
Picamera Documentation

Release 1.8

Dave Hughes

September 06, 2014

1	Links	3
2	Table of Contents	5
2.1	Python 2.7+ Installation	5
2.2	Python 3.2+ Installation	8
2.3	Quick Start	11
2.4	Basic Recipes	13
2.5	Advanced Recipes	27
2.6	Frequently Asked Questions (FAQ)	47
2.7	Camera Hardware	49
2.8	Deprecated Functionality	53
2.9	API Reference	59
2.10	Array Extensions	92
2.11	Change log	97
2.12	License	102
3	Indices and tables	105
	Python Module Index	107

This package provides a pure Python interface to the [Raspberry Pi camera](#) module for Python 2.7 (or above) or Python 3.2 (or above).

Links

- The code is licensed under the [BSD license](#)
- The [source code](#) can be obtained from GitHub, which also hosts the [bug tracker](#)
- The [documentation](#) (which includes installation, quick-start examples, and lots of code recipes) can be read on ReadTheDocs
- Packages can be downloaded from [PyPI](#), but reading the installation instructions is more likely to be useful

Table of Contents

2.1 Python 2.7+ Installation

There are several ways to install `picamera` under Python 2.7 (or above), each with their own advantages and disadvantages. Have a read of the sections below and select an installation method which conforms to your needs.

2.1.1 Firmware upgrades

The behaviour of the Pi's camera module is dictated by the Pi's firmware. Over time, considerable work has gone into fixing bugs and extending the functionality of the Pi's camera module through new firmware releases. Whilst the `picamera` library attempts to maintain backward compatibility with older Pi firmwares, it is only tested against the latest firmware at the time of release, and not all functionality may be available if you are running an older firmware. As an example, the `annotate_text` attribute relies on a recent firmware; older firmwares lacked the functionality.

To keep your firmware up to date use the following command, and reboot once successfully complete:

```
$ sudo rpi-update
```

You can determine the revision of your current firmware with the following command:

```
$ uname -a
```

The firmware revision is the number after the #:

```
Linux kermit 3.12.26+ #707 PREEMPT Sat Aug 30 17:39:19 BST 2014 armv6l GNU/Linux
                    /
                    /
firmware revision --+
```

2.1.2 Raspbian installation

If you are using the [Raspbian](#) distro, it is best to install `picamera` using the system's package manager: `apt`. This will ensure that `picamera` is easy to keep up to date, and easy to remove should you wish to do so. It will also make `picamera` available for all users on the system. To install `picamera` using `apt` simply:

```
$ sudo apt-get update
$ sudo apt-get install python-picamera
```

To upgrade your installation when new releases are made you can simply use `apt`'s normal upgrade procedure:

```
$ sudo apt-get update
$ sudo apt-get upgrade
```

If you ever need to remove your installation:

```
$ sudo apt-get remove python-picamera
```

Note: If you are using a recent installation of Raspbian, you may find that the python-picamera package is already installed (it is included by default in recent versions of [NOOBS](#)).

Note: The release of the picamera package on Raspbian lags behind the release of picamera packages on PyPI (see *System installation* below) by a few days. If you find yourself unable to upgrade immediately following a release announcement, you can either switch to a PyPI installation or wait a few days for the Raspbian packages to arrive in the repository.

2.1.3 User installation

This is the simplest (non-apt) form of installation, but bear in mind that it will only work for the user you install under. For example, if you install as the `pi` user, you will only be able to use picamera as the `pi` user. If you run python as root (e.g. with `sudo python`) it will not find the module. See *System installation* below if you require a root installation.

To install as your current user:

```
$ sudo apt-get install python-pip
$ pip install --user picamera
```

If you wish to use the classes in the `picamera.array` module then specify the “array” option which will pull in numpy as a dependency (be warned that building numpy takes a *long* time on a Pi):

```
$ pip install --user "picamera[array]"
```

Note that `pip` is **not** run with `sudo`; this is deliberate. To upgrade your installation when new releases are made:

```
$ pip install --user -U picamera
```

If you ever need to remove your installation:

```
$ pip uninstall picamera
```

2.1.4 System installation

A system installation will make picamera accessible to all users (in contrast to the user installation). It is as simple to perform as the user installation and equally easy to keep updated. To perform the installation:

```
$ sudo apt-get install python-pip
$ sudo pip install picamera
```

If you wish to use the classes in the `picamera.array` module then specify the “array” option which will pull in numpy as a dependency (be warned that building numpy takes a *long* time on a Pi):

```
$ sudo pip install "picamera[array]"
```

To upgrade your installation when new releases are made:

```
$ sudo pip install -U picamera
```

If you ever need to remove your installation:

```
$ sudo pip uninstall picamera
```

2.1.5 Virtualenv installation

If you wish to install picamera within a virtualenv (useful if you're working on several Python projects with potentially conflicting dependencies, or you just like keeping things separate and easily removable):

```
$ sudo apt-get install python-pip python-virtualenv
$ virtualenv sandbox
$ source sandbox/bin/activate
(sandbox) $ pip install picamera
```

If you wish to use the classes in the `picamera.array` module then specify the “array” option which will pull in numpy as a dependency (be warned that building numpy takes a *long* time on a Pi):

```
(sandbox) $ pip install "picamera[array]"
```

Bear in mind that each time you want to use picamera you will need to activate the virtualenv before running Python:

```
$ source sandbox/bin/activate
(sandbox) $ python
>>> import picamera
```

To upgrade your installation, make sure the virtualenv is activated and just use pip:

```
$ source sandbox/bin/activate
(sandbox) $ pip install -U picamera
```

To remove your installation simply blow away the virtualenv:

```
$ rm -fr ~/sandbox/
```

2.1.6 Development installation

If you wish to develop picamera itself, it is easiest to obtain the source by cloning the GitHub repository and then use the “develop” target of the Makefile which will install the package as a link to the cloned repository allowing in-place development (it also builds a tags file for use with vim/emacs with exuberant's ctags utility). The following example demonstrates this method within a virtual Python environment:

```
$ sudo apt-get install build-essential git git-core exuberant-ctags \
python-virtualenv
$ virtualenv sandbox
$ source sandbox/bin/activate
(sandbox) $ git clone https://github.com/waveform80/picamera.git
(sandbox) $ cd picamera
(sandbox) $ make develop
```

To pull the latest changes from git into your clone and update your installation:

```
$ source sandbox/bin/activate
(sandbox) $ cd picamera
(sandbox) $ git pull
(sandbox) $ make develop
```

To remove your installation blow away the sandbox and the clone:

```
$ rm -fr ~/sandbox/ ~/picamera/
```

For anybody wishing to hack on the project please understand that although it is technically written in pure Python, heavy use of `ctypes` is involved so the code really doesn't look much like Python - more a sort of horrid mish-mash of C and Python.

The project consists primarily of a class (`PiCamera`) which is a re-implementation of high-level bits of the `raspistill` and `raspivid` commands using the `ctypes` based `libmmal` header conversion, plus a set of

encoder classes which re-implement the encoder callback configuration in the aforementioned binaries. Various classes for specialized applications also exist (`PiCameraCircularIO`, `PiBayerArray`, etc.)

Even if you don't feel up to hacking on the code, I'd love to hear suggestions from people of what you'd like the API to look like (even if the code itself isn't particularly pythonic, the interface should be)!

2.1.7 Test suite

If you wish to run the picamera test suite, follow the instructions in *Development installation* above and then install the following additional dependencies (note: `avconv` is installed system-wide):

```
(sandbox) $ sudo apt-get install libav-tools
(sandbox) $ pip install Pillow pytest mock numpy
```

Finally, to run the test suite, execute the following command:

```
(sandbox) $ make test
```

Warning: The test suite takes a *very* long time to execute (at least 4 hours on an overclocked Pi). Depending on configuration, it can also lockup the camera requiring a reboot to reset, so ensure you are familiar with SSH or using alternate TTYs to access a command line in the event you need to reboot.

2.2 Python 3.2+ Installation

There are several ways to install picamera under Python 3.2 (or above), each with their own advantages and disadvantages. Have a read of the sections below and select an installation method which conforms to your needs.

2.2.1 Firmware upgrades

The behaviour of the Pi's camera module is dictated by the Pi's firmware. Over time, considerable work has gone into fixing bugs and extending the functionality of the Pi's camera module through new firmware releases. Whilst the picamera library attempts to maintain backward compatibility with older Pi firmwares, it is only tested against the latest firmware at the time of release, and not all functionality may be available if you are running an older firmware. As an example, the `annotate_text` attribute relies on a recent firmware; older firmwares lacked the functionality.

To keep your firmware up to date use the following command, and reboot once successfully complete:

```
$ sudo rpi-update
```

You can determine the revision of your current firmware with the following command:

```
$ uname -a
```

The firmware revision is the number after the #:

```
Linux kermit 3.12.26+ #707 PREEMPT Sat Aug 30 17:39:19 BST 2014 armv6l GNU/Linux
      /
     /
firmware revision --+
```

2.2.2 Raspbian installation

If you are using the Raspbian distro, it is best to install picamera using the system's package manager: `apt`. This will ensure that picamera is easy to keep up to date, and easy to remove should you wish to do so. It will also make picamera available for all users on the system. To install picamera using `apt` simply:

```
$ sudo apt-get update
$ sudo apt-get install python3-picamera
```

To upgrade your installation when new releases are made you can simply use apt's normal upgrade procedure:

```
$ sudo apt-get update
$ sudo apt-get upgrade
```

If you ever need to remove your installation:

```
$ sudo apt-get remove python3-picamera
```

Note: If you are using a recent installation of Raspbian, you may find that the python3-picamera package is already installed (it is included by default in recent versions of [NOOBS](#)).

Note: The release of the picamera package on Raspbian lags behind the release of picamera packages on PyPI (see [System installation](#) below) by a few days. If you find yourself unable to upgrade immediately following a release announcement, you can either switch to a PyPI installation or wait a few days for the Raspbian packages to arrive in the repository.

2.2.3 User installation

This is the simplest (non-apt) form of installation (though it's also complex to uninstall should you wish to do so later), but bear in mind that it will only work for the user you install under. For example, if you install as the `pi` user, you will only be able to use `picamera` as the `pi` user. If you run python as root (e.g. with `sudo python3`) it will not find the module. See [System installation](#) below if you require a root installation.

To install as your current user:

```
$ sudo apt-get install python3-pip
$ pip-3.2 install --user picamera
```

If you wish to use the classes in the `picamera.array` module then specify the "array" option which will pull in numpy as a dependency (be warned that building numpy takes a *long* time on a Pi):

```
$ pip-3.2 install --user "picamera[array]"
```

Note that `pip-3.2` is **not** run with `sudo`; this is deliberate. To upgrade your installation when new releases are made:

```
$ pip-3.2 install --user -U picamera
```

If you ever need to remove your installation:

```
$ pip-3.2 uninstall picamera
```

2.2.4 System installation

A system installation will make `picamera` accessible to all users (in contrast to the user installation). It is as simple to perform as the user installation and equally easy to keep updated. To perform the installation:

```
$ sudo apt-get install python3-pip
$ sudo pip-3.2 install picamera
```

If you wish to use the classes in the `picamera.array` module then specify the "array" option which will pull in numpy as a dependency (be warned that building numpy takes a *long* time on a Pi):

```
$ sudo pip-3.2 install "picamera[array]"
```

To upgrade your installation when new releases are made:

```
$ sudo pip-3.2 install -U picamera
```

If you ever need to remove your installation:

```
$ sudo pip-3.2 uninstall picamera
```

2.2.5 Virtualenv installation

If you wish to install `picamera` within a virtualenv (useful if you're working on several Python projects with potentially conflicting dependencies, or you just like keeping things separate and easily removable):

```
$ sudo apt-get install python3-pip python-virtualenv
$ virtualenv -p python3 sandbox
$ source sandbox/bin/activate
(sandbox) $ pip-3.2 install picamera
```

If you wish to use the classes in the `picamera.array` module then specify the “array” option which will pull in `numpy` as a dependency (be warned that building `numpy` takes a *long* time on a Pi):

```
(sandbox) $ pip-3.2 install "picamera[array]"
```

Bear in mind that each time you want to use `picamera` you will need to activate the virtualenv before running Python:

```
$ source sandbox/bin/activate
(sandbox) $ python
>>> import picamera
```

To upgrade your installation, make sure the virtualenv is activated and just use `easy_install`:

```
$ source sandbox/bin/activate
(sandbox) $ pip-3.2 install -U picamera
```

To remove your installation simply blow away the virtualenv:

```
$ rm -fr ~/sandbox/
```

2.2.6 Development installation

If you wish to develop `picamera` itself, it is easiest to obtain the source by cloning the GitHub repository and then use the “develop” target of the Makefile which will install the package as a link to the cloned repository allowing in-place development (it also builds a tags file for use with `vim/emacs` with `exuberant's` `ctags` utility). The following example demonstrates this method within a virtual Python environment:

```
$ sudo apt-get install build-essential git git-core exuberant-ctags \
python-virtualenv
$ virtualenv -p python3 sandbox
$ source sandbox/bin/activate
(sandbox) $ git clone https://github.com/waveform80/picamera.git
(sandbox) $ cd picamera
(sandbox) $ make develop
```

To pull the latest changes from git into your clone and update your installation:

```
$ source sandbox/bin/activate
(sandbox) $ cd picamera
(sandbox) $ git pull
(sandbox) $ make develop
```

To remove your installation blow away the sandbox and the clone:

```
$ rm -fr ~/sandbox/ ~/picamera/
```

For anybody wishing to hack on the project please understand that although it is technically written in pure Python, heavy use of `ctypes` is involved so the code really doesn't look much like Python - more a sort of horrid mish-mash of C and Python.

The project consists primarily of a class (`PiCamera`) which is a re-implementation of high-level bits of the `rastipstill` and `rastipvid` commands using the `ctypes` based `libmmal` header conversion, plus a set of *encoder classes* which re-implement the encoder callback configuration in the aforementioned binaries. Various classes for specialized applications also exist (`PiCameraCircularIO`, `PiBayerArray`, etc.)

Even if you don't feel up to hacking on the code, I'd love to hear suggestions from people of what you'd like the API to look like (even if the code itself isn't particularly pythonic, the interface should be)!

2.2.7 Test suite

If you wish to run the picamera test suite, follow the instructions in *Development installation* above and then install the following additional dependencies (note: `avconv` is installed system-wide):

```
(sandbox) $ sudo apt-get install libav-tools
(sandbox) $ pip install Pillow pytest mock numpy
```

Finally, to run the test suite, execute the following command:

```
(sandbox) $ make test
```

Warning: The test suite takes a *very* long time to execute (at least 4 hours on an overclocked Pi). Depending on configuration, it can also lockup the camera requiring a reboot to reset, so ensure you are familiar with SSH or using alternate TTYs to access a command line in the event you need to reboot.

2.3 Quick Start

Start a preview for 10 seconds with the default settings:

```
import time
import picamera

camera = picamera.PiCamera()
try:
    camera.start_preview()
    time.sleep(10)
    camera.stop_preview()
finally:
    camera.close()
```

Note that you should always ensure you call `close()` on the `PiCamera` object to clean up resources. The following example demonstrates that Python's `with` statement can be used to achieve this implicitly; when the `with` block ends, `close()` will be called implicitly:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.start_preview()
    time.sleep(10)
    camera.stop_preview()
```

The following example shows that certain properties can be adjusted “live” while a preview is running. In this case, the brightness is increased steadily during display:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.start_preview()
    try:
        for i in range(100):
            camera.brightness = i
            time.sleep(0.2)
    finally:
        camera.stop_preview()
```

The next example demonstrates setting the camera resolution (this can only be done when the camera is not recording) to 640x480, then starting a preview and a recording to a disk file:

```
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.start_preview()
    camera.start_recording('foo.h264')
    camera.wait_recording(60)
    camera.stop_recording()
    camera.stop_preview()
```

The camera's default resolution is the display's resolution. If the display has been disabled (e.g. with `tvservice -o`), then the default resolution is 1280x720.

Note: Note that `wait_recording()` is used above instead of `time.sleep()`. This method checks for errors (e.g. out of disk space) while the recording is running and raises an exception if one occurs. If `time.sleep()` was used instead the exception would be raised by `stop_recording()` but only after the full waiting time had run.

This example demonstrates starting a preview, setting some parameters and then capturing an image while the preview is running:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.start_preview()
    camera.exposure_compensation = 2
    camera.exposure_mode = 'spotlight'
    camera.meter_mode = 'matrix'
    camera.image_effect = 'gpen'
    # Give the camera some time to adjust to conditions
    time.sleep(2)
    camera.capture('foo.jpg')
    camera.stop_preview()
```

The following example customizes the Exif tags to embed in the image before calling `capture()`:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (2592, 1944)
    camera.start_preview()
    time.sleep(2)
    camera.exif_tags['IFD0.Artist'] = 'Me!'
    camera.exif_tags['IFD0.Copyright'] = 'Copyright (c) 2013 Me!'
```

```
camera.capture('foo.jpg')
camera.stop_preview()
```

See the documentation for `exif_tags` for a complete list of the supported tags.

The next example demonstrates capturing a series of images as a numbered series with a one minute delay between each capture using the `capture_continuous()` method:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.start_preview()
    time.sleep(1)
    for i, filename in enumerate(camera.capture_continuous('image{counter:02d}.jpg')):
        print('Captured image %s' % filename)
        if i == 100:
            break
        time.sleep(60)
    camera.stop_preview()
```

This example demonstrates capturing low resolution JPEGs extremely rapidly using the video-port capability of the `capture_sequence()` method. The framerate of the captures is displayed afterward:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.start_preview()
    start = time.time()
    camera.capture_sequence((
        'image%03d.jpg' % i
        for i in range(120)
    ), use_video_port=True)
    print('Captured 120 images at %.2ffps' % (120 / (time.time() - start)))
    camera.stop_preview()
```

This example demonstrates capturing an unencoded image in RGB format and producing a `numpy` array from the image:

```
import time
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiRGBArray(camera) as stream:
        camera.resolution = (1024, 768)
        camera.start_preview()
        time.sleep(2)
        camera.capture(stream, 'rgb')
        print(stream.array.shape)
```

2.4 Basic Recipes

The following recipes should be reasonably accessible to Python programmers of all skill levels. Please feel free to suggest enhancements or additional recipes.

2.4.1 Capturing to a file

Capturing an image to a file is as simple as specifying the name of the file as the output of whatever `capture()` method you require:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (1024, 768)
    camera.start_preview()
    # Camera warm-up time
    time.sleep(2)
    camera.capture('foo.jpg')
```

Note that files opened by picamera (as in the case above) will be flushed and closed so that when the capture method returns, the data should be accessible to other processes.

2.4.2 Capturing to a stream

Capturing an image to a file-like object (a `socket()`, a `io.BytesIO` stream, an existing open file object, etc.) is as simple as specifying that object as the output of whatever `capture()` method you're using:

```
import io
import time
import picamera

# Create an in-memory stream
my_stream = io.BytesIO()
with picamera.PiCamera() as camera:
    camera.start_preview()
    # Camera warm-up time
    time.sleep(2)
    camera.capture(my_stream, 'jpeg')
```

Note that the format is explicitly specified in the case above. The `BytesIO` object has no filename, so the camera can't automatically figure out what format to use.

One thing to bear in mind is that (unlike specifying a filename), the stream is *not* automatically closed after capture; picamera assumes that since it didn't open the stream it can't presume to close it either. However, if the object has a `flush` method, this will be called prior to capture returning. This should ensure that once capture returns the data is accessible to other processes although the object still needs to be closed:

```
import time
import picamera

# Explicitly open a new file called my_image.jpg
my_file = open('my_image.jpg', 'wb')
with picamera.PiCamera() as camera:
    camera.start_preview()
    time.sleep(2)
    camera.capture(my_file)
# At this point my_file.flush() has been called, but the file has
# not yet been closed
my_file.close()
```

Note that in the case above, we didn't have to specify the format as the camera interrogated the `my_file` object for its filename (specifically, it looks for a `name` attribute on the provided object). As well as using stream classes built into Python (like `BytesIO`) you can also construct your own *custom outputs*.

2.4.3 Capturing to a PIL Image

This is a variation on *Capturing to a stream*. First we'll capture an image to a `BytesIO` stream (Python's in-memory stream class), then we'll rewind the position of the stream to the start, and read the stream into a `PIL Image` object:

```
import io
import time
import picamera
from PIL import Image

# Create the in-memory stream
stream = io.BytesIO()
with picamera.PiCamera() as camera:
    camera.start_preview()
    time.sleep(2)
    camera.capture(stream, format='jpeg')
# "Rewind" the stream to the beginning so we can read its content
stream.seek(0)
image = Image.open(stream)
```

2.4.4 Capturing to an OpenCV object

This is another variation on *Capturing to a stream*. First we'll capture an image to a `BytesIO` stream (Python's in-memory stream class), then convert the stream to a numpy array and read the array with `OpenCV`:

```
import io
import time
import picamera
import cv2
import numpy as np

# Create the in-memory stream
stream = io.BytesIO()
with picamera.PiCamera() as camera:
    camera.start_preview()
    time.sleep(2)
    camera.capture(stream, format='jpeg')
# Construct a numpy array from the stream
data = np.fromstring(stream.getvalue(), dtype=np.uint8)
# "Decode" the image from the array, preserving colour
image = cv2.imdecode(data, 1)
# OpenCV returns an array with data in BGR order. If you want RGB instead
# use the following...
image = image[:, :, ::-1]
```

If you want to avoid the JPEG encoding and decoding (which is lossy) and potentially speed up the process, you can now use the classes in the `picamera.array` module. As OpenCV images are simply numpy arrays arranged in BGR order, one can use the `PiRGBArray` class and simply capture with the `'bgr'` format (given that RGB and BGR data is the same size and configuration, just with reversed color planes):

```
import time
import picamera
import picamera.array
import cv2

with picamera.PiCamera() as camera:
    camera.start_preview()
    time.sleep(2)
    with picamera.array.PiRGBArray(camera) as stream:
        camera.capture(stream, format='bgr')
```

```
# At this point the image is available as stream.array
image = stream.array
```

2.4.5 Capturing resized images

Sometimes, particularly in scripts which will perform some sort of analysis or processing on images, you may wish to capture smaller images than the current resolution of the camera. Although such resizing can be performed using libraries like PIL or OpenCV, it is considerably more efficient to have the Pi's GPU perform the resizing when capturing the image. This can be done with the *resize* parameter of the `capture()` methods:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (1024, 768)
    camera.start_preview()
    # Camera warm-up time
    time.sleep(2)
    camera.capture('foo.jpg', resize=(320, 240))
```

The *resize* parameter can also be specified when recording video with the `start_recording()` method.

2.4.6 Capturing consistent images

You may wish to capture a sequence of images all of which look the same in terms of brightness, color, and contrast (this can be useful in timelapse photography, for example). Various attributes need to be used in order to ensure consistency across multiple shots. Specifically, you need to ensure that the camera's exposure time, white balance, and gains are all fixed:

- To fix exposure time, set the `shutter_speed` attribute to a reasonable value, then set `exposure_mode` to `'off'`.
- To fix gains, set the `iso` attribute to an appropriate value (higher values imply higher gains).
- To fix white balance, set the `awb_mode` to `'off'`, then set `awb_gains` to a (red, blue) tuple of gains.

It can be difficult to know what appropriate values might be for these attributes. For `iso`, a simple rule of thumb is that 100 and 200 are reasonable values for daytime, while 400 and 800 are better for low light. To determine a reasonable value for `shutter_speed` you can query the `exposure_speed` attribute when `exposure_mode` is set to something other than `'off'`. This will tell you the camera's exposure speed as determined by the auto-exposure algorithm. Finally, to determine reasonable values for `awb_gains` simply query the property while `awb_mode` is set to something other than `'off'`. Again, this will tell you the camera's white balance gains as determined by the auto-white-balance algorithm.

The following script provides a brief example of configuring these settings:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.framerate = 30
    # Give the camera's auto-exposure and auto-white-balance algorithms
    # some time to measure the scene and determine appropriate values
    camera.iso = 200
    time.sleep(2)
    # Now fix the values
    camera.shutter_speed = camera.exposure_speed
    camera.exposure_mode = 'off'
    g = camera.awb_gains
    camera.awb_mode = 'off'
```

```
camera.awb_gains = g
# Finally, take several photos with the fixed settings
camera.capture_sequence(['image%02d.jpg' % i for i in range(10)])
```

2.4.7 Capturing timelapse sequences

The simplest way to capture long time-lapse sequences is with the `capture_continuous()` method. With this method, the camera captures images continually until you tell it to stop. Images are automatically given unique names and you can easily control the delay between captures. The following example shows how to capture images with a 5 minute delay between each shot:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.start_preview()
    time.sleep(2)
    for filename in camera.capture_continuous('img{counter:03d}.jpg'):
        print('Captured %s' % filename)
        time.sleep(300) # wait 5 minutes
```

However, you may wish to capture images at a particular time, say at the start of every hour. This simply requires a refinement of the delay in the loop (the `datetime` module is slightly easier to use for calculating dates and times; this example also demonstrates the `timestamp` template in the captured filenames):

```
import time
import picamera
from datetime import datetime, timedelta

def wait():
    # Calculate the delay to the start of the next hour
    next_hour = (datetime.now() + timedelta(hour=1)).replace(
        minute=0, second=0, microsecond=0)
    delay = (next_hour - datetime.now()).seconds
    time.sleep(delay)

with picamera.PiCamera() as camera:
    camera.start_preview()
    wait()
    for filename in camera.capture_continuous('img{timestamp:%Y-%m-%d-%H-%M}.jpg'):
        print('Captured %s' % filename)
        wait()
```

2.4.8 Capturing in low light

Using similar tricks to those in *Capturing consistent images*, the Pi's camera can capture images in low light conditions. The primary objective is to set a high gain, and a long exposure time to allow the camera to gather as much light as possible. However, the `shutter_speed` attribute is constrained by the camera's `framerate` so the first thing we need to do is set a very slow framerate. The following script captures an image with a 6 second exposure time (the maximum the Pi's camera module is currently capable of):

```
import picamera
from time import sleep
from fractions import Fraction

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    # Set a framerate of 1/6fps, then set shutter
    # speed to 6s and ISO to 800
```


```
camera.framerate = Fraction(1, 6)
camera.shutter_speed = 6000000
camera.exposure_mode = 'off'
camera.iso = 800
# Give the camera a good long time to measure AWB
# (you may wish to use fixed AWB instead)
sleep(10)
# Finally, capture an image with a 6s exposure. Due
# to mode switching on the still port, this will take
# longer than 6 seconds
camera.capture('dark.jpg')
```

In anything other than dark conditions, the image produced by this script will most likely be completely white or at least heavily over-exposed.

Note: The Pi's camera module uses a [rolling shutter](#). This means that moving subjects may appear distorted if they move relative to the camera. This effect will be exaggerated by using longer exposure times.

