.

7.2.5 Zone Counting

Overview

Zone counting creates different zones within a single image and counts the number of objects within each zone separately. This is useful in many applications, such as counting vehicles travelling on one side of a road, or counting the shoppers entering a mall.

See also:

To only count objects within a single zone and ignore all other objects, see the [Tracking People within a Zone](#) tutorial.

Zone counting is done by counting the number of objects detected by the object detection models that fall within the specified zones. For example, we can count the number of people in the blue and red zones, as shown in the GIF above. This is explained in the *How It Works* section.

Demo

To try our solution on your own computer, *install* and run PeekingDuck with the configuration file `zone_counting.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/zone_counting.yml>
```

How It Works

There are three main components to obtain the zone counts:

1. The detections from the object detection model, which are the bounding boxes,
2. The bottom midpoint of the bounding boxes, derived from the bounding boxes, and
3. The zones, which can be set in the `dabble.zone_count` configurable parameters.

1. Object Detection

We use an open source object detection model known as `YOLOv4` and its smaller and faster variant known as `YOLOv4-tiny` to identify the bounding boxes of objects we want to detect. This allows the application to identify where objects are located within the video feed. The location is returned as a pair of *x, y coordinates* in the form $[x_1, y_1, x_2, y_2]$, where (x_1, y_1) is the top left corner of the bounding box, and (x_2, y_2) is the bottom right. These are used to form the bounding box of each object detected. For more information on how to adjust the `yolo` node, check out its configurable parameters.

2. Bounding Box to Bottom Midpoint

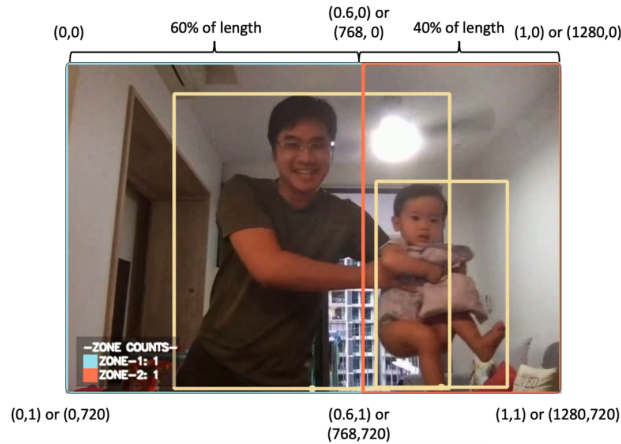
Given the top left (x_1, y_1) and bottom right (x_2, y_2) coordinates of each bounding box, the bottom midpoint (x_{bm}, y_{bm}) can be computed by taking lowest *y* coordinate $y_{bm} = y_2$, and the midpoint of the *x* coordinates $x_{bm} = (x_1 + x_2)/2$.

We found that using the bottom midpoint is the most efficient way to tell if something is in a specified zone. We attribute this to the use of the top-down or angled camera footages, which are commonly found in the use cases. The bottom midpoints of the bounding boxes usually correspond to the locations of the objects in these footages.

3. Zones

Zones are created by specifying the *x, y* coordinates of all the corner points that form the area of the zone **in a clockwise direction**. The coordinates can be in either fractions of the resolution or pixels. As an example, blue zone in the *zone counting GIF* was created using the following zone:

```
[[0, 0], [0.6, 0], [0.6, 1], [0, 1]]
```



Given a resolution of 1280 by 720, these correspond to the top-left of the image, 60% of the length at the top of the image, 60% of the length at the bottom of the image, and the bottom-left of the image. These points form the rectangular blue zone in a clockwise direction.

Note that because the x , y coordinates are fractions of the image resolution, the resolution config for `dabble.zone_count` needs to be set correctly.

For finer control over the exact coordinates, the pixel coordinates can be used instead. Using the same example, the blue zone can be created using the following zone configuration:

```
[[0, 0], [768, 0], [768, 720], [0, 720]]
```

When using pixel coordinates, the resolution is not needed. However, users should check to ensure that the pixel coordinates given fall within the image resolution so that the zone will work as intended.

Elaboration for this adjustment can be found the “4. *Adjusting Nodes*” section.

Note: Zones do not have to be rectangular in shape. They can be of any polygonal shape, dictated by the number and position of the x , y coordinates set in a zone.

4. Zone Counts

Given the bottom midpoints of all detected objects, we check if the points fall within the areas of the specified zones. If it falls inside any zone, an object count is added for that specific zone. This continues until all objects detected are accounted for, which gives the final count of objects in each specified zone.

Nodes Used

These are the nodes used in the earlier demo (also in `zone_counting.yml`):

```
nodes:
- input.visual:
  source: 0
- model.yolo:
  detect: ["person"]
- dabble.bbox_to_btm_midpoint
- dabble.zone_count:
  resolution: [1280, 720] # Adjust this to your camera's input resolution
  zones: [
```

(continues on next page)

(continued from previous page)

```

[[0, 0], [0.6, 0], [0.6, 1], [0, 1]],
[[0.6, 0], [1, 0], [1, 1], [0.6, 1]]
]
- draw.bbox
- draw.btm_midpoint
- draw.zone
- draw.legend:
  show: ["zone_count"]
- output.screen

```

1. Object Detection Node

By default, the node uses the YOLOv4-tiny model for object detection, set to detect people. Please take a look at the [benchmarks](#) of object detection models that are included in PeekingDuck if you would like to use a different model or model type better suited to your use case.

2. Bottom Midpoint Node

The bottom midpoint node is called by including `dabble.bbox_to_btm_midpoint` in the pipeline config declaration. This outputs the bottom midpoints of all detected bounding boxes. The node has no configurable parameters.

3. Zone Counting Node

The zone counting node is called by including `dabble.zone_count` in the run config declaration. This uses the bottom midpoints of all detected bounding boxes and outputs the number of object counts in each specified zone. The node configurable parameters can be found below.

4. Adjusting Nodes

The zone counting detections depend on the configuration set in the object detection models, such as the type of object to detect, etc. For the object detection model used in this demo, please see the `yolo` node documentation for adjustable behaviors that can influence the result of the zone counting node.

With regards to the zone counting node, some common node behaviors that you might need to adjust are:

- **resolution:** If you are planning to use fractions to set the coordinates for the area of the zone, the resolution should be set to the image/video/livestream resolution used.
- **zones:** Used to specify the different zones which you would like to set. The coordinates for each zone are given in a list in a clockwise order. See the [Nodes Used](#) section on how to properly configure multiple zones.

For more adjustable node behaviors not listed here, check out the [API Documentation](#).

<i>Zone Counting</i>	<i>Crowd Counting</i>
<i>Object Counting (Over Time)</i>	<i>People Counting (Over Time)</i>
<i>Object Counting (Present)</i>	

7.3 COVID-19 Prevention and Control

7.3.1 Face Mask Detection

Overview

Wearing of face masks in public places can help prevent the spread of COVID-19 and other infectious diseases. AI Singapore has developed a solution that checks whether or not a person is wearing a face mask. This can be used in places such as malls or shops to ensure that visitors adhere to the guidelines.

We have trained a custom YOLOv4 model to detect whether or not a person is wearing a face mask. This is explained in the *How It Works* section.

Demo

To try our solution on your own computer, *install* and run PeekingDuck with the configuration file `face_mask_detection.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/face_mask_detection.yml>
```

How It Works

The main component is the detection of face mask using the custom YOLOv4 model.

Face Mask Detection

We use an open source object detection model known as **YOLOv4** and its smaller and faster variant known as YOLOv4-tiny to identify the bounding boxes of human faces with and without face masks. This allows the application to identify the locations of faces and their corresponding classes (`no_mask = 0` or `mask = 1`) in a video feed. Each of these locations are represented as a pair of x, y coordinates in the form $[x_1, y_1, x_2, y_2]$, where (x_1, y_1) is the top-left corner of the bounding box, and (x_2, y_2) is the bottom right. These are used to form the bounding box of each human face detected.

The `model.yolo_face` node detects human faces with and without face masks using the YOLOv4-tiny model by default. The classes are differentiated by the labels and the colors of the bounding boxes when multiple faces are detected. For more information on how to adjust the `yolo_face` node, check out its configurable parameters.

Nodes Used

These are the nodes used in the earlier demo (also in `face_mask_detection.yml`):

```
nodes:
- input.visual:
  source: 0
- model.yolo_face
- draw.bbox:
  show_labels: true
- output.screen
```

1. Face Mask Detection Node

The `model.yolo_face` node is used for face detection and to classify if the face is masked or unmasked. To simply detect faces without needing to classify if the face is masked, you can also consider the `model.mtcnn` node.

2. Adjusting Nodes

Some common node behaviors that you might want to adjust are:

- `model_type`: This specifies the variant of YOLOv4 to be used. By default, the *v4tiny* model is used, but for better accuracy, you may want to try the *v4* model.
- `detect`: This specifies the class to be detected where `no_mask = 0` and `mask = 1`. By default, the model detects faces with and without face masks (default = [0, 1]).
- `score_threshold`: This specifies the threshold value. Bounding boxes with confidence score lower than the threshold are discarded. You may want to lower the threshold value to increase the number of detections.

7.3.2 Group Size Checking

Overview

As part of COVID-19 measures, the Singapore Government has set restrictions on the group sizes of social gatherings. AI Singapore has developed a vision-based [group size checker](#) that checks if the group size limit has been violated. This can be used in many places, such as in malls to ensure that visitors adhere to guidelines, or in workplaces to ensure employees' safety.

To check if individuals belong to a group, we check if the physical distance between them is close. The most accurate way to measure distance is to use a 3D sensor with depth perception, such as a RGB-D camera or a [LiDAR](#). However, most cameras such as CCTVs and IP cameras usually only produce 2D videos. We developed heuristics that are able to give an approximate measure of physical distance from 2D videos, addressing this limitation. This is further elaborated in the [How It Works](#) section.

Demo

To try our solution on your own computer, [install](#) and run PeekingDuck with the configuration file `group_size_checking.yml` as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/group_size_checking.yml>
```

How It Works

There are three main components to obtain the distance between individuals:

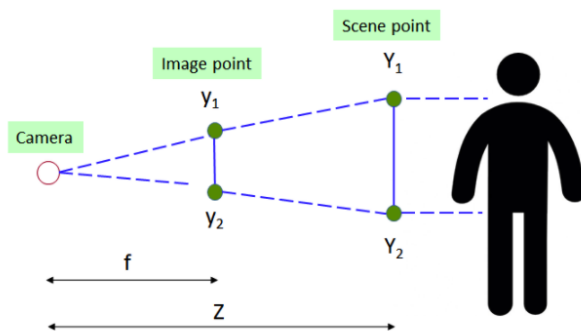
1. Human pose estimation using AI,
2. Depth and distance approximation, and
3. Linking individuals to groups.

1. Human Pose Estimation

We use an open source human pose estimation model known as [PoseNet](#) to identify key human skeletal points. This allows the application to identify where individuals are located within the video feed. The coordinates of the various skeletal points will then be used to determine the distance between individuals.

2. Depth and Distance Approximation

To measure the distance between individuals, we have to estimate the 3D world coordinates from the keypoints in 2D coordinates. To achieve this, we compute the depth Z from the x , y coordinates using the relationship below:



Using the similar triangle rule, we are able to compute Z .

$$\frac{y_1 - y_2}{Y_1 - Y_2} = \frac{f}{Z}$$

where:

- Z = depth or distance of scene point from camera
- f = focal length of camera

- y = y position of image point
- Y = y position of scene point

$Y_1 - Y_2$ is a reference or “ground truth length” that is required to obtain the depth. After numerous experiments, it was decided that the optimal reference length would be the average height of a human torso (height from human hip to center of face). Width was not used as this value has high variance due to the different body angles of an individual while facing the camera.

3. Linking Individuals to Groups

Once we have the 3D world coordinates of the individuals in the video, we can compare the distances between each pair of individuals. If they are close to each other, we assign them to the same group. This is a dynamic connectivity problem and we use the [quick find algorithm](#) to solve it.

Nodes Used

These are the nodes used in the earlier demo (also in [group_size_checking.yml](#)):