2.4.9 Capturing to a network stream

This is a variation of *Capturing timelapse sequences*. Here we have two scripts: a server (presumably on a fast machine) which listens for a connection from the Raspberry Pi, and a client which runs on the Raspberry Pi and sends a continual stream of images to the server. We'll use a very simple protocol for communication: first the length of the image will be sent as a 32-bit integer (in [Little Endian](#) format), then this will be followed by the bytes of image data. If the length is 0, this indicates that the connection should be closed as no more images will be forthcoming. This protocol is illustrated below:



Firstly the server script (which relies on PIL for reading JPEGs, but you could replace this with any other suitable graphics library, e.g. OpenCV or GraphicsMagick):

```
import io
import socket
import struct
from PIL import Image

# Start a socket listening for connections on 0.0.0.0:8000 (0.0.0.0 means
# all interfaces)
server_socket = socket.socket()
server_socket.bind(('0.0.0.0', 8000))
server_socket.listen(0)

# Accept a single connection and make a file-like object out of it
connection = server_socket.accept()[0].makefile('rb')
try:
    while True:
        # Read the length of the image as a 32-bit unsigned int. If the
        # length is zero, quit the loop
        image_len = struct.unpack('<L', connection.read(struct.calcsize('<L')))[0]
        if not image_len:
            break
        # Construct a stream to hold the image data and read the image
        # data from the connection
        image_stream = io.BytesIO()
```

```

        image_stream.write(connection.read(image_len))
        # Rewind the stream, open it as an image with PIL and do some
        # processing on it
        image_stream.seek(0)
        image = Image.open(image_stream)
        print('Image is %dx%d' % image.size)
        image.verify()
        print('Image is verified')
    finally:
        connection.close()
        server_socket.close()

```

Now for the client side of things, on the Raspberry Pi:

```

import io
import socket
import struct
import time
import picamera

# Connect a client socket to my_server:8000 (change my_server to the
# hostname of your server)
client_socket = socket.socket()
client_socket.connect(('my_server', 8000))

# Make a file-like object out of the connection
connection = client_socket.makefile('wb')
try:
    with picamera.PiCamera() as camera:
        camera.resolution = (640, 480)
        # Start a preview and let the camera warm up for 2 seconds
        camera.start_preview()
        time.sleep(2)

        # Note the start time and construct a stream to hold image data
        # temporarily (we could write it directly to connection but in this
        # case we want to find out the size of each capture first to keep
        # our protocol simple)
        start = time.time()
        stream = io.BytesIO()
        for foo in camera.capture_continuous(stream, 'jpeg'):
            # Write the length of the capture to the stream and flush to
            # ensure it actually gets sent
            connection.write(struct.pack('<L', stream.tell()))
            connection.flush()
            # Rewind the stream and send the image data over the wire
            stream.seek(0)
            connection.write(stream.read())
            # If we've been capturing for more than 30 seconds, quit
            if time.time() - start > 30:
                break
            # Reset the stream for the next capture
            stream.seek(0)
            stream.truncate()

        # Write a length of zero to the stream to signal we're done
        connection.write(struct.pack('<L', 0))
finally:
    connection.close()
    client_socket.close()

```

The server script should be run first to ensure there's a listening socket ready to accept a connection from the client script.

2.4.10 Recording video to a file

Recording a video to a file is simple:

```
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.start_recording('my_video.h264')
    camera.wait_recording(60)
    camera.stop_recording()
```

Note that we use `wait_recording()` in the example above instead of `time.sleep()` which we've been using in the image capture recipes above. The `wait_recording()` method is similar in that it will pause for the number of seconds specified, but unlike `time.sleep()` it will continually check for recording errors (e.g. an out of disk space condition) while it is waiting. If we had used `time.sleep()` instead, such errors would only be raised by the `stop_recording()` call (which could be long after the error actually occurred).

2.4.11 Recording video to a stream

This is very similar to *Recording video to a file*:

```
import io
import picamera

stream = io.BytesIO()
with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.start_recording(stream, format='h264', quality=23)
    camera.wait_recording(15)
    camera.stop_recording()
```

Here, we've set the *quality* parameter to indicate to the encoder the level of image quality that we'd like it to try and maintain. The camera's H.264 encoder is primarily constrained by two parameters:

- *bitrate* limits the encoder's output to a certain number of bits per second. The default is 17000000 (17Mbps), and the maximum value is 25000000 (25Mbps). Higher values give the encoder more "freedom" to encode at higher qualities. You will likely find that the default doesn't constrain the encoder at all except at higher recording resolutions.
- *quality* tells the encoder what level of image quality to maintain. Values can be between 1 (highest quality) and 40 (lowest quality), with typical values providing a reasonable trade-off between bandwidth and quality being between 20 and 25.

As well as using stream classes built into Python (like `BytesIO`) you can also construct your own *custom outputs*. This is particularly useful for video recording, as discussed in the linked recipe.

2.4.12 Recording over multiple files

If you wish split your recording over multiple files, you can use the `split_recording()` method to accomplish this:

```
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.start_recording('1.h264')
    camera.wait_recording(5)
    for i in range(2, 11):
        camera.split_recording('%d.h264' % i)
```

```

    camera.wait_recording(5)
    camera.stop_recording()

```

This should produce 10 video files named `1.h264`, `2.h264`, etc. each of which is approximately 5 seconds long (approximately because the `split_recording()` method will only split files at a key-frame).

The `record_sequence()` method can also be used to achieve this with slightly cleaner code:

```

import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    for filename in camera.record_sequence(
        '%d.h264' % i for i in range(1, 11)):
        camera.wait_recording(5)

```

Changed in version 1.3: The `record_sequence()` method was introduced in version 1.3

2.4.13 Recording to a circular stream

This is similar to *Recording video to a stream* but uses a special kind of in-memory stream provided by the picamera library. The `PiCameraCircularIO` class implements a *ring buffer* based stream, specifically for video recording. This enables you to keep an in-memory stream containing the last n seconds of video recorded (where n is determined by the bitrate of the video recording and the size of the ring buffer underlying the stream).

A typical use-case for this sort of storage is security applications where one wishes to detect motion and only record to disk the video where motion was detected. This example keeps 20 seconds of video in memory until the `write_now` function returns `True` (in this implementation, this is random but one could, for example, replace this with some sort of motion detection algorithm). Once `write_now` returns `True`, the script waits 10 more seconds (so that the buffer contains 10 seconds of video from before the event, and 10 seconds after) and writes the resulting video to disk before going back to waiting:

```

import io
import random
import picamera

def write_now():
    # Randomly return True (like a fake motion detection routine)
    return random.randint(0, 10) == 0

def write_video(stream):
    print('Writing video!')
    with stream.lock:
        # Find the first header frame in the video
        for frame in stream.frames:
            if frame.frame_type == picamera.PiVideoFrameType.sps_header:
                stream.seek(frame.position)
                break
        # Write the rest of the stream to disk
        with io.open('motion.h264', 'wb') as output:
            output.write(stream.read())

with picamera.PiCamera() as camera:
    stream = picamera.PiCameraCircularIO(camera, seconds=20)
    camera.start_recording(stream, format='h264')
    try:
        while True:
            camera.wait_recording(1)
            if write_now():
                # Keep recording for 10 seconds and only then write the
                # stream to disk
                camera.wait_recording(10)

```

```
        write_video(stream)
    finally:
        camera.stop_recording()
```

In the above script we use the threading lock in the `lock` attribute to prevent the camera's background writing thread from changing the stream while our own thread reads from it (as the stream is a circular buffer, a write can remove information that is about to be read). If we had stopped recording to the stream while writing we could eliminate the `with stream.lock` line in the `write_video` function.

Note: Note that *at least* 20 seconds of video are in the stream. This is an estimate only; if the H.264 encoder requires less than the specified bitrate (17Mbps by default) for recording the video, then more than 20 seconds of video will be available in the stream.

New in version 1.0.

2.4.14 Recording to a network stream

This is similar to *Recording video to a stream* but instead of an in-memory stream like `BytesIO`, we will use a file-like object created from a `socket()`. Unlike the example in *Capturing to a network stream* we don't need to complicate our network protocol by writing things like the length of images. This time we're sending a continual stream of video frames (which necessarily incorporates such information, albeit in a much more efficient form), so we can simply dump the recording straight to the network socket.

Firstly, the server side script which will simply read the video stream and pipe it to a media player for display:

```
import socket
import subprocess

# Start a socket listening for connections on 0.0.0.0:8000 (0.0.0.0 means
# all interfaces)
server_socket = socket.socket()
server_socket.bind(('0.0.0.0', 8000))
server_socket.listen(0)

# Accept a single connection and make a file-like object out of it
connection = server_socket.accept()[0].makefile('rb')
try:
    # Run a viewer with an appropriate command line. Uncomment the mplayer
    # version if you would prefer to use mplayer instead of VLC
    cmdline = ['vlc', '--demux', 'h264', '-']
    #cmdline = ['mplayer', '-fps', '25', '-cache', '1024', '-']
    player = subprocess.Popen(cmdline, stdin=subprocess.PIPE)
    while True:
        # Repeatedly read 1k of data from the connection and write it to
        # the media player's stdin
        data = connection.read(1024)
        if not data:
            break
        player.stdin.write(data)
finally:
    connection.close()
    server_socket.close()
    player.terminate()
```

Note: If you run this script on Windows you will probably need to provide a complete path to the VLC or mplayer executable. If you run this script on Mac OS X, and are using Python installed from MacPorts, please ensure you have also installed VLC or mplayer from MacPorts.

You will probably notice several seconds of latency with this setup. This is normal and is because media players buffer several seconds to guard against unreliable network streams. Some media players (notably mplayer in this

case) permit the user to skip to the end of the buffer (press the right cursor key in mplayer), reducing the latency by increasing the risk that delayed / dropped network packets will interrupt the playback.

Now for the client side script which simply starts a recording over a file-like object created from the network socket:

```
import socket
import time
import picamera

# Connect a client socket to my_server:8000 (change my_server to the
# hostname of your server)
client_socket = socket.socket()
client_socket.connect(('my_server', 8000))

# Make a file-like object out of the connection
connection = client_socket.makefile('wb')
try:
    with picamera.PiCamera() as camera:
        camera.resolution = (640, 480)
        camera.framerate = 24
        # Start a preview and let the camera warm up for 2 seconds
        camera.start_preview()
        time.sleep(2)
        # Start recording, sending the output to the connection for 60
        # seconds, then stop
        camera.start_recording(connection, format='h264')
        camera.wait_recording(60)
        camera.stop_recording()
finally:
    connection.close()
    client_socket.close()
```

It should also be noted that the effect of the above is much more easily achieved (at least on Linux) with a combination of netcat and the raspivid executable. For example:

```
server-side: nc -l 8000 | vlc --demux h264 -
client-side: raspivid -w 640 -h 480 -t 60000 -o - | nc my_server 8000
```

However, this recipe does serve as a starting point for video streaming applications. It's also possible to reverse the direction of this recipe relatively easily. In this scenario, the Pi acts as the server, waiting for a connection from the client. When it accepts a connection, it starts streaming video over it for 60 seconds. Another variation (just for the purposes of demonstration) is that we initialize the camera straight away instead of waiting for a connection to allow the streaming to start faster on connection:

```
import socket
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.framerate = 24

    server_socket = socket.socket()
    server_socket.bind(('0.0.0.0', 8000))
    server_socket.listen(0)

    # Accept a single connection and make a file-like object out of it
    connection = server_socket.accept()[0].makefile('wb')
    try:
        camera.start_recording(connection, format='h264')
        camera.wait_recording(60)
        camera.stop_recording()
    finally:
```

```
connection.close()
server_socket.close()
```

One advantage of this setup is that no script is needed on the client side - we can simply use VLC with a network URL:

```
vlc tcp/h264://my_pi_address:8000/
```

Note: VLC (or mplayer) will *not* work for playback on a Pi. Neither is (currently) capable of using the GPU for decoding, and thus they attempt to perform video decoding on the Pi's CPU (which is not powerful enough for the task). You will need to run these applications on a faster machine (though "faster" is a relative term here: even an Atom powered netbook should be quick enough for the task at non-HD resolutions).

2.4.15 Overlaying images on the preview

The camera preview system can operate multiple layered renderers simultaneously. While the picamera library only permits a single renderer to be connected to the camera's preview port, it does permit additional renderers to be created which display a static image. These overlaid renderers can be used to create simple user interfaces.

Note: Overlay images will *not* appear in image captures or video recordings. If you need to embed additional information in the output of the camera, please refer to *Overlaying text on the output*.

One difficulty of working with overlay renderers is that they expect unencoded RGB input which is padded up to the camera's block size. The camera's block size is 32x16 so any image data provided to a renderer must have a width which is a multiple of 32, and a height which is a multiple of 16. The specific RGB format expected is interleaved unsigned bytes. If all this sounds complicated, don't worry; it's quite simple to produce in practice.

The following example demonstrates loading an arbitrary size image with PIL, padding it to the required size, and producing the unencoded RGB data for the call to `add_overlay()`:

```
import picamera
from PIL import Image
from time import sleep

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.framerate = 24
    camera.start_preview()

    # Load the arbitrarily sized image
    img = Image.open('overlay.png')
    # Create an image padded to the required size with
    # mode 'RGB'
    pad = Image.new('RGB', (
        (img.size[0] + 31 // 32) * 32,
        (img.size[1] + 15 // 16) * 16,
    ))
    # Paste the original image into the padded one
    pad.paste(img, (0, 0))

    # Add the overlay with the padded image as the source,
    # but the original image's dimensions
    o = camera.add_overlay(pad.tostring(), size=img.size)
    # By default, the overlay is in layer 0, beneath the
    # preview (which defaults to layer 2). Here we make
    # the new overlay semi-transparent, then move it above
    # the preview
    o.alpha = 128
    o.layer = 3
```

```
# Wait indefinitely until the user terminates the script
while True:
    sleep(1)
```

Alternatively, instead of using an image file as the source, you can produce an overlay directly from a numpy array. In the following example, we construct a numpy array with the same resolution as the screen, then draw a white cross through the center and overlay it on the preview as a simple cross-hair:

```
import time
import picamera
import numpy as np

# Create an array representing a 1280x720 image of
# a cross through the center of the display. The shape of
# the array must be of the form (height, width, color)
a = np.zeros((720, 1280, 3), dtype=np.uint8)
a[360, :, :] = 0xff
a[:, 640, :] = 0xff

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.framerate = 24
    camera.start_preview()
    # Add the overlay directly into layer 3 with transparency;
    # we can omit the size parameter of add_overlay as the
    # size is the same as the camera's resolution
    o = camera.add_overlay(np.getbuffer(a), layer=3, alpha=64)
    try:
        # Wait indefinitely until the user terminates the script
        while True:
            time.sleep(1)
    finally:
        camera.remove_overlay(o)
```

Given that overlaid renderers can be hidden (by moving them below the preview's `layer` which defaults to 2), made semi-transparent (with the `alpha` property), and resized so that they don't fill the screen, they can be used to construct simple user interfaces.

New in version 1.8.

2.4.16 Overlaying text on the output

The camera includes a rudimentary annotation facility which permits up to 255 characters of ASCII text to be overlaid on all output (including the preview, image captures and video recordings). To achieve this, simply assign a string to the `annotate_text` attribute:

```
import picamera
import time

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.framerate = 24
    camera.start_preview()
    camera.annotate_text = 'Hello world!'
    time.sleep(2)
    # Take a picture including the annotation
    camera.capture('foo.jpg')
```

With a little ingenuity, it's possible to display longer strings:

```
import picamera
import time
```

```
import itertools

s = "This message would be far too long to display normally..."

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.framerate = 24
    camera.start_preview()
    camera.annotate_text = ' ' * 31
    for c in itertools.cycle(s):
        camera.annotate_text = camera.annotate_text[1:31] + c
        time.sleep(0.1)
```

And of course, it can be used to display (and embed) a timestamp in recordings (this recipe also demonstrates drawing a background behind the timestamp for contrast with the `annotate_bg` attribute):

```
import picamera
import datetime as dt

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.framerate = 24
    camera.start_preview()
    camera.annotate_bg = True
    camera.annotate_text = dt.datetime.now().strftime('%Y-%m-%d %H:%M:%S')
    camera.start_recording('timestamped.h264')
    start = dt.datetime.now()
    while (dt.datetime.now() - start).seconds < 30:
        camera.annotate_text = dt.datetime.now().strftime('%Y-%m-%d %H:%M:%S')
        camera.wait_recording(0.2)
    camera.stop_recording()
```

New in version 1.7.

2.4.17 Controlling the LED

In certain circumstances, you may find the camera module's red LED a hindrance. For example, in the case of automated close-up wild-life photography, the LED may scare off animals. It can also cause unwanted reflected red glare with close-up subjects.

One trivial way to deal with this is simply to place some opaque covering on the LED (e.g. blue-tack or electricians tape). Another method is to use the `disable_camera_led` option in the [boot configuration](#).

However, provided you have the [RPi.GPIO](#) package installed, and provided your Python process is running with sufficient privileges (typically this means running as root with `sudo python`), you can also control the LED via the `led` attribute:

```
import picamera

with picamera.PiCamera() as camera:
    # Turn the camera's LED off
    camera.led = False
    # Take a picture while the LED remains off
    camera.capture('foo.jpg')
```

Warning: Be aware when you first use the LED property it will set the GPIO library to Broadcom (BCM) mode with `GPIO.setmode(GPIO.BCM)` and disable warnings with `GPIO.setwarnings(False)`. The LED cannot be controlled when the library is in BOARD mode.

2.5 Advanced Recipes

The following recipes involve advanced techniques and may not be “beginner friendly”. Please feel free to suggest enhancements or additional recipes.

2.5.1 Unencoded image capture (YUV format)

If you want images captured without loss of detail (due to JPEG’s lossy compression), you are probably better off exploring PNG as an alternate image format (PNG uses lossless compression). However, some applications (particularly scientific ones) simply require the image data in numeric form. For this, the ‘yuv’ format is provided:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (100, 100)
    camera.start_preview()
    time.sleep(2)
    camera.capture('image.data', 'yuv')
```

The specific YUV format used is YUV420 (planar). This means that the Y (luminance) values occur first in the resulting data and have full resolution (one 1-byte Y value for each pixel in the image). The Y values are followed by the U (chrominance) values, and finally the V (chrominance) values. The UV values have one quarter the resolution of the Y components (4 1-byte Y values in a square for each 1-byte U and 1-byte V value). This is illustrated in the diagram below:

Single Frame YUV420:



Position in byte stream:



It is also important to note that when outputting to unencoded formats, the camera rounds the requested resolution. The horizontal resolution is rounded up to the nearest multiple of 32 pixels, while the vertical resolution is rounded up to the nearest multiple of 16 pixels. For example, if the requested resolution is 100x100, the capture will actually contain 128x112 pixels worth of data, but pixels beyond 100x100 will be uninitialized.

Given that the YUV420 format contains 1.5 bytes worth of data for each pixel (a 1-byte Y value for each pixel, and 1-byte U and V values for every 4 pixels), and taking into account the resolution rounding, the size of a 100x100 YUV capture will be:

$$\begin{aligned}
 & 128.0 \quad 100 \text{ rounded up to nearest multiple of } 32 \\
 \times & 112.0 \quad 100 \text{ rounded up to nearest multiple of } 16 \\
 \times & 1.5 \quad \text{bytes of data per pixel in YUV4:2:0 format} \\
 = & 21504.0 \quad \text{bytes}
 \end{aligned}$$

The first 14336 bytes of the data (128*112) will be Y values, the next 3584 bytes (128*112/4) will be U values, and the final 3584 bytes will be the V values.

The following code demonstrates capturing YUV image data, loading the data into a set of `numpy` arrays, and converting the data to RGB format in an efficient manner:

```
from __future__ import division

import time
import picamera
import numpy as np

width = 100
height = 100
stream = open('image.data', 'w+b')
# Capture the image in YUV format
with picamera.PiCamera() as camera:
    camera.resolution = (width, height)
    camera.start_preview()
    time.sleep(2)
    camera.capture(stream, 'yuv')
# Rewind the stream for reading
stream.seek(0)
# Calculate the actual image size in the stream (accounting for rounding
# of the resolution)
fwidth = (width + 31) // 32 * 32
fheight = (height + 15) // 16 * 16
# Load the Y (luminance) data from the stream
Y = np.fromfile(stream, dtype=np.uint8, count=fwidth*fheight).\
    reshape((fheight, fwidth))
# Load the UV (chrominance) data from the stream, and double its size
U = np.fromfile(stream, dtype=np.uint8, count=(fwidth//2)*(fheight//2)).\
    reshape((fheight//2, fwidth//2)).\
    repeat(2, axis=0).repeat(2, axis=1)
V = np.fromfile(stream, dtype=np.uint8, count=(fwidth//2)*(fheight//2)).\
    reshape((fheight//2, fwidth//2)).\
    repeat(2, axis=0).repeat(2, axis=1)
# Stack the YUV channels together, crop the actual resolution, convert to
# floating point for later calculations, and apply the standard biases
YUV = np.dstack((Y, U, V))[:height, :width, :].astype(np.float)
YUV[:, :, 0] = YUV[:, :, 0] - 16 # Offset Y by 16
YUV[:, :, 1:] = YUV[:, :, 1:] - 128 # Offset UV by 128
# YUV conversion matrix from ITU-R BT.601 version (SDTV)
#           Y         U         V
M = np.array([[1.164, 0.000, 1.596], # R
              [1.164, -0.392, -0.813], # G
              [1.164, 2.017, 0.000]]) # B
# Take the dot product with the matrix to produce RGB output, clamp the
# results to byte range and convert to bytes
RGB = YUV.dot(M.T).clip(0, 255).astype(np.uint8)
```

Note: You may note that we are using `open()` in the code above instead of `io.open()` as in the other examples. This is because `numpy.fromfile()` method annoyingly only accepts “real” file objects.

This recipe is now encapsulated in the `PiYUVArray` class in the `picamera.array` module, which means the same can be achieved as follows:

```
import time
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiYUVArray(camera) as stream:
```

```

camera.resolution = (100, 100)
camera.start_preview()
time.sleep(2)
camera.capture(stream, 'yuv')
# Show size of YUV data
print(stream.array.shape)
# Show size of RGB converted data
print(stream.rgb_array.shape)

```

Alternatively, see *Unencoded image capture (RGB format)* for a method of having the camera output RGB data directly.

Note: Capturing so-called “raw” formats ('yuv', 'rgb', etc.) does not provide the raw bayer data from the camera’s sensor. Rather, it provides access to the image data after GPU processing, but before format encoding (JPEG, PNG, etc). Currently, the only method of accessing the raw bayer data is via the *bayer* parameter to the `capture()` method. See *Raw Bayer data captures* for more information.

Changed in version 1.0: The `raw_format` attribute is now deprecated, as is the 'raw' format specification for the `capture()` method. Simply use the 'yuv' format instead, as shown in the code above.

Changed in version 1.5: Added note about new `picamera.array` module.

2.5.2 Unencoded image capture (RGB format)

The RGB format is rather larger than the YUV format discussed in the section above, but is more useful for most analyses. To have the camera produce output in RGB format, you simply need to specify 'rgb' as the format for the `capture()` method instead:

```

import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (100, 100)
    camera.start_preview()
    time.sleep(2)
    camera.capture('image.data', 'rgb')

```

The size of RGB data can be calculated similarly to YUV captures. Firstly round the resolution appropriately (see *Unencoded image capture (YUV format)* for the specifics), then multiply the number of pixels by 3 (1 byte of red, 1 byte of green, and 1 byte of blue intensity). Hence, for a 100x100 capture, the amount of data produced is:

$$\begin{aligned}
 & 128.0 \quad 100 \text{ rounded up to nearest multiple of } 32 \\
 \times & 112.0 \quad 100 \text{ rounded up to nearest multiple of } 16 \\
 \times & 3.0 \quad \text{bytes of data per pixel in RGB format} \\
 = & 43008.0 \quad \text{bytes}
 \end{aligned}$$

The resulting RGB data is interleaved. That is to say that the red, green and blue values for a given pixel are grouped together, in that order. The first byte of the data is the red value for the pixel at (0, 0), the second byte is the green value for the same pixel, and the third byte is the blue value for that pixel. The fourth byte is the red value for the pixel at (1, 0), and so on.

Loading the resulting RGB data into a `numpy` array is simple:

```

from __future__ import division

width = 100
height = 100
stream = open('image.data', 'w+b')
# Capture the image in RGB format
with picamera.PiCamera() as camera:
    camera.resolution = (width, height)

```

```
camera.start_preview()
time.sleep(2)
camera.capture(stream, 'rgb')
# Rewind the stream for reading
stream.seek(0)
# Calculate the actual image size in the stream (accounting for rounding
# of the resolution)
fwidth = (width + 31) // 32 * 32
fheight = (height + 15) // 16 * 16
# Load the data in a three-dimensional array and crop it to the requested
# resolution
image = np.fromfile(stream, dtype=np.uint8).\
    reshape((fheight, fwidth, 3))[:height, :width, :]
# If you wish, the following code will convert the image's bytes into
# floating point values in the range 0 to 1 (a typical format for some
# sorts of analysis)
image = image.astype(np.float, copy=False)
image = image / 255.0
```

This recipe is now encapsulated in the `PiRGBArray` class in the `picamera.array` module, which means the same can be achieved as follows:

```
import time
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiRGBArray(camera) as stream:
        camera.resolution = (100, 100)
        camera.start_preview()
        time.sleep(2)
        camera.capture(stream, 'rgb')
        # Show size of RGB data
        print(stream.array.shape)
```

Note: RGB captures from the still port do not work at the full resolution of the camera (they result in an out of memory error). Either use YUV captures, or capture from the video port if you require full resolution.

Changed in version 1.0: The `raw_format` attribute is now deprecated, as is the `'raw'` format specification for the `capture()` method. Simply use the `'rgb'` format instead, as shown in the code above.

Changed in version 1.5: Added note about new `picamera.array` module.

2.5.3 Rapid capture and processing

The camera is capable of capturing a sequence of images extremely rapidly by utilizing its video-capture capabilities with a JPEG encoder (via the `use_video_port` parameter). However, there are several things to note about using this technique:

- When using video-port based capture only the video recording area is captured; in some cases this may be smaller than the normal image capture area (see discussion in *Camera Modes*).
- No Exif information is embedded in JPEG images captured through the video-port.
- Captures typically appear “grainier” with this technique. Captures from the still port use a slower, more intensive denoise algorithm.

All capture methods support the `use_video_port` option, but the methods differ in their ability to rapidly capture sequential frames. So, whilst `capture()` and `capture_continuous()` both support `use_video_port`, `capture_sequence()` is by far the fastest method (because it does not re-initialize an encoder prior to each capture). Using this method, the author has managed 30fps JPEG captures at a resolution of 1024x768.

By default, `capture_sequence()` is particularly suited to capturing a fixed number of frames rapidly, as in the following example which captures a “burst” of 5 images:

```
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (1024, 768)
    camera.framerate = 30
    camera.start_preview()
    time.sleep(2)
    camera.capture_sequence([
        'image1.jpg',
        'image2.jpg',
        'image3.jpg',
        'image4.jpg',
        'image5.jpg',
    ])

```

We can refine this slightly by using a generator expression to provide the filenames for processing instead of specifying every single filename manually:

```
import time
import picamera

frames = 60

with picamera.PiCamera() as camera:
    camera.resolution = (1024, 768)
    camera.framerate = 30
    camera.start_preview()
    # Give the camera some warm-up time
    time.sleep(2)
    start = time.time()
    camera.capture_sequence([
        'image%02d.jpg' % i
        for i in range(frames)
    ], use_video_port=True)
    finish = time.time()
print('Captured %d frames at %.2ffps' % (
    frames,
    frames / (finish - start)))

```

However, this still doesn’t let us capture an arbitrary number of frames until some condition is satisfied. To do this we need to use a generator function to provide the list of filenames (or more usefully, streams) to the `capture_sequence()` method:

```
import time
import picamera

frames = 60

def filenames():
    frame = 0
    while frame < frames:
        yield 'image%02d.jpg' % frame
        frame += 1

with picamera.PiCamera() as camera:
    camera.resolution = (1024, 768)
    camera.framerate = 30
    camera.start_preview()
    # Give the camera some warm-up time
    time.sleep(2)

```

```
start = time.time()
camera.capture_sequence(filenamees(), use_video_port=True)
finish = time.time()
print('Captured %d frames at %.2ffps' % (
    frames,
    frames / (finish - start)))
```

The major issue with capturing this rapidly is firstly that the Raspberry Pi's IO bandwidth is extremely limited and secondly that, as a format, JPEG is considerably less efficient than the H.264 video format (which is to say that, for the same number of bytes, H.264 will provide considerably better quality over the same number of frames). At higher resolutions (beyond 800x600) you are likely to find you cannot sustain 30fps captures to the Pi's SD card for very long (before exhausting the disk cache).

If you are intending to perform processing on the frames after capture, you may be better off just capturing video and decoding frames from the resulting file rather than dealing with individual JPEG captures. Alternatively, you may wish to investigate sending the data over the network (which typically has more bandwidth available than the SD card interface) and having another machine perform any required processing. However, if you can perform your processing fast enough, you may not need to involve the disk or network at all. Using a generator function, we can maintain a queue of objects to store the captures, and have parallel threads accept and process the streams as captures come in. Provided the processing runs at a faster frame rate than the captures, the encoder won't stall:

```
import io
import time
import threading
import picamera

# Create a pool of image processors
done = False
lock = threading.Lock()
pool = []

class ImageProcessor(threading.Thread):
    def __init__(self):
        super(ImageProcessor, self).__init__()
        self.stream = io.BytesIO()
        self.event = threading.Event()
        self.terminated = False
        self.start()

    def run(self):
        # This method runs in a separate thread
        global done
        while not self.terminated:
            # Wait for an image to be written to the stream
            if self.event.wait(1):
                try:
                    self.stream.seek(0)
                    # Read the image and do some processing on it
                    #Image.open(self.stream)
                    #...
                    #...
                    # Set done to True if you want the script to terminate
                    # at some point
                    #done=True
                finally:
                    # Reset the stream and event
                    self.stream.seek(0)
                    self.stream.truncate()
                    self.event.clear()
                    # Return ourselves to the pool
                with lock:
                    pool.append(self)
```

```

def streams():
    while not done:
        with lock:
            if pool:
                processor = pool.pop()
            else:
                processor = None
        if processor:
            yield processor.stream
            processor.event.set()
        else:
            # When the pool is starved, wait a while for it to refill
            time.sleep(0.1)

with picamera.PiCamera() as camera:
    pool = [ImageProcessor() for i in range(4)]
    camera.resolution = (640, 480)
    camera.framerate = 30
    camera.start_preview()
    time.sleep(2)
    camera.capture_sequence(streams(), use_video_port=True)

# Shut down the processors in an orderly fashion
while pool:
    with lock:
        processor = pool.pop()
        processor.terminated = True
        processor.join()

```

2.5.4 Rapid capture and streaming

Following on from *Rapid capture and processing*, we can combine the video-port capture technique with *Capturing to a network stream*. The server side script doesn't change (it doesn't really care what capture technique is being used - it just reads JPEGs off the wire). The changes to the client side script can be minimal at first - just add `use_video_port=True` to the `capture_continuous()` call:

```

import io
import socket
import struct
import time
import picamera

client_socket = socket.socket()
client_socket.connect(('my_server', 8000))
connection = client_socket.makefile('wb')
try:
    with picamera.PiCamera() as camera:
        camera.resolution = (640, 480)
        camera.framerate = 30
        time.sleep(2)
        start = time.time()
        stream = io.BytesIO()
        # Use the video-port for captures...
        for foo in camera.capture_continuous(stream, 'jpeg',
                                           use_video_port=True):
            connection.write(struct.pack('<L', stream.tell()))
            connection.flush()
            stream.seek(0)
            connection.write(stream.read())
            if time.time() - start > 30:
                break

```

```
        stream.seek(0)
        stream.truncate()
    connection.write(struct.pack('<L', 0))
finally:
    connection.close()
    client_socket.close()
```

Using this technique, the author can manage about 10fps of streaming at 640x480 on firmware #685. One deficiency of the script above is that it interleaves capturing images with sending them over the wire (although we deliberately don't flush on sending the image data). Potentially, it would be more efficient to permit image capture to occur simultaneously with image transmission. We can attempt to do this by utilizing the background threading techniques from the final example in *Rapid capture and processing*:

```
import io
import socket
import struct
import time
import threading
import picamera

client_socket = socket.socket()
client_socket.connect(('spider', 8000))
connection = client_socket.makefile('wb')
try:
    connection_lock = threading.Lock()
    pool_lock = threading.Lock()
    pool = []

    class ImageStreamer(threading.Thread):
        def __init__(self):
            super(ImageStreamer, self).__init__()
            self.stream = io.BytesIO()
            self.event = threading.Event()
            self.terminated = False
            self.start()

        def run(self):
            # This method runs in a background thread
            while not self.terminated:
                # Wait for the image to be written to the stream
                if self.event.wait(1):
                    try:
                        with connection_lock:
                            connection.write(struct.pack('<L', self.stream.tell()))
                            connection.flush()
                            self.stream.seek(0)
                            connection.write(self.stream.read())
                    finally:
                        self.stream.seek(0)
                        self.stream.truncate()
                        self.event.clear()
                        with pool_lock:
                            pool.append(self)

count = 0
start = time.time()
finish = time.time()

def streams():
    global count, finish
    while finish - start < 30:
        with pool_lock:
            if pool:
```

```

        streamer = pool.pop()
    else:
        streamer = None
    if streamer:
        yield streamer.stream
        streamer.event.set()
        count += 1
    else:
        # When the pool is starved, wait a while for it to refill
        time.sleep(0.1)
    finish = time.time()

with picamera.PiCamera() as camera:
    pool = [ImageStreamer() for i in range(4)]
    camera.resolution = (640, 480)
    camera.framerate = 30
    time.sleep(2)
    start = time.time()
    camera.capture_sequence(streams(), 'jpeg', use_video_port=True)

# Shut down the streamers in an orderly fashion
while pool:
    streamer = pool.pop()
    streamer.terminated = True
    streamer.join()

# Write the terminating 0-length to the connection to let the server
# know we're done
with connection_lock:
    connection.write(struct.pack('<L', 0))

finally:
    connection.close()
    client_socket.close()

print('Sent %d images in %d seconds at %.2ffps' % (
    count, finish-start, count / (finish-start)))

```

On the same firmware, the above script achieves about 15fps. It is possible the new high framerate modes may achieve more (the fact that 15fps is half of the specified 30fps framerate suggests some stall on every other frame).

2.5.5 Capturing images whilst recording

The camera is capable of capturing still images while it is recording video. However, if one attempts this using the stills capture mode, the resulting video will have dropped frames during the still image capture. This is because images captured via the still port require a mode change, causing the dropped frames (this is the flicker to a higher resolution that one sees when capturing while a preview is running).

However, if the `use_video_port` parameter is used to force a video-port based image capture (see [Rapid capture and processing](#)) then the mode change does not occur, and the resulting video should not have dropped frames, assuming the image can be produced before the next video frame is due:

```

import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (800, 600)
    camera.start_preview()
    camera.start_recording('foo.h264')
    camera.wait_recording(10)
    camera.capture('foo.jpg', use_video_port=True)
    camera.wait_recording(10)
    camera.stop_recording()

```

The above code should produce a 20 second video with no dropped frames, and a still frame from 10 seconds into the video. Higher resolutions or non-JPEG image formats may still cause dropped frames (only JPEG encoding is hardware accelerated).

2.5.6 Recording at multiple resolutions

The camera is capable of recording multiple streams at different resolutions simultaneously by use of the video splitter. This is probably most useful for performing analysis on a low-resolution stream, while simultaneously recording a high resolution stream for storage or viewing.

The following simple recipe demonstrates using the *splitter_port* parameter of the `start_recording()` method to begin two simultaneous recordings, each with a different resolution:

```
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (1024, 768)
    camera.framerate = 30
    camera.start_recording('highres.h264')
    camera.start_recording('lowres.h264', splitter_port=2, resize=(320, 240))
    camera.wait_recording(30)
    camera.stop_recording(splitter_port=2)
    camera.stop_recording()
```

There are 4 splitter ports in total that can be used (numbered 0, 1, 2, and 3). The video recording methods default to using splitter port 1, while the image capture methods default to splitter port 0 (when the *use_video_port* parameter is also True). A splitter port cannot be simultaneously used for video recording and image capture so you are advised to avoid splitter port 0 for video recordings unless you never intend to capture images whilst recording.

New in version 1.3.

2.5.7 Recording motion vector data

The Pi's camera is capable of outputting the motion vector estimates that the camera's H.264 encoder calculates while generating compressed video. These can be directed to a separate output file (or file-like object) with the *motion_output* parameter of the `start_recording()` method. Like the normal *output* parameter this accepts a string representing a filename, or a file-like object:

```
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.framerate = 30
    camera.start_recording('motion.h264', motion_output='motion.data')
    camera.wait_recording(10)
    camera.stop_recording()
```

Motion data is calculated at the *macro-block* level (an MPEG macro-block represents a 16x16 pixel region of the frame), and includes one extra column of data. Hence, if the camera's resolution is 640x480 (as in the example above) there will be 41 columns of motion data ($((640 / 16) + 1)$, in 30 rows ($480 / 16$).

Motion data values are 4-bytes long, consisting of a signed 1-byte x vector, a signed 1-byte y vector, and an unsigned 2-byte SAD (*Sum of Absolute Differences*) value for each macro-block. Hence in the example above, each frame will generate 4920 bytes of motion data ($41 * 30 * 4$). Assuming the data contains 300 frames (in practice it may contain a few more) the motion data should be 1,476,000 bytes in total.

The following code demonstrates loading the motion data into a three-dimensional numpy array. The first dimension represents the frame, with the latter two representing rows and finally columns. A structured data-type is used for the array permitting easy access to x, y, and SAD values:

```
from __future__ import division

import numpy as np

width = 640
height = 480
cols = (width + 15) // 16
cols += 1 # there's always an extra column
rows = (height + 15) // 16

motion_data = np.fromfile(
    'motion.data', dtype=[
        ('x', 'i1'),
        ('y', 'i1'),
        ('sad', 'u2'),
    ])
frames = motion_data.shape[0] // (cols * rows * motion_data.dtype.itemsize)
motion_data = motion_data.reshape((frames, rows, cols))

# Access the data for the first frame
motion_data[0]

# Access just the x-vectors from the fifth frame
motion_data[4]['x']

# Access SAD values for the tenth frame
motion_data[9]['sad']
```

You can calculate the amount of motion the vector represents simply by calculating the [magnitude of the vector](#) with Pythagoras' theorem. The SAD ([Sum of Absolute Differences](#)) value can be used to determine how well the encoder thinks the vector represents the original reference frame.

The following code extends the example above to use PIL to produce a PNG image from the magnitude of each frame's motion vectors:

```
from __future__ import division

import numpy as np
from PIL import Image

width = 640
height = 480
cols = (width + 15) // 16
rows = (height + 15) // 16

m = np.fromfile(
    'motion.data', dtype=[
        ('x', 'i1'),
        ('y', 'i1'),
        ('sad', 'u2'),
    ])
frames = m.shape[0] // (cols * rows * m.dtype.itemsize)
m = m.reshape((frames, rows, cols))

for frame in range(frames):
    data = np.sqrt(
        np.square(m[frame]['x'].astype(np.float)) +
        np.square(m[frame]['y'].astype(np.float))
    ).clip(0, 255).astype(np.uint8)
    img = Image.fromarray(data)
```

```
filename = 'frame%03d.png' % frame
print('Writing %s' % filename)
img.save(filename)
```

You may wish to investigate the `PiMotionArray` class in the `picamera.array` module which simplifies the above recipes to the following:

```
import numpy as np
import picamera
import picamera.array
from PIL import Image

with picamera.PiCamera() as camera:
    with picamera.array.PiMotionArray(camera) as stream:
        camera.resolution = (640, 480)
        camera.framerate = 30
        camera.start_recording('/dev/null', format='h264', motion_output=stream)
        camera.wait_recording(10)
        camera.stop_recording()
        for frame in range(stream.array.shape[0]):
            data = np.sqrt(
                np.square(stream.array[frame]['x'].astype(np.float)) +
                np.square(stream.array[frame]['y'].astype(np.float))
            ).clip(0, 255).astype(np.uint8)
            img = Image.fromarray(data)
            filename = 'frame%03d.png' % frame
            print('Writing %s' % filename)
            img.save(filename)
```

Finally, the following command line can be used to generate an animation from the generated PNGs with `ffmpeg` (this will take a *very* long time on the Pi so you may wish to transfer the images to a faster machine for this step):

```
avconv -r 30 -i frame%03d.png -filter:v scale=640:480 -c:v libx264 motion.mp4
```

New in version 1.5.

2.5.8 Splitting to/from a circular stream

This example builds on the one in *Recording to a circular stream* and the one in *Capturing images whilst recording* to demonstrate the beginnings of a security application. As before, a `PiCameraCircularIO` instance is used to keep the last few seconds of video recorded in memory. While the video is being recorded, video-port-based still captures are taken to provide a motion detection routine with some input (the actual motion detection algorithm is left as an exercise for the reader).

Once motion is detected, the last 10 seconds of video are written to disk, and video recording is split to another disk file to proceed until motion is no longer detected. Once motion is no longer detected, we split the recording back to the in-memory ring-buffer:

```
import io
import random
import picamera
from PIL import Image

prior_image = None

def detect_motion(camera):
    global prior_image
    stream = io.BytesIO()
    camera.capture(stream, format='jpeg', use_video_port=True)
    stream.seek(0)
    if prior_image is None:
        prior_image = Image.open(stream)
```

```
    return False
else:
    current_image = Image.open(stream)
    # Compare current_image to prior_image to detect motion. This is
    # left as an exercise for the reader!
    result = random.randint(0, 10) == 0
    # Once motion detection is done, make the prior image the current
    prior_image = current_image
    return result

def write_video(stream):
    # Write the entire content of the circular buffer to disk. No need to
    # lock the stream here as we're definitely not writing to it
    # simultaneously
    with io.open('before.h264', 'wb') as output:
        for frame in stream.frames:
            if frame.frame_type == picamera.PiVideoFrameType.sps_header:
                stream.seek(frame.position)
                break
        while True:
            buf = stream.read1()
            if not buf:
                break
            output.write(buf)
    # Wipe the circular stream once we're done
    stream.seek(0)
    stream.truncate()

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    stream = picamera.PiCameraCircularIO(camera, seconds=10)
    camera.start_recording(stream, format='h264')
    try:
        while True:
            camera.wait_recording(1)
            if detect_motion(camera):
                print('Motion detected!')
                # As soon as we detect motion, split the recording to
                # record the frames "after" motion
                camera.split_recording('after.h264')
                # Write the 10 seconds "before" motion to disk as well
                write_video(stream)
                # Wait until motion is no longer detected, then split
                # recording back to the in-memory circular buffer
                while detect_motion(camera):
                    camera.wait_recording(1)
                print('Motion stopped!')
                camera.split_recording(stream)
    finally:
        camera.stop_recording()
```

This example also demonstrates writing the circular buffer to disk in an efficient manner using the `read1()` method (as opposed to `read()`).

Note: Note that `read1()` does not guarantee to return the number of bytes requested, even if they are available in the underlying stream; it simply returns as many as are available from a single chunk up to the limit specified.

New in version 1.0.

2.5.9 Custom outputs

All methods in the `picamera` library which accept a filename also accept file-like objects. Typically, this is only used with actual file objects, or with memory streams (like `io.BytesIO`). However, building a custom output object is extremely easy and in certain cases very useful. A file-like object (as far as `picamera` is concerned) is simply an object with a `write` method which must accept a single parameter consisting of a byte-string, and which can optionally return the number of bytes written. The object can optionally implement a `flush` method (which has no parameters), which will be called at the end of output.

Custom outputs are particularly useful with video recording as the custom output's `write` method will be called (at least) once for every frame that is output, allowing you to implement code that reacts to each and every frame without going to the bother of a full *custom encoder*. However, one should bear in mind that because the `write` method is called so frequently, its implementation must be sufficiently rapid that it doesn't stall the encoder (it must perform its processing and return before the next write is due to arrive).

The following trivial example demonstrates an incredibly simple custom output which simply throws away the output while counting the number of bytes that would have been written and prints this at the end of the output:

```
from __future__ import print_function

import picamera

class MyOutput(object):
    def __init__(self):
        self.size = 0

    def write(self, s):
        self.size += len(s)

    def flush(self):
        print('%d bytes would have been written' % self.size)

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.framerate = 60
    camera.start_recording(MyOutput(), format='h264')
    camera.wait_recording(10)
    camera.stop_recording()
```

The following example shows how to use a custom output to construct a crude motion detection system. We construct a custom output object which is used as the destination for motion vector data (this is particularly simple as motion vector data always arrives as single chunks; frame data by contrast sometimes arrives in several separate chunks). The output object doesn't actually write the motion data anywhere; instead it loads it into a numpy array and analyses whether there are any significantly large vectors in the data, printing a message to the console if there are. As we are not concerned with keeping the actual video output in this example, we use `/dev/null` as the destination for the video data:

```
from __future__ import division

import picamera
import numpy as np

motion_dtype = np.dtype([
    ('x', 'i1'),
    ('y', 'i1'),
    ('sad', 'u2'),
])

class MyMotionDetector(object):
    def __init__(self, camera):
        width, height = camera.resolution
        self.cols = (width + 15) // 16
        self.cols += 1 # there's always an extra column
```

```

self.rows = (height + 15) // 16

def write(self, s):
    # Load the motion data from the string to a numpy array
    data = np.fromstring(s, dtype=motion_dtype)
    # Re-shape it and calculate the magnitude of each vector
    data = data.reshape((self.rows, self.cols))
    data = np.sqrt(
        np.square(data['x'].astype(np.float)) +
        np.square(data['y'].astype(np.float))
    ).clip(0, 255).astype(np.uint8)
    # If there're more than 10 vectors with a magnitude greater
    # than 60, then say we've detected motion
    if (data > 60).sum() > 10:
        print('Motion detected!')
    # Pretend we wrote all the bytes of s
    return len(s)

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.framerate = 30
    camera.start_recording(
        # Throw away the video data, but make sure we're using H.264
        '/dev/null', format='h264',
        # Record motion data to our custom output object
        motion_output=MyMotionDetector(camera)
    )
    camera.wait_recording(30)
    camera.stop_recording()

```

You may wish to investigate the classes in the `picamera.array` module which implement several custom outputs for analysis of data with numpy. In particular, the `PiMotionAnalysis` class can be used to remove much of the boiler plate code from the recipe above:

```

import picamera
import picamera.array
import numpy as np

class MyMotionDetector(picamera.array.PiMotionAnalysis):
    def analyse(self, a):
        a = np.sqrt(
            np.square(a['x'].astype(np.float)) +
            np.square(a['y'].astype(np.float))
        ).clip(0, 255).astype(np.uint8)
        # If there're more than 10 vectors with a magnitude greater
        # than 60, then say we've detected motion
        if (a > 60).sum() > 10:
            print('Motion detected!')

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    camera.framerate = 30
    camera.start_recording(
        '/dev/null', format='h264',
        motion_output=MyMotionDetector(camera)
    )
    camera.wait_recording(30)
    camera.stop_recording()

```

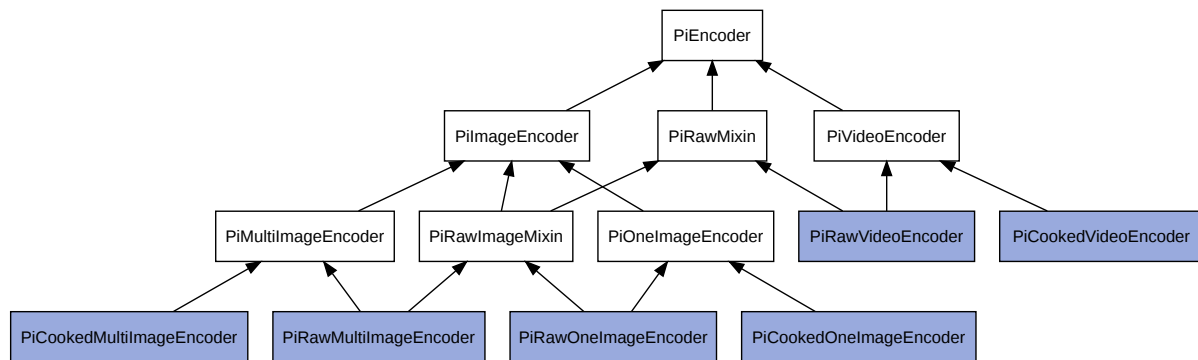
New in version 1.5.

2.5.10 Custom encoders

You can override and/or extend the encoder classes used during image or video capture. This is particularly useful with video capture as it allows you to run your own code in response to every frame, although naturally whatever code runs within the encoder’s callback has to be reasonably quick to avoid stalling the encoder pipeline.

Writing a custom encoder is quite a bit harder than writing a *custom output* and in most cases there’s little benefit. The only thing a custom encoder gives you that a custom output doesn’t is access to the buffer header flags. For many output formats (MJPEG and YUV for example), these won’t tell you anything interesting (i.e. they’ll simply indicate that the buffer contains a full frame and nothing else). Currently, the only format where the buffer header flags contain useful information is H.264. Even then, most of the information (I-frame, P-frame, motion information, etc.) would be accessible from the `frame` attribute which you could access from your custom output’s `write` method.

The encoder classes defined by picamera form the following hierarchy (shaded classes are actually instantiated by the implementation in picamera, white classes implement base functionality but aren’t technically “abstract”):



The following table details which PiCamera methods use which encoder classes, and which method they call to construct these encoders:

Method(s)	Call	Returns
<code>capture()</code>	<code>_get_image_encoder()</code>	<code>PiCookedOneImageEncoder</code>
<code>capture_continuous()</code>		<code>PiRawOneImageEncoder</code>
<code>capture_sequence()</code>		
<code>capture_sequence()</code>	<code>_get_images_encoder()</code>	<code>PiCookedMultiImageEncoder</code>
		<code>PiRawMultiImageEncoder</code>
<code>start_recording()</code>	<code>_get_video_encoder()</code>	<code>PiCookedVideoEncoder</code>
<code>record_sequence()</code>		<code>PiRawVideoEncoder</code>

It is recommended, particularly in the case of the image encoder classes, that you familiarize yourself with the specific function of these classes so that you can determine the best class to extend for your particular needs. You may find that one of the intermediate classes is a better basis for your own modifications.

In the following example recipe we will extend the `PiCookedVideoEncoder` class to store how many I-frames and P-frames are captured (the camera’s encoder doesn’t use B-frames):

```
import picamera
import picamera.mmal as mmal

# Override PiVideoEncoder to keep track of the number of each type of frame
class MyEncoder(picamera.PiCookedVideoEncoder):
    def start(self, output, motion_output=None):
        self.parent.i_frames = 0
        self.parent.p_frames = 0
        super(MyEncoder, self).start(output, motion_output)

    def _callback_write(self, buf):
        # Only count when buffer indicates it's the end of a frame, and
        # it's not an SPS/PPS header (..._CONFIG)
```

```

    if (
        (buf[0].flags & mmal.MMAL_BUFFER_HEADER_FLAG_FRAME_END) and
        not (buf[0].flags & mmal.MMAL_BUFFER_HEADER_FLAG_CONFIG)
    ):
        if buf[0].flags & mmal.MMAL_BUFFER_HEADER_FLAG_KEYFRAME:
            self.parent.i_frames += 1
        else:
            self.parent.p_frames += 1
    # Remember to return the result of the parent method!
    return super(MyEncoder, self)._callback_write(buf)

# Override PiCamera to use our custom encoder for video recording
class MyCamera(picamera.PiCamera):
    def __init__(self):
        super(MyCamera, self).__init__()
        self.i_frames = 0
        self.p_frames = 0

    def _get_video_encoder(
        self, camera_port, output_port, format, resize, **options):
        return MyEncoder(
            self, camera_port, output_port, format, resize, **options)

with MyCamera() as camera:
    camera.start_recording('foo.h264')
    camera.wait_recording(10)
    camera.stop_recording()
    print('Recording contains %d I-frames and %d P-frames' % (
        camera.i_frames, camera.p_frames))

```

Please note that the above recipe is flawed: PiCamera is capable of initiating *multiple simultaneous recordings*. If this were used with the above recipe, then each encoder would wind up incrementing the `i_frames` and `p_frames` attributes on the `MyCamera` instance leading to incorrect results.

New in version 1.5.

2.5.11 Raw Bayer data captures

The `bayer` parameter of the `capture()` method causes the raw Bayer data recorded by the camera's sensor to be output as part of the image metadata.


Note: The `bayer` parameter only operates with the JPEG format, and only for captures from the still port (i.e. when `use_video_port` is `False`, as it is by default).

Raw Bayer data differs considerably from simple unencoded captures; it is the data recorded by the camera's sensor prior to *any* GPU processing including auto white balance, vignette compensation, smoothing, down-scaling, etc. This also means:

- Bayer data is *always* full resolution, regardless of the camera's output `resolution` and any `resize` parameter.
- Bayer data occupies the last 6,404,096 bytes of the output file. The first 32,768 bytes of this is header data which starts with the string `'BRCM'`.
- Bayer data consists of 10-bit values, because this is the sensitivity of the `OV5647` sensor used by the Pi's camera. The 10-bit values are organized as 4 8-bit values, followed by the low-order 2-bits of the 4 values packed into a fifth byte.

	Bits								
	MSB	8	7	6	5	4	3	2	LSB
1	10	9	8	7	6	5	4	3	2
2	10	9	8	7	6	5	4	3	2
3	10	9	8	7	6	5	4	3	2
4	10	9	8	7	6	5	4	3	2
5	2	1	2	1	2	1	2	1	2

- Bayer data is organized in a BGGR pattern (a minor variation of the common Bayer CFA). The raw data therefore has twice as many green pixels as red or blue and if viewed “raw” will look distinctly strange (too dark, too green, and with zippering effects along any straight edges).



- To make a “normal” looking image from raw Bayer data you will need to perform [de-mosaicing](#) at the very least, and probably some form of [color balance](#).

This (heavily commented) example script causes the camera to capture an image including the raw Bayer data. It then proceeds to unpack the Bayer data into a 3-dimensional [numpy](#) array representing the raw RGB data and finally performs a rudimentary de-mosaic step with weighted averages. A couple of numpy tricks are used to improve performance but bear in mind that all processing is happening on the CPU and will be considerably slower than normal image captures:

```
from __future__ import (
    unicode_literals,
    absolute_import,
    print_function,
    division,
)
```

```
import io
import time
import picamera
```

```

import numpy as np
from numpy.lib.stride_tricks import as_strided

stream = io.BytesIO()
with picamera.PiCamera() as camera:
    # Let the camera warm up for a couple of seconds
    time.sleep(2)
    # Capture the image, including the Bayer data
    camera.capture(stream, format='jpeg', bayer=True)

# Extract the raw Bayer data from the end of the stream, check the
# header and strip it off before converting the data into a numpy array

data = stream.getvalue()[-6404096:]
assert data[:4] == 'BRCM'
data = data[32768:]
data = np.fromstring(data, dtype=np.uint8)

# The data consists of 1952 rows of 3264 bytes of data. The last 8 rows
# of data are unused (they only exist because the actual resolution of
# 1944 rows is rounded up to the nearest 16). Likewise, the last 24
# bytes of each row are unused (why?). Here we reshape the data and
# strip off the unused bytes

data = data.reshape((1952, 3264))[:1944, :3240]

# Horizontally, each row consists of 2592 10-bit values. Every four
# bytes are the high 8-bits of four values, and the 5th byte contains
# the packed low 2-bits of the preceding four values. In other words,
# the bits of the values A, B, C, D and arranged like so:
#
# byte 1   byte 2   byte 3   byte 4   byte 5
# AAAAAAAA BBBBBBBB CCCCCCCC DDDDDDDD AABBCDD
#
# Here, we convert our data into a 16-bit array, shift all values left
# by 2-bits and unpack the low-order bits from every 5th byte in each
# row, then remove the columns containing the packed bits

data = data.astype(np.uint16) << 2
for byte in range(4):
    data[:, byte::5] |= ((data[:, 4::5] >> ((4 - byte) * 2)) & 0b11)
data = np.delete(data, np.s_[4::5], 1)

# Now to split the data up into its red, green, and blue components. The
# Bayer pattern of the OV5647 sensor is BGGR. In other words the first
# row contains alternating green/blue elements, the second row contains
# alternating red/green elements, and so on as illustrated below:
#
# GBGBGBGBGBGBGB
# RGRGRGRGRGRGRG
# GBGBGBGBGBGBGB
# RGRGRGRGRGRGRG
#
# Please note that if you use vflip or hflip to change the orientation
# of the capture, you must flip the Bayer pattern accordingly

rgb = np.zeros(data.shape + (3,), dtype=data.dtype)
rgb[1::2, 0::2, 0] = data[1::2, 0::2] # Red
rgb[0::2, 0::2, 1] = data[0::2, 0::2] # Green
rgb[1::2, 1::2, 1] = data[1::2, 1::2] # Green
rgb[0::2, 1::2, 2] = data[0::2, 1::2] # Blue