```
nodes:
- input.visual:
  source: 0
- model.posenet
- dabble.keypoints_to_3d_loc:
  focal_length: 1.14
  torso_factor: 0.9
- dabble.group_nearby_objs:
  near_threshold: 1.5
- dabble.check_large_groups:
  group_size_threshold: 2
- draw.pose
- draw.group_bbox_and_tag
- output.screen
```

1. Pose Estimation Model

By default, we are using the PoseNet model with a ResNet backbone for pose estimation. Please take a look at the [benchmarks](#) of pose estimation models that are included in PeekingDuck if you would like to use a different model or model type better suited to your use case.

2. Adjusting Nodes

Some common node behaviors that you might need to adjust are:

- `focal_length` & `torso_factor`: We calibrated these settings using a Logitech c170 webcam, with 2 individuals of heights about 1.7m. We recommend running a few experiments on your setup and calibrate these accordingly.
- `near_threshold`: The maximum distance between 2 individuals, in meters, for them to be considered to be part of a group.
- `group_size_threshold`: The acceptable group size limit.

For more adjustable node behaviors not listed here, check out the [API Documentation](#).

3. Using Object Detection (Optional)

It is possible to use [object detection models](#) instead of pose estimation. To do so, replace the model node accordingly, and replace the node `dabble.keypoints_to_3d_loc` with `dabble.bbox_to_3d_loc`. The reference or “ground truth length” in this case would be the average height of a human, multiplied by a small factor.

You might need to use this approach if running on a resource-limited device such as a Raspberry Pi. In this situation, you'll need to use the lightweight models; we find lightweight object detectors are generally better than lightweight pose estimation models in detecting humans.

The trade-off here is that the estimated distance between individuals will be less accurate. This is because for object detectors, the bounding box will be compared with the average height of a human, but the bounding box height decreases if the person is sitting down or bending over.

Using with Social Distancing

To combat COVID-19, individuals are encouraged to maintain physical distance from each other. We've developed a social distancing tool that checks if individuals are too close to each other.

The nodes for social distancing can be stacked with group size checker, to perform both at the same time. Check out the [Social Distancing use case](#) to find out which nodes are used.

7.3.3 Social Distancing

Overview

To support the fight against COVID-19, AI Singapore developed a solution to encourage individuals to maintain physical distance from each other. This can be used in many places, such as in malls to encourage social distancing in long queues, or in workplaces to ensure employees' well-being. An example of the latter is [HP Inc.](#), which collaborated with us to deploy this solution on edge devices in its manufacturing facility in Singapore.

The most accurate way to measure distance is to use a 3D sensor with depth perception, such as a RGB-D camera or a [LiDAR](#). However, most cameras such as CCTVs and IP cameras usually only produce 2D videos. We developed heuristics that are able to give an approximate measure of physical distance from 2D videos, addressing this limitation. This is explained in the [How It Works](#) section.

Demo

To try our solution on your own computer, [install](#) and run PeekingDuck with the configuration file [social_distancing.yml](#) as shown:

Terminal Session

```
[~user] > peekingduck run --config_path <path/to/social_distancing.yml>
```

How It Works

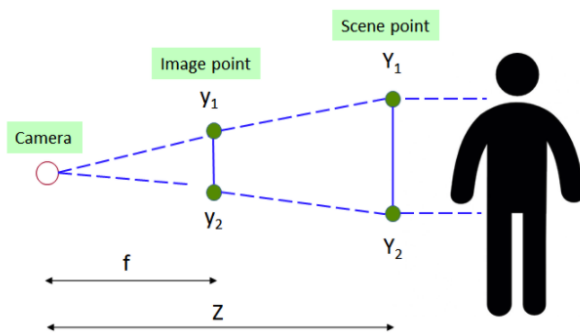
There are two main components to obtain the distance between individuals: #. Human pose estimation using AI, and #. Depth and distance approximation using heuristics.

1. Human Pose Estimation

We use an open source human pose estimation model known as [PoseNet](#) to identify key human skeletal points. This allows the application to identify where individuals are located within the video feed. The coordinates of the various skeletal points will then be used to determine the distance between individuals.

2. Depth and Distance Approximation

To measure the distance between individuals, we have to estimate the 3D world coordinates from the keypoints in 2D coordinates. To achieve this, we compute the depth Z from the x , y coordinates using the relationship below:



Using the similar triangle rule, we are able to compute Z .

$$\frac{y_1 - y_2}{Y_1 - Y_2} = \frac{f}{Z}$$

where:

- Z = depth or distance of scene point from camera
- f = focal length of camera
- y = y position of image point
- Y = y position of scene point

$Y_1 - Y_2$ is a reference or “ground truth length” that is required to obtain the depth. After numerous experiments, it was decided that the optimal reference length would be the average height of a human torso (height from human hip to center of face). Width was not used as this value has high variance due to the different body angles of an individual while facing the camera.

Once we have the 3D world coordinates of the individuals in the video, we can compare the distances between each pair of individuals and check if they are too close to each other.

Nodes Used

These are the nodes used in the earlier demo (also in `social_distancing.yml`):