# At this point we now have the raw Bayer data with the correct values

```

```
# and colors but the data still requires de-mosaicing and
# post-processing. If you wish to do this yourself, end the script here!
#
# Below we present a fairly naive de-mosaic method that simply
# calculates the weighted average of a pixel based on the pixels
# surrounding it. The weighting is provided by a byte representation of
# the Bayer filter which we construct first:

bayer = np.zeros(rgb.shape, dtype=np.uint8)
bayer[1::2, 0::2, 0] = 1 # Red
bayer[0::2, 0::2, 1] = 1 # Green
bayer[1::2, 1::2, 1] = 1 # Green
bayer[0::2, 1::2, 2] = 1 # Blue

# Allocate an array to hold our output with the same shape as the input
# data. After this we define the size of window that will be used to
# calculate each weighted average (3x3). Then we pad out the rgb and
# bayer arrays, adding blank pixels at their edges to compensate for the
# size of the window when calculating averages for edge pixels.

output = np.empty(rgb.shape, dtype=rgb.dtype)
window = (3, 3)
borders = (window[0] - 1, window[1] - 1)
border = (borders[0] // 2, borders[1] // 2)

rgb_pad = np.zeros((
    rgb.shape[0] + borders[0],
    rgb.shape[1] + borders[1],
    rgb.shape[2]), dtype=rgb.dtype)
rgb_pad[
    border[0]:rgb_pad.shape[0] - border[0],
    border[1]:rgb_pad.shape[1] - border[1],
    :] = rgb
rgb = rgb_pad

bayer_pad = np.zeros((
    bayer.shape[0] + borders[0],
    bayer.shape[1] + borders[1],
    bayer.shape[2]), dtype=bayer.dtype)
bayer_pad[
    border[0]:bayer_pad.shape[0] - border[0],
    border[1]:bayer_pad.shape[1] - border[1],
    :] = bayer
bayer = bayer_pad

# In numpy >=1.7.0 just use np.pad (version in Raspbian is 1.6.2 at the
# time of writing...)
#
#rgb = np.pad(rgb, [
#    (border[0], border[0]),
#    (border[1], border[1]),
#    (0, 0),
#    ], 'constant')
#bayer = np.pad(bayer, [
#    (border[0], border[0]),
#    (border[1], border[1]),
#    (0, 0),
#    ], 'constant')

# For each plane in the RGB data, we use a nifty numpy trick
# (as_strided) to construct a view over the plane of 3x3 matrices. We do
# the same for the bayer array, then use Einstein summation on each
# (np.sum is simpler, but copies the data so it's slower), and divide
```

```

# the results to get our weighted average:

for plane in range(3):
    p = rgb[..., plane]
    b = bayer[..., plane]
    pview = as_strided(p, shape=(
        p.shape[0] - borders[0],
        p.shape[1] - borders[1]) + window, strides=p.strides * 2)
    bview = as_strided(b, shape=(
        b.shape[0] - borders[0],
        b.shape[1] - borders[1]) + window, strides=b.strides * 2)
    psum = np.einsum('ijkl->ij', pview)
    bsum = np.einsum('ijkl->ij', bview)
    output[..., plane] = psum // bsum

# At this point output should contain a reasonably "normal" looking
# image, although it still won't look as good as the camera's normal
# output (as it lacks vignette compensation, AWB, etc).
#
# If you want to view this in most packages (like GIMP) you'll need to
# convert it to 8-bit RGB data. The simplest way to do this is by
# right-shifting everything by 2-bits (yes, this makes all that
# unpacking work at the start rather redundant...)

output = (output >> 2).astype(np.uint8)
with open('image.data', 'wb') as f:
    output.tofile(f)

```

This recipe is also encapsulated in the `PiBayerArray` class in the `picamera.array` module, which means the same can be achieved as follows:

```

import time
import picamera
import picamera.array
import numpy as np

with picamera.PiCamera() as camera:
    with picamera.array.PiBayerArray(camera) as stream:
        camera.capture(stream, 'jpeg', bayer=True)
        # Demosaic data and write to output (just use stream.array if you
        # want to skip the demosaic step)
        output = (stream.demosaic() >> 2).astype(np.uint8)
        with open('image.data', 'wb') as f:
            output.tofile(f)

```

New in version 1.3.

Changed in version 1.5: Added note about new `picamera.array` module.

2.6 Frequently Asked Questions (FAQ)

2.6.1 Can I put the preview in a window?

No. The camera module's preview system is quite crude: it simply tells the GPU to overlay the preview on the Pi's video output. The preview has no knowledge (or interaction with) the X-Windows environment (incidentally, this is why the preview works quite happily from the command line, even without anyone logged in).

That said, the preview area can be resized and repositioned via the `window` attribute of the `preview` object. If your program can respond to window repositioning and sizing events you can "cheat" and position the preview within the borders of the target window. However, there's currently no way to allow anything to appear on top of the preview so this is an imperfect solution at best.

2.6.2 Help! I started a preview and can't see my console!

As mentioned above, the preview is simply an overlay over the Pi's video output. If you start a preview you may therefore discover you can't see your console anymore and there's no obvious way of getting it back. If you're confident in your typing skills you can try calling `stop_preview()` by typing "blindly" into your hidden console. However, the simplest way of getting your display back is usually to hit `Ctrl+D` to terminate the Python process (which should also shut down the camera).

When starting a preview, you may want to set the *alpha* parameter of the `start_preview()` method to something like 128. This should ensure that when the preview is displayed, it is partially transparent so you can still see your console.

2.6.3 How much power does the camera require?

The camera requires 250mA when running. Note that simply creating a `PiCamera` object means the camera is running (due to the hidden preview that is started to allow the auto-exposure algorithm to run). If you are running your Pi from batteries, you should `close()` (or `destroy()`) the instance when the camera is not required in order to conserve power. For example, the following code captures 60 images over an hour, but leaves the camera running all the time:

```
import picamera
import time

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    time.sleep(1) # Camera warm-up time
    for i, filename in enumerate(camera.capture_continuous('image{counter:02d}.jpg')):
        print('Captured %s' % filename)
        # Capture one image a minute
        time.sleep(60)
        if i == 59:
            break
```

By contrast, this code closes the camera between shots (but can't use the convenient `capture_continuous()` method as a result):

```
import picamera
import time

for i in range(60):
    with picamera.PiCamera() as camera:
        camera.resolution = (1280, 720)
        time.sleep(1) # Camera warm-up time
        filename = 'image%02d.jpg' % i
        camera.capture(filename)
        print('Captured %s' % filename)
    # Capture one image a minute
    time.sleep(59)
```

Note: Please note the timings in the scripts above are approximate. A more precise example of timing is given in *Capturing timelapse sequences*.

If you are experiencing lockups or reboots when the camera is active, your power supply may be insufficient. A practical minimum is 1A for running a Pi with an active camera module; more may be required if additional peripherals are attached.

2.6.4 How can I take two consecutive pictures with equivalent settings?

See the *Capturing consistent images* recipe.

2.6.5 Can I use picamera with a USB webcam?

No. The picamera library relies on libmmal which is specific to the Pi's camera module.

2.6.6 How can I tell what version of picamera I have installed?

The picamera library relies on the `setuptools` package for installation services. You can use the `setuptools` `pkg_resources` API to query which version of picamera is available in your Python environment like so:

```
>>> from pkg_resources import require
>>> require('picamera')
[picamera 1.2 (/usr/local/lib/python2.7/dist-packages)]
>>> require('picamera')[0].version
'1.2'
```

If you have multiple versions installed (e.g. from `pip` and `apt-get`) they will not show up in the list returned by the `require` method. However, the first entry in the list will be the version that `import picamera` will import.

If you receive the error “No module named `pkg_resources`”, you need to install the `pip` utility. This can be done with the following command in Raspbian:

```
$ sudo apt-get install python-pip
```

2.6.7 How come I can't upgrade to the latest version?

If you are using Raspbian, firstly check that you haven't got both a PyPI (`pip`) and an `apt` (`apt-get`) installation of picamera installed simultaneously. If you have, one will be taking precedence and it may not be the most up to date version.

Secondly, please understand that while the PyPI release process is entirely automated (so as soon as a new picamera release is announced, it will be available on PyPI), the release process for Raspbian packages is semi-manual. There is typically a delay of a few days after a release before updated picamera packages become accessible in the Raspbian repository.

Users desperate to try the latest version make choose to uninstall their `apt` based copy (uninstall instructions are provided in the *installation instructions*, and install using *pip instead*. However, be aware that keeping a PyPI based installation up to date is a more manual process (sticking with `apt` ensures everything gets upgraded with a simple `sudo apt-get upgrade` command).

2.7 Camera Hardware

This chapter attempts to provide an overview of the operation of the camera under various conditions, as well as to provide an introduction to the low level software interface that picamera utilizes.

2.7.1 Camera Modes

The Pi's camera has a discrete set of input modes which are as follows:

Resolution	Aspect Ratio	Framerates	Video	Image	FoV
2592x1944	4:3	1-15fps	x	x	Full
1296x972	4:3	1-42fps	x		Full
1296x730	16:9	1-49fps	x		Full
640x480	4:3	42.1-60fps	x		Full
640x480	4:3	60.1-90fps	x		Full
1920x1080	16:9	1-30fps	x		Partial

Note: This table is accurate as of firmware revision #656. Firmwares prior to this had a more restricted set of modes, and all video modes had partial FoV. Please use `sudo rpi-update` to upgrade to the latest firmware.

Modes with full field of view (FoV) capture from the whole area of the camera's sensor (2592x1944 pixels). Modes with partial FoV only capture from the center 1920x1080 pixels. The difference between these areas is shown in the illustration below:



Which input mode is used cannot be *directly* controlled, but is selected based on the requested `resolution` and `framerate`. The rules governing which input mode is selected are as follows:

- The mode must be acceptable. Video modes can be used for video recording, or for image captures from the video port (i.e. when `use_video_port` is `True` in calls to the various capture methods). Image captures when `use_video_port` is `False` must use an image mode (of which only one currently exists).
- The closer the requested `resolution` is to the mode's resolution the better, but downscaling from a higher input resolution is preferable to upscaling from a lower input resolution.
- The requested `framerate` should be within the range of the input mode. Note that this is not a hard restriction (it is possible, but unlikely, for the camera to select a mode that does not support the requested framerate).
- The closer the aspect ratio of the requested `resolution` is to the mode's resolution, the better. Attempts to set resolutions with aspect ratios other than 4:3 or 16:9 (which are the only ratios directly supported by the modes in the table above) will choose the mode which maximizes the resulting FoV.

A few examples are given below to clarify the operation of this heuristic:

- If you set the `resolution` to 1024x768 (a 4:3 aspect ratio), and `framerate` to anything less than 42fps, the 1296x976 mode will be selected, and the camera will downscale the result to 1024x768.
- If you set the `resolution` to 1280x720 (a 16:9 wide-screen aspect ratio), and `framerate` to anything less than 49fps, the 1296x730 mode will be selected and downscaled appropriately.

- Setting `resolution` to 1920x1080 and `framerate` to 30fps exceeds the resolution of both the 1296x730 and 1296x976 modes (i.e. they would require upscaling), so the 1920x1080 mode is selected instead, although it has a reduced FoV.
- A `resolution` of 800x600 and a `framerate` of 60fps will select the 640x480 60fps mode, even though it requires upscaling because the algorithm considers the framerate to take precedence in this case.
- Any attempt to capture an image without using the video port will (temporarily) select the 2592x1944 mode while the capture is performed (this is what causes the flicker you sometimes see when a preview is running while a still image is captured).

2.7.2 Under the Hood

This section attempts to provide detail of what picamera is doing “under the hood” in response to various method calls.

The Pi’s camera has three ports, the still port, the video port, and the preview port. The following sections describe how these ports are used by picamera and how they influence the camera’s resolutions.

The Still Port

Firstly, the still port. Whenever this is used to capture images, it (briefly) forces the camera’s mode to the only supported still mode (see *Camera Modes*) so that images are captured using the full area of the sensor. It also appears to perform a considerable amount of post-processing on captured images so that they appear higher quality.

The still port is used by the various `capture()` methods when their `use_video_port` parameter is `False` (which it is by default).

The Video Port

The video port is somewhat simpler in that it never changes the camera’s mode. The video port is used by the `start_recording()` method (for recording video), and is also used by the various `capture()` methods when their `use_video_port` parameter is `True`. Images captured from the video port tend to have a “grainy” appearance, much more akin to a video frame than the images captured by the still port (this is due to the still port uses a slower, more aggressive denoise algorithm).

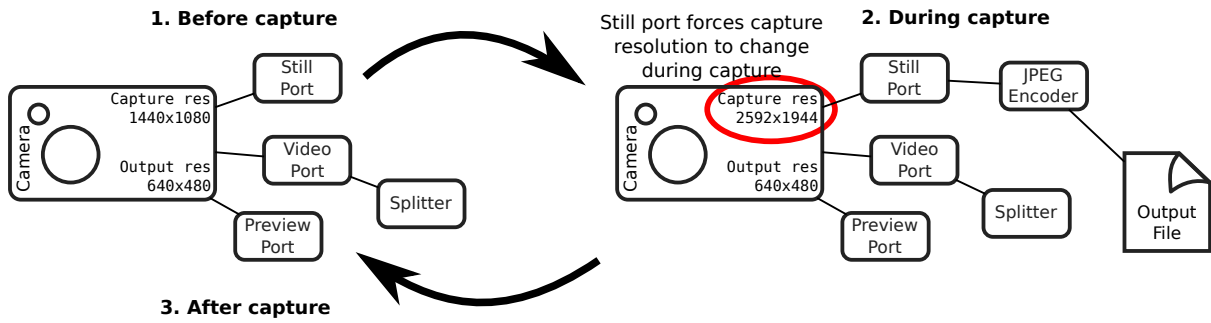
The Preview Port

The preview port operates more or less identically to the video port. As the preview port is never used for encoding we won’t mention it further in this section.

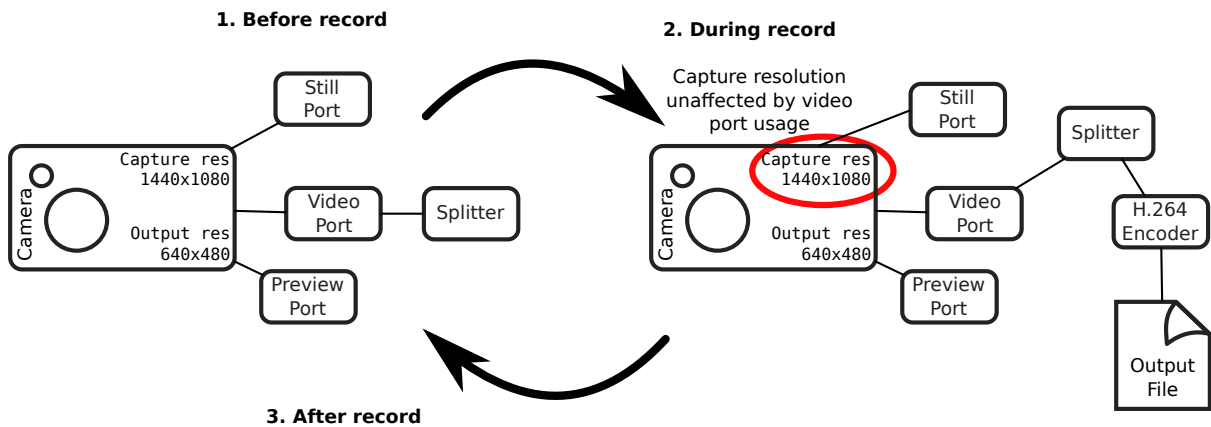
Encoders

The camera provides various encoders which can be attached to the still and video ports for the purpose of producing output (e.g. JPEG images or H.264 encoded video). A port can have a single encoder attached to it at any given time (or nothing if the port is not in use).

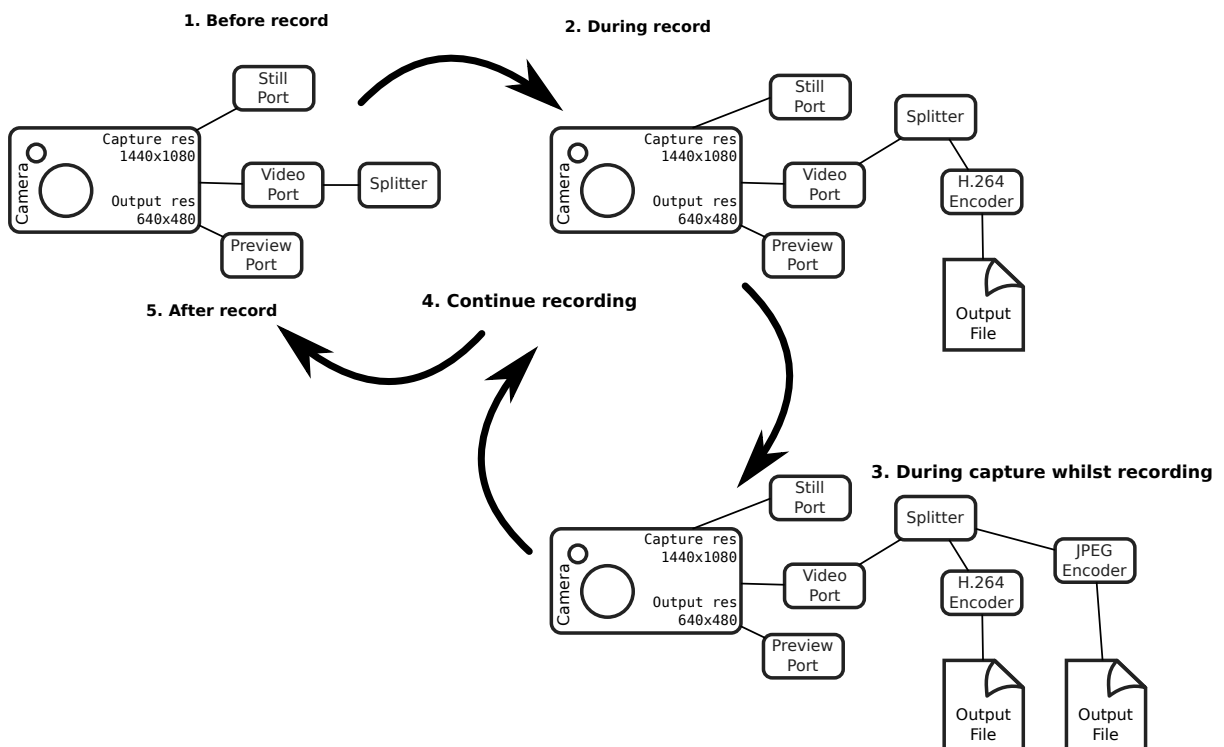
Encoders are connected directly to the still port. For example, when capturing a picture using the still port, the camera’s state conceptually moves through these states:



As you have probably noticed in the diagram above, the video port is a little more complex. In order to permit simultaneous video recording and image capture via the video port, a “splitter” component is permanently connected to the video port by picamera, and encoders are in turn attached to one of its four output ports (numbered 0, 1, 2, and 3). Hence, when recording video the camera’s setup looks like this:

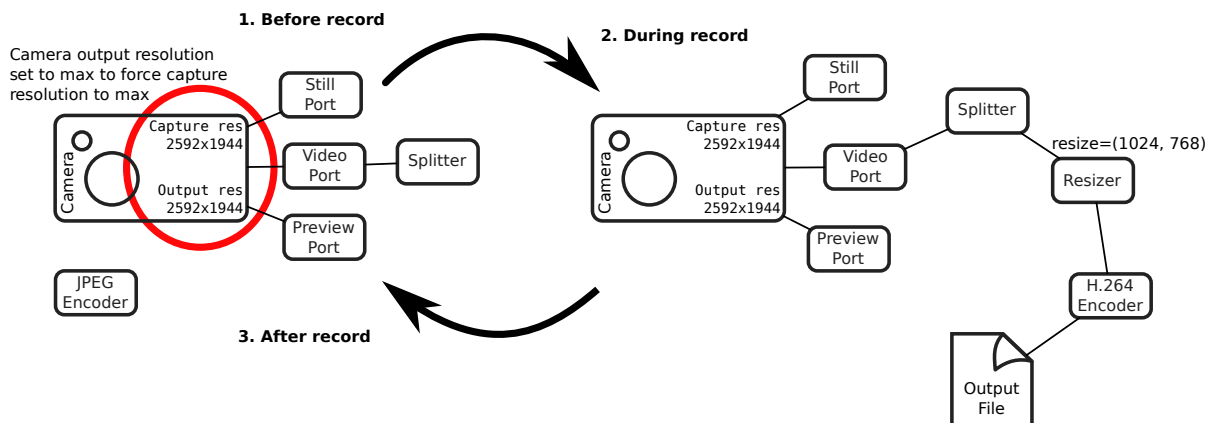


And when simultaneously capturing images via the video port whilst recording, the camera’s configuration moves through the following states:




When the `resize` parameter is passed to one of the aforementioned methods, a `resizer` component is placed between the camera’s ports and the encoder, causing the output to be resized before it reaches the encoder. This is particularly useful for video recording, as the H.264 encoder cannot cope with full resolution input. Hence, when

performing full frame video recording, the camera's setup looks like this:



Finally, when performing unencoded captures an encoder is (naturally) not required. Instead data is taken directly from the camera's ports. When raw YUV format is requested no components are attached to the ports at all (as all ports default to YUV output). Likewise, when capturing unencoded from the still port no encoders are used; instead the still port is reconfigured to output the required encoding (RGB, RGBA, BGR, etc.)

However, the video port is fixed in YUV mode (the video encoders expect YUV input and the splitter cannot convert encodings). When another raw format like RGBA is requested, a resizer is used (with its output resolution set to the input resolution, unless the `resize` option is specified with something different), and its output format is set to the requested raw format:



Please note that even the description above is almost certainly far removed from what actually happens at the camera's ISP level. Rather, what has been described in this section is how the MMAL library exposes the camera to applications which utilize it (these include the picamera library, along with the official *raspistill* and *raspivid* applications).

In other words, by using picamera you are passing through (at least) two abstraction layers which necessarily obscure (but hopefully simplify) the "true" operation of the camera.

2.8 Deprecated Functionality

The picamera library is (at the time of writing) nearly a year old and has grown quite rapidly in this time. Occasionally, when adding new functionality to the library, the API is obvious and natural (e.g. `start_recording()` and `stop_recording()`). At other times, it's been less obvious (e.g. unencoded captures) and my initial attempts have proven to be less than ideal. In such situations I've endeavoured to improve the API without breaking backward compatibility by introducing new methods or attributes and deprecating the old ones.

This means that, as of release 1.8, there's quite a lot of deprecated functionality floating around the library which it would be nice to tidy up, partly to simplify the library for debugging, and partly to simplify it for new users. To assuage any fears that I'm imminently going to break backward compatibility: I intend to leave a gap of at least a year between deprecating functionality and removing it, hopefully providing ample time for people to migrate their scripts.

Furthermore, to distinguish any release which is backwards incompatible, I would increment the major version number in accordance with [semantic versioning](#). In other words, the first release in which currently deprecated

functionality would be removed would be version 2.0, and as of the release of 1.8 it's at least a year away. Any future 1.x releases will include all currently deprecated functions.

Of course, that still means people need a way of determining whether their scripts use any deprecated functionality in the picamera library. All deprecated functionality is documented, and the documentation includes pointers to the intended replacement functionality (see `raw_format` for example). However, Python also provides excellent methods for determining automatically whether any deprecated functionality is being used via the `warnings` module.

2.8.1 Finding and fixing deprecated usage

As of release 1.8, all deprecated functionality will raise `DeprecationWarning` when used. By default, the Python interpreter suppresses these warnings (as they're only of interest to developers, not users) but you can easily configure different behaviour.

The following example script uses a number of deprecated functions:

```
import io
import time
import picamera

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.framerate = (24, 1)
    camera.start_preview()
    camera.preview_fullscreen = True
    camera.preview_alpha = 128
    time.sleep(2)
    camera.raw_format = 'yuv'
    stream = io.BytesIO()
    camera.capture(stream, 'raw', use_video_port=True)
```

Despite using deprecated functionality the script runs happily (and silently) with picamera 1.8. To discover what deprecated functions are being used, we add a couple of lines to tell the warnings module that we want “default” handling of `DeprecationWarning`; “default” handling means that the first time an attempt is made to raise this warning at a particular location, the warning's details will be printed to the console. All future invocations from the same location will be ignored. This saves flooding the console with warning details from tight loops. With this change, the script looks like this:

```
import io
import time
import picamera

import warnings
warnings.filterwarnings('default', category=DeprecationWarning)

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.framerate = (24, 1)
    camera.start_preview()
    camera.preview_fullscreen = True
    camera.preview_alpha = 128
    time.sleep(2)
    camera.raw_format = 'yuv'
    stream = io.BytesIO()
    camera.capture(stream, 'raw', use_video_port=True)
```

And produces the following output on the console when run:

```
/usr/share/pyshared/picamera/camera.py:149: DeprecationWarning: Setting framerate or gains as a tuple is deprecated; "
Setting framerate or gains as a tuple is deprecated; "
/usr/share/pyshared/picamera/camera.py:3125: DeprecationWarning: PiCamera.preview_fullscreen is deprecated
```

```
'PiCamera.preview_fullscreen is deprecated; '
/usr/share/pyshared/picamera/camera.py:3068: DeprecationWarning: PiCamera.preview_alpha is deprecated; use PiCamera.preview.alpha
'PiCamera.preview_alpha is deprecated; use '
/usr/share/pyshared/picamera/camera.py:1833: DeprecationWarning: PiCamera.raw_format is deprecated; use PiCamera.raw_format.required_format
'PiCamera.raw_format is deprecated; use required format '
/usr/share/pyshared/picamera/camera.py:1359: DeprecationWarning: The "raw" format option is deprecated; use PiCamera.raw_format.required_format
'The "raw" format option is deprecated; specify the '
/usr/share/pyshared/picamera/camera.py:1827: DeprecationWarning: PiCamera.raw_format is deprecated; use PiCamera.raw_format.required_format
'PiCamera.raw_format is deprecated; use required format '
```

This tells us which pieces of deprecated functionality are being used in our script, but it doesn't tell us where in the script they were used. For this, it is more useful to have warnings converted into full blown exceptions. With this change, each time a `DeprecationWarning` would have been printed, it will instead cause the script to terminate with an unhandled exception and a full stack trace:

```
import io
import time
import picamera

import warnings
warnings.filterwarnings('error', category=DeprecationWarning)

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.framerate = (24, 1)
    camera.start_preview()
    camera.preview_fullscreen = True
    camera.preview_alpha = 128
    time.sleep(2)
    camera.raw_format = 'yuv'
    stream = io.BytesIO()
    camera.capture(stream, 'raw', use_video_port=True)
```

Now when we run the script it produces the following:

```
Traceback (most recent call last):
  File "test_deprecated.py", line 10, in <module>
    camera.framerate = (24, 1)
  File "/usr/share/pyshared/picamera/camera.py", line 1888, in _set_framerate
    n, d = to_rational(value)
  File "/usr/share/pyshared/picamera/camera.py", line 149, in to_rational
    "Setting framerate or gains as a tuple is deprecated; "
DeprecationWarning: Setting framerate or gains as a tuple is deprecated; please use one of Python
```

This tells us that line 10 of our script is using deprecated functionality, and provides a hint of how to fix it. We change line 10 to use an int instead of a tuple for the framerate. Now we run again, and this time get the following:

```
Traceback (most recent call last):
  File "test_deprecated.py", line 12, in <module>
    camera.preview_fullscreen = True
  File "/usr/share/pyshared/picamera/camera.py", line 3125, in _set_preview_fullscreen
    'PiCamera.preview_fullscreen is deprecated; '
DeprecationWarning: PiCamera.preview_fullscreen is deprecated; use PiCamera.preview.fullscreen
```

Now we can tell line 12 has a problem, and once again the exception tells us how to fix it. We continue in this fashion until the script looks like this:

```
import io
import time
import picamera

import warnings
warnings.filterwarnings('error', category=DeprecationWarning)
```

```
with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    camera.framerate = 24
    camera.start_preview()
    camera.preview.fullscreen = True
    camera.preview.alpha = 128
    time.sleep(2)
    stream = io.BytesIO()
    camera.capture(stream, 'yuv', use_video_port=True)
```

The script now runs to completion, so we can be confident it's no longer using any deprecated functionality and will run happily even when this functionality is removed in release 2.0. At this point, you may wish to remove the `filterwarnings` line as well (or at least comment it out).

2.8.2 List of deprecated functionality

For convenience, all currently deprecated functionality is detailed below. You may wish to skim this list to check whether you're currently using deprecated functions, but I would urge users to take advantage of the warnings system documented in the prior section as well.

Unencoded captures

In very early versions of `picamera`, unencoded captures were created by specifying the `'raw'` format with the `capture()` method, with the `raw_format` attribute providing the actual encoding. The attribute is deprecated, as is usage of the value `'raw'` with the `format` parameter of all the capture methods. Hence, code like this:

```
camera.raw_format = 'rgb'
camera.capture('output.data', format='raw')
```

should be replaced with the following:

```
camera.capture('output.data', format='rgb')
```

Recording quality

The `quantization` parameter for `start_recording()` and `record_sequence()` is deprecated in favor of the `quality` parameter; this change was made to keep the recording methods consistent with the capture methods, and to make the meaning of the parameter more obvious to newcomers. The values of the parameter remain the same (i.e. 1-100 for MJPEG recordings with higher values indicating higher quality, and 1-40 for H.264 recordings with lower values indicating higher quality). Hence, the following code:

```
camera.start_recording('foo.h264', quantization=25)
```

should be replaced with:

```
camera.start_recording('foo.h264', quality=25)
```

Fractions as tuples

Several attributes in `picamera` expect rational (fractional) values. In early versions of `picamera`, these values could only be specified as a tuple expressed as `(numerator, denominator)`. In later versions, support was expanded to accept any of Python's numeric types including `int`, `float`, `Decimal`, and `Fraction`. Hence, the following code:

```
camera.framerate = (24, 1)
```

can be replaced with any of the following variations:

```

from decimal import Decimal
from fractions import Fraction

camera.framerate = 24
camera.framerate = 24.0
camera.framerate = Fraction(72, 3)
camera.framerate = Decimal('24')
camera.framerate = Fraction('48/2')

```

These attributes return a `Fraction` instance as well, but one modified to permit access as a tuple in order to maintain backward compatibility. This is also deprecated, hence code like this:

```

n, d = camera.framerate
print('The framerate is %d/%d fps' % (n, d))

```

should be updated to this:

```

f = camera.framerate
print('The framerate is %d/%d fps' % (f.numerator, f.denominator))

```

Although you may wish to simply convert the `Fraction` instance to a *float* for greater convenience:

```

f = float(camera.framerate)
print('The framerate is %0.2f fps' % f)

```

Preview functions

Release 1.8 introduced rather sweeping changes to the preview system to incorporate the ability to create multiple static overlays on top of the preview. As a result, the preview system is no longer “part of” the `PiCamera` class. Instead, it is represented by the `preview` attribute which is a separate `PiPreviewRenderer` instance when the preview is active. In turn this meant that `preview_alpha` was deprecated in favor of the `alpha` property of the new `preview` attribute. Similar changes were made to `preview_layer`, `preview_fullscreen`, and `preview_window`. Hence, the following code:

```

camera.start_preview()
camera.preview_alpha = 128
camera.preview_fullscreen = False
camera.preview_window = (0, 0, 640, 480)

```

should be replaced with:

```

camera.start_preview()
camera.preview.alpha = 128
camera.preview.fullscreen = False
camera.preview.window = (0, 0, 640, 480)

```

Not an enormous change, but the eagle-eyed will have noticed that the `preview` attribute is set to `None` when the preview is not active. This means that setting the `alpha` property *before* the preview is active is no longer possible. To accommodate this use-case, optional parameters were added to `start_preview()` to provide initial settings for the preview renderer. Therefore, the following code:

```

camera.preview_alpha = 128
camera.preview_fullscreen = False
camera.preview_window = (0, 0, 640, 480)
camera.start_preview()

```

should be replaced with:

```

camera.start_preview(
    alpha=128, fullscreen=False, window=(0, 0, 640, 480))

```

Finally, the `previewing` attribute is now obsolete (and thus deprecated) as its functionality can be trivially obtained by checking the `preview` attribute. Hence, the following code:

```
if camera.previewing:
    print('The camera preview is running')
else:
    print('The camera preview is not running')
```

can be replaced with:

```
if camera.preview:
    print('The camera preview is running')
else:
    print('The camera preview is not running')
```

Array stream truncation

In release 1.8, the base `PiArrayOutput` class was changed to derive from `io.BytesIO` in order to add support for seeking, and to improve performance. The prior implementation had been non-seekable, and therefore to accommodate re-use of the stream between captures the `truncate()` method had an unusual side-effect not seen with regular Python streams: after truncation, the position of the stream was set to the new length of the stream. In all other Python streams, the `truncate` method doesn't affect the stream position. The method is overridden in 1.8 to maintain its unusual behaviour, but this behaviour is nonetheless deprecated.

If you only need your code to work with the latest version of picamera you can replace calls like the following:

```
with picamera.array.PiYUVArray(camera) as stream:
    for i in range(3):
        camera.capture(stream, 'yuv')
        print(stream.array.shape)
        stream.truncate(0)
```

with this:

```
with picamera.array.PiYUVArray(camera) as stream:
    for i in range(3):
        camera.capture(stream, 'yuv')
        print(stream.array.shape)
        stream.seek(0)
        stream.truncate()
```

Unfortunately, this will not work if your script needs to work with prior versions of picamera as well (since such streams were non-seekable in prior versions). In this case, call `seekable()` to determine the correct course of action:

```
with picamera.array.PiYUVArray(camera) as stream:
    for i in range(3):
        camera.capture(stream, 'yuv')
        print(stream.array.shape)
        if stream.seekable():
            stream.seek(0)
            stream.truncate()
        else:
            stream.truncate(0)
```

Confusing crop

In release 1.8, the `crop` attribute was renamed to `zoom`; the old name was retained as a deprecated alias for backward compatibility. This change was made as `crop` was a thoroughly misleading name for the attribute (which actually sets the “region of interest” for the sensor), leading to numerous support questions. Hence, the following code:

```
camera.crop = (0.25, 0.25, 0.5, 0.5)
```

should be changed to:

```
camera.zoom = (0.25, 0.25, 0.5, 0.5)
```

Incorrect ISO capitalisation

In release 1.8, the `ISO` attribute was renamed to `iso` for compliance with [PEP-8](#) (even though it's an acronym this is still more consistent with the existing API; consider `led`, `awb_mode`, and so on). This means the following code:

```
camera.ISO = 100
```

should simply be replaced with:

```
camera.iso = 100
```

Frame types

Over time, several capabilities were added to the H.264 encoder in the GPU firmware. This expanded the number of possible frame types from a simple key-frame / non-key-frame affair, to a multitude of possibilities (P-frame, I-frame, SPS/PPS header, motion vector data, and who knows in future). Rather than keep adding more and more boolean fields to the `PiVideoFrame` named tuple, release 1.5 introduced the `PiVideoFrameType` enumeration used by the `frame_type` attribute and deprecated the `keyframe` and `header` attributes. Hence, the following code:

```
if camera.frame.keyframe:
    handle_keyframe()
elif camera.frame.header:
    handle_header()
else:
    handle_frame()
```

should be replaced with:

```
if camera.frame.frame_type == picamera.PiVideoFrameType.key_frame:
    handle_keyframe()
elif camera.frame.frame_type == picamera.PiVideoFrameType.sps_header:
    handle_header()
else:
    handle_frame()
```

Alternatively, you may find something like this more elegant (and more future proof as it'll throw a `KeyError` in the event of an unrecognized frame type):

```
handler = {
    picamera.PiVideoFrameType.key_frame: handle_keyframe,
    picamera.PiVideoFrameType.sps_header: handle_header,
    picamera.PiVideoFrameType.frame: handle_frame,
}[camera.frame.frame_type]
handler()
```

2.9 API Reference

This package primarily provides the `PiCamera` class which is a pure Python interface to the Raspberry Pi's camera module. Various ancillary classes are provided for usage with `PiCamera` including `PiVideoFrame` (for holding video frame meta-data), `PiCameraCircularIO` (for recording video to a ring-buffer), `PiEncoder` (an abstract base class for camera encoders), and the concrete encoder classes: `PiVideoEncoder`,

`PiCookedOneImageEncoder`, `PiCookedMultiImageEncoder`, `PiRawOneImageEncoder`, and `PiRawMultiImageEncoder`.

Note: In the documentation below several apparently “private” methods are documented (i.e. methods which have names beginning with an underscore). Most users can ignore these methods; they are intended for those developers that wish to override or extend the encoder implementations used by picamera.

Some may question, given that these methods are intended to be overridden, why they are declared with a leading underscore (which in the Python idiom suggests that these methods are “private” to the class). In the cases where such methods are documented, the author intends these methods to have “protected” status (in the idiom of C++ and Object Pascal). That is to say, they are not intended to be used outside of the declaring class, but are intended to be accessible to, and overridden by, descendent classes.

2.9.1 PiCamera

class `picamera.PiCamera`

Provides a pure Python interface to the Raspberry Pi’s camera module.

Upon construction, this class initializes the camera. The `camera_num` parameter (which defaults to 0) selects the camera module that the instance will represent. Only the Raspberry Pi compute module currently supports more than one camera, and this class has not yet been tested with more than one module.

The `stereo_mode` and `stereo_decimate` parameters configure dual cameras on a compute module for stereoscopic mode. These parameters can only be set at construction time; they cannot be altered later without closing the `PiCamera` instance and recreating it. The `stereo_mode` parameter defaults to `'none'` (no stereoscopic mode) but can be set to `'side-by-side'` or `'top-bottom'` to activate a stereoscopic mode. If the `stereo_decimate` parameter is `True`, the resolution of the two cameras will be halved so that the resulting image has the same dimensions as if stereoscopic mode were not being used.

Warning: Stereoscopic mode is untested in picamera at this time. If you have the necessary hardware, the author would be most interested to hear of your experiences!

The `resolution` of the camera is initially set to the display’s resolution unless the display has been disabled (e.g. with `tvservice -o`) in which case a default of 1280x720 is used.

No preview or recording is started automatically upon construction. Use the `capture()` method to capture images, the `start_recording()` method to begin recording video, or the `start_preview()` method to start live display of the camera’s input.

Several attributes are provided to adjust the camera’s configuration. Some of these can be adjusted while a recording is running, like `brightness`. Others, like `resolution`, can only be adjusted when the camera is idle.

When you are finished with the camera, you should ensure you call the `close()` method to release the camera resources:

```
camera = PiCamera()
try:
    # do something with the camera
    pass
finally:
    camera.close()
```

The class supports the context manager protocol to make this particularly easy (upon exiting the `with` statement, the `close()` method is automatically called):

```
with PiCamera() as camera:
    # do something with the camera
    pass
```

`_check_camera_open()`

Raise an exception if the camera is already closed

`_check_recording_stopped()`

Raise an exception if the camera is currently recording

`_get_image_encoder(camera_port, output_port, format, resize, **options)`

Construct an image encoder for the requested parameters.

This method is called by `capture()` and `capture_continuous()` to construct an image encoder. The `camera_port` parameter gives the MMAL camera port that should be enabled for capture by the encoder. The `output_port` parameter gives the MMAL port that the encoder should read output from (this may be the same as the camera port, but may be different if other component(s) like a splitter have been placed in the pipeline). The `format` parameter indicates the image format and will be one of:

- 'jpeg'
- 'png'
- 'gif'
- 'bmp'
- 'yuv'
- 'rgb'
- 'rgba'
- 'bgr'
- 'bgra'

The `resize` parameter indicates the size that the encoder should resize the output to (presumably by including a resizer in the pipeline). Finally, `options` includes extra keyword arguments that should be passed verbatim to the encoder.

`_get_image_format(output, format=None)`

Given an output object and an optional format, attempt to determine the requested image format.

This method is used by all capture methods to determine the requested output format. If `format` is specified as a MIME-type the “image/” prefix is stripped. If `format` is not specified, then `_get_output_format()` will be called to attempt to determine format from the `output` object.

`_get_images_encoder(camera_port, output_port, format, resize, **options)`

Construct a multi-image encoder for the requested parameters.

This method is largely equivalent to `_get_image_encoder()` with the exception that the encoder returned should expect to be passed an iterable of outputs to its `start()` method, rather than a single output object. This method is called by the `capture_sequence()` method.

All parameters are the same as in `_get_image_encoder()`. Please refer to the documentation for that method for further information.

`_get_output_format(output)`

Given an output object, attempt to determine the requested format.

We attempt to determine the filename of the `output` object and derive a MIME type from the extension. If `output` has no filename, an error is raised.

`_get_ports(from_video_port, splitter_port)`

Determine the camera and output ports for given capture options.

See *Under the Hood* for more information on picamera’s usage of camera, splitter, and encoder ports. The general idea here is that the capture (still) port operates on its own, while the video port is always connected to a splitter component, so requests for a video port also have to specify which splitter port they want to use.

`_get_settings()`

Returns the current camera settings as an MMAL structure.

This is a utility method for `_get_exposure_speed()`, `_get_analog_gain()`, etc. all of which rely on the `MMAL_PARAMETER_CAMERA_SETTINGS` structure to determine their values.

`_get_video_encoder(camera_port, output_port, format, resize, **options)`

Construct a video encoder for the requested parameters.

This method is called by `start_recording()` and `record_sequence()` to construct a video encoder. The `camera_port` parameter gives the MMAL camera port that should be enabled for capture by the encoder. The `output_port` parameter gives the MMAL port that the encoder should read output from (this may be the same as the camera port, but may be different if other component(s) like a splitter have been placed in the pipeline). The `format` parameter indicates the video format and will be one of:

- `'h264'`
- `'mjpeg'`

The `resize` parameter indicates the size that the encoder should resize the output to (presumably by including a resizer in the pipeline). Finally, `options` includes extra keyword arguments that should be passed verbatim to the encoder.

`_get_video_format(output, format=None)`

Given an output object and an optional format, attempt to determine the requested video format.

This method is used by all recording methods to determine the requested output format. If `format` is specified as a MIME-type the “video/” or “application/” prefix will be stripped. If `format` is not specified, then `_get_output_format()` will be called to attempt to determine format from the `output` object.

`add_overlay(source, size=None, **options)`

Adds a static overlay to the preview output.

This method creates a new static overlay using the same rendering mechanism as the preview. Overlays will appear on the Pi’s video output, but will not appear in captures or video recordings. Multiple overlays can exist; each call to `add_overlay()` returns a new `PiOverlayRenderer` instance representing the overlay.

The optional `size` parameter specifies the size of the source image as a `(width, height)` tuple. If this is omitted or `None` then the size is assumed to be the same as the camera’s current `resolution`.

The `source` must be an object that supports the `buffer protocol` which has the same length as an image in `RGB` format (colors represented as interleaved unsigned bytes) with the specified `size` after the width has been rounded up to the nearest multiple of 32, and the height has been rounded up to the nearest multiple of 16.

For example, if `size` is `(1280, 720)`, then `source` must be a buffer with length $1280 \times 720 \times 3$ bytes, or 2,764,800 bytes (because 1280 is a multiple of 32, and 720 is a multiple of 16 no extra rounding is required). However, if `size` is `(97, 57)`, then `source` must be a buffer with length $128 \times 64 \times 3$ bytes, or 24,576 bytes (pixels beyond column 97 and row 57 in the source will be ignored).

New overlays default to `layer 0`, whilst the preview defaults to `layer 2`. Higher numbered layers obscure lower numbered layers, hence new overlays will be invisible (if the preview is running) by default. You can make the new overlay visible either by making any existing preview transparent (with the `alpha` property) or by moving the overlay into a layer higher than the preview (with the `layer` property).

All keyword arguments captured in `options` are passed onto the `PiRenderer` constructor. All camera properties except `resolution` and `framerate` can be modified while overlays exist. The reason for these exceptions is that the overlay has a static resolution and changing the camera’s mode would require resizing of the source.

`capture(output, format=None, use_video_port=False, resize=None, splitter_port=0, **options)`

Capture an image from the camera, storing it in `output`.

If *output* is a string, it will be treated as a filename for a new file which the image will be written to. Otherwise, *output* is assumed to be a file-like object and the image data is appended to it (the implementation only assumes the object has a `write()` method - no other methods will be called).

If *format* is `None` (the default), the method will attempt to guess the required image format from the extension of *output* (if it's a string), or from the *name* attribute of *output* (if it has one). In the case that the format cannot be determined, a `PiCameraValueError` will be raised.

If *format* is not `None`, it must be a string specifying the format that you want the image written to. The format can be a MIME-type or one of the following strings:

- `'jpeg'` - Write a JPEG file
- `'png'` - Write a PNG file
- `'gif'` - Write a GIF file
- `'bmp'` - Write a Windows bitmap file
- `'yuv'` - Write the raw image data to a file in YUV420 format
- `'rgb'` - Write the raw image data to a file in 24-bit RGB format
- `'rgba'` - Write the raw image data to a file in 32-bit RGBA format
- `'bgr'` - Write the raw image data to a file in 24-bit BGR format
- `'bgra'` - Write the raw image data to a file in 32-bit BGRA format
- `'raw'` - Deprecated option for raw captures; the format is taken from the deprecated `raw_format` attribute

The *use_video_port* parameter controls whether the camera's image or video port is used to capture images. It defaults to `False` which means that the camera's image port is used. This port is slow but produces better quality pictures. If you need rapid capture up to the rate of video frames, set this to `True`.

When *use_video_port* is `True`, the *splitter_port* parameter specifies the port of the video splitter that the image encoder will be attached to. This defaults to 0 and most users will have no need to specify anything different. This parameter is ignored when *use_video_port* is `False`. See [Under the Hood](#) for more information about the video splitter.

If *resize* is not `None` (the default), it must be a two-element tuple specifying the width and height that the image should be resized to.

Warning: If *resize* is specified, or *use_video_port* is `True`, Exif metadata will **not** be included in JPEG output. This is due to an underlying firmware limitation.

Certain file formats accept additional options which can be specified as keyword arguments. Currently, only the `'jpeg'` encoder accepts additional options, which are:

- *quality* - Defines the quality of the JPEG encoder as an integer ranging from 1 to 100. Defaults to 85. Please note that JPEG quality is not a percentage and [definitions of quality](#) vary widely.
- *thumbnail* - Defines the size and quality of the thumbnail to embed in the Exif metadata. Specifying `None` disables thumbnail generation. Otherwise, specify a tuple of (*width*, *height*, *quality*). Defaults to (64, 48, 35).
- *bayer* - If `True`, the raw bayer data from the camera's sensor is included in the Exif metadata.

Note: The so-called “raw” formats listed above (`'yuv'`, `'rgb'`, etc.) do not represent the raw bayer data from the camera's sensor. Rather they provide access to the image data after GPU processing, but before format encoding (JPEG, PNG, etc). Currently, the only method of accessing the raw bayer data is via the *bayer* parameter described above.

Changed in version 1.0: The *resize* parameter was added, and raw capture formats can now be specified directly

Changed in version 1.3: The `splitter_port` parameter was added, and `bayer` was added as an option for the 'jpeg' format

`capture_continuous` (*output*, *format=None*, *use_video_port=False*, *resize=None*, *splitter_port=0*, *burst=False*, ***options*)

Capture images continuously from the camera as an infinite iterator.

This method returns an infinite iterator of images captured continuously from the camera. If *output* is a string, each captured image is stored in a file named after *output* after substitution of two values with the `format()` method. Those two values are:

- `{counter}` - a simple incremator that starts at 1 and increases by 1 for each image taken
- `{timestamp}` - a `datetime` instance

The table below contains several example values of *output* and the sequence of filenames those values could produce:

<i>output</i> Value	Filenames	Notes
'image{counter}.jpg'	image1.jpg, image2.jpg, image3.jpg, ...	
'image{counter:02d}.jpg'	image01.jpg, image02.jpg, image03.jpg, ...	
'image{timestamp}.jpg'	image2013-10-05 12:07:12.346743.jpg, image2013-10-05 12:07:32.498539, ...	1.
'image{timestamp:%H-%M-%S}.jpg'	image12-10-02-561527.jpg, image12-10-14-905398.jpg, image12-10-03-121002-001.jpg, ...	
'{timestamp:%H%M%S}-{counter:02d}.jpg'	121013-002.jpg, 121014-003.jpg, ...	2.

1. Note that because `timestamp`'s default output includes colons (:), the resulting filenames are not suitable for use on Windows. For this reason (and the fact the default contains spaces) it is strongly recommended you always specify a format when using `{timestamp}`.

2. You can use both `{timestamp}` and `{counter}` in a single format string (multiple times too!) although this tends to be redundant.

If *output* is not a string, it is assumed to be a file-like object and each image is simply written to this object sequentially. In this case you will likely either want to write something to the object between the images to distinguish them, or clear the object between iterations.

The *format*, *use_video_port*, *splitter_port*, *resize*, and *options* parameters are the same as in `capture()`.

If *use_video_port* is `False` (the default), the *burst* parameter can be used to make still port captures faster. Specifically, this prevents the preview from switching resolutions between captures which significantly speeds up consecutive captures from the still port. The downside is that this mode is currently has several bugs; the major issue is that if captures are performed too quickly some frames will come back severely underexposed. It is recommended that users avoid the *burst* parameter unless they absolutely require it and are prepared to work around such issues.

For example, to capture 60 images with a one second delay between them, writing the output to a series of JPEG files named `image01.jpg`, `image02.jpg`, etc. one could do the following:

```
import time
import picamera
with picamera.PiCamera() as camera:
    camera.start_preview()
    try:
        for i, filename in enumerate(camera.capture_continuous('image{counter:02d}.jpg')):
            print(filename)
```

```

        time.sleep(1)
        if i == 59:
            break
    finally:
        camera.stop_preview()

```

Alternatively, to capture JPEG frames as fast as possible into an in-memory stream, performing some processing on each stream until some condition is satisfied:

```

import io
import time
import picamera
with picamera.PiCamera() as camera:
    stream = io.BytesIO()
    for foo in camera.capture_continuous(stream, format='jpeg'):
        # Truncate the stream to the current position (in case
        # prior iterations output a longer image)
        stream.truncate()
        stream.seek(0)
        if process(stream):
            break

```

Changed in version 1.0: The *resize* parameter was added, and raw capture formats can now be specified directly

Changed in version 1.3: The *splitter_port* parameter was added

capture_sequence (*outputs*, *format='jpeg'*, *use_video_port=False*, *resize=None*, *splitter_port=0*, *burst=False*, ***options*)

Capture a sequence of consecutive images from the camera.

This method accepts a sequence or iterator of *outputs* each of which must either be a string specifying a filename for output, or a file-like object with a `write` method. For each item in the sequence or iterator of outputs, the camera captures a single image as fast as it can.

The *format*, *use_video_port*, *splitter_port*, *resize*, and *options* parameters are the same as in `capture()`, but *format* defaults to `'jpeg'`. The format is **not** derived from the filenames in *outputs* by this method.

If *use_video_port* is `False` (the default), the *burst* parameter can be used to make still port captures faster. Specifically, this prevents the preview from switching resolutions between captures which significantly speeds up consecutive captures from the still port. The downside is that this mode is currently has several bugs; the major issue is that if captures are performed too quickly some frames will come back severely underexposed. It is recommended that users avoid the *burst* parameter unless they absolutely require it and are prepared to work around such issues.

For example, to capture 3 consecutive images:

```

import time
import picamera
with picamera.PiCamera() as camera:
    camera.start_preview()
    time.sleep(2)
    camera.capture_sequence([
        'image1.jpg',
        'image2.jpg',
        'image3.jpg',
    ])
    camera.stop_preview()

```

If you wish to capture a large number of images, a list comprehension or generator expression can be used to construct the list of filenames to use:

```

import time
import picamera

```

```
with picamera.PiCamera() as camera:
    camera.start_preview()
    time.sleep(2)
    camera.capture_sequence([
        'image%02d.jpg' % i
        for i in range(100)
    ])
    camera.stop_preview()
```

More complex effects can be obtained by using a generator function to provide the filenames or output objects.

Changed in version 1.0: The *resize* parameter was added, and raw capture formats can now be specified directly

Changed in version 1.3: The *splitter_port* parameter was added

`close()`

Finalizes the state of the camera.

After successfully constructing a `PiCamera` object, you should ensure you call the `close()` method once you are finished with the camera (e.g. in the `finally` section of a `try...finally` block). This method stops all recording and preview activities and releases all resources associated with the camera; this is necessary to prevent GPU memory leaks.

`record_sequence(outputs, format='h264', resize=None, splitter_port=1, **options)`

Record a sequence of video clips from the camera.

This method accepts a sequence or iterator of *outputs* each of which must either be a string specifying a filename for output, or a file-like object with a `write` method.

The method acts as an iterator itself, yielding each item of the sequence in turn. In this way, the caller can control how long to record to each item by only permitting the loop to continue when ready to switch to the next output.

The *format*, *splitter_port*, *resize*, and *options* parameters are the same as in `start_recording()`, but *format* defaults to `'h264'`. The format is **not** derived from the filenames in *outputs* by this method.

For example, to record 3 consecutive 10-second video clips, writing the output to a series of H.264 files named `clip01.h264`, `clip02.h264`, and `clip03.h264` one could use the following:

```
import picamera
with picamera.PiCamera() as camera:
    for filename in camera.record_sequence([
        'clip01.h264',
        'clip02.h264',
        'clip03.h264']):
        print('Recording to %s' % filename)
        camera.wait_recording(10)
```

Alternatively, a more flexible method of writing the previous example (which is easier to expand to a large number of output files) is by using a generator expression as the input sequence:

```
import picamera
with picamera.PiCamera() as camera:
    for filename in camera.record_sequence(
        'clip%02d.h264' % i for i in range(3)):
        print('Recording to %s' % filename)
        camera.wait_recording(10)
```

More advanced techniques are also possible by utilising infinite sequences, such as those generated by `itertools.cycle()`. In the following example, recording is switched between two in-memory streams. Whilst one stream is recording, the other is being analysed. The script only stops recording when a video recording meets some criteria defined by the `process` function:

```

import io
import itertools
import picamera
with picamera.PiCamera() as camera:
    analyse = None
    for stream in camera.record_sequence(
        itertools.cycle((io.BytesIO(), io.BytesIO()))):
        if analyse is not None:
            if process(analyse):
                break
            analyse.seek(0)
            analyse.truncate()
        camera.wait_recording(5)
        analyse = stream

```

New in version 1.3.

remove_overlay (*overlay*)

Removes a static overlay from the preview output.

This method removes an overlay which was previously created by `add_overlay()`. The *overlay* parameter specifies the `PiRenderer` instance that was returned by `add_overlay()`.

split_recording (*output*, *splitter_port=1*, ***options*)

Continue the recording in the specified output; close existing output.

When called, the video encoder will wait for the next appropriate split point (an inline SPS header), then will cease writing to the current output (and close it, if it was specified as a filename), and continue writing to the newly specified *output*.

If *output* is a string, it will be treated as a filename for a new file which the video will be written to. Otherwise, *output* is assumed to be a file-like object and the video data is appended to it (the implementation only assumes the object has a `write()` method - no other methods will be called).

The *motion_output* parameter can be used to redirect the output of the motion vector data in the same fashion as *output*. If *motion_output* is `None` (the default) then motion vector data will not be redirected and will continue being written to the output specified by the *motion_output* parameter given to `start_recording()`. Alternatively, if you only wish to redirect motion vector data, you can set *output* to `None` and given a new value for *motion_output*.

The *splitter_port* parameter specifies which port of the video splitter the encoder you wish to change outputs is attached to. This defaults to 1 and most users will have no need to specify anything different. Valid values are between 0 and 3 inclusive.

Note that unlike `start_recording()`, you cannot specify format or other options as these cannot be changed in the middle of recording. Only the new *output* (and *motion_output*) can be specified. Furthermore, the format of the recording is currently limited to H264, and *inline_headers* must be `True` when `start_recording()` is called (this is the default).

Changed in version 1.3: The *splitter_port* parameter was added

Changed in version 1.5: The *motion_output* parameter was added

start_preview (***options*)

Displays the preview overlay.

This method starts a camera preview as an overlay on the Pi's primary display (HDMI or composite). A `PiRenderer` instance (more specifically, a `PiPreviewRenderer`) is constructed with the keyword arguments captured in *options*, and is returned from the method (this instance is also accessible from the `preview` attribute for as long as the renderer remains active). By default, the renderer will be opaque and fullscreen.

This means the default preview overrides whatever is currently visible on the display. More specifically, the preview does not rely on a graphical environment like X-Windows (it can run quite happily from a TTY console); it is simply an overlay on the Pi's video output. To stop the preview and reveal

the display again, call `stop_preview()`. The preview can be started and stopped multiple times during the lifetime of the `PiCamera` object.

All other camera properties can be modified “live” while the preview is running (e.g. `brightness`).

Note: Because the default preview typically obscures the screen, ensure you have a means of stopping a preview before starting one. If the preview obscures your interactive console you won’t be able to Alt+Tab back to it as the preview isn’t in a window. If you are in an interactive Python session, simply pressing Ctrl+D usually suffices to terminate the environment, including the camera and its associated preview.

start_recording (*output, format=None, resize=None, splitter_port=1, **options*)

Start recording video from the camera, storing it in *output*.

If *output* is a string, it will be treated as a filename for a new file which the video will be written to. Otherwise, *output* is assumed to be a file-like object and the video data is appended to it (the implementation only assumes the object has a `write()` method - no other methods will be called).

If *format* is `None` (the default), the method will attempt to guess the required video format from the extension of *output* (if it’s a string), or from the *name* attribute of *output* (if it has one). In the case that the format cannot be determined, a `PiCameraValueError` will be raised.

If *format* is not `None`, it must be a string specifying the format that you want the image written to. The format can be a MIME-type or one of the following strings:

- `'h264'` - Write an H.264 video stream
- `'mjpeg'` - Write an M-JPEG video stream
- `'yuv'` - Write the raw video data to a file in YUV420 format
- `'rgb'` - Write the raw video data to a file in 24-bit RGB format
- `'rgba'` - Write the raw video data to a file in 32-bit RGBA format
- `'bgr'` - Write the raw video data to a file in 24-bit BGR format
- `'bgra'` - Write the raw video data to a file in 32-bit BGRA format

If *resize* is not `None` (the default), it must be a two-element tuple specifying the width and height that the video recording should be resized to. This is particularly useful for recording video using the full area of the camera sensor (which is not possible without down-sizing the output).

The *splitter_port* parameter specifies the port of the built-in splitter that the video encoder will be attached to. This defaults to 1 and most users will have no need to specify anything different. If you wish to record multiple (presumably resized) streams simultaneously, specify a value between 0 and 3 inclusive for this parameter, ensuring that you do not specify a port that is currently in use.

Certain formats accept additional options which can be specified as keyword arguments. The `'h264'` format accepts the following additional options:

- *profile* - The H.264 profile to use for encoding. Defaults to `'high'`, but can be one of `'baseline'`, `'main'`, `'high'`, or `'constrained'`.
- *intra_period* - The key frame rate (the rate at which I-frames are inserted in the output). Defaults to `None`, but can be any 32-bit integer value representing the number of frames between successive I-frames. The special value 0 causes the encoder to produce a single initial I-frame, and then only P-frames subsequently. Note that `split_recording()` will fail in this mode.
- *inline_headers* - When `True`, specifies that the encoder should output SPS/PPS headers within the stream to ensure GOPs (groups of pictures) are self describing. This is important for streaming applications where the client may wish to seek within the stream, and enables the use of `split_recording()`. Defaults to `True` if not specified.
- *sei* - When `True`, specifies the encoder should include “Supplemental Enhancement Information” within the output stream. Defaults to `False` if not specified.

- *motion_output* - Indicates the output destination for motion vector estimation data. When `None` (the default), motion data is not output. If set to a string, it is assumed to be a filename which should be opened for motion data to be written to. Any other value is assumed to be a file-like object which motion vector is to be written to (the object must have a `write` method).

All formats accept the following additional options:

- *bitrate* - The bitrate at which video will be encoded. Defaults to 17000000 (17Mbps) if not specified. The maximum value is 25000000 (25Mbps). Bitrate 0 indicates the encoder should not use bitrate control (the encoder is limited by the quality only).
- *quality* - Specifies the quality that the encoder should attempt to maintain. For the 'h264' format, use values between 10 and 40 where 10 is extremely high quality, and 40 is extremely low (20-25 is usually a reasonable range for H.264 encoding). For the `mjpeg` format, use JPEG quality values between 1 and 100 (where higher values are higher quality). Quality 0 is special and seems to be a “reasonable quality” default.
- *quantization* - Deprecated alias for *quality*.

Changed in version 1.0: The *resize* parameter was added, and 'mjpeg' was added as a recording format

Changed in version 1.3: The *splitter_port* parameter was added

Changed in version 1.5: The *quantization* parameter was deprecated in favor of *quality*, and the *motion_output* parameter was added.

stop_preview()

Hides the preview overlay.

If `start_preview()` has previously been called, this method shuts down the preview display which generally results in the underlying display becoming visible again. If a preview is not currently running, no exception is raised - the method will simply do nothing.

stop_recording(*splitter_port=1*)

Stop recording video from the camera.

After calling this method the video encoder will be shut down and output will stop being written to the file-like object specified with `start_recording()`. If an error occurred during recording and `wait_recording()` has not been called since the error then this method will raise the exception.

The *splitter_port* parameter specifies which port of the video splitter the encoder you wish to stop is attached to. This defaults to 1 and most users will have no need to specify anything different. Valid values are between 0 and 3 inclusive.

Changed in version 1.3: The *splitter_port* parameter was added

wait_recording(*timeout=0, splitter_port=1*)

Wait on the video encoder for timeout seconds.

It is recommended that this method is called while recording to check for exceptions. If an error occurs during recording (for example out of disk space), an exception will only be raised when the `wait_recording()` or `stop_recording()` methods are called.

If *timeout* is 0 (the default) the function will immediately return (or raise an exception if an error has occurred).

The *splitter_port* parameter specifies which port of the video splitter the encoder you wish to wait on is attached to. This defaults to 1 and most users will have no need to specify anything different. Valid values are between 0 and 3 inclusive.

Changed in version 1.3: The *splitter_port* parameter was added

ISO

Retrieves or sets the apparent ISO setting of the camera.

Deprecated since version 1.8: Please use the `iso` attribute instead.

`_still_encoding`

Configures the encoding of the camera's still port.

This attribute controls the encoding of the camera's still port (see *Under the Hood* for more information). It is intended for internal use, but may be useful to developers wishing to implement *custom encoders*.

`analog_gain`

Retrieves the current analog gain of the camera.

When queried, this property returns the analog gain currently being used by the camera. The value represents the analog gain of the sensor prior to digital conversion. The value is returned as a `Fraction` instance.

New in version 1.6.

`annotate_background`

Controls whether a black background is drawn behind the annotation.

The `annotate_bg` attribute is a `bool` indicating whether or not a black background will be drawn behind the `annotation_text`. The background will appear in all output including image captures and video recording. The default is `False`.

`annotate_frame_num`

Controls whether the current frame number is drawn as an annotation.

The `annotate_frame_num` attribute is a `bool` indicating whether or not the current frame number is rendered as an annotation, similar to `annotate_text`. The default is `False`.

`annotate_text`

Retrieves or sets a text annotation for all output.

When queried, the `annotate_text` property returns the current annotation (if no annotation has been set, this is simply a blank string).

When set, the property immediately applies the annotation to the preview (if it is running) and to any future captures or video recording. Strings longer than 255 characters, or strings containing non-ASCII characters will raise a `PiCameraValueError`. The default value is `' '`.

`awb_gains`

Gets or sets the auto-white-balance gains of the camera.

When queried, this attribute returns a tuple of values representing the (*red*, *blue*) balance of the camera. The *red* and *blue* values are returned `Fraction` instances. The values will be between 0.0 and 8.0.

When set, this attribute adjusts the camera's auto-white-balance gains. The property can be specified as a single value in which case both red and blue gains will be adjusted equally, or as a (*red*, *blue*) tuple. Values can be specified as an `int`, `float` or `Fraction` and each gain must be between 0.0 and 8.0. Typical values for the gains are between 0.9 and 1.9. The property can be set while recordings or previews are in progress.

Note: This attribute only has an effect when `awb_mode` is set to `'off'`.

Changed in version 1.6: Prior to version 1.6, this attribute was write-only.

`awb_mode`

Retrieves or sets the auto-white-balance mode of the camera.

When queried, the `awb_mode` property returns a string representing the auto-white-balance setting of the camera. The possible values can be obtained from the `PiCamera.AWB_MODES` attribute, and are as follows:

- `'off'`
- `'auto'`
- `'sunlight'`

- 'cloudy'
- 'shade'
- 'tungsten'
- 'fluorescent'
- 'incandescent'
- 'flash'
- 'horizon'

When set, the property adjusts the camera's auto-white-balance mode. The property can be set while recordings or previews are in progress. The default value is 'auto'.

brightness

Retrieves or sets the brightness setting of the camera.

When queried, the `brightness` property returns the brightness level of the camera as an integer between 0 and 100. When set, the property adjusts the brightness of the camera. Brightness can be adjusted while previews or recordings are in progress. The default value is 50.

closed

Returns `True` if the `close()` method has been called.

color_effects

Retrieves or sets the current color effect applied by the camera.

When queried, the `color_effects` property either returns `None` which indicates that the camera is using normal color settings, or a `(u, v)` tuple where `u` and `v` are integer values between 0 and 255.

When set, the property changes the color effect applied by the camera. The property can be set while recordings or previews are in progress. For example, to make the image black and white set the value to `(128, 128)`. The default value is `None`.

contrast

Retrieves or sets the contrast setting of the camera.

When queried, the `contrast` property returns the contrast level of the camera as an integer between -100 and 100. When set, the property adjusts the contrast of the camera. Contrast can be adjusted while previews or recordings are in progress. The default value is 0.

crop

Retrieves or sets the zoom applied to the camera's input.

Deprecated since version 1.8: Please use the `zoom` attribute instead.

digital_gain

Retrieves the current digital gain of the camera.

When queried, this property returns the digital gain currently being used by the camera. The value represents the digital gain the camera applies after conversion of the sensor's analog output. The value is returned as a `Fraction` instance.

New in version 1.6.

drc_strength

Retrieves or sets the dynamic range compression strength of the camera.

When queried, the `drc_strength` property returns a string indicating the amount of `dynamic range compression` the camera applies to images.

When set, the attributes adjusts the strength of the dynamic range compression applied to the camera's output. Valid values are given in the list below:

- 'off'
- 'low'

- 'medium'
- 'high'

The default value is 'off'. All possible values for the attribute can be obtained from the `PiCamera.DRC_STRENGTHS` attribute.

New in version 1.6.

exif_tags

Holds a mapping of the Exif tags to apply to captured images.

Note: Please note that Exif tagging is only supported with the `jpeg` format.

By default several Exif tags are automatically applied to any images taken with the `capture()` method: `IFD0.Make` (which is set to `RaspberryPi`), `IFD0.Model` (which is set to `RP_OV5647`), and three timestamp tags: `IFD0.DateTime`, `EXIF.DateTimeOriginal`, and `EXIF.DateTimeDigitized` which are all set to the current date and time just before the picture is taken.

If you wish to set additional Exif tags, or override any of the aforementioned tags, simply add entries to the `exif_tags` map before calling `capture()`. For example:

```
camera.exif_tags['IFD0.Copyright'] = 'Copyright (c) 2013 Foo Industries'
```

The Exif standard mandates ASCII encoding for all textual values, hence strings containing non-ASCII characters will cause an encoding error to be raised when `capture()` is called. If you wish to set binary values, use a `bytes()` value:

```
camera.exif_tags['EXIF.UserComment'] = b'Something containing\x00NULL characters'
```

<p>Warning: Binary Exif values are currently ignored; this appears to be a libmmal or firmware bug.</p>
--

You may also specify datetime values, integer, or float values, all of which will be converted to appropriate ASCII strings (datetime values are formatted as `YYYY:MM:DD HH:MM:SS` in accordance with the Exif standard).

The currently supported Exif tags are:

Group	Tags
IFD0, IFD1	ImageWidth, ImageLength, BitsPerSample, Compression, PhotometricInterpretation, ImageDescription, Make, Model, StripOffsets, Orientation, SamplesPerPixel, RowsPerString, StripByteCounts, Xresolution, Yresolution, PlanarConfiguration, ResolutionUnit, TransferFunction, Software, DateTime, Artist, WhitePoint, PrimaryChromaticities, JPEGInterchangeFormat, JPEGInterchangeFormatLength, YcbCrCoefficients, YcbCrSubSampling, YcbCrPositioning, ReferenceBlackWhite, Copyright
EXIF	ExposureTime, FNumber, ExposureProgram, SpectralSensitivity, ISOSpeedRatings, OECF, ExifVersion, DateTimeOriginal, DateTimeDigitized, ComponentsConfiguration, CompressedBitsPerPixel, ShutterSpeedValue, ApertureValue, BrightnessValue, ExposureBiasValue, MaxApertureValue, SubjectDistance, MeteringMode, LightSource, Flash, FocalLength, SubjectArea, MakerNote, UserComment, SubSecTime, SubSecTimeOriginal, SubSecTimeDigitized, FlashpixVersion, ColorSpace, PixelXDimension, PixelYDimension, RelatedSoundFile, FlashEnergy, SpacialFrequencyResponse, FocalPlaneXResolution, FocalPlaneYResolution, FocalPlaneResolutionUnit, SubjectLocation, ExposureIndex, SensingMethod, FileSource, SceneType, CFAPattern, CustomRendered, ExposureMode, WhiteBalance, DigitalZoomRatio, FocalLengthIn35mmFilm, SceneCaptureType, GainControl, Contrast, Saturation, Sharpness, DeviceSettingDescription, SubjectDistanceRange, ImageUniqueID
GPS	GPSVersionID, GPSLatitudeRef, GPSLatitude, GPSLongitudeRef, GPSLongitude, GPSAltitudeRef, GPSAltitude, GPSTimeStamp, GPSSatellites, GPSStatus, GPSMeasureMode, GPSDOP, GPSSpeedRef, GPSSpeed, GPSTrackRef, GPSTrack, GPSImgDirectionRef, GPSImgDirection, GPSMapDatum, GPSDestLatitudeRef, GPSDestLatitude, GPSDestLongitudeRef, GPSDestLongitude, GPSDestBearingRef, GPSDestBearing, GPSDestDistanceRef, GPSDestDistance, GPSProcessingMethod, GPSAreaInformation, GPSDateStamp, GPSDifferential
EINT	InteroperabilityIndex, InteroperabilityVersion, RelatedImageFileFormat, RelatedImageWidth, RelatedImageLength

exposure_compensation

Retrieves or sets the exposure compensation level of the camera.

When queried, the `exposure_compensation` property returns an integer value between -25 and 25 indicating the exposure level of the camera. Larger values result in brighter images.

When set, the property adjusts the camera's exposure compensation level. Each increment represents 1/6th of a stop. Hence setting the attribute to 6 increases exposure by 1 stop. The property can be set while recordings or previews are in progress. The default value is 0.

exposure_mode

Retrieves or sets the exposure mode of the camera.

When queried, the `exposure_mode` property returns a string representing the exposure setting of the camera. The possible values can be obtained from the `PiCamera.EXPOSURE_MODES` attribute, and are as follows:

- 'off'
- 'auto'
- 'night'
- 'nightpreview'
- 'backlight'
- 'spotlight'
- 'sports'
- 'snow'
- 'beach'

- 'verylong'
- 'fixedfps'
- 'antishake'
- 'fireworks'

When set, the property adjusts the camera's exposure mode. The property can be set while recordings or previews are in progress. The default value is 'auto'.

exposure_speed

Retrieves the current shutter speed of the camera.

When queried, this property returns the shutter speed currently being used by the camera. If you have set `shutter_speed` to a non-zero value, then `exposure_speed` and `shutter_speed` should be equal. However, if `shutter_speed` is set to 0 (auto), then you can read the actual shutter speed being used from this attribute. The value is returned as an integer representing a number of microseconds. This is a read-only property.

New in version 1.6.

frame

Retrieves information about the current frame recorded from the camera.

When video recording is active (after a call to `start_recording()`), this attribute will return a `PiVideoFrame` tuple containing information about the current frame that the camera is recording.

If multiple video recordings are currently in progress (after multiple calls to `start_recording()` with different values for the `splitter_port` parameter), which encoder's frame information is returned is arbitrary. If you require information from a specific encoder, you will need to extract it from `_encoders` explicitly.

Querying this property when the camera is not recording will result in an exception.

Note: There is a small window of time when querying this attribute will return `None` after calling `start_recording()`. If this attribute returns `None`, this means that the video encoder has been initialized, but the camera has not yet returned any frames.

framerate

Retrieves or sets the framerate at which video-port based image captures, video recordings, and previews will run.

When queried, the `framerate` property returns the rate at which the camera's video and preview ports will operate as a `Fraction` instance which can be easily converted to an `int` or `float`.

Note: For backwards compatibility, a derivative of the `Fraction` class is actually used which permits the value to be treated as a tuple of (`numerator`, `denominator`).

Setting and retrieving framerate as a (`numerator`, `denominator`) tuple is deprecated and will be removed in 2.0. Please use a `Fraction` instance instead (which is just as accurate and also permits direct use with math operators).

When set, the property reconfigures the camera so that the next call to recording and previewing methods will use the new framerate. The framerate can be specified as an `int`, `float`, `Fraction`, or a (`numerator`, `denominator`) tuple. The camera must not be closed, and no recording must be active when the property is set.

Note: This attribute, in combination with `resolution`, determines the mode that the camera operates in. The actual sensor framerate and resolution used by the camera is influenced, but not directly set, by this property. See *Camera Modes* for more information.

hflip

Retrieves or sets whether the camera's output is horizontally flipped.

When queried, the `hflip` property returns a boolean indicating whether or not the camera's output is horizontally flipped. The property can be set while recordings or previews are in progress. The default value is `False`.

image_denoise

Retrieves or sets whether denoise will be applied to image captures.

When queried, the `image_denoise` property returns a boolean value indicating whether or not the camera software will apply a denoise algorithm to image captures.

When set, the property activates or deactivates the denoise algorithm for image captures. The property can be set while recordings or previews are in progress. The default value is `True`.

New in version 1.7.

image_effect

Retrieves or sets the current image effect applied by the camera.

When queried, the `image_effect` property returns a string representing the effect the camera will apply to captured video. The possible values can be obtained from the `PiCamera.IMAGE_EFFECTS` attribute, and are as follows:

- `'none'`
- `'negative'`
- `'solarize'`
- `'sketch'`
- `'denoise'`
- `'emboss'`
- `'oilpaint'`
- `'hatch'`
- `'gpen'`
- `'pastel'`
- `'watercolor'`
- `'film'`
- `'blur'`
- `'saturation'`
- `'colorswap'`
- `'washedout'`
- `'posterise'`
- `'colorpoint'`
- `'colorbalance'`
- `'cartoon'`
- `'deinterlace1'`
- `'deinterlace2'`

When set, the property changes the effect applied by the camera. The property can be set while recordings or previews are in progress, but only certain effects work while recording video (notably `'negative'` and `'solarize'`). The default value is `'none'`.

image_effect_params

Retrieves or sets the parameters for the current `effect`.

When queried, the `image_effect_params` property either returns `None` (for effects which have no configurable parameters, or if no parameters have been configured), or a tuple of numeric values up to six elements long.

When set, the property changes the parameters of the current `effect` as a sequence of numbers, or a single number. Attempting to set parameters on an effect which does not support parameters, or providing an incompatible set of parameters for an effect will raise a `PiCameraValueError` exception.

The effects which have parameters, and what combinations those parameters can take is as follows:

Effect	Parameters	Description
'solarize'	<code>yuv, x0, y1, y2, y3, x0, y0, y1, y2</code>	<code>yuv</code> controls whether data is processed as RGB (0) or YUV(1). Input values from 0 to <code>x0 - 1</code> are remapped linearly onto the range 0 to <code>y0</code> . Values from <code>x0</code> to 255 are remapped linearly onto the range <code>y1</code> to <code>y2</code> . Same as above, but <code>yuv</code> defaults to 0 (process as RGB).
'colorpoint'	<code>yuv, quadrant</code>	Same as above, but <code>x0, y0, y1, y2</code> default to 128, 128, 128, 0 respectively. <code>quadrant</code> specifies which quadrant of the U/V space to retain chroma from: 0=green, 1=red/yellow, 2=blue, 3=purple. There is no default; this effect does nothing until parameters are set.
'colorbalance'	<code>lens, r, g, b, u, v</code>	<code>lens</code> specifies the lens shading strength (0.0 to 256.0, where 0.0 indicates lens shading has no effect). <code>r, g, b</code> are multipliers for their respective color channels (0.0 to 256.0). <code>u</code> and <code>v</code> are offsets added to the U/V plane (0 to 255).
'colorbalance'	<code>lens, r, g, b</code>	Same as above but <code>u</code> are defaulted to 0.
'colorbalance'	<code>lens, r, b</code>	Same as above but <code>g</code> also defaults to 1.0.
'colorswap'	<code>dir</code>	If <code>dir</code> is 0, swap RGB to BGR. If <code>dir</code> is 1, swap RGB to BRG.
'posterize'	<code>steps</code>	Control the quantization steps for the image. Valid values are 2 to 32, and the default is 4.
'blur'	<code>size</code>	Specifies the size of the kernel. Valid values are 1 or 2.
'film'	<code>strength, u, v</code>	<code>strength</code> specifies the strength of effect. <code>u</code> and <code>v</code> are offsets added to the U/V plane (0 to 255).
'watercolor'	<code>u, v</code>	<code>u</code> and <code>v</code> specify offsets to add to the U/V plane (0 to 255). No parameters indicates no U/V effect.

New in version 1.8.

iso

Retrieves or sets the apparent ISO setting of the camera.

When queried, the `iso` property returns the ISO setting of the camera, a value which represents the *sensitivity of the camera to light*. Lower values (e.g. 100) imply less sensitivity than higher values (e.g. 400 or 800). Lower sensitivities tend to produce less “noisy” (smoother) images, but operate poorly in low light conditions.

When set, the property adjusts the sensitivity of the camera. Valid values are between 0 (auto) and 1600. The actual value used when `iso` is explicitly set will be one of the following values (whichever is closest): 100, 200, 320, 400, 500, 640, 800.

The attribute can be adjusted while previews or recordings are in progress. The default value is 0 which means automatically determine a value according to image-taking conditions.

Note: You can query the `analog_gain` and `digital_gain` attributes to determine the actual gains being used by the camera. If both are 1.0 this equates to ISO 100. Please note that this capability requires an up to date firmware (#692 or later).

Note: With iso settings other than 0 (auto), the `exposure_mode` property becomes non-functional.

Note: Some users on the Pi camera forum have noted that higher ISO values than 800 (specifically up to 1600) can be achieved in certain conditions with `exposure_mode` set to `'sports'` and `iso` set to 0. It doesn't appear to be possible to manually request an ISO setting higher than 800, but the picamera library will permit settings up to 1600 in case the underlying firmware permits such settings in particular circumstances.

led

Sets the state of the camera's LED via GPIO.

If a GPIO library is available (only RPi.GPIO is currently supported), and if the python process has the necessary privileges (typically this means running as root via sudo), this property can be used to set the state of the camera's LED as a boolean value (`True` is on, `False` is off).

Note: This is a write-only property. While it can be used to control the camera's LED, you cannot query the state of the camera's LED using this property.

meter_mode

Retrieves or sets the metering mode of the camera.

When queried, the `meter_mode` property returns the method by which the camera determines the exposure as one of the following strings:

- `'average'`
- `'spot'`
- `'backlit'`
- `'matrix'`

When set, the property adjusts the camera's metering mode. All modes set up two regions: a center region, and an outer region. The major difference between each mode is the size of the center region. The `'backlit'` mode has the largest central region (30% of the width), while `'spot'` has the smallest (10% of the width).

The property can be set while recordings or previews are in progress. The default value is `'average'`. All possible values for the attribute can be obtained from the `PiCamera.METER_MODES` attribute.

overlays

Retrieves all active `PiRenderer` overlays.

If no overlays are current active, `overlays` will return an empty iterable. Otherwise, it will return an iterable of `PiRenderer` instances which are currently acting as overlays. Note that the preview renderer is an exception to this: it is *not* included as an overlay despite being derived from `PiRenderer`.

preview

Retrieves the `PiRenderer` displaying the camera preview.

If no preview is currently active, `preview` will return `None`. Otherwise, it will return the instance of `PiRenderer` which is currently connected to the camera's preview port for rendering what the camera sees. You can use the attributes of the `PiRenderer` class to configure the appearance of the preview. For example, to make the preview semi-transparent:

```
import picamera

with picamera.PiCamera() as camera:
    camera.start_preview()
    camera.preview.alpha = 128
```

New in version 1.8.

preview_alpha

Retrieves or sets the opacity of the preview window.

Deprecated since version 1.8: Please use the `alpha` attribute of the `preview` object instead.

preview_fullscreen

Retrieves or sets full-screen for the preview window.

Deprecated since version 1.8: Please use the `fullscreen` attribute of the `preview` object instead.

preview_layer

Retrieves or sets the layer of the preview window.

Deprecated since version 1.8: Please use the `layer` attribute of the `preview` object instead.

preview_window

Retrieves or sets the size of the preview window.

Deprecated since version 1.8: Please use the `window` attribute of the `preview` object instead.

previewing

Returns `True` if the `start_preview()` method has been called, and no `stop_preview()` call has been made yet.

Deprecated since version 1.8: Test whether `preview` is `None` instead.

raw_format

Retrieves or sets the raw format of the camera's ports.

Deprecated since version 1.0: Please use `'yuv'` or `'rgb'` directly as a format in the various capture methods instead.

recording

Returns `True` if the `start_recording()` method has been called, and no `stop_recording()` call has been made yet.

resolution

Retrieves or sets the resolution at which image captures, video recordings, and previews will be captured.

When queried, the `resolution` property returns the resolution at which the camera will operate as a tuple of (`width`, `height`) measured in pixels. This is the resolution that the `capture()` method will produce images at, and the resolution that `start_recording()` will produce videos at.

When set, the property reconfigures the camera so that the next call to these methods will use the new resolution. The resolution must be specified as a (`width`, `height`) tuple, the camera must not be closed, and no recording must be active when the property is set.

The property defaults to the Pi's currently configured display resolution unless the display has been disabled (with `tvservice -o`) in which case it defaults to 1280x720 (720p).

Note: This attribute, in combination with `framerate`, determines the mode that the camera operates in. The actual sensor framerate and resolution used by the camera is influenced, but not directly set, by this property. See *Camera Modes* for more information.

rotation

Retrieves or sets the current rotation of the camera's image.

When queried, the `rotation` property returns the rotation applied to the image. Valid values are 0, 90, 180, and 270.

When set, the property changes the rotation applied to the camera's input. The property can be set while recordings or previews are in progress. The default value is 0.

saturation

Retrieves or sets the saturation setting of the camera.

When queried, the `saturation` property returns the color saturation of the camera as an integer between -100 and 100. When set, the property adjusts the saturation of the camera. Saturation can be adjusted while previews or recordings are in progress. The default value is 0.

sharpness

Retrieves or sets the sharpness setting of the camera.

When queried, the `sharpness` property returns the sharpness level of the camera (a measure of the amount of post-processing to reduce or increase image sharpness) as an integer between -100 and 100. When set, the property adjusts the sharpness of the camera. Sharpness can be adjusted while previews or recordings are in progress. The default value is 0.

shutter_speed

Retrieves or sets the shutter speed of the camera in microseconds.

When queried, the `shutter_speed` property returns the shutter speed of the camera in microseconds, or 0 which indicates that the speed will be automatically determined by the auto-exposure algorithm. Faster shutter times naturally require greater amounts of illumination and vice versa.

When set, the property adjusts the shutter speed of the camera, which most obviously affects the illumination of subsequently captured images. Shutter speed can be adjusted while previews or recordings are running. The default value is 0 (auto).

Note: You can query the `exposure_speed` attribute to determine the actual shutter speed being used when this attribute is set to 0. Please note that this capability requires an up to date firmware (#692 or later).

Note: In later firmwares, this attribute is limited by the value of the `framerate` attribute. For example, if framerate is set to 30fps, the shutter speed cannot be slower than 33,333 μ s (1/fps).

vflip

Retrieves or sets whether the camera's output is vertically flipped.

When queried, the `vflip` property returns a boolean indicating whether or not the camera's output is vertically flipped. The property can be set while recordings or previews are in progress. The default value is `False`.

video_denoise

Retrieves or sets whether denoise will be applied to video recordings.

When queried, the `video_denoise` property returns a boolean value indicating whether or not the camera software will apply a denoise algorithm to video recordings.

When set, the property activates or deactivates the denoise algorithm for video recordings. The property can be set while recordings or previews are in progress. The default value is `True`.

New in version 1.7.

video_stabilization

Retrieves or sets the video stabilization mode of the camera.

When queried, the `video_stabilization` property returns a boolean value indicating whether or not the camera attempts to compensate for motion.

When set, the property activates or deactivates video stabilization. The property can be set while recordings or previews are in progress. The default value is `False`.

Note: The built-in video stabilization only accounts for `vertical and horizontal motion`, not rotation.

zoom

Retrieves or sets the zoom applied to the camera's input.

When queried, the `zoom` property returns a `(x, y, w, h)` tuple of floating point values ranging from 0.0 to 1.0, indicating the proportion of the image to include in the output (this is also known as the “Region of Interest” or ROI). The default value is `(0.0, 0.0, 1.0, 1.0)` which indicates that everything should be included. The property can be set while recordings or previews are in progress.

2.9.2 PiCameraCircularIO

```
class picamera.PiCameraCircularIO(camera, size=None, seconds=None, bitrate=17000000,
                                  splitter_port=1)
```

A derivative of `CircularIO` which tracks camera frames.

`PiCameraCircularIO` provides an in-memory stream based on a ring buffer. It is a specialization of `CircularIO` which associates video frame meta-data with the recorded stream, accessible from the `frames` property.

Warning: The class makes a couple of assumptions which will cause the frame meta-data tracking to break if they are not adhered to:

- the stream is only ever appended to - no writes ever start from the middle of the stream
- the stream is never truncated (from the right; being ring buffer based, left truncation will occur automatically)

The `camera` parameter specifies the `PiCamera` instance that will be recording video to the stream. If specified, the `size` parameter determines the maximum size of the stream in bytes. If `size` is not specified (or `None`), then `seconds` must be specified instead. This provides the maximum length of the stream in seconds, assuming a data rate in bits-per-second given by the `bitrate` parameter (which defaults to 17000000, or 17Mbps, which is also the default bitrate used for video recording by `PiCamera`). You cannot specify both `size` and `seconds`.

The `splitter_port` parameter specifies the port of the built-in splitter that the video encoder will be attached to. This defaults to 1 and most users will have no need to specify anything different. If you do specify something else, ensure it is equal to the `splitter_port` parameter of the corresponding call to `PiCamera.start_recording()`. For example:

```
import picamera

with picamera.PiCamera() as camera:
    with picamera.PiCameraCircularIO(camera, splitter_port=2) as stream:
        camera.start_recording(stream, format='h264', splitter_port=2)
        camera.wait_recording(10, splitter_port=2)
        camera.stop_recording(splitter_port=2)
```

frames

Returns an iterator over the frame meta-data.

As the camera records video to the stream, the class captures the meta-data associated with each frame (in the form of a `PiVideoFrame` tuple), discarding meta-data for frames which are no longer fully stored within the underlying ring buffer. You can use the frame meta-data to locate, for example, the first keyframe present in the stream in order to determine an appropriate range to extract.

2.9.3 CircularIO

```
class picamera.CircularIO(size)
```

A thread-safe stream which uses a ring buffer for storage.

`CircularIO` provides an in-memory stream similar to the `io.BytesIO` class. However, unlike `BytesIO` its underlying storage is a `ring buffer` with a fixed maximum size. Once the maximum size is reached, writing effectively loops round to the beginning to the ring and starts overwriting the oldest content.

The `size` parameter specifies the maximum size of the stream in bytes. The `read()`, `tell()`, and `seek()` methods all operate equivalently to those in `io.BytesIO` whilst `write()` only differs in the wrapping

behaviour described above. A `read1()` method is also provided for efficient reading of the underlying ring buffer in write-sized chunks (or less).

A re-entrant threading lock guards all operations, and is accessible for external use via the `lock` attribute.

The performance of the class is geared toward faster writing than reading on the assumption that writing will be the common operation and reading the rare operation (a reasonable assumption for the camera use-case, but not necessarily for more general usage).

getvalue()

Return `bytes` containing the entire contents of the buffer.

read(n=-1)

Read up to `n` bytes from the stream and return them. As a convenience, if `n` is unspecified or `-1`, `readall()` is called. Fewer than `n` bytes may be returned if there are fewer than `n` bytes from the current stream position to the end of the stream.

If 0 bytes are returned, and `n` was not 0, this indicates end of the stream.

read1(n=-1)

Read up to `n` bytes from the stream using only a single call to the underlying object.

In the case of `CircularIO` this roughly corresponds to returning the content from the current position up to the end of the write that added that content to the stream (assuming no subsequent writes overwrote the content). `read1()` is particularly useful for efficient copying of the stream's content.

readable()

Returns `True`, indicating that the stream supports `read()`.

readall()

Read and return all bytes from the stream until EOF, using multiple calls to the stream if necessary.

seek(offset, whence=0)

Change the stream position to the given byte `offset`. `offset` is interpreted relative to the position indicated by `whence`. Values for `whence` are:

- `SEEK_SET` or 0 – start of the stream (the default); `offset` should be zero or positive
- `SEEK_CUR` or 1 – current stream position; `offset` may be negative
- `SEEK_END` or 2 – end of the stream; `offset` is usually negative

Return the new absolute position.

seekable()

Returns `True`, indicating the stream supports `seek()` and `tell()`.

tell()

Return the current stream position.

truncate(size=None)

Resize the stream to the given `size` in bytes (or the current position if `size` is not specified). This resizing can extend or reduce the current stream size. In case of extension, the contents of the new file area will be NUL (`\x00`) bytes. The new stream size is returned.

The current stream position isn't changed unless the resizing is expanding the stream, in which case it may be set to the maximum stream size if the expansion causes the ring buffer to loop around.

writable()

Returns `True`, indicating that the stream supports `write()`.

write(b)

Write the given bytes or bytearray object, `b`, to the underlying stream and return the number of bytes written.

lock

A re-entrant threading lock which is used to guard all operations.

size

Return the maximum size of the buffer in bytes.

2.9.4 PiVideoFrameType

class `picamera.PiVideoFrameType`

This class simply defines constants used to represent the type of a frame in `PiVideoFrame.frame_type`. Effectively it is a namespace for an enum.

frame

Indicates a predicted frame (P-frame). This is the most common frame type.

key_frame

Indicates an intra-frame (I-frame) also known as a key frame.

sps_header

Indicates an inline SPS/PPS header (rather than picture data) which is typically used as a split point.

motion_data

Indicates the frame is inline motion vector data, rather than picture data.

New in version 1.5.

2.9.5 PiVideoFrame

class `picamera.PiVideoFrame` (*index, frame_type, frame_size, video_size, split_size, timestamp*)

This class is a namedtuple derivative used to store information about a video frame. It is recommended that you access the information stored by this class by attribute name rather than position (for example: `frame.index` rather than `frame[0]`).

index

Returns the zero-based number of the frame. This is a monotonic counter that is simply incremented every time the camera returns a frame-end buffer. As a consequence, this attribute cannot be used to detect dropped frames. Nor does it necessarily represent actual frames; it will be incremented for SPS headers and motion data buffers too.

frame_type

Returns a constant indicating the kind of data that the frame contains (see `PiVideoFrameType`). Please note that certain frame types contain no image data at all.

frame_size

Returns the size in bytes of the current frame.

video_size

Returns the size in bytes of the entire video up to the current frame. Note that this is unlikely to match the size of the actual file/stream written so far. Firstly this is because the frame attribute is only updated when the encoder outputs the *end* of a frame, which will cause the reported size to be smaller than the actual amount written. Secondly this is because a stream may utilize buffering which will cause the actual amount written (e.g. to disk) to lag behind the value reported by this attribute.

split_size

Returns the size in bytes of the video recorded since the last call to either `start_recording()` or `split_recording()`. For the reasons explained above, this may differ from the size of the actual file/stream written so far.

timestamp

Returns the presentation timestamp (PTS) of the current frame as reported by the encoder. This is represented by the number of microseconds (millionths of a second) since video recording started. As the frame attribute is only updated when the encoder outputs the end of a frame, this value may lag behind the actual time since `start_recording()` was called.

Warning: Currently, the video encoder occasionally returns “time unknown” values in this field which picamera represents as `None`. If you are querying this property you will need to check the value is not `None` before using it.

Changed in version 1.5: Deprecated `header` and `keyframe` attributes and added the new `frame_type` attribute instead.

header

Contains a bool indicating whether the current frame is actually an SPS/PPS header. Typically it is best to split an H.264 stream so that it starts with an SPS/PPS header.

Deprecated since version 1.5: Please compare `frame_type` to `PiVideoFrameType.sps_header` instead.

keyframe

Returns a bool indicating whether the current frame is a keyframe (an intra-frame, or I-frame in MPEG parlance).

Deprecated since version 1.5: Please compare `frame_type` to `PiVideoFrameType.key_frame` instead.

position

Returns the zero-based position of the frame in the stream containing it.

2.9.6 PiEncoder

class `picamera.PiEncoder` (*parent, camera_port, input_port, format, resize, **options*)

Base implementation of an MMAL encoder for use by PiCamera.

The *parent* parameter specifies the `PiCamera` instance that has constructed the encoder. The *camera_port* parameter provides the MMAL camera port that the encoder should enable for capture (this will be the still or video port of the camera component). The *input_port* parameter specifies the MMAL port that the encoder should connect to its input. Sometimes this will be the same as the camera port, but if other components are present in the pipeline (e.g. a splitter), it may be different.

The *format* parameter specifies the format that the encoder should produce in its output. This is specified as a string and will be one of the following for image encoders:

- 'jpeg'
- 'png'
- 'gif'
- 'bmp'
- 'yuv'
- 'rgb'
- 'rgba'
- 'bgr'
- 'bgra'

And one of the following for video encoders:

- 'h264'
- 'mjpeg'

The *resize* parameter is either `None` (indicating no resizing should take place), or a (*width, height*) tuple specifying the resolution that the output of the encoder should be resized to.

Finally, the *options* parameter specifies additional keyword arguments that can be used to configure the encoder (e.g. bitrate for videos, or quality for images).

The class has a number of attributes:

camera_port

A pointer to the camera output port that needs to be activated and deactivated in order to start/stop capture. This is not necessarily the port that the encoder component's input port is connected to (for example, in the case of video-port based captures, this will be the camera video port behind the splitter).

encoder

A pointer to the MMAL encoder component, or None if no encoder component has been created (some encoder classes don't use an actual encoder component, for example `PiRawImageMixin`).

encoder_connection

A pointer to the MMAL connection linking the encoder's input port to the camera, splitter, or resizer output port (depending on configuration), if any.

event

A `threading.Event` instance used to synchronize operations (like start, stop, and split) between the control thread and the callback thread.

exception

If an exception occurs during the encoder callback, this attribute is used to store the exception until it can be re-raised in the control thread.

format

The image or video format that the encoder is expected to produce. This is equal to the value of the *format* parameter.

input_port

A pointer to the MMAL port that the encoder component's input port should be connected to.

output_port

A pointer to the MMAL port of the encoder's output. In the case no encoder component is created, this should be the camera/component output port responsible for producing data. In other words, this attribute **must** be set on initialization.

outputs

A mapping of key to (output, opened) tuples where output is a file-like object, and opened is a bool indicating whether or not we opened the output object (and thus whether we are responsible for eventually closing it).

outputs_lock

A `threading.Lock` instance used to protect access to `outputs`.

parent

The `PiCamera` instance that created this `PiEncoder` instance.

pool

A pointer to a pool of MMAL buffers.

resizer

A pointer to the MMAL resizer component, or None if no resizer component has been created.

resizer_connection

A pointer to the MMAL connection linking the resizer's input port to the camera or splitter's output port, if any.

_callback (*port*, *buf*)

The encoder's main callback function.

When the encoder is active, this method is periodically called in a background thread. The *port* parameter specifies the MMAL port providing the output (typically this is the encoder's output port, but in the case of unencoded captures may simply be a camera port), while the *buf* parameter is an MMAL buffer header pointer which can be used to obtain the data to write, along with meta-data about the current frame.

This method *must* release the MMAL buffer header before returning (failure to do so will cause a lockup), and should recycle buffers if expecting further data (the `_callback_recycle()` method can be called to perform the latter duty). Finally, this method must set `event` when the encoder has finished (and should set `exception` if an exception occurred during encoding).

Developers wishing to write a custom encoder class may find it simpler to override the `_callback_write()` method, rather than deal with these complexities.

`_callback_recycle` (*port, buf*)

Recycles the buffer on behalf of the encoder callback function.

This method is called by `_callback()` when there is a buffer to recycle (because further output is expected). It is unlikely descendent classes will have a need to override this method, but if they override the `_callback()` method they may wish to call it.

`_callback_write` (*buf, key=0*)

Writes output on behalf of the encoder callback function.

This method is called by `_callback()` to handle writing to an object in `outputs` identified by *key*. The *buf* parameter is an MMAL buffer header pointer which can be used to obtain the length of data available (`buf[0].length`), a pointer to the data (`buf[0].data`) which should typically be used with `ctypes.string_at()`, and meta-data about the contents of the buffer (`buf[0].flags`). The method is expected to return a boolean to indicate whether output is complete (`True`) or whether more data is expected (`False`).

The default implementation simply writes the contents of the buffer to the output identified by *key*, and returns `True` if the buffer flags indicate end of stream. Image encoders will typically override the return value to indicate `True` on end of frame (as they only wish to output a single image). Video encoders will typically override this method to determine where key-frames and SPS headers occur.

`_close_output` (*key=0*)

Closes the output associated with *key* in `outputs`.

Closes the output object associated with the specified *key*, and removes it from the `outputs` dictionary (if we didn't open the object then we attempt to flush it instead).

`_create_connections` ()

Creates all connections between MMAL components.

This method is called to connect the encoder and the optional resizer to the input port provided by the camera. It sets the `encoder_connection` and `resizer_connection` attributes as required.

`_create_encoder` ()

Creates and configures the MMAL encoder component.

This method only constructs the encoder; it does not connect it to the input port. The method sets the `encoder` attribute to the constructed encoder component, and the `output_port` attribute to the encoder's output port (or the previously constructed resizer's output port if one has been requested). Descendent classes extend this method to finalize encoder configuration.

Note: It should be noted that this method is called with the initializer's `option` keyword arguments. This base implementation expects no additional arguments, but descendent classes extend the parameter list to include options relevant to them.

`_create_pool` ()

Allocates a pool of MMAL buffers for the encoder.

This method is expected to construct an MMAL pool of buffers for the `output_port`, and store the result in the `pool` attribute.

`_create_resizer` (*width, height*)

Creates and configures an MMAL resizer component.

This is called when the initializer's `resize` parameter is something other than `None`. The *width* and *height* parameters are passed to the constructed resizer. Note that this method only constructs the re-

sizer - it does not connect it to the encoder. The method sets the `resizer` attribute to the constructed resizer component.

`_open_output` (*output*, *key=0*)

Opens *output* and associates it with *key* in `outputs`.

If *output* is a string, this method opens it as a filename and keeps track of the fact that the encoder was the one to open it (which implies that `_close_output()` should eventually close it). Otherwise, *output* is assumed to be a file-like object and is used verbatim. The opened output is added to the `outputs` dictionary with the specified *key*.

`close` ()

Finalizes the encoder and deallocates all structures.

This method is called by the camera prior to destroying the encoder (or more precisely, letting it go out of scope to permit the garbage collector to destroy it at some future time). The method destroys all components that the various create methods constructed and resets their attributes.

`start` (*output*)

Starts the encoder object writing to the specified output.

This method is called by the camera to start the encoder capturing data from the camera to the specified output. The *output* parameter is either a filename, or a file-like object (for image and video encoders), or an iterable of filenames or file-like objects (for multi-image encoders).

`stop` ()

Stops the encoder, regardless of whether it's finished.

This method is called by the camera to terminate the execution of the encoder. Typically, this is used with video to stop the recording, but can potentially be called in the middle of image capture to terminate the capture.

`wait` (*timeout=None*)

Waits for the encoder to finish (successfully or otherwise).

This method is called by the owning camera object to block execution until the encoder has completed its task. If the *timeout* parameter is `None`, the method will block indefinitely. Otherwise, the *timeout* parameter specifies the (potentially fractional) number of seconds to block for. If the encoder finishes successfully within the timeout, the method returns `True`. Otherwise, it returns `False`.

`active`

Returns `True` if the MMAL encoder exists and is enabled.

2.9.7 PiVideoEncoder

`class picamera.PiVideoEncoder` (*parent*, *camera_port*, *input_port*, *format*, *resize*, ***options*)

Encoder for video recording.

This derivative of `PiEncoder` configures itself for H.264 or MJPEG encoding. It also introduces a `split()` method which is used by `split_recording()` and `record_sequence()` to redirect future output to a new filename or object. Finally, it also extends `PiEncoder.start()` and `PiEncoder._callback_write()` to track video frame meta-data, and to permit recording motion data to a separate output object.

`_callback_write` (*buf*, *key=0*)

Extended to implement video frame meta-data tracking, and to handle splitting video recording to the next output when `split()` is called.

`_create_encoder` (*bitrate=17000000*, *intra_period=None*, *profile='high'*, *quantization=0*, *quality=0*, *inline_headers=True*, *sei=False*, *motion_output=None*)

Extends the base `_create_encoder()` implementation to configure the video encoder for H.264 or MJPEG output.

`split` (*output*, *motion_output=None*)

Called to switch the encoder's output.

This method is called by `split_recording()` and `record_sequence()` to switch the encoder's output object to the `output` parameter (which can be a filename or a file-like object, as with `start()`).

start (*output, motion_output=None*)

Extended to initialize video frame meta-data tracking.

2.9.8 PiImageEncoder

class `picamera.PiImageEncoder` (*parent, camera_port, input_port, format, resize, **options*)

Encoder for image capture.

This derivative of `PiEncoder` extends the `_create_encoder()` method to configure the encoder for a variety of encoded image outputs (JPEG, PNG, etc.).

_create_encoder (*quality=85, thumbnail=(64, 48, 35), bayer=False*)

Extends the base `_create_encoder()` implementation to configure the image encoder for JPEG, PNG, etc.

2.9.9 PiRawMixin

class `picamera.PiRawMixin` (*parent, camera_port, input_port, format, resize, **options*)

Mixin class for “raw” (unencoded) output.

This mixin class overrides the initializer of `PiEncoder`, along with `_create_resizer()` and `_create_encoder()` to configure the pipeline for unencoded output. Specifically, it disables the construction of an encoder, and sets the output port to the input port passed to the initializer, unless resizing is required (either for actual resizing, or for format conversion) in which case the resizer's output is used.

_callback_write (*buf, key=0*)

Overridden to strip alpha bytes when required.

_create_connections ()

Overridden to skip creating an encoder connection; only a resizer connection is required (if one has been configured).

_create_encoder ()

Overridden to skip creating an encoder. Instead, this class simply uses the resizer's port as the output port (if a resizer has been configured) or the specified input port otherwise.

_create_resizer (*width, height*)

Overridden to configure the resizer's output with the required encoding.

2.9.10 PiCookedVideoEncoder

class `picamera.PiCookedVideoEncoder` (*parent, camera_port, input_port, format, resize, **options*)

Video encoder for encoded recordings.

This class is a derivative of `PiVideoEncoder` and only exists to provide naming symmetry with the image encoder classes.

2.9.11 PiRawVideoEncoder

class `picamera.PiRawVideoEncoder` (*parent, camera_port, input_port, format, resize, **options*)

Video encoder for unencoded recordings.

This class is a derivative of `PiVideoEncoder` and the `PiRawMixin` class intended for use with `start_recording()` when it is called with an unencoded format.

Warning: This class creates an inheritance diamond. Take care to determine the MRO of super-class calls.

2.9.12 PiOneImageEncoder

class `picamera.PiOneImageEncoder` (*parent, camera_port, input_port, format, resize, **options*)
Encoder for single image capture.

This class simply extends `_callback_write()` to terminate capture at frame end (i.e. after a single frame has been received).

2.9.13 PiMultiImageEncoder

class `picamera.PiMultiImageEncoder` (*parent, camera_port, input_port, format, resize, **options*)
Encoder for multiple image capture.

This class extends `PiImageEncoder` to handle an iterable of outputs instead of a single output. The `_callback_write()` method is extended to terminate capture when the iterable is exhausted, while `PiEncoder._open_output()` is overridden to begin iteration and rely on the new `_next_output()` method to advance output to the next item in the iterable.

`_next_output` (*key=0*)

This method moves output to the next item from the iterable passed to `start()`.

2.9.14 PiRawImageMixin

class `picamera.PiRawImageMixin` (*parent, camera_port, input_port, format, resize, **options*)
Mixin class for “raw” (unencoded) image capture.

The `_callback_write()` method is overridden to manually calculate when to terminate output.

`_callback_write` (*buf, key=0*)

Overridden to manually calculate when to terminate capture (see comments in `__init__()`).

2.9.15 PiCookedOneImageEncoder

class `picamera.PiCookedOneImageEncoder` (*parent, camera_port, input_port, format, resize, **options*)
Encoder for “cooked” (encoded) single image output.

This encoder extends `PiOneImageEncoder` to include Exif tags in the output.

2.9.16 PiRawOneImageEncoder

class `picamera.PiRawOneImageEncoder` (*parent, camera_port, input_port, format, resize, **options*)
Single image encoder for unencoded capture.

This class is a derivative of `PiOneImageEncoder` and the `PiRawImageMixin` class intended for use with `capture()` (et al) when it is called with an unencoded image format.

Warning: This class creates an inheritance diamond. Take care to determine the MRO of super-class calls.

2.9.17 PiCookedMultiImageEncoder

class `picamera.PiCookedMultiImageEncoder` (*parent, camera_port, input_port, format, resize, **options*)

Encoder for “cooked” (encoded) multiple image output.

This encoder descends from `PiMultiImageEncoder` but includes no new functionality as video-port based encodes (which is all this class is used for) don’t support Exif tag output.

2.9.18 PiRawMultiImageEncoder

class `picamera.PiRawMultiImageEncoder` (*parent, camera_port, input_port, format, resize, **options*)

Multiple image encoder for unencoded capture.

This class is a derivative of `PiMultiImageEncoder` and the `PiRawImageMixin` class intended for use with `capture_sequence()` when it is called with an unencoded image format.

Warning: This class creates an inheritance diamond. Take care to determine the MRO of super-class calls.

2.9.19 PiRenderer

class `picamera.PiRenderer` (*parent, layer=0, alpha=255, fullscreen=True, window=None, crop=None, rotation=0, vflip=False, hflip=False*)

Base implementation of an MMAL video renderer for use by PiCamera.

The *parent* parameter specifies the `PiCamera` instance that has constructed this renderer. The *layer* parameter specifies the layer that the renderer will inhabit. Higher numbered layers obscure lower numbered layers (unless they are partially transparent). The initial opacity of the renderer is specified by the *alpha* parameter (which defaults to 255, meaning completely opaque). The *fullscreen* parameter which defaults to `True` indicates whether the renderer should occupy the entire display. Finally, the *window* parameter (which only has meaning when *fullscreen* is `False`) is a four-tuple of (*x*, *y*, *width*, *height*) which gives the screen coordinates that the renderer should occupy when it isn’t full-screen.

This base class isn’t directly used by `PiCamera`, but the two derivatives defined below, `PiOverlayRenderer` and `PiPreviewRenderer`, are used to produce overlays and the camera preview respectively.

close()

Finalizes the renderer and deallocates all structures.

This method is called by the camera prior to destroying the renderer (or more precisely, letting it go out of scope to permit the garbage collector to destroy it at some future time).

alpha

Retrieves or sets the opacity of the renderer.

When queried, the `alpha` property returns a value between 0 and 255 indicating the opacity of the renderer, where 0 is completely transparent and 255 is completely opaque. The default value is 255. The property can be set while recordings or previews are in progress.

crop

Retrieves or sets the area to read from the source.

The `crop` property specifies the rectangular area that the renderer will read from the source as a 4-tuple of (*x*, *y*, *width*, *height*). The special value (0, 0, 0, 0) (which is also the default) means to read entire area of the source. The property can be set while recordings or previews are active.

For example, if the camera’s resolution is currently configured as 1280x720, setting this attribute to (160, 160, 640, 400) will crop the preview to the center 640x400 pixels of the input.

Note that this property does not affect the size of the output rectangle, which is controlled with `fullscreen` and `window`.

Note: This property only affects the renderer; it has no bearing on image captures or recordings (unlike the `zoom` property of the `PiCamera` class).

fullscreen

Retrieves or sets whether the renderer appears full-screen.

The `fullscreen` property is a bool which controls whether the renderer takes up the entire display or not. When set to `False`, the `window` property can be used to control the precise size of the renderer display. The property can be set while recordings or previews are active.

hflip

Retrieves or sets whether the renderer's output is horizontally flipped.

When queried, the `vflip` property returns a boolean indicating whether or not the renderer's output is horizontally flipped. The property can be set while recordings or previews are in progress. The default is `False`.

Note: This property only affects the renderer; it has no bearing on image captures or recordings (unlike the `hflip` property of the `PiCamera` class).

layer

Retrieves or sets the layer of the renderer.

The `layer` property is an integer which controls the layer that the renderer occupies. Higher valued layers obscure lower valued layers (with 0 being the "bottom" layer). The default value is 2. The property can be set while recordings or previews are in progress.

rotation

Retrieves or sets the current rotation of the renderer.

When queried, the `rotation` property returns the rotation applied to the renderer. Valid values are 0, 90, 180, and 270.

When set, the property changes the rotation applied to the renderer's output. The property can be set while recordings or previews are active. The default is 0.

Note: This property only affects the renderer; it has no bearing on image captures or recordings (unlike the `rotation` property of the `PiCamera` class).

vflip

Retrieves or sets whether the renderer's output is vertically flipped.

When queried, the `vflip` property returns a boolean indicating whether or not the renderer's output is vertically flipped. The property can be set while recordings or previews are in progress. The default is `False`.

Note: This property only affects the renderer; it has no bearing on image captures or recordings (unlike the `vflip` property of the `PiCamera` class).

window

Retrieves or sets the size of the renderer.

When the `fullscreen` property is set to `False`, the `window` property specifies the size and position of the renderer on the display. The property is a 4-tuple consisting of (`x`, `y`, `width`, `height`). The property can be set while recordings or previews are active.

2.9.20 PiOverlayRenderer