```
nodes:
- input.visual:
  source: 0
- model.posenet
- dabble.keypoints_to_3d_loc:
  focal_length: 1.14
  torso_factor: 0.9
- dabble.check_nearby_objs:
  near_threshold: 1.5
  tag_msg: "TOO CLOSE!"
- draw.pose
- draw.tag:
  show: ["flags"]
- output.screen
```

1. Pose Estimation Model

By default, we are using the PoseNet model with a ResNet backbone for pose estimation. Please take a look at the [benchmarks](#) of pose estimation models that are included in PeekingDuck if you would like to use a different model or model type better suited to your use case.

2. Adjusting Nodes

Some common node behaviors that you might need to adjust are:

- `focal_length` & `torso_factor`: We calibrated these settings using a Logitech c170 webcam, with 2 individuals of heights about 1.7m. We recommend running a few experiments on your setup and calibrate these accordingly.
- `tag_msg`: The message to show when individuals are too close.
- `near_threshold`: The minimum acceptable distance between 2 individuals, in meters. For example, if the threshold is set at 1.5m, and 2 individuals are standing 2.0m apart, `tag_msg` doesn't show as they are standing further apart than the threshold. The larger this number, the stricter the social distancing.

For more adjustable node behaviors not listed here, check out the [API Documentation](#).

3. Using Object Detection (Optional)

It is possible to use [object detection models](#) instead of pose estimation. To do so, replace the model node accordingly, and replace the `dabble.keypoints_to_3d_loc` node with `dabble.bbox_to_3d_loc`. The reference or “ground truth length” in this case would be the average height of a human, multiplied by a small factor.

You might need to use this approach if running on a resource-limited device such as a Raspberry Pi. In this situation, you'll need to use the lightweight models, and we find that lightweight object detectors are generally better than lightweight pose estimation models in detecting humans.

The trade-off here is that the estimated distance between individuals will be less accurate. This is because for object detectors, the bounding box will be compared with the average height of a human, but the bounding box height decreases if the person is sitting down or bending over.

Using with Group Size Checker

As part of COVID-19 measures, the Singapore Government has set restrictions on the group sizes of social gatherings. We've developed a [group size checker](#) that checks if the group size limit has been violated.

The nodes for group size checker can be stacked with social distancing, to perform both at the same time. Check out the [Group Size Checking use case](#) to find out which nodes are used.

<i>Social Distancing</i>	<i>Group Size Checking</i>
<i>Face Mask Detection</i>	

FAQ AND TROUBLESHOOTING

8.1 How can I post-process and visualize model outputs?

The common output of all `model` nodes is *bboxes*. *bboxes* can be used for subsequent actions like counting (`dabble.bbox_count`), drawing (`draw.bbox`), tagging (`draw.tag`), etc. You can also create custom nodes which take *bboxes* as an input to visualize your results.

8.2 How can I dynamically use all prior outputs as the input at run time?

Specifying “*all*” as the input allows the node to receive all prior outputs as the input. This is used by nodes such as `draw.legend` and `output.csv_writer`.

8.3 How do I debug custom nodes?

You can add code in custom nodes to print the contents of their inputs. For more info, please see the tutorial on *debugging*.

8.4 Why does `input.visual` progress stop before 100%?

`input.visual` provides progress information if it is able to get a total frame count for the input. This number is obtained using `opencv`’s `CV_CAP_PROP_FRAME_COUNT` API, which attempts to estimate the total frame count using the input media’s metadata duration and FPS. However, the total frame count is only an estimate. It is not guaranteed to be accurate because it is affected by potential errors, such as frame corruption, video decoder failure, inaccurate FPS, and rounding errors.

8.5 Why does the output screen flash briefly and disappear on my second run?

If you are running PeekingDuck on the Windows Subsystem for Linux (WSL), this erroneous behavior may be caused by a WSL bug where the key buffer is not flushed. Please refer to this [GitHub issue](#) for more details.

GLOSSARY

The following are built-in data types recognized by PeekingDuck nodes. Users can define custom data types when working with custom nodes.

(input) all (Any)

This data type contains all the outputs from preceding nodes, granting a large degree of flexibility to nodes that receive it. Examples of such nodes include `draw.legend`, `dabble.statistics`, and `output.csv_writer`.

bboxes (numpy.ndarray)

A NumPy array of shape $(N, 4)$ containing normalized bounding box coordinates of N detected objects. Each bounding box is represented as (x_1, y_1, x_2, y_2) where (x_1, y_1) is the top-left corner and (x_2, y_2) is the bottom-right corner. The order corresponds to *bbox_labels* and *bbox_scores*.

bbox_labels (numpy.ndarray)

A NumPy array of shape (N) containing strings representing the labels of detected objects. The order corresponds to *bboxes* and *bbox_scores*.

bbox_scores (numpy.ndarray)

A NumPy array of shape (N) containing confidence scores $[0, 1]$ of detected objects. The order corresponds to *bboxes* and *bbox_labels*.

btm_midpoint (List[Tuple[int, int]])

A list of tuples each representing the (x, y) coordinates of the bottom middle of a bounding box for use in zone analytics. The order corresponds to *bboxes*.

count (int)

An integer representing the number of counted objects.

cum_avg (float)

Cumulative average of an attribute over time.

cum_max (float | int)

Cumulative maximum of an attribute over time.

cum_min (float | int)

Cumulative minimum of an attribute over time.

density_map (numpy.ndarray)

A NumPy array of shape (H, W) representing the number of persons per pixel. H and W are the height and width of the input image, respectively. The sum of the array is the estimated total number of people.

filename (str)

The filename of video/image being read.

fps (float)

A float representing the Frames Per Second (FPS) when processing a live video stream or a recorded video.

img (numpy.ndarray)

A NumPy array of shape $(height, width, channels)$ containing the image data in BGR format.

keypoints (numpy.ndarray)

A NumPy array of shape $(N, K, 2)$ containing the (x, y) coordinates of detected poses where N is the number of detected poses, and K is the number of individual keypoints. Keypoints with low confidence scores (below threshold) will be replaced by -1 .

keypoint_conns (numpy.ndarray)

A NumPy array of shape $(N, D'_n, 2, 2)$ containing the (x, y) coordinates of adjacent keypoint pairs where N is the number of detected poses, and D'_n is the number of valid keypoint pairs for the n -th pose where both keypoints are detected.

keypoint_scores (numpy.ndarray)

A NumPy array of shape (N, K) containing the confidence scores of detected poses where N is the number of detected poses and K is the number of individual keypoints. The confidence score has a range of $[0, 1]$.

large_groups (List[int])

A list of integers representing the group IDs of groups that have exceeded the size threshold.

masks (numpy.ndarray)

A NumPy array of shape (N, H, W) containing N detected binarized masks where H and W are the height and width of the masks. The order corresponds to [bbox_labels](#).

(input) none

No inputs required.

(output) none

No outputs produced.

obj_3D_locs (List[numpy.ndarray])

A list of N NumPy arrays representing the 3D coordinates (x, y, z) of an object associated with a detected bounding box.

obj_attrs (Dict[str, Any])

A dictionary of attributes associated with each bounding box, in the same order as [bboxes](#). Different nodes that produce this [obj_attrs](#) output type may contribute different attributes.

pipeline_end (bool)

A boolean that evaluates to True when the pipeline is completed. Suitable for operations that require the entire inference pipeline to be completed before running.

saved_video_fps (float)

FPS of the recorded video, upon filming.

zones (List[List[Tuple[float, ...]])

A nested list of Z zones. Each zone is described by 3 or more points which contains the (x, y) coordinates forming the boundary of a zone. The order corresponds to [zone_count](#).

zone_count (List[int])

A list of integers representing the count of a pre-selected object class (for example, “person”) detected in each specified zone. The order corresponds to [zones](#).

Deprecated since version 1.2.0: [obj_tags \(List\[str\]\)](#) is deprecated and now subsumed under [obj_attrs](#). [dabble.check_nearby_objs](#) now accesses this attribute by using the `flags` key of [obj_attrs](#). [draw.tag](#) has been refactored for more drawing flexibility by accepting [obj_attrs](#) as input.

Deprecated since version 1.2.0: [obj_groups \(List\[int\]\)](#) is deprecated and now subsumed under [obj_attrs](#). Affected nodes ([dabble.group_nearby_objs](#), [dabble.check_large_groups](#), and [draw.group_bbox_and_tag](#)) now access this attribute by using the `groups` key of [obj_attrs](#).

API DOCUMENTATION

<i>input</i>	Reads data from a given input.
<i>augment</i>	Performs image processing.
<i>model</i>	Deep learning model nodes for computer vision.
<i>dabble</i>	Algorithms that perform calculations/heuristics on the outputs of <i>model</i> .
<i>draw</i>	Draws results/outputs to an image.
<i>output</i>	Writes/displays the outputs of the pipeline.

10.1 input

Description

Reads data from a given input.

Deprecated since version 1.2.0: `input.live` and `input.recorded` are deprecated. They have been replaced by the `input.visual` node.

Modules

10.2 augment

Description

Performs image processing. This can be done before or after the model.

Modules

10.3 model

Description

Deep learning model nodes for computer vision.

Modules

10.4 dabble

Description

Algorithms that perform calculations/heuristics on the outputs of `model`.

Modules

10.5 draw

Description

Draws results/outputs to an image.

Deprecated since version 1.2.0: `draw.image_processor` is deprecated, and replaced by the nodes `augment.brightness` and `augment.contrast`.

Changed in version 2.0.0: `draw.poses` has been renamed to `draw.pose`.

Modules

10.6 output

Description

Writes/displays the outputs of the pipeline.

Modules

PYTHON MODULE INDEX

a

augment, 135

d

dabble, 136

draw, 136

i

input, 135

m

model, 136

o

output, 137

Symbols

(input) all, [133](#)
 (input) none, [134](#)
 (output) none, [134](#)

A

augment
 module, [135](#)

B

bbox_labels, [133](#)
 bbox_scores, [133](#)
 bboxes, [133](#)
 btm_midpoint, [133](#)

C

count, [133](#)
 cum_avg, [133](#)
 cum_max, [133](#)
 cum_min, [133](#)

D

dabble
 module, [136](#)
 density_map, [133](#)
 draw
 module, [136](#)

F

filename, [133](#)
 fps, [133](#)

I

img, [134](#)
 input
 module, [135](#)

K

keypoint_conns, [134](#)
 keypoint_scores, [134](#)
 keypoints, [134](#)

L

large_groups, [134](#)

M

masks, [134](#)
 model
 module, [136](#)
 module
 augment, [135](#)
 dabble, [136](#)
 draw, [136](#)
 input, [135](#)
 model, [136](#)
 output, [137](#)

O

obj_3D_locs, [134](#)
 obj_attrs, [134](#)
 output
 module, [137](#)

P

pipeline_end, [134](#)

S

saved_video_fps, [134](#)

Z

zone_count, [134](#)
 zones, [134](#)