```
class picamera.PiOverlayRenderer (parent, source, size=None, layer=0, alpha=255,
                                fullscreen=True, window=None, crop=None, rotation=0,
                                vflip=False, hflip=False)
```

Represents an MMAL renderer with a static source for overlays.

This class descends from `PiRenderer` and adds a static source for the MMAL renderer. The optional `size` parameter specifies the size of the source image as a `(width, height)` tuple. If this is omitted or `None` then the size is assumed to be the same as the parent camera's current `resolution`.

The `source` must be an object that supports the *buffer protocol* which has the same length as an image in RGB format (colors represented as interleaved unsigned bytes) with the specified `size` after the width has been rounded up to the nearest multiple of 32, and the height has been rounded up to the nearest multiple of 16.

For example, if `size` is `(1280, 720)`, then `source` must be a buffer with length `1280 x 720 x 3` bytes, or 2,764,800 bytes (because 1280 is a multiple of 32, and 720 is a multiple of 16 no extra rounding is required). However, if `size` is `(97, 57)`, then `source` must be a buffer with length `128 x 64 x 3` bytes, or 24,576 bytes (pixels beyond column 97 and row 57 in the source will be ignored).

The `layer`, `alpha`, `fullscreen`, and `window` parameters are the same as in `PiRenderer`.

update (*source*)

Update the overlay with a new source of data.

The new `source` buffer must have the same size as the original buffer used to create the overlay. There is currently no method for changing the size of an existing overlay (remove and recreate the overlay if you require this).

2.9.21 PiPreviewRenderer

```
class picamera.PiPreviewRenderer (parent, source, layer=2, alpha=255, fullscreen=True,
                                window=None, crop=None, rotation=0, vflip=False,
                                hflip=False)
```

Represents an MMAL renderer which uses the camera's preview as a source.

This class descends from `PiRenderer` and adds an MMAL connection to connect the renderer to an MMAL port. The `source` parameter specifies the MMAL port to connect to the renderer.

The `layer`, `alpha`, `fullscreen`, and `window` parameters are the same as in `PiRenderer`.

2.9.22 PiNullSink

```
class picamera.PiNullSink (parent, source)
```

Implements an MMAL null-sink which can be used in place of a renderer.

The `parent` parameter specifies the `PiCamera` instance which constructed this null-sink. The `source` parameter specifies the MMAL port which the null-sink should connect to its input.

The null-sink can act as a drop-in replacement for `PiRenderer` in most cases, but obviously doesn't implement attributes like `alpha`, `layer`, etc. as it simply dumps any incoming frames. This is also the reason that this class doesn't derive from `PiRenderer` like all other classes in this module.

close ()

Finalizes the null-sink and deallocates all structures.

This method is called by the camera prior to destroying the null-sink (or more precisely, letting it go out of scope to permit the garbage collector to destroy it at some future time).

2.9.23 Exceptions

exception `picamera.PiCameraWarning`

Base class for PiCamera warnings.

exception `picamera.PiCameraError`

Base class for PiCamera errors.

exception `picamera.PiCameraValueError`

Raised when an invalid value is fed to a `PiCamera` object.

exception `picamera.PiCameraRuntimeError`

Raised when an invalid sequence of operations is attempted with a `PiCamera` object.

exception `picamera.PiCameraClosed`

Raised when a method is called on a camera which has already been closed.

exception `picamera.PiCameraNotRecording`

Raised when `stop_recording()` or `split_recording()` are called against a port which has no recording active.

exception `picamera.PiCameraAlreadyRecording`

Raised when `start_recording()` or `record_sequence()` are called against a port which already has an active recording.

exception `picamera.PiCameraMMALError` (*status*, *prefix*='')

Raised when an MMAL operation fails for whatever reason.

`picamera.mmal_check` (*status*, *prefix*='')

Checks the return status of an mmal call and raises an exception on failure.

The *status* parameter is the result of an MMAL call. If *status* is anything other than `MMAL_SUCCESS`, a `PiCameraMMALError` exception is raised. The optional *prefix* parameter specifies a prefix message to place at the start of the exception's message to provide some context.

2.10 Array Extensions

The `picamera.array` module provides a set of classes which aid in constructing n-dimensional `numpy` arrays from the camera output. In order to avoid adding a hard dependency on `numpy` to `picamera`, the module is not automatically imported by the main `picamera` package and must be explicitly imported.

The following classes are defined in the module:

2.10.1 PiArrayOutput

class `picamera.array.PiArrayOutput` (*camera*, *size*=None)

Base class for capture arrays.

This class extends `io.BytesIO` with a `numpy` array which is intended to be filled when `flush()` is called (i.e. at the end of capture).

truncate (*size*=None)

Resize the stream to the given size in bytes (or the current position if size is not specified). This resizing can extend or reduce the current file size. The new file size is returned.

In prior versions of `picamera`, truncation also changed the position of the stream (because prior versions of these stream classes were non-seekable). This functionality is now deprecated; scripts should use `seek()` and `truncate()` as one would with regular `BytesIO` instances.

2.10.2 PiRGBArray

class `picamera.array.PiRGBArray` (*camera*, *size=None*)

Produces a 3-dimensional RGB array from an RGB capture.

This custom output class can be used to easily obtain a 3-dimensional numpy array, organized (rows, columns, colors), from an unencoded RGB capture. The array is accessed via the `array` attribute. For example:

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiRGBArray(camera) as output:
        camera.capture(output, 'rgb')
        print('Captured %dx%d image' % (
            output.array.shape[1], output.array.shape[0]))
```

You can re-use the output to produce multiple arrays by emptying it with `truncate(0)` between captures:

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiRGBArray(camera) as output:
        camera.resolution = (1280, 720)
        camera.capture(output, 'rgb')
        print('Captured %dx%d image' % (
            output.array.shape[1], output.array.shape[0]))
        output.truncate(0)
        camera.resolution = (640, 480)
        camera.capture(output, 'rgb')
        print('Captured %dx%d image' % (
            output.array.shape[1], output.array.shape[0]))
```

If you are using the GPU resizer when capturing (with the `resize` parameter of the various `capture()` methods), specify the resized resolution as the optional `size` parameter when constructing the array output:

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    with picamera.array.PiRGBArray(camera, size=(640, 360)) as output:
        camera.capture(output, 'rgb', resize=(640, 360))
        print('Captured %dx%d image' % (
            output.array.shape[1], output.array.shape[0]))
```

2.10.3 PiYUVArray

class `picamera.array.PiYUVArray` (*camera*, *size=None*)

Produces 3-dimensional YUV & RGB arrays from a YUV capture.

This custom output class can be used to easily obtain a 3-dimensional numpy array, organized (rows, columns, channel), from an unencoded YUV capture. The array is accessed via the `array` attribute. For example:

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiYUVArray(camera) as output:
```

```
camera.capture(output, 'yuv')
print('Captured %dx%d image' % (
    output.array.shape[1], output.array.shape[0]))
```

The `rgb_array` attribute can be queried for the equivalent RGB array (conversion is performed using the ITU-R BT.601 matrix):

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiYUVArray(camera) as output:
        camera.resolution = (1280, 720)
        camera.capture(output, 'yuv')
        print(output.array.shape)
        print(output.rgb_array.shape)
```

If you are using the GPU resizer when capturing (with the `resize` parameter of the various `capture()` methods), specify the resized resolution as the optional `size` parameter when constructing the array output:

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    camera.resolution = (1280, 720)
    with picamera.array.PiYUVArray(camera, size=(640, 360)) as output:
        camera.capture(output, 'yuv', resize=(640, 360))
        print('Captured %dx%d image' % (
            output.array.shape[1], output.array.shape[0]))
```

2.10.4 PiBayerArray

`class picamera.array.PiBayerArray(camera)`

Produces a 3-dimensional RGB array from raw Bayer data.

This custom output class is intended to be used with the `capture()` method, with the `bayer` parameter set to `True`, to include raw Bayer data in the JPEG output. The class strips out the raw data, constructing a 3-dimensional numpy array organized as (rows, columns, colors). The resulting data is accessed via the `array` attribute:

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiBayerArray(camera) as output:
        camera.capture(output, 'jpeg', bayer=True)
        print(output.array.shape)
```

Note that Bayer data is *always* full resolution, so the resulting array always has the shape (1944, 2592, 3); this also implies that the optional `size` parameter (for specifying a resizer resolution) is not available with this array class. As the sensor records 10-bit values, the array uses the unsigned 16-bit integer data type.

By default, `de-mosaicing` is **not** performed; if the resulting array is viewed it will therefore appear dark and too green (due to the green bias in the Bayer pattern). A trivial weighted-average demosaicing algorithm is provided in the `demosaic()` method:

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiBayerArray(camera) as output:
```

```
camera.capture(output, 'jpeg', bayer=True)
print(output.demosaic().shape)
```

Viewing the result of the de-mosaiced data will look more normal but still considerably worse quality than the regular camera output (as none of the other usual post-processing steps like auto-exposure, white-balance, vignette compensation, and smoothing have been performed).

2.10.5 PiMotionArray

class `picamera.array.PiMotionArray` (*camera*, *size=None*)

Produces a 3-dimensional array of motion vectors from the H.264 encoder.

This custom output class is intended to be used with the `motion_output` parameter of the `start_recording()` method. Once recording has finished, the class generates a 3-dimensional numpy array organized as (frames, rows, columns) where `rows` and `columns` are the number of rows and columns of **macro-blocks** (16x16 pixel blocks) in the original frames. There is always one extra column of macro-blocks present in motion vector data.

The data-type of the array is an (x, y, sad) structure where `x` and `y` are signed 1-byte values, and `sad` is an unsigned 2-byte value representing the **sum of absolute differences** of the block. For example:

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    with picamera.array.PiMotionArray(camera) as output:
        camera.resolution = (640, 480)
        camera.start_recording(
            '/dev/null', format='h264', motion_output=output)
        camera.wait_recording(30)
        camera.stop_recording()
        print('Captured %d frames' % output.array.shape[0])
        print('Frames are %dx%d blocks big' % (
            output.array.shape[2], output.array.shape[1]))
```

If you are using the GPU resizer with your recording, use the optional `size` parameter to specify the resizer's output resolution when constructing the array:

```
import picamera
import picamera.array

with picamera.PiCamera() as camera:
    camera.resolution = (640, 480)
    with picamera.array.PiMotionArray(camera, size=(320, 240)) as output:
        camera.start_recording(
            '/dev/null', format='h264', motion_output=output,
            resize=(320, 240))
        camera.wait_recording(30)
        camera.stop_recording()
        print('Captured %d frames' % output.array.shape[0])
        print('Frames are %dx%d blocks big' % (
            output.array.shape[2], output.array.shape[1]))
```

Note: This class is not suitable for real-time analysis of motion vector data. See the `PiMotionAnalysis` class instead.

2.10.6 PiAnalysisOutput

class `picamera.array.PiAnalysisOutput` (*camera*, *size=None*)

Base class for analysis outputs.

This class extends `io.IOBase` with a stub `analyse()` method which will be called for each frame output. In this base implementation the method simply raises `NotImplementedError`.

analyse (*array*)

Stub method for users to override.

2.10.7 PiRGBAnalysis

class `picamera.array.PiRGBAnalysis` (*camera*, *size=None*)

Provides a basis for per-frame RGB analysis classes.

This custom output class is intended to be used with the `start_recording()` method when it is called with *format* set to `'rgb'` or `'bgr'`. While recording is in progress, the `write()` method converts incoming frame data into a numpy array and calls the stub `analyse()` method with the resulting array (this deliberately raises `NotImplementedError` in this class; you must override it in your descendent class).

Warning: Because the `analyse()` method will be running within the encoder's callback, it must be **fast**. Specifically, it needs to return before the next frame is produced. Therefore, if the camera is running at 30fps, `analyse` cannot take more than 1/30s or 33ms to execute (and should take considerably less given that this doesn't take into account encoding overhead). You may wish to adjust the framerate of the camera accordingly.

The array passed to `analyse()` is organized as (rows, columns, channel) where the channels 0, 1, and 2 are R, G, and B respectively (or B, G, R if *format* is `'bgr'`).

2.10.8 PiYUVAnalysis

class `picamera.array.PiYUVAnalysis` (*camera*, *size=None*)

Provides a basis for per-frame YUV analysis classes.

This custom output class is intended to be used with the `start_recording()` method when it is called with *format* set to `'yuv'`. While recording is in progress, the `write()` method converts incoming frame data into a numpy array and calls the stub `analyse()` method with the resulting array (this deliberately raises `NotImplementedError` in this class; you must override it in your descendent class).

Warning: Because the `analyse()` method will be running within the encoder's callback, it must be **fast**. Specifically, it needs to return before the next frame is produced. Therefore, if the camera is running at 30fps, `analyse` cannot take more than 1/30s or 33ms to execute (and should take considerably less given that this doesn't take into account encoding overhead). You may wish to adjust the framerate of the camera accordingly.

The array passed to `analyse()` is organized as (rows, columns, channel) where the channel 0 is Y (luminance), while 1 and 2 are U and V (chrominance) respectively. The chrominance values normally have quarter resolution of the luminance values but this class makes all channels equal resolution for ease of use.

2.10.9 PiMotionAnalysis

class `picamera.array.PiMotionAnalysis` (*camera*, *size=None*)

Provides a basis for real-time motion analysis classes.

This custom output class is intended to be used with the *motion_output* parameter of the `start_recording()` method. While recording is in progress, the `write` method converts incoming motion data into numpy arrays and calls the stub `analyse()` method with the resulting array (which deliberately raises `NotImplementedError` in this class).

Warning: Because the `analyse()` method will be running within the encoder's callback, it must be **fast**. Specifically, it needs to return before the next frame is produced. Therefore, if the camera is running at 30fps, `analyse` cannot take more than 1/30s or 33ms to execute (and should take considerably less given that this doesn't take into account encoding overhead). You may wish to adjust the framerate of the camera accordingly.

The array passed to `analyse()` is organized as (rows, columns) where `rows` and `columns` are the number of rows and columns of **macro-blocks** (16x16 pixel blocks) in the original frames. There is always one extra column of macro-blocks present in motion vector data.

The data-type of the array is an (x, y, sad) structure where `x` and `y` are signed 1-byte values, and `sad` is an unsigned 2-byte value representing the **sum of absolute differences** of the block.

An example of a crude motion detector is given below:

```
import numpy as np
import picamera
import picamera.array

class DetectMotion(picamera.array.PiMotionAnalysis):
    def analyse(self, a):
        a = np.sqrt(
            np.square(a['x'].astype(np.float)) +
            np.square(a['y'].astype(np.float))
        ).clip(0, 255).astype(np.uint8)
        # If there're more than 10 vectors with a magnitude greater
        # than 60, then say we've detected motion
        if (a > 60).sum() > 10:
            print('Motion detected!')

with picamera.PiCamera() as camera:
    with DetectMotion(camera) as output:
        camera.resolution = (640, 480)
        camera.start_recording(
            '/dev/null', format='h264', motion_output=output)
        camera.wait_recording(30)
        camera.stop_recording()
```

You can use the optional `size` parameter to specify the output resolution of the GPU resizer, if you are using the `resize` parameter of `start_recording()`.

2.11 Change log

2.11.1 Release 1.8 (2014-09-05)

1.8 consists of several new features and the usual bug fixes:

- A new chapter on detecting and correcting deprecated functionality was added to the docs (#149)
- Stereoscopic cameras are now tentatively supported on the Pi compute module. Please note I have no hardware for testing this, so the implementation is possibly (probably!) wrong; bug reports welcome! (#153)
- Text annotation functionality has been extended; up to 255 characters are now possible, and the new `annotate_frame_num` attribute adds rendering of the current frame number. In addition, the new `annotate_background` flag permits a dark background to be rendered behind all annotations for contrast (#160)
- Arbitrary image overlays can now be drawn on the preview using the new `add_overlay()` method. A new recipe has been included demonstrating overlays from PIL images and numpy arrays. As part of this

work the preview system was substantially changed; all older scripts should continue to work but please be aware that most preview attributes are now deprecated; the new `preview` attribute replaces them (#144)

- Image effect parameters can now be controlled via the new `image_effect_params` attribute (#143)
- A bug in the handling of framerates meant that long exposures (>1s) weren't operating correctly. This *should* be fixed, but I'd be grateful if users could test this and let me know for certain (Exif metadata reports the configured exposure speed so it can't be used to determine if things are actually working) (#135)
- A bug in 1.7 broke compatibility with older firmwares (resulting in an error message mentioning "mmal_queue_timedwait"). The library should now on older firmwares (#154)
- Finally, the confusingly named `crop` attribute was changed to a deprecated alias for the new `zoom` attribute (#146)

2.11.2 Release 1.7 (2014-08-08)

1.7 consists once more of new features, and more bug fixes:

- Text overlay on preview, image, and video output is now possible (#16)
- Support for more than one camera on the compute module has been added, but hasn't been tested yet (#84)
- The `exposure_mode 'off'` has been added to allow locking down the exposure time, along with some new recipes demonstrating this capability (#116)
- The valid values for various attributes including `awb_mode`, `meter_mode`, and `exposure_mode` are now automatically included in the documentation (#130)
- Support for unencoded formats (YUV, RGB, etc.) has been added to the `start_recording()` method (#132)
- A couple of analysis classes have been added to `picamera.array` to support the new unencoded recording formats (#139)
- Several issues in the `PiBayerArray` class were fixed; this should now work correctly with Python 3, and the `demosaic()` method should operate correctly (#133, #134)
- A major issue with multi-resolution recordings which caused all recordings to stop prematurely was fixed (#136)
- Finally, an issue with the example in the documentation for custom encoders was fixed (#128)

Once again, many thanks to the community for another round of excellent bug reports - and many thanks to 6by9 and jamesh for their excellent work on the firmware and official utilities!

2.11.3 Release 1.6 (2014-07-21)

1.6 is half bug fixes, half new features:

- The `awb_gains` attribute is no longer write-only; you can now read it to determine the red/blue balance that the camera is using (#98)
- The new read-only `exposure_speed` attribute will tell you the shutter speed the camera's auto-exposure has determined, or the shutter speed you've forced with a non-zero value of `shutter_speed` (#98)
- The new read-only `analog_gain` and `digital_gain` attributes can be used to determine the amount of gain the camera is applying at a couple of crucial points of the image processing pipeline (#98)
- The new `drc_strength` attribute can be used to query and set the amount of dynamic range compression the camera will apply to its output (#110)
- The `intra_period` parameter for `start_recording()` can now be set to 0 (which means "produce one initial I-frame, then just P-frames") (#117)

- The *burst* parameter was added to the various `capture()` methods; users are strongly advised to read the cautions in the docs before relying on this parameter (#115)
- One of the advanced recipes in the manual (“splitting to/from a circular stream”) failed under 1.5 due to a lack of splitter-port support in the circular I/O stream class. This has now been rectified by adding a *splitter_port* parameter to the constructor of `PiCameraCircularIO` (#109)
- Similarly, the *array_extensions* introduced in 1.5 failed to work when resizers were present in the pipeline. This has been fixed by adding a *size* parameter to the constructor of all the custom output classes defined in that module (#121)
- A bug that caused picamera to fail when the display was disabled has been squashed (#120)

As always, many thanks to the community for another great set of bug reports!

2.11.4 Release 1.5 (2014-06-11)

1.5 fixed several bugs and introduced a couple of major new pieces of functionality:

- The new `picamera.array` module provides a series of custom output classes which can be used to easily obtain numpy arrays from a variety of sources (#107)
- The *motion_output* parameter was added to `start_recording()` to enable output of motion vector data generated by the H.264 encoder. A couple of new recipes were added to the documentation to demonstrate this (#94)
- The ability to construct custom encoders was added, including some examples in the documentation. Many thanks to user Oleksandr Sviridenko (d2rk) for helping with the design of this feature! (#97)
- An example recipe was added to the documentation covering loading and conversion of raw Bayer data (#95)
- Speed of unencoded RGB and BGR captures was substantially improved in both Python 2 and 3 with a little optimization work. The warning about using alpha-inclusive modes like RGBA has been removed as a result (#103)
- An issue with out-of-order calls to `stop_recording()` when multiple recordings were active was resolved (#105)
- Finally, picamera caught up with raspistill and raspivid by offering a friendly error message when used with a disabled camera - thanks to Andrew Scheller (lurch) for the suggestion! (#89)

2.11.5 Release 1.4 (2014-05-06)

1.4 mostly involved bug fixes with a couple of new bits of functionality:

- The *sei* parameter was added to `start_recording()` to permit inclusion of “Supplemental Enhancement Information” in the output stream (#77)
- The *awb_gains* attribute was added to permit manual control of the auto-white-balance red/blue gains (#74)
- A bug which cause `split_recording()` to fail when low framerates were configured was fixed (#87)
- A bug which caused picamera to fail when used in UNIX-style daemons, unless the module was imported *after* the double-fork to background was fixed (#85)
- A bug which caused the *frame* attribute to fail when queried in Python 3 was fixed (#80)
- A bug which caused raw captures with “odd” resolutions (like 100x100) to fail was fixed (#83)

Known issues:

- Added a workaround for full-resolution YUV captures failing. This isn’t a complete fix, and attempting to capture a JPEG before attempting to capture full-resolution YUV data will still fail, unless the GPU memory split is set to something huge like 256Mb (#73)

Many thanks to the community for yet more excellent quality bug reports!

2.11.6 Release 1.3 (2014-03-22)

1.3 was partly new functionality:

- The `bayer` parameter was added to the 'jpeg' format in the capture methods to permit output of the camera's raw sensor data (#52)
- The `record_sequence()` method was added to provide a cleaner interface for recording multiple consecutive video clips (#53)
- The `splitter_port` parameter was added to all capture methods and `start_recording()` to permit recording multiple simultaneous video streams (presumably with different options, primarily `resize`) (#56)
- The limits on the `framerate` attribute were increased after firmware #656 introduced numerous new camera modes including 90fps recording (at lower resolutions) (#65)

And partly bug fixes:

- It was reported that Exif metadata (including thumbnails) wasn't fully recorded in JPEG output (#59)
- Raw captures with `capture_continuous()` and `capture_sequence()` were broken (#55)

2.11.7 Release 1.2 (2014-02-02)

1.2 was mostly a bug fix release:

- A bug introduced in 1.1 caused `split_recording()` to fail if it was preceded by a video-port-based image capture (#49)
- The documentation was enhanced to try and full explain the discrepancy between preview and capture resolution, and to provide some insight into the underlying workings of the camera (#23)
- A new property was introduced for configuring the preview's layer at runtime although this probably won't find use until OpenGL overlays are explored (#48)

2.11.8 Release 1.1 (2014-01-25)

1.1 was mostly a bug fix release:

- A nasty race condition was discovered which led to crashes with long-running processes (#40)
- An assertion error raised when performing raw captures with an active `resize` parameter was fixed (#46)
- A couple of documentation enhancements made it in (#41 and #47)

2.11.9 Release 1.0 (2014-01-11)

In 1.0 the major features added were:

- Debian packaging! (#12)
- The new `frame` attribute permits querying information about the frame last written to the output stream (number, timestamp, size, keyframe, etc.) (#34, #36)
- All capture methods (`capture()` et al), and the `start_recording()` method now accept a `resize` parameter which invokes a resizer prior to the encoding step (#21)
- A new `PiCameraCircularIO` stream class is provided to permit holding the last n seconds of video in memory, ready for writing out to disk (or whatever you like) (#39)
- There's a new way to specify raw captures - simply use the format you require with the capture method of your choice. As a result of this, the `raw_format` attribute is now deprecated (#32)

Some bugs were also fixed:

- `GPIO.cleanup` is no longer called on `close()` (#35), and GPIO set up is only done on first use of the `led` attribute which should resolve issues that users have been having with using picamera in conjunction with GPIO
- Raw RGB video-port based image captures are now working again too (#32)

As this is a new major-version, all deprecated elements were removed:

- The continuous method was removed; this was replaced by `capture_continuous()` in 0.5 (#7)

2.11.10 Release 0.8 (2013-12-09)

In 0.8 the major features added were:

- Capture of images whilst recording without frame-drop. Previously, images could be captured whilst recording but only from the still port which resulted in dropped frames in the recorded video due to the mode switch. In 0.8, `use_video_port=True` can be specified on capture methods whilst recording video to avoid this.
- Splitting of video recordings into multiple files. This is done via the new `split_recording()` method, and requires that the `start_recording()` method was called with `inline_headers` set to `True`. The latter has now been made the default (technically this is a backwards incompatible change, but it's relatively trivial and I don't anticipate anyone's code breaking because of this change).

In addition a few bugs were fixed:

- Documentation updates that were missing from 0.7 (specifically the new video recording parameters)
- The ability to perform raw captures through the video port
- Missing exception imports in the encoders module (which caused very confusing errors in the case that an exception was raised within an encoder thread)

2.11.11 Release 0.7 (2013-11-14)

0.7 is mostly a bug fix release, with a few new video recording features:

- Added `quantisation` and `inline_headers` options to `start_recording()` method
- Fixed bugs in the `crop` property
- The issue of captures fading to black over time when the preview is not running has been resolved. This solution was to permanently activate the preview, but pipe it to a null-sink when not required. Note that this means rapid capture gets even slower when not using the video port
- LED support is via RPi.GPIO only; the RPIO library simply doesn't support it at this time
- Numerous documentation fixes

2.11.12 Release 0.6 (2013-10-30)

In 0.6, the major features added were:

- New `'raw'` format added to all capture methods (`capture()`, `capture_continuous()`, and `capture_sequence()`) to permit capturing of raw sensor data
- New `raw_format` attribute to permit control of raw format (defaults to `'yuv'`, only other setting currently is `'rgb'`)
- New `shutter_speed` attribute to permit manual control of shutter speed (defaults to 0 for automatic shutter speed, and requires latest firmware to operate - use `sudo rpi-update` to upgrade)

- New “Recipes” chapter in the documentation which demonstrates a wide variety of capture techniques ranging from trivial to complex

2.11.13 Release 0.5 (2013-10-21)

In 0.5, the major features added were:

- New `capture_sequence()` method
- `continuous()` method renamed to `capture_continuous()`. Old method name retained for compatibility until 1.0.
- `use_video_port` option for `capture_sequence()` and `capture_continuous()` to allow rapid capture of JPEGs via video port
- New `framerate` attribute to control video and rapid-image capture frame rates
- Default value for `ISO` changed from 400 to 0 (auto) which fixes `exposure_mode` not working by default
- `intraperiod` and `profile` options for `start_recording()`

In addition a few bugs were fixed:

- Byte strings not being accepted by `continuous()`
- Erroneous docs for `ISO`

Many thanks to the community for the bug reports!

2.11.14 Release 0.4 (2013-10-11)

In 0.4, several new attributes were introduced for configuration of the preview window:

- `preview_alpha`
- `preview_fullscreen`
- `preview_window`

Also, a new method for rapid continual capture of still images was introduced: `continuous()`.

2.11.15 Release 0.3 (2013-10-04)

The major change in 0.3 was the introduction of custom Exif tagging for captured images, and fixing a silly bug which prevented more than one image being captured during the lifetime of a PiCamera instance.

2.11.16 Release 0.2

The major change in 0.2 was support for video recording, along with the new `resolution` property which replaced the separate `preview_resolution` and `stills_resolution` properties.

2.12 License

Copyright 2013,2014 Dave Hughes

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The *bayer pattern diagram* in the documentation is derived from [Bayer_pattern_on_sensor.svg](#) which is copyright (c) Colin Burnett (User:Cburnett) on Wikipedia, modified under the terms of the GPL:

This work is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or any later version. This work is distributed in the hope that it will be useful, but without any warranty; without even the implied warranty of merchantability or fitness for a particular purpose. See version 2 and version 3 of the GNU General Public License for more details.

The *YUV420 planar diagram* in the documentation is [Yuv420.svg](#) created by Geoff Richards (User:Qef) on Wikipedia, released into the public domain.

Indices and tables

- *genindex*
- *modindex*
- *search*

p

`picamera`, 59

`picamera.array`, 92